



Digital Commons@

Loyola Marymount University
LMU Loyola Law School

LMU/LLS Theses and Dissertations

2022

Exploring Algorithmic Literacy for College Students: An Educator's Roadmap

Susan Gardner Archambault
Loyola Marymount University

Follow this and additional works at: <https://digitalcommons.lmu.edu/etd>



Part of the [Artificial Intelligence and Robotics Commons](#), [Higher Education Commons](#), and the [Library and Information Science Commons](#)

Recommended Citation

Archambault, Susan Gardner, "Exploring Algorithmic Literacy for College Students: An Educator's Roadmap" (2022). *LMU/LLS Theses and Dissertations*. 1160.
<https://digitalcommons.lmu.edu/etd/1160>

This Dissertation is brought to you for free and open access by Digital Commons @ Loyola Marymount University and Loyola Law School. It has been accepted for inclusion in LMU/LLS Theses and Dissertations by an authorized administrator of Digital Commons@Loyola Marymount University and Loyola Law School. For more information, please contact digitalcommons@lmu.edu.



Digital Commons@

Loyola Marymount University
LMU Loyola Law School

LMU/LLS Theses and Dissertations

2022

Exploring Algorithmic Literacy for College Students: An Educator's Roadmap

Susan Gardner Archambault

Follow this and additional works at: <https://digitalcommons.lmu.edu/etd>



Part of the [Artificial Intelligence and Robotics Commons](#), [Higher Education Commons](#), and the [Library and Information Science Commons](#)

This Dissertation is brought to you for free and open access by Digital Commons @ Loyola Marymount University and Loyola Law School. It has been accepted for inclusion in LMU/LLS Theses and Dissertations by an authorized administrator of Digital Commons@Loyola Marymount University and Loyola Law School. For more information, please contact digitalcommons@lmu.edu.

LOYOLA MARYMOUNT UNIVERSITY

Exploring Algorithmic Literacy for College Students:
An Educator's Roadmap

by

Susan Gardner Archambault

A dissertation presented to the Faculty of the School of Education,

Loyola Marymount University,

in partial satisfaction of the requirements for the degree

Doctor of Education

2022

Exploring Algorithmic Literacy for College Students:

An Educator's Roadmap

Copyright © 2022

by

Susan Gardner Archambault


**Loyola Marymount University
School of Education
Los Angeles, CA 90045**

This dissertation written by Susan Archambault, under the direction of the Dissertation Committee, is approved and accepted by all committee members, in partial fulfillment of requirements for the degree of Doctor of Education.

April 12, 2022

Date

Dissertation Committee



Philip Molebash, Ph.D., Dissertation Chair



[\(Apr 12, 2022 13:12 PDT\)](#)

Anna Bargagliotti, Ph.D., Committee Member



[\(Apr 12, 2022 13:40 PDT\)](#)

Maryann Krikorian, Ph.D., Committee Member

ACKNOWLEDGEMENTS

Countless hours were spent on this research project, and I want to acknowledge the people who guided me throughout the process:

To the sixteen participants of this study, thank you for your time, your thoughtful reflections, your candidness, and your leadership as scholars and teachers working in critical algorithm studies and related fields. Without you, there would have been no study. I know that the wisdom you shared during this process will inspire future practitioners and researchers.

To my Chair, Dr. Philip Molebash, thank you for your mentorship and encouragement. Thank you for wanting to take on this challenging topic even during the middle of a pandemic. Thank you for your constructive feedback that always resulted in a better dissertation, and for having my best interests at heart.

To my Committee Members, Dr. Anna Bargagliotti and Dr. Maryann Krikorian, thank you for your expertise and thoughtful feedback. Thank you for setting me up on the path to success early on with your suggested revisions and probing questions—they had a lasting impact on the shape of the study.

To the Reference and Instruction Department in the William H. Hannon Library, thank you for your encouragement and willingness to occasionally fill in for me so I could take time off to work on this research. I appreciate you. To my boss, Dean Kristine Brancolini, thank you for encouraging me from start to finish on this journey and sharing my enthusiasm for this research topic. Thank you for even bringing it into the Dean's Book Club.

To Cohort 16, thank you for creating such a welcoming community. You are truly an inspiring group of future leaders, and I have learned from each of you.

DEDICATION

This dissertation is dedicated to my parents, who taught me the value of a good education and have always believed in me. Dad, I wish you were still here to see this. It is also dedicated to my husband Scott and dog Leo for their support, patience, and sacrifices during the three years of my doctoral program. Finally, I also dedicate this to my sister Beth.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iii
DEDICATION	iv
LIST OF TABLES	ix
LIST OF FIGURES	x
ABSTRACT	xi
CHAPTER 1: INTRODUCTION	1
Background.....	5
Statement of the Problem	10
History of Algorithms.....	13
Algorithmic Justice.....	15
Purpose of the Study.....	18
Research Questions	19
Significance of the Study.....	19
Conceptual Framework	21
Research Design and Methodology.....	21
Limitations and Delimitations	22
Definitions of Key Terms.....	23
Organization of Dissertation.....	29
CHAPTER 2: REVIEW OF THE LITERATURE	30
Introduction	30
Information Literacy.....	30
Critical Information Literacy Movement	33
ACRL Framework	34
Related Frameworks	38
Conceptual Framework	40
Recent Studies on Information Literacy.....	47
Beginning Steps: Algorithms Embedded in Information Literacy Instruction	49
Additional Teaching Tools and Strategies From Other Fields	55
Algorithms.....	58
Technical Aspects of Algorithms	59
Google Search Engine Algorithms	64
YouTube Recommender Algorithm	67
Instagram	69
Twitter Timeline Algorithm	71
Facebook Ranking Algorithm	72
Algorithms for News	74
Amazon Item-Based Collaborative Filtering.....	76
Social Aspects of Algorithms	78
Power of Algorithms	80
Algorithmic Bias and Ethics.....	83

Filter Bubbles in News Feeds.....	88
Fake News	91
User Studies on Algorithms.....	95
Social Media User Studies.....	95
Awareness.....	95
Folk Theories.....	98
Behavior and Attitudes	100
Curated News User Studies	102
Search Engine User Studies.....	104
User Studies in Other Domains	104
Algorithmic Bias User Studies	106
User Studies on College Students.....	107
Measurements.....	110
Knowledge and Awareness	110
Attitudes, Perceptions, and Behaviors.....	113
Conclusion.....	115
CHAPTER 3: METHODOLOGY	117
Research Questions	118
Method.....	118
Participants	119
Setting.....	122
Data Collection.....	122
Analysis Plan.....	124
List of Codes in Alphabetical Order Across All Themes.....	126
Dedoose Analysis	131
Limitations and Delimitations	132
Validity and Trustworthiness	133
Conclusion.....	134
CHAPTER 4: RESULTS AND MAIN FINDINGS.....	136
Background.....	136
Introductory Summary of Findings	139
Theme 1: Knowledge Components of Algorithmic Literacy	143
General Characteristics of Algorithms	143
Black Box	144
Common Misconceptions.....	147
Algorithm Strengths	151
Machine Learning Versus Rule Based	154
Key Algorithmic Operations	158
Human-Algorithm Interplay.....	161
Algorithmic Reasoning Versus Human Reasoning.....	164
Infrastructure	168
Key Domains in Everyday Life Using Algorithms	169

Social Services.....	171
Social Media	174
Search	177
Health Care.....	179
Human Resources	180
Recreation and Retail	181
Finance	182
Education	183
Ethical Considerations.....	185
Bias Causes.....	185
Privacy of Personal Data	189
Accountability	192
Oppression, Privilege, Power	195
Commodification	197
Theme 2: Student Behaviors and Coping Strategies	199
Gaming the System.....	200
Broader Advocacy	205
Critical Evaluation.....	207
Skepticism Anarchy.....	209
Transfer of Knowledge.....	212
Theme 3: Pedagogy	213
Algorithmic Awareness (lower level)	214
Algorithmic Literacy (higher level).....	217
Challenges of Assessment	219
Placement in the Curriculum	224
Tactile (building)	228
Folk Theories.....	231
Everyday Life	235
Data Transparency Tools.....	237
Mass Media	240
Speculative Futuring.....	241
Contextualized Stories.....	243
Perspective Taking	244
Expectancy Violations.....	246
Conclusion.....	248
CHAPTER 5: DISCUSSION AND IMPLICATIONS.....	251
Findings	251
Study Background	251
Discussion of Findings	252
Knowledge Components	252
General Characteristics of Algorithms	252
Key Domains in Everyday Life Using Algorithms	257
Ethical Considerations for the Use of Algorithms	261

Student Coping Behaviors	266
Pedagogy	272
Algorithmic Awareness Versus Algorithmic Literacy	272
Challenges of Assessment	274
Placement in the Curriculum	276
Teaching Strategies	277
Limitations.....	285
Future Research	286
Implications and Recommendations.....	287
Theoretical Implications and Recommendations	287
Practical Implications and Recommendations	297
Policy Implications and Recommendations	307
Conclusion.....	312
APPENDIX A.....	314
APPENDIX B.....	315
APPENDIX C.....	317
APPENDIX D.....	319
REFERENCES	320

LIST OF TABLES

Table	Page
1. Select Advocacy and Research Groups	17
2. Interview Participant Demographic Data	121
3. Focus Group Participant Demographic Data.....	122
4. Code Frequencies Across All Themes	141
5. Code Frequencies: General Characteristics of Algorithms	144
6. Code Frequencies: Key Domains in Everyday Life	171
7. Code Frequencies: Ethical Considerations	185
8. Code Frequencies: Coping Behavior.....	200
9. Code Frequencies: Pedagogy	214
10. Expanding Frame 1: Authority is Constructed and Contextual.....	291
11. Expanding Frame 2: Information Creation as a Process	293
12. Expanding Frame 3: Information has Value	295
13. Expanding Frame 6: Searching as Strategic Exploration	297
14. Selected Activities for Students Related to Algorithmic Literacy	301

LIST OF FIGURES

Figure	Page
1. Algorithmic Literacy Across Frames 1-3 and Frame 6	43
2. Final Codes Organized by Theme	139
3. Expansion of Framework for Information Literacy for Higher Education.....	289
4. Dimensions of Teaching Algorithmic Literacy to Non-Computer Science Majors	298

ABSTRACT

Exploring Algorithmic Literacy for College Students: An Educator's Roadmap

by

Susan Gardner Archambault

Research shows that college students are largely unaware of the impact of algorithms on their everyday lives. Also, most university students are not being taught about algorithms as part of the regular curriculum. This exploratory, qualitative study aimed to explore subject-matter experts' insights and perceptions of the knowledge components, coping behaviors, and pedagogical considerations to aid faculty in teaching algorithmic literacy to college students. Eleven individual, semi-structured interviews and one focus group were conducted with scholars and teachers of critical algorithm studies and related fields. Findings suggested three sets of knowledge components that would contribute to students' algorithmic literacy: general characteristics and distinguishing traits of algorithms, key domains in everyday life using algorithms (including the potential benefits and risks), and ethical considerations for the use and application of algorithms. Findings also suggested five behaviors that students could use to help them better cope with algorithmic systems and nine teaching strategies to help improve students' algorithmic literacy. Suggestions also surfaced for alternative forms of assessment, potential placement in the curriculum, and how to distinguish between basic algorithmic awareness

compared to algorithmic literacy. Recommendations for expanding on the current Association of College and Research Libraries' *Framework for Information Literacy for Higher Education* (2016) to more explicitly include algorithmic literacy were presented.

CHAPTER 1

INTRODUCTION

There is growing concern over how much of everyday life is increasingly mediated through algorithmic decision-making (Head et al., 2020). Picture a typical day in the life of college student “CJ” and the monumental data exhaust trail they leave behind, as their personal information accidentally gets generated using computers, cell phones, and behavior surveillance. In the morning, CJ unlocks their iPhone through Face ID and reads news stories heavily filtered through the Google News aggregator (<https://news.google.com>) based on their presumed interests and geographical location. Meanwhile, they get nudged to buy all kinds of pet products thanks to a previous Google search on best puppy names. In the afternoon, CJ submits a resume for a job that confuses the screening algorithm in the Human Resources department because it cannot infer the gender or race of the applicant. CJ’s online Statistics 101 test is auto proctored, and thanks to getting interrupted by their roommate they are flagged for cheating because of the presence of a second person in the room. In the evening, CJ logs into Tinder and looks at a list of potential romantic partners based on the algorithm’s prediction of who they will like. Throughout all of this, CJ’s smartwatch has tracked every heartbeat. In a typical day, CJ will “release hundreds of megabytes of data to digital platforms with little understanding of how it will be used to persuade and manipulate them” (Hobbs, 2020, p. 530).

The information literacy skills that students learn in college within an academic context may not always carry over to their personal and civic lives. Information literacy is defined as “the set of integrated abilities encompassing the reflective discovery of information, the understanding of how information is produced and valued, and the use of information in creating

new knowledge and participating ethically in communities of learning” (Association of College and Research Libraries, 2016, p. 7). This definition incorporates not only elements from traditional library instruction with a focus on finding and evaluating information, but also newer elements of media literacy, digital literacy, news literacy, and critical thinking (Head et al., 2020). The information landscape is increasingly complex, and it is important that students understand the underlying structures at play in the information systems they use for both their academic and personal lives. To be “information literate” now, students need a new set of information skills, the most pressing of which is learning how information works in the age of algorithms (Head et al., 2020). Experts warn that algorithms put too much control in the hands of corporations and governments, perpetuate bias, create filter bubbles, and limit personal choices (Rainie & Anderson, 2017). Recent studies from the Pew Research Center found that algorithmically generated content platforms such as Facebook played a large role in how Americans 18 and over got their news, but many did not understand how the news feeds work or were unaware that they target advertising based on their individual online behaviors (Smith, 2019). Similarly, a 2017 survey of college students found that most were unaware that the news they got from Google and through Facebook was filtered using algorithms (Powers, 2017).

Most information literacy instruction students receive in college is limited to helping them write a traditional research paper. A content analysis of course-related research assignment handouts across 28 U.S. colleges found that faculty expected students to mostly use library resources, scholarly databases, and course readings as sources (Head & Eisenberg, 2010). Another study found that students frequently used sources uploaded by their instructor; when they did look for outside sources, they favored the top search results from search engines which

made them “over-reliant on, and passive users of, the decisions of popular search engines” (Bhatt & MacKenzie, 2019, p. 315). A focus group study of 103 students across eight different colleges and universities found that algorithms barely, if ever, were taught in the classroom (Head et al., 2020). For the most part, the information students needed for school assignments had nothing to do with life beyond college, and they got largely outdated advice in their courses about the internet such as trusting .org and .edu domains (Head et al., 2020). The same study interviewed 37 faculty and found that the majority encouraged students to use peer-reviewed research and do close reading and textual analysis, but they did not discuss the role that algorithms played in the current information landscape (Head et al.).

Another study that surveyed and interviewed students across 11 different colleges and universities found that students’ news engagement behaviors were very different for academic versus personal purposes (Head et al., 2019). For personal purposes, 77% used social media as a pathway to news while only 2% used it for academic purposes, but 83% used library databases for academic purposes largely because news was not accepted by their professors as valid information or had no bearing on the assignment. The authors went on to recommend more professors integrate news into the classroom to help students question what they are not seeing in the news they receive, what stories are prominent, and which are ignored. Students need better evaluation skills to navigate free information platforms, where algorithms are interacting behind the scenes with human behavioral data to give dynamic rather than static personalized output. Users of online platforms need algorithmic knowledge to better analyze the context and various factors that influenced their search results in order to make independent and evaluative judgements as a critical consumer (Cotter & Reisdorf, 2020). Bakke (2020) noted that

information literacy instruction is often limited to a “checklist” approach for source evaluation skills, but most checklists have features that can easily be manipulated or faked in the post-truth era. Also, checklists encourage students to gloss over the outside context. One of the top 20 trends in 2020 for academic libraries is machine learning and artificial intelligence, with a call for libraries to use information literacy instruction as a vector to introduce ethical considerations in algorithmic concepts (Association of College and Research Libraries Research Planning and Review Committee, 2020).

Algorithms are sets of instructions or sequences of logical steps for a computer to use on a body of data in order to accomplish a task, such as organizing search results by relevance (Gillespie, 2014). For example, the task of giving a user the most relevant search results might involve calculating the combined values of pre-weighted objects in the index database in order to make it likely the user clicks on one of the first five results (Gillespie, 2016a). Another way to describe an algorithm is as a recipe, or a step-by-step guide that prescribes how to obtain a certain goal, given specific parameters (Bucher, 2018). Engineers choose algorithms based on values such as how quickly they return the result and the load they impose on the system’s available memory (Gillespie, 2016a). Algorithms are “trained” on a corpus of known data that has been certified in some way by the designer or past user behavior (Gillespie). Algorithms based on machine learning rewrite themselves as they work rather than repeatedly processing a stable set of instructions (Brogan, 2016). The algorithms that power consumer online platforms serve as invisible lines of coding that are used to personalize content to match users likes and dislikes (Head et al., 2020). For example, algorithms are used in Netflix film recommendations (www.netflix.com), Amazon purchase recommendations (www.amazon.com), and to filter our

social media news feeds. This sometimes has the positive effect of making our lives more efficient and productive. However, people knowingly or unknowingly delegate slices of authority to algorithms, resulting in the potential for a narrowing of their information landscape and a yielding of their agency (Lloyd, 2019). A global study conducted by Pegasystems Inc. (2019) found that although 84% of participants used one or more AI-powered devices such as email spam filters, predictive search terms, virtual assistants, Facebook-recommended news, and online shopping recommendations during the past year; only 34% thought they had interacted with some sort of AI technology. This gap between use and awareness must be eliminated.

Background

Shoshana Zuboff (2019) cautioned in *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* that only some of the behavioral data collected are applied to product or service improvement, while the rest are labeled as “behavioral surplus” and fed into prediction products fueled by machine algorithms. The prediction products get traded into what she called the “behavioral futures markets.” The development that has allowed technologists and systems to collect and process data in real time on a large scale is called “big data” (Head et al., 2020). The three distinctive traits of big data are large volume, velocity, and a scope so large that it approaches whole populations rather than only a sample (Kitchin & Lauriault, 2015). In some cases, such as with visual pattern recognition, voice recognition, or geolocation tracking, the United States federal government partners with third party companies to exploit user data scraped from the internet, leading to a surveillance state. A good example of this is federal use of facial recognition software, which is known to misidentify people with darker skin tones or disabled people more often as “outsiders” or “intruders” because the sets of

faces the algorithms were trained on were not diverse enough (Buolamwini & Gebru, 2018). Such misidentification could lead to unwarranted detainment and arrest. One of the most egregious examples of misidentification was when Google Photos automatically labelled images of two Black friends as “gorillas,” its algorithm digging up years of institutionalized racism (Reidsma, 2019). Another example was when a Google Maps search on the [n-word] led to a map of the White House during Obama’s presidency (Noble, 2018).

Algorithmic literacy is a growing subset of information literacy, and it teaches “critical awareness of what algorithms are, how they interact with human behavioral data in information systems, and social and ethical issues related to their use” (Head et al., 2020, p. 49). A recent study reported that 68% of adults felt it was unacceptable for companies to collect large quantities of data about individuals because these programs violate people’s privacy, are unfair, and/or do not work as well as decisions made by humans (Smith, 2018). Unfortunately, current regulatory laws in the United States do little to protect user privacy. Safiya Noble felt the public “must have significantly more algorithmic literacy” (Noble, 2018, p. 25). College students, in particular, need to understand the wide-reaching social impact of algorithm-driven platforms like Google, Facebook, and Amazon that they use in their daily life. In *The Googlization of Everything: (And Why We Should Worry)*, Siva Vaidhyanathan (2011) warned that Google’s core business model is about selling access to our attention to advertisers, not facilitating searches. The goal of these information systems is to make money, not be a common good or a shared community (Lindh & Nolin, 2016). To understand and use information most effectively, students must recognize how personal data is being manipulated to control the information they see and paint an unbalanced picture.

Safiya Noble (2018) argued in *Algorithms of Oppression: How Search Engines Reinforce Racism* that algorithms powering search engines promote bias against women and people of color. Noble presented screenshots from Google searches of “Black girls” to show the pornification of Black women as sex objects and noted that “Google’s dominant narratives reflect the kind of hegemonic framework and notions that are often resisted by women and people of color” (Noble, 1998, p. 24). Along with pointing out the psychological damage done to those misrepresented in information systems, Noble also pointed out that racism and sexism are profitable. As Sara Wachter-Boettcher (2017) noted, systems using algorithms are designed to be used by people like their creators, usually privileged cisgender heterosexual White males. Technology products carry the imprint of the dominant culture since the industry has very little racial or gender diversity (and is particularly lacking in Latinx or Black employees) (Wachter-Boettcher). The algorithms make choices about which results to filter, which to show on the first page, and which to exclude, and these decisions reflect the underlying systemic racism and sexism prevalent in society.

Algorithms even impacted recent elections; for example, Cambridge Analytica used algorithms to develop a targeted advertising campaign to Facebook (www.facebook.com) profiles that influenced users’ voting habits (Rosenberg et al., 2018). The data collected by Cambridge Analytica (<http://cambridgeanalytica.org/>, but this company no longer exists) was sourced through a seemingly benign online personality quiz called “This is Your Digital Life” that profiled individuals based on short measures that calculated openness, conscientiousness, extroversion, agreeableness, and neuroticism as described by Donnellan et al. (2006). By accepting usage of the app, Cambridge Analytica was able to pair this with a user’s Facebook

data, including their entire network of friends and their private messages (Cohen, 2018). The majority of fake news stories from the 2016 election were engineered purely to make money and revenue from online advertising networks, particularly Google's, and designed without any political motive (Graham, 2017). The creators often would reinvest part of their earnings to buy fake Facebook profiles and paid Facebook directly to promote their pages until they became viral (Graham). The recent COVID-19 pandemic has increased the use of digital platforms and led to an even greater need for deeper understanding of third-party data collection practices. The rapid spread of both disinformation and misinformation on social media surrounding the pandemic has also added unnecessary confusion and fear, and it has adversely impacted peoples' judgment (Bellingcat Investigation Team, 2020). Four types of false information were spread: 1) generic information, such as videos of crowds hoarding food items purportedly because of the pandemic but really from a different time period; 2) disinformation from credible sources, such as government officials insinuating COVID-19 was a bioweapon; 3) fake sources attempting to appear credible, such as someone on Twitter using a fake Massachusetts Institute of Technology identity; and 4) profiteers such as people selling bogus products claiming to help ward off sickness (Bellingcat Investigation Team, 2020).

This points to a growing concern over the ways in which algorithms can unduly influence us, divide us, and create the potential for everyone to experience different realities. Algorithms have a strong social justice component, and students need to understand how algorithms use personal data in underlying information architectures which have a wide-reaching influence on both individuals and the public sphere. Algorithms usually lack data for the behaviors they're interested in, so they substitute stand-in data, or proxies, to correlate with the behavior (O'Neil,

2016). Cathy O’Neil (2016) explained how algorithms often rely on incomplete, inaccurate, or unfair data, and this perpetuates existing bias and societal prejudices on a large scale. She explained that racism is analogous to algorithms because both are “powered by haphazard data gathering and spurious correlations, reinforced by institutional inequities, and powered by confirmation bias” (O’Neil, 2016, p. 23). Both make the assumption that because certain types of people have behaved a certain way in the past, this results in “a binary prediction that all people” with similar characteristics will behave that same way (O’Neil, 2016, p. 23). One of her key examples was how algorithms were used to determine credit scores—sometimes a person’s zip code was a variable (their IP address was matched with real estate data) as well as their Facebook friends’ credit scores; in both cases, people in affluent neighborhoods with more affluent friends will get a higher score, further perpetuating the wealth gap. In our largely segregated cities, wealth is usually a proxy for race, and the poor in the riskier demographics are offered less available credit and higher interest rates which codifies past injustices. Nopper (2019) described how many financial technology (FinTech) companies now use social media information (for example, quantity of social media contacts and frequency of interactions) or search engine history to determine your “digital character,” a proxy for credibility, stability, and reliability in terms of paying back a loan.

Similarly, algorithms used in the criminal justice system to deploy police use zip codes as a variable to determine criminal hot spots likely to meet quotas, with the end result being that it serves as a proxy for racial profiling resulting in more arrests in low-income areas which further perpetuates the cycle (O’Neil, 2016). This results in the sorting of winners and losers through class and race (O’Neil). In another example, prison sentences are partially based on predictions

of recidivism by algorithms, and one variable used for this prediction is when the person had their first encounter with law enforcement. Since people of color and Blacks in particular are more likely to be in neighborhoods where surveillance is greater, they are more likely to have had encounters sooner (O’Neil, 2016). One groundbreaking case study found that the risk scores assigned to more than 7,000 people arrested in Broward County, Florida, were grossly inaccurate in predicting whether people would commit new violent crimes—only 20% of the people it predicted to commit crimes really went on to do it (Angwin et al., 2016). Furthermore, it was biased against Black people; when the effect of race was isolated, it accounted for Black defendants being 77% more likely to be pegged as having a higher risk of committing a future violent crime. Algorithms “leap from one field to the next” and “slam doors in the face of millions of people, often for the flimsiest of reasons, and offer no appeal” (O’Neil, 2016, p. 31). Other examples where algorithms quantify traits based on biased or unfair data include teacher assessment (e.g., student test scores are a stand-in for teacher effectiveness); hiring decisions (e.g., names that sound non-White are penalized by automated Human Resources (HR) screening or credit score is a factor in hiring); and college rankings (e.g., reputation is a huge factor which is subjective) (Marshall, 2018; O’Neil, 2016). College students need to learn about algorithms and the large role they play in governing our lives and arguably sometimes violating civil rights laws. For information literacy skills to be a means of personal empowerment for lifelong learning, it is important that the curriculum be expanded to include algorithmic literacy.

Statement of the Problem

College students’ research habits make them vulnerable to the misinformation and personalization often caused by algorithmic filtering in online information systems. For example,

students often receive incomplete information by clicking on and relying on the top search results even when they are not credible or relevant (Bhatt & MacKenzie, 2019; Head, 2012, 2016; Wineburg & McGrew, 2017). This reflects the priorities of algorithmic filters rather than the best available information, and it illustrates that students delegate a lot of authority to algorithmic filters. Bakke (2020) stated “information literacy instruction must not just address the credibility of individual sources, but also *how* those sources are found and selected . . . greater algorithmic literacy leads to more credible research” (p. 5). Also, a recent study suggested college students are vulnerable to fake news and disinformation: over two-thirds of undergraduate students failed to identify a “news” story as satirical, and 95% never located the Public Relations (PR) firm behind a supposedly “nonpartisan” website (Wineburg et al., 2020). Algorithmic literacy needs to be part of the mainstream university curriculum, but it is not currently being taught much in the classroom (Head et al., 2020). Little has been published yet on how to teach algorithmic literacy to college students since it is a relatively new field of study.

Information literacy instruction in higher education is correlated with academic achievement (Bowles-Terry, 2012) and critical thinking (Albitz, 2007), and most accrediting organizations include it as a core competency (Saunders, 2007). Students need greater awareness of how much trust they are placing in algorithms to select sources for them in both academic and personal research, because awareness could lead to more deliberate and reflective search habits over the prioritization of convenience. Even though algorithms can help limit information overload, they also limit user agency by making decisions about what information to display and filter out (Swart, 2021). There is a recent call in the literature to improve students’ algorithmic awareness and algorithmic literacy skills (Bakke, 2020; Bhatt & MacKenzie, 2019; Clark, 2018;

Cohen, 2018; Head et al., 2020; Hobbs, 2020; Lloyd, 2019; Nichols & Stornaiuolo, 2019; Ridley & Pawlick-Potts, 2021; Valtonen et al., 2019), and several researchers suggested integrating algorithmic literacy into existing information literacy instruction (Bakke, 2020; Clark, 2018; Head et al., 2020; Ridley & Pawlick-Potts, 2021). Both for the sake of their personal learning and as citizens, students need to understand how algorithm-driven information platforms like Google feed on their personal data to target ads and manipulate what content gets displayed in attempts to influence their behavior (Powers, 2017). It may be a challenge to communicate algorithmic concepts to learners without extensive backgrounds in math or computer science; Powers (2017) stated “it is unrealistic to expect young adults to learn what personalization is, how it works, and why it matters without specific exposure to the topic in school curricula . . . given the complexity of the subject matter, it is logical to address news personalization in college” (p. 1316). Bhatt and MacKenzie (2019) added:

If students are restricted in what they can know because they are unaware of how exogenous actors (algorithms) actually work, and how they guide their choices and shape their experiences online, then it becomes important to educate them to be critically aware during their digital searches for information, research and critical argument, and to educate them to be reflective about their ritualized practices with digital literacy. (p. 315)

New competencies surrounding algorithmic literacy would empower students to participate in the sharing of information in a responsible and ethical way and make informed decisions through the critical analysis of balanced information from all sides, not just the information that aligns with their preexisting beliefs (Head et al., 2020). Carmi and Yates (2020) argued that users need to understand the various factors influencing who and what reaches them in search results and social media feeds in order to make rational judgments about the information they encounter. Otherwise, they tend to believe the information they search and

encounter on different platforms is objective and neutral. Furthermore, a lack of this knowledge makes them more likely to unwittingly internalize search results without questioning and critiquing them (Carmi & Yates). Lloyd (2019) added that students need to question search results and automated decisions, and key concepts such as bias, trust, credibility, opacity, diversity, and social justice need to supplement the traditional lessons around “search” in information literacy pedagogy. Greater algorithmic literacy skills would shed light on how the underlying structure of the information systems used in wider society across a variety of sectors are being influenced by algorithms, creating the potential for unfairness and hidden bias. Students need to build greater awareness of the scope and reach of algorithms, and the potential harms that they can bring to individuals and the public sphere. This will allow them to develop a sense of personal agency and exercise greater influence over algorithmic decision-making.

History of Algorithms

The word “algorithm” was indirectly coined by ninth-century Persian mathematician Abdullah Muhammad bin Musa al-Khwarizmi; his name was rendered “Algorithmi” when translated into Latin (Bucher, 2018). His scripts described the basic methods of arithmetic for computer processors (Bucher). During the 17th century, the German philosopher Gottfried Wilhelm Leibniz provided groundwork for Boolean algebra and using if/then conditionals for calculating truth (Bucher). Algorithms come from the artificial intelligence (AI) branch of computer science. AI involves developing ways for computers to simulate human-like intelligent behavior, interpreting and absorbing new information for improved problem-solving, and recognizing patterns (Head et al., 2020). Early computers in the 1930s-1950s such as the Electronic Numerical Integrator and Computer (ENIAC) lay the foundation for artificial

intelligence, but AI didn't have commercial applications until the 1980s-1990s (Marshall, 2018). One famous example was the International Business Machines (IBM) Deep Blue supercomputer, which defeated the world chess champion in 1997 (Marshall).

Beginning in the early 2000s, AI began appearing in consumer products. We now live in the age of algorithms, and algorithms have the potential to shape our lives in expansive and significant ways (Head et al., 2020). According to a large group of technology experts, computers using machine learning might match or even exceed human intelligence and capabilities in areas like complex decision-making, reasoning and learning, sophisticated analytics and pattern recognition, visual acuity, speech recognition, and language translation (Anderson et al., 2018). Algorithms have the potential for social good, such as AI technology that optimizes traffic lights in real time by continuously adapting its behavior based on current traffic patterns; algorithms used to protect endangered wildlife species; algorithms used to combat climate change; and algorithms used to target HIV-education at homeless youths (Hager et al., 2019; Rolnick et al., 2019; Shi et al., 2020). However, a great tension exists between the potential innovation brought on by algorithms on one hand, and the serious privacy and civil rights violations that occur when algorithms reflect the racism that gets programmed into them on the other hand. As Meredith Broussard (2019) argued in *Artificial Unintelligence: How Computers Misunderstand the World*, social decisions are about more than just calculations, and it is dangerous to rely on data alone to make social and value judgements. The recent COVID-19 pandemic also changed established purchasing patterns, banking patterns, and health symptoms to such a large degree that many algorithms malfunctioned because they relied heavily on historical data patterns no longer true; examples included less accurate health diagnostics for

sepsis or cancer and fraud detection patterns incorrectly flagging low-income individuals and minorities (Parikh & Navathe, 2022).

Algorithmic Justice

An algorithm is usually modeled, trained on a body of data, and then adjusted as the results are examined. Since algorithms are always self-modifying as they gain experience and find patterns, they may eventually become inscrutable to their creators (Head et al., 2020). Despite being processed by computers, algorithms are not neutral or value free. They are influenced by decisions made by the humans who design them and the preexisting data on which they are trained (Head et al.). Algorithms used in machine learning solve problems by looking at the ways that people have already solved it in the past (their inputs). They look at statistical correlations, or what correlates with aggregate past behaviors, without looking at causality or meaning (O’Neil, 2016). The problem is that algorithms are largely hidden and lack transparency; Frank Pasquale (2015) used the metaphor of a black box to explain how the workings of algorithms are largely mysterious. Algorithms are often proprietary trade secrets, which van Dijck (2013) equated with patents or other types of intellectual property. Virginia Eubanks (2018) argued that the digital tracking now used on the poor results in a low rights environment with little political accountability or transparency. These collective concerns have inspired the algorithmic justice movement. Algorithmic justice is defined as “the application of principles of social justice and applied ethics to the design, deployment, regulation, and ongoing use of algorithmic systems so that the potential for harm is reduced,” along with the promotion of “awareness and sensitivity among coders and the general public about how data collection

practices, machine learning, AI, and algorithms may encode and exacerbate inequality and discrimination” (Head et al., 2020, p. 49).

There are a growing number of advocacy groups and institutes promoting algorithmic justice (see Table 1).

Table 1*Select Advocacy and Research Groups*

Organization	Website	About
Ada Lovelace Institute	https://www.adalovelaceinstitute.org/	Research institute with a mission to ensure data and AI work for people and society
A.I. For Anyone	https://www.aiforanyone.org/	Works to educate citizens about AI
AI Now Institute	https://ainowinstitute.org	Research center at New York University dedicated to understanding the social implications of AI
Algorithm & Data Literacy Project	https://algorithmliteracy.org/	Raises awareness and educates kids about the presence of algorithms and how they influence our digital experiences
American Civil Liberties Union (ACLU)	https://www.aclu.org	Advocates of privacy and technology to ensure that civil liberties are enhanced rather than compromised by technology innovations
Big Brother Watch (UK)	https://bigbrotherwatch.org.uk	Independent civil liberties group fighting to reclaim privacy and defend freedoms during technological change
Carceral Tech Resistance Network	https://www.carceral.tech/	A coalition of organizers who are campaigning against the design and experimentation of technologies by police, prisons, border enforcement, and commercial partnerships
Data & Society	https://datasociety.net	Nonprofit research group looking at the impact of AI and automation on labor, health, and online disinformation
Electronic Frontier Foundation	https://www.eff.org/	Nonprofit organization defending digital privacy and free speech
Library Freedom Project	https://libraryfreedom.org/	Teaches librarians about surveillance threats, privacy rights, and digital tools to thwart surveillance
Our Data Bodies	https://www.odbproject.org	Human rights and data justice organization
Surveillance Technology Oversight Project (S.T.O.P.)	https://www.stopspying.org/	Fights to end discriminatory surveillance, primarily in New York

A few are focused on the social impacts of artificial intelligence; this includes the AI Now Institute (n.d.) at New York University; the Surveillance Technology Oversight Project (Surveillance Technology Oversight Project [S.T.O.P.], 2022), also in New York; and the Ada

Lovelace Institute (n.d.). Several groups are focused on privacy and civil rights: the American Civil Liberties Union (ACLU) (2022) works to ensure that civil liberties are enhanced rather than compromised by technology innovations; in the UK the Big Brother Watch (UK) (2021) is an independent civil liberties group that fights to reclaim privacy and defend freedoms during technological change; the Electronic Frontier Foundation (n.d.) defends digital privacy and free speech; Our Data Bodies (n.d.) is a data justice organization; the Library Freedom Project (n.d.) teaches librarians about surveillance threats, privacy rights, and digital tools to thwart surveillance; and the Carceral Tech Resistance Network (2020) campaigns against surveillance technologies. Rounding out the list is The Algorithm & Data Literacy Project (Kids Code Jeunesse, 2022), a partnership between Kids Code Jeunesse (KCJ) and the Canadian Commission for the United Nations Educational, Scientific and Cultural Organization (CCUNESCO) that features videos, discussion questions, and other resources aimed at children; A.I. for Anyone (n.d.), working to educate citizens about AI; and Data & Society (n.d.), a nonprofit research group looking at the impact of AI and automation on labor, health, and online disinformation. This is only a small sampling of advocacy groups related to the algorithmic justice movement.

Purpose of the Study

The purpose of this exploratory, qualitative study was to do the following: 1) determine the key knowledge components for teaching algorithmic literacy to college students; and 2) understand the coping behaviors that would contribute to students' algorithmic literacy. Existing studies about college students and algorithmic literacy were scarce, but they suggested students were not fully prepared to navigate the rapidly changing information landscape increasingly

shaped by algorithms. Therefore, determining what the key knowledge components of algorithmic literacy would be, what student coping behaviors would help them interact with algorithmic systems more effectively, and any pedagogical considerations on how to teach algorithmic literacy would be a first step towards equipping educators with a way to integrate this important, emerging field of study into the mainstream curriculum.

Research Questions

The following research questions guided this exploratory, qualitative study:

1. What components of algorithmic literacy are specific to college students?
2. What behavior and knowledge contribute to students' algorithmic literacy?

Significance of the Study

This study was significant because algorithmic literacy is an emerging area of research; as Cotter and Reisdorf (2020) noted, “At present, there is no widely accepted operationalization of algorithmic knowledge” (Cotter & Reisdorf, 2020, p. 753). The study began to address the gap in the college curriculum for algorithmic literacy and algorithmic awareness by providing educators with a better sense of whether there was a growing consensus among subject-matter experts on what constitutes algorithmic literacy for college students. Furthermore, it addressed whether there is consensus on what behavior contribute to students' algorithmic literacy. Findings will empower educators to develop appropriate teaching strategies that would help students—as both consumers and creators of information—understand how algorithms interact with data behind the scenes to limit diversity and democracy. Hargittai and Micheli (2019) listed “awareness of how algorithms influence what people see” as one of 10 critical dimensions of Internet skills (p. 113). Hartman-Caverly and Chisholm (2020) noted that students are “the future

of these technologies . . . they will be both the consumers and adopters, as well as designers, programmers, marketers, investors, and administrators who make decisions on how and whether to execute these systems” (p. 13). As such, students need to develop greater awareness of the underlying impact of largely unregulated algorithms on both their personal data and within larger society. Students need to understand how seemingly invisible algorithms and automated decision-making have the power to influence and further deepen existing structural inequalities.

This study also suggested how to expand the current *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries [ACRL], 2016) to more explicitly include algorithmic literacy. Information literacy emerges during interaction with search systems, but modern system designs hide or render unworkable the contextual information needed for the judgment processes of information literacy (Smith & Matteson, 2018). Students cannot realistically avoid using products like Google and social media platforms and being subject to the algorithmic processes influencing their search results during their educational experiences. Therefore, they need to better understand how algorithms work and realize the potential for bias (Noble, 2018). When search engines unfairly promote certain sites to make a profit, this bias limits the diversity and democracy inherent to the information (Granka, 2010). The information systems we rely on are shaped by cultural hegemony and are used to “undermine our trust in news, politics, and each other” (Head et al., 2020, p. 6). The current *Framework for Information Literacy for Higher Education* (ACRL, 2016) did not explicitly include algorithmic literacy, but there are many opportunities to integrate it (Clark et al., 2017).

Conceptual Framework

My conceptual framework pulled from the *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries, 2016). The *Framework* is organized around six threshold concepts, or key ideas about the rapidly evolving information ecosystem, that are central to information literacy and result in enlarged understanding and new ways of thinking and practicing. Threshold concepts represent bottlenecks for novice researchers, the mastery of which amount to “a-ha” moments that allow them to better understand research in that discipline (Meyer et al., 2010). Each threshold concept also has a related set of knowledge practices (proficiencies or abilities) and dispositions (tendency to act or think in a particular way). The six frames are a cluster of interconnected core concepts, with flexible options for implementation, rather than a rigid list of skills. This paper expanded on four out of the six core concepts of the *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries, 2016) to more explicitly include algorithmic literacy. For more information on my conceptual framework, see Chapter 2.

Research Design and Methodology

This study was designed using a qualitative and exploratory approach. This was appropriate because the purpose of qualitative research is to “gain insights into a particular phenomenon of interest” (Mills & Gay, 2019, p. 7). The qualitative research methods chosen were semi-structured interviews and a focus group. Semi-structured interviews involved “a set of predetermined topics” but also allowed for flexibility to tailor the questions to each interviewee, which was appropriate since the sample was diverse (Crano et al., 2014, p. 287). Also, the dialogue was allowed to evolve naturally beyond the existing structured questions (Ahlin, 2019).

When researchers lack substantial scholarly knowledge on a subject, they often choose semi-structured interviews with subject matter experts as the research method (Ahlin, 2019). The second research method, focus groups, are group interviews where you try to collect a shared understanding of your research question by allowing participants to take turns answering questions, so they all give input (Mills & Gay, 2019). Focus groups are often used “in the formative stages in a new line of research or to obtain preliminary evidence to learn about an understudied topic” as well as “to establish the content validity of a construct” (Crano et al., 2014, p. 298). Focus groups were chosen as a second data collection method in order to establish data triangulation and enhance the trustworthiness of the findings. Purposive sampling was done to recruit participants identified as scholars and teachers in the field of algorithm studies from the literature review. Potential participants were invited to participate through email by the researcher, and all interviews and the focus group were recorded using the Zoom videoconferencing software (Version 5.8.3, <https://zoom.us/>). Participants were asked to fill out an Informed Consent Form, and the researcher followed the developed interview protocol and focus group protocol. All portions of this study were submitted to Loyola Marymount University’s (LMU) Institutional Review Board (IRB) for approval prior to any data collection.

Limitations and Delimitations

Ideally, qualitative research focuses on participants in their natural field setting (Creswell, 2009), but this wasn’t possible due to the COVID-19 pandemic as well as the wide variety in participants’ geographical locations. The study took place through the Zoom video conferencing platform, where body language was harder to observe (Morgan, 2019). Not all participants in the study were equally articulate (Creswell, 2009), perspectives of marginalized

participants were lacking, and participants in the focus group may also have felt pressure to conform with the majority opinion (Crano et al., 2014). The researcher's presence serving as focus group moderator and interviewer may have biased the results (Crano et al., 2014; Creswell, 2009). There was a sampling bias because sampling was purposive and not random, and due to the small sample size for both the interviews and focus group the results were not necessarily transferable to other situations.

Definitions of Key Terms

- **Algorithm:** The set of logical rules used to organize and act on a body of data to solve a problem or accomplish a goal that is usually carried out by a machine (Head et al., 2020, p. 49).
- **Algorithmic audit:** Assessments of the algorithm's negative impact on the rights and interests of stakeholders (Brown et al., 2021, p. 2).
- **Algorithmic awareness:** The degree to which users know what algorithms can be used for and in what online contexts algorithms are actually used (Dogruel et al., 2021, p. 5).
- **Algorithmic fairness:** Addresses the idea that algorithmic decisions should not produce biased, discriminatory, or unfair results (Shin et al., 2020).
- **Algorithmic justice:** The application of principles of social justice and applied ethics to the design, deployment, regulation, and ongoing use of algorithmic systems so that the potential for harm is reduced; along with the promotion of awareness and sensitivity among coders and the general public about how data collection practices, machine learning, AI, and algorithms may encode and exacerbate inequality and discrimination (Head et al., 2020, p. 49).

- Algorithmic literacy:** A subset of information literacy that entails a critical awareness of what algorithms are, how they interact with human behavioral data in information systems, and an understanding of the social and ethical issues related to their use (Head et al., 2020, p. 49); also defined as “the capacity and opportunity to be aware of both the presence and impact of algorithmically-driven systems on self- or collaboratively-identified goals, and the capacity and opportunity to crystalize this understanding into a strategic use of these systems to accomplish said goals” (DeVito, 2021, p. 339:3); also defined as a set of capabilities used to organize and apply algorithmic curation, control, and active practices relevant when managing one’s AI environment along with an understanding of how algorithms reconstruct realities and are expressions of broader systems of power (Shin et al., 2021); also defined as “*being aware* of the use of algorithms in online applications, platforms, and services, *knowing* how algorithms work, being able to *critically evaluate* algorithmic decision-making as well as having the skills to *cope with* or even *influence* algorithmic operations” (Dogruel et al., 2021, p. 4); also critical algorithmic literacy defined as recognizing knowledge as situated, constructed within and in relation to the discursive landscape of social worlds and involving the cultivation of a critical consciousness through recognizing and responding to algorithms as expressions of broader systems of power (Cotter, 2019).
- Artificial intelligence (AI):** A branch of computer science that develops ways for computers to simulate human-like intelligent behavior, able to interpret and absorb new information for improved problem-solving, and recognize patterns (e.g., speech recognition or facial recognition) (Head et al., 2020, p. 49).

- **Artificial intelligence literacy (AI literacy):** a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace (Long & Magerko, 2020, p. 2).
- **Attention economy:** Companies (both platforms and people who use the platforms to sell, entertain, or persuade) try to engage and keep people's (limited) attention; this rewards clickbait and influences the design of algorithms and platforms to maximize time spent online (Head et al., 2020, p. 49).
- **Bias:** The inclination or prejudice for or against one person or group, especially in a way that is considered to be unfair (Bobkowski & Younger, 2018).
- **Big data:** A set of technological capabilities developed in recent years which, when used in combination, allows for the continuous gathering and processing of large volumes of fine-grained and exhaustive data drawn from multiple sources to be combined and analyzed continuously (Head et al., 2020, p. 49).
- **Black box:** A system that can be viewed only through its inputs and outputs, not its internal process (7th Media Empire, 2020, p. 10).
- **Code:** The technical language used to write algorithms and other computer programs (7th Media Empire, p. 10).
- **Data exhaust:** Information accidentally generated as people use computers, carry cell phones, or have their behavior captured through surveillance which becomes valuable when acquired, combined, and analyzed in great detail at high velocity (Head et al., 2020, p. 49).

- **Data literacy:** The component of information literacy that enables individuals to access, interpret, critically assess, manage, handle and ethically use data (Prado & Marzal, 2013, p. 126).
- **Digital inequality:** The ways that offline socioeconomic inequalities related to demographic categories, like educational attainment, income, ethnicity, and age, are reproduced in online settings (Warshaw et al., 2016, p. 273).
- **Digital redlining:** The practice of creating and perpetuating inequities between already marginalized groups specifically through the use of digital technologies (“Digital redlining,” 2020).
- **Ethics washing:** The practice of fabricating or exaggerating a company’s interest in equitable AI systems that work for everyone (Johnson, 2019).
- **Fake News:** A piece of information that includes content that is not fact-based but presented by the content creator as if it were true (Wang & Fussell, 2020).
- **Filter bubble:** When people in an online environment are exposed only to opinions and information that conform to their existing beliefs (Pariser, 2011).
- **Folk theories:** Intuitive, informal theories that individuals develop to explain the outcomes, effects, or consequences of technological systems, which guide reactions to and behavior towards said systems (DeVito et al., 2017, p. 3165).
- **Inference literacy:** A subset of digital literacy, this describes the beliefs and misconceptions people have about how companies collect and make inferences from their data (Warshaw et al., 2016, p. 271).

- **Information literacy:** The set of integrated abilities encompassing the reflective discovery of information, the understanding of how information is produced and valued, and the use of information in creating new knowledge and participating ethically in communities of learning (Association of College and Research Libraries, 2016, p. 7).
- **Machine learning:** The use of algorithms, data sets, and statistical modeling to build models that can recognize patterns to make predictions and interpret new data (enabling computers to learn from data with little human intervention) (Head et al., 2020, p. 49-50).
- **Metaliteracy:** A unified construct that supports the acquisition, production, and sharing of knowledge in collaborative online communities and challenges traditional skills-based approaches to information literacy by recognizing related literacy types and incorporating emerging technologies (Mackey & Jacobson, 2011).
- **Modeling:** Building a simplified or optimized reality of the world, often using data and a process that predicts, ranks, associates, or classifies (Diakopoulos & Kolisha, 2017, p. 818).
- **Personalization:** Process (used in targeted digital advertising) of displaying search results or modifying the behavior of an online platform to match an individual's expressed or presumed preferences, established through creating digital profiles and using that data to predict whether and how an individual will act on algorithmically selected information (Head et al., 2020, p. 50).
- **Platform:** Software deployed online (e.g., Google, YouTube, Instagram, Facebook) to provide a service, such as web search, video sharing, shopping, or social interaction. Often these systems use proprietary algorithms to mediate the flow of information while

enabling third parties to develop apps, advertising, and content, thus becoming digital spaces for the individual performance of identity online, data driven persuasion, and group formation through social interaction (Head et al., 2020, p. 50).

- **Post-truth:** Pertaining to an era where factual accuracy is no longer significant or relevant (“Post-truth,” 2021).
- **Pseudo-code:** A plain language description of the steps in an algorithm (Clarke & Kaptanian, 2018).
- **Search engine optimization (SEO):** The process of using a range of techniques, including augmenting HTML (hypertext markup language) code, web page copy editing, site navigation, or linking campaigns, in order to improve how well a site or page gets listed in search engines for particular search topics (Noble, 2018, p. 47).
- **Surveillance capitalism:** A new economic order that claims human experience as free raw material for hidden commercial practices of extraction, prediction, and sales; a parasitic economic logic in which the production of goods and services is subordinated to a new global architecture of behavioral modifications (Zuboff, 2019, preface).
- **Technochauvinism:** Bias that says the technological solution is always better than every other solution; the belief that technology is always the solution (Broussard, 2019).
- **Truth decay:** A set of four related trends: increasing disagreement about facts and analytical interpretations of facts and data; a blurring of the line between opinion and fact; an increase in the relative volume, and resulting influence, of opinion and personal experience over fact; and declining trust in formerly respected sources of factual information (Kavanagh & Rich, 2018).

- **XAI (explainable AI):** Enables human users to understand, appropriately trust, and effectively manage the emerging generation of artificially intelligent partners and to deploy AI systems that have the ability to explain their rationale, characterize their strengths and weaknesses, and convey an understanding of how they will behave in the future (Ridley & Pawlick-Potts, 2021, p. 8).

Organization of Dissertation

Chapter 1 provided an overview of the research study on algorithmic literacy. Background information was provided for algorithms in general, and the need for algorithmic literacy in college students more specifically. The purpose of the research and its significance were described. Chapter 2 will provide a review of literature related to the research focus, and Chapter 3 will outline the study's methodology. Chapter 4 will present the key themes that emerged from the interviews and focus group. Lastly, Chapter 5 will discuss the findings and recommendations.

CHAPTER 2

LITERATURE REVIEW

Introduction

This study aimed to explore algorithmic literacy for college students. Algorithmic literacy as a subset of information literacy was the conceptual frame explored in this study, and as such, this literature review began by tracing the history of the term “information literacy” (IL) within the context of higher education. It explained evolving IL definitions and IL standards from the Association of College and Research Libraries (ACRL), as well as reflected on the fact that information literacy instruction needs to reflect recent changes in the information landscape. It documented the few instances where algorithmic literacy instruction was taught to students either as part of information literacy instruction or in other fields. Next, the literature review explored the technical and social aspects of algorithms, with a special focus on search engines, social media, algorithmic power, algorithmic bias, filter bubbles, and fake news. Finally, the review ended with user studies on algorithms, including how algorithmic knowledge, attitudes, perceptions, and behaviors were measured.

Information Literacy

Librarians have a long history of teaching information literacy and viewing information literacy as a social justice issue because it empowers people from all walks of life to seek the truth. Being able to understand, evaluate, and use information is a human right, and it is essential for full participation in a democratic society (Zurkowski, 1974). Paul Zurkowski (1974) coined the term “information literacy” (IL) in 1974, and he described it as the ability to find what is known on any subject using the right tools and techniques. He felt information literacy should

extend to all segments of society. In 1989, the American Library Association (ALA) defined information literacy as the ability to locate, evaluate, and use information effectively. They urged schools and colleges to integrate information literacy concepts into their curriculum because these skills would allow for lifelong learning and help balance social and economic inequalities. Furthermore, information literacy was a means of “personal empowerment” because it would allow people to “verify or refute expert opinion and to become independent seekers of truth” (American Library Association, 1989, para. 6).

Within higher education, the most influential definition of information literacy came from ALA’s Association of College and Research Libraries (ACRL). In 2000, ACRL published the *Information Literacy Competency Standards for Higher Education* (American Library Association, Association of College & Research Libraries, 2000). The document offered five standards for assessing competence in information literacy, and each standard had corresponding performance indicators with more details (22 indicators in total). Each performance indicator, in turn, had possible outcomes to make it easier to assess (87 outcomes in total). The five standards were as follows:

- the information literate student determines the nature and extent of the information need;
- the information literate student accesses needed information effectively and efficiently;
- the information literate student evaluates information and its sources critically and incorporates selected knowledge into his or her knowledge base and value system;

- the information literate student uses information effectively to accomplish a specific purpose; and
- the information literature student understands the economic, legal, and social implications of information.

The standards included both higher and lower-level thinking skills, and they were useful to guide the practice of information literacy instruction.

In 2005, the Alexandria Proclamation declared that information literacy “lies at the core of lifelong learning” because it empowers people to use information effectively “to achieve their personal, social, occupational and educational goals” (International Federation of Library Associations and Institutions, 2005, para. 2). The proclamation went on to say that IL was a basic human right and promotes social inclusion of all nations (para. 2). Within higher education, information literacy instruction is correlated with academic achievement (Bowles-Terry, 2012). It has also been linked to critical thinking (Albitz, 2007). The National Leadership Council for Liberal Education and America’s Promise (LEAP) (2007) included information literacy as one of the essential learning outcomes for a liberal education. Also, five out of the six regional higher education accrediting organizations—the Middle States Commission on Higher Education, New England Association of Schools and Colleges, Southern Association of Colleges and Schools, Northwest Commission on Colleges and Universities, and Western Association of Schools and Colleges—included information literacy in their standards to a significant degree (Saunders, 2007). Megan Oakleaf (2011) noted that many students were overconfident in their information literacy skills, despite evidence to the contrary.

Critical Information Literacy Movement

There is a growing “critical information literacy” movement that focuses on how students can be encouraged to advocate for social change by identifying and acting upon oppressive power structures in information (Tewell, 2015). The term “critical information literacy” emerged just a few years after ACRL introduced the term information literacy (Luke & Kapitzke, 1999). It was inspired by the critical pedagogy framework of educators like Paulo Freire, who advocated for students to develop a “critical consciousness” and act on their findings of social and political inequalities. Even before the rise of search engines and algorithms, library search tools such as a library catalogs using Library of Congress classification had a long history of being biased and classifying minority groups as problem people (Noble, 2018). For example, Asian Americans were once classified as Yellow Peril, Jewish people were once classified as the Jewish question, African Americans were once classified as Negroes, and the term “illegal aliens” is still being used (Noble). Harris (2010) explained that the main objective in critical information literacy was not predetermined learning outcomes, as outlined in the ACRL competency standards, but rather was to foster the development of a critical consciousness about information and information sources in students. This experience would vary widely depending on the learner and his community and therefore outcomes could not be predetermined. Harris listed guided reflection and dialogue as two strategies that would facilitate this. Swanson (2010) further added that some students buy into Freire’s banking concept and view using information as a rote skill. Therefore, in these cases information systems were “structures that tend to reinforce the societal status quo by supporting the values behind it and limiting access to media outlets and mass distribution of information” (p. 265). Elmborg (2006) also felt that critical

information literacy allowed students to take control of their learning—they could engage with significant problems in the world that matter to them rather than merely acquiring new knowledge content. They could address their own needs and questions independently later in life. Gregory and Higgins (2013) more recently posited that critical information literacy differed from standard definitions of information literacy because it “takes into consideration the social, political, economic, and corporate systems that have power and influence over information production, dissemination, access, and consumption” (p. 4).

ACRL Framework

The ACRL standards were rescinded in 2016 by ACRL and replaced with the *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries [ACRL], 2016). This was partially, but not entirely, a result of criticism from critical information literacy proponents. The *Framework* was developed by the ACRL Information Literacy Competency Standards for Higher Education Task Force starting in 2012 and spanning several years as they circulated three rounds of drafts for public comment and gathered feedback through public meetings and one online questionnaire (ACRL, 2016). The *Framework* stated, “information literacy is the set of integrated abilities encompassing the reflective discovery of information, the understanding of how information is produced and valued, and the use of information in creating new knowledge and participating ethically in communities of learning” (p. 7). As discussed in chapter 1, the *Framework* is organized into six frames, each consisting of a concept central to information literacy, a set of knowledge practices, and a set of dispositions. The *Framework* is grounded in the idea of “threshold concepts,” which originated with economists Meyer and Land (2003), who wanted to identify especially challenging concepts for

students in their field. Threshold concepts have also been applied to other disciplines as a way of thinking about the recurring difficulties experienced by students who are novices to that discipline and must pass through a portal of struggle (“the threshold”) in order to develop an understanding of the discipline (Gibson & Jacobson, 2014). Often, threshold concepts embody the way that experts in the field think, but the concepts are not explicitly taught as part of the curriculum. Threshold concepts are transformative (permanently change how a student thinks); irreversible (not likely to be forgotten); integrative (reveal how previously unrelated knowledge is related); bounded (set disciplinary boundaries); and potentially troublesome (represent a new way of understanding) (Meyer & Land, 2006). Threshold concepts were first applied to information science in 2011 (Townsend et al., 2011), and the Framework Task Force were further inspired to apply it to the new *ACRL Framework*.

The *Framework* (ACRL, 2016) was and still is controversial since it is very difficult to assess due to its broadness (Oakleaf, 2014). What is “troublesome” depends on the learner, so it cannot be standardized (Saracevic, 2014). Also, Reed (2015) argued that the frames were too theoretical and difficult to teach in a community college environment. Fulkerson et al. (2017) argued that there was a lack of metaliteracy (expanded list of abilities needed by an information literate person) and metacognition (connecting the affective, cognitive, and behavioral learning domains through critical self-reflection) in the document. The *Framework* (ACRL, 2016) was heavily influenced by a Delphi research study that engaged expert practitioners on the topic of threshold concepts (Townsend et al., 2016). This method is a qualitative approach in which a small group of experts anonymously answer questions in writing through several rounds in order to generate consensus. The panel of experts consisted of 19 participants across four rounds, and

they were chosen based on their knowledge of and active participation in the field of information literacy and library instruction, as operationalized through publication, teaching, or leadership in professional organizations. The results of the study showed agreement that the threshold concept approach was useful for information literacy instruction, and over 50 potential threshold concepts were proposed but this got reduced to six final concepts. The methodology used to inform the *ACRL Framework* (ACRL, 2016) resulted in some criticism because it relied on only a small number of experts, thus leaving out important viewpoints (Battista et al., 2015).

Cunningham and Dunaway (2017) analyzed the *Framework* (ACRL, 2016) dispositions and, using exploratory factor analysis, concluded that all of the dispositions could really be condensed into these four: mindful self-reflection, productive persistence, responsibility to community, and toleration of ambiguity. Critics of the old ACRL standards commended the *Framework* for addressing the sociocultural complexities of information. However, some felt that it still did not go far enough in connecting information literacy to social justice. Specifically, critics highlighted that it did not go far enough in challenging the status quo and encouraging action; it failed to recognize the importance of culture (and privileging the culture of academic research); it failed to look at the mechanisms that established the power and authority structures within the scholarly conversation; and it lacked civic engagement in the document, which implied an acceptance of the existing structure of information and knowledge (Battista et al., 2015). One extreme view was that the *Framework* (ACRL, 2016), by minimizing its social justice stance, accepted the status quo of neoliberal education and the perpetuation of existing race, class, and gender structures (Saunders, 2017). However, more recently, Heffernan (2020) argued that the *Framework* (ACRL, 2016) could be viewed through a Diversity, Equity and

Inclusion (DEI) lens because it prompts us to ask questions and pursue answers from unheard or systemically silenced voices. Furthermore, when configured as a guide for question-asking rather than static learning outcomes, it encouraged us to challenge elitism, racism, sexism, ableism, and biases within the entire information ecosystem and thus allowed us to expand information literacy to include algorithmic literacy (as well as other literacies like racial literacy and digital literacy).

Two assessment instruments for the *Framework* (ACRL, 2016) were published. The first was *The Threshold Achievement Test for Information Literacy* (TATIL) (Carrick Enterprises, 2018), a fixed-choice test that measured knowledge and dispositions based on the *Framework*. There are currently four modules available for the constructs of evaluating process and authority; strategic searching; research and scholarship; and the value of information. In 2019, the first valid and reliable scale for measuring ACRL framework-based knowledge practices and dispositions for seven constructs of information literacy was published (Doyle et al., 2019). The 36-item scale was validated on graduate students and measured how students perceived their own information literacy knowledge practices and dispositions. Students rated their own skill level on a seven-point Likert scale: Novice (1), Advanced (2), Emerging (3), Advanced Emerging (4), Developing (5), Advanced Developing (6), and Expert (7). The scale measured all six of the ACRL *Framework* (ACRL, 2016) threshold concepts (each had its own subscale), with the “search as strategic exploration” frame broken up into two different sections, one for tools and tasks, and another for exploration-mindset. There were four scale questions that could be applied to algorithmic literacy. Under the “information has value” frame, the questions were “I understand that my personal information has value online, and make informed choices to manage

my preferences for how this information is used” and “I feel comfortable as an active creator in the information economy, rather than as a passive consumer.” Under the “searching as strategic exploration” frame, the questions that could apply to search algorithms were “I recognize the ways in which search tools organize information” and “I remain persistent when faced with a challenging search.”

Related Frameworks

As discussed previously, algorithmic literacy is a subset of information literacy (Head et al., 2020). It also incorporates elements of data literacy, defined by Prado and Marzal (2013) as “the component of information literacy that enables individuals to access, interpret, critically assess, manage, handle and ethically use data” (p. 126). Other literacies that are also closely related include media literacy, new media literacy, computational thinking, platform literacy, digital literacy, and artificial intelligence literacy (DeVito, 2021; Dogruel, 2021b; Head et al., 2020; Long & Magerko, 2020; Ridley & Pawlick-Potts, 2021). Dogruel (2021b) proposed four dimensions for algorithm literacy for lay internet users based on existing concepts in media literacy research. The first dimension was factual awareness or knowledge consisting of defining algorithms, being aware of areas of application, knowing how they work and what the effects are, and regulatory/legal measures. The second dimension was critical evaluation consisting of the opportunities and risks, individual and societal effects, and privacy implications. The third dimension was coping behaviors consisting of privacy-related measures, result-related measures (e.g., consult different search engines), and critical communication and activism. The fourth dimension was creation design consisting of programming skills, coding skills, and deliberate manipulation.

There are no widely accepted competencies or standards for algorithmic literacy yet for the K-12 or higher education curriculum, but Long and Magerko (2020) published a conceptual framework for “AI literacy,” defined as “a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace” (p. 2). The “AI literacy” framework proposed five overarching themes (along with 17 competencies and 15 design considerations): 1) What is AI? 2) What can AI do? 3) How does AI work? 4) How should AI be used? and 5) How do people perceive AI (Long & Magerko, 2020)? Zhou et al. (2020) evaluated 49 K-12 AI education works based on this framework and developed a checklist to show which core competencies were applied in each learning activity. Van Brummelen et al. (2021) developed a series of student workshops based on eight of these competencies and evaluated them by coding student responses to questions such as “describe in your own words what AI is” and “describe the differences between rule-based AI and machine learning based AI.” Also, the “AI for K12” working group published a draft set of standards for K-12 classrooms around AI literacy encompassing five big ideas: 1) Perception (computers perceive the world using sensors); 2) Representation & Reasoning (computers maintain representation of the world and use them for reasoning); 3) Learning (computers can learn from data); 4) Natural interaction (intelligent agents require many types of knowledge to interact naturally with humans); and 5) Societal Impact (artificial intelligence can impact society in both positive and negative ways) (Touretzky et al., 2019). Kim et al. (2021) proposed a detailed AI curriculum to achieve the competencies of AI knowledge, AI skill, and AI attitude for elementary school. The International Society for Technology in Education (ISTE) currently has “ISTE Standards for Students,” and there is a

digital citizen component that includes standard 1.2.d, “students manage their personal data to maintain digital privacy and security and are aware of data-collection technology used to track their navigation online,” as well as standard 1.2.a, “students cultivate and manage their digital identity and reputation and are aware of the permanence of their actions in the digital world.” (International Society for Technology in Education [ISTE], 2016).

Lao (2020) proposed a new machine learning education framework rubric for college students that included knowledge, skills, and attitudes. Similarly, Sulmont (2019) proposed a taxonomy for machine learning goals for non-majors. Druga et al. (2021) provided a family AI literacy framework called “the 4As” consisting of ask (interact fluently with AI application to identify AI bias), adapt (modify or customize the application by trying to trick it), author (create a new AI application), and analyze (analyze the architecture and modify). Dezuanni (2021) argued that “it is important for teachers to develop new ways to introduce critical thinking about algorithms into the classroom” as part of digital media literacy education in Australia (p. 882). Finally, algorithms were one of the “big ideas of computer science” listed in the *CS Principles Project*, which outlined a curriculum framework for a college-level Advanced Placement course in computer science (Astrachan & Briggs, 2012).

Conceptual Framework

This study used the *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries, 2016) as its conceptual framework. The six core “threshold concepts” that anchored the frames are:

- **Frame 1:** Authority Is Constructed and Contextual (authority depends on information need and context; different communities recognize different authority types; authority

is not infallible, and users should remain skeptical of the systems that have elevated that authority and the information created by it);

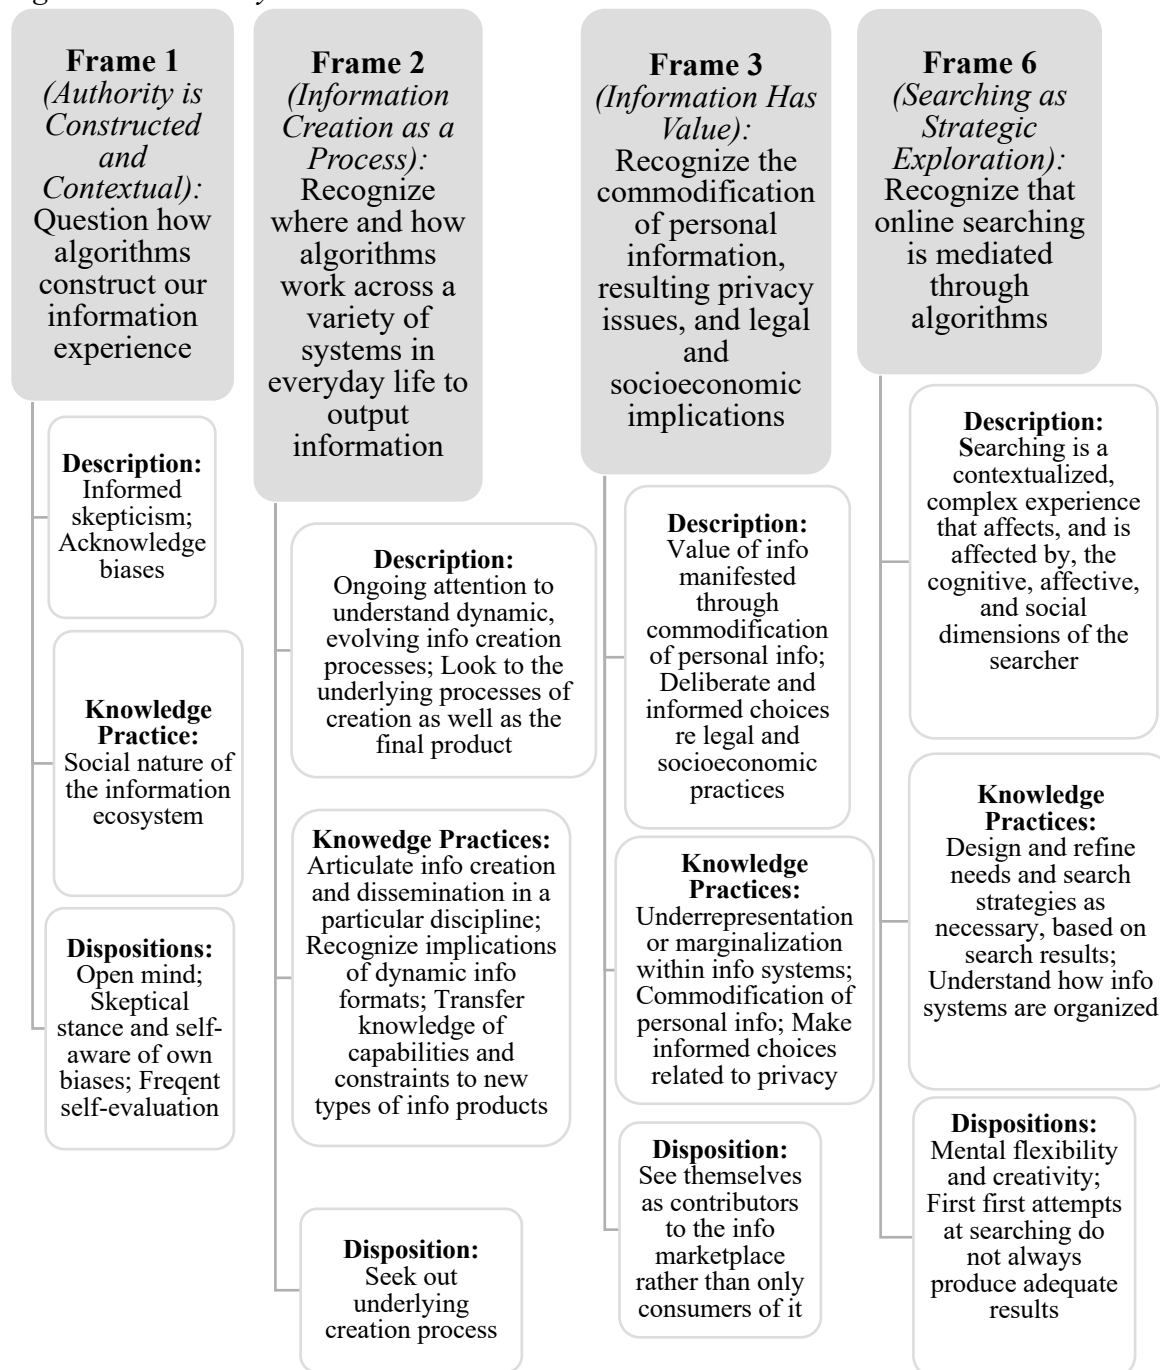
- **Frame 2:** Information Creation as a Process (since the online environment obscures the format of information, users need to recognize the underlying process that went into producing the information and sharing the information);
- **Frame 3:** Information Has Value (the value of information is manifested in various contexts, including publishing practices, access to information, the commodification of personal information, and intellectual property laws; information users have responsibilities as consumers and creators; information production is impacted by political, social, legal, and economic issues);
- **Frame 4:** Research as Inquiry (research is iterative; users gain knowledge through exploring ideas, which always leads to more questions);
- **Frame 5:** Scholarship as Conversation (scholarship is a discursive practice where users and creators come together to negotiate meaning; knowledge is always shifting and it does not stand still);
- **Frame 6:** Searching as Strategic Exploration (locating information is difficult and requires diverse search strategies and a revision of those search strategies) (ACRL, 2016).

My conceptual framework focused on expanding Frames 1-3 and Frame 6 to explicitly include algorithmic literacy. Frame 4, Research as Inquiry, and Frame 5, Scholarship as Conversation, were not included because they were not related to the construct of algorithmic literacy. Figure 1

maps out the expansion of these four frames to include algorithmic literacy, which would strengthen the framework's overall integration of social justice and metaliteracy.

Figure 1

Algorithmic Literacy Across Frames 1-3 and Frame 6



Note: Adapted from “Framework for Information Literacy for Higher Education,” by Association of College & Research Libraries, 2016 (<http://www.ala.org/acrl/standards/ilframework>), copyright 2016 by the Association of College & Research Libraries.

Figure 1 outlines that Frame 1, Authority Is Constructed and Contextual (ACRL, 2016), could be expanded to include questioning how algorithms construct our information experience. The existing description “experts view authority with an attitude of informed skepticism and an openness to new perspectives, additional voices, and changes in schools of thought” (ACRL, 2016, p. 12) could include questioning the authority delegated to algorithms, while the existing description “experts understand the need to determine the validity of the information created by different authorities and to acknowledge biases that privilege some sources of authority over others, especially in terms of others’ worldviews, gender, sexual orientation, and cultural orientation” could include algorithm bias and structural forces that exclude certain groups of people. The existing knowledge practice “understand the increasingly social nature of the information ecosystem where authorities actively connect with one another and sources develop over time” could be applied to the collective influence of all internet users on algorithms (for example, popularity ranking). The existing dispositions “develop and maintain an open mind when encountering varied and sometimes conflicting perspectives,” “develop awareness of the importance of assessing content with a skeptical stance and with a self-awareness of their own biases and worldview,” and “are conscious that maintaining these attitudes and actions requires frequent self-evaluation” could all be applied to looking for algorithmic decisions that sometimes produce bias reflecting their creators’ priorities. One consequence of this is filter bubbles, or when people in an online environment are exposed only to opinions and information that conform to their existing beliefs (Pariser, 2011). Another consequence is fake news, disinformation, and clickbait often perpetuated by algorithms.

Frame 2, Information Creation as a Process (ACRL, 2016), could be expanded to include recognizing where and how algorithms work across a variety of systems in everyday life to output information. The existing description “the dynamic nature of information creation and dissemination requires ongoing attention to understand evolving creation processes” (ACRL, 2016, p. 14) could be applied to understanding that algorithms are dynamic and unstable, while the description “recognizing the nature of information creation, experts look to the underlying processes of creation as well as the final product to critically evaluate the usefulness of the information” could be applied to understanding the wider infrastructure that algorithms are embedded in. The existing knowledge practices “articulate the traditional and emerging processes of information creation and dissemination in a particular discipline” and “transfer knowledge of capabilities and constraints to new types of information products” could be applied to expanding knowledge about algorithmic systems beyond search and social media to other domains in everyday life, while the knowledge practice “recognize the implications of information formats that contain static or dynamic information” could reinforce that algorithms are constantly changing. The existing disposition “are inclined to seek out characteristics of information products that indicate the underlying creation process” could apply to understanding how algorithms generally work and that they are dependent on the data they are processing.

Frame 3, Information Has Value (ACRL, 2016), could be broadened to include recognizing the commodification of personal information, resulting privacy issues, and legal and socioeconomic implications. The existing description “the value of information is manifested in various contexts, including . . . the commodification of personal information” (ACRL, 2016, p. 16) could apply to knowing that the revenue of online platforms using algorithmic systems

comes from advertising. The description “experts also understand that the individual is responsible for making deliberate and informed choices about when to comply with and when to contest current legal and socioeconomic practices concerning the value of information” could be applied to how algorithms interact with human behavioral data to engender privacy issues. The existing knowledge practice “understand how and why some individuals or groups of individuals may be underrepresented or systematically marginalized within the systems that produce and disseminate information” could be applied to the legal aspects of algorithmic harms, and the knowledge practice “make informed choices regarding their online actions in full awareness of issues related to privacy and the commodification of personal information” could be applied to big data collection practices. The existing disposition “see themselves as contributors to the information marketplace rather than only consumers of it” could refer to the feedback loop between algorithms and humans.

Frame 6, Searching as Strategic Exploration (ACRL, 2016), could be extended to include recognizing that online searching is mediated through algorithms. The existing description “experts realize that information searching is a contextualized, complex experience that affects, and is affected by, the cognitive, affective, and social dimensions of the searcher” could be expanded to explain that search results are personalized through not only search queries but also digital identity and the actions of others, which results in an endless feedback loop. The existing knowledge practice “design and refine needs and search strategies as necessary, based on search results” could refer to manipulating your search tactics to improve algorithmic outputs, while “understand how information systems are organized in order to access relevant information” could refer to understanding key variables impacting your search results. The existing

dispositions “exhibit mental flexibility and creativity” and “understand that first attempts at searching do not always produce adequate results” could refer to maneuvering your actions to improve search results.

Recent Studies on Information Literacy

Some of the biggest studies examining the state of information literacy in higher education came from Project Information Literacy (PIL) (Project Information Literacy [PIL], 2022), a nonprofit research institute that conducts ongoing, national studies on what it is like being a student in the digital age. In one PIL study, they found that college students relied largely on Google, Wikipedia, and course readings for their information (Head & Eisenberg, 2009). This approach reflected conceptualizing information-seeking as something learned by repetition and done with efficiency and utility rather than as an opportunity for deep learning (Head & Eisenberg). Along similar lines, another study found that first-year students believed freely available Internet resources were sufficient for academic work, they considered Google to be a sufficient search tool, and they assumed that accessibility was an indicator of quality (Hinchliffe et al., 2018). Two more studies found that undergraduate students clicked on higher-ranked results even if less credible or relevant (Schultheiß et al., 2018; Wineburg & McGrew, 2017).

Another PIL study looked at how students engaged with news—68% of the college students found news overwhelming, and only 14% felt very confident that they could tell fake news apart from real news (Head et al., 2018). More than 80% of the students participating in a 2016 Stanford History Education Group Study failed to question the legitimacy of a photograph of wilting flowers accompanied by the caption “Fukushima Nuclear Flowers” (Wineburg et al., 2016). The photo was posted on a photo sharing website and was framed to imply the flowers

had birth defects caused by the Fukushima Daiichi nuclear disaster, but in reality wilting is a common mutation in daisies (Wineburg et al., 2016). A later study by the same research group surveyed undergraduate students and found two-thirds of them failed to identify a “news story” that came from a satirical website; Ninety-five percent of them never located the PR firm behind a supposedly “nonpartisan” website; and many attributed undue weight to easily manipulated signals of credibility, such as an organization’s non-profit status, its links to authoritative sources, or its appearance (Wineburg et al., 2020). Barshaba et al. (2020) found that college students overestimated the self-reported extent they used independent, external sources to evaluate the credibility of online information when compared to their observed behavior. Along similar lines, McGrew et al. (2018) measured students’ civic online reasoning—the ability to effectively search for, evaluate, and verify social and political information online—and found that they struggled. Haider and Sundin (2020) found that late teens in Sweden fell under four categories of evaluators: the non-evaluator (both low trust and low agency, leading to inaction); the naïve evaluator (high trust and low agency, leading to blind acceptance), the skeptical evaluator (low trust and high agency, leading to suspicion of everything including the knowledge production system itself), and the confident evaluator (high trust and high agency, leading to mistrust and caution but giving the person agency). The confident evaluator was most helpful for information evaluation in the age of algorithms because it moved away from a binary “checklist” approach and offered flexibility and adaptiveness depending on the context (Haider & Sundin, 2020). Awareness did not always lead to avoidance of the activity; for example, even when participants had awareness of the role of the technical infrastructure (e.g., YouTube or Facebook) in things like the spreading of conspiracy theories, this did not necessarily stop them from

participating or getting stuck in the cycle. However, at least the awareness played a role in evaluation, and teens acknowledged a shared responsibility as they purposefully interacted with the information intermediary (Haider & Sundin, 2020).

Beginning Steps: Algorithms Embedded in Information Literacy Instruction

In a syllabi content analysis, Fiesler et al. (2020) analyzed 115 university technology ethics courses in the United States and found the majority of home departments were in computer science, followed by information science, philosophy, and communication. Fifty-five of the courses explored the topic of “AI & algorithms” and covered concepts like black box algorithms, machine learning, algorithmic fairness, bias, profiling, and the need for transparency despite the challenge of auditing algorithms. However, there are only a small number of published instances of algorithms being embedded into the information literacy curriculum. The INFO 101 8-week credit-bearing information literacy course at Pierce College was designed to align with the six frames from the ACRL *Framework for Information Literacy for Higher Education* (ACRL, 2016; Fisher, 2017). In the module covering Frame 6, “Searching as Strategic Exploration,” students read a newspaper article from *The Washington Post* (Guarino, 2016) about Google’s algorithm bias on display when a Google Images search for “three White teenagers” yielded happy Caucasians compared to a search for “three Black teenagers” which yielded mug shots. Students were asked the same question posed by Safiya Noble (2018), “If Google is not responsible for its algorithm, who is?” (Noble, 2018, p. 80). Some students felt that Google reflected the racial prejudice of its users (and thus was not directly responsible for seemingly racist results), while others felt Google needed to be more proactive (Fisher, 2017). Students reported the article made them think about Google in new ways. The course learning

outcome most closely related to this activity was to identify the ideas and perspectives behind current information issues such as evolving information technology in order to recognize the role of information in society.

Gardner (2019) described the curriculum for a public university credit-bearing information literacy course on algorithmic bias. The curriculum consisted of learning how Google indexed the web, algorithmic bias, an assignment comparing search results between Google and DuckDuckGo (using the searches climate change and immigration to the United States), identifying advertisements in the results, and analyzing bias within Google Image search results (searching for either God, working mom, Mexican, lead singer, doctor, handsome, beautiful women, cute children, happy family, professor, or expert). Students did well at deciphering the Google image search results but had trouble identifying advertisements. Students also wrote algorithmic directions for the specific task of how to make a peanut butter and jelly sandwich for someone with no prior knowledge of this using “if, then” statements. Then, the instructor attempted to make the sandwiches according to the directions. Other activities and subsequent discussion included Joy Buolamwini’s TED talk (Buolamwini, 2016) on her experience with racial bias in facial recognition software, algorithms used in parole and recidivism decisions, public opinions on algorithms, and browser extension tools like Panopticlick (Toner et al., 2017). The latter is a research project of the Electronic Frontier Foundation that tests online privacy protection; it is now called Cover Your Tracks (<https://coveryourtracks.eff.org/>).

Learning outcomes for the course were linked with four different frames from the *Framework for Information Literacy for Higher Education* (Association of College & Research

Libraries, 2016). Discussing the ways that search systems privileged some perspectives and communities while presenting barriers to others was tied to Frame 5, “Scholarship as Conversation,” to recognize that systems privileged authorities and that not having a fluency in the language and process of a discipline disempowered their ability to participate and engage. Articulating that search results were not a natural reflection of the world was linked to Frame 2, “Information Creation as a Process,” to recognize the implications of information formats that contained static or dynamic information. Questioning how algorithms shaped your personal choices and perspectives on a given topic was linked with Frame 1, “Authority is Constructed and Contextual,” to develop awareness of the importance of assessing content with a skeptical stance and with a self-awareness of their own biases and worldview. Finally, recognizing that search engine revenue came from advertising, not information retrieval services was linked to Frame 3, “Information Has Value,” to understand how the commodification of their personal information and online interactions affected the information they received.

Clark et al. (2017) received an Institute of Museum and Library Services (IMLS) grant entitled *RE:Search—Unpacking the Algorithms That Shape Our UX* to create an algorithmic awareness open educational resource course for librarians. Members of the project team developed the *Algorithmic Awareness Action Handbook* and a syllabus with activities (Clark, 2018). The learning outcomes included contextualizing basic mathematical formulas and codes within algorithms; explaining foundational concepts such as the Binary Search Tree algorithm and Dijkstra’s algorithm; the effects of algorithms on popular systems like Google search, Netflix, Facebook, and Amazon; the implications of gender and racial diversity and other ethical ramifications within algorithmic systems; and the importance of algorithmic awareness (Clark,

2018). One exercise provided a theorized example of the workings of an algorithm using pseudocode (an informal high-level description of the algorithm's operating principle augmented with natural language to make it easier to understand). Participants worked through adding values and making coding decisions for navigating the library (e.g., what would a patron ask and where would they go), then answered discussion questions as a group. Another component was a teaching tool that demonstrated an “x-ray” of the variables affecting their search query, such as the weather and location data. Yet another activity involved downloading personal data profiles (Clark, 2018). Clark explained algorithms in the context of Frame 1 of the *Framework for Information Literacy for Higher Education* (ACRL, 2016), “Authority is Constructed and Contextual,” because algorithms could be viewed as a “creator” informed by human decisions (Clark et al., 2018). Clark et al. (2017) argued for “algorithmic awareness” as a digital competency.

Mooney (2019) outlined a syllabus proposal for a course entitled “The Sociology of Information Disorder.” The content was intended to expand on core information literacy concepts in order to help understand the proliferation of “alternative truths” stemming from the current information infrastructure. Learning outcomes included analyzing the role of information and information technology in a democratic society, recognizing bias in information creators, interpreters, and processes of dissemination, and reflecting on personal information behaviors and awareness strategies. Section D was entitled “the information landscape” and included a lesson “how do algorithms impact what we know?” intended to explore the question of how algorithms powered the list of results from a search engine or a social media site (Mooney, 2019, p. 129). Awareness of algorithms was necessary to critically examine the tools that guide us to

sources. Another lesson in this section was “what is the extent of digital surveillance and how does it impact information-seeking and online behavior?” (Mooney, 2019, p. 130). It touched on the social impacts of big data and the surveillance-for-service process, which can lead to loss of control, behavior modification, and machine-mediated decisions. A third lesson explored “what is the impact of social media on public information?” with a focus on Facebook (Mooney, 2019, p. 132).

Hartman-Caverly and Chisholm (2020) advocated for librarians to teach privacy literacy, which they aligned with Frame 3 of the *Framework for Information Literacy for Higher Education* (ACRL, 2016), “Information has Value,” operationalized as recognizing how personal data and metadata were collected along with the potential implications; assessing how personal data was shared and making informed, intentional choices to safeguard privacy; identifying privacy issues facing our society; and describing the positive case for privacy as a human right fundamental to individual well-being. The article discussed the privacy paradox, a phenomenon where people’s actual behaviors often contradict their stated privacy values due to cognitive factors such as information asymmetries between the system or service provider and the user, lack of user knowledge of system design as it pertains to data flows, and lack of the technical and legal literacy needed by the user to understand privacy-related terms of service. Büchi et al. (2017), working from the field of Communication and Media, found that the more skilled Internet users were, the more they would engage in privacy protection. They defined privacy protection behaviors as frequently changing settings so that content was only visible to specific people; using fake information online such as a fake name; blocking, deleting, or deactivating cookies; and monitoring which information was available about you online. Additional

researchers Trepte et al. (2015) within the Communication Science field defined privacy literacy as having six dimensions: knowledge of the practices of online service providers, institutions, and organizations; technical aspects of online privacy and data protection; knowledge of potential privacy threats and risks; knowledge of laws and legal aspects; strategies for individual online privacy control; and strategies for dealing with privacy threats.

Hartman-Caverly and Chisholm (2020) stated that structural factors also accounted for the privacy paradox, such as system defaults which violate privacy, data aggregation and brokerage, and increased transaction fees for opting into privacy protections. There were also personal factors at play such as the time needed to make good privacy decisions. The article concluded with suggested active learning activities for a privacy literacy workshop, including students visiting privacy stations and answering these questions: Where have you left data tracks today; What data do you think is collected about you regularly; What apps do you use daily or weekly; What steps do you already take to protect your data; and What does privacy mean to you. Another activity asked students to independently use a list of interactive websites to explore data tracking and their personal advertisement behavioral profiling in real time to reveal backend algorithmic processes; websites included ClickClickClick (a browser-based game on online profiling available at <https://clickclickclick.click/>), Webkay (a demonstration of all the data your browser knows about you, available at <https://webkay.robinlinus.com/>), and personal Google, Facebook, Twitter, and Instagram advertising profiles. Also, small groups read case studies from an assigned category of location services, health data, criminal justice, or consumer profiling to explore the impacts of automatically monitored data on society. The last activity involved a

personal reflection on the benefits and risks of technology use in order to develop a personal data plan, which included a privacy checklist and additional resources.

Additional Teaching Tools and Strategies From Other Fields

Garrett et al. (2020) analyzed 31 standalone AI ethics classes and 20 artificial intelligence (AI) and machine learning (ML) technical courses from U.S. universities to understand which ethics-related topics instructors included in courses. They found bias was on 87% of the syllabi, usually taught through current event examples such as the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) recidivism algorithm, use of facial analysis to predict sexual orientation, and Google mislabeling African Americans as “gorillas.” Furthermore, they found the consequences of algorithms was on 45% of syllabi and focused on filter bubbles, recommender systems, propaganda, online advertising, confirmation bias, civil rights, and democracy. Current event examples used included the use of predictive analytics by child protection agencies and in criminal sentencing. Thirty-two percent of syllabi included privacy, and used examples such as Facebook’s “people you may know” feature, the Cambridge Analytica scandal, data brokering companies like Acxiom (<https://www.acxiom.com/>), and the use of analytics and surveillance to hire employees. Gallagher (2017) suggested activities to teach writing for algorithmic audiences in a composition class. For example, students attempted to manipulate Facebook’s algorithmically driven timeline to enable their writing to be read more widely; students managed metadata through YouTube such as keywords, description, tags, location, and the enabling or disabling of comments and ratings in order to better understand the processes and procedures by which YouTube prioritized their videos; students attempted to have their websites ranked highly by Google’s PageRank; students wrote an algorithm narrative about

an algorithm such as PageRank or Facebook's EdgeRank using terms of service agreements, white papers, or corporate statements to identify the values and ideologies of those who design algorithms; students tracked their website analytics (e.g., page views, duration) and wrote a weekly progress report to their peers, noting how changes impacted search results.

Several studies suggested activities related to Google searching. Bakke (2020) suggested a "Search Reflection" assignment in the context of health and medical topics where students record themselves searching on Google for their topics using the think aloud protocol, analyze the recording for key patterns in their search process, and finally evaluate their search process for strengths and weaknesses. It was suggested this assignment could be followed by a standard annotated bibliography assignment so students see the connection of source credibility across both assignments. Hobbs (2020) suggested a group activity where students signed into their Google accounts and conducted searches on countries like Finland, United States, or Germany on the same day. Students are then asked to post screenshots of their results on a digital bulletin board. Despite their similar ages, demographics, and geographic location, students' search engine results will vary greatly. Simon and Swartz (2012) created Image Atlas, a tool that lets you display search results from search engines in different countries at the same time to highlight cultural differences. Along similar lines, Ochigame and Ye (2021) developed Search Atlas, a tool to highlight how search results for the same query differed across different countries in order for users to reflect on the way results are shaped by cultural and political bias. They noted that creating such tools was challenging because "Google intentionally lacks an API [application programming interface] for web search results, and deploys various tactics to block scrapers" (Ochigame & Ye, 2021, p. 1981). Suggested topics from Powers (2017) warranting more

instruction for college students included an overview of algorithms, the implications of news personalization, the specific types of user data collected, the human judgments that go into programming algorithms and evaluating news sources, and customization options in portal sites such as Facebook and Google News. Brodsky et al. (2020) recommend a video on how algorithms worked to improve knowledge of online search.

Shapiro et al. (2020) described a series of lesson plans called “mapping self in society” (formerly Re-Shape) to teach students about their personal geography and physical movement data using the free ViewRanger application. Long et al. (2017) described LuminAI, an interactive installation at a local arts event that allowed humans to interact with artificial intelligence agents to improvise dance movements and learn more about AI. Gandolfi and Ferdig (2018) described the use of video games to show differences in algorithm visibility and to demonstrate hidden biases in some game algorithms. Fouquaert and Mechant (2021) developed an online tool called Instawareness (<https://instawareness.ugent.be/>) that functioned as a visual feedback tool to increase awareness of Instagram’s curation algorithms. Melsión et al. (2021) described a bias visualization tool to help children understand gender bias in AI. The University of Helsinki (2019) developed “Elements of AI,” a free online course. Brazil developed an introductory course called “Machine Learning for AI” (von Wangenheim et al., 2020). A 30-hour summer workshop for middle school students on AI literacy was developed through a collaboration between Massachusetts Institute of Technology (MIT), Personal Robots Group (PRG), and Boston College (Lee et al., 2021). A seven-hour AI literacy course was designed for Chinese university students based on the textbook *AI for Everyone* published by AI4kids (Kong et al., 2021). Marques et al. (2000) described 30 instructional units aimed at teaching machine

learning (ML) concepts in K-12 schools. Williams et al. (2021) developed a curriculum for middle school teachers called “How to Train Your Robot: AI and Ethics Curriculum.” Rodríguez-García et al. (2020) developed LearningML, a platform to teach and learn machine learning. Lindner et al. (2019) described several unplugged activities about AI developed for K-12 that involved physical objects to foster kinaesthetic engagement. Burton et al. (2018) discussed teaching computer ethics through science fiction, while Fiesler (2018) described using the television show *Black Mirror* (Bathurst & Kinnear, 2018) to do speculative futuring (encouraging speculation about the negative implications of technology in the future) with students taking an information ethics class.

Algorithms

Some major challenges to studying algorithms include the fact that the source code is often hidden, the fact that algorithms are constantly changing, and the complexity of the decisions made by predictive algorithms (Kitchin, 2017). Algorithms are “embedded within complex socio-technical assemblages made up of a heterogeneous set of relations including potentially thousands of individuals, data sets, objects, apparatus, elements, protocols, standards, laws, etc. that frame their development” (Kitchin, 2017, p. 20). There are many reasons for a general lack of algorithmic transparency, including a fear that revealing trade secrets may lead to people gaming the system or be a bad business decision, the fact that it is often impossible to determine when the outputs of algorithms may be biased or erroneous, concern that revealing input may violate privacy rights, and a fear of overwhelming the end user with too much information (Association for Computing Machinery US Public Policy Council, 2017; Diakopoulos & Koliska, 2017). Also, the end user may lack the technical knowledge or literacy

required to make sense of transparency (Willson, 2017). Platforms are strategic in sharing information about their algorithms and only make some information visible to certain actors (Cotter, 2019). Additional challenges include the lack of established baseline about how algorithms operate and gaps in vocabulary among users to describe their algorithmic encounters (Swart, 2021). Nonetheless, it is still possible to develop an introductory overview of the workings of common commercial algorithms for college students and how they influence everyday life in order to develop greater algorithmic literacy and awareness. Cotter and Reisdorf (2020) posited that knowledge about algorithms was best obtained through a platform-specific approach rather than a general approach, with direct experience of platform algorithms the best way to gain insight. Therefore, this section explored technical elements of common, everyday platforms college students were likely to use.

Technical Aspects of Algorithms

Algorithms are a set of instructions that make machines do something automatically according to a condition. Matthew Reidsma (2019) claimed algorithms should be viewed as a large set of binary choices that create a “decision tree” based upon many variables that are not necessarily apparent to the user. At the most basic level, algorithms must be defined to control for all possible circumstances. They start with an “if–then statement,” telling a program to execute a section of code only if the condition is “true.” (Bucher, 2018). To accommodate a “false” condition, the statements also need to include an “else” statement or alternative path (Bucher). The algorithm indicates what should happen when, which is called the “flow of control” and implemented in the source code or pseudocode (Bucher). Algorithms are independent of programming languages and specific machines. To be operational, algorithms are

based on particular representations and structures of data—they work in tandem with data structures to perform operations on them—as well as other elements like data types, databases, compilers, hardware, and CPU (Bucher). Donald Knuth (1998) stated that algorithms had five broadly defined properties: finiteness, definiteness, input, output, and effectiveness. Algorithmic orderings were never neutral because they come with “specific affordances that both enable and constrain . . . they come with certain assumptions and values about the world on which they are acting” (Bucher, 2018, p. 23).

There are deterministic (immutable) algorithms and machine learning algorithms; the latter learns to predict outputs based on previous examples of relationships between input data and outputs (training examples) and do not necessarily have one correct result (MacCormick & Bishop, 2013). During training, a computer is given a large set of labeled data, or samples that have already been classified, so that later it can use this to guess the right class when it is presented with an unlabeled sample (MacCormick & Bishop). Machine learning algorithms lack human oversight, and they change their behavior to enhance performance through experience (Bucher, 2018). Often, machine learning algorithms fall somewhere between supervised learning and unsupervised learning (i.e., no data about the desired outputs) (Bucher, 2018). In social media, machine learning algorithms learn to recognize patterns in the data and predict the likelihood of user actions and preferences. Models are constructed out of the set of discovered relationships and consistencies, which then can be used to automate things like classification, estimation, or prediction (Bucher). A task like filtering email spam is binary because it is classified as either spam or not spam. However, most tasks used in social media like ranking and recommending are more complicated.

Diakopoulos (2014) discussed four different types of decisions algorithms made. The first was prioritization, which brings attention to some things at the expense of others; an example is predictive policing which assigns risk and draws police attention to what has a higher risk. Prioritization has the potential to be discriminatory. The second was classification, which involves categorizing an entity as part of a given class by considering key characteristics; examples are anyone with a GPA above 95% is classified as being on the honor roll, or the Content ID algorithm that scans YouTube videos and classifies them as either violating copyright or falling under fair use. However, classification can make false positives and false negatives and runs the risk of being biased, especially if it has been tuned to privilege one over the other. The third was association, or creating relationships between entities (e.g., “related to” or “similar to”); for example, associating a potential foster parent with his brother who is a convicted felon. The last was filtering, or including and excluding information according to various rules or criteria, inputs to filtering often rely on decisions made during prioritizing, classification, or association. For example, news personalization apps filter news according to how the news has been categorized, associated to a person’s interests, and prioritized for that person, with the end result being either over-emphasis or censoring certain information (Diakopoulos, 2014). Latzer and Festic (2019) suggested algorithmic applications online encompassed nine different functions: search, aggregation, observation/surveillance, prognosis/forecast, filtering, recommendation, scoring, content production, and allocation. They also suggested the four main domains of everyday life where algorithms were most influential were social and political orientation, recreation, commercial transactions, and socializing (Latzer & Festic). Diakopoulos (2014) emphasized the human influences that were embedded in

algorithms such as criteria choices, training data and its potential bias, semantics, and interpretation.

Some common algorithmic techniques for information ordering that underpin many running systems on the Internet include logistic regression models, based on the concept of probability, and the Naïve Bayes classifier, which assumes that the presence of a particular feature in a class is unrelated to the presence of any other feature. The purpose is to rank variables with the most impact on the desired outcome, expressed as a value or weight (O'Neil, 2016). Rieder (2017) highlighted the Naïve Bayes classifier as a good example of the social power of algorithms because the technique was probabilistic (classifications are not binary but with degrees of certainty), adaptive (they 'learn' from experience), well suited for personalization, and could apply to almost everything, including emails, people, situations, or ideas. It was content based in that it used text or properties of documents (such as a list of words it contained as well as their frequency) as a stand-in for meaning, and it predicted membership probabilities for each class by calculating the probability that a given record or data point belonged to a particular class. Very frequent or very rare words were excluded (Rieder, 2017). A similar technique is the term frequency-inverse document frequency (TF-IDF), used to retrieve relevant documents from a collection of documents by looking at words that appeared very often as assigned to a certain category in training but were rare for others and seeing these as a strong indicator for that category (Clark, 2018). It was calculated by multiplying the term frequency (how often a particular term appears in the document) by the inverse of the document frequency of that term (how many documents contain at least one instance of that term) (Clark, 2018).

Google's proprietary PageRank algorithm had roots in TF-IDF, and it ranks documents for a search term and identifies search term value matches (Clark, 2018). PageRank also had roots in Binary search; it is a search algorithm that finds the position of a target value within a sorted array by comparing the target value to the middle element of the array (Clark). If they are unequal, the half in which the target does not lie is eliminated and the search continues on the remaining half until it is successful—if not ever successful, the target is eliminated (Clark). Other common algorithmic techniques include the nearest neighbors, which finds the nearest neighbor to that sample in the training data and then uses the class of this nearest neighbor as its prediction (MacCormick & Bishop, 2013). A variation on this is the K-nearest neighbors, where K is a small number like three or five, in which case you choose the class that is most popular among these neighbors (Bucher, 2018). In Facebook, nearest-neighbors is applied by finding the image most similar to it in Facebook's database of labeled photos to determine if it is a face (Bucher). Another algorithm used in Facebook to find the shortest path between friends and associates or create weights that define closeness of a friendship is Dijkstra's algorithm (Clark, 2018). Dijkstra's algorithm finds the shortest path between nodes in a graph; a variant fixes a single node as the "source" node and finds shortest paths from the source to all other nodes, producing a shortest-path tree (Clark).

Additional common algorithmic techniques include support vector machines, which analyze data for classification or regression analysis; and neural networks, which specify the depth of hidden layers to process information similar to how the human brain uses neurons, inputs, and outputs and adjusts to changing input (Bucher, 2018). Also, there are decision trees, which use a series of sequential decisions to reach a specific result similar to the game "20

questions”; and random forests, which leverage the output of multiple (randomly created) decision trees to specify the number of trees and splits to try in each tree (Bucher, 2018; MacCormick & Bishop, 2013). Merge sort is another general-purpose, comparison-based sorting algorithm used to compare user’s Netflix video ranking and preferences (Clark, 2018). It is a divide and conquer algorithm, which means that it is broken into subproblems which are repeatedly solved, then these solutions are combined to solve the original problem (Cormen, 2013). In this case, the array is divided and sub-sorted, then brought back together, with the input order of equal elements being preserved in the implementation of sorted output (Clark, 2018). Lastly, the Bloom filter is a type of collaborative filtering designed to test whether an element is present in a set (Clark, 2018). It is fast, but can turn up false positives—a query returns either “possibly in set” or “definitely not in set.” It is used on Amazon to set up user recommendations based on current and past browsing or purchasing patterns (i.e., only show users what they have not purchased previously) (Clark).

Google Search Engine Algorithms

According to Lewandowski (2017), a search engine result consisted of at least three different ranking functions: (a) ranking of the results lists from the Web index; (b) ranking of vertical results within collections such as news, images, etc.; and (c) ranking of Universal Search containers (boxes presenting top results from vertical search engines) within the general Search Engine Results Pages (SERP). Ranking within collections was based on several factors. The first was text, or linguistic cues that included word frequency to identify the webpages on which the user’s query terms seemed most important; the identification of synonyms such as vocalist versus singer; relevant phrases such as differentiating a *hot dog* from a hot *dog*; and web

metadata (Granka, 2010). The second factor was popularity, and popularity cues included PageRank (Google’s proprietary search algorithm), or the number of links to a site equating to “votes” for that site along with the relative “authoritativeness” of the site (Granka). PageRank tries to measure relevance based on popularity, but hyperlinks are not stable and therefore Google search rankings are always changing. Part of popularity is also how often a site has been clicked on by other users (Lewandowski, 2017). Other factors for this include user clicks to average the most frequently clicked on results across all users who have issued that query, reading time on webpages, and patterns of query reformulation (Granka, 2010). The third factor was freshness or how new the content was (Lewandowski, 2017). The fourth was locality, or the geographical location of the user. The fifth was personalization, or the interests of the user including past behavior. The last was technical rankings factors such as how fast a server could process requests and how reliable it was (Lewandowski).

Search engines also present text-based, contextual advertisements on the SERPs that often say “sponsored link” (Lewandowski, 2017). They often appear on the right-hand side of the search or at the top of the search results themselves. Any content in the featured snippet is coming from a third-party website, and not Google itself (California State University, Dominguez Hills, 2019). Featured snippets tend to be YouTube (owned by Google) videos, lists, or short paragraphs of content. Another influence on the results is search engine optimization (SEO); websites can be optimized by their creators to be featured in a snippet (California State University, Dominguez Hills, 2019). Google first began to personalize its search engine in 2009, using 57 variables from users’ behavior to decide what results to display after the keywords were entered (Hobbs, 2020). Today Google delivers different results to different users based on over

200 criteria, including language, geographical location, gender, time spent on a website, click rates, time of day, and previous search history (Dean, 2021). Given how dominant Google is, it now has access to a huge amount of personal information from most users including Gmail, Google Maps, YouTube views, Google Docs, Google Calendar, and Android phone data.

Google argued that its right to prioritize how it presented information to users was a First Amendment, freedom of speech right (Souto-Otero & Beneito-Montagut, 2013). The truth is that the goal of Google is to make money, not be a common good or a shared community (Lindh & Nolin, 2016). The weighting of each criterion changes and is re-established every time a query is searched (Graham, 2017). Google's PageRank algorithm establishes "a hegemonic power structure that prevents users from accessing the web without Google's influence" (Graham, 2017, p. 8). PageRank can be abused; for example, Google bombing was the practice of excessively hyperlinking to a website to cause it to rise to the top of PageRank and co-opting content or a term and redirecting it to unrelated content (Noble, 2018). An example was when George W. Bush became associated with the search terms "miserable failure." Search engine optimization is the general process of using a range of techniques, including augmenting HTML code, web page copy editing, site navigation, linking campaigns and more, in order to improve how well a site or page gets listed in search engines for particular search topics (Noble). Many SEO companies have been developed to influence the process of moving search results up. Students actively, though often unknowingly, participate and shape the content that is shown in search results. In this respect, they are both consumers of and creators of search algorithms contributing to result ranking. It is important to increase awareness of the commercial interests that highly influence their search results.

Google has two components to its ad business which influence the results: AdWords and AdSense. Google AdWords is part of “Google properties” and hosts advertisements built into its own products (e.g., its search engine and Gmail) (Graham, 2017). AdWords is an auction process that Google operates to add sponsored results to search engine queries, which sit separately on top or to the side of organic (unpaid) results (California State University, Dominguez Hills, 2019). An auction occurs every time a search is conducted. Google gives a “quality score” to each company placing a bid. If a user clicks on one of these sponsored advertisements, the company being advertised pays Google; if not, no money is exchanged. The Google AdSense program is part of the “Google Network Members’ properties” brokerage service, and it lets Google link relevant third-party content (like blogs or news sites) with advertisements on third-party websites (non-Google sites) (Graham, 2017). Put another way, AdSense are ads purchased through Google that appear on other websites (not the search engine). These banner ads are often at the top of other websites, and they may relate to something previously searched (California State University, Dominguez Hills, 2019).

YouTube Recommender Algorithm

Google bought YouTube in 2006 for 1.65 billion (Associated Press, 2006). Koul (2019) described the five places a video can appear. The first place was on the Home page, where videos are displayed based on viewer’s subscriptions or based on videos watched by those with similar interests and watch patterns. The second place was the search box, which matches up with metadata such as title, tags, and thumbnails. The third was suggested videos, which show up on the side panel or as “watch next” if autoplay is turned on, and these are based on prior activity from the viewer’s watch history and topically related content. The fourth was trending, which

shows new and popular videos based on the viewer's country or geography and the growth rate of views. The last was subscriptions. The algorithm behind YouTube recommendations or "up next" videos was based on deep learning and scalability embedded in its algorithm ratios (Covington et al., 2016).

YouTube employed Google Brain (also open sourced as TensorFlow), the deep neural network architecture system. As discussed earlier, neural networks are machine-learning technologies that act similar to the way a human brain thinks. Covington et al. (2016) explained that the recommendation system consisted of two neural networks: the first of which was candidate generation. The candidate generation network provided broad personalization through collaborative filtering that takes input from user activity history and narrows down the potential video corpus from millions of potential videos. A nearest neighbor lookup was performed to generate hundreds of candidate video recommendations. This model was trained on the implicit feedback of user watch history, where a user completing a video is a positive example and a user not clicking on a video is a negative example. Also, the age of the training example was used as an input feature to represent the time-dependent behavior of popularity to avoid an implicit bias towards the past because of being trained to predict future behavior from historical examples. Other signals used include search history and user demographics like geography, device, gender, age, and whether logged in (Covington et al.).

The second neural network used was ranking to eventually assign an independent score to each video, with the end goal being expected watch time (found to be a better capture of engagement than click-through rate) (Covington et al., 2016). Some factors that impacted ranking included video quality (view count, ratings, commenting, favoriting, sharing), user

specificity (user view count of and interaction with the seed video and time of watch), and diversification (removing similar videos) (Davidson et al., 2010). There are hundreds of feedback data points in total (Shin, 2020). If a user was recommended a video but did not watch it, it got demoted on the next page load (Covington et al., 2016). Koul (2019) reported that YouTube signals in algorithms are now based less on authority (e.g., number of views, subscribers, and watch time) and more on viewer behavior such as videos watched, click through rates, average time spent watching, and engagement (such as likes, dislikes, comments, adding to a playlist, subscribing, ratings). Live A/B testing is used to evaluate the performance of the recommendation system: “live traffic is diverted into distinct groups where one group acts as the control or baseline and the other group is exposed to a new feature, data, or user interface” (Davidson et al., 2010, p. 296). The two groups are then compared against one another over a set of predefined metrics and possibly swapped for another time period. Rieder et al. (2018) studied YouTube’s “relevance” search algorithm, which provided an ordered list of videos in response to a user query, and found that YouTube’s search function was highly reactive to attention cycles and the dominance of YouTube-native contents.

Instagram

Instagram launched as a free iPhone application in 2010 primarily as a place to post and share photos with custom filters (Laestadius, 2017). In 2012, Facebook purchased Instagram (Dinger, 2020). Instagram expanded its services beyond photo sharing to include video in the Feed up to 60 seconds, transitory Stories (images and videos that disappear after 24 hours), and the video platform IGTV (Instagram television) for longer videos (Instagram, 2021). It also expanded advertising to include not just sponsored posts and advertisements, but also direct

shopping in the platform. The data policy in Instagram's help center stated, "we may use information about what you do on Instagram and Facebook as well as your activity on third-party sites and apps you use" as well as warned they may also use information "provided by businesses outside of Instagram or Facebook Company Products to decide which ads to show you" (Instagram, 2021). Examples listed included people you follow and posts you like, information and interests on Facebook, websites and apps you visit, or information that advertisers, their partners, and Instagram marketing partners share (Instagram, 2021). Instagram had "influencers" or power users that were often hired by brands to advertise products and lifestyles in a seemingly organic fashion (Davies, 2019). A user could follow someone even if they do not follow the user back, follow hashtags, or "mute" to hide posts or stories from a certain account without unfollowing that account (Instagram, 2021); also, you could tag others in photos and videos, share posts on Facebook and Twitter, write captions, and add location information and tags (Dinger, 2020). Methods to track user interactions varied depending on the mode: posts had likes and comments, but stories had view counts, while Direct Messages (DM) had likes (Instagram, 2021). The Instagram algorithmic feed for displaying posts began in 2016 and included the following factors: past behavior on similar content (including what type of content you like to engage with such as video versus photo), recency, relationship with the person sharing (previously viewed, liked, tagged together in photos, saved posts, DM replies, or commented on posts), frequency of logging in and engaging, how many people you follow, and usage time (Constine, 2018). Also, the engagement the post had received in general was a factor (Davies, 2019). A post was first shown to a select group of follows to gauge engagement (Davies).

Twitter Timeline Algorithm

Twitter started in 2006 and was a niche social network, with the goal of helping people find out what was happening right now (Davies, 2019; Nemeth, 2020). It was favored by the media and political elites and thus played a key role in driving the news agenda. In 2016, Twitter introduced its new timeline algorithm which had previously been displayed in reverse chronological order. The hashtag #RIPTwitter started trending because Twitter users were outraged with the way their timelines were rearranged (DeVito et al., 2017). In 2017, Twitter increased its limit from 140 characters per tweet to 280 characters (Huddleston, 2017). Nemeth (2020) explained that the Twitter algorithm used deep neural networks and had three sections. The first was ranked tweets which appear at the top of the timeline and were based on a relevance model plus recency; the second was “in case you missed it” which showed tweets ordered by relevance; and the last was reverse-chronological order tweets (after seeing all of the Tweets selected by the algorithm as highly relevant, the remaining would show up in reverse chronological order).

The algorithm used the following factors to make predictions of relevancy: a tweet’s recency; the presence of media card attachments (image or video) in the Tweet and if the user preferred that type; weighting of the total interactions with the Tweet including clicks (like = 1; reply = 2, and retweet = 3 points), the Tweet’s author and your past interactions with them, strength of connection with them, and origin of your relationship; tweets you found engaging in the past; and how often and how heavily you used Twitter (Davies, 2019; Koumchatzky & Andryeyev, 2017). Additional factors included profile credibility (i.e., avoid being spammy, having broken links, and automation software); whether the content was native versus linked to

other sites (native took precedence); time spent reading someone's Tweets or visiting their profile; use of trending hashtags; and timing of the Tweet (when more users were online to see it) (Davies, 2019). A tweet was first served to a small percentage of users to measure initial engagement. A/B testing was done to gauge usage and enjoyment of Twitter (Koumchatzky & Andryeyev, 2017). Users had the option to switch between the algorithmic and chronological timelines, but the algorithmic timeline was the default (Davies, 2019; Nemeth, 2020). Twitter advertisements were naturally integrated into the feeds of relevant followers and there were three types: promoted accounts, which appeared in the user's "Who to Follow" sections; promoted trends, which showed up on top of a user's trending topics; and promoted Tweets, which showed up near the top of the timeline (Barnhart, 2017).

Facebook Ranking Algorithm

In 2012, Facebook purchased the Instagram social photo sharing application for \$1 billion dollars (Upbin, 2012). Both Facebook and Instagram curate users' content and feeds based on similar algorithms. Facebook used a set of algorithms called Ranking, which used to be called The EdgeRank algorithm. Ranking used signals, or data points, to organize stories for the Facebook news feed and decide what posts people see and in what order (Cooper, 2020). Key factors for making the decisions included four elements: (a) the available inventory of stories (stories available to you that you have not yet seen); (b) the signals, or data points that can inform ranking decisions such as recency of a story, who posted it, do you typically interact with this person or have mutual interests, how fast is your Internet connection, what kind of phone, type of media in the post such as video or photo and your past interactions with that type, popularity of a post (how many others have posted the same thing) and meaningful interactions

with a post from other users like commenting on and sharing, and complaints or negative feedback from other users on a post); (c) the predictions made, including how likely you were to comment on a story, share with a friend, hide a story, or report a story; and (d) a relevancy score for each story representing a weighted score of how interested you were in any given story (Mosseri, 2018).

There were even more factors at play. Back in 2015, the Facebook News Feed curation algorithm employed 100,000 factors to customize the feed of any given user, doing more to omit content rather than include it (Eslami et al., 2015). The personalization algorithm that created these unique platform media environments were multi-nodal sourcing from the user history, user preferences, location data, level of novelty, personal networks and the advertisers (Bozdag, 2013). DeVito (2017) discussed the differences between Facebook story selection and traditional news values, finding a set of nine News Feed values that drove story selection: friend relationships, explicitly expressed user interests, prior user engagement, implicitly expressed user preferences, post age, platform priorities, page relationships, negatively expressed preferences, and content quality. Also, there was evidence that friend relationships acted as an overall influence on all other story selection values. Chin et al. (2015) found that because of conformity motivation, when a large number of Facebook users had previously “liked” a post, this made someone more likely to also click the “like” button and go along with the crowd. It was also possible to reorder your feed by date.

In 2019, the “why am I seeing this post?” feature was introduced (Cooper, 2020). Kitchin (2017) wrote that the Facebook algorithm “does not act from above in a static, fixed manner, but rather works in concert with each individual user, ordering posts dependent on how one interacts

with ‘friends.’ Its parameters then are contextually weighted and fluid” (p. 21). Bozdag (2013) added that when Facebook promoted recommendations from the user’s most active friends and demoted actions from less active friends, the system controlled both the users’ information and who they could reach. Facebook also injected ads into the News Feed and used a similar but separate ranking algorithm to determine the degree of interest in a page or business’ ads (Constine, 2016). Facebook limited the number of ads you saw, and therefore wanted to maximize the likelihood that the ones it showed you would resonate with you or get you to click on them so it could earn more money (Constine).

Algorithms for News

The rise of social media platforms for the dissemination of news has contributed to the decrease in popularity of traditional and established modes of journalism (Head et al., 2020). Shin (2020) described three approaches to algorithms in news recommendation systems. The first was collaborative filtering, which predicted user preferences as a linear (constructed by adding the results for each term), weighted combination of other user preferences. It recommended items based on the behavior of other users in terms of ratings, shopping and purchase information, and transaction data—in other words, it leveraged the experiences of users with similar profiles. Also, the number of news articles to be considered for recommendation was used as a parameter along with the number of steps needed to reach a given state. The second approach was content-based filtering, which made recommendations based on user preferences for product features, considering both user and product. Using the content of shared attribute space (a compact, abstract representation with a set of intervals and values), a user-

profile was constructed as well as an item-profile, and then the two were compared. The third approach was a hybrid approach, combining both collaborative filtering and content filtering.

Weber and Kosterich (2018) did a content analysis of the code in 59 open source mobile news apps to study the role of algorithmic judgment in news filtering and distribution. They noted that designers of news applications had different priorities than journalists and mainly focused on code, personalization algorithms, and relational databases rather than journalistic standards or relative importance. The algorithmic process of news production was described, with an initial information collection phase that often included recent items “liked” on social media and the country and language. This was followed by a pre-processing stage of data cleaning (making it compatible for a mobile app) and a training set for the algorithms specifying priorities such as most popular or keywords of articles liked. The actual algorithms used include Naïve Bayes (for probability filtering), support vector machine (SVM) for categorization based on dimensions, and k-nearest neighbors. The algorithm then predicted what the user would like and placed it at the top of the feed (Weber & Kosterich).

Diakopoulos and Koliska (2017) developed guidelines for algorithmic transparency in news algorithms based on focus groups with journalists. They came up with factors that offered opportunities for more disclosure across four layers or stages in the algorithmic process related to automatically generated algorithms, algorithmically enhanced curation such as when algorithms select, curate, recommend, or personalize Facebook newsfeeds, and dissemination such as prediction models. The first layer was when data is fed into the system, and opportunities for greater transparency lie in the quality of the information such as its accuracy, uncertainty or error margins, timeliness, and completeness. Other opportunities were in the sampling method,

definitions of variables, sources (public versus private), volume of training data used in machine learning, assumptions/limitations of the data collection, and inclusion of personally identifiable information. The second layer was the model and how it was transformed, and opportunities for greater transparency included the input variables and features, the target variables for optimization, feature weightings, name or type of model (linear versus nonlinear), software modeling tools used, source code or pseudo-code, ongoing human influence and updates, and explicitly embedded rules or thresholds for interpreting a value. The third layer was inference or output (e.g., classification, predicted scores, recommendations) and opportunities for greater transparency included existence and types of inferences made, benchmarks for accuracy, error analysis including protocols for fixing them, and confidence values or other uncertainty information. The fourth layer was interface for the end-user, and opportunities for more transparency included integrating algorithm information directly into the user interface that the user could interact with through some kind of algorithmic presence signal such as frequently asked questions (FAQ) or icons, on/off options for overriding algorithms, and tweakability of inputs and weights.

Amazon Item-Based Collaborative Filtering

Amazon began as a bookstore in 1994 but is now one of the most popular commercial sites on the web. Amazon launched item-based collaborative filtering, also known as item-to-item collaborative filtering, in 1998 (Linden et al., 2003). The process involved finding related items for each item in the catalog, with related defined as “people who buy one item are unusually likely to buy the other” (Smith & Linden, 2017, p. 13). Items were combined based on a customer’s current context and previous interests, then items already seen were filtered out.

The goal of relatedness was perceived quality and usefulness for the user. The processing of pairing similar items was done offline, allowing it to happen in real-time. The algorithm received information about which recommendations were effective in terms of resulting in a purchase, like, or click and learned from this to make improvements. Other factors considered included the order in which things were purchased in terms of time directionality (e.g., memory cards are usually purchased after a camera); newer items needed an explore/exploit since they had not yet had much opportunity for purchase; and some items had a lifecycle so immediate intent was important. For example, with one-time purchases, recurring purchases like toothpaste versus evolving purchases like a baby rattle that should change over time as the baby grows versus uncertain intent that should allow for discovery and serendipity. Since Amazon had so much data, the false signals tended to die off (Smith & Linden). Recommendations appeared in several places: the homepage had recommendations based on past purchases and items browsed or rated; the search result pages showed related items; the shopping cart recommendations showed recommended items; and more recommendations showed up at the end of an order suggesting items for later (Smith & Linden). Amazon has had claims of algorithmic bias: one example was when Amazon de-ranked gay romance books due to an algorithmic alteration that recategorized them as not family-friendly, which is an example of how algorithms have the power to reassemble culture (Lloyd, 2019). Another example was differential or dynamic pricing—“online shoppers who live in wealthy neighborhoods are offered products at different prices than those in less affluent neighborhoods” (Hobbs, 2020, p. 523).

Social Aspects of Algorithms

There is a growing field of “critical algorithm studies” that goes beyond the technical considerations of algorithms to consider their socio-cultural contexts and the large influence algorithms have on everyday life and society (Beer, 2017; Gillespie, 2014, 2016b; Kitchin, 2017; Seaver, 2017; Striphas, 2015; Willson, 2017). Machine learning algorithms should not be viewed as only objective computational processes limited to the output of lists, processes, or organized steps. They are socially constructed and relational in that they are dependent on interaction with humans to construct meaning, resulting in a unique combination of the technical and the social (Lloyd, 2019). Rieder (2017) described algorithms as “the outcome of situated encounters between computing environments, algorithmic techniques, and local requirements” (p. 102). Algorithms are “always *in becoming*” since events are not static and users always feed algorithms new data (Bucher, 2018, p. 28). They change with the event and also have the ability to change the event—when users offer up their data, algorithms rewrite themselves in an endless feedback loop (Bucher, 2018). Seaver (2017) argued that the best way to study algorithms was ethnographically as sociotechnical systems that were part of culture, using ethnographic techniques such as scavenging and fieldwork. Seaver (2019) argued that algorithms were social constructions that formalized informal qualities such as your taste in music but were unstable; the constant interplay between algorithms and humans ensured there were always many permutations and you would never log into the same Facebook twice, for example. Therefore, if you approached algorithms as an outsider, your discoveries were often “partial, temporary, or contingent” (Seaver, 2019, p. 415).

Willson (2017) noted that we delegate tasks and everyday activities to algorithms, and “the way it is . . . engaged with in turn impacts upon those things, people and processes that it interacts with—with varying consequences” (p. 139). Furthermore, “algorithms are embedded in complex amalgams of political, technical, cultural and social interactions” (Willson, 2017, p. 141). Beer (2017) insisted that algorithms were modelled on visions of the social world, and with outcomes in mind influenced by commercial interests and agendas. To study algorithms as separate from their social ecology was a mistake; algorithms exist both in code and in the social consciousness as a concept that is frequently used to stand for something. Rieder (2017) pointed out that someone, a human, had to make decisions on what and how to formalize algorithmic techniques so that the best possible input/output equation was used to get the best results—derived from an encounter between data, a purpose, and a feedback mechanism. These critical studies cut across a wide range of disciplines, including new media studies, platform studies, information science, human-computer interaction, philosophy, sociology, law, and computer science. Algorithmic systems make predictions that impact people’s lives through choices in classification, sorting, ordering, and ranking; these decisions in turn shape what people know, who they know, what is visible, and what they experience. Power is “operationalised through the algorithm, in that the algorithmic output cements, maintains or produces certain truths” (Beer, 2017, p. 8). Users’ understanding of the world and themselves is changed through algorithmic interaction.

Gillespie (2014) discussed six dimensions of “public relevance” algorithms and the socio-political ramifications for the resulting shifts in public knowledge. The six dimensions were:

- patterns of inclusion (the selection of some information and the exclusion of other information);
- cycles of anticipation (inferences about users);
- evaluation of relevance (algorithms determined what is relevant, correct, or legitimate knowledge);
- promise of algorithmic objectivity (algorithms presented themselves as impartial and exempt of human intervention);
- entanglement with practice (algorithms influenced users to change their practices); and
- the production of calculated publics (algorithmically generated groups or outputs made claims about the “public” they purported to represent and the users’ place amidst them, which in turn shaped the public’s sense of self). Gillespie (2014)

He concluded that a sociological analysis of algorithms should look at both the process of how algorithms chose information for users as well as the social process by which this turned into a legitimate and trusted source.

Power of Algorithms

Pasquale (2015) asserted that since it was unknown how algorithms weigh and calculate our data during Internet use, we cannot know the value of our actions online. The protection granted to algorithms as trade secrets perpetuated a black box society where the values embedded in the codes remained hidden. As a result, people do not realize the large amount of information being collected about them, how it is used, or what the consequences might be; nonetheless, a large amount of authority is increasingly expressed algorithmically. Smith and

Linden (2017) noted “recommendations and personalization live in the sea of data we all create as we move through the world” (p. 18). Willson (2017) noted that Google ran continual secret experiments on its users (as of 2013, they were running 10,000 experiments a year in search and ads) and made constant updates to its algorithms that shift priorities and boost certain results over others (such as privileging mobile-optimized sites in mobile technology search results). Lewandowski (2017) performed a brief comparison of Google’s statements with its actual practices and found that Google operated on statements that were at least in part contrary to their actual practices. Our increasing dependence on algorithms for everyday activities has made the power of companies like Google normal and expected (Willson, 2017).

According to Ip (2018), there were three types of data on the Internet: (a) data that you consciously gave companies (e.g., your name, email, date of birth); (b) data that was automatically monitored (e.g., where you logged in from, what time you did it, where else you visited on the web); and (c) data that was modeled or predicted from other data (e.g., your quantified attractiveness or trustworthiness). Personal data from the Internet has been used to create personality profiles, or psychographs, that organized user data into behaviors that fit in the “big five” of traits: openness, conscientiousness, extraversion, agreeableness, and neuroticism (also known as OCEAN) (Cohen, 2018). User data may also include posts, comments, text messages, and uploads, as well as media that was deleted before it was posted; online and offline data gets merged into a profile that is sold to advertisers (Ip, 2018). Cheney-Lippold (2011) noted that the resulting user categorizations constructed by algorithms constituted a form of surveillance, and “[use] the capacity of suggestion to softly persuade users towards models of normalized behavior and identity” (p. 177). Newer forms of algorithmic personalization that take

on the role of surveillance include facial recognition software monitoring emotions, and sensors that track movements. For example, a smartwatch or wearable device monitors heart rate or movement, and home assistants analyze tone of voice and vocabulary in order to combine real-time emotion data with user targeting (Hobbs, 2020). This is scary because it has the potential to use a person's emotional state as a weapon to sell them a fake news story or ad they will be vulnerable to. Also, human elements are now translated into reductive and inhuman data points in products that blend the offline with the online—for example, in wearable technologies where “the process from biology (heart) and practice (walking) to data becomes unquestioned, normalized, and invisible” (Willson, 2017, p. 148).

Boyd and Crawford (2012) gave six warnings about the potential harms of big data from algorithms used to uncover large-scale patterns in human behavior:

- Big data changes the definition of knowledge.
- Claims to objectivity and accuracy are misleading.
- Bigger data are not always better data.
- Taken out of context, big data loses its meaning.
- Just because it is accessible does not make it ethical.
- Limited access to big data creates new digital divides.

Graham (2017) stated “It is important that web users understand that the web is structured around financial incentives and that, collectively, the actions of following links and sharing pages are intrinsically economic and carry significant consequences for the future of the global information ecology” (p. 17). Also, as discussed in Chapter 1, Zuboff (2019) used the term surveillance capitalism to describe the collection and commodification of user behavioral data

for both persuasive and predictive purposes. Companies gather “data exhaust” from our online interactions, defined as the information “incidentally generated as people use computers, carry cell phones, or have their behavior captured through surveillance which becomes valuable when acquired, combined, and analyzed in great detail at high velocity” (Head et al., 2020, p. 49). Head et al. (2020) also used the term “attention economy” to describe how platforms and their advertisers try to engage and keep people’s attention online as long as possible to make money, which encouraged algorithmic design similar to clickbait (designed to make the user click, often at the expense of accuracy or quality).

Algorithmic Bias and Ethics

According to Pessach and Shmueli (2020), there were four main causes of biases or unfairness in machine learning algorithms. The first was when bias was already included in the datasets used for training, often due to historical biases. The second was bias caused by missing data because the sample was not representative of the population. The third was bias that stemmed from algorithmic objectives aimed at having high prediction accuracy at the expense of minority groups, and the fourth was bias when “proxy” attributes were substitutes for sensitive attributes that exploited gender, age, and race. As discussed in Chapter 1, much has been written about algorithmic bias in many different areas of society, including search engine results (Noble, 2018), credit scores (Nopper, 2019); the criminal justice system (O’Neil, 2016); and in things like drop-down boxes that lack options for the multiracial (Wachter-Boettcher, 2017). Baker and Potts (2013) reported how the auto-complete search algorithm offered by Google produced suggested terms which could be viewed as racist, sexist or homophobic. During the same year, Google showed bias in algorithmically presented results of personal name searching; more than

2,000 searches of racially associated names across two websites revealed that names suggestive of Black people yielded ads suggestive of an arrest most of the time, compared to names suggestive of White people that did not (Sweeney, 2013).

In the social services domain, Virginia Eubanks (2018) documented the bias in an algorithm used by a county child welfare agency that predicted how much at risk a child was for maltreatment. The algorithm used community calls to a child abuse hotline (reflecting the bias of neighbors) and child placement in foster care (reflecting the bias of caseworkers) as proxies for child maltreatment, along with predictive variables like use of public services and other measurements of poverty. This was biased towards the poor because public service use data did not include equivalent data on wealthier parents who used nannies or private drug treatment, for example. Even though the algorithm was supposed to support rather than replace human decision-making, the algorithm scores influenced child abuse investigators and made them question their own judgment. Also, accepting help from the County came with strings—parents were under even more scrutiny and had to comply with a number of extra duties or they would lose their children, leading to even more stress and ironically sometimes creating the abuse it was seeking to prevent.

Buolamwini and Gebru (2018) studied facial recognition software and found darker-skinned females were the most misclassified group (with error rates of up to 34.7%), compared to the maximum error rate for lighter-skinned males at 0.8%. Gillespie (2016a) offered this explanation for how bias can creep into algorithms for facial recognition: “sometimes new phenomena emerge that the training data simply did not include and could not have anticipated; just as often, something important was overlooked as irrelevant, or was scrubbed from the

training data in preparation for the development of the algorithm” (p. 21). For example, if an algorithm was trained on selfies from an online service used primarily by people of only particular races, then it would not be as accurate when applied to a more diverse population. Algorithms are usually programmed for efficiency, but the impact of the coders’ biases or the impact of our culture’s bias are often not taken into account. As Cathy O’Neil (2016) put it, “a model’s blind spots reflect the judgments and priorities of its creators” (p. 21).

Other researchers reporting algorithmic bias included Caliskan et al. (2017), who found cultural stereotypes embedded in popular online translation systems. Chen et al. (2018) investigated the rankings of candidates produced by three different resume search engines with respect to inferred gender and found overall, men ranked higher than women with equivalent features. Edelman and Luca (2014) tested for racial discrimination against landlords in the online rental marketplace Airbnb and found that non-Black hosts charged approximately 12% more than Black hosts for the equivalent rental. Turow (2011) explored how digital profiling produced different offers and access to premium brands based on a person’s financial attractiveness or social risk. On the opposite side, predatory ad targeting by for-profit universities to low-income residents was often partnered with “too good to be true” loan options (O’Neil, 2016). Vulnerable populations often exposed their pain points unknowingly through Internet use (e.g., through Google searches) and this data would get mined for profit (O’Neil). Marginalized groups were especially vulnerable because they often lacked the power and privilege to be able to fight back against algorithms built on hegemonic (White, cis, male) perspectives micro-targeted to exploit their weaknesses. As one student commented during a user study on algorithms, a predictive algorithm was “just a fancy technology-enabled form of stereotyping and discrimination, which

is inherently problematic, but it's easier for us to overlook because it's happening online and we're not seeing it" (Head et al., 2020, p. 16). Obermeyer et al. (2019) performed an algorithm audit of a widely used healthcare algorithm that targeted patients for high-risk care management programs by predicting who would benefit the most and found a large amount of racial bias when comparing White and Black patients. The algorithmic risk scores used predicted health costs as a proxy for real illness and health needs; this was biased because Black people generate fewer medical expenses due to things like doctor bias in recommending less preventive care for Black patients or less frequent use of health care services due to competing demands of child care or jobs. The researchers suggested that using better algorithm training labels like number of chronic conditions would yield less biased results.

Keyes (2018) conducted a content analysis of academic papers in the area of automatic gender recognition technology to show how research in this area treats gender as binary, immutable, and with a physiological component the majority of the time; this illustrated that gender recognition technology is not trans-inclusive and does not rely on self-disclosed gender information. In another paper, Keyes et al. (2021b) argued for gender to be treated as multiplicitous in gender classifier algorithms and provided a list of sensitizing questions. Keyes et al. (2019) called for the human computer-interaction (HCI) community to focus less on universalism, efficiency, and a one-size-fits-all approach; instead, they should "interact with other communities on their terms, with an expectation that their members are those best-equipped to define and describe the difficulties being faced" (p. 5). Keyes et al. (2021a) argued that when artificial intelligence is used to classify identity (for example, diagnose autism or identify homosexuality), it is "being used to reinforce existing stereotypes of disability and

sexuality, further legitimizing their violence” (p. 170). In the case of using facial recognition to diagnose autism, often inconsistent data sets were used as inputs due to hidden disabilities and lack of disclosure (Bennett & Keyes, 2019). There was also the problem of pre-existing bias around gender, race and ethnicity, class, and geography in autism diagnostic tools used in the medical field brought in by training data; by adding technical and scientific authority, this gave even more weight to decisions that had a life-changing impact but were not always correct (Bennett & Keyes, 2019).

There is growing attention given to the concept of algorithmic fairness and ethics in machine learning algorithms. For example, The Institute of Electrical and Electronics Engineers (IEEE) and the European Union (EU) both released ethical principles of AI (Independent High-Level Expert Group on Artificial Intelligence, 2019; Institute of Electrical and Electronics Engineers [IEEE] Global Initiative on Ethics of Autonomous and Intelligent Systems, 2019). A recent study found that even though governments had sought to achieve algorithmic accountability in the public sector through various outlets, including principles and guidelines, prohibitions and moratoria, public transparency, impact assessments, audits and regulatory inspection, external/independent oversight bodies, rights to hearing and appeal, and procurement conditions, evidence about the impact and effectiveness of algorithmic accountability policies was currently limited (Ada Lovelace Institute et al., 2021). Latzer and Festic (2019) emphasized that “accounting for the case-specific scope and context of algorithmic selections is . . . highly relevant for appropriate policy conclusions” because, for example, identical algorithms applied for recommending books versus recommending medical treatments “call for very different policies due to the disparity of risks of these automated algorithmic selections” (p. 6). Some

technology companies have created ethics groups, such as Microsoft’s research group called FATE which stands for Fairness, Accountability, Transparency, and Ethics in AI (Microsoft, 2021). Unfortunately, the concept of “fairness” is difficult to quantify in the context of algorithms and means different things to different people (Mitchell et al., 2021). There is also often a tension or trade-off between accuracy and fairness (Pessach & Shmueli, 2020).

Collaboration with and input from all stakeholders is a key element in fairness, as illustrated by new frameworks such as the “Accountability Agency Framework,” intended to facilitate a critical discourse through collaboration between different stakeholders of the system including domain experts, advocacy groups, academics, and policy experts (Cech, 2021). Another study that involved 48 practitioners across a variety of AI domains found that organizational culture was often a barrier to AI fairness efforts with its focus on a fast-paced development and deployment cycle (Madaio et al., 2020). The practitioners co-designed a new “AI Fairness Checklist” with the researchers that moved beyond a binary yes/no checklist approach to act as a framework for team members at different decision points to check for AI fairness within the entire lifecycle of a product (i.e., collaborating with stakeholders during the envision, prototype, build, launch, and evolve phases) (Madaio et al., 2020).

Filter Bubbles in News Feeds

Algorithms in social media news feeds determine what news reaches which individuals, and they are primarily concerned with relevance and giving users what they want based on their personal interests rather than objectivity. Algorithms highlight sources “based on engagement potential rather than quality” (Head et al., 2020, p. 19). Companies like Facebook give advertisers the tools to micro-target by collecting data on attributes like ethnicity, income,

political orientation, and hobbies; they enable extremists to locate and foster potential allies (Head et al.). Eli Pariser (2011) coined the phrase “filter bubble,” and he warned of the dangers of having your own personal, unique universe of information where you do not see what information gets included or edited out, or what criteria determines this. Filtering algorithms are editorial in nature rather than objective, and largely driven by commercialization and corporate bias. Corporate bias is when an information agency is biased towards the interests of its ownership or financial backing such as an advertiser rather than the user (Bobkowski & Younger, 2018). For example, if you were a liberal, you could expect to see only progressive links that you will already agree with, which would serve to confirm your pre-existing beliefs but limit your exposure to any other viewpoints. Bobkowski and Younger (2018) posited that this leads to several forms of cognitive bias for the user. Cognitive bias itself is an error in judgment as the result of your own implicit (unconscious) or explicit (conscious) bias. The filter bubble phenomenon can lead to the “hostile media effect,” or the tendency of those with strong opinions or beliefs to assume that the mass media is against them, in favor of the counter point of view. It can also lead to the “Dunning Kruger effect,” or the tendency of those with low ability or knowledge of a topic to overestimate their competency in that topic. And finally, it can lead to “confirmation bias,” when you only seek out and trust sources of information that confirm your opinions.

Without a good flow of information, this limits the democratic exchange of ideas and leaves everyone isolated, with a distorted sense of reality in which they overestimate the importance of certain ideas and underestimate others due to media bias caused by algorithms (Head et al., 2020). This accounts for some of the current political polarization in the United

States, because there is no shared frame of reference anymore. Vaidhyanathan (2018) argued that Facebook's "embedded" consultants played a large role in crafting Donald Trump's online advertising during the 2016 presidential race, which led the campaign toward inflammatory, visually striking messages that gained attention and were widely shared throughout the network. Facebook profited by selling more ads, and Trump profited by attracting more votes. Lewis and McCormick (2018) studied recommended videos on YouTube before the 2016 election and found that 643 out of 1000 political videos had an obvious bias, with 86% favoring Trump while only 14% favored Clinton. Further, more than 513,000 Twitter accounts had tweeted links to at least one video, and the most active 19 accounts were bots that cited the videos more than 1000 times. In another example, White nationalist Dylann Roof's murder of nine African Americans at a church in South Carolina was allegedly influenced by his Google search results for "Black on White crime" (Hersher, 2017), which retrieved sites from places like the White supremacist organization Council of Conservative Citizens (Noble, 2018).

Newman et al. (2020) noted that there was concern over social media and algorithms encouraging echo-chambers and pushing communities apart. Head et al. (2019) surveyed and did telephone interviews with students from 11 different colleges and universities about how they engaged with news and found that 71% used Facebook as a pathway to news weekly, 54% used YouTube, 51% used Instagram, and 42% used Twitter. Another study asserted that young adults spent significant time using social media platforms like Facebook and Twitter, and they were frequent consumers of social media news (Wang & Fussell, 2020). In a comparison of nine countries, the United States showed a preference for more partial news along with Spain, France, and Italy (Newman et al., 2020). To see the impact of the filter bubble, *The Wall Street Journal*

created a tool called “Blue Feed, Red Feed” where users could click on a topic such as “guns” or “healthcare” to see the difference in Facebook filter feeds between political affiliations (Keegan, 2016). Newsfeeds have largely replaced traditional news organizations, who are now reliant on being picked up by algorithms for an audience. Social media platforms have led to “the disaggregation of published information and its redistribution” through personalization, meaning we do not all see the same information and the original context is missing, making it difficult to evaluate (Head et al., 2020, p. 7). Reviglio and Agosti (2020) added that since personalization relies on the comparability of the user with others (through collaborative filtering), it paradoxically denied individual uniqueness. Personalization of media content limited the diversity of information, from exposure to discovery, which for the individual resulted in filter bubbles and for groups resulted in echo chambers (the formation of groups of like-minded users framing and reinforcing a shared narrative).

Fake News

Fake news consists of both misinformation (content that is mistakenly or inadvertently created) and disinformation (content that is deliberately created to deceive) (Affelt, 2019). Bots are “social media accounts that are operated entirely by computer programs and are designed to generate posts and engage with content on a particular platform” (Wardle, 2018, para. 6). In disinformation campaigns, bots can be used to draw attention to misleading narratives, to hijack platforms’ trending lists, or to create the deception of public interest and approval. Manufactured amplification can occur when the reach or spread of information is boosted through artificial means (either by humans or machines), such as the promotion of certain links or hashtags on social media (Wardle, 2018, para. 18). There are online price lists for different types of

amplification, including prices for generating fake votes and signatures in online polls and petitions, and the cost of lowering the ranking of certain content in search engine results (Wardle).

Vosoughi et al. (2018) investigated true and false news stories distributed on Twitter from 2006 to 2017. News was defined as any story or claim with an assertion to it, and rumors were defined as the social act of spreading news through the Twitter network. The researchers studied 126,000 stories tweeted by three million people and classified them as true or false using six independent fact-checking organizations. They found that false news circulated significantly farther, faster, deeper, and more broadly than true news, with the effects more pronounced for false political news. False news inspired replies expressing greater surprise and disgust, which they hypothesized was because false news was more novel. They also found that robots accelerated the spread of true and false news at the same rate. AI contributed to the problem of fake news being spread over social media like Twitter and Facebook because of how easy it was to manipulate visual imagery, thus enabling disinformation campaigns that could be targeted and personalized (Wardle, 2018, para. 3). The Defense Advanced Research Projects Agency (DARPA) developed a checklist of five things to look for in bot accounts (fake, automated accounts): (a) no photo or biography in the user profile; (b) formulaic, repetitive language and a lack of understanding of context; (c) single-mindedness and only posting about one topic repeatedly; (d) tweeting quickly, at odd times, or with a viewpoint that was inconsistent over time; and (e) following few accounts, but may be followed by a lot of other bots (often bots that have seemingly nothing in common) (Affelt, 2019). Also, there were plug-ins to mark social bots, although users were often dissatisfied with the results (Wang & Fussell, 2020). Another

thing to look for was a Twitter handle consisting of words or names, followed by random numbers (Affelt, 2019). In 2018, Twitter tried to cut down on the use of bots by issuing a statement saying the form of automation (including scheduling) to post identical or substantially similar content would not be permitted (Davies, 2019).

Friggeri et al. (2014) traced a population of rumors being inserted via photo uploads in Facebook, using a corpus of 4,761 rumors with 45% being false, 26% being true, and 29% being mixed or undetermined. The false rumors were revealed when someone posted a reply comment with a link to Snopes, a website documenting memes and urban legends (Snopes Media Group Inc., 2022; www.snopes.com). This research revealed that false rumors were frequently loaded and frequently called out, and that the majority of reshares happened after the item was called out as being false in a comment. This implied that the people sharing either did not notice the comment/call out, or intentionally ignored it. Lewandowsky et al. (2012) similarly discussed examples of the persistence of mistaken beliefs in false rumors despite extensive corrective efforts and evidence to the contrary. For example, the rumor that Barrack Obama had been born outside the United States, and the rumor of a relationship between a common childhood vaccine for the measles and autism. The spread of the latter rumor online contributed to measles outbreaks in 2018 in several major U.S. cities (Affelt, 2019).

Misinformation is also often facilitated by sources like YouTube videos; Lewandowsky et al. (2012) stated that YouTube videos often included information about health issues that contradicted official reference material. “Deepfakes” is the term “currently being used to describe fabricated media produced using artificial intelligence” (Wardle, 2018, para. 10). By integrating different elements of existing multimedia files, AI enables ways to create new

deceptive content. The likelihood that people will pass on information is based strongly on the likelihood of its eliciting an emotional response in the recipient, rather than how true it is (Lewandowsky et al., 2012). Also, evidence that threatens a person's worldview can strengthen initially held beliefs. Therefore, disinformation that is designed to provoke an emotional reaction can prosper in online spaces when algorithms detect that a user is more likely to engage with or react to similar content (Wardle, 2018, para. 1). Lewandowsky et al. (2012) suggest that framing evidence in a way that endorsed the same values as the audience and trying to foster healthy skepticism about information sources were good strategies to fight against fake news.

Several tools to help identify fake news are freely available (Affelt, 2019). One tool for analyzing Amazon consumer reviews is [reviewmeta.com](https://www.reviewmeta.com), which leverages algorithms and data science to look for unnatural patterns (Clark, 2019). A tool for analyzing healthcare news is [HealthNewsReview.org](https://www.healthnewsreview.org), which assesses the quality of the stories and news releases using a standardized rating system. There are several fact-checking sites for evaluating information including Snopes (Snopes Media Group Inc., 2022), [politifact.com](https://www.politifact.com) (Poynter Institute, n.d.), and [factcheck.org](https://www.factcheck.org) (Annenberg Public Policy Center of The University of Pennsylvania, 2022). A Washington, DC-based organization, News Literacy Project (NLP), created the checklist "Ten Questions for Fake News Detection" as well as a virtual platform called Checkology (<https://get.checkology.org/>) with different lessons; one lesson features social-media algorithms (News Literacy Project, 2022). Also, there are tools for detecting whether social media accounts are bots, such as Botometer (Yang et al., 2019), and whether images on social media were manipulated, such as TinEye (<https://tineye.com/>). Another teaching strategy is to have students practice interpreting advertising and digital propaganda using frameworks such as Herman and

Chomsky's (1988) five-filter model of news as propaganda, then make the connection for how biased search results from algorithmic personalization can serve as a type of coercive propaganda if people are unaware of how the filtering has shaped their results.

User Studies on Algorithms

There is a growing body of literature that examines users' understanding and experiences of algorithms in everyday life. There was variation in people's level of algorithmic awareness as well as behavior, attitudes, and perceptions towards algorithms across research studies in the domains of social media, curated news, search, and others. There was also variation in the way that algorithmic knowledge, attitudes, perceptions, and behaviors were measured.

Social Media User Studies

Awareness

In 2015, Eslami et al. performed a groundbreaking three-phase mixed methods study examining 40 Facebook users' awareness and perceptions of the Facebook News Feed curation algorithm. Phase one consisted of a pre-test questionnaire and interview to measure algorithmic awareness, where they discovered 62.5% were unaware of the News Feed curation. In phase two, participants used the FeedVis application (a Facebook application designed for the study but not publicly available) to visualize omitted material from their news feeds and view rankings of how often certain friends' stories were displayed while being interviewed; and in phase three participants answered email questions between two and six months later to evaluate the consequences of any insight gained by observing the algorithm outputs. The study shed light on the fact that some users, through misunderstanding filters, wrongly attributed the curation of their feeds to the habits or intents of their friends and family and drew false conclusions about

their interpersonal relationships based on the recommendations. Also, it showed the promise of building greater awareness of algorithms in order to give users more agency over their feeds because the educational intervention led to changes in behavior for 83% of the participants.

Similarly, Alvarado and Waern (2018) found that greater algorithmic awareness from user workshops led to changes in the usage of Facebook to cater more to user interests. On the other hand, Rader (2017) found much greater user awareness of the Facebook news feed filter in a survey than did Eslami et al. (2015): Seventy-three percent of respondents did not think they were seeing every post in their Facebook News Feed created by their Facebook friends. Furthermore, 36% felt Facebook used computer programs or algorithms to automatically choose what stories to show them and another 50% felt “maybe” this was true. Lu (2020) also found that most survey respondents did know that their news feeds were filtered, but still only 30% proactively curated their news feed. Socioeconomic factors positively associated with curation were age (younger), gender (female), race (White), education (highly educated), and political affiliation (left leaning). Income had no association. Perceived user controllability, Facebook use, and perceived knowledge were all predictors of proactive curation. Gran et al. (2021) did a survey in Norway to gauge user algorithm awareness on YouTube. Sixty-one percent of the respondents had either no or low algorithm awareness. The researchers classified users as having six levels of algorithmic awareness: unaware; uncertain; affirmative; neutral; sceptic; and critical. The critical type was the most algorithmically aware. They found a significant negative correlation between awareness and age, and positive correlations between awareness and education, males, and living in urban locations.

Proferes (2017) conducted a survey to gauge Twitter user knowledge and discovered many users were unaware of the ways Twitter tracked users across the Web and that Twitter sold access of real-time tweets to third parties. For example, a majority were unsure whether Twitter tailored the advertisements they saw based on information collected from third parties, and a majority were unsure whether or not Twitter received user browsing behavior information on third-party websites that had Twitter buttons or widgets if the user did not interact with the button or widget. Other findings revealed misunderstanding and uncertainty regarding how defaults on the site were set, especially for Direct Messages and whether it captured Global Positioning System (GPS) information in Tweets by default. Also, there was confusion over what information was public on Twitter and how protecting an account changed information availability. Warshaw et al. (2016) tested user knowledge about the data collection and inferencing capabilities of social media companies and found a gap between beliefs in what companies were doing with their data and algorithmic practices. Users fell into two categories: a “market research” cluster who believed companies collected data by asking users directly for their personal information and then made common sense inferences, and a “data mining” cluster who believed companies collected users’ online behavior data and made recommendations based on a user’s past behavior. The market research cluster were more ethnically diverse and had a lower sense of personal agency in the data economy, while the data mining cluster had higher income. The researchers speculated that there were previous histories of being a victim of stereotyping and feeling powerless that marginalized groups or those with lower income brought that contributed to their sense of low agency.

Fouquaert and Mechant (2021) found the Instawareness intervention tool helped raise algorithm awareness, although it did not change user feelings or critical concerns towards Instagram. They also found support for the Algorithm Paradox: “When people know the existence of curation algorithms, they claim to be bothered by them while not acting accordingly (e.g., by altering settings, changing interaction modes or visiting other profiles)” (p. 17). Eslami et al. (2019) found that 87% of interviewees lacked awareness of the Yelp review filtering algorithm, and upon disclosure theories for the filtering included how active the user was on Yelp, number of friends, level of completeness of the profile, whether reviews received feedback, extremity in viewpoint, level of detail, and length. Transparency resulted in changes in attitude such as wanting to write reviews differently or wanting to leave the Yelp system entirely (Eslami et al., 2019).

Folk Theories

Some user studies focused on folk theories, or qualitative studies exploring user ideas and mental models about how algorithms worked. Hamilton et al. (2014) declared in a 2014 paper that “mental models are now often thought inevitable and are evaluated in terms of utility for the user, rather than their verisimilitude” (p. 638). DeVito et al. (2017) defined folk theories as “intuitive, informal theories that individuals develop to explain the outcomes, effects, or consequences of technological systems, which guide reactions to and behavior towards said systems” (p. 3165). Rader and Slaker (2017) noted that folk theories explained cause and effect and were developed out of people’s everyday experiences with technologies. Bucher (2017) defined this as the “algorithmic imaginary,” or “the way in which people imagine, perceive and experience algorithms” (p. 31). In 2016, Eslami et al. explored user “folk theories” of automated

curation in Facebook news feeds and found the top theories attributed curation to the amount of personal interaction with a friend, the format of a post (text versus image), settings, how much friends post versus don't post, Facebook being all-powerful and perceptive, and their own similarity to a friend. DeVito et al. (2017) explored folk theories of Twitter users in reaction to Twitter's switch to a timeline that was algorithmically curated; resistance was a result of expectation violations, and more negative reactions resulted in more detailed explanations or folk theories. Bucher (2017) found that Facebook users had a wide variety of feelings about Facebook algorithms, including a feeling of being classified and profiled by algorithms in ways that didn't always identify with their own sense of self; whoa moments when people became aware of being "found" by an algorithm; annoyance over faulty predictions or failure to understand a human context (e.g., showing a year in review memory of someone deceased); and strategizing to have posts gain greater popularity such as posting at a certain time or structuring the post in specific ways. Also, like Eslami et al. (2016) found, some users felt algorithms ruined friendships by making certain people disappear from view and sometimes suddenly reappear.

Dogruel (2021a) found support among German internet users for four "folk theories" that were operationalized through statements related to economic orientation theory (platforms' economic interests), personal interaction theory (personal interaction on social network sites feeds curation), popularity theory (the general popularity of content impacts its ranking), and categorization theory (algorithms assign users to certain categories). DeVito (2021) studied the folk theories of 25 social media users and found that their ability and willingness to adapt to platform changes depended on the amount of information they had (including both information gleaned from outside sources and their own experimentation), along with their perception of

whether the core purpose of the platform was violated by the change. Haider and Sundin (2021) interviewed 17–19-year-olds in Sweden and Denmark and found users' engagement with information intermediaries was shaped partly by their previous experiences with how systems responded, which in turn shaped their anticipation of future outcomes and shaped their strategies to anticipate traps. Also, when content diverged from what was expected, this opened up new ways to engage with the infrastructure such as gaming the algorithm (e.g., watching different genres of videos, liking different posts, or adjusting searches in anticipation of feeding it different data to get better results).

Behavior and Attitudes

Min (2019) implemented a survey and discovered four categories of users in terms of the extent to which they practiced algorithmic engagement with social media filtering: (a) the disengaged; (b) the negative curator who block and unfollow; (c) the positive curator who add and follow; and (d) the activist. Min (2019) also found similar correlations to Lu (2020): a correlation between disengagement and increased age, less internet skills, less education, a smaller social media network size, less social media use, and a lower level of political efficacy. There were also correlations between females and the negative curator category, and Blacks and both the positive curator category and the activist category. Burrell et al. (2019) analyzed Tweets and discovered three strategies by users attempting to exert control over Twitter algorithms: personalization against platform defaults, working around platform errors and limitations (e.g., recognizing context and nuance), and virtual assembly work (e.g., configuring inclusion and exclusion signals such as shared blocklists).

Dogruel (2019) studied Facebook user attitudes towards online behavioral advertising and found that presenting users with interpretable and concrete information that wasn't too detailed about why an ad was algorithmically selected led to greater acceptance. Also, some user distrust in advertising did not seem to spill over to influence their trust of the platform as the carrier of the message. Eslami et al. (2018) discovered that users wanted clear ad explanations that were non-creepy, and a personal approach that linked to their identity. Although there was belief in algorithmic authority, seeing incorrect inferences led to algorithmic disillusionment. A common theme in these studies was advocating for more transparency and disclosure of how algorithms worked. Eslami et al. (2016) suggested including seams or visible hints disclosing certain aspects of automation operations, while Alvarado and Waern (2018) suggested presenting more transparency in Facebook profiling, greater ability to manage profiles, more user control and more ability to turn things on or off or report negative feedback, and the ability to un-follow memories. Eslami et al. (2018) suggested increasing the visibility of users' algorithmic self or integrating more explanations into advertisements.

Rao et al. (2016) found mismatches between user expectations of privacy and services' actual privacy practices, and they suggested a need for simplified privacy notices and more transparency in areas such as the collection of contact information and financial information without an account, collection of location information, and data deletion options. Rader et al. (2020) studied user reactions to their personal ad settings and preferences on Facebook and Google, and most users found the inferences made about them either plausible, implausible but rationalized (i.e., related to family members' activities, outdated, or aggregated generic demographic data), or completely irrelevant. In a study that went beyond social media, Human

and Cech (2021) found evidence of “dark patterns” (e.g., user interfaces that tricked users into doing things and do not have user best interests in mind) such as asymmetric design, hidden information, and unclear statements from a cognitive walkthrough by web designers of cookie and privacy consent mechanisms on the websites of Google, Amazon, Facebook, Apple, and Microsoft.

Curated News User Studies

Several studies looked at user attitudes towards algorithms that curated news more broadly, not just within social media. Fletcher and Nielsen (2018) found the majority of media users frequently failed to understand exactly how the information they received was filtered, but they also did not uncritically accept it because they were skeptical of all forms of selection—including both editorial and algorithmic. Those with interest in soft news topics like entertainment, arts, and sports were more likely to approve of algorithmic selection on the basis of what their friends had consumed. Approval for algorithmic selection was stronger amongst younger people across two studies (Fletcher & Nielsen, 2018; Thurman et al., 2019). Shin (2020) found that satisfaction and trust with news recommendation systems were influenced by perceptions of transparency and accuracy. Similarly, Shin and Park (2019) found higher trust as well as perceptions of fairness, accountability, and transparency led to greater satisfaction with algorithms. Thurman et al. (2019) found in an international survey of 26 countries that algorithmic news selection based on a user’s past consumption behavior was considered slightly better than selection by editors and journalists or selection based on a friend’s past consumption. Similar to this, Logg et al. (2019) found in a more general study that people used algorithmic advice more than human advice, despite blindness to algorithm’s processes. Thurman et al.

(2019) found as trust in news organizations and media's political independence fell, people were less likely to agree that selection by editors was a good way to get news. As trust in media's commercial independence fell, privacy concerns increased, concerns about missing important information as a result of more personalized news increased, or concerns about missing challenging viewpoints increased—people were less likely to agree that any form of news selection was a good thing. To an extent, users did not associate the operation of automated news personalization with the operation of news organizations, believing that the technology was immune from untrustworthy news media.

Monzer et al. (2020) identified through a study of focus groups in Germany tension between journalistic integrity and the increasing revenue of personalization algorithms. Users perceived the benefits to be usefulness from higher relevance and the saving of time, but had concerns over being stereotyped, manipulation, data protection, filter bubbles, and untrustworthy news organizations. User strategies for more agency included deleting cookies and avoidance, and they desired more transparency to create more trust, more well-known brands, more user control over settings, and feedback opportunities. Heuer and Breiter (2020) found that when users in Germany rated news stories on an 11-point trust scale, they were mostly able to distinguish trustworthy recommendations of quality news from untrustworthy when the stories were presented alone. However, untrustworthy news stories benefited from appearing in a trustworthy context (e.g., when one untrustworthy story was combined with four trustworthy stories). Siles et al. (2020) explored folk theories about Spotify and found that users linked their experience with other algorithms to their folk theories of how Spotify's algorithm worked; for example, they presumed that the Netflix and Facebook algorithms shared a common logic with

the Spotify algorithms and they compared the features used by competitors such as Apple Music to Spotify.

Search Engine User Studies

Schmidt et al. (2019) found that most users of Google search were, to a certain extent, aware of the existence of and the (approximate) mechanisms of algorithmic personalization, but not all of them connected prioritization and personalization directly with their own user behavior (almost none of them noticed patterns of personalization based on previous search queries). Some, but not all, were concerned with data collection. Thomas et al. (2019) found from surveys and interviews that users mentioned the following concepts to explain search engine results the most often: popularity, wording, commercial interests, and personalization. They did not mention authority, recency, type of result, SEO, or diversity very often. Bakke (2020) found in a think-aloud usability study that reputation was a criterion for selecting Google search results, and that Google's rankings influenced what users chose more than they were aware of. Cotter and Reisdorf (2020) found that education and age were correlated with algorithmic knowledge, as well as frequency and breadth of use of online search. Income was also correlated with frequency and breadth of use of online search. Gender and race/ethnicity were not associated.

User Studies in Other Domains

Costello and Floegel (2020) interviewed users of mental health apps owned by technology companies, and users described a feeling of compromise when using apps for mood tracking and mental health diagnostics that did not fully meet their needs. App users raised concerns about the profit motives behind predictive advertising and did not trust the app with their health data, yet they felt resigned to the ubiquity of big tech companies. Rader and Slaker

(2017) found the folk theories of 30 activity tracker users were based on both information provided in the interface as well as from users' own perceptions of their activities. They suggested that interfaces encouraging folk theories about how sensor data were produced and used may encourage more informed privacy self-management choices. Hargittai et al. (2020) interviewed adult internet users across five countries about algorithms used in the areas of voice assistants, products and services, video sites, and maps/GPS tracking and discussed the results for awareness, understanding of how algorithmic systems processed information, and attitudes.

Dogruel et al. (2020) interviewed 30 German internet users and found users defined algorithms according to experience-based definitions, technical definitions, or were uninformed, and they learned about algorithms through education-related and work-related contexts, news media, or fictional media. Users had greater awareness of algorithmic decision-making in the areas of advertising, online shopping, and media use, but lesser awareness in the areas of job searching, online dating, news, and navigation systems. Furthermore, the factors of personalization and subjective control (e.g., users who interacted with undesired content such as advertisements they felt they did not actively control or solicit) influenced user perception of how algorithms impacted their internet use, and users perceived algorithms' impact on their autonomy as either enhancing, restricting, or challenging (Dogruel et al., 2020). Shin et al. (2021) studied algorithmic literacy in the context of personalized media recommender systems (such as Netflix and Amazon) and concluded that perceived fairness, accountability, transparency, and explainability all led to an increase in user trust, which led to a willingness to share additional data and resulted in more accurate and personalized data, which in turn led to a positive attitude and positive intention to keep using the service. Algorithmic literacy was

defined in the study as an understanding of how algorithms curated and processed information, recommended social connections, and reconstructed realities and were expressions of broader systems of power.

Algorithmic Bias User Studies

Woodruff et al. (2018) looked at how traditionally marginalized populations perceived algorithmic fairness through workshops and interviews. They found that the concept of algorithmic fairness was largely unfamiliar but learning about it elicited some negative feelings and had the potential to impact trust in a product or company. However, participants perceived algorithmic systems as small in scope, and they most often attributed algorithmic unfairness to either (a) a non-diverse population of programmers; (b) prejudiced online behavior by members of society; or (c) the news media. Plane et al. (2017) presented users with discriminatory online advertising scenarios, and users found it problematic when discrimination occurred as a result of explicit demographic targeting by race rather than in response to online behavior. However, they did not care whether a human or an algorithm was responsible for the discrimination. Araujo et al. (2020) found domain-specific knowledge, belief in equality, and online self-efficacy were associated with more positive attitudes about the usefulness, fairness, and risk of decisions made by AI, but increased levels of privacy concerns had a negative association. Grgić-Hlača et al. (2020) measured how people perceived the unfairness of using different features (e.g., personality or criminal history of friends and family) to describe a defendant to algorithmically predict the defendant's risk of engaging in criminal activity in the near future but found mixed results. Many additional user studies looking at user perceptions of algorithmic fairness were described in the systematic review by Starke et al. (2021). Register and Ko (2020) measured

users' ability to self-advocate against harmful machine learning models, presented in scenarios, by analyzing their critiques in hypothetical letters to the enforcer of the model. Analysis was done by looking for the presence or absence of coded criteria such as construct validity, outliers, causality, ways the model could be gamed, or consequences.

User Studies on College Students

A handful of user studies on algorithms looked specifically at college students. In one of the first, Powers (2017) conducted semi-structured interviews and a survey of college students to determine their awareness of news personalization and how the Facebook and Google platforms selected and prioritized news. Students from the interviews most often began a news search at Google/Google News, Yahoo/Yahoo News, CNN, Facebook, and Twitter. When using news portals, or solely aggregated content, most students knew their digital media habits were tracked but not how the data was collected. They also did not know personalization was based on past interactions and preferences. When using news outlets (where there was some original content), students were less aware that their digital media habits were tracked to personalize their news and less aware of the impact of other web users on their news personalization. Only 24% of the students taking the survey were aware that Facebook prioritized certain posts and hid others (Powers, 2017, p. 1325). Most students knew they could adjust their preferences and to unfollow people to hide posts, but few were aware that actions by users they did not follow and the actions of Facebook engineers, editors, or curators influenced their news selection and ranking. In terms of Google search, only a minority said they would likely get different search results from someone else on Google. Few knew they could make adjustments that would affect the news they saw, that actions taken by other users impacted their results, that details about their online

session impacted the results, or that actions taken by Google engineers, editors, and curators impacted their results. Powers concluded “college students’ lack of awareness of specific gatekeeping functions that Google and Facebook employ keeps them from fully evaluating the potential introduction of bias into news selection and prioritization, including whose interests are being prioritized by algorithms and human editors” (p. 1330).

Later studies on college students found mixed results. Brodsky et al. (2020) found in a survey that only 51.4% of students were algorithm aware in terms of Facebook not showing them all the stories friends posted, and that awareness was associated with higher algorithm engagement but not associated with usage or media literacy (Brodsky et al., 2020, p. 47). Students indicated more awareness for how online shopping worked than for how online search worked, and there was only a positive association between media literacy and online shopping algorithmic awareness (not online search algorithmic awareness). Watching an intervention video about algorithms did increase knowledge of how online searches worked, but many students did not transfer algorithmic awareness across platforms not covered in the video. On the other hand, Head et al. (2020) found through focus groups that students were somewhat aware of how algorithms influenced their online experiences, such as collecting data to target advertisements and personalize, even if they couldn’t explain how they worked. Students were ambivalent and skeptical about algorithm-driven platforms and used defensive practices to protect their privacy such as using ad blockers, clearing cookies, using Firefox instead of Chrome, using a virtual private network (VPN,) and having multiple user accounts for different interests to exercise digital code switching. Science, technology, engineering, and mathematics (STEM) majors or students with roommates who were STEM majors used the most strategies.

Students' top concerns about computer algorithms were that platforms were "listening" across devices, that algorithms and automated decision-making reinforced inequalities, that platforms were shaping individual content and ads they saw, and that online users were not seeing the same reality. Despite a general lack of trust in traditional authority figures, there was still trust placed in Google as "the arbiter of truth" (Head et al., 2020, p. 2).

Similarly, Koenig (2020) found through a content analysis of student journals that students across a wide variety of ages did have a basic awareness of algorithms (how algorithms worked), but they did not always engage critically (why algorithms functioned as they do and the embedded values) or rhetorically (recognition of the interconnected components of humans and algorithms through rhetorical moves to influence the algorithm) in the platform's inputs or outputs. After students were asked to reflect through journal writing, the act of reflection led to more critical and rhetorical awareness of the algorithm's influence (Koenig, 2019, 2020). Wang and Fussel (2020) found through interviews that college students considered the news source, personal relevance, and personal interest as well as the entertainment value, helpfulness to others, and their own reputation when deciding to share or fact-check on Facebook or Twitter. Swart (2021) found young people aged 16-26 made sense of news curation on social media through the specific platform, comparing features, comparing type of content, looking at the experiences of other users, and popular media reporting. Their emotions about algorithms ranged from feeling they were rational, useful, recognizing their commercial nature, and feeling censored. They were hesitant of algorithmically curated news due to fears of missing out, concerns around surveillance, and because they wanted elements of surprise in journalism. They felt their own role in shaping algorithms was limited, and they were passive with engaging in

news. Even though they did have some knowledge around algorithms, this did not necessarily impact their behavior. Pierson (2018) surveyed undergraduate computer science students on ethical dilemmas with algorithms and found gender differences in beliefs about algorithmic fairness, as well as evidence that views on algorithmic fairness could be changed by discussion.

Measurements

Knowledge and Awareness

Measuring algorithmic knowledge and skills is challenging because often there is “no ground truth for knowing how algorithms actually function in many cases” and what correct responses would be to questions (Hargittai et al., 2000, p. 767). Nonetheless, researchers have attempted to measure this. A fixed choice approach was found most often in the literature when researchers captured algorithmic knowledge or awareness. Eslami et al. (2015), Rader (2017), and Powers (2017) measured algorithmic awareness of the Facebook News Feed filter through multiple choice true-false-uncertain scales. Brodsky et al. (2020) measured Facebook algorithm awareness and Facebook news feed adjustment with a series of yes-no questions adapted from Eslami et al. (2015). Powers (2017) measured algorithmic awareness of the Google search engine by asking multiple choice yes-no-uncertain questions about whether two people would likely get the same results if they entered the same search terms for news at the same time on Google; whether Google News allowed users to make adjustments that affected the news they saw and in what order; and whether five types of actions affected either the type of news participants would find on Google or the order in which results are shown. Proferes (2017) tested respondents’ beliefs about Twitter by asking them if they believed various statements were accurate-inaccurate-unsure about the data/metadata, algorithms, protocols, interfaces, defaults,

users, and informational content. Also, statements were presented about Twitter’s socioeconomic facets that users identified as accurate or inaccurate. Latzer et al. (2020) measured Swiss internet users’ awareness of algorithmic selection and related risks through a series of true or false statements, and they measured attitudes about algorithms by asking to what extent users agreed or disagreed with statements about algorithmic selection. Zhu et al. (2020) used a five-point Likert scale ranging from Strongly Disagree (1) to 5 Strongly Agree (5) to measure social media competence in college students.

Dogrueel et al. (2021) validated the “Algorithm Literacy Scale for Internet Users” that measured awareness of algorithms (to what extent internet users were aware of the areas in which algorithms were used and in which particular applications or devices) by asking binary questions (it is used, it is not used, don’t know) about whether various media products such as computer games and various media functions such as creating weather forecasts use algorithms to function. The same scale also measured knowledge about algorithms (i.e., how algorithmic systems on the internet personalize and customize information based on information collected, what type of information gets processed, and how this impacts the content) through one multiple choice question and a series of true/false/don’t know questions (Dogruel et al., 2021). The same study also developed three scenarios with more concrete examples of algorithms in use asked as true/false questions to measure algorithmic-specific decision-making. A combination of Likert scale and multiple-choice questions were used to measure Canadians’ familiarity with artificial intelligence and its capabilities (Government of Canada, 2021). Similarly, multiple choice questions were used to gauge Americans’ knowledge of everyday technologies that use artificial intelligence and machine learning (Zhang & Dafoe, 2019). Kong et al. (2021) measured AI

literacy concepts through a multiple-choice test with four choices per question; for example, “Which algorithm for supervised learning involves the concept of ‘birds of a feather flocking together’”? with the correct answer being k-nearest neighbors was a sample question.

Some studies defined the assessment of algorithmic awareness or knowledge on a continuum, ranging from basic awareness to a more sophisticated understanding. One such continuum defined basic awareness for students as knowing that an algorithm was a complex, computational procedure (Koenig, 2019, 2020). This included an understanding of how algorithms could affect their education and research goals, the social conventions that algorithmic outputs created, specialized discourse associated with algorithms, a basic technological understanding of how algorithms functioned, and having technological curiosity (Koenig, 2019, 2020). Koenig defined higher levels of algorithmic awareness as “critical awareness” of why algorithms function the way they do, including the dominant perspective that shapes algorithmic design cultures, seeing use contexts as an inseparable aspect of computers, understanding the institutional forces that shape algorithmic outputs and the influence of the outputs, and scrutinizing popular representations of algorithms. Also, Koenig defined higher levels of awareness as “rhetorical awareness” of the interrelationship between people and algorithms, meaning students actively made rhetorical moves that enacted change in algorithmic outputs including persuasion, complexity, reflection or critical assessment, and agency. Another study defined critical algorithmic literacy as contextual, heterogeneous, always partial, and inseparable from questions of power (Cotter, 2020). Specifically, Cotter (2020) described two overlapping types of knowledge about social media algorithms—“technical knowledge,” which included design knowledge of the algorithm’s goals and methods knowledge of how the

algorithm would reach those goals, and “practical knowledge,” which generated tactics true to one’s sense of self within the social world while acknowledging the regulatory role of algorithms. DeVito (2021) developed four levels at which users theorized about algorithmically-driven social platforms in the context of self-presentation; the lower “functional” levels consisted of basic awareness (awareness of the presence of an algorithmic process) and causal powers (awareness that the algorithmic process has some specific causal effect), while the higher “structural” levels consisted of mechanistic fragments (awareness there are multiple criteria used during the algorithmic process to make decisions) and mechanistic ordering (awareness of not only multiple criteria used to make decisions but also causal ordering within this criteria).

Attitudes, Perceptions, and Behaviors

Most studies measuring attitudes or self-reported perceptions used a Likert scale. Al-Zahrani (2015) used a five-point Likert scale ranging from Strongly Agree (5) to Strongly Disagree (1) to measure student experience, use, expertise, attitudes, and self-efficacy that contributed to their level of digital citizenship. Similarly, Takavarasha et al. (2018) measured quantifiable aspects of digital citizenship using a four-point Likert scale ranging from agree strongly to disagree strongly. The researchers excluded the neutral option “in order to discourage students from using it to avoid answering survey questions” (Takavarasha et al., 2018, p. 7). Shin et al. (2020) used a seven-point Likert scale ranging from Strongly Disagree (1) to Strongly Agree (7) to measure “algorithm acceptance” or more specifically the variables of fairness, accountability, transparency, personalization, accuracy, trust, use behavior, and attitude. Shin (2021) added measurements for the variables of explainability, credibility, and information seeking in the context of judging the credibility of chatbots, and Shin et al. (2021) added a

measurement for the variable of intention in the context of media recommender systems. Shin et al. (2022) added measurements for self-disclosure, privacy concerns, and self-efficacy in the context of personal privacy. Gran et al. (2021) measured self-reported algorithm awareness using a five-point Likert scale ranging from No Awareness (1) to Very High Awareness (5) as well as attitudes towards algorithmic functions using a five-point Likert scale ranging from Very Negative (1) to Very Positive (5). Also, Dogruel (2019) used a five-point Likert scale to measure trust in Facebook, Lu (2020) used a four-point Likert scale to measure perceived knowledge of Facebook, and Kong et al. (2021) used a five-point Likert scale asking for level of agreement to measure perceptions of levels of AI literacy and sense of empowerment to deal with AI (e.g., I understand various applications of AI in everyday life; I think of myself as someone who can do well in an AI literacy course). A systematic review on perceptions of algorithmic fairness across 39 studies found most measurements of fairness used Likert scales on vignettes or scenarios (Starke et al., 2021).

Likert scales were used to gauge Americans' perceptions of the impact of bias in various domains such as hiring, criminal justice, health, autonomous vehicles, surveillance, and fake news (Zhang & Dafoe, 2019). Zarouali et al. (2021) developed a 13-item "Algorithmic Media Content Awareness Scale (AMCA-scale)" to measure users' perception of their algorithmic awareness across the dimensions of content filtering, automated decision-making, human-algorithm interplay, and ethical considerations. A five-point rating scale was used, ranging from Not at All Aware (1) to Completely Aware (5) for various statements about algorithms in media content (Zarouali et al., 2021). The scale was validated in the Netherlands for the platforms of Facebook, YouTube, and Netflix. Min (2019) measured algorithmic engagement with social

media through a series of yes/no questions for the following dimensions: algorithmic awareness, algorithmic engagement with news consumption, and algorithmic engagement with news posting. Rao et al. (2016) used several different Likert scales to measure user behavior around privacy activities, and measured website privacy expectations using a series of true/false questions. Latzer and Festic (2019) suggested five measurements to assess algorithmic governance in everyday life: usage of automated algorithm selection applications, subjective significance assigned to the applications, user awareness, user awareness of related risks, and user practices to cope with risks.

Conclusion

This literature review established that college students needed to improve their algorithmic literacy skills, and that these skills were not being taught as part of the regular curriculum. It explored algorithmic literacy as a subset of information literacy by looking at the history of information literacy and the beginning steps taken to integrate algorithmic literacy into existing information literacy instruction. It illustrated that the *Framework for Information Literacy for Higher Education* (ACRL, 2016) could be expanded to more explicitly include algorithmic literacy. The review also covered the technical and social aspects of algorithms, and it looked at user studies related to algorithms measuring knowledge, attitudes, perceptions, and behaviors. Taken together, these elements of the literature review were explored to ascertain what knowledge components and behaviors might contribute to students' algorithmic literacy skills. Algorithmic literacy is still a growing field, and there is not yet agreement on how to define competencies. The goal of this study was to add to the literature on how algorithmic literacy could be taught to college students, specifically looking at the knowledge components

and coping behaviors that would help students interact with algorithmic systems more effectively. This information would allow educators to begin integrating this important, emerging field of study into the mainstream curriculum. The next chapter will present the methodology for the study.

CHAPTER 3

METHODOLOGY

As discussed in the literature review presented in the previous chapter, algorithmic literacy is a growing subset of information literacy, and it teaches critical awareness of what algorithms are, how they interact with human behavioral data in information systems, and social and ethical issues related to their use (Head et al., 2020). The literature review discussed the current *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries, 2016) and other frameworks related to algorithmic literacy, along with the technical, social, and ethical aspects of algorithms and established that it is not being taught in the mainstream college curriculum. User studies established that both the general public and college students need to improve their algorithmic awareness and algorithmic literacy skills, but more research is needed to determine what the key components of algorithmic literacy consist of for college students and how this could be taught in the curriculum.

This study expanded upon previous work by focusing on the experiences and opinions of scholars and teachers of critical algorithm studies or related fields. Research indicated there was little information available about how algorithmic awareness and algorithmic literacy were being taught to college students, including curriculum content and placement, teaching methods, coping strategies, and methods of assessment. As such, the purpose of this research was to explore these topics from the perspective of subject matter experts to determine to what extent there was consensus and how the information could inform expanding the *Framework for Information Literacy for Higher Education* (Association of College & Research Libraries, 2016) to include algorithmic literacy.

Research Questions

In order to learn more about the key components of algorithmic literacy and how this could be taught in the college curriculum, the following research questions guided this qualitative study:

1. What components of algorithmic literacy are specific to college students?
2. What behavior and knowledge contribute to students' algorithmic literacy?

Method

To answer the research questions, the researcher used a qualitative approach. This was appropriate because the purpose of qualitative research is to “gain insights into a particular phenomenon of interest” (Mills & Gay, 2019, p. 7). The primary qualitative research method chosen was semi-structured interviews with scholars and teachers of critical algorithm studies or related areas across a variety of fields of study. These involved “a set of predetermined topics” but also allowed for flexibility to tailor the questions to each interviewee, which is appropriate when the sample is diverse (Crano et al., 2014, p. 287). Also, the dialogue was allowed to evolve naturally beyond the existing structured questions (Ahlin, 2019). In semi-structured interviews, the interviewer asks participants a series of predetermined but open-ended questions (Ayres, 2008). When researchers lack substantial scholarly knowledge on a subject, they often choose semi-structured interviews with subject matter experts as the research method (Ahlin, 2019). This matched the scenario the researcher faced when trying to learn more about the emerging field of algorithmic literacy.

The secondary qualitative research method chosen was one focus group with scholars and teachers in critical algorithm studies or related fields. Focus groups are group interviews where

you try to collect a shared understanding of your research question by allowing participants to take turns answering questions, so they all give input (Mills & Gay, 2019). A focus group “involves a moderator who obtains information by encouraging verbal discussion and interaction about a focused topic among a group of participants” (Crano et al., 2014, p. 297). Focus groups are often used “in the formative stages in a new line of research or to obtain preliminary evidence to learn about an understudied topic” as well as “to establish the content validity of a construct” (Crano et al., 2014, p. 298). Focus groups were appropriate for this research because there is no agreed-upon operationalization of algorithmic literacy yet (Cotter & Reisdorf, 2020; Hargittai et al., 2020; Kitchin, 2017). Also, the researcher wanted to use a second data collection method to establish data triangulation by comparing data from both methods to enhance the trustworthiness of the findings. In the sections that follow, details are provided about the participants, setting, data collection, and analysis plan.

Participants

This study used purposive sampling to identify participants. Purposive sampling is defined as “selecting a sample that is believed to be representative of a given population” (Mills & Gay, 2019, p. 159) in order to best help understand the research problem. It is a useful data collection procedure when participants cannot be directly observed (Mills & Gay, 2019); this was the case in this study since critical algorithm scholars and teachers are a geographically dispersed population. A total of 39 critical algorithm studies scholars and teachers were identified through the literature review and solicited via email by the researcher to participate in the study (see Appendix A). Inclusion criteria consisted of the authoring of at least one major source of information from the literature review, with preference given to participants who also

teach college students, while balancing the need for interdisciplinary representation from scholars working in the social sciences, science and technology, and library and information science. Demographic characteristics such as gender, age, and ethnicity were not considered and did not play a role in the selection process. The original email provided a general description of the research project, the research objectives, the institution, how participants were identified, and an assurance that participants would have anonymity in the public dissemination of data as recommended by best practices (Crano et al., 2014). The subjects were told in the initial recruitment email that the study's purpose was to explore the components of algorithmic awareness relevant to college students, along with the corresponding behavior and knowledge that contributes to this awareness. Subjects were asked to participate in either a one-hour interview or a one-hour focus group videorecorded over Zoom and to give the researcher their availability if they wanted to participate. Once participation was confirmed, a copy of the Informed Consent Form was emailed to them in advance and collected at the time of their participation. They were also sent a Zoom meeting invitation by the researcher at the specified date and time.

Eleven of the subjects expressed a willingness to be interviewed, and the researcher selected all 11 of them as interview participants. Another six of the subjects expressed a willingness to participate in a focus group, and five of them were selected because it was possible to coordinate meeting at the same time. Twenty-three of the subjects either did not respond or indicated that they did not want to participate in the study. The researcher sent one follow-up email to all respondents who did not respond to the initial email. All subjects in the semi-structured interviews and focus groups were given pseudonyms to protect anonymity. The

semi-structured interview participants consisted of five females, four males, one who preferred not to say, and one non-binary third gender. Ten participants identified as White and one as other for race/ethnicity. Four participants were in the 25–34-year age range, five participants were in the 35–44-year age range, and two participants were in the 45–54-year age range. Nine participants were from the United States, one was from Ireland, and one was from Austria. For a complete list of demographics for interview participants, see Table 2.

Table 2

Interview Participant Demographic Data

Name	Gender	Race/Ethnicity	Age	Job Title	Country
Avery	F	White	35-44	Assistant Professor	USA
Diane	F	White	25-34	Research Scientist	USA
Evelyn	F	White	45-54	Associate Professor	USA
Foster	M	White	35-44	Assistant Professor/Research Assistant	Austria
Jess	Non-binary	White	25-34	PhD Candidate	USA
Kenzie	F	White	35-44	Assistant Professor	USA
Kris	Prefer not to say	White	34-44	Assistant Professor	USA
Mason	M	White	25-34	Postdoctoral Researcher	USA
Monica	F	Other	25-34	Assistant Professor	USA
Nolen	M	White	34-44	Assistant Professor	USA
Richard	M	White	45-54	Professor	Ireland

Note. Pseudonyms were given to all participants. The rest of the demographic data was provided by the participants.

The focus group participants consisted of three females and two males who all identified as White. One participant was in the 25–34-year age range, one participant was in the 35–44-year age range, one participant was in the 45–54-year age range, one participant was in the 55–64-year age range, and one participant was in the 65–74-year age range. Four participants were from the United States and one participant was from Canada. For a complete list of demographics for focus group participants, see Table 3.

Table 3*Focus Group Participant Demographic Data*

Name	Gender	Race/Ethnicity	Age	Job Title	Country
Becca	F	White	65-74	Scholar in Residence Information	USA
Charlotte	F	White	25-34	Literacy Coordinator	USA
Jordan	M	White	45-54	Lead for Research Informatics	USA
Madeline	F	White	55-64	Professor Emerita	Canada
Norman	M	White	35-44	Assistant Professor	USA

Note. Pseudonyms were given to all participants. The rest of the demographic data was provided by the participants.

Setting

Interviews and the focus group were conducted and recorded using the Zoom videoconferencing technology. Informed consent was obtained from all participants at the very start of each scheduled session by sending participants a link to the Informed Consent Form in Qualtrics (2021). All participants checked off “I agree to participate in this study” followed by typing their name. The next page of the Informed Consent Form asked for the voluntary demographics of gender identity, race/ethnicity, age, job, country, and highest level of school completed. Interviews ranged from 44 minutes to 78 minutes, with an average of 54 minutes. The focus group was 54 minutes.

Data Collection

Questions from the interview protocol were used to guide a discussion to address the research questions focused on what components of algorithmic literacy were specific were to college students and what behavior and knowledge contribute to students’ algorithmic literacy (see Appendix B). Specifically, nine open-ended questions were asked with additional prompts

as needed. The first question, “What led you to become actively involved in algorithm studies?”, was the most general and easily answered, as recommended in the literature, in order to build rapport (Crano et al., 2014). This was followed by three additional questions about the impact of algorithmic decision-making, its strengths and weaknesses, and key domains where it had the greatest impact. After that, the remaining questions focused more directly on how to teach algorithms to college students or adult learners, including “If you were creating a list of the key elements to capture the idea of ‘algorithmic awareness’ or ‘algorithmic literacy’ in adult learners or college students, what content areas would you identify as most important?”, “What teaching tools or teaching strategies can aid in learners’ understanding of algorithms?”, “What behaviors reflect high levels of algorithmic awareness?”, and “How can we measure or assess someone’s basic level of algorithmic awareness or algorithmic literacy?” For the complete list of questions, see Appendix B.

Questions from the focus group protocol were used to guide a group discussion to also address the research questions focused on the components of algorithmic literacy specific to college students and what behavior and knowledge contribute to students’ algorithmic literacy (see Appendix C). Specifically, five open-ended questions were asked with additional prompts as needed. Questions were open-ended to align with the best practice of asking focus groups open-ended questions to generate richer conversation (Cyr, 2019; Kamberelis & Dimitriadis, 2013). Also, “why” questions were not asked to avoid inadvertently putting participants on the defensive. All five questions were almost identical to questions asked during the semi-structured interviews, and those five questions were chosen in order to triangulate and confirm data from the interviews deemed important for answering this study’s research questions. Specifically,

participants were asked about and discussed key domains using algorithmic decision-making having a high impact, the key elements for algorithmic awareness or algorithmic literacy, teaching tools or teaching strategies to aid in learners' understanding of algorithms, what behaviors reflected high levels of algorithmic awareness, and how algorithmic awareness and algorithmic literacy could be measured. For the complete list of questions, see Appendix C.

Analysis Plan

The interviews and focus group were video and audio recorded and transcribed by Zoom. Then, a content analysis of the transcripts was conducted, where the researcher “systematically reviews qualitative unstructured data and classifies them according to themes, characteristics, and patterns considered to be meaningful in addressing research questions” (Crano et al., 2014, p. 303). First, a review of the literature allowed for deductive coding or a set of a priori codes; in deductive coding, the researcher determines the coding categories drawn from previous research prior to conducting the coding (Crano et al., 2014). Next, the transcripts were reviewed initially with notes written in the margins to get a general sense of the meaning and to allow for inductive coding or emergent coding. In inductive coding, the researcher allows coding categories to emerge from the data (Crano et al., 2014). The researcher’s final set of codes included both a priori and emergent codes and consisted of 39 codes and three child codes organized around three broad themes. The three broad themes were 1) Algorithmic Literacy Knowledge Components; (2) Student Behaviors and Coping Strategies; and (3) Pedagogy.

The coding software program Dedoose (Dedoose software tool version 9.0.46, 2021; www.dedoose.com/) was used to aid in coding. The audio recordings and the transcribed text were uploaded as files, and descriptor sets were created for each participant with demographic

information that linked to their audio and text files. All 39 codes and three child codes were set up in Dedoose. Then, the researcher listened to each interview or focus group while simultaneously reading the transcript, going line by line to determine what the relevant excerpts were, highlighted them, and applied any relevant codes in Dedoose. Often, the researcher applied multiple codes to an excerpt; the maximum number of codes applied to any excerpt was four. During the coding process, two codes were merged with existing codes because there was too much overlap. Specifically, the code “public attitude” was merged with the existing code “common misconceptions,” and the code “discrimination” was merged with the existing code “oppression, privilege, power.” Also, one new code was added as a teaching method for algorithms later after all the other codes had been finalized: “everyday life.” Finally, the original code “other personalization” was changed to “recreation and retail” after reviewing the interview transcripts. In order for a code or child code to remain, it had to be mentioned by at least three different participants. During this process, the researcher also edited excerpts to correct for errors in transcription made by Zoom to ensure that all quotes would reflect exactly what was stated by the participants. Although subjects were videotaped over Zoom, the recording was only used for research purposes. Only the principal investigator had access to the recording. Participant identity was only disclosed throughout the analysis in the context of a subject’s pseudonym and voluntary demographic information. Real names will never be used in any public dissemination of these data (publications, presentations, etc.). All research materials and consent forms were stored on a secure server and permanently deleted after this research project was completed. Confidentiality cannot be guaranteed in a focus group setting; however, the researcher asked all participants to respect other participant’s privacy and keep all information shared confidential.

List of Codes in Alphabetical Order Across All Themes

The finalized codes were as follows:

Accountability: This code referred to the process of assigning responsibility for harm when algorithmic decision-making resulted in discriminatory and inequitable outcomes, as well as regulatory/legal structures.

Algorithm strengths: This code captured helpful characteristics of algorithms and situations where algorithms were described as beneficial.

Algorithmic awareness: This code included discussion about the basic degree to which users know when algorithms are being used and what they can be used for. This is lower on the taxonomy of learning than algorithmic literacy.

Algorithmic literacy: This code included discussion about a higher level of learning than basic algorithmic awareness, including a more critical awareness of what algorithms are, how they interact with human behavioral data in information systems, and an understanding of the social and ethical issues related to their use.

Algorithmic reasoning versus human reasoning: This code pertained to the differences in how algorithms reason compared to humans, with an emphasis on the way computational thinking abstracts reality in order to make it fit within limited computational constructs.

Bias causes: This code addressed the idea that algorithmic decisions often produce biased results, usually due to either historical bias (human biases included in training datasets), incomplete or unrepresentative data, or proxies used for sensitive attributes such as race that become part of feedback loops.

Black box: This code alluded to the fact that an algorithmic system is mysterious because it can be viewed only through its inputs and outputs, not its internal process.

Broader advocacy: This blanket code referred to applying principles of social justice and applied ethics to the design, deployment, regulation, or ongoing use of algorithmic systems so that the potential for harm is reduced.

Challenges of assessment: This code outlined challenges in assessing algorithmic awareness or algorithmic knowledge, often stemming from the “black box,” secretive nature of the actual algorithmic working, the fact that algorithms are constantly changing, or the fact that they vary from platform to platform.

Commodification: This code pertained to the underlying financial incentives and motivations of companies using algorithms; for example, internet algorithms are often designed to maximize time spent online in order to generate more advertising revenue.

Common misconceptions: This code included popular stereotypes or misunderstandings of AI and algorithms, such as being neutral or superior to humans.

Contextualized stories: This code dealt with contextualizing stories of algorithmic bias in order to make them more human.

Critical evaluation: This code referred to actively questioning algorithmic design, decisions, and outputs and seeking alternatives or other perspectives if applicable.

Data transparency tools: This code alluded to the incorporation of information into user interfaces to visibly show how algorithms work under the hood.

Education: This domain code referred to algorithms used in schools, colleges, or universities.

Everyday life: This code pertained to helping users recognize the algorithms they use and encounter often in daily life.

Expectancy violations: This code alluded to situations when a user is surprised by algorithmic output, often leading to good insight in how they work.

Finance: This domain code included algorithms used in banking, loans, and credit.

Folk theories: This code pertained to intuitive, informal theories that individuals develop to explain the outcomes, effects, or consequences of algorithms, which guide reactions to and behavior towards algorithmic systems whether accurate or inaccurate.

Gaming the system: This code applied to the manipulation, experimentation, or hacking of algorithmic systems (e.g., customizing settings, creating fake profiles, changing the inputs) to gain better user control.

Health care: This domain code included algorithms used in healthcare and medicine.

Human resources: This domain code included algorithms used in personnel recruitment and management.

Human-algorithm interplay: This code referred to the role that humans play in programming algorithms, and algorithms actively making decisions for and with people.

Infrastructure: This code alluded to the broader sociotechnical system that algorithms are embedded in. Algorithms never work in isolation—they are always dependent on other things and part of larger systems dependent on interactions with people.

Key algorithmic operations: This code included key algorithmic operations mentioned by participants including filtering (isolating what's important, e.g., when social media sites filter

stories that relate to your known interests to design your own personalized feed), classification (picking a category, often used in advertising), and ranking (making an ordered list).

Machine learning versus rule based: This code distinguished between rule-based algorithmic systems that rely on explicitly stated, predictable, and static models of a domain versus machine learning systems that create their own models, are not predictable, and have more autonomy.

Mass media: This code included sources of information about algorithms from the mass media including books, documentaries, films, news, podcasts, etc.

Oppression, privilege, power: This blanket code referred to situations when an agent group, whether knowingly or unknowingly, abused a target group through algorithms. This is rooted historically and maintained through individual and systematic discrimination, bias, and social prejudice, resulting in a condition of privilege for the agent group at the expense of the target group. Privileged social groups have access to resources that enhance one's chances of getting what one needs in order to lead a comfortable, productive, and safe life.

Perspective taking: This code referred to exposure to someone else's personalized feed or algorithmic outputs to trigger the realization that algorithms work differently on different people.

Placement in the curriculum: This code referred to where and how algorithmic literacy concepts can be embedded into the mainstream curriculum for college students. It included the child code *information literacy instruction* to tease out instances where participants mentioned specifically information literacy instruction or libraries.

Privacy of personal data: This code referred to data about oneself (data exhaust generated through the use of computers, cell phones, or surveillance behavior) that is automatically available to other individuals and organizations through big data collection practices. The individual is often not able to exercise a substantial degree of control over those data and their use.

Recreation and retail: This domain code included algorithms used in recreation and retail delivering personalized content in order to display results to match an individual's expressed or presumed preferences. Examples included entertainment algorithms such as Netflix and retail algorithms such as Amazon.

Search: This domain code included programs that search for information on the World Wide Web such as Google.

Skepticism anarchy: This code referred to being overly cautious and suspicious of algorithms resulting in a feeling of determinism.

Social media: This domain code included websites and applications that enabled users to create and share content or to participate in social networking. Examples of social media platforms included Facebook, Twitter, TikTok, and YouTube. Also included were the child codes *filter bubble*, which referred to when people in an online environment were exposed only to opinions and information that conform to their existing beliefs; and *disinformation/misinformation*, which referred to information that was false, whether deliberately intended to cause harm or not.

Social services: This domain code included algorithms used in a range of public services run by the government intended to provide support and assistance towards particular groups, including criminal justice and welfare.

Speculative futuring: This code referred to speculating about the harmful impacts and ethical considerations of algorithms in the future.

Tactile (building): This code pertained to hands-on engagement in the building of algorithms in order to better understand how they work.

Transfer of knowledge: This code referred to when a user applied algorithmic knowledge about one domain to a new domain.

Dedoose Analysis

After all the transcripts were coded, the researcher used the “Analyze” feature of Dedoose to create several charts. Under qualitative charts, the “code application” chart was used to get frequency totals for occurrences of a code across all participants along how many times specific participants mentioned a code. The researcher used the “code presence” chart to see the presence of a code by each participant rather than number of occurrences. The “code co-occurrence” chart was also used to see frequencies for situations when the same codes were applied to the same excerpt or overlapping excerpts (see Appendix D for a sample Dedoose chart of “code co-occurrences”). Under mixed methods charts, the researcher used the “codes x descriptor” chart to look at each code by demographic data such as gender and field of study to see if there were patterns. However, there were either no noticeable patterns or there were not enough use cases to run a chi-square test to look for significant differences between participants when grouped demographically.

Limitations and Delimitations

One drawback with online, video-based focus groups and interviews was that nonverbal cues were harder to observe (Morgan, 2019). Also, for focus groups, turn taking was more problematic than focus groups held face-to-face (Morgan, 2019). Another problem was the contrived research setting of an online Zoom meeting rather than a natural setting, despite its necessity due to the COVID-19 pandemic as well as the wide variety in participants' geographical locations. Ideally, qualitative research focuses on participants in their natural field setting (Creswell, 2009). It also wasn't possible to offer the interviews in any language other than English despite having international participants. This was somewhat mitigated by the shared enthusiasm for the research topic amongst both participants and the researcher. Another limitation was the lack of moderating experience on the part of the researcher for the focus group; it is a best practice to choose individuals with previous moderating experience (Cyr, 2019). This helps offset the inherent lack of control over the flow of talk by the moderator that is another potential limitation (Kamberelis & Dimitriadis, 2013). Along similar lines, interviewers need to be careful about unintentionally reinforcing content through nonverbal cues like smiles and nods, which could have biasing effects on the rest of the interview (Crano et al., 2014). In general, the researcher's presence during the interviews may bias responses (Creswell, 2009). Participants in focus groups may also feel pressure to conform with the majority opinion, although the researcher tried to offset this by providing a safe environment where everyone felt free to express their true opinion (Crano et al. 2014). Also, not all participants were equally articulate or talkative (Creswell, 2009).

The small sample size used for both the interviews and focus group means the results were not necessarily transferable to other situations. Also, the lack of racial diversity amongst participants in the study was a limitation because all but one participant identified as White. Members of marginalized communities have the positionality to notice issues around the use of algorithms related to oppression and structural inequities that members of a privileged class might miss. Along similar lines, the researcher's interpretation of the results of this study was informed by her positionality, which was largely one of privilege: she was a White, able-bodied, cisgender female. Her native language was English, and she grew up as part of the middle class. As such, the perspective of marginalized community members was not represented in the interpretation of the results. Another limitation with both the interviews and focus group was that this only provided indirect information filtered through the views of the participants (Creswell, 2009). Member checks were also not done to verify the overall codes or themes with the study's participants; this would have increased credibility of the study (Seidman, 2013). Finally, although it is a best practice for single researchers to find another person to cross-check their codes to reach intercoder agreement, this was also not done (Creswell, 2009). To mitigate this, the researcher reviewed five randomly selected transcripts after coding was complete to double check for her own reliability in coding consistently, and the researcher also used peer debriefing with the dissertation committee because this "adds validity to an account" (Creswell, 2009, p. 192).

Validity and Trustworthiness

One step taken to ensure the data was valid and trustworthy was to not ask leading questions during the semi-structured interviews and focus groups that might limit the participant

into providing only one kind of answer (Ayres, 2008). To increase the trustworthiness of both the interviews and focus group, the researcher exercised patience and good listening skills, and also double checked all transcripts to correct for any transcription errors. Triangulation was used to compare data sources before drawing conclusions: the literature review, semi-structured interviews, and focus group were cross-checked to test the validity of all conclusions from multiple sources. To establish dependability, the processes used for data collection (the interview and focus group protocols and Zoom recordings), data analysis, and data interpretation were rigorously documented to increase validity and leave a solid audit trail if future researchers wanted to apply this study to a new context. The researcher also presented information that ran counter to the study's main themes when applicable to ensure all participant perspectives were represented. During the analysis, primary data with direct quotes was included to provide richer descriptions and increase validity. A codebook with clear definitions of all codes was also created. The researcher followed the protocols for ensuring ethical research standards such as having participants sign an informed consent form approved by the IRB, protecting the data, and ensuring confidentiality.

Conclusion

In summary, a qualitative study employing both a focus group and semi-structured interviews was conducted with scholars and teachers in fields related to critical algorithm studies. This method was selected to take beginning steps towards determining if there was a growing consensus among experts in terms of what components of algorithmic literacy were specific to college students and what behavior and knowledge contributed to students' algorithmic literacy. As discussed previously, there was no widely accepted operationalization of

algorithmic knowledge yet for this emerging field of study. This study contributed to the research on how algorithmic literacy can be taught in the college curriculum. A second intention of this study was to discover ways that the *Framework for Information Literacy for Higher Education* (ACRL, 2016) could be expanded to include algorithmic literacy.

CHAPTER 4

FINDINGS

Background

This qualitative, exploratory study sought to investigate the topic of algorithmic literacy for college students in order to better understand the elements and characteristics of a successful curriculum for algorithmic literacy. Through conversations with scholars and teachers in critical algorithm studies across different disciplines, the researcher sought to determine the extent to which these subject matter experts agreed on what a curriculum for algorithmic literacy would look like and how students' understanding of algorithms could be improved. Eleven individual semi-structured interviews and one focus group consisting of five people were conducted with scholars and teachers of critical algorithm studies. From this research emerged approximately 648 minutes of recorded transcripts. Thematic pattern analysis was conducted on the transcripts, and this resulted in a total of 39 codes and three child codes that were applied to 815 excerpts through the Dedoose software approximately 883 times.

Codes consisted of a blend of a priori codes and emergent codes from the data. The researcher developed a codebook with a set of initial a priori codes based on the research questions, conceptual framework, literature review, and interview questions. The original a priori codes for the first theme "Algorithmic Literacy Knowledge Components" that were general characteristics of algorithms were black box, algorithm strengths, machine learning versus ruled based, key algorithmic operations, and human-algorithm interplay. The emergent codes the researcher added as general characteristics after reviewing the transcripts from the interviews were common misconceptions, algorithmic reasoning versus human reasoning, and

infrastructure. The idea of common misconceptions was unanticipated. The researcher expected algorithmic reasoning versus human reasoning to be included in the code key algorithmic operations rather than its own concept and did not fully grasp the importance of the entire infrastructure, or socio-technical network behind an algorithm, as a key trait until after the interviews. The original a priori codes that were key domains in everyday life using algorithms were search, social media (with child codes filter bubble and disinformation/misinformation), social services, health care, and recreation and retail (formerly called *other personalization* because it was expected for news aggregation to show up in this code). The emergent domain codes the researcher added after reviewing the transcripts were human resources, finance, and education. These were not originally included because they were not as dominant in the literature review. The original a priori codes for ethical considerations for the use of algorithms were bias causes, privacy of personal data, accountability, oppression, privilege, power, and commodification. No new codes were added as ethical considerations after reviewing the transcripts.

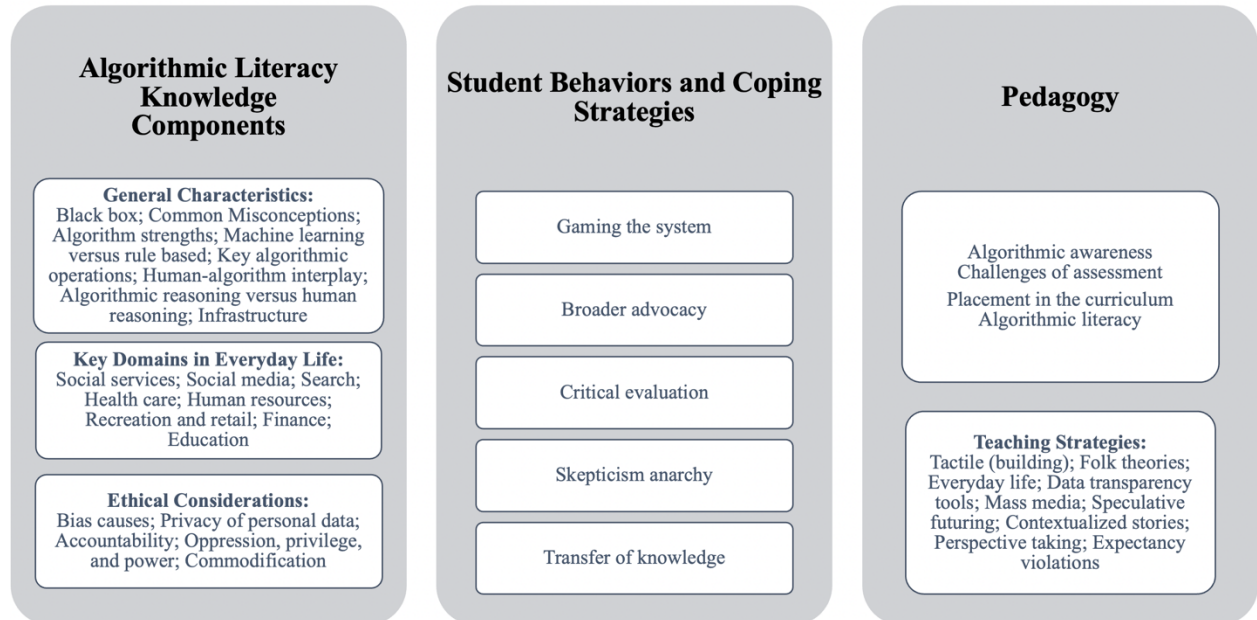
The original a priori codes for the second theme “Student Behaviors and Coping Strategies” were gaming the system and critical evaluation. The emergent behavior codes the researcher added after reviewing the transcripts were broader advocacy, skepticism anarchy, and transfer of knowledge. There were so few user studies that involved college students specifically, that the researcher could not predict ahead of time what most of the behavior codes would be. The original a priori codes for the third theme “Pedagogy” were algorithmic awareness, challenges of assessment, algorithmic literacy, placement in the curriculum, and mass media. The emergent codes the researcher added after reviewing the transcripts were tactile (building),

folk theories, everyday life, data transparency tools, speculative futuring, contextualized stories, perspective taking, transfer of knowledge, and expectancy violations. Most of these concepts did show up somewhere in the literature review, but often in the context of working with K-12 students or within the context of general user studies (not with college students). The researcher did not foresee the connection between these strategies and college students specifically until participant after participant brought them up during the interviews.

The focus group was used to triangulate and confirm data found from the semi-structured interviews, and all of the codes in the codebook emerged at least once during the focus group with the exception of the teaching strategy codes transfer of knowledge and contextualized stories. As such, a decision was made to merge code frequency counts and individual quotes from the focus group with the interview data analysis. The codes were organized around three main themes (see Figure 2). The results in this chapter were organized around these three primary themes: (1) Knowledge Components of Algorithmic Literacy; (2) Student Behaviors and Coping Strategies; and (3) Pedagogy. The first theme, Knowledge Components of Algorithmic Knowledge, was further divided into three subcategories: (a) General Characteristics and Distinguishing Traits of Algorithms; (b) Key Domains in Everyday Life Using Algorithms; and (c) Ethical Considerations for the Use of Algorithms. This thematic analysis was in service of the following two research questions: what components of algorithmic literacy are specific to college students; and what behavior and knowledge contribute to students' algorithmic literacy.

Figure 2

Final Codes Organized by Theme



Introductory Summary of Findings

Dedoose allowed the researcher to count the number of times something was coded in an excerpt and the percentage of participants who discussed a particular code to determine the most dominant codes, as well as set a minimum threshold for each code to be applied three times or more in order to be included. Also, Dedoose allowed the researcher to see frequencies for situations when the same codes were applied to the same excerpt or overlapping excerpts to see which codes cut across multiple aspects of algorithmic literacy. The top 10 code applications overall across all three themes when organized by depth, or total frequency discussed across all participants, were the social services domain (55 times), the social media domain (50 times), algorithmic awareness (45 times), challenges of assessment (39 times), bias causes (35 times),

black box (33 times), common misconceptions (30 times), algorithm strengths (30 times), placement in the curriculum (29 times), and algorithmic literacy (29 times). For the complete list of codes organized by frequency across all three themes, see Table 4.

Table 4*Code Frequencies Across All Themes*

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Social services	55	75%
Social media	50	88%
Algorithmic awareness (lower levels)	45	81%
Challenges of assessment	39	75%
Bias causes	35	81%
Black box	33	81%
Common misconceptions	30	63%
Algorithm strengths	29	69%
Placement in curriculum	29	69%
Algorithmic literacy (higher levels)	28	69%
Machine learning versus rule based	27	69%
Privacy of personal data	27	69%
Accountability	26	69%
Tactile (building)	25	94%
Key algorithmic operations	25	75%
Folk theories	24	81%
Everyday life	24	75%
Human-algorithm interplay	23	81%
Data transparency tools	22	56%
Mass media	21	63%
Gaming the system	19	69%
Broader advocacy	19	50%
Critical evaluation	18	56%
Algorithmic reasoning versus human reasoning	18	50%
Oppression, privilege, power	17	63%
Commodification	17	56%
Search	16	69%
Infrastructure	16	50%
Health care	16	44%
Human resources	14	63%
Speculative futuring	13	50%
Recreation and retail	13	44%
Contextualized stories	12	38%
Finance	12	38%
Perspective taking	11	56%

Table 4 (continued)

Code Frequencies Across All Themes

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Skepticism anarchy	10	44%
Transfer of knowledge	9	38%
Expectancy violations	9	31%
Education	7	25%

As discussed in Chapter 3, often the researcher applied multiple codes to the same excerpt because there was overlap between concepts, both within the same theme (knowledge components, student behaviors and coping strategies, and pedagogy) and across all three themes. The codes with the most co-occurrences, or overlap with other codes, were (in order of most frequent): algorithmic awareness, bias causes, social services, privacy of personal data, black box, everyday life, machine learning versus rule based, and key algorithmic operations. Specifically, the top four pair matches when all or part of an excerpt was coded with the same two codes were everyday life and algorithmic awareness (11 times); bias causes and machine learning versus rule based (seven times); gaming the system and algorithmic awareness (seven times); and privacy of personal data and algorithmic awareness (seven times). The overlap in codes was indicative of how interrelated the codes were and the fact that these concepts should be used concurrently in the classroom. For example, a lesson on the key domain of health care algorithms might also touch on the knowledge components of algorithmic reasoning versus human reasoning and bias causes; meanwhile, the professor might use the teaching strategy of contextualized stories to present a case study of an individual situation with a health care algorithm that inspires the behavior of broader advocacy in students.

Theme 1: Knowledge Components of Algorithmic Literacy

During the interviews and focus group, a key set of knowledge components emerged as teaching topics for algorithmic literacy. The knowledge components were divided into three subcategories. The first subcategory of knowledge components consisted of general characteristics and distinguishing traits of algorithms that would be helpful for a learner to understand in order to become more algorithmically literate. The second subcategory of knowledge components consisted of key domains in everyday life using algorithms, with a focus on how they operated and the potential benefits and risks. These would help students develop awareness of the variety of use cases where algorithms are impacting society the most. The third subcategory of knowledge components consisted of general ethical considerations for the use and application of algorithms. Some, but not all, participants associated the third category with a more sophisticated level of knowledge beyond just basic algorithmic awareness. There was overlap between the three subcategories, and they could be taught concurrently. This section was organized by the three subcategories, and within each subcategory the codes were listed and discussed by order of frequency (going from most frequent to least frequent). To see the frequency totals again for all codes, refer to Table 4.

General Characteristics of Algorithms

The first subcategory of knowledge components consisted of general characteristics and distinguishing traits of algorithms that would be helpful for building algorithmic literacy. See Table 5 for all the code frequencies within this category.

Table 5*Code Frequencies: General Characteristics of Algorithms*

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Black box	33	81%
Common misconceptions	30	63%
Algorithm strengths	29	69%
Machine learning versus rule based	27	69%
Key algorithmic operations	25	75%
Human-algorithm interplay	23	81%
Algorithmic reasoning versus human reasoning	18	50%
Infrastructure	16	50%

Black Box

The concept that algorithmic systems were often mysterious black boxes, or systems that can be viewed only through inputs and outputs but not through the internal process, came up frequently. Kris stated, “Algorithms are so heady and seem very confusing because they’re happening all inside the computer and we don’t know what’s going on.” Jess said that an algorithm “outputs cryptic information, but nobody knows how it works.” Kenzie also made a similar comment:

We don’t actually know what algorithms are doing in many cases, there’s sort of a limit, there’s a ceiling to what we can objectively know about algorithms, so it makes it difficult to measure knowledge about them. (Interview)

Foster explained why this poses a major challenge to teaching algorithms:

And every aspect of algorithms you learn more about teaches you about all of the aspects that you have no idea about, that are black boxes, some of them, as I said, intrinsically so, and we will never be able to coherently and completely explain them. (Interview)

During the focus group, Norman felt understanding the black box nature of algorithms was a key element for learning about them. He stated, “Being able to actually reflect on what you don’t know about algorithms in relationship to what you know is probably pretty important.” Norman

added that even though he studies algorithms, he himself cannot be completely sure he knows how they work:

I think I know how—and I, you know I studied this stuff—I think I know roughly how Facebook’s timeline makes decisions about what I see and what I don’t, but I don’t actually really know 100%. (Focus Group)

Along similar lines, Monica explained the difficulties of verifying how algorithms really work. She stated, “You try to map it with what actually is out there, and the system says they do, but most of the systems are not open, so you actually don’t know, to justify it.” Monica discussed how algorithms lack transparency and are a ‘black box’ not only in how they operate, but also in whether they exist at all, how they impact users, and what happens in machine learning. In the focus group, Becca also mentioned that sometimes the mere existence of algorithms can be a black box:

There’s a lot of stuff happening in the background that we don’t even know about that’s operating on our world, so there could be something massive [using algorithmic decision making] out there that I’m not even aware of. (Focus Group)

Jess discussed the mystery that sometimes surrounds whether or not an algorithm exists:

We don’t know, one of the big issues with understanding algorithms is a lot of the times, we do not know if the thing is an algorithm. Sometimes it could be that there is an algorithm in the system and selling it in a way that doesn’t mention algorithms at all, sometimes it could be that they’re selling it as having algorithms but they actually haven’t got the algorithm to work, and so they’re having sweatshop workers in Malaysia answer their questions for 20 cents a pop, we’ve seen examples of both of these phenomena happening. (Interview)

Mason discussed the unknown impacts of algorithms on users as another aspect of the “black box”:

Algorithms are influencing many parts of our lives and often in unseen ways or ways that may not be visible to the people impacted by them. Maybe—because they are used to inform decisions that then those decisions impact people’s lives, whether that’s residents of the city or students—maybe the algorithm that shaped those decisions or informed a

decision maker, that algorithm itself may not be visible to the downstream stakeholders impacted by them. (Interview)

Several participants discussed reasons why algorithms are so secretive. Evelyn commented, “A lot of these things are trade secrets,” and Mason alluded to “ways that organizations obfuscate algorithms,” such as if the “city agency that created them doesn’t want to disclose them.”

Monica felt that “sometimes the algorithm is a scapegoat for [companies] to just say it’s the system, I can’t do anything.” Frank Pasquale (2015) helped popularize the term “black box” in relationship to algorithms through his book *The Black Box Society: The Secret Algorithms That Control Money and Information*. Nolen explained why the book is useful:

[Frank Pasquale] talks a lot about how algorithms are related to corporations and that, a lot of the black boxness is legal black boxness, it’s about the legal protections afforded to corporations in a particular context, and I think that’s really useful, because then you realize that oh, algorithms aren’t secret because of their inner nature, they’re secret because companies, this is the way that companies are designed—if companies were designed differently, they work differently. (Interview)

Richard pointed out that sheer complexity was to blame for part of the black box nature of algorithms: “On all of these systems you’re never interfacing with one algorithm right you’re interfacing with tens of thousands of them linked together in complicated systems.” Becca felt that part of the reason algorithms were difficult to pin down was because many of them they are dynamic and always changing:

You know algorithms are, these systems change in response to all kinds of things . . . obviously we won’t know [what the change is] because they’re doing it in secret, because they’re trying to outfox the SEO or whatever or who knows what they’re responding to, they aren’t going to tell us. (Focus Group)

Nolen expressed a similar opinion about the instability of algorithms:

The systems are not very stable and what I mean by that . . . a social media news feed algorithm, a personalization system, a recommender system, something like that, and those systems are under continuous development at software companies right, the

Facebook news feed algorithm has been changed continuously. And so that's not stable, not in any absolute sense, it might be sort of relatively stable at given moments in time, but it's always changing in small and big ways. (Interview)

Jordan explained that because algorithms are largely impenetrable and continuously changing, it is important to be transparent with students when you are teaching them about algorithms so there is a shared understanding of how you are defining an algorithm:

I think transparency, I think, coming to a shared understanding early in any teaching session or any part of when you're writing scholarship about what you're defining as an algorithm in that moment, and not mathematically this is a binary heap algorithm and this is how it works, that's not—you're up above that level. And so what you're asking them to think about is here's what we agree an algorithm is—a set of rules, followed by a computer or something like that, and then kind of pulling on that transparency and saying okay well, what does that mean. (Focus Group)

Common Misconceptions

There are popular stereotypes or misunderstandings surrounding algorithms and artificial intelligence that students need to be disabused of. Participants pointed out that since algorithms are currently a trending topic, many people think they are new, but they have actually been around a long time. Also, they discussed people's tendency to have strong opinions about algorithms that are extreme in either direction. Foster elaborated:

We are currently in an incredibly hyped phase of algorithms, this insane buzzword that will either solve everything or is going to ruin society entirely, and it's always one of the two never somewhere in the middle ground. And if you actually know what most of these algorithmic systems do, and you have the knowledge of math and statistics, you realize that the majority of methods that are involved in that, have been around for sometimes literally centuries, but at least 50 years. The only difference is that we have now the computational resources, the processing speeds, to have the same methods operate on data sets that are by a factor of 10, 100, or 1000, or a million larger than what was possible, even just 20 years ago. (Interview)

Jess felt that an important teaching strategy was to show students examples of algorithms “in places that people don’t think of” outside of the “high profile explicit” places. They explained why that was a good strategy:

Because that I think to my eyes slightly denaturalizes the claim of, like, oh, this is the brand new thing that’s going to change the world, well we’ve had a lot of them since the 80s and—spoiler alert—things haven’t got better since the 80s. (Interview)

Evelyn also commented on the fact that ruled-based processing isn’t new, despite many people conceiving of it as new:

Human organizations have been following processes to make decisions for a really long time. So the fact that there is now an automated process to make a decision, I don’t actually . . . we have to start asking ourselves well what’s really different about that, and there are two things right, one of them is scale and speed, so the decisions can be made a lot faster than a human process, and the other one is the availability of data to use as input to the decision. (Interview)

Nolen remarked that it’s helpful to pick apart different aspects of algorithms and show how they have been used before in order to demystify the idea that they are new:

I find it really useful to think of algorithms not as being their own brand new kind of technology, but as being part of a variety of other sets of technology—another classifying scheme, we’ve had lots of those before. They’re a kind of bureaucracy, we’ve had those before—they are just technologies straight up, they are like sharpened sticks that you use to stab something with right, they are that much technology. So thinking about them in these other contexts, outside of any of this gee whiz algorithms are so modern right now frames, I find really, really helpful because there’s so much mystification around them. (Interview)

Kris expressed a similar opinion to Foster in terms of people’s extreme reactions to algorithms:

Technological determinism versus utopianism, so getting across the understanding that algorithms aren’t good or bad, right, they aren’t one or the other, and I think that also feeds into that question of binary stuff that I was talking about earlier. Algorithmic literacy really is just a type of critical thinking. And so understanding that algorithms aren’t just bad things that are out to ruin things and make things really horrible for everyone, and they’re also not going to make everything amazing and flowers and everything’s perfect now that we have these magic machines that make things work, in ways that we couldn’t possibly conceive of. (Interview)

Another theme that came up was people believing that AI and algorithms were similar to how they are portrayed in science fiction. Mason commented:

A lot of people sometimes . . . they think of algorithms, they think of AI, and they go right to killer robots. And I think we're a long way, well rather there are many ways closer to home, in which algorithms are already impacting our lives, and I think having grounded examples from local communities would be really helpful. (Interview)

Along similar lines, Diane also felt that showing people how algorithms were impacting their daily lives was a good way to demystify them:

Being able to differentiate AI in the real world from sort of science fiction AI, people come in with a lot of preconceptions of what AI is or what it could be based on popular media—both the news as well as TV and movies and books. And I think helping people to sort of distinguish what is and is not possible in the current realm of AI and what might be possible in the future is important. . . . There's also an issue with people not recognizing that AI is in their mundane social media news feed, for example, because they're expecting it to look like a humanoid robot, so I think the mismatch between what they think it is and then where it actually appears in their daily life can be strong. (Interview)

Science fiction also tends to give AI and algorithms human qualities, and the problem of anthropomorphizing was discussed. In the focus group, Madeline said it's important "not to personalize [algorithms] too much" or you run the risk of getting "the robot rights crowd involved, and you don't want them." Nolen discussed how people sometimes believed an algorithm was a person going through their social media posts: "When someone thinks people are involved in the algorithm, they might think that people are really involved in the algorithm, like someone's looking at their posts—and being like this person picked this thing." He gave another example:

It's like the conspiracy theory about how iPhones are listening to you because you see an ad for something that you talked about near your phone, or something like that. And that's not a very plausible theory about how advertising works right. It would be so much data for every iPhone in the world to be constantly recording and sending audio data back

to Apple that was somehow going to make it into other companies that weren't Apple in order to give them ads when it's so much easier for them to just profile you. (Interview)

Several participants discussed the opposite as another common misconception: that algorithms and AI are working completely free from any human influence, and therefore neutral and unbiased. Foster described it as “preconceived notions of the objectivity of technical systems and artifacts.” Diane commented:

The role that humans play, I think, is really big—the fact that people are actively involved and curating these data sets or training the AI, and it's actually a lot of human work that goes into building and customizing and maintaining a machine learning algorithm, for example. I think that's really important, and I think a lot of people don't understand that. I think it's just sort of like the AI is doing its own thing, it's unbiased, it's solving these problems on its own. Well, no, a lot of human input went into that, I think that's something that's really important, understanding at a high level how some of these commonly used algorithms work. (Interview)

Nolen also commented:

There's a sort of broadly popular idea that algorithms are impartial right, that they're going to be consistent and they're going to give you the same results over and over. So if you're dealing with the kind of domain where having sort of regularized decision making is valuable, for instance, in government assistance decisions where you don't want it to be down to the arbitrariness of a potentially racist desk clerk, or something like that. That can be good in theory, the problem of course is that as many, many, many people have shown, these systems are not impartial. (Interview)

Monica stated, “Most people don't think these systems make mistakes,” and she went on to describe her observations of how strongly users trusted Facebook's algorithms during her research:

Participants followed the lead for whatever the algorithm created, because they felt it knows better. . . . The algorithm played a major role for telling people what their friend groupings mean even though it was supposed to be objective. . . . I wondered why do people just trust these algorithms, that sometimes bias their minds. (Interview)

Jess alluded to the “imaginary use or myths of AI, the idea that algorithms are just better than us,” and they felt we lived “in a cultural environment but treats AI as this futuristic thing that's

just better than us, that we should defer to,” an attitude they felt also contributes to techno solutionism and less interrogation over algorithms. They elaborated:

AI isn't God, AI is a technology built by some people somewhere. . . . It is a very human process. . . . We don't know exactly how a particular system works, but a lot of the times we know how systems more generally work, which is some people got together and hacked it together till it looked like it worked And I feel like taking that mindset is likely to give people a lot better a sense of how reliable these systems are and whether they should be stuck in charge of vital parts of infrastructure. (Interview)

Nolen stressed the importance of reminding people that humans are ultimately behind algorithmic decision-making:

These technologies aren't what's impacting society, it's the social decisions of the people who are designing, building, and maintaining them that are impacting the society. So, the frame changes from technology over here impact society over there to we're all in society—this part of society is impacting this other part of society and maybe themselves, and I find that useful, because then it sort of moves the frame around a little bit. . . . It brings humans back in. (Interview)

Diane brought up the stereotype of AI and algorithms as being “something stodgy or cold or associated with nerdy people.” She had some ideas for how to combat this stereotype with students:

Making learning opportunities culturally or personally relevant to people can be a really useful inroad, so tying into students' personal interests and creating learning interventions that sort of consider cultural background and motivation can be useful . . . [for example] sample based music compositions, like electronic music or hip hop music, and create those using a code—and it's been shown to be really popular amongst women and students of color because it sort of allows you to create something personally meaningful and expressive, and this idea of computing as something a little bit more creative. (Interview)

Algorithm Strengths

Participants felt that algorithmic reasoning was not as good as human reasoning when it came to decisions requiring context or emotion or where the answer is not straightforward.

Kenzie's assertion that "In the social sciences, you have lots of biases so that's riskier" was typical. Richard elaborated:

The [algorithmic] systems that the inventors would say are most suitable are strong rule-based systems, so where there are strong kind of yes and no's as to how the system should work and so where the data is relatively clean and definitive, because the problem with a lot of the algorithms is that they're working on problems that are very complicated particularly so a lot of them work quite well on kind of hard scientific, maybe more kind of natural science, physics kinds of questions. (Interview)

Jess indicated that algorithms were good in situations where "it's pretty unambiguous like it's either right or wrong, and whether it's right or wrong, there are a set range of possibilities it could be." They added that algorithms shined in cases where "it's a really firm, grounded task, so situations where it's not a matter of, there's any discretion or any contextual knowledge that's necessary." Kris also stated that algorithms are good with "factual information needs where there is a clear answer to the question that you're plugging (in) . . . there's a right answer." Mason felt algorithms were useful in "cases where the outcomes are fairly predictable." In addition to being good at making decisions in controlled environments with less ambiguity, algorithms are effective at speed and efficiency. Foster noted, "The current computer systems can do it at 10,000 times faster than 15-20 years ago." Richard gave as a positive example the screen that shows you when the next bus is coming when you are using public transportation. Evelyn gave as an example anti-lock brakes on a car:

My car has algorithms in it to decide when my anti-lock brakes are going to come on. . . . There are so many ways that increases in processing power are helping make people's lives better and safer and more efficient. So, you can't really lump algorithms into a bucket and say they're all good or they're all bad—it's a tool for people and companies. (Interview)

Other algorithm strengths mentioned by participants were doing repetitive tasks faster than humans and detecting patterns in large sets of data. Diane stated, "Detecting patterns in

large amounts of data or being able to process large amounts of data quickly and efficiently, is something that we as people are not necessarily very good at.” Mason also mentioned algorithms were useful when “you are trying to anticipate unexpected events or rare event prediction . . . [and] given enough data, you can start to learn some of these patterns.” For these reasons, medical data or data in the hard sciences or natural sciences were mentioned by several participants as a good match for algorithms because they can improve discovery or innovation. Kenzie asserted, “Medical data, trying to come up with different treatments for things like cancer or other really significant elements that, you can see patterns and really, really big data sets.” Jess felt exploratory science such as developing “new drugs, or to fold proteins” could be a good match, while Diane stated “Potentially in a positive way, I think AI is being used in a lot of healthcare settings, and I think people see a lot of promise there.” Evelyn cited mammogram algorithms that flag things for a doctor to look at as being beneficial. Kenzie mentioned climate change as another area where algorithms were beneficial. Foster felt that genetics research, or identifying problematic genes or doing better gene sequencing would be a good match because of the potential for faster scenario-based modeling. Foster also pointed out that this was a highly regulated field, making the algorithms less dangerous because they had been regulated for ethics, compared to algorithms used in other domains where there was no required regulation.

The domain of internet searching was mentioned by several participants as a positive example of algorithms when users were seeking factual information. Evelyn explained:

Search algorithms have had a huge positive impact on people’s ability to learn about things that they wouldn’t have been able to learn about before, and build on the knowledge of others in ways that—I mean the undergrads that I have in my classes just have no recollection of what it was like to actually have to go look up stuff in a card catalog or find someone’s phone number in the phone book, the speed with which people

are able to connect to the information that they need to do things that are important to them is a huge positive. You know, it has had huge positive benefits. (Interview)

Avery also elaborated on why search engines were a positive example of algorithms:

There's so much data available, there's so much information available we have to use that we need these computer processes, we need Google, there's no way to search out the vast amount of information that we would want to find—the one pizza place down the street, without these amazingly powerful algorithms, so I think for helping, for sharing information, they're really excellent. . . . I imagine COVID would have been really different without having access to all this technology and information. (Interview)

In the focus group, Jordan shared that he thinks it is important to teach students about algorithm strengths alongside their weaknesses to achieve more balance:

I haven't really experimented with this as much but just kind of flipping the script as you're teaching and saying, are there benefits to algorithms. . . . Because I think we can kind of get down the road of there's bias, it's monetized, but something I think I'm interested in just visiting more in instruction is, are there ways that this is helpful to humanity? And it might be, maybe there aren't as many opportunities for that, but just understanding that there are probably times where this can be helpful, just putting that out there . . . that might be another way that would help with teaching so you're not always on the sort of working through the negativity, the negative experiences of algorithms. (Focus Group)

Machine Learning Versus Rule Based

The majority of participants felt it was important to distinguish between rule based algorithmic systems that rely on clearly stated and stable models of a domain, versus machine learning systems that use statistical modeling to create their own models. Diane felt that this also fell under the umbrella of data literacy because a lot of the data literacy standards “relate to machine learning.” Rule-based algorithms involve instructions constructed by a human that are direct and less ambiguous, whereas machine learning builds models that learn from new data with little human intervention. Kris explained, “When we talk about algorithms, a lot of times it can get conflated because any set of rules is an algorithm.” Along similar lines, Foster noted that

algorithms have a broad definition, because technically “any rule-based system that takes an input and produces some kind of output, with some kind of processing happening in between” is an algorithm. Mason defined algorithms as “sets of rules or heuristics or a procedure for doing something,” which could be as simple as an algorithm for cooking a meal, but he emphasized “It’s probably worth distinguishing between rule-based algorithms and predictive algorithms or what might be referred to as learning algorithms.” Nolen also expressed a similar opinion: “Everyone should probably have a rough understanding of machine learning in principle, in the take prior data, learn about it, anticipate things” way. In the focus group, Norman articulated it this way:

I would say that the learning part of it’s really important from the standpoint of feedback loops and how machine learning is often integrated as part of algorithmic systems, so it’s not just that you have a program, it’s doing X thing, it’s that you have a program that’s doing X thing and using previous iterations to then improve relevance or whatever it is to feed you more information—so understanding that feedback loop process I think is really important. (Focus Group)

Notably, Evelyn disagreed and felt that an understanding of machine learning is not necessary for someone to have the algorithmic awareness they need for everyday life. She stated:

But understanding what machine learning is and how it works—I don’t think that people need to understand that stuff. I just think that they need to know there’s a process acting on their data, and that it produces different outcomes for different people. (Interview)

There was overlap seven times between the code *machine learning versus ruled based* and the code *bias causes*. Nolen explained that it is sometimes hard to justify using machine learning systems: “Given what we know about how machine learning systems encode bias in them in many of the places that they’re used now for, those supposed guarantees of objectivity.” He elaborated:

Machine learning systems, let's say in particular that are trained on prior data, are going to be biased in ways that are shaped by that data and they're going to be biased on purpose, because of that prior data—that's how they are made, that is the way that they work is bias on purpose. Not in the sense of like unfair bias, but it's in the sense of, if you want to say that one kind of person is more likely to do one thing than another person, that's the bias you're trying to build into the system. (Interview)

Nolen also described the ambiguity of trying to use training data to make predictions for things that are more subjective like music recommender algorithms, and he felt a good lesson for students was to have them participate in making training data by classifying music:

If you're a big software company, and you have a bunch of expert tags on data—like audio data, like music—you have a machine learning training set, and so they of course have done that, they've tried to learn a machine learning, machine listening model to try to automatically apply those things to songs and they apparently do that to some mixed degree. . . . One way to think about sort of how training data gets made is just participate in [music classification] and recognizing some of the arbitrariness in it, because that scales, what gets built into machine learning systems, I think that's useful. (Interview)

Participants also believed that rule-based systems were more transparent than complex, data driven machine learning systems and therefore preferable. Mason stated, “A rule-based decision tree is at least able to be inspected and contest by stakeholders” which is “preferable to more complex data driven ones where it's less clear why the algorithm produced a particular output.” Monica described machine learning models that use training data like this:

They are by nature non understandable by humans and. . . This causes limitations to transparency, by itself, for example, why can't we say how this algorithm works, because it's a neural network it learns, it combines features, its combined features are not any more understandable than the raw input of the system—but [more transparency] doesn't need to be very detailed, but at a high level, some technicality that would help users understand rather than say, no one knows, but why no one knows, or what if there are these technical limitations, what are the mitigation approaches we can use to actually get some understanding. (Interview)

Mason gave concrete examples to help differentiate between a rule-based system and a machine learning system. One example he gave to distinguish between the two was of recidivism

algorithms for parole sentencing, and he emphasized that these are more transparent than machine learning:

Not the predictive algorithms, but deciding parole sentencing was often done with these spreadsheet formulas of like okay, if the crime was this type and if the defendant had these types of prior records and things like that. Now, there's many flaws there, but at least that is able to be inspected and contested—citizens could say okay in a town hall or request for commentary or something like that—people could look at that and say okay well actually I don't agree that this particular record should be used, or this particular feature or variable should be used in this type of way. But in machine learning algorithms and in neural networks, these deep learning algorithms, computationally code heavy algorithms, it's actually really hard to get that kind of transparency and say oh this variable is being used, and it's being used in this type of way to group by these factors, or to say if you're above this, in the same way. In a decision tree, you can always trace that back. . . . There are many ways of thinking about algorithms, but at the very least thinking about ones that are rule based decision trees or these handcrafted excel formulas, versus more data driven or ones that use machine learning to update their rules [is helpful]. (Interview)

A second example Mason offered to illustrate a rule-based algorithm was one used by hospitals to “determine which workers received access to the [COVID-19] vaccines first.” He explained how this was done: “By assigning weights to particular types of workers, like if you are above the age of X, then you are a higher priority, or if you're in this type of role you're a higher priority.” It did not end up being successful:

What it turned out in effect to lead to is that younger frontline workers like nurses received the vaccine later, were deprioritized, were lower priority, than older hospital administrators who often did not have direct contact with patients and there was a huge protest. . . . They ended up overturning it. . . . That's I think an example of something where there's no computation involved. But it's still an algorithm, it was still some set of variables, linear, some weights assigned to those, you could think of that as a decision tree and decision trees are—there's computational decision tree algorithms that could be data driven. There are machine learning versions of decision trees, but there's also just sort of rule-based if you are above this age or with this role, or you live in this location and you can assign some outcome there. (Interview)

Key Algorithmic Operations

Most participants felt it was important for students to have a basic awareness of how algorithms functioned, and 75% of all participants referenced at least one algorithmic operation. One algorithmic operation discussed among participants the most was filtering, or isolating what was important such as when social media sites filter out news stories that are not a match to presumed interests so that you do not see them. Filtering was mentioned as problematic in the context of the platforms Facebook, TikTok, and Instagram because the algorithm decides what content you are looking at. Monica mentioned that in her research, she noticed people were puzzled when friends were discussing Facebook posts they had not seen because they didn't understand there were filtering algorithms. Avery described the way these platforms filtered content this way: "Like a gatekeeper in society, they show you X or they show you Y and you never see Y, if you see X, and I think that is very problematic." Avery elaborated:

Just using Facebook as an example . . . it hides people that you don't talk to, but also when you don't like somebody's posts and things and you write negative things you see their posts just as much as when you like positive things, I think that is something that people don't really realize, that that moves them up in your feed too so it's not just like oh, I see my friends more. I also see content that I wait too long to scroll through and all of a sudden now that's showing up in my feed, so I think there should be some level of understanding of how it works. (Interview)

Other participants also expressed concerns that filtering was disempowering autonomy by not letting people choose for themselves what content to pay attention to. For example, Kenzie commented:

There are concerns about autonomy—so individuals being able to exercise free will, in the face of systems that create these structures that we have to work within, and that narrow our range of decision making and potentially obscure possible options [and] decisions that we can make, and make them more difficult to see, so in that way sort of limiting the range of action, limiting autonomy potentially. (Interview)

Another algorithmic operation discussed among participants was classification, or categorizing data into a class or category. Kris stated, “Classification or ontologies . . . how an algorithm works is that it classifies things into buckets and we count them in these buckets so classification I think is a big piece [of algorithmic awareness].” Nolen commented that a lot of people don’t understand how classification based on their demographic data works:

People [are] saying, how is it possible that I’ve been targeted for this extremely demographically appropriate to me thing, that just got released and that I’m talking about, because it is new and that is also being advertised, and people just don’t believe they can be profiled. So I want people to have a sense of how profiling works and sort of doesn’t work, because then they’ll have a better sense of what kinds of critiques make sense. (Interview)

In the focus group, Becca recommended starting a discussion on classification with students by looking at library classification and subject headings information, then moving on to apply this to other systems. She explained:

You can look at library classification and subject headings information to talk about systems and where bias comes in and, what is this telling us, and then. . . . You can’t see inside Google to figure out how they’re doing things but I could see using that as a departure point to say look, all of these things are systems, they all have certain design features that come with biases, and think about that when you’re doing searches. (Focus Group)

Kris described an activity to help students better understand Internet classification algorithms:

They classify their search schemes (from Google search logs) or their browser history, entries for a week or a month or however long, and they create their own visualization and then they write a little paper about the picture that they’ve created and how that actually relates to their understanding of who they are, what they’re doing. So what’s the difference between this kind of very weird digital shadow that we’ve created—this little slice of your digital shadow—and then your reality, and how are those different or similar, what story does this data tell that you couldn’t have otherwise told, that is in there. . . . It makes it personal but it also really gets us away from worrying about things like coding. (Interview)

Monica discussed her research on the grouping of friendship networks on Facebook, such as classifying friends by family, acquaintance, stranger, etc. and then sharing posts deliberately with certain groups. Since Facebook both makes automatic categorizations and allows users to create custom groups, this could be a potential activity for students to learn more about classification. Nolen talked about classification algorithms used in music, and why he felt they were “shaping how people discover new music, what music they listen to, changing sort of what it means to have taste in music in the present.” Nolen noted that automation and music were an interesting combination because it was both technological and expressive. He stated, “It violates a lot of assumptions about how humans and machines are supposed to be different from each other.” He described a class activity geared towards helping students understand how classification works in music recommender systems:

One exercise that I do with my students, I have them pretend to be data labelers in class, so Pandora radio is a long-standing music recommender service in the U.S. And they famously have this thing called the music genome project where people with undergraduate music degrees, they hire them to hand code songs before they play on the system, so things about what the music is like. That data historically was the data that they used to represent the music, so you would find similar songs based on these expert codings. . . . What I do in class is I play three different song snippets and I give them a little worksheet that has an abridged version of some of the categories that you might classify things on, if you were being one of the coders, and so they sort of get to encounter edge cases—things where it’s not obvious what the right answer should be. Then they pair up and they compare their answers to each other, and they try to explain why their answers are different from each other because they’re never the same.
(Interview)

Classification was also discussed as problematic in the context of facial recognition software and in particular the way it misclassifies darker skinned females. Mason’s comment was representative of the concerns expressed by participants:

Facial recognition and facial detection more generally, there's often biases in society that get baked into the kinds of data available, and then those systems end up performing worse for people of color and women of color, in particular. (Interview)

A third algorithmic operation discussed by participants was ranking, usually in the context of search algorithms. Google was mentioned as a good platform to teach this on. Kris stated that even though Google ranking algorithms are not completely transparent, the basic idea of PageRank is documented enough to use it as a teaching tool. Specifically:

PageRank on Google to at least kind of understand the more people that link to X, the more X is going to move up on the page, we do have information from Google about some of the rudimentary parts of the way it works. (Interview)

Along similar lines, Kenzie suggested this:

Ask people about different data inputs that matter for search ranking, so if you have a sense of these are the different signals that are used to—in this case, to rank content in search results, you can ask people to sort of indicate whether they think those things matter for how search results are arranged on the page for them. (Interview)

In the focus group, Charlotte described a slightly different approach involving students designing their own search algorithm:

I have students try and design what the ideal library discovery algorithm would be so you type in your keywords, what kinds of things do you want at the top, and spoiler alert it's not book reviews, even though that's what happens with our discovery, so we talk about what kinds of things would you want there, what kinds of decisions are you making, and because it's about library database systems, something they kind of feel less passionate about and don't have as much personal experience with, they are able to kind of do more of that computational thinking about what kinds of things would I weight or privilege or move up the results here. So just kind of working backwards, and designing a hypothetical algorithm. (Focus Group)

Human-Algorithm Interplay

The interdependence between algorithms and people emerged as an important concept. Avery alluded to this interdependence when she commented: "I think any digital platform that

we engage with is somehow, we are affecting them, but they are affecting us, at the same time.”

Diane explained it like this:

One of the things that stands out to me about algorithmic systems as opposed to other types of technology, is that it’s actively sort of making decisions for and with people already in the world. . . . I think that sort of differentiates it a little bit from other types of technology in the sense that it’s more important for people to be able to sort of understand what’s going on behind the scenes in a way that it’s not super important for them to understand how their telephone works, for example. (Interview)

Kris described the mutuality between humans and algorithms in this way:

Understanding how we think about people and how we think about technologies as separate things, but really they’re the same right, so we could classify that as actor networks or cyborgs or cyborg feminism may be right, I think that that’s another piece of understanding how these work is that, everything is mediated technically, and many, many things are also now mediated algorithmically right. And so understanding that the human is not just—it’s not people and tech, but that it is the interrelationships and context that’s really the thing that matters the most when we’re asking these kinds of questions about how algorithms work and what they are. (Interview)

Diane explained that during her research with interactive installation pieces allowing people to engage with AI agents, she realized how little people understood how to do so:

It was really clear that people, they enjoyed interacting with it, but they didn’t really understand that they were even interacting with an AI—and they certainly didn’t understand all the different sort of ways it could respond or what it was doing behind the scenes—and sort of seeing that and seeing how it limited the creative collaboration between the human and AI . . . [led to her interest in how to] create experiences or activities that can get people to understand a little bit more. (Interview)

The role that humans play in programming algorithms came up as an important concept, especially to what degree the algorithm makes decisions independently versus to what degree the algorithm and the human work together to make decisions. Richard explained, using the domain of killer drones as his example, that there were three levels for automation:

There’s a human “on the loop,” “in the loop,” and “off the loop.” And so “off the loop” would be you just let the algorithm decide whether they can fire the rocket to kill people, and you don’t have people involved in the decision making at all it is fully automated,

you let the algorithm do facial recognition on the target and then it makes a decision whether to fire the rocket or not. The “on the loop” is it will make that decision, but it asks you to confirm it or it will make a decision, and it will act on a decision, but you can intervene. And the “in the loop” is that you have to make the decision, so it identifies what it thinks is the target, but the person makes the decision. . . . [If you have watched] any sci-fi movie over the last five years, “off the loop” is bad, people need to be involved in that decision making process. (Richard)

Richard emphasized: “There are lively debates across all kinds of domains around the degree to which you cede power to the algorithm.” Richard used the intelligent transportation system to show different examples—an example of “on the loop” was where “algorithms made decisions in real time how long the traffic lights stay on the different phases,” but operators in a control room can override or intervene when they want to; whereas an example of “off the loop” was traffic fines based on the camera network where a camera takes pictures if you go too fast and automatically issues a fine. Richard explained the difference:” The algorithm has been given the power to act on it with no human intervention.” Other domains mentioned by various participants around this issue included whether a surgeon or doctor should be involved in decisions about medical procedures, whether police should be involved in predictive policing decisions, to what extent the algorithm should factor into a judge’s parole decision, whether the bank manager should be involved in decisions to give people loans, and whether human resources should rely exclusively on algorithms to filter out resumes. Most participants felt that off the loop decisions were dangerous for many high stakes domains; as Avery clearly put it: “I think when decisions are made without questioning the algorithm that’s what’s problematic.” Diane felt it was important to use algorithm strengths (e.g., efficiency, speed, pattern detection) “as complimentary skills, rather than sort of just taking it at face value” and to consider the

context. Mason explained that when he worked with urban fire risk modeling algorithms, he made sure the algorithm was not making the final decision:

We worked closely with the domain experts to make sure that, this was something they used to inform their decisions it didn't replace it, this wasn't the only thing that they used to inform their decisions, because algorithms that have some transparency or explainability to them [are better]. (Interview)

Jess also stressed the importance of algorithms and people working together, and taking advantage of the strengths from both. They stated:

[Algorithmic decision-making is] not beneficial because the AI is smarter than us, it's useful because it can do a thing that we could do, but much, much faster, and it can then present the results to the researchers who can then say okay, this might be a thing, is it? (Interview)

Algorithmic Reasoning Versus Human Reasoning

Participants discussed the differences in how algorithms reasoned compared to humans. Several participants felt that understanding the computational thinking process of algorithms was one important component of algorithmic literacy; Foster described this as the way that large complex tasks get compartmentalized into smaller parts and connected to each other in a structured way. Jordan indicated logic models could be helpful for teaching algorithmic reasoning:

Stealing some techniques from logic, because programming is sort of about—if you do this, then you kind of follow this but also [complicating it is], not only is it going to follow these rules, it's going to learn and kind of do some of this on its own. So I want to say logic models are useful and going back to my philosophy 101 class in college. (Focus Group)

Jess cited algorithmic decision making for OCR text as an example of algorithmic reasoning that worked well. They elaborated:

[It] doesn't need to know what it's doing, the point of it is that it is a lot faster than people, it is taking a task that is human but boring and making it human and fast. But, in

both cases, the thing I sort of tend to highlight is speed because I think that's ultimately what it's about—it's not about AI can do something that we can't, it's AI can do a very narrow range of things that we can already do but AI can do those things faster.
(Interview)

Kris stated, "In cases where we want to be efficient and we just want to do something quickly and get something pretty good, an algorithm is fine, but human decisions require human hearts."

Avery echoed the same sentiment when she stated, "I think when human beings' lives are connected to decisions that algorithms make, it can be a negative." Diane felt "AI is less well suited for problems that require human empathy." Richard used the example of prison sentencing to explain that an algorithm would make its decision based on looking at previous cases that are similar and what kinds of sentences were given, whereas a human being might make a different decision that could be more lenient or stricter based on more nuanced information. Other participants also agreed that algorithmic reasoning does not work well in situations that are not so black and white. Foster noted:

The closer we move to making decisions about human beings, the more we run into the danger of abstracting away from reality in a way that is simply not accurate enough anymore, and then we run into the problems of bias, transparency, discrimination.
(Interview)

Richard further explained:

Part of the problem is they tend to simplify things so they take complicated, real world systems and they abstract, they generalize, they reduce, they simplify, because that's, to get algorithms that will work requires that kind of stuff, and it's the same with the data, the data is always abstracted and generalized and so on. . . . They are not so great on social science and humanities things because there's a lot more going on in them, there's a lot more wicked problems, there's a lot more variability, there's a lot of politics and complexity, there's a lot of contestation, so what you think is good, what I think is good, and what should be the appropriate outcome might be completely different, so anything where there's a lot of politics with a big P and a small P algorithms struggle. (Interview)

Kris expressed a similar viewpoint:

There are these kind of more exploratory questions where there's a large level of human element that we might need, where back and forth has to happen, because incognizance is occurring right, so we might not know what we don't know—and so articulating our information need, or our question, can be very challenging and with things that are either more exploratory or very human centered, I think that algorithms need to be applied very judiciously, we need to be very careful about thinking about being able to come up with a one size fits all rule list. I mean an algorithm is just a list of instructions to follow right—if then, if then—it can get very complex or not very complex. And the more complex it gets, and the more it relies on training data to then make its decision making, the less we as humans can read and understand and think about it legibly. . . . I think, in situations where there are humans that are involved, where we're talking about life or death, where we're talking about actually having to understand the variables, where oppression is likely—that's when we need to be really, really careful. (Interview)

Jess gave a concrete example from the health care domain of how algorithmic reasoning would work on psychiatric diagnosis, and how this could differ from the way a doctor would arrive at a different decision that would be better for the patient:

The belief that psychiatric diagnosis is the kind of problem . . . it follows, there are a set series of outcomes like it's do this thing, and if this thing happens then do that thing, but, of course, and all you're really doing is adding speed and objectivity to a human process that as a result of being human is slow and fallible. But in practice that's not the case, because in practice there's a big gap between the rules that are written down and what diagnosis actually looks like, and there are sometimes very, very good reasons for these rules, these differences. For example, if you have an algorithmic system that knows perfectly how to diagnose people with bipolar disorder, it will produce different outcomes from doctors, not because it is more accurate, but because if you are a doctor and you are a thoughtful doctor, and you run into someone with [bipolar], which is a highly stigmatized condition, and that person—you know their insurance does not cover the medication for [bipolar], but you know that the insurance does cover that exact same medication for major depression, and also that major depression is much less socially stigmatized—then you are going to deliberately misdiagnose them, not because you are wrong or stupid or fallible, but because in the broader context which you can see, which an algorithm can't, that is the decision that makes the most sense. And an AI can't do that because an AI cannot see that context and an AI cannot see that subjective, sometimes the best outcome is to be deliberately wrong. (Interview)

Another domain that came up as an area where algorithmic reasoning may be inferior to human reasoning was self-driving automobiles. Evelyn explained:

The idea of self-driving cars scares the crap out of me, because I feel like driving is something where you encounter a lot of varied situations, where very quick decisions have to be made and you know, I'm not a technologist, so I can't say how close we are to having cars that can make those kinds of decisions. But I don't think a human brain needs the same kind of data to make decisions as a computer does, and I think that having the necessary data sets to drive a car safely and under all the circumstances that people encounter when they're driving it's just a really hard problem. (Interview)

Foster had his students play *The Moral Machine* (Massachusetts Institute of Technology [MIT] Media Lab, 2020) game, which is an algorithmic system that generates random ethical dilemmas for a self-driving car that detects an oncoming accident. You are given some information about the people you have to choose between killing—for example, are they a convicted criminal, are there elderly people, are there children, etc. He used the game as a metaphor for the way that automation designers have to make certain decisions. Foster further elaborated how he used the game:

To explain to my students the design of these algorithmic systems is never disconnected from society, whatever technical issue you are working on. You have to think a little bit further, you have to think about where the system is being used and how you could, by creating a system like that, run into ethical problems that are so much larger than the single person can answer. . . . It teaches how to deal with a problem that has very abstract, set rules that you can't really change and impact, and it also tells you that some of these decisions are incredibly ethically laden and burdened with moral ambiguity. (Interview)

Diane also discussed the importance of context when making decisions that impact human lives, and how this context was lacking in off the loop algorithmic decision-making. She described it this way:

We understand when people make decisions, we're sort of able to develop a theory of mind and think about how they're coming to that decision, and in the context of what we know about their personality, but we don't have that same context, when we're working with an AI or when an AI is suggesting something to us, and sometimes I think that can be dangerous, because people take those suggestions at face value. (Interview)

Infrastructure

Half of the participants alluded to the importance of understanding the broader sociotechnical system that algorithms were embedded in, or what some referred to as the infrastructure. Algorithms never work in isolation—they are always dependent on other things and part of larger systems dependent on interaction with people. In the focus group, Madeline referred to it this way: “A bit more about what’s going on underneath . . . at least understanding that there is a system underneath,” including social and ethical implications. Kenzie explained it like this: “Breaking it down into sort of data, model, output, and these kinds of things—giving a basic schematic I guess of what algorithmic processes look like.” Jess described it this way:

To me, I guess this isn’t a technical concept, I guess it is from a social science lens, but the question of infrastructure, I think, is always important to me. In the sense of it’s never just an algorithm it’s always dependent on other things—like a facial recognition system is dependent on CCTV [closed-circuit television] cameras, and is dependent on people checking the alerts, and police officers to follow up on the alerts. And a fraud detection system is dependent on people inputting the information and the customer service agents, when you phone up irate because your phone got cut off, and the reason I bring all this up is I think it’s vital to make clear to people that, even if the algorithm is as perfect as things could be, that doesn’t mean that the outcome of it being there is going to be good, if it’s plugged into an otherwise dangerous system that is just sped up. (Interview)

Richard expressed it similarly:

If I was to try and advise somebody on learning about algorithms, you want to learn about both bits [the rules, or guidance, or user requirements that shape the algorithm plus the algorithm itself]—you want to learn about the context and the guidance that fed into how the system was designed and built, and then you also want to understand the actual system as it’s built on, what work it does, how it does that work, how it makes its decisions and choices and then what the consequences of those decisions and choices are. And so there’s a bunch of things to think about, but I would tend to put them under kind of politics and praxes is kind of how I phrase it so that’s understanding kind of the social, technical context of what’s happening, and then how the thing actually functions in practice, how does it work. So what shapes the system, and how the system functions, and then trying to unpack that and understand what kinds of relations produce it and are produced by it. . . . There’s a lot of moving parts here. And, none of this is created outside of context so the reasons why the system is being built and developed, the logics,

the systems of thought, the regulations, the governmentalities around it, the marketplace around it, even the finance can make a big difference. (Interview)

In the focus group, Norman emphasized the importance of understanding the broader social context that an algorithm is situated in:

One of the things that I run into when talking about algorithmic awareness and algorithmic literacy is that, I want to capture the context because I think that the social context is actually much more important in many ways, than the algorithm itself in the sense that these algorithms always exist within the platform or within the system and having a broader understanding of the different elements of that system. . . . And so understanding that sort of broader context, if it's a platform that might need platform literacy, if it's within the context of a work environment, having a literacy for understanding the politics of the workplace, I think is really important. I would also add that I think that we get a little bit trapped in the term algorithm itself again because of the ways in which these are also reliant on data structures, on the way the defaults are set, and so again that broader set of literacy is, I think, really important. (Focus Group)

Nolen also felt it was impossible to distinguish algorithms from their broader historical context, and he wanted students to understand this:

One of the issues we have here is distinguishing between what are sort of intrinsic qualities of algorithms and data, let's say, and what are qualities that are sort of historically specific accidents. . . . It's worth picking apart how the systems work in different places, if you want to know things about an algorithm you really need to know more about the sort of institutional context that the algorithm operates in and where it's coming from. . . . [Students need] to have a more sophisticated understanding of the socio-technicality of algorithms if they were going to be algorithmically literate, because I think we're beyond the point now where algorithmic literacy should just mean know that algorithms exist and be aware that they're trained. (Interview)

Key Domains in Everyday Life Using Algorithms

The second subcategory of knowledge components consisted of key domains where algorithms were problematic. See Table 6 for all the code frequencies within this second subcategory. Kris articulated that one important component of teaching algorithmic awareness was to look at everyday life. Kris suggested, "Different kinds of places where algorithms intersect with our daily life . . . so I'm thinking about health care, thinking about education,

thinking about how algorithms operate in different use cases.” Mason explained why it would be important to introduce students to a variety of domains:

Part of this idea of algorithmic awareness is understanding the role that algorithms play in our daily lives and how we are both affected by and mutually impacting the algorithms . . . generating data. . . . And I think having learners understand the breadth of domains and application types that involve algorithms, I think would be really important, and this is where yes, you can draw on examples from education, from finance, from healthcare, the public government, decision making in the public sector. (Interview)

Foster stressed the importance of having students learn about algorithms across different areas of society:

Algorithmic systems are incredibly pervasive—they’re ubiquitous and they’re just about everywhere. . . . I always ask my students [to think of a] single societal area that doesn’t have, in any way, shape, or form, contact with algorithms and no one can come up with anything. [Interview]

Kris commented, “Power is a really big key idea in algorithmic awareness, understanding how power relations are imbued in and animated by the use of algorithms in daily life.” Nolen also expressed the idea that learning about the different domains where algorithms were impactful was an important way to address their power:

There’s a sort of famous set of them . . . there’s different domains, which is the standard way we talk about algorithmic powers. . . . And then the other axis is sort of what level of control do they have, do they just shape things, do they actively determine aspects of your life chances? (Interview)

Several participants also grouped algorithmic domains into two categories based on level of control. The first level included more benign algorithms that gently nudged and shaped the collective thinking such as social media algorithms, which Kenzie explained this way:

Any impact that they have is sort of at a more micro level I guess, the kind of impact that they have is in the moment and is more minimal [but there is] an accretion over time, so, many people doing things with algorithmic systems and being nudged in different ways might add up to a lot. (Interview)

The second group of algorithms had a more severe and direct impact on human lives, such as predictive policing. Monica explained:

In my mind, there are two types of different systems—private sectors and public sectors—that use these algorithmic systems, and they both have great impact on society but in different ways. . . . When public sectors like governments or public agencies use algorithms to make high stakes decisions such as predictive policing, child welfare . . . these are the systems that have the highest impact because they sometimes even decide life or death. . . . Usually their impacts are more immediate. But, at the same time, there are a lot of sociotechnical systems, usually led by the public sector, like Google, Facebook, Amazon—and one of them might be a recommendation for shopping—it probably would not be a life or death problem or high stakes, but these things can accumulate and shape society’s behavior. . . . Those systems usually act as a collective . . . which, in the long term, can be also very impactful so, as I said, I think both are very high impact, but in a different nature. (Interview)

Table 6

Code Frequencies: Key Domains in Everyday Life

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Social services	55	75%
Social media	50	88%
Search	16	69%
Health care	16	44%
Human resources	14	63%
Recreation and retail	13	44%
Finance	12	38%
Education	7	25%

Social Services

A range of public services run by the government intended to provide support and assistance towards particular groups were brought up as troublesome, with a special emphasis on surveillance. Richard explained:

A lot of these systems are designed to manage or to regulate or to modulate how people behave, so there’s a kind of governmentality embedded in—a governmentality is a set of logics and rationales for how society is managed—and these technologies are designed to get you to act in certain ways, they either discipline you or they are forms of control and capture. (Interview)

Kenzie also discussed why policing algorithms and sentencing algorithms used in criminal justice were problematic:

Algorithms used in the criminal justice system and in policing and things like that where an algorithm making a decision about sentencing somebody for some kind of alleged crime has enormous implications for that person's life. And certainly, where they send police based on modeling of risk also has a very significant impact on the life opportunities, the life chances of people in communities where the police are sent. (Interview)

The codes *social services* and *bias causes* co-occurred five times, which reflected participants' association of social services algorithms with bias. For example, Jess commented on the discrimination that often happened in policing and surveillance:

When there are people engaging in machine learning in policing and surveillance, for example, the harms are twofold: the first one, is that the systems are often discriminatory, and even when they're not discriminatory, reinforce a surveillance dragnet. (Interview)

Jess named automatic fraud checks in welfare and facial recognition as additional examples where algorithmic systems "directly inform and shape who we can be, what we can do, what we have access to." Evelyn explained why it was important to teach people about facial recognition:

I think facial recognition . . . a lot of people are using it even sometimes if they don't know that they're using it, because facial recognition is increasingly being built into just surveillance technologies that are out in the world, and I feel like there is a really big disconnect between when people feel like they're anonymous and when they actually are, and that technology isn't as good as people would want it to be, for the way that they're using it. (Interview)

Foster named "risk assessment for recidivism" and "profiling systems that formalize bureaucratic processes including . . . who gets which social services" as problematic. Mason named recidivism too, but he also brought up municipal services:

In the government, in the public sector, in criminal justice, algorithms are being used to determine the risk of recidivism for parolees so they would be used to inform judges' decisions about parole sentencing or bail decisions, or in maybe less high stakes, but still

important decisions in the public sector, algorithms are being used to inform decisions about municipal services—including things like 311 calls—different responses to 311 calls for pothole repair or watermain repair. (Interview)

The codes *social services* and *mass media* co-occurred five times, and four of the five times reflected when a participant cited Virginia Eubanks' (2018) book *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* as a good way to illustrate problems with welfare algorithms. Nolen indicated the book illustrated “people relatively disempowered, seeking assistance with the government” and trying to get benefits. During the focus group, Norman also referred to the book, calling it “a really good treatise on the ways in which algorithmic decision making is used to essentially restrict access to government benefits and things like that and punish the poor.” Kenzie explained why she had her students read and discuss the book:

In Virginia Eubanks' [2018] *Automating Inequality*, I like the case study of the system in Pittsburgh that was for child welfare and basically alerting when there is a possible case of child abuse or neglect, and I like that case study for use in classes, because I think it's a good, clear example of how bias can enter into an algorithmic system in the real world, a good, clear example so usually talking about how people tend to report Black and Latinx people more often for child abuse and neglect than White parents so you have this biased data set, it doesn't reflect the reality of abuse and neglect being concentrated among different races, there's no basis for that, but then it's being integrated into the system. (Interview)

Kris commented on another aspect of welfare—algorithms used to allocate food stamps:

Do we need an algorithm to be making these decisions about allocating food stamps to people, or do we just need to hire and support more social workers who can make these very human decisions as people, and can actually look at the variables for centering care, rather than centering efficiency? (Interview)

Social Media

The domain of social media, defined here as websites and applications that enabled users to create and share content or participate in social networking, was cited often by participants as problematic. Nolen felt that social media played “a major part in how lots and lots of people encounter information.” Kenzie alluded to the filter bubble—the way people in an online environment are exposed only to opinions and information that conform to their existing beliefs—when she described the way that social media forms our view of reality:

[Social media algorithms] have a really important role in shaping our view of reality because of the ways that they are regulating the visibility of different people and ideas online, which is how most of us get a lot of our information now, and they kind of give us a particular view of reality that isn't everything—that is curated in some way, and it follows particular rules about what should be visible and not, usually that the platforms sort of dictate. And that then sort of sets the agenda for the kinds of ideas and people that we are collectively thinking about, and then in our smaller subpopulations that might get brought into our orbit in the different spaces that we go into online—so I think about it as creating structure and nudging. (Interview)

In the focus group, Norman described “the way that Twitter serves us things that are trending, things that are tailored” as having a similar impact, while Charlotte brought up “YouTube recommendation algorithms and how that has such a profound way in shaping how [students] understand the world and move through it.” Avery reiterated the idea of filter bubbles caused by social media during her interview:

The idea of the filter bubble, the Facebook and Instagram filter bubble, is very real. I know that I live it every day, nobody talks about the other side, and unless I go searching out for it, so I think that's a huge effect on people who think that everybody is like minded to them, which is already a psychological phenomenon that we all think everybody's just like us. But then when you actually—it's reinforced every day, what do you mean, I have 500 contacts they are all the same, we all think the same thing—I think that's really detrimental to society, I think when we're not exposed to different people's ideas, I mean I think we're living with what happened, what happens when we're not exposed to other people's ideas. (Interview)

Monica discussed both the way that the Facebook news feed hides posts from you and the way that Yelp sometimes hides reviews as examples of misleading the user. She elaborated on why this was the case for the Yelp platform filtering algorithm:

Some people think the final rating of a product is the average rating of the product, but that's not true, there are algorithms . . . changing these so even just giving this information is very important. On Yelp, there's a review filtering algorithm that will decide if your review is legit or not, and will put it either in the filtered reviews, or they call it non recommended or recommended. But the thing is that if you logged in for example, you'd write a review for Pizza Hut, and then you later go to their page and you're logged in with your username you would see your review at the top of the reviews for that under recommended, but if it is filtered, other people will not see it under recommended, actually, so there is this bias, even in the existence—so I don't know if that's a deception, but something happened. . . .Yelp would show you as recommended, even if it's filtered. (Interview)

A few participants discussed online dating platforms as impactful because they control what choices you have. Richard explained, “If you're on a dating site, [the algorithms control] who is placed in front of you as to who you might date, so your future partner.” Avery revealed that she taught business students about LinkedIn, which has many of the same filtering properties as other social media, to “give them some background on, it's not just about your profile it's about all the ways, on the back end.” Avery also talked to business students about Indeed to try to make it relevant to them.

Participants discussed the political ramifications of the filter bubble. Multiple participants alluded to the Cambridge Analytica scandal, which Monica explained like this:

For example, on Facebook, this news curation tool that decides what goes on your feed or not, previous research has shown that it can actually decide, these social media platforms can decide the result of an election, so you see how this could be impactful. (Interview)

Foster stated, “The way that content is ranked and presented to people has huge implications and impacts on the political landscape all over the world.” Becca commented in the focus group:

Certainly, the algorithms we encounter through social media have been very influential, and you can see that happening in the world in terms of amplifying things that are out there and kind of distorting our general conversations around the world. . . . Facebook and Instagram, YouTube, TikTok, these great propaganda engines, and people are really savvy about that—people who we wish would shut up have been able to figure out how to get their message out, I mean they're very good at that and amplifying it. (Focus Group)

Mason offered a similar viewpoint:

Recently, algorithms that shape how we access and share information . . . are having an incredible impact on our democracy right now, on our politics, and the way that misinformation gets produced and shared and amplified at scales that are unprecedented, that other types of information or communication technologies just couldn't offer. . . . [They] can amplify hate speech, can amplify political misinformation, propaganda. . . . They optimize for engagement, but engagement has particular political effects that are not always thought about and incorporated into the design. (Interview)

Mason gave a specific example of how YouTube recommendation algorithms have led to increased radicalism:

The recommendation algorithms on YouTube have led to increased radicalization because they amplify, and this is broadly true for these recommendation algorithms more generally, that by which I mean click through rate is one metric that gets used for optimization, right, how many times people click on a link or for YouTube, the rate at which people click on particular videos—but that optimizes for sensationalism or extremism. and there have been studies that have shown that the optimization function, this metric they use to train the recommendation algorithms for YouTube, have led to increased radicalism, increased extremism, particularly towards the Alt right, towards White nationalism, towards White supremacy. So I think that's a major, major risk to democracy right now, to having political discourse, and it certainly has amplified misinformation about COVID vaccines. (Interview)

Avery also felt that misinformation on social media had contributed to less people getting vaccinated for COVID-19:

Half the country is not getting vaccinated because of these anti science campaigns that are all over social media, and Biden himself said, this is dangerous and it's bad for the public. . . . I think we are where we are right now, because in part because of that, we should not be here again. (Interview)

Kris felt that the fact people get their news on social media now means that it becomes inseparable from entertainment, with drastic results. Kris explained:

Real systemic change has to happen, I think, also we have to really overhaul ways that people get information, ways that people. . . . Information and entertainment are also so intertwined now right, we have all of this emotional charge when we interact with information. . . . It isn't new, but it is really heightened, especially now. (Interview)

Search

Search engines (defined here as a program for finding particular sites on the World Wide Web such as Google) were a domain that came up in both positive and negative contexts.

Positive examples were coded as *algorithm strengths* and discussed in the context of that code.

Avery felt search engines “function to influence us to buy things and to read certain websites versus other websites.” Kenzie felt information algorithms played a gatekeeper role:

I was thinking about all these different ways that algorithms have a shaping role in our lives, and particularly in our information environments, which is for me, the key to how we think about everything, the kind of gatekeeping, the idea of gatekeeping and the way that we think about the reality around us. (Interview)

In the focus group, Norman stated, “Information gathering and seeking tools and the algorithms that inform those have a tremendous impact on the way that we manifest and tackle the Web essentially, so like Google’s PageRank.” Jordan responded with his opinion about search engines and how much they have changed:

[Search] prioritizes certain forms of knowledge, certain forms of knowledge that have been optimized in the correct way, the way that the search engines themselves are starting to select and feature parts of their own content in the display and the results. I mean, I got into the field in 2003, and Google was basically 10 blue links that’s what it looked like. All of us know that that’s not anywhere, I mean it’s hard to find those 10 blue links inside of that search experience, which is algorithmically generated and featured and tailored and curated for a particular user—so search engines for me has been obviously a point of research, but also it’s one of these ubiquitous moments where algorithms can be kind of very influential because it does kind of speak to what is, what is the right answer for this query. (Focus Group)

Nolen commented that he feels search engines are both useful and troubling simultaneously:

I think search engines are useful, and I think part of the problem with a lot of these systems isn't that they're bad, it's that they are so useful and problematic, at the same time right, there's always these tensions and sort of a search engine is super handy. And because it's so useful, all those other problems, no one who's being a critic of Google is going to be like, we should live our lives without any kind of search engine. Because that would be a real pain in the butt. (Interview)

Evelyn also alluded to a tension in search engines between usefulness and unintended consequences:

There are downsides, like I've also noticed people just being unable to cope when they can't find the information that they need, the stuff that's easy to find shapes our activities and our actions and in ways that are really hard to quantify. (Interview)

Kris felt it was important to teach undergraduates about how search engines really work:

I think that we take for granted this idea that people can just use Google, but no one knows what they're doing on Google, we teach ourselves, and if you're lucky, you know how to teach yourself, and if you're not, you're doing something weird, people end up in really weird places on the Internet all the time. And we need to teach people search literacy and search skills, and I think that this would be a natural piece of that class because search is 1,000% algorithmically mediated. It is a great case study or example for how algorithms work. (Interview)

In the focus group, both Jordan and Becca also advocated for teaching search algorithms to students. Jordan commented that search is so ubiquitous we often forget algorithms are involved:

I think for us in the information environment, for teaching with our students, that search experience has been best for teaching, but also, I think it's one of the more—because it seems like it's sort of a natural act at this point, it's invisible. (Focus Group)

Becca also felt that since everyone uses search, it is an easy domain to start teaching with, and then it allows for expanding to other domains:

I think it could be a really good sort of jumping off spot, to start with search to start with yeah those creepy ads they're following you around I get it and then. . . . When you branch out from that and start talking about the social impact of algorithms in all kinds of ways [like labor] then I think they can see that there's justice issues involved and there

are things that are affecting lives in a way that they don't think about with this kind of consumerist search idea. . . . I think it's good to [also] talk about those as part of information literacy. (Focus Group)

Finally, Safiya Noble's (2018) book *Algorithms of Oppression: How Search Engines Reinforce Racism* was discussed by four different participants as a great way to teach about bias in search engines. Here was a representative comment from Mason:

Sofia Noble's [2018] *Algorithms of Oppression* . . . is good to show how search engines reinforce racism through stereotyping based on both predictive text of search queries, but then also what actually gets returned, right, like if you search for Black women or search for CEO and it's all White men. (Interview)

Health Care

The health care domain elicited both positive and negative comments. The positive aspects were coded as *algorithm strengths* and discussed in that section. Jess felt that medical diagnosis was a high stakes domain. Jess stated, "Anything that is based on automatic determinations that have material impacts on people is a terrible idea." They added:

Once AI is involved, that strength in exploratory science and menial labor suddenly becomes a weakness, because now, instead of getting something right very, very fast it's getting something wrong very, very fast, which means that suddenly the number of people impacted shoots up astronomically. (Interview)

Kris also cautioned against letting algorithms make decisions in the health care domain without the input of doctors:

When we're talking about plugging numbers into a system that then does black box calculations that, as a doctor, they may not actually understand why it's offering these particular treatment options to then offer to the patient—as an example, that's where we really need to be careful and we really need to be thinking critically. (Interview)

Several participants associated health care algorithms with racial bias. Mason commented:

In healthcare, algorithms are being used to inform how hospitals and clinicians prioritize care for different patients, and so the recent paper by Obermeyer [Obermeyer et al., 2019]

finding that Black patients were rated as low risk, more often than White patients [is an example of] . . . performance disparities in healthcare” (Interview)

Kris also cited bias in recommended treatment options produced by algorithms:

Are these algorithmic systems telling doctors to communicate the wide variety of treatment options to everybody [in prostate cancer decisions]? Because one of the things that we see is that Black men are offered different kinds of treatment options, just as a matter of course, right, and some of these are more invasive or don't fit with what people actually want out of how they define a quality of life. And so anytime we introduce a sociotechnical system into a very human problem like health care for example . . . it's going to impact the people who are already being punished or oppressed by systems. (Interview)

Avery brought up another troubling aspect of algorithms in this domain-wearable technology such as the Fitbit:

Now we've got all our wearables, and that's affecting us in ways, so I read about how insurance companies are using your wearables to up your insurance, and you know I think the Internet of Things is very, very scary. (Interview)

Human Resources

The field of personnel recruitment, retention, and management emerged as problematic.

In the focus group, Norman expressed this opinion:

One of the biggest algorithmic influences on and that has a direct impact on society is automated scheduling actually for labor, and things are happening in the labor market that are completely distorting and changing the way that people structure their lives and structure their time. (Focus Group)

Kenzie echoed some of the same ideas when describing labor systems:

The ways algorithm sort of take over a lot of the bureaucratic processes that we have had in place for a long time, but then sort of automating those processes, so they create this kind of control structure for workers to work within. And sort of measuring and evaluating work that creates incentives to do things in a certain way, and in a certain time frame, and things like that. (Interview)

Algorithms used for performance evaluations were mentioned as a problem because they were so cryptic; Avery explained, “Several companies that are doing evaluations via these algorithmic

programs, people get fired and they don't know why." Algorithms used during the hiring process were also discussed as worrisome if a hiring manager relied exclusively on them to filter out candidates. Diane described this situation:

Imagining a hiring manager that uses AI to filter resumes, for example, that as opposed to sort of taking the AI's result at face value, they might look at the AI's result, but then also recognize some of the issues with those hiring algorithms that sometimes—you know, they prioritize existing biases that were already in hiring. (Interview)

Diane added that hiring algorithms trying to detect emotions are a bad idea:

People are trying to detect emotions . . . in a hiring call, trying to detect sort of emotional responses of the hiring candidate—I think those types of things are not the greatest use of AI because they're really nuanced issues. And it's really nuanced to be able to tell if someone's angry and what that means, like anger may not necessarily be a negative emotion, depending on the context. So I think those are maybe less appropriate uses for AI or areas where AI is not strong at doing what we as humans are quite good at doing. (Interview)

Kenzie described a game called *Survival of the Best Fit* (Csapo et al., 2019) that helped illustrate how hiring algorithms make decisions and illustrated how bias can creep into the decision-making. She elaborated:

There was [a game] that was related to hiring algorithms, and I think the idea of that one was trying to demonstrate how bias can kind of enter into those outcomes or the way that they prioritize different candidates. And that one was really good at sort of putting you in this role-playing scenario where you're thinking about actually employing this system and using it to aid in your own decision making, and it was sort of demonstrating how things can go awry. . . . You use [candidate variables] as inputs to think about who would I hire and then thinking about how that factors into algorithmic decision-making. (Interview)

Recreation and Retail

This domain included areas of entertainment and retail that delivered personalized content to match presumed user preferences. Diane categorized this group of domains as making “suggestions about what people should purchase or how they should entertain themselves.” The

platforms Amazon and Netflix were named frequently by participants as being the most influential and shaping people's behavior. Avery alluded to the superficiality of these platform recommendations: "There's 800 of them right, and they all give you recommendations, and are those recommendations really catered to you or they really just catered to the general public?"

Richard explained why it was important to teach students how algorithms in this domain work:

They know that Netflix is giving them recommendations about what TV program they should watch next or what movie they should watch next, and they know they're being nudged to do that, but they don't really know the kind of mechanics about how those decisions are being made and so on, so we might talk about those kinds of things . . . to illustrate the ways in which their lives are digitally mediated and then how that mediation is working in practice. (Interview)

Avery felt that Amazon's price differential based on geographical profiling was cause for getting upset:

I think when you shop on Amazon, and Amazon is showing you products that have higher, that cost more because you're of a certain geographical area, or income level—I think that's really, I think if people knew about that they'd be really upset, and people do not know that that happens. (Interview)

Finance

Participants articulated concerns over financial services algorithms that were being used to determine who gets access to loans. For example, credit-scoring algorithms often estimated the probability applicants would default on a loan by comparing them to people who shared similar characteristics (and whether those people who were "similar" paid late or defaulted in the past). In the focus group, Charlotte commented, "Financial algorithms like credit scores, and determining sort of who is worthy of credit and purchases, have such an impact and perpetuate inequality, at least in the United States, in really profound ways." Mason noted that users often lacked awareness of financial algorithms making these consequential decisions. He said, "You

apply for a loan, you get denied for that loan, but it's not always clear that it was an algorithm that did that." Kenzie stressed the importance of people learning how credit score algorithms worked:

If I know the basics of how a credit score algorithm works, then I may feel that I'm more empowered to behave in ways that will suit my interests and that I can in some ways, to some degree, ensure that I'm not getting a very, very low score that will close up opportunities for me. (Interview)

Jess also warned about the dangers of bank fraud checking algorithms:

Bank fraud checking algorithms . . . if you are someone who is, as the phrase goes, data poor, someone who only appears patchily in various systems, which is a lot of the population, particularly in the U.S., where we have so much income inequality—then what you look like to one of these systems is, by default, a crazy person. And they might just arbitrarily—from your perspective, arbitrarily—put fraud stops on your cards, and that's not a small thing. . . . That might be I couldn't pay my rent and I got dinged with an extra 50 bucks that I don't have in costs because it was a late check, or I got thrown out, or so on and so forth. So that kind of system, I think, is really interesting because it's one that we don't often think of because it is in the background, banks don't bill themselves as using algorithmic systems to do this, but they do, even if they're not the shiny whiz bang ones that we think of. (Interview)

Education

Mason commented on the potential for harm of many types of education technologies using algorithms and the need to think more broadly about their ethical impacts. He highlighted "Disparities in how these types of systems worked for learners, or the surveillance that data driven ed tech motivated and enabled in the classroom." He explained how automated proctoring could cause bias and oppression:

During COVID-19, the interest in students taking tests at home exploded, and the companies providing [automated proctoring services] turned to more algorithmic approaches to this using facial recognition, using gaze detection, to notice where students might be looking on the screen or off the screen. And posture detection, like are they turning away from the screen, things like that, to supposedly monitor whether students were cheating or not—that's the stated purpose of these. But with facial recognition and facial detection more generally, there's often biases in society that get baked into the

kinds of data available, and then those systems end up performing worse for people of color and women of color, in particular, not to mention that just by definition those automated proctoring systems define what is normal and what is abnormal while taking a test, and so disabled students, neuro divergent students, learners who have children or have childcare responsibilities or eldercare responsibilities, don't always look like other students in the way that they might learn. And so they've been flagged—students in any of these different groups are flagged for academic integrity violations at significantly higher rates. (Interview)

Mason also criticized intelligent tutoring systems for not being student centered:

These intelligent tutoring systems, these applications that track learners' performance on different types of content and provide them with assessments and content at increasingly higher levels of difficulty—which, at face value sounds great, right, you want instruction to be personalized, you want it to be adaptive to learners. But then, as you dig a little deeper and it's not always clear how teachers are actually involved in using these in the classroom or in designing, say the curricular progression of them, or what it means for students to actually have mastered content or what types of content get incorporated into their design, things like that, which beliefs about what modes of teaching are valuable right. So often they reinforce this very assessment focused or instructionist model of teaching versus maybe a more student centered or more dialogic or Socratic modes of teaching. (Interview)

Participants also commented on the impact of algorithms on college admissions. In the focus group, Madeline shared that she thinks this was having a great impact:

The ones that serve up who gets to go to which college, the algorithms behind admissions and education, and education pervading down into the K-12 and God knows, maybe in preschool, right, in streaming children and in streaming content and in homogenizing educational experience, to a large extent, building towards a homogenized product at the end of K-12. (Focus Group)

Avery relayed that she shared an article with her students about the role that social media plays in the college admissions process:

[The article is] all about how certain universities are now accepting students based on their social media profiles, and it's not so much about your SAT scores or your GPA, but it's how social are you on social media, because that's reflective of what kind of person you're going to be on campus, and the more social you are on campus, the more likely you are to graduate. (Interview)

Avery also discussed the role that AI-powered writing assistants such as Grammarly can play in the admissions process:

There was another article that came out during the pandemic, in lieu of doing standardized tests for some international students, they had them submit papers. . . . It ran through almost like a Grammarly type of program. And all these international students were not accepted, because their English is not you know, it's not Grammarly style English, and so we talked about Grammarly a little bit too, I think that's a really good example—a lot of them use it, they use it just to check their papers and it's great. But then my question to them is well if it changed your sentence would you just go with what it said, or would you have some autonomy and say no, I think this sentence is better. (Interview)

Ethical Considerations

The third subcategory of knowledge components consisted of ethical considerations surrounding the design and use of algorithms. See Table 7 for the code frequencies within this subcategory.

Table 7

Code Frequencies: Ethical Considerations

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Bias causes	35	81%
Privacy of personal data	27	69%
Accountability	26	69%
Oppression, privilege, power	17	63%
Commodification	17	56%

Bias Causes

Eighty-one percent of all participants discussed the causes of bias in algorithms, elaborating on the ways that algorithms could rely on incomplete, inaccurate, or unfair data and how this perpetuated existing bias and societal prejudices on a large scale. Richard gave an overview of the ways bias was embedded into either the data itself or the code:

Even if you have a perfect algorithm, if the data is garbage, then it's garbage in, garbage out—it doesn't matter whether the algorithm is perfect. If you've got messy data or dirty data or biased data or gamed data, any kind of data with problems, the system will work imperfectly and actually will throw out all kinds of errors and create all kinds of issues. There are a whole bunch of issues around how algorithmic systems—on the one side, there's issues around the writing of the algorithms and the rule sets and what biases and prejudices and values and opinions and so on get embedded into the code—and then there's also the problem around the data and the quality of the data, the veracity of the data, fidelity and so on, so those are I guess the main kinds of issues. And that's why you get a lot of contestation around those kinds of issues. (Interview)

In the focus group, Becca also explained how the source of bias could be coming from either the code itself or the data:

The computers often [are] being programmed in ways that fail to recognize the weaknesses or the biases of the people writing the code, but then also the data sets that come in, the training sets that are so biased and just amplifies all of that, but I think that makes it kind of interesting because, even if you don't feel like—you know I don't really understand what's going on under the hood—you can talk about the social, kind of what you see socially, going on in response to algorithmic systems. (Focus Group)

Evelyn believed the quality of the data being fed into the algorithm was the main source of bias:

For anything that there is an appropriate data set, an algorithm should be fine right, and so the problem is not necessarily in writing the decision-making part of it, it's how good are the data that the decision making is based on for the problem that the algorithm is trying to solve. And I think that's where a lot of this stuff runs into problems is the data sets are just incomplete or they're biased in some way. (Interview)

Monica expressed a similar opinion and emphasized how harmful it can be when biased data scales:

Training data . . . when we talk about fairness and bias, how this happens is usually data is one of the main sources that this happens, like if the training data is biased—and we do have tons of examples of those—then the system will be biased. . . . People don't know how harmful these mistakes could be to some people, and how they can inflict harm particularly on marginalized communities. . . . These algorithms could be very influential and they are scaling up the society biases so they can, if there is one racist police officer deciding one thing for one person, these algorithms can scale that back up to hundreds of people, and that's very challenging. (Interview)

Jess also expressed disgust for the way data is currently being used:

The data scientific practices were fucked . . . I [left that field] for the same reason that a meatpacker might become a vegan, because I know exactly how gross the sausage manufacturing process is, and I felt that it was a moral obligation to do better. (Interview)

Kris and Richard talked about the ways that institutionalized racism could get embedded into algorithmic systems and offered specific examples. Kris cited predictive policing: “If you go somewhere to observe crime and you keep going to that place, you’re going to observe more crime there because you’re going there more.” Richard cited using previous arrest records as training data:

If your training data set is the history of arrest records, then you’re basically training the system to recognize crimes by Black people, that’s your training data, that’s what you’re going to get, your input will lead to a certain kind of output. (Interview)

Kenzie discussed the way that bias can come through crowd sourcing social media moderation, resulting in censorship:

Examples of social media moderation gone bad are really useful because it’s another way of demonstrating how algorithms can be biased. . . . So that might be talking about Black Lives Matter activists that are talking about different racial justice issues and getting their content taken down or maybe—usually what’s happened is there are a lot of people that report the content as being problematic in some way. Not that it is, but just because usually there is some kind of ideological stance against it, or maybe some kind of implicit bias, I guess, that’s going on. That expression of or that use of reporting, flagging, can sort of quickly signal to algorithms that there’s something wrong and that the content needs to be taken down, and ultimately looks like some kind of censorship and, basically is, without the intention of censorship on the platform’s part usually. So, again I think that’s really useful if you’re teaching college students, those kinds of examples in the online space, I think, are always useful because they’re very familiar to people and oftentimes they have encountered something like it before. (Interview)

Mason used a similar example of questioning what the demographics were of the people doing the labeling for what constitutes “hate speech” on social media:

It can be helpful to know that you have to show this algorithm or use this training data to show this algorithm some set of examples, and some of those examples are given labels saying, this is an example of hate speech, and this is not an example of hate speech. And

that could be helpful to know so that the algorithm can learn from those examples some pattern of what constitutes hate speech or not, and then given some new example from the millions of tweets or Facebook posts, make some guess, prediction, over whether this is hate speech or not. And then from there, well, actually who's doing that labeling and what are the demographics of the people who are doing that labeling? What kind of experiences and knowledge do they have about the language used in these tweets and Facebook posts and things like that? (Interview)

Nolen also emphasized the importance of teaching students to look carefully at the source of machine learning data to assess it for bias:

Is the system being trained on tagged examples that humans have had to tag, and now we can think about the human tagging process? Or is it somehow trying to identify patterns out of data that are just in there, in which case the human bias comes in, in a slightly different way than it would in a sort of tagging example. (Interview)

In the focus group, Madeline commented, "People talk about the bias of algorithms, but it's not the algorithm necessarily that's biased, it's encoding the bias of the people who created it."

Avery alluded to the fact that coders were a largely homogeneous group who reinforced institutional "norms" such as gender binaries:

I think that's the key, right if you could get women and people of color into the STEM field to become coders a lot of these issues, probably wouldn't exist with the way these algorithms function, it seems like it's such a male, seemingly toxic type of field. (Interview)

Nolen warned that algorithmic fairness can mean "a bajillion different" things to different people, some of which are mutually incompatible with each other. He was pessimistic about the idea of ever being able to completely extricate algorithmic systems from bias:

We often have a sense that, oh if it worked, it should be fair, if it worked in this obvious to me way, then it would be fair, that would be the right way for it to work, and so the other thing I'd want people to sort of realize is that that's probably not true. That the thing that you think is obviously fair is not going to work or it's going to be unfair in some way that you haven't quite anticipated. . . . I think one of the things you'll realize, if you think about these systems or interact with them a lot, is that that's not true, your fantasy version of the system where it's going to be not biased anymore, no, it's definitely going to be biased still. (Interview)

Privacy of Personal Data

Participants discussed the lack of user control over personal data anytime they use computers or cell phones. Foster noted, “The saying ‘if it’s free, you’re paying for it with your own data and your own privacy’ is very true.” Evelyn explained the pervasive nature of this “data exhaust”:

There’s a quote in that book [*Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World* written by Bruce Schneier, 2015] that says data is the exhaust of computing which I really love, because one of the fundamental concepts, I think, is just that anytime you interact with the technical system, there’s data that’s generated by necessity, these things don’t operate without data. Even if you just want to go to a web page, well under the hood there has to be an IP address and that IP address is unique to you and your location and. . . . There’s this aspect of—you can’t call somebody without your location being revealed, because your phone has to talk to a cell tower and your cell tower is in a physical location, and so there’s this I guess infrastructural way of looking at it, that nothing you do with technology is data free that, I think, sort of opens people’s eyes a little bit maybe, I think that’s a really fundamental concept. (Interview)

Kris commented on the widespread impact of our digital footprint caused by data exhaust:

We create a digital self, right, a kind of a representation of ourselves, through our digital footprint. And then that is then used to push information at us, to sell us things, to create pictures of how we are, how we feel, who we are as people. (Interview)

Two participants advocated for data visualization to make it more understandable to users.

Richard described “data physicalization” for climate change data, where you knit a scarf that has the climate change data in rows that are either blue or red depending on if it got hotter or colder.

Mason described “data days” for community members at a library:

They have tangible artifacts to have people demonstrate the ways that they generate data in the world, whether that’s through GPS pings or emails or Facebook likes or if they have a Fitbit—you know, health data and things like that—and there’s [also] data dating . . . where they had people create a dating app representation of yourself, but only using data that you generate so not saying anything about yourself, but sketch out your sleep data, or your walking habit data, or you’re Yelp—or things like that, trying to drive that home. (Interview)

There was consensus among participants that data literacy is an important component of algorithmic awareness, and that students were not fully aware of the extent to which their data was being used. Diane commented, “Understanding that they [platforms] are looking at your personal data and how they are using that personal data to make decisions about what to show you and who they’re sharing that data with, that’s important.” Evelyn similarly described data literacy like this: “Trying to help people understand and think about what data are collected about them and what systems can do with that information.” She offered further details about what data collection knowledge was needed:

Being aware of the kinds of data that are generated when you use systems, so knowing a little bit about how the Internet works, understanding when the systems you’re using have sensors in them and when they don’t, understanding a little bit more about what those sensors are and what they’re collecting, how things are instrumented. (Interview)

Kenzie indicated that understanding how personal data gets used in algorithmic modeling was important:

Having some idea of the way that data is collected about us and how those things act as sort of vectors for different models is really useful, because that gives a sense of or allows us to sort of imagine what’s possible with that data, how we might be identified and assessed based on that data. And so I take that as sort of the tip of the iceberg kind of literacy stuff. (Interview)

Mason stressed the importance of understanding how your actions produce data that both gets fed back to you through personalized recommendations and also gets used as training data:

Be able to identify how our actions generate data that get used as part of algorithms. . . . The potential, at least for your likes and clicks to produce recommendations or personalizations or just to train their larger model of—speaking of Facebook or YouTube or Netflix—train their larger models. (Interview)

Richard commented on his students’ lack of data awareness and lack of caution:

When I do lectures to [my students] about things like data brokers, they’re always amazed at the amount of data that’s being held about them by different companies and

the extent to which decisions are being made related to that. They know that it's there, and they know the adverts that they're getting are influenced by it, but they don't necessarily know the extent or the kind of the size of the data held by these companies, or how the data is being monetized and the consequences to them from it. . . . They are much more open than I would be. . . . I would be a lot more cautious about how I would use or buy into some of these technologies, than they would be. A lot of them are just, you know they like the convenience, they like what these technologies do for them, and they don't often question what's going on, on the other side of the equation. (Interview)

Foster also commented on attitudes of complacency when it came to personal data:

The problem isn't that we're being oppressed by a tyrannical regime of data hoarders, the problem is that we're very happy and willingly giving up all of these rights and all of our data for convenience's sake, because the services offer something for us. (Interview)

Avery offered the opinion that people often don't feel autonomy over their data anymore, but that we could teach students to regain some power over it:

A question of ownership of our data, we don't really think that we're owners of it anymore, because we just put it out there for people to use. And don't realize that, at times, it can get used against us, so I think, maybe autonomy is a good phrase that I thought of to use with students that—you know, you can take back your data if you know how. (Interview)

In the focus group, Madeline shared that she believed the Canadian “Do Not Track”

documentary series (see <https://donottrack-doc.com/en/>) was a “really useful algorithmic tool for teaching, frightening but useful.” Becca also advocated for discussing privacy protection strategies with students:

Talking about what you can do to protect your privacy is something that a lot of us have done, which in some ways can be misleading because then it kind of makes it seem like it's up to you to protect your privacy individually, and you don't want to do that. And yet, students often are doing things to protect their privacy, and they can share those things and talk about why is that important to me or what happens if I don't have this, or what would happen to certain classes of people if their privacy was invaded in certain ways, so you could maybe branch from that into some of these bigger social questions. But that's certainly something that I've done with students in the past, try installing a couple of these things. (Focus Group)

Diane cautioned that more awareness does not always equate to a change in behavior when it comes to protecting privacy:

You have more awareness about your personal data being used online, for example, it seems to indicate, you might be more cautious about sharing your personal data, but I know that's not always the case, sometimes people have a lot of information about privacy, and they still choose to share their data. . . . It's not necessarily the case that it would change behavior, but people would at least be aware. . . . I think it's one thing for someone to make a decision to share their information when they know that they're sharing it versus when they don't. (Interview)

In the focus group, Charlotte also cautioned that awareness about privacy may not correlate with the behavior you would expect:

Maybe it doesn't inspire a particular action, and I also think that because, as much as I know about privacy as a librarian, I am also lazy—I use the same password for everything, so. . . . Sometimes behavior doesn't always correlate to knowledge. (Focus Group)

Evelyn stressed that worrying about privacy is a secondary task and doesn't come into play for people to simply use systems in the way they want to, which makes it a tough sell:

I think the casual user who is using their computer for work and their smartphone and whatever they don't really need to think about it that way, in order to use these systems to get done what they need to get done. So I think that's one of the really big challenges with privacy or algorithmic literacy or whatever is in security, privacy, we talk about how they're always secondary tasks security and privacy are a secondary task. You use the system to do a thing, and security is a thing you also have to do, but it's not the reason you have to use the thing. And I kind of think algorithmic literacy sits in that same kind of bucket which is I may be using Google to find information or I may be using Facebook to keep in touch with people, and then there's this extra task of having to understand how the system works, so you understand what you're seeing and people don't want to do extra work, they just really don't. (Interview)

Accountability

Participants discussed the issue of accountability for harm caused by algorithmic decision-making in the context of current regulatory structures. Evelyn commented, “I guess the

assumption that I'm building on is that our current regulatory structure is not going to force companies to be transparent with people." She elaborated:

There needs to be a trusted source of information about how [algorithms] could affect them, that's what people really care about. But we don't have any kind of infrastructure in this country to really do that, there's—sector-based auto pilots for planes have a different regulatory agency than behavioral advertising right. And so I think this is a case where the technology has just moved so much faster than our ability to understand the implications of regulating it. So personally, I don't think the answer is to make everyone an expert on algorithms, I think it's to have people who are in authority have a better way to communicate about the implications of those things to people. (Interview)

Jess attributed the lack of regulation partly to assumptions that technology can perform better than humans. They stated, "Proposals that involved deploying new technologies often face insufficient scrutiny and analysis as a result of this cultural imaginary that just assumes that they're probably going to be . . . better at it." They referred to the European Union General Data Protection Regulation of 2018 (EUGDPR, 2018), which protects citizens in the European Union and has stricter controls on privacy than any law in the United States:

"I think the big bit is, well frankly, why we don't have European style regulation [European Union General Data Protection Regulation of 2018]. I mean, I know the answer to that question and it's a depressing answer, but we should. And it's kind of bizarre to me that we don't, that we're still debating should Facebook regulate itself, or should we do something, and we haven't even meaningfully specified what that something might be, it is such an obvious thing. (Interview)

Avery also believed there needed to be change at the policy level to spark more accountability:

Honestly, the biggest change will happen in my mind in policy. I mean, unfortunately the Senate and Congress have so many other things to worry about, but they have been trying to make strides to get some of these big platforms to be more transparent. And I think that's what it's going to take to make change—it's going to be on a policy level. (Interview)

Evelyn voiced the opinion that the government had tools to regulate online platforms if it were made a priority:

I feel like the government has the tools to deal with it, if it wanted to. And there are political and social reasons why it's a difficult problem for our government to regulate and deal with, but the tools exist to regulate it. (Evelyn)

Two participants offered examples of how accountability gets shirked. Mason talked about it within the social media domain:

I actually think even just platforms as a concept is something that I think we should interrogate critically, that these social media platforms and the way that they work at scale enables particular speeds, particular ways that misinformation gets amplified, but also gets used as a defense for why the companies can't do anything about it. So the argument at least goes that well we're just a platform, we just repost these, we don't have the resources to actually vet the content or to support content moderation, and Facebook now after a lot of protest and scandal, they've hired content moderators, but not at anything close to the scale that they need. (Interview)

Richard talked about the way algorithms were used to replace humans as drone control operators because this made them less liable for moral responsibility:

There are lots of studies now that show that, even though the drone control operators who are flying the drones in Afghanistan are based in a bunker outside of Seattle, they have really high levels of PTSD [post-traumatic stress disorder]. Because if you're flying a bomber and you throw the bombs out the back, you don't know if they hit the targets or whatever. If you're the drone operator in a bunker, you have high-definition CCTV cameras, and you literally see the people being blown up, so you know that you kill people. And so they want to move to automated systems because they want to reduce PTSD, and they want to absolve responsibility. And also, you can get operators who won't press the button. (Interview)

Several participants offered paths to greater algorithmic accountability outside of policy change. Foster believed it would be more fruitful to focus on the user perspective:

We sort of need to go away from this large policy thinking, how do we design ethics guidelines, and how do we design laws that regulate algorithmic accountability at large, and really think about the concrete moment that the process is in, in time. What does the job seeker do that wants to know why their rating is so low or so high, and how does the person working for the government agency who's using that system—how does that person respond, and what is their actual agency to do that. If you think about accountability in these very specific terms, then some of the solutions that can help us increase the accountability of the systems become blatantly obvious right away, and

much, much more clear than if you think of it in these very large and abstract terms. (Interview)

Mason also stressed the importance of allowing stakeholders to give input into algorithmic design:

People who are designing AI, how do they think about fairness and ethics and how do they get those perspectives from the people impacted by the systems. . . . [It is important] for people impacted by algorithms to have some say in their design, in how they get developed, what uses they're put to, how they get used in context. (Interview)

Nolen noted that in his own work, investigating how technology is impacting people was an important missing piece:

I research engineers, and how they think about what they're building, and I think that's important and it's not, it doesn't exhaust the question. And so, studying engineering does not actually tell you about how their technology is affecting people. . . . You can learn about what the engineers think and how they try to understand it. . . . How did the engineers try to interpret the results that they're getting and change things, but I don't want to pretend that the engineers actually know everything because they certainly don't. (Interview)

Jess expressed the opinion that traditional algorithmic audits often did not work, because they were missing the emotional aspects of how the algorithm had impacted users. Audits should also include this:

Actually engaging with the affective and emotional aspects of algorithms, and reflective practice about not just how could the algorithm be better, but how could the developers be better, so that they stopped producing the same mistakes. (Interview)

Oppression, Privilege, Power

This code referred to situations resulting in conditions of oppression, privilege, or power for groups of people as a result of algorithms. Kris suggested that technological progress was sometimes a value that was put before all others when it comes to algorithms. Kris stated, "We really do need to ask that classic just because we can doesn't mean we should question a lot

when oppression is highly likely.” Several participants commented on the power of algorithms to abuse social groups who were already hurting, thus perpetuating systematic discrimination and social prejudice. Kris commented, “Different people are impacted differently at different times by algorithmic systems, [but] they cause injury and harm at a disproportionately high rate to people who experience more oppression socially.” Kris offered this opinion:

The biggest threat or problem with algorithmic systems is that they can magnify oppression, so people who are already oppressed are going to be oppressed even harder and even greater, right, when we use more efficient systems to oppress them, which is what the algorithm is doing. (Interview)

Foster offered a similar opinion when discussing algorithms in the social services domain:

The people affected by these [algorithmic governance] systems are always in some way already disenfranchised by society, they’re already in trouble. They’re struggling with some kind of the bureaucratic state be that the criminal justice system or social security or health care, for instance, and these are the people who are then also struggling the most to get some kind of accountability out of these systems. It’s a different case if I as a White male and as a researcher in academia get struck by a system or feel like I’m treated unfairly by the system, I have different ways of sort of getting my case heard or objecting than someone who’s, for instance, long term unemployed and desperately in need of some government support. (Interview)

Richard discussed the difference in how sorting algorithms worked on people depending on their social class or ethnicity:

Sorting algorithms have a big influence because they socially sort you as to whether you do get certain services or what quality of services or what choices, and they’re on the basis of decisions that were made about you. Whether that’s related to marketing or advertising or whether it’s related to service delivery and so on . . . it makes a difference to who you are, so if I’m of a particular ethnicity or a particular race or of a particular political persuasion, then different sets of algorithms will be working on me in a different way than if I’m White and rich and of a certain social class and so on. It is differentiated on the basis of who you are as to . . . I’m much less likely to come up in a security thing or in predictive policing or as being flagged as a credit risk or any of those kinds of things, and in fact I’m more likely to be the other way and to get special offers and be treated as a valued customer and get some kind of privilege, so the same algorithm will work differently on me than it would if I was a Black Muslim. (Interview)

Evelyn conveyed the same idea of systematic discrimination when talking about facial recognition software:

[Facial recognition is] way better for White people than it is for people of color, for example, and it has the potential to really, really mess up somebody's life if they are misidentified using this technology. So I mean, I think biometric identification is something that has affected people's lives negatively, and as I look towards the future, I think that's going to be really, really hard to regulate. . . . But I think that when you get into questions of safety and security, the water gets a lot muddier and I think that very reasonable arguments can be made for using biometric identification algorithms to catch the bad guys that sort of make it harder for the people who aren't the bad guys to go about their lives the way they're used to, or the way they think they are, so I think that that has a really big potential to be a problem for us in the future. (Interview)

Finally, Mason discussed how algorithms discriminated when it came to flagging hate speech:

Hate speech is unfortunately always changing, the nature of what it looks like, and the algorithms themselves—research has shown that machine learning algorithms that supposedly detect hate speech or toxic language often flag Black, African American speakers as speaking or using hate speech more often or queer or LGBTQ speakers as using toxic language more often than straight speakers, so algorithms are not going to solve the problem that algorithms caused, unfortunately. (Interview)

Commodification

Participants discussed the underlying financial incentives and motivations of platforms using algorithms. Kris described the typical capitalist mindset: “Efficiency and productivity are the things that we care about, damn everything else to hell.” Evelyn added, “Companies . . . have a fiduciary interest in their bottom line, and a lot of these things are trade secrets.” Foster commented on the way that the capitalism mindset translates human experience into free raw material for behavioral data:

The way that the digital landscape is governed by . . . surveillance capitalism, that is something that has a long lasting and really foundational impact on the way that our economic systems work, the way that the digital landscape works, and also the way that we communicate and consume information. (Interview)

Evelyn described business models that were reliant on keeping people's attention in order to maximize their own profits, not to serve the user:

I think when it comes to for-profit companies whose business model is based on shaping people's attention so that they can make money—I think that they're, the incentives are sort of a little out of whack there, because they're not doing this to try to help people, they're doing it to try to sort of prop up the company's bottom line, and I think that that's where we are getting into a lot of problems with the way some algorithms are being used. (Interview)

Richard stressed the importance of analyzing how algorithmic systems were financed, because it would have very different motivations if it was funded by a venture capitalist versus a philanthropist:

The finance can make a big difference, like if you're building a system and using philanthropy money, that's very different from building a system funded by venture capital, where venture capitalists have got milestones around what kind of return they're going to get or what kind of points, which makes big decisions about what kind of system you'll use and how it's going to monetize the data that's run through the algorithms. But it's very different to kind of an open data, open software, open science kind of approach onto a system, so even what seems like a relatively trivial thing around financing can have huge consequences for what's actually built and how the thing operates. (Interview)

Richard also gave an example of how something more neutral like automatic license plate recognition could easily become less neutral if configured differently:

If you have automatic number plate recognition cameras, and you're then using it to socially sort the flow of cars or cars that need to be checked or whatever, or linked in to the police in that way, then it becomes a different type of system. And so it depends on how the system is configured. (Interview)

Nolen stressed the importance of looking at the underlying motivations of the people who created the algorithm:

Algorithms aren't one thing, and they don't have one kind of logic. And they are designed by people, they're maintained by people, and they're updated by people who are working under all sorts of weird constraints and ideas. People who have ideas about the way the world works and how should it work, what is my company trying to do, does my manager have a weird particular vendetta against some other manager who runs some

other bit of the product, and we are trying to beat that part of the product, all these things happen, they're all parts of how algorithmic systems actually exist in the world. (Interview)

In the focus group, both Norman and Charlotte discussed the importance of teaching students to look at underlying financial motivations. Norman commented:

Algorithms may do one thing, but if they exist in a context where they're being used to drive profit generation that's going to be really important to know and to understand, for someone to understand that oh, this algorithm is serving me information in a particular way, because these are the motivations of the people that are putting the algorithm in place and feeding this algorithm to me. (Focus Group)

Charlotte also observed:

The role of financial influence and capital in many of the algorithms that we interact with as a primary motivator to their existence, I think that is something that is lost on many of the students I work with, and once they start to understand that it just really changes fundamentally how they relate to these systems, so I think it's obscured by many of them, and I think it's important that it's made explicit as a content item that's taught. (Focus Group)

Richard shared that he used the book *How to Run a City Like Amazon, and Other Fables* (Graham et al., 2019) to help his students imagine what a city would look like if it was run using the business models of various companies such as Google or Amazon. He further described the activity:

It's about using the story as a way of looking at the way these different technical systems that are developed by these companies, how the different business models are then embedded in the technologies that they develop. You know, how do they influence how everyday life works, so it kind of draws them out and exaggerates them. (Interview)

Theme 2: Student Behaviors and Coping Strategies

The second theme contained behaviors that college students could use to help them cope with algorithmic systems. See Table 8 for the code frequencies within this category. The behavior skepticism anarchy was portrayed as a negative coping strategy that students needed to

move beyond. The other four behaviors of gaming the system, transfer of knowledge, critical evaluation, and broader advocacy were discussed in the positive context of helping students assess the benefits and risks of algorithms and actively shape them to meet their needs. The codes are listed and discussed by order of frequency (going from most frequent to least frequent).

Table 8

Code Frequencies: Coping Behaviors

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Gaming the system	19	69%
Broader advocacy	19	50%
Critical evaluation	18	56%
Skepticism anarchy	10	44%
Transfer of knowledge	9	38%

Gaming the System

Participants associated the code *gaming the system*, or manipulating algorithmic systems to get better outputs and gain more user control, with the code *algorithmic awareness*. This was reflected in the seven co-occurrences between the two codes. Foster noted that “gamers, people who play a lot of video games” seem to have a higher level of algorithmic awareness. Foster stated, “Every computer game in and of itself is also nothing else than an algorithmic system with some pretty graphics on top of it.” He further explained, “By adapting to these systems, to learning the requirements of the game or what’s necessary to win or to compete, you train your own mind in computational thinking.” Avery believed that monitoring students’ “technological curiosity . . . about how [algorithms] function” might be a way to gauge algorithmic awareness. Richard defined the term “gaming the system” here:

Some of [my students] might not necessarily know how the algorithm works, but they know how to game it or to find shortcuts through it, so they'll know game beats or they'll know how to kind of feed it data that will lead to the kind of outcome that they want and so on, so they're used to kind of interfacing with these systems all the time even if they don't necessarily know the exact mechanics of how they work, they have some intuition around it. (Interview)

Kenzie described similar behaviors involving testing and experimenting as helpful for learning about algorithms:

A big part of learning about algorithms is having some experiences with them and interacting with them in ways where you can kind of observe, test, experiment, and do a little bit of reverse engineering to kind of understand what's happening, and if I behave in this way, what happens, if I behave another way what happens, those kinds of things produce really useful insight for wrapping your head around algorithms. (Interview)

In the focus group, Charlotte used the phrase "hacking the system" to describe the same concept:

So much information evaluation is really focused on habituation and ritual, and that doesn't really work very well in an algorithmically driven environment, so really even just teaching some of those I hate this word soft skills, but these sort of dispositions about hacking the system, trying to understand what it's trying to get you to do, just thinking a little bit more creatively about what you're being served content wise, and yeah just trying to peek behind that box, a little bit and having that sort of disposition to the information that you are receiving. (Focus Group)

Monica shared that her research on social media algorithms revealed people with higher levels of algorithmic awareness were either "very actively engaged" or "they have something at stake that keeps them there on the path to awareness." She found that the aware users of Facebook were more actively posting to the system and using it actively on a daily basis, rather than being passive readers of the system. She labeled this difference as "active engagement" versus "passive engagement." Some users had been on Facebook for more than 10 years yet were unaware of filtering because they were only passively engaging with the system. Also, there didn't seem to be a correlation between awareness and technical knowledge because the

computer science students in the study were less aware. She further explained the behaviors of someone actively engaged:

Because these systems were invisible in their structure, what happened was that those who were aware are usually more engaged with the system, they were actively trying to look at signals, for example, some of them had business pages, and Facebook gives you a report like how many people saw your post, and then they can try to understand—okay, why did this post get a lot of people liking it, but not this one, was it only the content, and then they tried to figure out . . . when you post a lot you get a sense of some of my posts get more attention than others, you engage in discussions with your friends, that gave them some information, and some of them played a lot with the news feed adjustment or other settings, those are the behaviors of more aware people compared to passive engagement. (Interview)

Nolen reluctantly agreed that efforts to play with an algorithm at least superficially were an indicator of algorithmic literacy, but he pointed out the limitations of correlating the two:

Certain efforts to play with the algorithm I think are the form of vernacular algorithmic literacy that we have now, and I think they're interesting, just as cultural practices—they're not always exactly the epistemic practices you might want, the ideal things for learning, but I don't think that always matters in an everyday, on a day to day thing like do I care if someone is properly thinking of algorithms—no, it's none of my business . . . that at least gives me a sense that, they understand that they're tangling with some kind of algorithmic entity. (Interview)

In the focus group, Madeline commented that her research team found through focus groups with students that one specific behavior they used to try to game the system. She continued: “Having more than one Instagram account to be more than one kind of person on Instagram” Nolen also gave the example of using a different account on a website to try to trick the algorithm into giving you different results. He explained why this both was and was not an indicator of algorithmic literacy:

Using incognito mode in a browser or using a different account on a website to try to get some different result, that in some ways is kind of algorithm literacy, just because it indicates that you have a sense that you're dealing with kind of a fickle agent on the other end right, the problem with those is . . . you don't know anything from doing that, like you don't know anything absolute from doing that, people will be like look, I made an

account, when I set up the demographic sign up information [I set it to] a different gender, and now my results are different—a million other things are also different about those two accounts, and . . . for all you know, it's not using that demographic stuff, it's very hard to know what thing you've done is having consequences. (Interview)

Richard also described the behavior of using false data in the context of an indicator of awareness:

Often some of the ways that people [try to control algorithms] is just they put in dummy data or false data like, if you're going to connect to a Wi-Fi network in an airport and they want you to give an email address, a lot of people will put in a false email address because they don't want to be tracked. (Interview)

Richard gave further examples of gaming the system through cheap workarounds or more serious intentional fraud:

How you measure things makes a difference to how people use the system and how they will try and game the system. . . . So there's still fraud that happens across these systems, there's still kind of cheap, some workarounds that people will use. Some of them are really stupid things [for example] I was at a meal where the person next to me was [waving their arm] for like 20 minutes . . . fooling the Fitbit into thinking [they were] walking . . . they didn't want to mess up their statistics . . . of not having gone through their daily 10,000 step walk so they found a way of kind of gaming the data. . . . There's a lot of "number theater" that can kind of go on within different organizations around how you calculate different things so again, I try and use those kinds of examples [such as manipulation or fraud with police data and crime statistics, housing data, breathalyzer tests, COVID-19 data], those kind of lying with statistics, as a way to illustrate to students how you can configure the parameters of the algorithm to get different answers out. How you set up the algorithm, how you weight different parameters in the algorithm or exclude different parameters, will make a difference to what comes out the other end. So how you tweak the algorithm or how you tweak the data, you can tweak both of them, can lead to radically different outcomes. (Interview)

In the focus group, Becca described trying to game Internet algorithms as a "social push and pull":

When people are putting something online, they've tried to figure out what is, what's going to make my link go to the top, and then they'll adjust for that, the SEO thing, you'll see that at work in trying to figure out YouTube algorithms or TikTok algorithms, so there's always this social kind of push and pull . . . so kind of figuring out that this isn't just computers . . . this is people using computers to do things. (Focus Group)

In the same focus group, Jordan also alluded to people trying to work the TikTok algorithm as a form of algorithmic awareness because it signaled that they believed they had agency over the algorithm and a level of control:

In terms of metrics for what is an algorithmically aware person, it's probably at a first level it's anybody who thinks that they are an agent in any of these online experiences, or in the interactions they have with their software, that's a first step. I mean it doesn't end there, there are lots of levels of agency, but those people who are really trying to work the TikTok algorithm—that's a form of awareness, that's probably the first level of awareness that you're looking for even as you are aware, there are elements of that that are proprietary and hidden . . . it has to start with being able to see yourself and place yourself and maybe find ways to be an agent in your interactions with online platforms and software. (Focus Group)

Kenzie also brought up the importance of agency and feeling like you can exert power over algorithms even though it may be limited:

I try to emphasize the degree of autonomy or agency that we can still exercise in those spaces, so there's sort of like structure, but we're still able to sort of operate within that structure, to an extent. (Interview)

In the focus group, participants discussed two assignments for students designed to help them game the system. Madeline described one involving Twitter profiles and the feedback loop:

A colleague of mine had students create profiles and analyze the responses to profiles, and he was using Twitter for a group of students who would probably have to use it professionally at some point anyway, and just talking about how does the content of what you see change, based on what you do, to get at . . . the feedback loop, right, to see what are the impacts of this and how is the machine learning [deciding] what to serve you or what to do with your content, and what to do with you as content to some extent, so again having the students see impacts of their own actions, I think, is a really good one. (Focus Group)

Norman described another assignment designed to help students think about ways to hack the system:

Going back to the idea of hacking, I think, is also really important and how, one of the things I asked students, and this is more for an ethics class is, how could you have used

this or how could a nefarious actor have used this, and that's a really interesting and kind of fun way to get them thinking about holes in the system, how the system works, how it could be angled, how it could be biased. So that's kind of a fun way to do it too. (Focus Group)

Broader Advocacy

Participants discussed the behavior of broader advocacy, defined as expanding the application of social justice principles to algorithmic systems beyond only yourself to include the greater community. A wide range of behaviors discussed in this category were civic-related, including complaining to your Senators or Congressmen about algorithmic bias or the way an algorithm works, getting involved in activist groups and movements, voting for people who will advocate for more transparency in algorithmic platforms, submitting a *Freedom of Information Act* (1967) request, and donating to organizations such as the Electronic Frontier Foundation. Monica also recommended *AI Fairness 360* (International Business Machines [IBM], n.d.), “an open source toolkit . . . to help people examine, report, and mitigate discrimination and bias” as well as *Algorithmic Equity Toolkit* (American Civil Liberties Union of Washington, 2020), developed by American Civil Liberties Union of Washington that focuses on surveillance and decision-making technologies used by governments. Mason described broader advocacy in the context of “what we might do about” algorithmic injustices, and he gave a long list of examples:

Thinking about okay, what policy measures should we put in place, what would regulation look like, what would institutional responses to algorithms look like, I guess, what should they look like, what are institutional responses to algorithms in higher education, versus what should institutions, whether that's your school or your healthcare system or institutions like regulatory agencies, what should the Consumer Financial Protection Bureau do, what should the Federal Trade Commission do, how actually should algorithms be, given their impact in society and in shaping, once you've walked the students through how they shape every aspect of their lives, okay now what do I do, what do the learners do about this, can they lobby their representatives, can they show up at a university board meeting to say no, I don't want these proctoring systems or organize in response against these, or develop counter masks for facial recognition. (Interview)

Diane brought up a paper she had read where they evaluated students' ability to self-advocate against harmful machine learning models:

[The paper was] assessing students by assessing their self-advocacy skills, so I thought this was really interesting—they taught students about machine learning using a couple of different methods, and then they gave them a scenario and they asked them to essentially write a letter sort of advocating on their behalf, to the company that had created an AI. And they sort of laid out a specific scenario, and people were able to write better letters that were well supported by evidence under one condition than the other. But I thought it was an interesting assessment method, if the goal is to enable people to self-advocate on AI related issues, that could be an interesting way to do it. (Interview)

This blanket code contained behaviors that signaled a higher level of emotional involvement, critical awareness, and caring. Mason expressed the opinion that broader advocacy does not require detailed technical knowledge of how algorithms work:

Although that's useful knowledge, [it shouldn't be] a precursor to people impacted by AI having some opinion about how it's designed, or giving meaningful feedback into how it gets used in their lives right, so students shouldn't need to know what training data is or the fact that a facial recognition system is a classification model . . . in order to say no actually, I don't want to be surveilled by my school system, or I feel anxious when I'm surveilled while taking this high stakes test, or I am neurodivergent and it doesn't recognize my behaviors as just my normal behaviors because they're abnormal for electronics right, all of that shouldn't be conditioned on students knowing how the algorithms work. (Interview)

Evelyn pointed out that broader advocacy was coupled with a belief that you can make a difference with your actions, because algorithmic decisions are made by people and therefore these people can be influenced:

The idea that all of these things are made, they're designed and made by people, and so there is a whole range of ways these things could be designed and implemented. And so, one of the important concepts I wanted to convey [in my class] was we can actually influence the way these things work, because people are making them—it's not like a given, it's not a force of nature, these are things that are designed and made by people, and so we can demand that they're made differently or we can make suggestions about how they should be used differently . . . it's not set in stone. (Interview)

Broader advocacy also included thinking more extensively about how algorithms impacted communities (not only at the individual level). Mason explained it like this: “Being able to broaden out from just yourself to thinking about what role does this play in your community, in society more generally, and what are the harms of those.” Kenzie discussed the work that the human rights and data justice organization Our Data Bodies (n.d.) had done with communities:

Our Data Bodies have done a lot of really good work around different communities . . . they’ve done a lot of in person workshops with community members to think about more on the data side, but I think they’ve come up with some really interesting activities that are usually a mixture of moving around, reflection, discussion, and not usually at a very technical level, but putting it into practical terms and then giving people greater context. I think the idea of involving the community and different people coming together and thinking and talking together is really, really valuable, especially because in this scenario it’s possible to think about these systems that maybe have a more significant impact on day-to-day life for the communities. (Interview)

In the focus group, Becca suggested having students do a community-based project:

I’d rather have it be a little more based in experience and the agency piece, I think, is really good and maybe having—you know one way of seeing what they can do is asking them to explain it or to develop a program to teach it to their elders, or to a community group or something. And before you turn them loose, maybe making sure they haven’t made any huge mistakes or misleading people or whatever, but that might be a way to get a sense of what do they know, what do they care about, how are they able to conceptualize this stuff. (Focus Group)

Critical Evaluation

Critical evaluation, or actively questioning algorithmic decisions and seeking alternatives when appropriate, surfaced as a behavior or coping strategy. Avery stated, “What’s problematic is when we start to rely on [algorithms] as gospel, as well the output said this, so it must be what it is, without doing our own research on it.” Kris elaborated on the definition of critical evaluation:

Part of being a good citizen . . . is knowing how to think critically and how to ask questions about who, what, where, when, why. We're also in a very fractured time where a lot of people think they are thinking critically, but they are not. . . . Algorithms make things more efficient, and so we need to use critical thinking skills if we're talking about algorithmic literacy, to ask does this process need to be more efficient, why and how, what are we prioritizing when we prioritize using algorithms in this context, and is that a thing that we want and need to prioritize, or should we be prioritizing other things and using other methods of getting at this problem. (Interview)

Several different examples of behaviors reflecting critical evaluation were given. These entailed thoughtfully seeking alternatives rather than either completely desisting with the use of something out of fear or passively continuing to use something out of helplessness. In the focus group, Madeline named using VPNs as something students did based on her research. Avery gave using DuckDuckGo (<https://duckduckgo.com/>) instead of Google as an example "because it doesn't capture all this data, or I don't just passively accept all the rules and regulations." Diane suggested that tagging Facebook posts as "misinformation" was a helpful behavior that inspired critical evaluation in others. Diane also offered the example of hiring managers critically reflecting on the recommendations made by an algorithm and trying to mitigate for bias rather than blindly accepting them:

Making an extra effort to for example, look for more women candidates or whatever to sort of counterbalance that flaw of the AI system, I think that type of human in the loop interaction, where the human is sort of working with AI, taking the information that's brought to them, but then critically engaging with that and seeking out more information where it's needed, I feel like that would definitely indicate a high level of algorithmic awareness. (Interview)

In the focus group, Charlotte suggested having students reflect specifically on what they were going to do with the information they had learned about algorithms. She stated, "Because it may be, I'm going to consider it in the evaluation of it, maybe it doesn't inspire a particular action." This again hints at the paradox that algorithmic awareness may not always lead to actual

change in behavior, as discussed previously. Avery noted that students need steps that they can take to be more empowered of their data. She commented, “Students [need] a basic awareness of how algorithms can affect their education and research goals . . . [such as] saying when you Google that and you click the first link, what happens, why don’t you go to the fifth link.” In the focus group, Madeline suggested asking students what workarounds they used and also suggested a tool to help them see beyond only the first search result:

I would ask the students, if you’re aware of algorithms what do you do to get around it, what do you do to see the things that you don’t normally see. . . . There’s a search engine *Million Short* [<https://millionshort.com/>] where you can actually see the search results after the first one million hits, or 10,000 hits, which can kind of raise that and bring that kind of thing up as well. (Focus Group)

Several participants believed critical evaluation of platform settings and apps was a useful behavior. Foster suggested, “Trying to be very deliberate about the services that they use where they have the choice. Jess suggested, “Being thoughtful about your platform settings and trying to be thoughtful about what platforms you sign up for.” Evelyn elaborated on this idea:

There’s no reasonable way to really know where those data are going, and so the point at which you can exert some control is before you download it and install it. And, especially if it’s free, then you have to know that you’re paying with your data and, is this really somebody that I want to trust with information about how much I use my phone when I’m on the toilet, or something like that. . . . How much of an impact is that going to have, I have no idea, but I do think that that’s a small, lightweight thing that people could integrate into their daily lives, just when they want to download something wait half an hour and think about it, do I really need this or not, can I do this some other way that’s not going to actually cause me to give away information I might not be comfortable with. (Interview)

Skepticism Anarchy

Another coping strategy that surfaced was skepticism anarchy, or being overly cautious and suspicious of algorithms after developing a little bit of awareness, resulting in a feeling of powerlessness. Foster commented, “I do think there are a lot of people who are afraid and have a

high level of anxiety about the sort of digitalization and technology that is starting to rule our world.” In particular, Foster noted that the black box aspect of algorithms “leads to a certain anxiety and fear of the future, because it is so easy to move in a direction of a dystopian future, as we know it from science fiction.” Avery described how her students often believed they had no control over algorithms:

Algorithmic literacy practices are very difficult because we feel very like we have no control over it, and we have no power, we just feel powerless was the word I was looking for. So I think that’s the biggest takeaway for me when I was researching this was that, everybody I talked to all the students I talked to just at the end felt like there’s nothing they could do about it. (Interview)

In the focus group, Becca echoed the same thing but offered a hint of hope for moving students beyond this phase:

We found doing focus groups with college students that they felt resigned, like there’s nothing I can do about it, I have to use these systems and I can’t influence them. But when we started talking about these other algorithmic processes that work in society they got really engaged and fired up, and I think it kind of gave them permission to challenge this stuff in a way that they just didn’t feel empowered to do. (Focus Group)

Kenzie also expressed that people often feel a lack of autonomy when they have a little bit of knowledge about algorithmic systems:

There’s also then sort of concerns about autonomy, so individuals being able to exercise free will, in the face of systems that create these structures that we have to work within and that narrow our range of decision making and potentially obscure possible options that we could, decisions that we could make—make them more difficult to see, so in that way limiting the range of action, limiting autonomy potentially. (Interview)

Nolen commented, “Having that base level of technical knowledge actually might be a liability more than something useful.” For example, he continued, if you only had a beginner’s level understanding of deep learning, “you’ll be, like, oh my God, there’s going to be super intelligences that are going to do whatever.” Nolen gave the analogy of media literacy under the

Trump presidency, which resulted in the “fake news” phenomena implying that a person should be wary of all information:

This is the same shape as media literacy, what these people are doing—questioning where messages come from, and they end up in this disorganized skepticism anarchy situation. You know, QAnon and all of that. So . . . I would be curious to see where that line of thought gets you if you apply it to algorithmic literacy as well, what would be the potential negative consequences of pursuing algorithmic literacy education, because there certainly must be some. (Interview)

Behaviors associated with the “skepticism anarchy” coping strategy included (usually expressed tongue-in-cheek) getting off Facebook and taking it off your phone, shutting off Alexa because it’s tracking everything, not signing up for and downloading any new apps, throwing away your computer, or not using Google products anymore. However, participants pointed out that these were not effective behaviors. Nolen explained why not using Gmail anymore wouldn’t really protect your data:

Protecting your data is one thing, but there are a lot of ways to sort of ostensibly, like you’re protecting your data, but not really be making a difference, or people who sort of conspicuously . . . don’t use Gmail because I don’t want Google to harvest my email, except that more than half the people that I email with are using Gmail, and so they have my emails because it’s in the other person’s inbox right. So those kinds of questions which I think are similar to what we talked about a minute ago, about how a certain level of algorithmic literacy might be worse than none, because you get a kind of confidence you’re like, I know what I’m doing. But you don’t necessarily, so I think I come from a position of skepticism about a lot of those things just because I think that the things that people think indicate that they know what they’re doing might actually bite them in the ass later on. (Interview)

Along similar lines, Jess explained why throwing up your hands and saying you won’t use any algorithmic systems as a way to protect your privacy was not a realistic option:

This might be just me and my students, but one of the difficulties I often run into is there are a lot of people who talk about it through this lens of—particularly around privacy and around filtering and sorting—like oh well, you should have known that, before you signed up for the thing, just don’t use it. And honestly, a really useful exercise for those people is once you’ve done the work of identifying all the different places that an

algorithm is being used, have them not use Facebook for a week, have them try and not use their bank account for a week, see if they maybe get to a better degree that you can't just call it consent, because theoretically, [do] you have a choice if it drastically makes your life harder. (Interview)

Avery also brought up the idea of convenience as a reason for people to continue using algorithms, even in the face of learning more about how they work:

I think the problem is people don't care, because it's so convenient just to let my life be dictated by these algorithms. But once you start sharing with people what they actually do on the back end and what, you know, how they influence you, I think people change their minds about it, whether they actually do anything is different, I don't know if anybody does anything about it. (Interview)

Jess explained that someone who is cynical to the point of throwing away their computer and refusing to engage was not really exhibiting a behavior that indicates a high level of algorithmic literacy:

Internally and cynically, I want to be like well did they throw that computer away, but practically I'm like no, that's kind of not the point almost, what you want is not someone who throws their computer away, what you want is someone who simultaneously recognizes that they're engaged in a Faustian choice and that they have a civic responsibility to de-Faustify it. (Interview)

Transfer of Knowledge

Transfer of knowledge referred to when a user was able to apply algorithmic knowledge from one domain to a new domain. This was a code discussed less often, and there was disagreement over the extent to which users could effectively do this. Two participants discussed support for this based on their research. Monica reported seeing similar behaviors of "gaming the system" across different social media platforms in her research. Kenzie stated, "The more exposure we have to different algorithmic systems, the more we get a sense of what an algorithm is, I guess, and what it's capable of." She added, "Hearing about different outcomes that are problematic, or different issues that arise can be also useful for thinking sort of across different

systems and how they may also exhibit the same kinds of issues.” She explained how she observed transfer of knowledge in her social media research:

I think having information about other systems, what I’ve seen in my own studies is that people do often kind of talk across systems, they can sort of see similarities and draw conclusions in that way about how systems work, because they’ve seen them in action in other contexts usually, because I’m thinking about online algorithms, platform algorithms, all those algorithms are very similar in ways. I mean there are different kinds of systems that serve different purposes, but there is a lot of similarity across those systems, so it’s easy to make an inference by thinking about what I do, how a recommendation algorithm on a shopping site functions versus how YouTube’s recommendation algorithm—it’s kind of easy to make those jumps and probably easier than thinking about what that might mean in some other settings like doling out public benefits and things like that. (Interview)

Mason suggested that being able to transfer knowledge between algorithmic domains reflected a higher level of algorithmic awareness:

I actually think being able to apply that definition then into their lives, like being able to notice when algorithms are at work in a new setting that wasn’t one of the examples given, if they come to class with new examples or examples from new domains, I think that would be quite high for our awareness. (Interview)

Nolen, on the other hand, did not think it was possible to have much transfer of knowledge because each algorithm is so unique:

You’re not going to be like oh, I’ve learned some of the algorithms in general, therefore, I can understand how the COMPAS algorithm works and how the Facebook newsfeed algorithm works—that’s not really going to help because they’re different, they work in different ways . . . the sort of similarity among these systems is overstated, and their diversity is something that we actually really need to be thinking about. (Interview)

Theme 3: Pedagogy

The third theme, pedagogy, included opinions on how to define basic algorithmic awareness compared to more advanced algorithmic literacy, along with the challenges of assessing algorithmic knowledge and potential placement in the curriculum. Also, nine specific teaching strategies to use with college students designed to improve their algorithmic literacy

skills were discussed. The codes are listed and discussed by order of frequency (going from most frequent to least frequent), except for “algorithmic literacy,” which was placed immediately after “algorithmic awareness” for easier comparison. To see the frequency totals for all codes in this section, refer to Table 9.

Table 9

Code Frequencies: Pedagogy

Code	Total Times Coded Across All Participants	Percentage of All Participants Who Discussed a Code
Algorithmic awareness (lower levels)	45	81%
Challenges of assessment	39	75%
Placement in curriculum	29	69%
Algorithmic literacy (higher levels)	28	69%
Tactile (building)	25	94%
Folk theories	24	81%
Everyday life	24	75%
Data transparency tools	22	56%
Mass media	21	63%
Speculative futuring	13	50%
Contextualized stories	12	38%
Perspective taking	11	56%
Expectancy violations	9	31%

Algorithmic Awareness (lower level)

Participants overwhelmingly believed that a big component of algorithmic awareness was having a very basic understanding of what algorithms are, what they do, and being able to recognize them and notice them in daily life. This was reflected by the high number (11) of co-occurrences between the code “algorithmic awareness” and the code “everyday life.” Diane described awareness as being able to “recognize AI in the real world” and “being able to critically engage with the AI that you use on a daily basis.” Avery described awareness in the context of having a basic idea how algorithms functioned on platforms people were comfortable

with, such as social media. Monica explained that recognizing the existence of algorithms in everyday life and having a basic idea of what was happening was not as simple as it sounds due to their lack of transparency. Both Mason in an interview and Madeline in the focus group emphasized that recognizing when something is an algorithm versus not an algorithm was a component of awareness. Nolen suggested knowledge of what algorithms you interact with and how they work were components of awareness, although he added being aware “that algorithms are happening, I don’t think that’s a real problem nowadays” for college students. Some participants also suggested having an idea of the wide range of domains where algorithms can exist, even outside of the ones you use in daily life, was a component of awareness. Jess identified “using generic examples people are aware of, also pointing out existence in different domains they may not be familiar with” as components of awareness, while Nolen identified “a recognition that algorithms in different domains are in fact different from each other” as a component of awareness. Similarly, in the focus group, Charlotte suggested that “identification of the range of different types of algorithms out there” was a component of awareness. This idea of applying knowledge about algorithms across different domains was reflected by the four co-occurrences between the code “algorithmic awareness” and the code “transfer of knowledge.” There were several challenges of assessment articulated for measuring algorithmic awareness across different domains, including the need for different assessment schemes, the fact that algorithms were always changing, and their secretive nature.

There was agreement among some participants that awareness included a recognition of the fact that there is an underlying process acting on your data without needing to fully understand all the mechanics behind it. Foster described it as “the general awareness that

algorithms and computer systems are doing some kind of processing on us,” while Avery described it as knowing there are calculations involved for how platforms process information. Some participants agreed that another key component of algorithmic awareness was perceiving some of the potential dangers and risks to privacy involved with the collection of personal data. This was reflected in the seven co-occurrences between the code “algorithmic awareness” and the code “privacy of personal data.” Diane described the same concept as recognizing how your data was used “to make decisions about what to show you,” while Mason used the term “data traces.” Others used the term “profiling” or “surveillance.” Kris commented, “Power is a really big kind of key idea in algorithmic awareness, understanding how power relations are imbued in and animated by the use of algorithms in daily life.”

A few participants believed recognition of the role that humans played in designing algorithms and encoding bias also belonged as a component of algorithmic awareness, but others believed this was beyond just basic awareness. A few participants indicated that machine learning and the closely related concepts of training data and feedback loops should fall under basic algorithmic awareness, while others thought these concepts were more advanced. There were seven co-occurrences with the behavior code of “gaming the system” and the code “algorithmic awareness,” as discussed previously, including feelings of self-efficacy and empowerment. There were also four co-occurrences between the code “folk theories” and the code “algorithmic awareness”; for example, Jess described one component of algorithmic awareness as “thinking about algorithms, thinking about their mental models for how they think the systems work.” A few participants also included affect measures as part of algorithmic awareness. Kris commented that part of awareness included attitudes:

Their general attitudes, are they good or bad type of thing, and you can even ask something like do you think, on the whole, that algorithms are better for society or worse for society or about equal, right, and that might be proxy to some level of awareness, or some level of value or belief. (Interview)

Algorithmic Literacy (higher level)

There was not a lot of agreement among participants who tried to distinguish between lower levels of algorithmic awareness and a higher level of algorithmic literacy. Also, some participants used the terms algorithmic awareness and algorithmic literacy interchangeably. However, a few participants believed that a higher level of technical knowledge about how algorithms worked would fall under algorithmic literacy. Kenzie classified “specific processes such as Bayes classifier” under more advanced knowledge beyond just awareness. Richard referred to it as “all the mechanics behind it.” Mason suggested knowledge of how our actions generated data that trained the larger models in machine learning would fall under algorithmic literacy. Foster suggested knowing how algorithms worked and being a developer, practitioner, or domain expert were at the highest levels of algorithmic literacy. He elaborated on some of the skills necessary:

The way that we conceptualize what information is and what data is, the way that we conceptualize what the processing and transformation of this data can mean, some foundational knowledge of statistics is definitely a part of algorithmic literacy as well. To understand and explain why, for instance, certain deep learning and machine learning methods are intrinsically opaque by design, by the way that they work. (Interview)

Diane disagreed that a high level of technical knowledge was necessary for higher levels of knowledge about AI:

I think a lot of approaches to AI education for people without a CS background are sort of starting from the place of trying to get people to program an AI at some level, and I don't necessarily think people necessarily need to know how to program even at a basic level to be able to understand some key, high level ideas about AI. (Interview)

A few participants explicitly listed the ethical considerations of algorithms' impact on individuals and society at a higher level of learning beyond basic algorithmic awareness. In the focus group, Charlotte suggested, "An understanding [of] ethical considerations for algorithms, so going back to . . . the social, to me that would show a sophisticated understanding." Kenzie indicated "thinking about the problems that can arise, the unintended consequences that can come with their application in different contexts" was a higher level and "having some awareness of the issue of bias and discrimination is really valuable." Foster described the ethical considerations this way and believed it was at a level beyond awareness:

How technologies that scale impact societies . . . the connection between technology and society, the impact that technology can have, in general, but more specifically algorithmic systems can have, is also a really important part of algorithmic literacy and of course it's always a spectrum. (Interview)

Avery believed there were different types of algorithmic literacy, but a critical awareness and an idea of what behaviors can be used when engaging with algorithms are at a higher level:

I think there are three levels, there is a basic awareness of it, like I understand what's happening, there is a critical awareness, like I don't like this so much, and then there is take action and actually do something about it—I think those are three different levels. (Interview)

Mason equated algorithmic literacy with ethics and the behavior of broader advocacy. He shared:

I think the higher levels would be thinking about downstream harms, impacts, not only harms but possibly benefits, I guess, we could be a little generous here. And then, what we might do about it. (Interview)

Kenzie articulated something similar that equated broader advocacy with a higher level beyond just algorithmic awareness:

The most important part of it for most people is not that technical piece, but having a sense of what algorithms mean for me as an individual, for the communities that I'm a part of, for the country that I live in, and for the world at these different levels. (Interview)

Finally, there were a few co-occurrences between the code “algorithmic literacy” and the behavior code “critical thinking.” Kris went so far as to suggest, “Algorithmic literacy really is just a type of critical thinking.” Critical thinking was discussed in the context of understanding the root causes of bias and oppression. It was also discussed in the context of understanding the critical role that humans need to take when interacting with algorithmic systems in order to avoid taking outputs at face value before investigating the context.

Challenges of Assessment

Participants outlined several challenges in trying to assess knowledge about algorithms. First of all, there are many different components to algorithmic literacy, and they cannot all be assessed in the same way. Foster explained the different components:

There are so many different aspects and perspectives that you can have on algorithmic literacy, as we just discussed before from this sort of critical standpoint, a general understanding (the social science side of things), to the mathematical one—all of these would need very different and very specific schemes for assessment. (Interview)

Foster believed the technical side was easier to assess because it cuts across platforms and can measure broader computational thinking and machine learning. Evelyn also believed that assessing the technical aspect of algorithms was more straightforward:

Is the goal to have students know about these uses and the potential impacts that those uses have, or is the goal to help them understand the technology? I think those are completely separate. And so, if you want to know if they understand the technology, you can just ask them right, that’s just basically a test about that stuff, but if you want to know if they understand that these things are being used out in the world and the implications of them, then I think you would have to be a little bit more scenario based. (Interview)

Diane described a way to assess computational thinking through examples without being platform specific:

Ask them to sort of work through those problems in certain ways, I feel like that could also be a way of doing it in a little bit more of a platform independent way—asking them to solve a problem in the way that a certain AI algorithm would, I feel like you could potentially use that to assess. (Interview)

Kris felt that context was important for assessment. Kris commented, “When I think about algorithmic literacy, I think about how literacy in general is a contextual thing.” Monica also noted, “You need to give people some context, otherwise it would be very abstract.” Foster believed that if you wanted to assess people’s understanding of the impact of algorithmic systems, you need to consider the domain. He stated, “The potential negative or positive impact of an algorithmic system in social security or the criminal justice system requires really different understanding, than if you’re talking about a different problem domain.” Diane believed you should not separate out technical questions from the specific platforms they related to, because this would be confusing and may punish people who don’t understand technical terminology but still understand the underlying concept. She stated:

Part of the reason I think about asking about specific platforms is because, people are familiar with them, and a lot of the terminology with AI gets complicated if you sort of detach it from an algorithm. So you can ask someone to describe how they think Facebook works or your Facebook news feed algorithm, but then, if you ask them, could you describe how adversarial neural network works, most people wouldn’t be able to do that. So I think you have to figure out the right, I think something I’ve struggled with in making assessments is figuring out how to ask questions that don’t use the language of AI developers, and I think you can sort of get around that by asking people about their own experiences or asking them to describe how they think a piece of technology that they use works. (Interview)

Monica believed it was generally helpful to specify the domain of the algorithm you were assessing:

Asking if you’re worried about an algorithm predicting your credit score versus an algorithm helping you buy a shoe on Amazon, they have very different awareness levels, risk levels, concern levels . . . usually I’m more towards more concrete context situations,

but I do think you also need some of those abstract, general questions you can apply in different domains. (Interview)

In the focus group, Charlotte also offered the opinion that any assessment scheme would have to be tailored to the specific algorithm:

Some of the specifics depend on the particular algorithm, right, so I almost feel like whatever measure or scale, to some extent would need to be tailored to either a particular category of algorithms like let's say personalization algorithms or a particular platform almost. (Focus Group)

Also in the focus group, Norman described a platform-specific approach to assessing how people believed algorithms worked on Twitter from his own research but pointed out its limitations:

By mapping every single public statement that Twitter has basically ever made about its algorithms, essentially using that as a ground truth which you know, obviously the systems change over time, which is also a really tricky part about this, and then creating basically a quiz system that I surveyed people with, and there were obviously some real, significant limitations of that approach, but the trick of it is really trying to establish some kind of ground truth that someone's awareness or beliefs about the platform can be measured against. (Focus Group)

Kenzie commented that the "black boxing of algorithms" made it difficult to assess "the signals that contribute to ranking of content online." She continued, "We don't actually know what algorithms are doing in many cases, there's sort of a limit, there's a ceiling to what we can objectively know about algorithms, so it makes it difficult to measure knowledge about them." Kris pointed out that trying to create a rating scale for a platform "cements us at a point in time when this is something that's kind of ever changing." Also, Kris pointed out that creating a rating scale to assess an algorithm would ironically be "an algorithmic act." She stated, "By creating a scale, which is essentially one rule, of how to classify something or a bunch of different rules, we are kind of chasing our own tail." Kris elaborated on how to more authentically assess algorithmic literacy:

It wouldn't be multiple choice, it would probably be vignettes and ethics questions and questions where you really need to kind of consider a bunch of different positions and choose the one that works the best for the situation you're in type of thing, so very context heavy questions. (Interview)

Diane also believed that specific examples that cut across platforms might work to evaluate the ethics component of algorithms:

Asking people also about the ethics, asking people to critically reflect on algorithms as well, or if they think there's any problems or challenges, I feel like getting into that sort of critical inquiry building is also an important component, and that's definitely platform independent—I feel like you can give people examples. (Interview)

In the focus group, Becca cautioned against assessment that was not genuine. She described how faculty reacted to the *Information Literacy Competency Standards for Higher Education* (American Library Association, Association of College and Research Libraries, 2000):

We did a week long workshop with faculty in 2000, and it culminated with reading the new at that time *Standards for Information Literacy* which I thought were kind of nifty, because we hadn't had anything like that, and they were looking at them going like oh, my God, I'd fail all of these things, but somehow I'm able to write scholarly books, what the. . . . And I could see a student passing, proving that they could do these things, and still couldn't do a decent job of research, so I guess I've ever since been somewhat biased against the, like this would show that you can do this thing, I'd rather have it be a little more based in experience and the agency piece. (Focus Group)

In the focus group, Madeline also remarked on the difficulty of measuring authentic algorithmic literacy:

If we're measuring it so that we can stamp them as being algorithmically literate and send them on their way, then answering some questions in hypothetical situations may work. If what we want to do is build algorithmic awareness so that they can be more conscious agents in their interactions with information, then the evaluation would have to look quite different. (Focus Group)

Self-reported knowledge was discussed as one way to assess algorithmic literacy, but it had drawbacks. Monica commented, "For measuring awareness, your question should not be leading, because people might not know, but they will think they know." In the focus group,

Norman remarked, “That connection between sort of belief, and actual knowledge is really tricky to navigate.” Kenzie stressed, “There are people who underestimate or overestimate their skills, so I imagine that would also be the case for algorithms, so my guess is that these kinds of self-reports are not the best measure.” Diane suggested measuring perceived knowledge both before and after an instructional activity; this could be accomplished by “asking them to assess before and after the activity, how much they felt they knew about this” activity topic. Kenzie suggested one strategy used in digital literacy for minimizing the unreliability of self-report measures:

You ask people about a series of terms related to digital technologies and ask people to rate their familiarity with them, and then you throw in a fake one just to gauge how much people are maybe fudging their answers—so that’s one way just to get at something that’s maybe a little bit more reliable than a straight self-report of how knowledgeable are you about algorithms. (Interview)

Kenzie also suggested observing people perform a task and “finding ways to measure whether they have successfully demonstrated the skill” by looking at things like the timing, completion rate, and “maybe even self-reports about how difficult the task was.”

Participants also discussed self-efficacy as a measure rather than perceived knowledge. Kenzie suggested that self-efficacy might be a more appropriate measurement because it captures the degree of empowerment someone feels. She explained that a high level of knowledge could equate with when someone feels “that they can confidently and successfully sort of work within algorithmically mediated spaces and get what they need.” She continued:

If there is that foundation that I can accomplish, I can be successful, I can satisfy my needs and interests with the knowledge that I have, then that’s at least an adequate level—but because we all have different needs and interests, it’s difficult to create a continuum of knowledge, where how somebody behaves would be the model for what we want for people to be algorithmically literate. (Interview)

In the focus group, Charlotte commented that behavior doesn't always correlate with knowledge and awareness, so measuring the frequency of a behavior may not be a good measure of knowledge. She gave the example of a user who reported every ad they received as inappropriate:

[Giving an example of user behavior] I just hit report for every single ad thinking that they'll go away, but you just end up with weirder and weirder and weirder ads, so if it was just about frequency of behaviors like muting, reporting, or trying to manipulate the algorithm, I think that doesn't always correlate to sometimes that analysis. (Focus Group)

There was disagreement about whether it is a good idea to assess for developmental levels of algorithmic literacy or awareness. Kenzie shared, "I try to resist the instinct to create different levels of knowledge." On the other hand, Avery suggested you could create an assignment or prompt and "develop a rubric with you're a novice, you're an expert," and Mason believed it was possible to distinguish between developmental levels akin to moving up Bloom's taxonomy. Foster also discussed "a certain taxonomy that starts with I have no idea that algorithms exist" all the way up to higher levels of being aware how they work and being a developer, practitioner, or domain expert. In the focus group, Charlotte also suggested creating a developmental rubric to assess algorithmic literacy:

I think a criterion reference system very similar to TATIL [the Threshold Achievement Test for Information Literacy, Carrick Enterprises, 2018] would work really well because that was designed basically with experts doing exactly like what we're doing—what would a high level of sophistication look like, and then looking at that developmentally across the student samples. I also think potentially using rubrics to evaluate student work for particular traits that we've identified would work. (Focus Group)

Placement in the Curriculum

Participants offered several plans for where and how algorithmic literacy concepts could be embedded into the curriculum for college students. Foster urged that these concepts should

come as early as possible. He stated, “It’s so foundational about a certain way of thinking about the world and technology . . . it’s very good for them to have that knowledge in the back of their head, while they are learning about other things.” Most participants suggested that the present situation was not standardized in terms of students receiving algorithmic literacy instruction, and that it happened only haphazardly within certain programs of study that did not reach all students or was through “optional” or “elective” courses. Richard explained that at the university where he works, aside from science and technology studies degree programs that included it, there were a set of “tech society courses” within some other degree programs such as anthropology, geography, sociology, and business that touched on algorithmic literacy, but they were not formally organized or connected:

They link to particular academic programs, so you might have somebody in law teaching about privacy and surveillance, you might have somebody in sociology teaching about kind of digital games and online sociability and web communities and so on, and then you might have me who’s doing more kind of smart city and kind of algorithmic governance courses and so on, but the only students who do those courses are within my degree program. (Interview)

Diane suggested a similar approach because “it’s so hard to get a new class into the curriculum.” She described an approach to infuse algorithmic literacy into existing courses across the curriculum:

I think incorporating algorithmic literacy into courses that we already have could be a really useful way . . . algorithmic literacy, it comes up in all sorts of contexts and incorporating a unit on algorithmic literacy in classes, in the context of journalism or sociology or I feel like you can learn about algorithms in all sorts of contexts because they’re using all sorts of contexts—political science, I feel like that can be a really useful inroad . . . integrating it is as a lessons across different sort of interdisciplinary curricular could be a really successful way to reach a broad audience. (Interview)

Avery described a similar approach of infusing it into different disciplines or domains wherever it made sense rather than proposing a new course, and she suggested training the faculty on how to do this:

What I've been thinking about trying to propose is a workshop for faculty, so if your discipline is music, how do we incorporate this into music, or your discipline is science, how do you . . . there's many ways to incorporate it into science, and in my college we have a computer science major, but they don't even talk about it in the computer science major which is crazy to me, we don't even have an ethics class it's crazy. . . . Because it wouldn't have to be, faculty are so protective of their courses and their time, but a half a class conversation, I think that would really help. . . . I taught a race and ethnicity class that had nothing to do with algorithms but I brought in this topic of policing and the criminal justice system and how they use these algorithms, and even that little lecture on it, I think really influenced students, because a lot of them wrote papers on it later. (Interview)

Three participants suggested algorithmic literacy could be integrated into part of one new required course for college students, but not as the primary focus. Nolen suggested it could be part of a required general science and technology studies course focusing on the social aspects of technology, including its design, deployment, and use:

I think that algorithmic literacy is one part of what should be a more general science and technology studies requirement for college students . . . when I get students in my intro STS [Science, Technology, and Society] class, they are just mind boggled that you might think about science as being a thing that humans do in particular contexts, or that technologies are not just about increasing efficiency and applying scientific knowledge to the world to make life better or whatever. And so I think within the context of that, algorithmic literacy makes a lot of sense. (Interview)

Jess suggested algorithmic literacy would fit into a required course on infrastructures:

I think it needs to be broader than just algorithmic literacy and could maybe be described as 'infrastructural literacy'. That is: embedding an awareness of the world we take for granted, and sensitizing people to the work that goes into them and how dependent we are on them. (Interview)

Kris suggested algorithmic literacy would fit with a required course for undergraduates on search skills:

I think that every undergraduate student should absolutely have to take a search skills class, and I think it should be covered there . . . we need to teach people search literacy and search skills, and I think that this would be a natural piece of that class because search is 1,000% algorithmically mediated. (Interview)

Richard suggested a required course on basic data literacy where “it really wouldn’t necessarily just single out algorithms, it will be a wider thing around kind of digital literacy, data literacy, maybe algorithmic literacy.” He elaborated:

I think that things around basic data literacy is something that every student should know, they should know something about how to handle data . . . around that, you can fit digital literacy in relation to that. . . . They should all have some basic stuff around data security and around ransomware and around phishing and around all that kind of stuff because they potentially have high consequences for them. And I think within that you can fit in a wider kind of reflexive—even if it’s just one class—around how these kinds of systems are consequential for them in different ways, so how digital tech is mediated in their life and what kinds of things they might want to look out for, if they get interested in that, then they can go off and do the specialized courses. (Interview)

Other participants also suggested a new required course that could include algorithmic literacy, but one that was taught by a cross-disciplinary team to represent both the technical and ethical side. Foster explained, “You need domain knowledge from both sides (and) interdisciplinary teaching teams.” He further explained why both sides were needed: social scientists and scientists from the humanities understand “the problematic nature of algorithms” but “are really running into that wall of algorithmic literacy when it comes to some of the technical underpinning;” and on the other hand domain experts in computer science “have an incredibly strong algorithmic literacy on the sort of technology and foundational mathematics side of things” but are “completely and utterly oblivious of any connections to the social sciences or humanities on what these technologies do, how technologies that scale impact societies.” Evelyn suggested a new required general course on privacy that could include elements of algorithmic literacy, but that also took an interdisciplinary approach to teaching:

I feel like something like that [general course on privacy] could be implemented at the scale of a university or multiple universities, but it requires, I think that if a CS person were teaching the class the focus would be way more on getting the technology right, and so it requires I think a multi-disciplinary approach to communicating about it. . . . The idea, I think, at least for me, would be to get students to engage with discussions about how the way the technology is designed impacts how it is used and how it affects people. (Interview)

Finally, there was some support for integrating algorithmic literacy into information literacy instruction led by the library. Kenzie suggested that this would be a good fit:

A natural fit would be through information literacy instruction and libraries, but now all the libraries have been shut down and so that's a problem [in K-12 . . . that would be a really natural fit where the curriculum could occur. (Interview)

Richard described a set of “optional student orientation courses” run by the library at his institution around information literacy and digital literacy where algorithmic literacy could be included. Along similar lines, Mason remarked, “Why algorithmic literacy and not data literacy or information literacy?” Avery commented, “I think algorithmic awareness fits right in with information literacy,” and she suggested a required freshman course “through the library” or even “workshops or seminars” that students were forced to take. She noted, “I think different libraries are starting to accept it's a big part of information literacy,” but she added “I know most universities wouldn't force something like that” to make it required.

Tactile (building)

Almost all participants in this study (94%) discussed the benefits of tactile, hands-on engagement with algorithms as a teaching strategy. Foster explained, “One of the most prevalent ways of teaching computational thinking is teaching it by doing it.” He continued, “It teaches a certain understanding for how we need to abstract reality in order to make it fit within these computational constructs that algorithms are.” Diane advocated for “low-tech, hands-on

activities” such as the ones in *AI Unplugged* (<https://aiunplugged.lmc.gatech.edu/>), which has “some existing activities that people can do that I think are really—that can be a fun way and, a not so intimidating way, to learn about AI concepts.” Jess explained the benefits of applied learning, using the example of building a computer vision system and how this would help debunk the idea that algorithms were foolproof:

I think something really powerful is to have people actively engaged in the work of building a system, it could be something as simple as a computer vision system that recognizes cans of soup, but I think that that kind of applied material learning is really useful. First, because obviously it’s a different form of learning and so it’s good for engaging people who may not enjoy reading 300 papers. But second and, most importantly, because of that whole cultural myth of we should defer to AI because AI is smarter than us, well I haven’t found any way better to convince people that that’s wrong than to build an AI system. . . . I think a big part of convincing people that algorithms can be fallible is having them build them and see how easily these things break, and also how much uncertainty and ambiguity and just plain sweat gets packed into building even a simple system and then being like okay . . . you have a system to detect cans of soup, it’s only right 70% of the time, and it took five of you a week, imagine if it had to detect every single object, would you trust it? I think that’s a big part of it is demystifying it. (Interview)

Diane stated, “This idea of being able to program or teach something and then test it out is important, the programming need not involve actual code programming.” She remarked that research illustrated having people do “physical embodied simulations with algorithms” lowered the intimidation factor. Diane described an activity for families where participants built a semantic network with wooden tiles, then trained an AI chat bot to answer questions based on the connections and relationships in their network. The activity helped them realize the limitations of algorithms:

I had them sort of construct semantic networks, which was objects connected to the ideas—like cat has a tail, dog has a tail. And you know dog and cat are mammals, and you sort of connect these concepts with their relationships . . . so they built their network with wooden tiles on a play mat and then took a picture of it, and then there was an AI chat bot that could then answer questions, using the network. So you could ask it what is

a cat? and it would be like a cat is a mammal—if that’s what they put on their board—or they could have put a cat is a cow, and it would have said that. . . . Some people were kind of surprised, there’s this chat bot, there was this board, and making the connection that, oh, this chatbot actually literally only knows the things that I have put on this board, it doesn’t actually know anything else. (Interview)

Several participants gave examples of physical enactments of algorithms to make them more concrete. Mason described children acting out algorithms by making a peanut butter and jelly sandwich as a good way to teach without involving code:

You have children act out an algorithm, for example, teach me how to make a peanut butter jelly sandwich with an algorithm, and the kid is like you put the peanut butter on the bread, and then the person grabs a scoop of peanut butter with their hands and is like oh, you didn’t tell me to pick up the knife, to try to get explicit about steps. (Interview)

Avery did a similar thing with her college-level students, but used a grilled cheese sandwich as the example to illustrate computational logic. She also used whiteboards for a second activity to illustrate algorithms’ rapid speed:

This is an example I just did in class the other day, I showed them the steps to make a grilled cheese, and it’s like here’s a mathematical function, I start with bread, then I put the cheese on it, I fry it, and now I’ve done this—to kind of show them that it’s really just a step by step process; another thing I’ve done with them is I will get whiteboards and have them all write random numbers, and then I will say in 10 seconds get in order, and they have to figure out an order to get into. And then we can say get in a different order, and sort of explain it like yeah, that didn’t take you very long, but imagine if there was eight million of you, and I said get in order in 10 seconds, how would you possibly do that—you couldn’t, that’s why we need these algorithms to do it for us. (Interview)

Other examples of hands-on learning were shared by Richard, who took his students on a physical field trip to the control room to help them understand intelligent transport system algorithms, and Evelyn, who brought in an old laptop from IT, took it apart in front of students, and passed around the hard drive while discussing “the places in a computer where data lives and is stored . . . we talked about all the different ways that technology can determine your location.” Diane suggested interactive museum exhibitions to communicate basic AI concepts. Two more

participants recommended using Google's *Quick, Draw!* (Jongejan et al., n.d.) to let people test out how to train a machine learning algorithm. Monica explained, "People could play with that, and that was a very useful tool, because they would learn about training data, you draw a few shoes and it learns what is a shoe." Jess suggested machine learning projects on *Scratch* (<https://scratch.mit.edu/>) although they warned "I suspect they're pretty abstract." In the focus group, Jordan described a less abstract, pseudo code exercise designed to teach how if-then statements worked and also get students to reflect on biases that creep in as you are encoding the rules:

We do a short pseudo code exercise called "Programming the Library" where they're just putting together a series of if-then statements about—you enter the library, where do you go from here, how do you get from there, and then at the end of that exercise you've kind of got this idea of oh, this is how you might write for a computer program. And then you're able to pull apart okay, why did you decide that you went to the second floor, and it's because that's the social floor, and this is a group that really wanted to connect with people, so that's a bias that gets played out in the way that they encoded a rule about entering the library. So transparency, defining the initial concept and then, just what does a statement look like when you make a statement for a machine that it will follow, what does that mean, what does that look like, how are there biases that come forward in that moment. (Focus Group)

Folk Theories

Folk theories, or informal theories that people develop to explain how algorithms work, were discussed by participants. Jess described the concept as "thinking about how we think about algorithms" or "reflective practices." They noted, "One of the things that's really interesting is that a lot of the algorithmic awareness literature doesn't look at algorithmic awareness directly so much as it looks at folk theories, how people think about these systems." Kenzie described it as "mental models" and noted it was "a way to kind of work around getting at awareness." She gave the example of people attributing the fact that they don't see all social media posts from their

friends to their own inattentiveness as an example of a folk theory. Kris explained further the concept of folk theories:

We don't generally know how these things work it's all proprietary, so we come up with folklore . . . people just, we make up stories for why is my computer doing this weird thing to me right now, we make up a story that kind of fits right. We think it's the algorithm . . . it has become part of the discourse where people are like oh my algorithm is doing this, and it's like is your algorithm doing that? What do you mean and what do you think an algorithm is? But it has kind of entered into the popular discourse in a lot of ways, especially in systems like TikTok or Instagram or places where the algorithm does decide what content you're looking at. (Interview)

Mason discussed folk theories as one way to measure algorithmic awareness, and he referenced literature on the folk theories of social media super users:

Some researchers who look at this idea of folk theories of algorithms, so [for example a researcher has been] looking at how content creators, whatever that means—like people on YouTube or TikTok, or social media influencers on Instagram—how they think the algorithms are working to either recommend their content or to move them higher in somebody's sort of visibility . . . but that would be another way [at getting at algorithmic awareness], okay, why do you think that you saw this particular type of content, and to sort of back out their mental models and folk theories. (Interview)

Monica also discussed folk theories in the context of measuring operational awareness of algorithms:

One way that we [measure operational awareness] is through folk theories, where we try to give scenarios to people, you cannot generally ask how do you think the algorithm works, and one way to measure the level of awareness is that you try to map it with what actually is out there, and the system says they do, but most of the systems are not open, so you actually don't know, to justify it—but one way is to measure the level of, we call it perceived awareness—you ask people, you give people scenarios and ask them why do you think it is happening, . . . we tell them to look at this story, this story did or did not appear, what do you think is happening, and that helps them to try to talk about the theories they have. (Interview)

Nolen expressed frustration with researchers who try to improve algorithmic awareness through the testing of various folk theories because he maintained the systems have too many unknowns to be able to claim correlations for how they work:

Lazy, semi-quasi experimental efforts to understand how these systems work, so some methods that involved trying to Google for different things, or trying to get certain results on Spotify and being like see, look at the results that I have. Because most people who are trying to do those things don't have a very vigorous theory about what's happening on the other side to cause those results. And so they're very, very, very constrained in ways that I think the people who do them don't often appreciate. So it's a tricky space, because I think, in many cases, the people who are doing the studies are working from actual real concerns . . . there are severe limits to what you can learn, I think, from that kind of approach. (Interview)

Several other participants pointed out that people's folk theories are often incorrect. Evelyn noted, "I started noticing these issues that people were having with making assumptions about how these systems worked that didn't really seem like how they actually worked." Diane commented, "People have folk theories sort of about how algorithms work that are not, and don't necessarily match up with reality, but people do develop a theory of how they think it works." Jess emphasized that having people reflect on the possibility that their folk theories were inaccurate and what the implications of these inaccuracies would be for how they interacted with algorithmic systems was key to using this technique to further algorithmic awareness. They explained:

Have people think about the fact that their understanding may or may not be accurate, but whether it is or not it's going to shape their perspective on these systems in certain ways, and they should think about what they're thinking about . . . that gets back to less about the technical and more about the cultural and also the self, so thinking about how I understand AI and thinking about how I engage with AI as a result. So what I'd look at is almost sort of reflective practices . . . the sort of active reflective exercises, the thinking about like okay, I don't know how the algorithm works, but I know how I think it works, and I should really think about how that informs how I use it and what the consequences are if I'm wrong, I think that's an important part of it. (Interview)

Other participants discussed folk theories directly in the context of a teaching strategy with college students. Richard articulated this idea when he stated:

I would probably just have conversations with [students] about their interactions with digital systems and how they understand what those systems are doing, and their kind of

general sense of how they work, and what the consequences of the outcomes might be. (Interview)

In the focus group, three participants also described teaching strategies that would probe at students' folk theories. Charlotte advocated for students to engage in perceived knowledge of how algorithms worked through journaling and critical reflection:

Because so many students do engage with these systems, they have knowledge about how they perceive them, how their feeds are customized, also engaging students in critical reflection or journaling on their own search habits or information habits, when you sit down and think about it, I think, students have more knowledge there to share, so I think engaging in critical reflection is also a great strategy. (Focus Group)

Madeline discussed the challenge of trying to measure students' agency over algorithmic systems since it is largely internalized, but concluded that reflective processes such as writing about how they worked or having students teach others about algorithmic awareness might be good strategies:

So it's that difficulty that we have with all sorts of information literacy evaluation, is what are we evaluating and why and can we evaluate authentically what we think is most important, which is . . . building agency . . . I'm not sure that there is a way of effectively doing that. I think reflective processes, I think some of the critical reflections on use, I think having students do some sort of writing or explanation or teaching others about algorithmic awareness, I think that to me gets a little bit more at an interiorized knowledge and perhaps agency than an exteriorized knowledge. (Focus Group)

Norman agreed with Madeline that explaining to others how algorithms worked, or applying their folk theories to explanations of how they worked, was a good teaching strategy:

I do think that being able to explain to others, is something that—again setting aside the issue of the gap between belief and actual accurate knowledge, but explaining to others is really a critical, in my view, demonstration of knowledge, correct or incorrect. (Focus Group)

Everyday Life

Many participants stressed the importance of helping users recognize the algorithms they use and come into contact within everyday life, and they felt this was an important component of algorithmic awareness (reflected by the 11 co-occurrences between the codes “everyday life” and “algorithmic awareness”). Foster believed that the lower level of the taxonomy for algorithmic awareness would include “I’m aware of some examples of algorithms that affect me in my personal life.” Mason advised, “You can draw on examples from people’s daily lives and the kinds of systems that they interact with, maybe these are social media examples or entertainment like Netflix or eBay, or recommendation algorithms.” Diane commented that showing people where AI is already working in their daily lives makes for “super relatable entry points to start talking about AI.” In the focus group, Jordan explained one of the reasons he likes to teach about search algorithms was because people “practice that, they do that day in, day out.” Jess recommended “using generic examples that people are aware of, so reasoning through together how Facebook sorts things and what the likely consequences might be” along with “the user walkthrough, the guided user walkthrough of a platform they’re familiar with to spot the places where an algorithm surfaces.” Richard stressed that you can teach people a lot about “how these algorithmic systems affect their everyday life without them knowing anything about what the algorithm actually does or functions.” Kenzie also remarked, “Most people don’t need a terrible amount of complexity to get a basic sense of what these systems are doing enough to help them in their daily lives, basically we don’t all need to be computer scientists.” Avery explained that using algorithms students interacted with every day was an effective teaching strategy because they will care more:

You'd have to direct them to certain platforms that they are comfortable with, because I think that's a big part of it, it's like, if I don't use it, you know if I don't use Wikipedia, I don't really care that there are robots that edit Wikipedia and are changing words that shouldn't be changed, but everybody uses social media . . . I'd be hard pressed to find a person who's not at least tried it out for a little while. (Interview)

Several participants stressed that highlighting everyday algorithms was so important because they are often invisible to users. Richard stated, "Algorithms are increasingly mediating everything, and they can be backgrounded" or "indirectly rather than directly mediating your everyday interactions and behavior." Richard gave the example of tap water as something that may be indirectly mediated by "a technical system behind it that's measuring water pressure, and water quality, and water distribution." Monica also expressed concern that people were often unaware of algorithms:

Right now, the design of many of these systems are so opaque and invisible it's like we have an invisible structure that people don't even realize that they are consuming algorithmic content. . . . You're getting affected by these algorithms every day either public or private, and we don't know. (Interview)

Diane also expressed the same opinion:

People don't necessarily realize that they're interacting with AI, when they're on social media or scrolling through their Netflix . . . giving them some examples and asking them if they use it and then talking about how those examples use AI . . . [they may be] sort of surprised, they're like oh, I didn't think about this, I didn't think that this necessarily was using AI. (Interview)

Participants offered class activity ideas around everyday algorithms. Richard shared his audit assignment:

We get students to do audits of all the digital technology and data systems that they come into contact with from when they wake up to when they go to bed, and pretty much everything is digitally mediated even if you're not necessarily aware that it's been digitally mediated. (Interview)

Mason suggested having students write about their everyday encounters with algorithms:

Have people create a diary or a journal of encounters that they have with algorithms or what they think are algorithms and then contextualizing that, because some of this might be more apparent, more transparent, to the learners as the end users of technology if it's on your phone, maybe it has the face ID, like it recognizes your face or something like that, but even predictive text right or smart home devices, those might be sort of low hanging fruit ways of getting at early experiential knowledge . . . or diary entries of any systems they interact with. (Interview)

Nolen suggested a class discussion around everyday algorithms:

I would ask my students what algorithms they've interacted with in the last week and see what they say. Because I think that would tell you a lot about what they think algorithms are, what they don't think algorithms are, and then that would probably be a starting point, then, for a discussion about how do those algorithms work, are they good or bad, how might they be changing, how might they be different from other ways of answering similar problems. (Interview)

Data Transparency Tools

Participants advocated for the incorporation of information into user interfaces to visibly show what's happening under the hood for how algorithms work. Kenzie explained why this was helpful:

Algorithms are really abstract in a lot of ways, so activities that I think can make them easier to visualize . . . at least having that mental model in your head of some kind of basic schematic of what's going on, can be really useful—so activities or instruction that allow for that can give you more concrete objects to hold onto. (Interview)

Evelyn offered more details about what data transparency tools looked like and their purpose:

What I've been doing in my research is trying to figure out ways that companies could incorporate lightweight information into their actual user interfaces to try to help people understand a little bit more about what's happening under the hood . . . ways to incorporate little bits of information through people's normal use of these technologies to kind of help them understand . . . what kinds of signals can we give them that help convey a little bit more about what's going on . . . so ways to help people, and this was inspired by a study that I did, where I interviewed people about activity tracker use Fitbits and smart watches and stuff like that, and the things that really sort of encouraged people to start speculating about what the heck was this thing it was actually counting that's like a step, right, because it's not actually counting steps. It's taking accelerometer measurements and then processing them using some kind of algorithm to infer whether a step has taken place or not . . . how can we trigger people to speculate about what's going

on behind the hood because, if we can get their attention, then they're motivated to think about it. (Interview)

Participants gave several additional concrete examples of data transparency tools. In the focus group, Charlotte shared that the *Instawareness* tool (Fouquaert & Mechant, 2021) showing the invisible curation algorithm of Instagram “is always a hit with my students.” Kenzie described a facial recognition tool called *How Normal Am I?* (Schep, 2020) that assessed your face:

You can grant it access to the camera, and then it will do assessments of how attractive you are, maybe personality traits, things like that, that can sort of be computed or inferred by facial analysis but that are . . . the point was that there is lots of imprecision there, that it's not a straightforward, objective process. And so again, putting you in the place of being in front of this actual system and seeing the outcomes that it might produce, and you having as your face—being the data input—it's a good way to sort of give some insight about what's going on as aided by these little instructional cues and things that give greater context for what's happening and what you're seeing. (Interview)

Diane described a tool for emotion recognition called *EnableX* (n.d.;

<https://www.enablex.io/cpaas/faceai/>):

I saw a presentation recently, which is maybe part of why I'm thinking about it, on emotion recognition where they did some demos there, some free demos online of emotion recognition AI that really sort of highlights how limited it is in terms of being able to recognize emotions from looking at your face. (Interview)

Kris described interactive resources like *The Myth of the Impartial Machine* (Feng & Wu, 2019) that allowed you to change inputs and see the outputs visualized in domains such as predictive policing:

There are some really lovely interactive web resources . . . where you can play with data, and it will give you visualizations that show you exactly how, for example, with predictive policing . . . you can actually slide number of visits and you can see the changes in a data set that happen as a result of changing some of these different inputs in an algorithmic kind of program. That kind of thing I think is really lovely because it makes it really clear what is happening to people. (Interview)

Monica described a tool she built for social media that helped her gauge people's lack of algorithmic awareness:

I built a tool to show what I can see versus what I cannot see because of algorithmic filtering . . . and I compared it and said wow I don't see a lot, and then later I just gave it to people to play with, and what I realized was that people even didn't know there is an algorithm at work, so [I discovered] this lack of awareness about the presence of algorithms and also the impact they have on people's attitudes or approaches towards the systems. (Interview)

In the focus group, Jordan described having students look at their ad personalization profiles on Google:

Even just having them look at what kind of profile—you can use some of these tools, you know, if you go and download your data and really take a look at what demographically they've scoped you as, that can be one of these moments where the students just go I had no idea, I had no idea Google thought, and sometimes the accuracy of those profiles, it shakes you right. You look at it and you think they got this completely right. . . . Using some of those data transparency tools that they've built into their systems, out of necessity, and sometimes policy that they have to be transparent, those are great tools. (Focus Group)

Despite general enthusiasm for data transparency tools, Monica warned about the challenges they pose. She asserted, "Most of these tools have these challenges, so they don't live for a long time." She observed that the area of "algorithmic transparency, interpretable AI, or explainable AI" was exploding, but one challenge was "what are the things you can say, what do the algorithms want to accomplish, even the developers don't know how they work." She also described challenges with building the tools because they change so rapidly and can present legal difficulties in adhering to terms of service agreements:

Either you use API, or you have to scrape, and either has—one is sometimes not following the terms of service of a system—which [a recent lawsuit] actually allowed researchers to do auditing and scraping for auditing . . . but sometimes, even if you don't break any rules but you use API, they can just shut down some features of the API and then you can't work anymore . . . this is a black box, you need to access some information to build such tools to bring awareness, and that black box—either you don't

have access to, or even you get access to some part of it, they can shut it down at any moment. (Interview)

Mass Media

Several books and documentaries related to algorithms have already been discussed in previous sections (e.g., *Algorithms of Oppression*, *Automating Inequality*, *The Black Box Society*, *Coded Bias*). Cathy O’Neil’s (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* was brought up by two participants as a book that could be comprehensible to the general public. Avery commented that she wished there were more mainstream articles about algorithms. She said, “I think the general person doesn’t want to sit and read *The Black Box Society*, so I feel like there has got to be more mainstream articles about it.” Participants discussed using news stories, current events, or other mass media related to algorithms as a vehicle for discussion. Diane advocated for science communication outlets offering podcasts that “explain at sort of at a higher level for a broad audience how algorithms or other scientific or technology concepts work, I think that’s a great resource.” Kenzie suggested a news story about the Microsoft “Tay” Twitter chat bot “that quickly became horrendously misogynistic, racist, White nationalist, after interacting with Twitter for just a few hours,” and she thought this story was especially effective if you’re dealing with a population very familiar with online culture. Monica recommended the *Propublica* series (“Machine Bias: Investigating Algorithmic Injustice”, 2015-2022) that covers examples of bias in different systems to engage people in discussion. Evelyn described how she used news stories in her data class:

I also looked up new stories, so we could talk about sort of okay, what’s going on in this news story so that . . . several years ago when there was the terrorist shooting in California and there was that whole controversy over whether the government could force Apple to break the encryption on the iPhone—so that happened while I was teaching this class, and that gave us a really great opportunity to talk about encryption because it was

in the news and no one knows what encryption is, it's really complicated math and stuff, but you don't need to understand the math to understand what encryption does. And so that was how I tried to do it when I was teaching this class. (Interview)

Diane expressed the opinion that having the ability to discuss news stories involving algorithms was a higher form of algorithmic awareness:

Being aware when you hear about something related to algorithms or AI in the news—and being able to sort of comment on that and relate that to your own experiences or develop an educated opinion on that information would indicate a high level of algorithmic awareness. (Interview)

Several participants also suggested self-guided tutorials freely available on the Internet as a good source of information about algorithms that could be tailored to the appropriate age level, specifically Massachusetts Institute of Technology (MIT) Media Lab's *An Ethics of Artificial Intelligence Curriculum for Middle School Students* (Payne, 2019), Kids Code Jeunesse's (2022) *The Algorithm and Data Literacy Project*, and University of Helsinki's (2019) *Elements of AI*.

Speculative Futuring

Speculative futuring, or speculating about the harmful impacts of technology in the future, was brought up by several participants as a teaching strategy for imagining alternative algorithms that were less harmful. Speculative futuring often involved selecting popular fiction (often science fiction), a movie, or TV show with algorithms in it and critiquing the algorithmic systems that were in place. Kris gave the specific examples of the film *Minority Report* (Molen et al., 2002), Ted Chiang's (2019) science fiction short stories, Ruja Benjamin's (2016) *Designer and Discarded Genomes: An Experiment with Speculative Methods Over Time* book, and the TV show *Black Mirror* (Bathurst & Kinnear, 2018). Mason also discussed *Black Mirror* as an example:

Casey Fiesler at The University of Colorado Boulder, she has done some work putting together materials for this idea of a *Black Mirror* writer’s room where she has her students speculate about the harmful effects of technology. I think she’s focused a little bit more on tech ethics than algorithmic literacy specifically, but with a little bit of adaptations of those kind of materials—she has students take a current technology and extrapolate out as if you were writing an episode of the Netflix show *Black Mirror* and envision how could this go wrong. . . . I think that technique actually is a really cool way to ground the harms of technology but again, not explicitly about algorithms or algorithmic literacy so it may need a little bit of adaptation. (Interview)

Speculative futuring was brought up in the focus group. Norman referred to it as “design fictions,” while Madeline referred to it as “flash fiction,” defined as “having the students write a flash fiction piece about algorithms going terribly wrong or terribly right.” Madeline offered her own example of a video called *Epic 2015* (Sloan & Thompson, 2004) made back in 2004, that made some dystopian predictions about the future. She elaborated:

It was looking forward to the future, when things like Google and the yellow pages merged, and so I used it as a consciousness raising thing because that hadn’t happened yet. And maybe when this video came out, but I can see it being used as a how many of these things are now true in this future forecast video. (Focus Group)

Richard described some speculative fiction writing exercises assigned to his students and explained how it worked in a similar way to science fiction:

How science fiction works, it uses both estrangement and defamiliarization, so it takes something you’re familiar with and kind of alters it slightly, so you get a fresh view on it, which is often what science fiction does right. The science fiction kind of shows you an alternative future or as a way of actually reflecting on the present, and kind of what’s happening, so getting students to actually write those speculative pieces is a way of getting them to reflect on and think through what some of these systems might be doing. (Interview)

Kris brought up another example of speculative futuring, *The Oracle for Trans-feminist Technologies* (Coding Rights, 2021). It included a card exercise for collective brainstorming on alternative imaginaries surrounding technologies. The cards list objects, values (e.g., accountability), and situations (e.g., support the right to privacy and control over personal data)

and allowed for “imagining a world where efficiency is not the value that we are centering, where consuming is not the way that we’re thinking about things.” Kris elaborated on the purpose of speculative futuring:

Very hands-on speculative future is really, really important in questions where we’re trying to untangle or change the balance of power . . . these kinds of exercises allow us to think more broadly about what might be possible, both good and bad. . . . Thinking outside the box and being creative is also really important and being able to think of possible futures where things aren’t this way, but they are a different way. (Interview)

Nolen also expressed the importance of getting people to imagine alternative ways that algorithms could work in the future, and he called this ability a form of literacy:

Algorithms work this way now. Is it not because they are algorithms but just because they’ve been built in this way? And there could be algorithms that work in some other way, and given that we live in a world with algorithms, helping people get to a spot where they can imagine other ways for algorithms to work, I think, is the useful thing for us to do, right—that level of literacy, I think, would be great because if we only get to the level of algorithms have these consequences now, and therefore they always have these consequences, then we’re stuck. (Interview)

Contextualized Stories

Several participants discussed the importance of contextualizing stories of algorithmic bias to make it less abstract for students. Kenzie explained why this was a useful strategy:

I may never encounter, or the people that are important to me in my life may never encounter, a situation where they would have a judge that uses an algorithm to help them decide how to sentence them for an alleged crime. So what I know about those [algorithmic systems I use] is helpful in thinking about other systems, but not necessarily for my own life—it’s at a more abstract level I guess. (Interview)

Richard discussed how he used real stories to teach his students about algorithms:

I use a lot of real examples . . . I can tell stories of real projects where we’ve been handling the data and writing the algorithms, and we’ve been dealing with all the politics of working with stakeholders and institutions and what they’re trying to do . . . I try and do it through, I try and reveal all of those kinds of issues, and issues around you know, maybe biases or values or data quality or how the systems are then used to make decisions about you, or whatever I try to do it all with real examples. (Interview)

Kris also suggested using storytelling to illustrate the oppression that algorithmic systems can cause:

Storytelling, being able to come up with vignettes or cases where we can really clearly explain in a human way to people, so Virginia Eubanks (2018) has some really lovely vignettes in her book or Safiya Noble's (2018) *Algorithms of Oppression*, lovely vignettes. . . . One of the things we really need to do here is do more storytelling and do more vignettes [and] cases, these are human problems, and so we really need to contextualize them narratively. (Interview)

Mason suggested using Ruja Benjamin's *Race After Technology* (2019) because it had "a number of examples" to contextualize bias in a more human way. He also suggested discussing the following as case studies:

[Examples of] speech recognition systems not understanding speakers of African American vernacular English or other dialects generally, there's examples of some facial recognition, I think, is a classic one, there's that Netflix documentary [*The Social Dilemma*, Rhodes, 2020] and *Coded Bias* [Kantayya et al., 2020] from Joy Buolamwini and some others, Safia Noble's *Algorithms of Oppression* [Noble, 2018]. (Interview)

Monica pointed out that personal experience with oppression by algorithms would impact people's awareness, and that we should capitalize on this when we think about teaching others to be more aware of the ethical impacts:

Personal exposure and experience affects people's awareness. . . . If you are from a marginalized community, you will probably, the chance of you getting harmed by an algorithmic system is higher and therefore your awareness is higher . . . that's something we need to think about how we can utilize—maybe marginalized community members could be teachers of these algorithms or show the significance of the problem, if there is lack of awareness. (Interview)

Perspective Taking

Participants suggested perspective taking, or exposure to someone else's personalized feed or algorithmic outputs to aid in the realization that algorithms worked differently on different people. Evelyn explained the concept more:

I think that one of the really hard things for typical computer users who use their iPhones and don't really think much about how it works—perspective taking is always hard for people just as human beings and thinking about how somebody else, looking at the same web page, is not going to see the same thing as you. Having that common frame of reference, that's just gone right—that's not a thing that exists online anymore. But unless you can see something through someone else's eyes or see it from their perspective, you don't really get that sort of sense of how different people's online worlds are, and so it's been really hard for people to extrapolate from that—that there is some process that's operating on some kind of data that is making what I see different from what somebody else sees. And I don't, I mean I'm still struggling with how to get that across to people. (Interview)

Kenzie suggested having people compare their ad personalization profiles with each other:

The aspect of comparing experiences [I like], because that allows us to extend our own insights a bit. And I think it makes it more concrete . . . one of the useful things about experiential learning with algorithmic systems is that, I know a lot about my own—the data that I'm generating and the inputs that are going in, which is why, looking at those ad categories, it's pretty easy in many cases to guess why you're classified in a certain way, and then so if you can think about what you know about somebody else, you might be able to make more guesses or have more hypotheses, I guess, about how systems are functioning or how they might function or explanations for different outcomes and outputs that you see, so sharing experiences, I think, makes the information more concrete, and maybe stickier" (Interview)

Several participants suggested having students compare Google search results with each other. In the focus group, Becca recommended students search the same general topic and compare results in order to discuss “what shaped this, how much of this is about location, how much of this is actually just about the words you chose to search.” Kris suggested having students do different Google searches and “talking to each other about how, when I searched for X this was my auto complete, what was your auto complete.” Jess named the tool Search Atlas (Ochigame & Ye, 2021) to highlight how search results for the same query in Google differ across different countries; They noted this showed it was “heavily weighted by what they assume you care about based on where you are.” Jess also suggested the tool Split Screen (Keegan, 2021) to illustrate the filter bubble on Facebook based on things like political persuasion or

gender. Both Charlotte and Norman discussed Red Feed, Blue Feed (Keegan, 2016) in the focus group to show the political filter bubble. Norman elaborated:

Being able to look at someone else's feed like the Red Feed Blue Feed is a really, really good example of how, this thing works totally differently for someone else than it does for me, and I'm able to put sort of myself in the information world or informational shoes of someone else, is a really powerful moment to realize that we're all not living in sort of that same world. (Focus Group)

Avery suggested teaching students about "the social conventions that algorithmic outputs can create" (e.g., Amazon recommending items based on connecting your interests to those of past customers). Evelyn also discussed the importance of helping people understand the interdependence that has to happen in order for models to be personalized. She wanted people to speculate about "what information do they need about me, and what information do they need about everyone" in order to be able to tailor results individually because "in order to know stuff about me, the system needs to have information about everyone." Monica suggested having users log in on Yelp, then compare their list of recommended reviews with someone else's list to realize it was different. She explained how they can do this:

Log out, and log in as a guest or another user, and go to each site you wrote a review for, and check all the recommended reviews and check to see if yours are there. Because if you're logged in as yourself, Yelp will show all your reviews as recommended, even if they are filtered. (Interview)

Expectancy Violations

Expectancy violations referred to when a user was surprised by algorithmic output, and this often triggered insight into how they worked. Kenzie described it like this: "If something goes wrong, or if you have an expectation for the system and it is violated in some way, those moments are really generative of new insight." She added, "That kind of experiential learning is really, really important." Evelyn pointed out that the algorithmic process was "supposed to be

invisible, they're supposed to be a part of the infrastructure that's invisible, and so we only really notice them when they break down." Several participants suggested looking at your own ad profiles for Facebook and Google as a trigger for learning through expectancy violations. Evelyn described a study she did that looked at the inferences that Facebook and Google made about you:

People would look at their list, and it would say that they were into parenting, when they were like well I don't have any kids, why does it think I'm a parent, and so just what kinds of ways can we catch people's attention like that, and try to help them sort of think about why am I getting this. (Interview)

In the focus group, Norman commented "Students get really freaked out when it's super wrong because it's like wait, why does it think this, right. . . . I've had a lot of joy for those kinds of moments." Kenzie described an assignment her students did involving the ad categories that Facebook had sorted them into:

The exercise revolves around thinking about what are some of the labels that seem to ring true for you, which ones seem to accurately reflect who you think you are. And how did they know that about you, what are some possible ways that this could have been inferred. And then looking at the different ad categories that are wrong, that are maybe errors or there's some misjudgment that's happened, and again thinking about why that misjudgment, what might have happened, and because it's sort of in the context of individuals, the person's life, they're better able to judge . . . what behavior am I doing what data inputs am I producing that are giving insight into this kind of classification that's kind of building knowledge about me. (Interview)

Other examples described by participants that may solicit expectancy violations included giving a platform such as Yelp a new algorithm, reflecting on any different behavior by the platforms you used, and being suspicious that your Fitbit is cheating. Diane described a paper about Yelp users who were surprised to learn Yelp had introduced a new filtering algorithm:

I guess Yelp introduced an algorithm that sort of ranks reviews and filters out reviews and people didn't realize for a really long time that they had introduced this algorithm until their reviews sort of started disappearing and they didn't know why, and then they

started asking, a lot of people were really upset because they were like we didn't, we had no idea that the reviews were ranked using this algorithm now, it used to just be based on I don't know what was most recent or popular, but I thought that was an interesting thing, as well as that, sometimes people are just not aware that there is an algorithm. (Interview)

In the focus group, Becca recommended engaging students in discussions about whether they have observed any platforms they use regularly starting to behave differently-which could signal changes in their algorithms. She elaborated:

Maybe it would be interesting to talk to students about have you seen something you use start to behave differently, what's going on there, what do you think is going on . . . that might be a space where you could talk about oh, I'm observing something different going on here. That might spark some conversation maybe not knowledge but knowledge of what I don't know. (Focus Group)

In the last example, Evelyn described how a Fitbit user (in one of her studies) became suspicious it was cheating steps, and this led to more speculation about how it really worked:

The people who are really sort of interested in why the heck is this thing [Fitbit] working the way it does were the ones who felt like their steps were not being counted accurately, so there was an example of a woman who was wearing her wrist worn device, who went to a show and it counted steps when she was applauding for the performers of the show, and she's like wow this thing isn't doing what I thought it was doing. (Interview)

Conclusion

The primary research questions attempted to gauge what knowledge components of algorithmic literacy were specific to college students and what behavior would contribute to students' algorithmic literacy. Overall, the findings indicated that there were three subcategories of knowledge that would contribute to students' algorithmic literacy. The first subcategory consisted of general characteristics and distinguishing traits of algorithms that a learner would need to understand in order to become more algorithmically literate. The codes in this subcategory were "black box," "common misconceptions," "algorithm strengths," "machine learning versus rule based," "key algorithmic operations," "human-algorithm interplay,"

“algorithmic reasoning versus human reasoning,” and “infrastructure.” The second subcategory consisted of key domains in everyday life where algorithms were influential, with an emphasis on how they operated and the associated benefits and risks. The codes in this subcategory were “social services,” “social media” (along with child codes “filter bubble” and “disinformation/misinformation”), “search,” “health care,” “human resources,” “recreation and retail,” “finance,” and “education.” The third subcategory consisted of ethical considerations surrounding the application of algorithmic decision-making. The codes in this subcategory were “bias causes,” “privacy of personal data,” “accountability,” “oppression, privilege, and power,” and “commodification.”

The findings indicated there were five behaviors that students could use as coping strategies when learning about algorithms. The codes in this category were “gaming the system,” “broader advocacy,” “critical evaluation,” “skepticism anarchy,” and “transfer of knowledge.” The findings also indicated recommended pedagogy for the teaching of algorithms to college students. Participants had different opinions on how to define basic algorithmic awareness compared to a more advanced understanding, discussed in the codes “algorithmic awareness” and “algorithmic literacy.” Participants also discussed different ideas for where algorithmic literacy should be taught in the curriculum in the code “placement in the curriculum,” and they discussed alternative strategies for assessing algorithmic knowledge in the code “challenges of assessment.” Participants discussed nine teaching strategies designed to build students’ knowledge about algorithms in the codes “tactile (building),” “folk theories,” “everyday life,” “data transparency tools,” “mass media,” “speculative futuring,” “contextualized stories,” “perspective taking,” and “expectancy violations.” This study proposed to expand the

Framework for Information Literacy for Higher Education (Association of College & Research Libraries, 2016) to include algorithmic literacy. A plan for doing so will be discussed in chapter 5, along with other implications for practice.

CHAPTER 5

DISCUSSION

Findings

Study Background

This qualitative study, consisting of 11 semi-structured interviews and one focus group with scholars and teachers of critical algorithm studies, investigated the extent there was consensus among this group of experts on the content and behaviors that contribute to students' algorithmic literacy. Further, it explored instructional pedagogy surrounding the teaching of algorithmic literacy to college students. Previous research has demonstrated that college students are either lacking in algorithmic awareness or failing to engage critically with information they find online or through social media that has been heavily influenced by algorithmic filters (Barshaba et al., 2020; Bhatt & MacKenzie, 2019; Brodsky et al., 2020; Head, 2012, 2016; Head et al., 2018; Hinchliffe et al., 2018; Koenig, 2020; Powers, 2017; Schultheiß et al., 2018; Wineburg et al., 2016, 2020; Wineburg & McGrew, 2017). Algorithmic literacy is not currently being taught as part of the mainstream curriculum for college students (Head et al., 2020). There is a call in the literature, therefore, to improve students' algorithmic awareness and algorithmic literacy skills (Bakke, 2020; Bhatt & MacKenzie, 2019; Clark, 2018; Cohen, 2018; Head et al., 2020; Hobbs, 2020; Lloyd, 2019; Nichols & Stornaiuolo, 2019; Ridley & Pawlick-Potts, 2021; Valtonen et al., 2019). Critical algorithm studies are still a relatively new field, so there is not yet agreement on how to define algorithmic awareness or algorithmic literacy competencies, and more research needs to be done (Cotter & Reisdorf, 2020; Dogruel, 2021b; Ridley & Pawlick-Potts, 2021; Shin et al., 2022). This study sought to contribute to the literature on algorithmic

literacy for college students by exploring potential knowledge components, behavioral dimensions, and pedagogy.

Discussion of Findings

The findings suggested key knowledge components to include as part of the content curriculum for teaching algorithms to college students. These consisted of three sub-components: general characteristics and distinguishing traits of algorithms, key domains in everyday life using algorithms including how they work and the benefits and risks, and ethical considerations for the use of algorithms. The findings also suggested five behavioral dimensions associated with how students interact and cope with algorithms, along with suggested pedagogical practices for the classroom.

Knowledge Components

General Characteristics of Algorithms

Findings in this study indicated eight general traits or characteristics that were helpful for a learner to know in order to become more algorithmically literate. These attributes were as follows:

- black box or secretive nature of algorithms,
- common misconceptions often attributed to algorithms,
- algorithm strengths,
- difference between machine learning versus ruled based algorithms,
- key algorithmic operations such as filtering, classification, and ranking,
- human-algorithm interplay, defined as the role that humans play in programming algorithms,

- difference between algorithmic reasoning compared to human reasoning, and
- importance of the underlying infrastructure, or larger sociotechnical system, that algorithms are embedded in.

There was overlap between the general characteristics of algorithms brought up by participants in this study and key ideas in K-12 Artificial Intelligence (AI) literacy frameworks. Included in the big ideas published by the “AI for K12” working group were Perception (computers perceive the world using sensors), Representation & Reasoning (computers maintain representation of the world and use them for reasoning), and Learning (computers can learn from data); these overlapped with the concepts of algorithmic reasoning versus human reasoning and machine learning versus rule based from the current study (Touretzky et al., 2019). Similarly, Long and Magerko’s (2020) conceptual framework for “AI literacy” included competencies or design considerations for “understanding intelligence” (the difference between human and machine intelligence), “decision-making” (recognizing how computers reasoned and made decisions), “machine learning steps,” and “learning from data.” Also included were “human role in AI,” “AI’s strengths & weaknesses,” and “acknowledging preconceptions,” similarly articulated in this study as human-algorithm interplay, algorithm strengths, and common misconceptions (Long & Magerko, 2020).

The “black box,” hidden nature of algorithmic calculations, especially those used in machine learning, was documented in the literature along with concerns that these were considered trade secrets (Diakopoulos & Koliska, 2017; van Dijck, 2013; Head et al., 2020; Kitchin, 2017; Pasquale, 2015). This concept is not only fundamental for learning about algorithms, but it makes it challenging to teach about algorithms because the inner nature of how

they work is often obscure. Common misconceptions about algorithms corroborate work by Boyd and Crawford (2012) and Gillespie (2014), who suggested perceptions of algorithmic objectivity and accuracy were misleading, along with Broussard's (2019) research on technochauvinism, or the feeling that the technological solution is always better. Seaver (2017) called algorithms a modern myth "attributed with great significance and power" (p. 2). The key strengths of algorithms identified by participants in this study—speed and efficiency were acknowledged throughout the literature (for example, see Head et al., 2020; Kitchin, 2017), along with some similar examples of their potential benefits when used for traffic, climate change, and endangered wildlife (Hager et al., 2019; Rolnick et al., 2019). Also, user perceptions of the helpfulness of algorithms to prevent information overload and improve efficiency were documented by Dogruel et al. (2020), while Monzer et al. (2020) reported users felt algorithms saved them time and improved relevance, and Head et al. (2020) reported that college students felt algorithms made things more convenient.

Machine learning algorithms were identified by participants as an important concept, and likewise it was acknowledged in the literature as important because of the lack of human oversight and unstable, changing behavior based on the interpretation of new data (Brogan, 2016; Bucher, 2018; Head et al., 2020; MacCormick & Bishop, 2013). Participants in the study acknowledged a substantial overlap between machine learning and data literacy because of training data and a learner's need to recognize their personal data was used to train machine learning algorithms. Prado and Marzel (2013) argued data literacy, in turn, was an important component of information literacy because it entailed critically assessing data and its sources along with using data ethically. Fiesler et al. (2020) found through a syllabi analysis of 115

university technology ethics courses in the United States that 55 of the courses explored the topic of “AI & algorithms” and specifically covered machine learning along with black box algorithms. Lao (2020) proposed a new “Machine Learning Education Framework” for college students that covered how machine learning knowledge was created, how bias could occur in machine learning systems, and the resulting societal implications. Sulmont (2019) also proposed a taxonomy for machine learning goals for non-majors. Taken together, these two frameworks could be useful resources for teaching students about machine learning versus ruled-based algorithms. Other possible resources could include adapting material aimed at K-12 students from Marques et al. (2020), Rodríguez-García et al. (2020), and von Wangenheim et al. (2020).

The identification of key algorithmic operations as an important fundamental concept for building algorithmic literacy skills was echoed by Dogruel (2021b), who emphasized that lay users’ lack of understanding of algorithmic operations contributed to the opaqueness of algorithms. Diakopoulos (2014) discussed the same algorithmic operations as participants within the context of decision-making, Latzer and Festic (2019) offered an extended list of different algorithmic applications online in the context of nine different functions or algorithmic operations that included filtering, and Zarouali et al. (2021) included content filtering as a dimension of algorithmic awareness in the context of algorithmic awareness of media content. Zarouali et al. (2022) identified prioritization, association, classification, and filtering as the key algorithmic techniques most related to algorithmic persuasion online. Understanding human-algorithm interplay, or how algorithms were socially constructed and relational, was another important component for being more algorithmically literate, which coincided with Lloyd’s (2019) emphasis on the interdependency between algorithms and humans and algorithms’ unique

combination of technical and social. Bucher (2018) noted that when users offered up their data, algorithms rewrote themselves in an endless feedback loop, while Beer (2017) and Seaver (2019) discussed the instability and constant interplay between algorithms and humans and how users' understanding of the world and themselves was changed through algorithmic interaction. Additionally, "human-algorithm interplay" was one dimension to measure users' perceptions of their algorithmic awareness across the platforms of Facebook, YouTube, and Netflix (Zarouali et al., 2021), while Koenig (2019, 2020) categorized the interrelationship between people and algorithms as one dimension of college students' rhetorical awareness that could be measured through moves made to influence the algorithm.

Another knowledge component identified in this study, algorithmic reasoning versus human reasoning, corresponded with what Sulmont (2019) identified as human thinking versus computer processing, and she noted it was a barrier for students to understand. Sulmont (2019) recommended the tactic of simulating algorithms to help students understand this concept better. Diakopoulos and Kolisha (2017) described algorithmic reasoning as building a simplified or optimized reality of the world, while O'Neil (2016) warned that algorithmic reasoning looked at statistical correlations for past behaviors without looking at causality or meaning, and Broussard (2019) warned data alone could not make the same social and value judgements as humans. The last knowledge component in the category of general characteristics and traits of algorithms, infrastructure, was similarly described by Willson (2017) as "complex amalgams of political, technical, cultural and social interactions" (p. 141) as well as by Kitchin (2017) as "complex socio-technical assemblages" (p. 20). Gregory and Higgins (2013) used similar words to argue that critical information literacy needed to take into consideration the "social, political,

economic, and corporate systems that have power and influence over information production, dissemination, access, and consumption” (p. 4)—in other words, the entire infrastructure.

Key Domains in Everyday Life Using Algorithms

Findings in this study revealed eight key domains in everyday life using algorithms that were important for students to develop awareness of:

- social services,
- social media (along with the related concepts of filter bubbles and disinformation/misinformation),
- search,
- health care,
- human resources,
- recreation and retail,
- finance, and
- education.

There was support in the literature that this type of domain knowledge was considered part of algorithmic literacy; Dogruel et al. (2021) included questions about the domains of search, social media, banking, health wearables, personalized advertising, and GPS navigation as part of the “Algorithm Literacy Scale for Internet users.” Similarly, Latzer and Festic (2019) suggested algorithmic governance in everyday life could best be achieved by examining the areas of social and political orientation, recreation, commercial transactions, and socializing. Long and Magerko (2020) listed “interdisciplinarity,” or identifying a variety of technologies that use AI, as a competency for “AI literacy” in K-12 students. There was also evidence that some of these

domains were lacking in general algorithmic awareness more than others. Dogruel et al. (2020) found general users had greater awareness of algorithmic decision-making in the areas of advertising, online shopping, and media use but less awareness in the areas of job searching, online dating, news, and navigation systems; Zhang and Dafoe (2019) found Americans perceived AI used for surveillance and digital manipulation more important and impactful than bias found in hiring, disease diagnosis, or criminal justice bias; and Brodsky et al. (2020) found in college students greater awareness of online shopping algorithms than online search.

Looking at the domains more specifically, there was concern in the literature over surveillance algorithms used in social services such as facial recognition software that misidentified people with darker skin tones, disabled people, or people with a non-binary gender identity (Bennett & Keyes, 2019; Buolamwini & Gebru, 2018; Keyes, 2018; Keyes et al. 2021a, 2021b), along with recidivism algorithms used in criminal justice (Angwin, et al., 2016; O’Neil, 2016). Also, Virginia Eubanks (2018) argued that the digital tracking used on the poor resulted in a low rights environment with little political accountability or transparency, and her work was cited by multiple participants in the study as a good example of a case study for the welfare domain. Much was written in the literature about the technical aspects of social media algorithms, including what signals were most important (Bozdog, 2013; Bucher, 2018; Clark, 2018; Constine, 2018; Cooper, 2020; Covington et al., 2016; Davidson et al., 2010; Davies, 2019; Eslami et al., 2015; Instagram, 2021; Koul, 2019; Nemeth, 2020; Rieder et al., 2018). Participants’ concerns over social media algorithms that prioritized engagement over quality was also confirmed by recent research, with a special focus on the resulting filter bubble effects and the rapid spread of false information (Affelt, 2019; Bobkowski & Younger, 2018; Chin et al.,

2015; Head et al., 2020; Kavanagh & Rich, 2018; Newman et al., 2020; Pariser, 2011; Rosenberg et al., 2018; Wardle, 2018). The literature confirmed that both the general public (Eslami et al., 2015, 2016, 2019; Gran et al., 2021; Proferes, 2017; Warshaw et al., 2016) and college students specifically (Brodsky et al., 2020; Head et al., 2018; McGrew et al., 2018; Powers, 2017; Wineburg et al., 2016, 2020) need to improve their algorithmic awareness in the domain of social media; college students were especially vulnerable to misinformation or disinformation. Garrett et al. (2020) found evidence that both the social services and social media domains were being taught in AI ethics and technical classes in the United States. Items taught included the COMPAS recidivism algorithm, use of predictive analytics by child protection agencies and in criminal sentencing, the use of facial analysis to predict sexual orientation, filter bubbles, the Cambridge Analytica scandal, and confirmation bias (Garrett et al., 2020). Also, Gardner (2019) taught about algorithms in the social services domain as part of information literacy instruction.

The domain of search is possibly the most relevant for college students because it helps them meet their informational goals. The technical aspects of how search algorithms work was reflected in the literature review (California State University, Dominguez Hills, 2019; Graham, 2017; Granka, 2010; Hobbs, 2020; Lewandowski, 2017) along with how much power they exerted over anyone trying to access the web (Graham, 2017). There was evidence that both the general public (Bakke, 2020; Schmidt et al., 2019; Thomas et al., 2019; Zhang & Dafoe, 2019) and college students specifically (Head et al., 2020; Hinchliffe et al., 2018; McGrew et al., 2018; Powers, 2017; Schultheiß et al., 2018; Wineburg & McGrew, 2017) were somewhat lacking in awareness of the ways that algorithms influenced their search results. There was support in the

literature for information literacy pedagogy to be expanded to include the online search domain and the questioning of search results driven by algorithmic decision-making (Bakke, 2020; Clark, 2018; Head et al., 2020; Heffernan, 2020; Lloyd, 2019; Ridley & Pawlick-Potts, 2021). There was also evidence that including algorithms as part of the search domain was beginning to be included as part of the information literacy instruction (Clark, 2018; Fisher, 2017; Gardner, 2019; Mooney, 2019), often with the help of using Safiya Noble (2018) as a case study for how algorithms promoted bias against women and people of color. Some of the same concerns that participants had with algorithms used in the health care domain coincided with research by Obermeyer et al. (2019), suggesting a large amount of racial bias when comparing White and Black patients, and Bennett and Keyes (2019), suggesting bias around gender, race and ethnicity, class, people with disabilities, and geography in health diagnostic tools. Also, wearable health technology came up as a concern by participants, and this was explored by Costello and Floegel (2020) who found that users of mental health apps owned by technology companies did not fully trust the app with their health data or predictive advertising but felt resigned to the ubiquity of big technology companies.

Concerns participants had over algorithms used in the human resources domain were similar to those articulated by O'Neil (2016), who suggested names that sounded non-White were penalized by automated human resources screening, and Chen et al. (2018), who suggested rankings of candidates by resume search engines contained gender bias. Human resources analytics used in hiring was one of the content areas taught in standalone artificial intelligence ethics and technical classes at the university level (Garrett et al., 2020). The recreation and retail domain consisted mostly of comments about Netflix and Amazon algorithms, and the technical

aspects of how those algorithms worked were described by Clark (2018), Clark and Kaptanian (2018), Linden et al. (2003), and Smith and Linden (2017). Hobbs (2020) and Turow (2011) described the dynamic pricing practices used by online shopping sites like Amazon depending on digital profiling and factors like how affluent your neighborhood was and your perceived financial attractiveness, while Lloyd (2019) explained Amazon's impact on culture through its categorizations. Zarouali et al. (2021) asked about awareness of Netflix's content filtering in their algorithmic media awareness scale. Participants' comments on algorithms used in the finance domain coincided with research by O'Neil (2016) and Nopper (2019), who discussed how credit score algorithms used variables such as zip codes or Facebook friends' credit scores and how these proxies for digital character further perpetuated the wealth gap. Finally, concerns voiced by participants over algorithms used in the education domain coincided with research by O'Neil (2016), who described how schools would cheat to get to the top of university rankings once they figured out what proxies were used in the ranking algorithm as stand-ins for excellence. For example, some colleges made admitted students retake the SAT in order to boost their test scores, and other colleges lowered their academic standards to increase their graduation rate. Also, O'Neil (2016) described teacher assessment where algorithms used student test scores are a stand-in for teacher effectiveness, without considering any external factors.

Ethical Considerations for the Use of Algorithms

Findings in this study indicated five ethical considerations were important knowledge components to improve algorithmic literacy:

- bias causes, defined as the cause of bias in algorithmic outputs,
- privacy of personal data,

- accountability for algorithmic harm,
- oppression/privilege/power, or when a target group is abused through algorithms, and
- commodification, defined as the underlying financial incentives of companies using algorithms.

Findings corroborate recent research by Dogruel (2021b), who proposed that one of the four dimensions for algorithmic literacy for internet users was critical evaluation, consisting of opportunities and risks, individual and societal effects, and privacy implications. Similarly, Zarouali et al. (2021) included ethical considerations on a scale measuring algorithmic awareness of media content, with questions specifically about bias and privacy, and Dogruel et al. (2021) included questions about discrimination and bias in the “Algorithm Literacy Scale” for internet users. Long and Magerko (2020) listed “ethics” as a competency for “AI literacy” in K-12 students, specifically citing privacy, diversity, bias, transparency, and accountability. The “AI for K12” working group also listed “societal impact” as one of the five big ideas for AI literacy (Touretzky et al., 2019), while the International Society for Technology in Education (ISTE) included management of personal data and digital identity as competency standards for digital citizenship (International Society for Technology in Education [ISTE], 2016). Koenig (2019, 2020) defined higher levels of algorithmic awareness for college students as “critical awareness” of why algorithms function the way they do, including the dominant perspective that shapes algorithmic design cultures and understanding the institutional forces that shape algorithmic outputs. Fiesler et al. (2020) found evidence that university technology ethics courses in the United States explored algorithmic fairness, bias, profiling, and the need for transparency,

despite the challenge of auditing algorithms, while Garrett et al. (2020) found evidence of bias and privacy being taught.

The causes of bias articulated by participants corresponded with those articulated by Cathy O’Neil (2016) and Pessach and Shmueli (2020): unfair data, missing or incomplete data, algorithmic objectives that disadvantaged minority groups, and “proxy” attributes that substituted for sensitive attributes such as race or gender and exploited them on a large scale. The precariousness of keeping personal data private while using online platforms was also described by Head et al. (2020) and Ip (2018), along with big data practices that acquired, combined, and analyzed the data (Boyd & Crawford, 2012; Cheney-Lippold, 2011). Human and Cech (2021) found evidence of “dark patterns” or user interfaces that tricked users into doing things for cookie and privacy consent on Google, Amazon, Facebook, Apple, and Microsoft. Warshaw et al. (2016) found there was a substantial gap between what people believed companies were doing with their data and the current reality, while Rao et al. (2016) found mismatches between user expectations of privacy and services’ actual privacy practice. Head et al. (2020) found evidence that college students were somewhat aware of and concerned about automated decision-making reinforcing inequalities and causing privacy violations but felt the convenience was worth giving up their privacy for. Hartman-Caverly and Chisholm (2020) taught privacy literacy concepts in the context of information literacy instruction with learning outcomes for recognizing how personal data and metadata were collected, assessing how data was shared, making choices to safeguard privacy, and describing privacy as a human right. Trepte et al. (2015) defined privacy literacy in the Communication Science field along six dimensions: data collection practices, technical aspects, potential privacy threats and risks, laws

and legal aspects, strategies for privacy control, and strategies for dealing with threats. These resources would be helpful for teaching students about privacy in the context of algorithmic literacy.

Participants attributed part of the lack of algorithmic accountability by big technology companies over algorithmic harms to the lack of regulation in the United States, especially when compared to *The European Union General Data Protection Regulation of 2018* (European Union General Data Protection Regulation [EUGDPR], 2016) used in Europe. This coincided with research concluding the impact and effectiveness of algorithmic accountability policies were currently limited (Ada Lovelace Institute et al., 2021), and there was often a trade-off between accuracy and fairness (Pessach & Shmueli, 2020) and speed and fairness (Madaio et al., 2020). Participants' call for more emphasis on the impact to users was echoed by Cech (2021), Keyes et al. (2019), Latzer and Festic (2019), and Madaio et al. (2020), who called for more collaboration with all stakeholders during the design phase and more emphasis on the potential risks involved for the end user. Participants' concerns over the oppression/privilege/power that algorithms could exert over people who were already oppressed and had no recourse were corroborated by research from O'Neil (2016), who described many ways algorithms took advantage of vulnerable, marginalized populations who lacked the privilege and power to fight back as their weaknesses were exploited. Examples included predatory advertising of for-profit colleges, subprime lending, hiring decisions, credit scores, recidivism algorithms, and predictive policing. Bennett and Keyes (2019) offered the extreme example of computer vision facial recognition programmed to diagnose autism at the earliest stage possible, leading to autistic children being murdered as mercy killings. There was some evidence of a digital divide around the

sociodemographic factors of age, education, gender, race, income, or location correlating with algorithmic awareness and behavior. Cotter and Reisdorf (2020), Gran et al. (2021), Lu (2020), Min (2019), and Warshaw et al. (2016) all found evidence of a digital divide, but Dogruel et al. (2020) did not find evidence of a digital divide.

The last ethical consideration, the commodification and commercialization of personal information by big technology companies, was exemplified by descriptions of Google's core business model from Graham (2017) and Vaidhyathan (2011), while Head et al. (2020) described the attention economy designed to maximize attention and time spent online, and Zuboff (2019) described how behavioral data was fed into prediction products fueled by machine algorithms. Despite a general lack of trust in traditional authority figures, there was still trust placed in Google by college students (Head et al., 2020). Similarly, Dogruel (2019) found user distrust in advertising did not seem to spill over to influence their trust of the platform as the carrier of the message. DeVito (2021) found evidence that users' perceptions of what a platform was used for, its mission, and its values impacted their willingness to adapt to platform changes. Dogruel (2019) found that perceptions of transparency led to greater acceptance, while additional studies found perceptions of transparency led to trust (Monzer et al., 2020; Shin, 2020; Shin & Park, 2019). Meanwhile, Shin et al. (2021) found that perceived fairness, accountability, transparency, and explainability all led to an increase in trust. Trust, in turn, led to a willingness to share more data, which led to greater accuracy in algorithmic outputs, which led to greater satisfaction (Shin et al., 2021). Research confirmed that companies like Google try to portray themselves as a common good or a shared community (Lindh & Nolin, 2016). However, Woodruff et al. (2018) found increased awareness about algorithmic fairness led to a decrease in

trust, and Eslami et al. (2019) found transparency about how the Yelp algorithm really worked correlated with a change in attitude. This suggested greater awareness of a company's underlying business motives to manipulate users in ways that benefit their bottom line may lead to a decrease in trust and result in less opportunity for exploitation of a user's trust. Gallagher (2017) suggested having students research an algorithm such as PageRank using terms of service agreements, white papers, and corporate statements to identify their values and ideologies.

Student Coping Behaviors

Findings in this study revealed five coping behaviors students may deploy when learning about algorithms. The five behaviors were as follows:

- gaming the system, defined as the manipulation, experimentation, or hacking of algorithmic systems to gain better user control;
- broader advocacy, defined as applying principles of social justice to reduce the potential for algorithmic harm;
- critical evaluation, defined as actively questioning algorithmic design, decisions, and outputs and seeking alternatives or other perspectives if applicable;
- skepticism anarchy, defined as being overly cautious and suspicious of algorithms resulting in a feeling of determinism; and
- transfer of knowledge, defined as applying algorithmic knowledge about one domain to a new domain.

Behaviors are strongly related to algorithmic literacy because of the way that algorithms work in collaboration with humans to produce output (the human-algorithm interplay); people continuously shape and re-shape the algorithmic process through their own behaviors and

interactions. Dogruel (2021b) theorized that coping strategies, or abilities to use existing algorithms competently, were connected to users' knowledge and evaluation skills. In order to act competently, a user must first be aware of when there is an algorithm, then evaluate the effects of their possible actions on the algorithm, and finally act out the behavior (Dogruel, 2021b). Similarly, Latzer and Festic (2019) stated “without awareness, users cannot accurately assess potential benefits and risks” (p. 8). Latzer and Festic (2019) described user practices as “generally aimed at coping with risks that companies induce through their data collection and analysis strategies” (p. 10). However, there were conflicting findings over whether awareness and knowledge impacted behavior. Eslami et al. (2015) found awareness positively correlated with more active manipulation and exploration, DeVito (2021) found awareness positively correlated with more adaptability, Brodsky et al. (2020) found awareness positively correlated with higher algorithmic engagement, and Cotter and Reisdorf (2020) found knowledge positively correlated with frequency and breadth of use. On the other hand, Lu (2020), Fouquaert and Mechant (2021), Swart (2021), and Haider and Sundin (2020) found that awareness did not significantly change behavior. One possible explanation for this was the “privacy paradox,” described by Hartman-Caverly and Chisholm (2020) as a phenomenon where people's actual behaviors often contradicted their stated privacy values. Haider and Sundin (2020) suggested that although awareness was not always enough to change behavior, at least awareness played a role in the evaluation and resulted in more purposeful interaction with the information intermediary.

There was overlap between the “Family AI Literacy Framework” by Druga et al. (2021) and the idea of “gaming the system”—one of the four As was to adapt by modifying or customizing the application in order to trick it. The disposition of “productive persistence” that

Cunningham and Dunaway (2017) identified in the *ACRL Framework for Information Literacy for Higher Education* (2016) and described as employing alternative strategies and learning from mistakes was also similar to gaming the system. Evidence of users “gaming the system,” or manipulating, experimenting, and hacking algorithmic systems to gain more control, was prevalent in user studies on algorithms. Bucher (2017) reported Facebook users strategized to have posts gain greater popularity such as posting at a certain time or structuring the post in specific ways, and DeVito (2021) also described social media users who adjusted their self-presentation tactics based on awareness of the algorithmic process such as reposting. Haider and Sundin (2021) described users watching different genres of videos, liking different posts, and feeding it different data, while both Büchi et al. (2017) and Dogruel (2021b) reported people using fake information online as a coping strategy. Head et al. (2020) also found that college students used multiple user accounts for different interests to exercise digital code switching. The following variables were positively correlated with active engagement with algorithms: platform use (Haider & Sundin, 2021; Lu, 2020; Min, 2019), perceived user control (Lu, 2020), perceived knowledge (Lu, 2020), and adaptability (DeVito, 2021). Dogruel et al. (2020) found a small group of users who believed they needed to take an active part in influencing algorithms and that it was up to them to regain their autonomy, while Cotter and Reisdorf (2020) believed that direct experience of platform algorithms was the best way to gain insight. Gallagher (2017) suggested the activity of having students try to manipulate Facebook’s algorithmically driven timeline to enable their writing to be read more widely.

There was overlap between the behavior of broader advocacy and the ethical consideration of accountability (six co-occurrences where an excerpt was coded as both). There

was also overlap with the concept of “algorithmic justice,” defined as applying principles of social justice and applied ethics to the design, use, and regulation of algorithmic systems to reduce potential harm, inequality, and discrimination (Head et al., 2020). Broader advocacy also mirrored the critical information literacy movement because the focus was on encouraging students to advocate for social change by acting on oppressive power structures in information (Tewell, 2015). Students could advocate for better regulation like the European Union General Data Protection Regulation (EUGDPR, 2016) or put pressure on policymakers to support recent legislation (for example, the Algorithmic Accountability Act (2022). The Freedom of Information Act (1967) could be another potential avenue, although a loophole was that one of the protected categories is trade secrets (Diakopoulos, 2015). Finally, the disposition of “responsibility to community” that Cunningham and Dunaway (2017) identified in the *ACRL Framework for Information Literacy for Higher Education* (2016) replicated broader advocacy; it was described as demonstrating a sense of responsibility to the community when evaluating sources of information and being conscientious how they invoked authority to gain credibility.

Critical evaluation entailed actively questioning algorithmic design, decisions, and outputs and sometimes seeking alternatives. Information literacy has been linked to critical thinking previously (Albitz, 2007). Further, the dispositions of “mindful self-reflection,” defined as questioning assumptions about authority when evaluating sources of information, and “toleration of ambiguity,” defined as treating authority as subjective because it is based on the context of the information need, that were identified as present in the *ACRL Framework for Information Literacy for Higher Education* (2016) overlapped with the behavior of critical thinking (Cunningham & Dunaway, 2017). Long and Magerko (2020) listed “critical thinking”

as a design consideration in AI literacy for K-12 educators, specifically encouraging learners to be critical consumers of AI technology by questioning their intelligence and trustworthiness. Gran et al. (2021) found a correlation between higher levels of algorithmic awareness and “the critical,” while Haider and Sundin (2020) described one group of users as “the confident evaluators” who felt agency over algorithms and had the flexibility to adapt their behavior depending on the context and to evaluate information in a way that moved away from a binary “checklist” approach. Several researchers reported defensive practices to protect privacy that mirrored those given by participants in the study. Burrell et al. (2019) found Twitter users personalized platform defaults by configuring inclusion and exclusion signals like blocklists, Monzer et al. (2020) reported deleting cookies, Dogruel (2021b) reported modifying settings, Büchi et al. (2017) reported changing visibility settings and blocking, deleting, or deactivating cookies, and Head et al. (2020) reported that college students tried to protect privacy by using ad blockers, clearing cookies, using Firefox instead of Chrome, or using a VPN. Gardner (2019) reported having college students do search exercises comparing Google and DuckDuckGo as part of information literacy instruction.

Carmi and Yates (2020) argued that users needed to understand the various factors influencing who and what reached them in search results and social media feeds in order to make rational judgments about the information they encountered. Several tools for fact-checking and evaluation were identified in the literature, including ReviewMeta for uncovering biased or fake Amazon reviews (Clark, 2019), Health News Review for rating news stories in the medical field (<https://www.healthnewsreview.org/>), Snopes for checking if news was misinformation (Snopes Media Group Inc., 2022), and Botometer for social media bots (Yang et al., 2019). Register and

Ko (2020) measured users' ability to self-advocate against harmful machine learning models, presented in scenarios, by analyzing their critiques in hypothetical letters to the enforcer of the model. Analysis was done by looking for the presence or absence of coded criteria such as construct validity, outliers, causality, and consequences (Register and Ko, 2020).

The skepticism anarchy behavior consisted of being overly cautious and suspicious of algorithms to the point of feeling a sense of determinism. This coincided with research by Haider and Sundin (2020), who categorized a group of users as "the skeptical evaluators" because they were suspicious of everything, including the knowledge production system itself. Somewhat similarly, Monzer et al. (2020) found one user strategy for more agency over personalization news algorithms included avoidance, and Swart (2021) found users who were passive with engaging in news because they felt their own role in shaping algorithms was limited. Dogruel et al. (2020) and Dogruel (2021b) found users who felt personalization algorithms were restricting their autonomy and manipulative, and their coping strategy was to ignore them. The last behavior, transfer of knowledge, involved transferring algorithmic knowledge from one domain to another. This was replicated by Siles et al. (2020) who explored folk theories about Spotify and found that users linked their experience with other algorithms to their folk theories of how Spotify's algorithm worked. On the other hand, Brodsky et al. (2020) found many students did not transfer algorithmic awareness across new platforms not covered in a video about how online searches worked.

Pedagogy

Algorithmic Awareness versus Algorithmic Literacy

There was support from participants in this study for classifying a basic understanding of what an algorithm does (a process acting on your data), recognizing algorithms at work in everyday life across domains, and recognizing privacy risks associated with your personal data as lower levels of algorithmic awareness. Also, the coping behaviors of skepticism anarchy or avoidance, gaming the system for more empowerment, and transfer of knowledge across different domains were associated with lower levels of algorithmic awareness. This coincided with the definition from Dogruel et al. (2021) of algorithmic awareness as the degree to which users know what algorithms can be used for and in what online contexts algorithms are actually used. In terms of behavioral components, DeVito (2021) defined algorithmic literacy as “the capacity and opportunity to be aware of both the presence and impact of algorithmically-driven systems on self- or collaboratively-identified goals, and the capacity and opportunity to crystalize this understanding into a strategic use of these systems to accomplish said goals” (DeVito, 2021, p. 339:3). More specifically, DeVito (2021) developed a classification system with four levels to describe the extent that social media users manipulated the algorithm. The lower “functional” levels consisted of basic awareness (awareness of the presence of an algorithmic process) and causal powers (awareness that the algorithmic process had some specific causal effect), while the higher “structural” levels consisted of mechanistic fragments (awareness there were multiple criteria used during the algorithmic process to make decisions) and mechanistic ordering (awareness of not only multiple criteria used to make decisions but also causal ordering within this criteria).

There was less consensus from participants in the study over what constituted a higher level beyond basic algorithmic awareness to approach algorithmic literacy. To the extent there was any consensus, there was support for placing more sophisticated technical knowledge of how algorithms worked at a higher level. However, technical knowledge was not the only important component mentioned by participants. There was also some support for placing ethical considerations such as the consequences and impact of algorithmic bias and oppression on greater society at a higher level. The behavioral components of (a) actively avoiding taking algorithmic outputs at face value through critical evaluation and (b) broader advocacy, or taking action to reduce the harm of algorithmic systems at the community, political, or societal level, were associated with higher levels of algorithmic literacy beyond basic awareness. This somewhat dovetailed with the definition of algorithmic literacy for college students by Head et al. (2020) as a subset of information literacy that entailed a critical awareness of what algorithms are, how they interact with human behavioral data in information systems, and an understanding of the social and ethical issues related to their use (Head et al., 2020, p. 49). There was also overlap with Shin et al. (2021), who defined algorithmic literacy for platform users as having capabilities for algorithmic curation, control, and active practices to manage your environment along with an understanding of how algorithms reconstruct realities and are expressions of broader systems of power, while Cotter (2020) emphasized the importance of recognizing knowledge as situated and constructed in order to build a critical consciousness to respond to algorithms as expressions of broader systems of power. Dogruel et al. (2021) defined algorithmic literacy for internet users as awareness of the use of algorithms in online applications, platforms,

and services, along with knowing how algorithms work, being able to critically evaluate algorithmic decision-making, and having coping skills to influence algorithmic operations.

Challenges of Assessment

Participants identified the different components of algorithmic literacy (technical versus social) and the different domains where they resided as major challenges for assessment because they may require separate scales of assessment. This was reflected in the fact that at the time of this writing, there were only two validated scales in the literature measuring perceived algorithmic awareness (Zarouali et al., 2021) and algorithmic awareness and knowledge (Dogruel et al., 2021), and both were not comprehensive because they were limited to apply to only one category of algorithms. Participants also cautioned against a fixed-choice approach to assessment because of the difficulties in establishing a “right” answer due to the black box, secretive nature of algorithms and the fact that algorithms are unstable and always changing (resulting in quickly outdated information). Also, they warned that technical jargon for non-computer science majors could get in the way of conceptual knowledge, and they believed capturing critical thinking or the “why” of a mental model with fixed choice responses would be difficult, especially considering how contextual algorithmic literacy was in the sense that algorithms were always embedded in and dependent on a broader sociotechnical system. The same sentiments were expressed by Cotter (2020), who emphasized how strategic and nontransparent platforms were in sharing information about how their algorithms worked, and Bucher (2018), who emphasized that algorithms were always in the process of changing as users feed algorithms new data. Hargittai et al. (2020) and Swart (2021) emphasized the lack of ground truth for knowing how algorithms actually functioned and therefore what the correct responses

would be, while Swart (2021) also discussed how gaps in user vocabulary to describe their algorithmic encounters could be a barrier. Similarly, Hargittai et al. (2020) recommended discussing people's everyday experiences with systems they used often to help focus on relevant domains without being explicit about the emphasis on algorithms.

DeVito (2021), Koenig (2019, 2020), and Register and Ko (2020) offered more qualitative alternatives to assessing algorithmic literacy as recommended by the study participants. As discussed previously, there were conflicting findings in the literature over whether awareness and knowledge impacted behavior, especially when it came to privacy, but this study largely supported that the two were correlated. Further, some participants believed that self-efficacy, or feelings of empowerment for being able to manipulate the algorithms towards one's best interests, was correlated with algorithmic awareness and algorithmic literacy. Also, some participants believed that attitudes or emotions needed to come into play (strong feelings of I don't like this) in order to drive students to the behavior of broader advocacy. Hargittai et al. (2020) stated that attitudes could be signals of awareness and understanding, Araujo et al. (2020) found a positive correlation between self-efficacy and positive attitudes towards artificial intelligence, while Shin et al. (2022) found a positive correlation between perceptions of fairness, explainability, accountability, and transparency and efficacy. In terms of other attempts to measure attitudes, Al-Zahrani (2015) measured self-efficacy for digital citizenship, Kong et al. (2021) measured sense of empowerment to deal with AI, Doyle et al. (2019) measured self-efficacy for information literacy, Shin et al. (2020) measured algorithmic acceptance, Gran et al. (2021) measured attitudes towards algorithmic functions, Zhang and Dafoe (2019) measured perceptions of bias, and Latzer et al. (2020) measured attitudes about algorithms.

Placement in the Curriculum

Participants confirmed that algorithmic literacy was not taught at the college level in any standardized way to reach all students, and this finding coincides with research by Head et al. (2020), who reported algorithms barely, if ever, were taught in the classroom. Participants suggested either incorporating it into existing courses or building a new course that it could be part of. They stressed that most subject areas could find a connection since algorithmic literacy is so interdisciplinary, and even touching on it in just one unit or class period for that subject domain could be helpful. The subject areas suggested where it could be integrated beyond just one class period were science and technology (with a focus on the social aspects), infrastructural literacy, search skills, data literacy, digital literacy, privacy, or information literacy. The literature also suggested a mixture of existing and suggested subject areas where it could be deployed. Fiesler et al. (2020) found for university technology ethics courses in the United States, the majority of home departments were in computer science, followed by information science, philosophy, and communication. Bhatt and Makenzie (2019), Hobbs (2020), and Nichols and Stornaiuolo (2019) suggested embedding algorithmic literacy into digital literacy, Cohen (2018) and Valtonen et al. (2019) suggested putting it in with media literacy, Bakke (2020), Clark (2018), Head et al. (2020), and Ridley and Pawlick-Potts (2021) suggested integrating it into information literacy, and Prado and Marzal (2013) suggested information literacy and data literacy should be taught together. Clark (2018), Fisher (2017), Gardner (2019), and Mooney (2019) reported teaching algorithmic bias as part of information literacy instruction on Google searching, and Gardner (2019) also taught about bias in the areas of facial recognition

software and recidivism algorithms. Hartman-Caverly and Chisholm (2020) taught privacy literacy as an extension of information literacy.

Teaching Strategies

Findings in this study revealed nine teaching strategies to help students develop greater algorithmic literacy skills. The nine strategies were:

- tactile (building), defined as hands-on engagement in the building of algorithms;
- folk theories, defined as theories to explain algorithmic outputs;
- everyday life, defined as recognizing the algorithms users come into frequent contact with;
- data transparency tools, defined as the incorporation of information into user interfaces to visibly show how algorithms work;
- mass media, defined as sources of information about algorithms from the media;
- speculative futuring, defined as speculating about the harmful impacts of algorithms in the future;
- contextualized stories, defined as making instances of algorithmic bias more human;
- perspective taking, defined as exposure to someone else's algorithmic outputs to trigger the realization that algorithms work differently on different people; and
- expectancy violations, defined as situations when a user is surprised by algorithmic outputs often leading to better insight for how they work.

The first strategy, tactile (building), was described by participants in the form of applied learning activities to make algorithms less intimidating and help users realize their limitations and

imperfections. Examples given included building a computer vision system, a semantic network from wooden tiles, a peanut butter and jelly sandwich, a grilled cheese sandwich, taking apart a laptop, and using Google's *Quick, Draw!* (Jongejan, et al., n.d.) to train a machine learning drawing algorithm. The tactile (building) teaching strategy overlapped with Long and Magerko (2020), who listed "embodied interaction" and "opportunities to program" as design considerations for educators to create learner-centered artificial intelligence in K-12. Druga et al. (2021) included "author" (creating new AI applications) in the "Family AI Literacy Framework," and DeVito et al. (2017) characterized building algorithmic literacy as a matter of "learning by doing." Long et al. (2017) described *LuminAI*, an interactive installation at a local arts event where humans improvised dance movements with artificial intelligence agents and learned more about AI in the process. Lindner et al. (2019) described several unplugged activities about AI developed for K-12 that involved physical objects to foster kinaesthetic engagement. Two specific examples of the tactile (building) teaching strategy were used in information literacy instruction. Gardner (2019) reported having students in a credit-bearing information literacy course write algorithmic directions for how to make a peanut butter and jelly sandwich using "if, then" statements, while Clark (2018) reported an activity for librarians using pseudocode where they added values and made coding decisions for navigating the library.

The second teaching strategy, folk theories (perceptions about how algorithms worked), was discussed by participants as one way to measure algorithmic awareness. Despite the gap between belief and knowledge as well as limitations of this method caused by the black box nature of algorithms, participants suggested having students engage in reflective practices on how they use algorithmic systems, such as journaling. Folk theories surfaced in the literature

review in the context of user studies. DeVito et al. (2017) defined folk theories as intuitive, informal theories that individuals developed to explain how technological systems worked. These theories shaped how a user interacted with and received a system whether factually correct or not. Folk theories were also referred to as mental models, and they required qualitative methodologies to explain how people perceived cause and effect (Rader & Slaker, 2017). Researchers explored user folk theories for how social media platforms worked (Bucher, 2017; DeVito et al., 2017; DeVito, 2021; Eslami et al., 2016, 2019; Siles et al., 2020), general internet user folk theories (Dogruel, 2021a), folk theories for how companies collected and made inferences from user data (Warshaw et al., 2016), and folk theories regarding the data that activity trackers with sensors collected (Rader and Slaker, 2017). Examples of this applied to college students came from Bakke (2020), who suggested a “Search Reflection” activity for students using the think aloud method to help identify and reflect on their practices, and Koenig (2019, 2020), who found after students were asked to reflect through journal writing that the act of reflection led to more critical and rhetorical awareness of the algorithm’s influence. Given the importance of folk theories in shaping the user experience of algorithmic systems, it is important to help students examine these and correct inaccurate mental models.

The third teaching strategy was helping users recognize algorithms they used in everyday life and how they were affected by them. Participants emphasized this was a key component of algorithmic awareness and that it was important because many everyday algorithms were invisible to users. Also, it was a good teaching strategy because students would care more about something they used, and it would appear less abstract. Suggested activities for students employing this teaching strategy included a cognitive walkthrough of how a common everyday

algorithm such as Facebook worked, an audit of all systems they have come into contact with that use algorithms, or a class discussion or diary describing all their encounters with algorithms. The everyday life teaching strategy overlapped with Long and Magerko (2020), who listed “leverage learners’ interests” including everyday experiences as a design consideration for teaching artificial intelligence in K-12. Also, several researchers included users’ perceptions of the impact of different algorithms on their everyday life as a component of algorithmic awareness (Dogruel et al., 2020, 2021; Government of Canada, 2021; Latzer & Festic, 2019; Zhang & Dafoe, 2019). Hartman-Caverly and Chisholm (2020) mirrored the idea of students performing an algorithm audit with activities asking college students to describe all the places they had left data tracks today and to develop a personal data plan.

The fourth teaching strategy was data transparency tools (the incorporation of additional information into user interfaces to illustrate how algorithms worked under the hood). Ridley and Pawlick-Potts (2021) advocated for XAI, or explainable AI, to provide more transparency of how AI worked and an understanding of likely future behavior for the end user. Diakopoulos and Koliska (2017) called for more transparency in news algorithms by integrating algorithm information directly into the user interface through signals like frequently asked questions (FAQ), icons, on/off options, and tweakability of inputs and weights. Eslami et al. (2016) suggested including seams or visible hints disclosing certain aspects of automation operations, while Alvarado and Waern (2018) and Eslami et al. (2018) suggested presenting more transparency in social media ad profiling. Smith and Matteson (2018) noted that information literacy emerges during interaction with search systems, but modern system designs hide or render unworkable the contextual information needed for the judgment processes of information

literacy. Participants in this study advocated for using data transparency tools with students to help them better visualize how algorithmic systems worked, despite a tool's potentially short life span due to the dynamic nature of algorithms and the difficulties of tool builders to gain the access they need. Ochigame and Ye (2021) also noted that Google lacked an API for web search results and used tactics to block scraping, posing difficulties for tool builders.

Data transparency tools participants recommended to use with students included *Instawareness* for Instagram (Fouquaert & Mechant, 2021), *How Normal Am I?* for facial recognition (Schep, 2020), *Enablex* for emotion recognition (Enablex, n.d.), and *The Myth of the Impartial Machine* for predictive policing (Feng & Wu, 2019). Also, they suggested having students download personal data from something like Google Ad Settings or Facebook Ad Settings to see how they were classified. The data transparency tools teaching strategy overlapped with Long and Magerko's (2020) "explainability" and "promote transparency" design considerations for AI educators. Also, several data transparency tools were discussed in the literature as teaching tools (Clark, 2018; Eslami et al., 2015; Melsión et al., 2021; Wang & Fussell, 2020). Although Friggeri et al. (2014) and Lewandowsky et al. (2012) discussed the persistence of disinformation and misinformation on social media despite efforts at corrective transparency tools like relabeling items as "false," other researchers found positive benefits from using data transparency tools in other contexts such as increased algorithmic awareness (Fouquaert & Mechant, 2021), greater adaptability (DeVito, 2021), and greater acceptance of advertising (Dogruel, 2019).

The fifth teaching strategy was drawing on sources of information about algorithms from the mass media. Participants cited books by Ruja Benjamin (2019), Virginia Eubanks (2018),

Safiya Noble (2018), Cathy O’Neil (2016), and Frank Pasquale (2015), indicating these would likely be required reading for improving students’ algorithmic literacy. Also, the film *Coded Bias* (Kantayya et al., 2020) was cited along with the *Propublica* Machine Bias series (“Machine Bias: Investigating Algorithmic Injustice [2015 to 2022 series]”), and the additional strategies of using current events, podcasts, and adapting free self-guided tutorials from MIT Media Lab (Payne, 2019), Kids Code Jeunesse (2022), and University of Helsinki (2019). These findings coincide with Fisher (2017), who reported using a newspaper article to teach algorithm bias, as well as Garrett et al. (2020), who reported faculty using the current event of the Cambridge Analytica scandal in AI ethics classes, and Gardner (2019), who reported using Ted Talks, podcasts, and *Propublica* to teach about algorithm bias. Also, Ridley and Pawlick-Potts (2021) suggested libraries could use tutorials and resources from Kids Code Jeunesse (2022), University of Finland (2019), and A.I. for Anyone (n.d.) to teach algorithmic literacy concepts. Additional educational resources available for adaptation identified in the literature included those described by Kong et al. (2021), Lee et al. (2021), and Williams et al. (2021). Relying on the mass media to learn more about algorithms was a useful strategy not limited to college students; Dogruel et al. (2020) interviewed internet users and found users learned about algorithms through news media and fictional media.

The sixth teaching strategy was speculative futuring (speculating about the harmful impacts and ethical considerations of algorithms in the future). Participants emphasized the importance of this strategy to get students thinking creatively outside the box about alternative ways for algorithms to work, which could include writing speculative fiction or critiquing algorithms from popular science fiction. Specific examples good for speculative futuring

activities included the film *Minority Report* (Molen et al., 2002), Ted Chiang’s short stories (Chiang, 2019), Ruja Benjamin’s series of speculative field notes on genetic engineering (Benjamin, 2016), the TV show *Black Mirror* (Bathurst & Kinnear, 2018), where many episodes offered a speculative look at where technology was headed, the video *Epic 2015* (Sloan and Thompson, 2004), and the *Oracle for Trans-feminist Tech-nologies* card exercise (Coding Rights, 2021). The speculative futuring teaching strategy coincided with Long and Magerko (2020), who listed “imagine future AI” as a competency for AI literacy in K-12. Also, Fiesler (2018), described using *Black Mirror* to do speculative futuring with students taking an information ethics class, and Burton et al. (2018) discussed teaching computer ethics through science fiction.

The seventh teaching strategy was the use of contextualized stories of algorithmic bias in order to make bias less abstract and more human. Participants cited work by Benjamin (2019), Eubanks (2018), and Noble (2018); the documentaries *Coded Bias* (Kantayya et al., 2020) and *The Social Dilemma* (Rhodes, 2020); and examples of bias in facial recognition software and speech recognition software as good examples of contextualized stories to use with students. Also, one participant commented on the promise of having marginalized community members teach others about algorithmic bias in order to contextualize the harm they experienced. Using contextualized stories as a teaching strategy was supported by Hartman-Caverly and Chisholm (2020), who reported using case studies to teach privacy literacy. Also, Plane et al. (2017) found that when users were presented with discriminatory online advertising scenarios, they found it problematic when discrimination occurred as a result of explicit demographic targeting by race.

Similarly, Pierson (2018) found that college students' beliefs about algorithmic fairness could be changed through discussion.

The eighth teaching strategy was perspective taking, or exposure to someone else's personalized feed or algorithmic outputs to trigger the realization that algorithms worked differently on different people. Participants suggested activities to show students how different everyone's online experience was due to personalization such as comparing ad profiles, comparing Google search results for the same thing, comparing Google autocomplete for the same thing, and comparing Yelp featured reviews when logged in versus logged out. Also, the tools Search Atlas (Ochigame & Ye, 2021), Split Screen (Keegan, 2021), and Blue Feed, Red Feed (Keegan, 2016) were suggested to teach about the filter bubble effect. This teaching strategy coincided with Hobbs (2020), who suggested having students conduct searches on different countries on the same day using their Google accounts and posting screenshots to a digital bulletin. Research by Brodsky et al. (2020) and Powers (2017) found college students were largely unaware of personalization algorithms in Google and Facebook filtering their results, which suggests they would benefit from this teaching strategy.

The ninth teaching strategy was expectancy violations, or situations when a user was surprised by algorithmic output leading to better insight into how it worked. Participants stressed that when a user thought something went wrong or they felt misjudged by the algorithm, this was a rare opportunity for insight because usually the algorithmic process was invisible if it worked as planned. Example activities for students employing this teaching strategy were looking at their own Google or Facebook ad profile and ad categories, or asking students to look for any platform behavior changes. This teaching strategy surfaced in the research in the context of user

studies. Rader et al. (2020) studied user reactions to their personal ad settings and preferences on Facebook and Google and asked if anything was surprising, inaccurate, uncomfortable, missing, or confusing; they found expectancy violations caused users to develop sensemaking theories. Similarly, DeVito et al. (2017) explored folk theories of Twitter users in reaction to Twitter's switch to a timeline that was algorithmically curated and found resistance was a result of expectation violations, and that more negative reactions resulted in more detailed explanations or folk theories. In addition, Eslami et al. (2018) found that seeing incorrect inferences led to algorithmic disillusionment; Swart (2021) found users' sensemaking strategies of algorithms were context-specific and triggered by expectancy violations; Haider and Sundin (2021) found when content diverged from what was expected, this opened up new ways to engage with the infrastructure; and Dogruel et al. (2020) found users who encountered undesired content they did not actively select became more aware of algorithms.

Limitations

While the study's findings shed light on the knowledge components, student coping behaviors, and pedagogy surrounding algorithmic literacy, some limitations to the research design should be considered. Specifically, the research setting for the interviews and focus group was the contrived online setting of Zoom rather than a natural environment where nonverbal cues were harder to observe and turn taking for the focus group was more problematic. The presence of the researcher serving as focus group moderator and interviewer may have biased the responses, along with the fact that the researcher was inexperienced in this role. As with any focus group, participants may have felt pressure to conform with a majority opinion. Also, not all participants in the focus group or interview were equally articulate. Another limitation was the

fact that all participants except for one identified as White, therefore perspectives from marginalized participants were limited. There was sampling bias because the sample was purposive and not random. Also, the small sample size limited transferability of the results to apply to universities and colleges in all settings. Member checks were not done with participants to confirm the overall coding themes, which would have increased credibility. Also, the transcripts were only single coded by the researcher and not subjected to interrater reliability with a second researcher, leading to less objectivity and consistency in coding. Finally, due to the constantly changing nature of algorithms, any research on this topic is only a snapshot in time and instantly out of date.

Future Research

Considering the limited research on algorithmic literacy both generally and for college students specifically, there are many opportunities for future studies. Future studies should include representation from marginalized scholars and teachers in the area of critical algorithm studies to explore the impact of race on the topics covered in this research. Also, researchers could experiment with further validating and operationalizing the dimensions of algorithmic literacy pertaining to knowledge components, coping behaviors, and teaching strategies for college students suggested by this study. For example, an in-depth syllabi content analysis could be undertaken for university-level courses covering algorithms. More user studies specifically looking at college students' level of algorithmic awareness and algorithmic literacy are needed due to the current limited research and mixed results. The literature showed no consensus on how "algorithmic awareness" and "algorithmic literacy" were differentiated and defined, and this

study, while adding to the scholarly conversation on this topic, also found similar inconsistencies when applied to college students. As such, this is an area for future research.

Concrete behaviors and their correlation with algorithmic awareness is another important area for future study due to the strong human-algorithm interplay. Although a user must first be aware of when there is an algorithm in order to evaluate the effects of their possible actions on the algorithm, there were conflicting findings over whether awareness and knowledge are strong enough factors to impact behavior in a way that lets users interact with algorithms more competently. The extent to which user awareness is influenced by or influences both algorithm behavior and algorithm attitudes such as perceived autonomy, perceived trust, and perceived knowledge needs further research. Also, some evidence of a digital divide around the sociodemographic factors of age, education, gender, race, income, or location correlating with algorithmic awareness and behavior surfaced in the literature review of this study; this is one area for further research. Finally, the only published, validated measurements for algorithmic awareness were validated on general users outside of the United States and limited to only specific domains. Since nothing has been validated yet for college students or for general users in the United States, these are additional research opportunities.

Implications and Recommendations

In addition to future research, there are several theoretical, practical, and policy implications and recommendations from the current study.

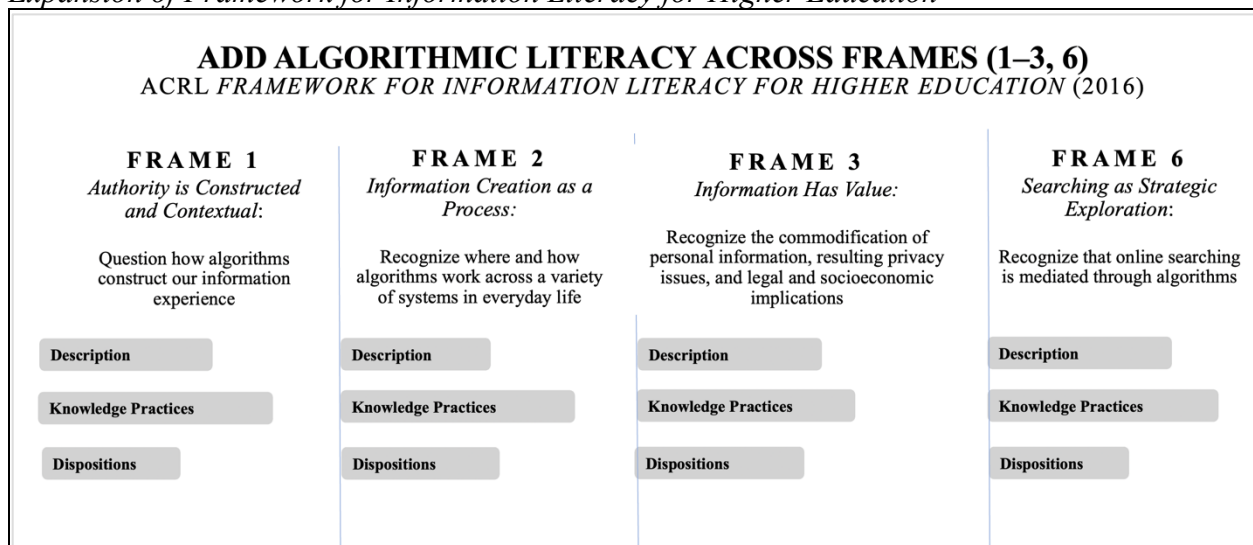
Theoretical Implications and Recommendations

This study offered a path to expand on the current *Framework for Information Literacy for Higher Education* (ACRL, 2016) to include algorithmic literacy (see Figure 3) more

explicitly. Because the everyday online information flow is now shaped by algorithms, new competencies for developing awareness of how algorithms influence what people see and the resulting individual and societal impacts are justified. Proposed expansions to the *Framework* are presented here for Frames 1-3 and Frame 6 that apply dimensions of algorithmic literacy related to knowledge components, coping behaviors, and pedagogy from the findings in this study. Each frame has a general description along with accompanying knowledge practices (proficiencies or abilities) and dispositions (tendency to act or think in a particular way). For each frame, the proposed extensions will be described in detail. Also, see Tables 10 through 13 for a list of existing parts of each frame verbatim that are relevant to algorithmic literacy along with the proposed new additions. The tables exclude existing components of each frame that are not related to algorithmic literacy; note that the frames are not presented in their entirety in this paper.

Figure 3

Expansion of Framework for Information Literacy for Higher Education



Note: Adapted from “Framework for Information Literacy for Higher Education,” by Association of College and Research Libraries, 2016 (<http://www.ala.org/acrl/standards/ilframework>), copyright 2016 by the Association of College and Research Libraries.

Frame 1, “Authority is Constructed and Contextual,” can be expanded to include the overall concept of questioning how algorithms construct our information experience (see Table 10). The description for this frame emphasized that users should remain skeptical of the systems that elevated the authority and the information created by it, and users should acknowledge biases that privilege some sources of authority over others. The following could be added to the description of Frame 1 to specifically include skepticism for the authority delegated to algorithms and an acknowledgement of their bias: “Experts understand the need to question the authority that is delegated to algorithms given their power to influence and manipulate and their propensity to produce biased results.” Three new knowledge practices could be added to Frame 1 as related proficiencies to help students reach the expert level: “understand that algorithms are informed by human decisions rather than being neutral,” “understand how algorithmic systems can produce bias from training data,” and “understand the ways that algorithms can artificially

amplify information to make it appear that certain views are widely shared or trustworthy.” And finally, the existing dispositions to maintain an open mind, assess content with a skeptical stance, be self-aware of your own biases and worldview, and exercise frequent self-evaluation could be enhanced by adding four additional dispositions to help students question algorithmic authority: “reflect on the ethical ramifications of algorithms that amplify existing structural patterns,” “develop awareness that algorithms work differently on different people,” “critically evaluate algorithmic design and decision making,” and “reflect on the role of information technology in a democratic society and the potential for political manipulation.”

Table 10*Expanding Frame 1: Authority is Constructed and Contextual: Question how Algorithms Construct our Information Experience*

Existing Descriptions, Knowledge Practices, and Dispositions that Could Apply to Algorithms	Suggested New Descriptions, Knowledge Practices, and Dispositions to Include Algorithms
<ul style="list-style-type: none"> • Experts view authority with an attitude of informed skepticism and an openness to new perspectives, additional voices, and changes in schools of thought (D) • Experts understand the need to determine the validity of the information created by different authorities and to acknowledge biases that privilege some sources of authority over others, especially in terms of others' worldviews, gender, sexual orientation, and cultural orientation (D) • Understand the increasingly social nature of the information ecosystem where authorities actively connect with one another and sources develop over time (KP) • Develop and maintain an open mind when encountering varied and sometimes conflicting perspectives (Disp.) • Develop awareness of the importance of assessing content with a skeptical stance and with a self-awareness of their own biases and worldview (Disp.) • Are conscious that maintaining these attitudes and actions requires frequent self-evaluation (Disp.) 	<ul style="list-style-type: none"> • Experts understand the need to question the authority that is delegated to algorithms given their power to influence and manipulate and their propensity to produce biased results (D) • Understand that algorithms are informed by human decisions rather than being neutral (KP) • Understand how algorithmic systems can produce bias from training data (KP) • Understand the ways that algorithms can artificially amplify information to make it appear that certain views are widely shared or trustworthy (KP) • Reflect on the ethical ramifications of algorithms that amplify existing structural patterns (Disp.) • Develop awareness that algorithms work differently on different people (Disp.) • Critically evaluate algorithmic design and decision making (Disp.) • Reflect on the role of information technology in a democratic society and the potential for political manipulation (Disp.)

Note: D = description, KP = knowledge practice, and Disp. = disposition.

Frame 2, “Information Creation as a Process,” can be expanded to include the overall concept of recognizing where and how algorithms work across a variety of systems in everyday life to output information (see Table 11). The description for this frame emphasized that since

the online environment obscures the format of information, users should look at the underlying process that went into creating the information. Furthermore, some information formats contain dynamic information. The following could be added to the description to specifically address the underlying creation process of algorithms and that they are dynamic: “Experts understand that algorithms never work in isolation, they are embedded in a broader sociotechnical system and dependent on the data they process.” The existing knowledge practices that could apply to recognizing where and how algorithms work in everyday life—articulating the processes of information creation and dissemination in a particular discipline, recognizing the implications of information formats that contain dynamic information, and transferring knowledge of capabilities and constraints to new types of information products—could be complemented by three new knowledge practices. The new knowledge practices are “recognize various applications of algorithmic decision making in everyday life,” “develop theories to explain algorithmic outputs in everyday life,” and “understand that machine learning algorithms are dynamic and unpredictable because they rewrite themselves as they work.” Finally, the existing disposition to seek out characteristics of information products that indicate the underlying creation process could be expanded to include these four new dispositions: “transfer knowledge of algorithmic decision making to new domains,” “reflect on the power of algorithms to mediate everyday life and shape the flow of information,” “develop awareness of the instability of algorithms as they engage in endless feedback loops and actively engage in the manipulation of algorithms to exert more user control,” and “speculate on the potential benefits versus future harmful impacts related to the use of algorithms in different domains.”

Table 11*Expanding Frame 2: Information Creation as a Process: Recognize Where and How Algorithms Work Across a Variety of Systems in Everyday Life to Output Information*

Existing Descriptions, Knowledge Practices, and Dispositions that Could Apply to Algorithms	Suggested New Descriptions, Knowledge Practices, and Dispositions to Include Algorithms
<ul style="list-style-type: none"> • The dynamic nature of information creation and dissemination requires ongoing attention to understand evolving creation processes (D) • Recognizing the nature of information creation, experts look to the underlying processes of creation as well as the final product to critically evaluate the usefulness of the information (D) • Articulate the traditional and emerging processes of information creation and dissemination in a particular discipline (KP) • Recognize the implications of information formats that contain static or dynamic information (KP) • Transfer knowledge of capabilities and constraints to new types of information products (KP) • Are inclined to seek out characteristics of information products that indicate the underlying creation process (Disp.) 	<ul style="list-style-type: none"> • Experts understand that algorithms never work in isolation, they are embedded in a broader sociotechnical system and dependent on the data they process (D) • Develop theories to explain algorithmic outputs in everyday life (KP) • Recognize various applications of algorithmic decision making in everyday life (KP) • Understand that machine learning algorithms are dynamic and unpredictable because they rewrite themselves as they work (KP) • Transfer knowledge of algorithmic decision making to new domains (Disp.) • Reflect on the power of algorithms to mediate everyday life and shape the flow of information (Disp.) • Develop awareness of the instability of algorithms as they engage in endless feedback loops, and actively engage in the manipulation of algorithms to exert more user control (Disp.) • Speculate on the potential benefits versus future harmful impacts related to the use of algorithms in different domains (Disp.)

Note: D = description, KP = knowledge practice, and Disp. = disposition.

Frame 3, “Information has Value,” can be expanded to enhance the overall concept of recognizing the commodification of personal information, resulting privacy issues, and legal and socioeconomic implications (see Table 12). The description for this frame emphasized the commodification of personal information and how the individual is responsible for making deliberate and informed choices about when to comply with and when to contest current legal and socioeconomic practices concerning the value of information. The following could be added

to the description to make it clearer how personal information is monetized online: “Experts understand that online platforms often use proprietary algorithms that allow third parties to access, aggregate, or sell personal data.” The existing knowledge practices of understanding how and why some individuals or groups of individuals may be underrepresented or systematically marginalized within the systems that produce and disseminate information and making informed choices regarding online actions in full awareness of issues related to privacy and the commodification of personal information could be complemented by three new knowledge practices that state the following: “understand that the revenue of free online platforms comes from advertising,” “identify general data collection and use practices of free online platforms,” and “understand the laws and legal aspects related to privacy protection in order to assign responsibility for harm caused by big data collection practices.” Finally, the existing disposition to see themselves as contributors to the information marketplace rather than only consumers of it can be complemented by these three new dispositions: “develop awareness of the ways individual digital profiles classify you and others based on presumed preferences or characteristics,” “recognize the value of critically evaluating the potential benefits versus the potential risks of free online platforms before using them,” and “recognize the value of privacy as a human right.”

Table 12*Expanding Frame 3: Information has Value: Recognize the Commodification of Personal Information, Resulting Privacy Issues, and Legal and Socioeconomic Implications*

Existing Descriptions, Knowledge Practices, and Dispositions that Could Apply to Algorithms	Suggested New Descriptions, Knowledge Practices, and Dispositions to Include Algorithms
<ul style="list-style-type: none"> • The value of information is manifested in various contexts, including . . . the commodification of personal information (D) • Experts also understand that the individual is responsible for making deliberate and informed choices about when to comply with and when to contest current legal and socioeconomic practices concerning the value of information (D) • Understand how and why some individuals or groups of individuals may be underrepresented or systematically marginalized within the systems that produce and disseminate information (KP) • Make informed choices regarding their online actions in full awareness of issues related to privacy and the commodification of personal information (KP) • See themselves as contributors to the information marketplace rather than only consumers of it (Disp.) 	<ul style="list-style-type: none"> • Experts understand that online platforms often use proprietary algorithms that allow third parties to access, aggregate, or sell personal data (D) • Understand that the revenue of free online platforms comes from advertising (KP) • Identify general data collection and use practices of free online platforms (KP) • Understand the laws and legal aspects related to privacy protection in order to assign responsibility for harm caused by big data collection practices (KP) • Develop awareness of the ways individual digital profiles classify you and others based on presumed preferences or characteristics (Disp.) • Recognize the value of critically evaluating the potential benefits versus the potential risks of free online platforms before using them (Disp.) • Recognize the value of privacy as a human right (Disp.)

Note: D = description, KP = knowledge practice, and Disp. = disposition.

Frame 6, “Searching as Strategic Exploration,” can be expanded to include the overall concept of recognizing that online search is mediated through algorithms (see Table 13). The description for this frame emphasized recognizing that information searching is a contextualized, complex experience that affects, and is affected by, the cognitive, affective, and social dimensions of the searcher. The following could be added to the description: “Experts recognize that online searching across free platforms is mediated through invisible algorithms, and that

there is the potential for algorithm bias.” The existing knowledge practices of designing and refining needs and search strategies as necessary based on search results and understanding how information systems are organized in order to access relevant information can be enhanced by two new knowledge practices: “understand key data points that inform ranking and filtering decisions in online search results” and “identify content when searching online that has been subject to algorithmic curation.” Finally, the existing dispositions of exhibiting mental flexibility and creativity and understanding that first attempts at searching do not always produce adequate results can be expanded to include two new dispositions: “develop awareness of the ways that the online searcher is mutually affected by and impacting their search results through usage behavior” and “develop awareness that search results are not objective and contain inherent biases.”

Table 13*Expanding Frame 6: Searching as Strategic Exploration: Recognize that Online Search is Mediated through Algorithms*

Existing Descriptions, Knowledge Practices, and Dispositions that Could Apply to Algorithms	Suggested New Descriptions, Knowledge Practices, and Dispositions to Include Algorithms
<ul style="list-style-type: none"> • Experts realize that information searching is a contextualized, complex experience that affects, and is affected by, the cognitive, affective, and social dimensions of the searcher (D) • Design and refine needs and search strategies as necessary, based on search results (KP) • Understand how information systems are organized in order to access relevant information (KP) • Exhibit mental flexibility and creativity (Disp.) • Understand that first attempts at searching do not always produce adequate results (Disp.) 	<ul style="list-style-type: none"> • Experts recognize that online searching across free platforms is mediated through invisible algorithms, and that there is the potential for algorithm bias (D) • Understand key data points that inform ranking and filtering decisions in online search results (KP) • Identify content when searching online that has been subject to algorithmic curation (KP) • Develop awareness of the ways that the online searcher is mutually affected by and impacting their search results through usage behavior (Disp.) • Develop awareness that search results are not objective and contain inherent biases (Disp.)

Note: D = description, KP = knowledge practice, and Disp. = disposition.

Practical Implications and Recommendations

The research findings have implications and recommendations for higher education teachers, libraries, college students, and K-12. These are discussed in the following subsections below.

Higher Education Teachers. The findings in this study yield recommendations to help faculty teach algorithmic literacy to college students (see Figure 4). The first dimension, knowledge components, was further divided into three subcategories. The first category of knowledge components consisted of general characteristics and distinguishing traits of algorithms, the second category consisted of key domains in everyday life using algorithms, and the third category consisted of general ethical considerations for the use and application of algorithms. There was overlap between the subcategories, and they could be taught concurrently

as part of algorithmic literacy instruction. There was not consensus from participants in the study over what constituted a higher level of learning beyond basic algorithmic awareness to qualify as algorithmic literacy. There was some support for categorizing either a sophisticated technical knowledge of how algorithms function as “literacy” or a sophisticated understanding of the social and ethical issues related to their use as “literacy.” Therefore, the knowledge components are only intended to be described as “algorithmic literacy” for non-computer science majors who have a less firm grasp of the technical aspects of algorithms.

Figure 4

Dimensions of Teaching Algorithmic Literacy to Non-Computer Science Majors

Knowledge Components	Student Behaviors and Coping Strategies	Pedagogy
<ul style="list-style-type: none"> • General characteristics and distinguishing traits of algorithms • Key domains in everyday life using algorithms • Ethical considerations for algorithms, including bias causes, oppression, privacy, commodification, and accountability 	<ul style="list-style-type: none"> • Broader advocacy • Critical evaluation • Gaming the system • Transfer of knowledge • Skepticism anarchy 	<ul style="list-style-type: none"> • Teaching strategies • Placement in the curriculum • Algorithmic awareness versus algorithmic literacy • Alternative assessment

The findings in this study also suggested five student coping behaviors that were correlated with learning more about algorithms. Given the strong human-algorithm interplay and the fact that algorithms are always actively making decisions for and with people, empowering students with behaviors and management strategies for how to impact algorithmic decision-making and protect their data is an important component of algorithmic literacy. It was suggested that faculty could help students move beyond the less helpful behavior of “skepticism anarchy,”

or avoidance, into the more productive behaviors of “gaming the system” to gain more agency over algorithmic decision-making and a “transfer of knowledge” across different systems that use algorithms in everyday life. It was also suggested that faculty could help students use critical evaluation rather than taking algorithmic outputs at face value. The last teaching strategy was for faculty to encourage students to practice broader advocacy by applying algorithmic knowledge to address some of the social justice problems caused by algorithms.

Additionally, the findings in this study yield recommendations for pedagogical practice among faculty. To help offset the abstract nature of algorithms, it is important to connect the elements of algorithmic literacy to students’ everyday life and have them engage in hands-on learning. It is also beneficial for students to reflect on their mental models for how platforms they engage with work and when their expectations of these platforms are violated. Additional teaching strategies faculty can employ to help strengthen students’ algorithmic literacy include using data transparency tools to shed light on how algorithms work under the hood, engaging students in perspective taking to realize algorithms work differently on different people, exposing students to case studies to contextualize algorithmic bias, and supplementing the curriculum with examples of algorithms in the mass media. Faculty may also engage students in conversations that speculate about the harmful impacts and ethical considerations of algorithms in the future.

Algorithmic literacy is not currently being taught as part of the mainstream curriculum for college students in a standardized way, and this study offered suggestions for where it may be placed in the curriculum. Due to its interdisciplinary nature, it overlaps with many other literacies in addition to information literacy, including data literacy, privacy literacy, media

literacy, computational thinking, platform literacy, digital literacy, infrastructural literacy, and artificial intelligence literacy. Algorithmic literacy could be included in courses where these other literacies are already being taught in the curriculum, and it could be intentionally added to any formal frameworks or standards that exist for these overlapping literacies. Participants in this study also suggested algorithmic literacy could be included in science and technology courses that focus on social or ethical aspects or in courses on search skills. The literature also revealed it was being taught in information science, philosophy, and communication courses (Fiesler et al., 2020). Algorithmic literacy opens up the opportunity for cross-collaboration between different academic disciplines and team-teaching. Since most subject areas can easily find a connection with algorithmic literacy due to its interdisciplinarity, faculty may find ways to touch on it during a single class period across all subject domains. Towards this end, another recommendation is for faculty to begin sharing practical examples with each other of how to apply algorithmic literacy into different disciplines through online communities of practice. This would be a time-saver and prevent faculty from having to reinvent the wheel. They could contribute to an existing open educational resource (OER) such as Community of Online Research Assignments (Loyola Marymount University, 2022), or they could create a new OER specifically for algorithmic literacy as a place to share teaching pedagogy on algorithmic literacy grounded in real-life examples. See Table 14 as a starting point with practical examples from this study. Universities can also invest in more training for faculty to develop their own algorithmic literacy skills in order to better teach students.

Table 14*Selected Activities for Students Related to Algorithmic Literacy*

Domain or Concept	Description	Citation
Algorithmic Reasoning versus Human Reasoning	<ul style="list-style-type: none"> • <i>Moral Machine</i> lets you judge moral dilemmas in a driverless car. 	Massachusetts Institute of Technology Media Lab (2020)
Human Resources, Machine Learning versus Rule Based	<ul style="list-style-type: none"> • <i>Survival of the Best Fit</i> lets you make hiring decisions that are then used to train an algorithm. 	Csapo et al. (2019)
Commodification, Speculative Futuring	<ul style="list-style-type: none"> • Explore what a city would look like if it were run on the business models of various companies. 	Graham et al. (2019)
Commodification	<ul style="list-style-type: none"> • Research an algorithm such as PageRank using terms of service, white papers, & corporate statements to identify its values. 	Gallagher (2017)
Classification, Data Transparency Tools, Expectancy Violations, Perspective Taking, Search	<ul style="list-style-type: none"> • Look at your Google ad settings to see how you were categorized and compare with others. 	Google (n.d.)
Filtering, Social Media	<ul style="list-style-type: none"> • Compare recommended reviews on Yelp when you are logged in versus not. 	Eslami et al. (2019)
Broader Advocacy	<ul style="list-style-type: none"> • Write hypothetical self-advocacy letters to the enforcer of a harmful machine learning model. 	Register & Ko (2020)
Disinformation/ Misinformation, Social Media	<ul style="list-style-type: none"> • Use the <i>Botometer</i> to identify social media bots. <ul style="list-style-type: none"> • Check for news misinformation with <i>Snopes</i>. • Verify images on social media with <i>TinEye</i>. 	Yang et al. (2019) Snopes Media Group Inc. (2022) TinEye (2022)

Table 14 (continued)*Selected Activities for Students Related to Algorithmic Literacy*

Domain or Concept	Description	Citation
Speculative Futuring	<ul style="list-style-type: none"> • <i>Black Mirror</i> writer's room. 	Fiesler (2018)
	<ul style="list-style-type: none"> • <i>Oracle for Trans-feminist Tech-nologies</i> card exercise for collective brainstorming on alternatives 	Coding Rights (2021)
	<ul style="list-style-type: none"> • <i>Epic 2015</i> flash forward video 	Sloan & Thompson (2004)
Data Transparency Tools, Social Media	<ul style="list-style-type: none"> • <i>Instawareness</i> tool to show invisible curation of Instagram 	Fouquaert & Mechant (2021)
Data Transparency Tools, Privacy of Personal Data	<ul style="list-style-type: none"> • <i>How Normal am I?</i> for facial recognition on your own face 	Schep (2020)
Data Transparency Tools	<ul style="list-style-type: none"> • <i>Enablex</i> for emotion recognition on your own face 	Enablex (n.d.)
Bias, Data Transparency Tools	<ul style="list-style-type: none"> • <i>Myth of the Impartial Machine</i> lets you change data inputs to see the impact on outputs 	Feng & Wu (2019)
Filter Bubble, Perspective Taking, Social Media	<ul style="list-style-type: none"> • <i>Split Screen</i> shows Facebook feeds from different perspectives 	Keegan (2021)
Classification, Commodification, Expectancy Violations, Folk Theories, Perspective Taking	<ul style="list-style-type: none"> • Journal prompts for Amazon, social media friends/followers, Google autocomplete, and Facebook ad profile. 	Koenig (2020)
Machine Learning versus Rule Based, Tactile (building)	<ul style="list-style-type: none"> • Use <i>Quick, Draw!</i> to have a neural network guess what you are drawing. 	Jongejan et al. (n.d.)

Table 14 (continued)*Selected Activities for Students Related to Algorithmic Literacy*

Domain or Concept	Description	Citation
Tactile (building)	<ul style="list-style-type: none"> Pseudo code exercise “Programming the Library” 	Clark (2018)

Note: The activities in this table are adapted from “Moral Machine,” by Massachusetts Institute of Technology Media Lab, 2020 (<https://www.media.mit.edu/projects/moral-machine/overview/>), copyright 2020 by Massachusetts Institute of Technology Media Lab; “Survival of the Best Fit,” by G. Csapo, J. Kim, M. Klasinc, & A. Elkattan, 2019 (<https://www.survivalofthebestfit.com>), copyright 2019 by Survival of the Best Fit; and How to Run a City Like Amazon, by M. Graham, R. Kitchin, S. Mattern, and J. Shaw, 2019, Meatspace Press, Creative Commons BY-NC-SA; “Writing for Algorithmic Audiences,” by J. R. Gallagher, 2017, *Computers and Composition*, 45, pp. 25–35 (<https://doi.org/10.1016/j.compcom.2017.06.002>), copyright 2019 by Elsevier Inc.; Google Ad Settings, by Google, n.d. (<https://adssettings.google.com/>), copyright n.d. by Google LLC; “User Attitudes Towards Algorithmic Opacity and Transparency in Online Reviewing Platforms,” by M. Eslami, K. Vaccaro, M. K. Lee, A. Elazari Bar On, E. Gilbert, and K. Karahalios, 2019, *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–14 (<https://doi.org/10.1145/3290605.3300724>), copyright 2019 by Association for Computing Machinery; “Learning Machine Learning with Personal Data Helps Stakeholders Ground Advocacy Arguments in Model Mechanics,” by Y. Register and A. J. Ko, 2020, *Proceedings of the 2020 ACM Conference on International Computing Education Research*, pp. 67–78, copyright 2020 by Association for Computing Machinery; “Arming the Public with Artificial Intelligence to Counter Social Bots,” by K. C. Yang, O. Varol, C. A. Davis, E. Ferrara, A. Flammini and F. Menczer, 2019, *Human Behavior and Emerging Technologies*, 1(1), pp. 48–61 (<https://doi.org/10.1002/hbe2.115>), copyright 2019 Wiley Periodicals, Inc.; Snopes, by Snopes Media Group, 2022 (<https://www.snopes.com/>), copyright 2022 by Snopes Media Group Inc.; TinEye, by TinEye, 2022 (<https://tineye.com/>), copyright 2022 by TinEye; Black Mirror, Light Mirror: Teaching Technology Ethics Through Speculation, by C. Fiesler, 2018 (<https://cfiesler.medium.com/the-black-mirror-writers-room-teaching-technology-ethics-through-speculation-fla9e2deccf4>), CC BY-NC-SA 4.0; The Oracle for Transfeminist Technologies, by Coding Rights, 2021 (<https://www.transfeministech.codingrights.org>), copyright 2021 by Coding Rights; Epic 2015, by R. Sloan and M. Thompson, 2004 (<https://www.albinoblacksheep.com/flash/epic2015>), copyright 2004 by Albino Blacksheep; “Making Curation Algorithms Apparent: A Case Study of ‘Instawareness’ as a Means to Heighten Awareness and Understanding of Instagram’s Algorithm,” by T. Fouquaert and P. Mechant, 2021, *Information, Communication & Society*, pp. 1–21, copyright 2021 by Taylor & Francis Group; How Normal Am I, by T. Schep, 2020 (<https://www.sherpapièces.eu/overview/how-normal-am-i>), copyright 2020 by SHERPA; Enablex Face AI Demo, by Enablex, 2022 (<https://www.enablex.io/cpaas/faceai/>), copyright 2022 by vCloudx; “The Myth of the Impartial Machine,” by A. Feng and S. Wu, 2019, *Parametric Press*, 1 (<https://parametric.press/issue-01/the-myth-of-the-impartial-machine/>), copyright 2019 by Parametric Press; Split Screen, by J. Keegan, 2021 (<https://themarkup.org/citizen-browser/2021/03/11/introducing-split-screen>), copyright 2021 by The Markup; “The Algorithms Know Me and I Know Them: Using Student Journals to Uncover Algorithmic Literacy Awareness,” by A. Koenig, 2020, *Computers and Composition*, 58, pp. 1–14 (<https://doi.org/10.1016/j.compcom.2020.102611>), copyright 2020 by Elsevier Inc.; and Quick, Draw!, by J. Jongejan, H. Rowley, T. Kawashima, J. Kim, and N. Fox-Gieg, n.d. (<https://quickdraw.withgoogle.com/>), copyright n.d. by Google; and Algorithmic Awareness, by J. A. Clark, 2018, ([jasonclark/algorithmic-awareness](https://jasonclark.com/algorithmic-awareness/)), CC by 4.0.

This study also suggested that faculty should experiment with nontraditional, qualitative forms of assessment when evaluating students’ algorithmic literacy skills. Due to many of the general characteristics and distinguishing traits of algorithms—including their black box and secretive nature, instability, difficulty verifying a “ground truth” for how they really work, abstractness, technical complexity, different ways of functioning across different domains, and contextual nature dependent on an underlying infrastructure—it is not easy to assess algorithms through traditional quantitative measures such as fixed choice tests. Instead, this study suggests faculty would be better served using a qualitative approach that allows for assessment of critical thinking, behaviors, and attitudes. Assessment prompts should not only be qualitative but also

context-heavy and provide concrete details of the exact domain the algorithm is in. A qualitative approach would avoid the pitfalls of punishing students for not understanding technical terminology or not being allowed to explain their mental models. As discussed previously, whether mental models are accurate or not, they are an important component in helping students become more conscious agents in their algorithmic encounters. Examples of alternative assessment given by participants included evaluating a self-advocacy letter to the company that had created an AI in a given scenario and evaluating students on how well they taught algorithmic awareness concepts to their elders.

This study suggested that higher education institutions may be relying on educational technology that is biased, and they need to think more critically about the ethical impacts of using this technology. For example, automated test proctoring systems that use facial recognition often unfairly flag students of color and disabled students for academic integrity violations that they did not commit. The recent COVID-19 pandemic increased usage of such technologies, giving students little choice but to use them. Another example is the college admissions process, which can be biased against international students if a program like Grammarly scores their essays.

Libraries. As discussed in the “Theoretical Implications and Recommendations” section, there is the opportunity for libraries to embed algorithmic literacy more explicitly into the existing *Framework for Information Literacy for Higher Education* (ACRL, 2016). There was support in the literature for integrating algorithmic literacy into information literacy (Bakke, 2020; Clark et al., 2018; Head et al., 2020; Ridley & Pawlick-Potts, 2021), and some librarians are already doing this (Clark, 2018; Fisher, 2017; Gardner, 2019; Hartman-Caverly & Chisholm,

2020; Mooney, 2019). Academic librarians can draw from the suggested knowledge components, coping behaviors, and pedagogical considerations for teaching algorithmic literacy presented in this study (see Figure 4) and share lesson plans with each other through open educational resources and communities of practice. Librarians who serve as liaisons to academic departments can curate collections from the mass media related to algorithms and target them to specific courses, thus helping to infuse algorithmic literacy more widely on campus (Ridley & Pawley-Potts, 2021). There is also the opportunity for public libraries to help the general public improve their algorithmic literacy through educational programming such as the “Data Days” for community members described by one of the participants in this study. Public libraries can draw on the teaching strategies and data transparency tools presented in this study to integrate active learning, and they can also draw on some of the free self-guided educational resources discussed in this study such as from University of Helsinki (2019) and A.I. for Anyone (n.d.). Ridley and Pawley-Potts (2021) suggested public libraries can also partner with K-12 schools, AI research and industry groups, and non-profit advocacy and training groups to develop algorithmic literacy training and advocate for making algorithms more understandable to citizens of all ages.

College Students. There is a call in the literature to improve students’ algorithmic awareness and algorithmic literacy skills (Bakke, 2020; Bhatt & MacKenzie, 2019; Clark, 2018; Cohen, 2018; Head et al., 2020; Hobbs, 2020; Lloyd, 2019; Nichols & Stornaiuolo, 2019; Ridley & Pawlick-Potts, 2021; Valtonen et al., 2019) because these skills need improvement (Barshaba et al., 2020; Bhatt & MacKenzie, 2019; Brodsky et al., 2020; Head, 2012, 2016; Head et al., 2018; Hinchliffe et al., 2018; Koenig, 2020; Powers, 2017; Schultheiß et al., 2018; Wineburg et al., 2016, 2020; Wineburg & McGrew, 2017). If higher education faculty and libraries carried

out some of the recommendations presented in this study to address the gap, college students would benefit because their algorithmic skills would improve. In particular, the recommendations could help students develop more awareness around the invisible structures influencing the information systems they use for both academic and personal purposes. They could interact with algorithmic systems more deliberately and reflectively and become more acutely aware of the dangers involved when we rely on data alone to make social and value judgments. They could question more how algorithms construct their everyday experiences, and they could be empowered to use tactics that impact algorithmic decision-making in ways that are beneficial to both them and the greater community.

K-12. Given the large amount of overlap between the findings in this study and K-12 Artificial Intelligence (AI) literacy frameworks (Druga et al., 2021; Long & Magerko, 2020; Touretzky et al., 2019) this suggests that K-12 teachers and students could benefit from the main findings in this study. K-12 teachers could use the knowledge components, behaviors, and pedagogy for teaching algorithmic literacy (see Figure 4) and adapt it for a K-12 audience. Somewhat outside the scope of this paper, K-12 students need privacy protection from both the invasion of school surveillance technologies and the big data collection practices of EdTech companies that mine personal student digital data for profit. Given the growth of educational technology since the COVID-19 pandemic and the implications for increased misuse of student data, more students were potentially exposed to privacy and equity risks. Better privacy protection needs to be in place along with more training for teachers in the area of privacy. Low-income students dependent on school-issued computers for homework are especially vulnerable.

Policy Implications and Recommendations

This study suggested certain educational policy changes that could help students become more algorithmically literate. Foundational skills such as privacy of personal data and the recognition of mis- and disinformation need to start earlier than college; they should become part of the regular K-12 curriculum given their increased importance in navigating everyday life. Policymakers need to encourage the development of these skills through bills that get successfully enacted and funded. A failed attempt at such a bill was the Digital Citizenship and Media Literacy Bill (2017), which would require the California superintendent of public instruction to develop best practice guidelines for teachers on digital citizenship, Internet safety, and media literacy through consultation with an advisory committee. Existing draft K-12 frameworks such as the “AI Literacy Framework” (Long & Magerko, 2020) or the framework drafted by the “AI for K12” working group (Touretzky et al., 2019) outlining similar skills need to be formally adapted and implemented by school systems. Also, although the Family Educational Rights and Privacy Act (1974) (FERPA) is the oldest and most important law protecting student privacy, it needs to be updated and strengthened to better protect student privacy in the digital age. For example, FERPA’s definition of “educational record” is too broad and makes it unclear whether students’ personal data via school surveillance technologies is part of that record (Rhoades, 2020). Also, a 2011 FERPA amendment allowed schools to disclose data to third parties as educational partners (Fedders, 2019).

All college students who are not majoring in computer science or mathematics should be required to take a course entitled “Information Literacies” built into their general education requirements. The course should embody an expanded definition of information literacy to also

include related literacies such as data literacy, privacy literacy, algorithmic literacy, and artificial intelligence literacy. The course should be worth one credit, and the class should meet approximately one hour per week as a large group. Two hours of additional guided online activities per week should be required incorporating teaching strategies and activities similar to those discussed in this study (see Table 14 and Figure 4). The course could follow a framework like the expanded *Framework for Information Literacy for Higher Education* (ACRL, 2016) proposed in this study (see Figure 3), and it should feature an interdisciplinary team of guest speakers to touch on key aspects of the course. For example, different faculty with expertise in privacy literacy and big data, search engines, statistics and data bias, and artificial intelligence ethics from various domains should all be enlisted to share their expertise as guest speakers. All college students who are computer science and mathematics majors should be required to take a different one credit course as part of their general education requirements entitled “Technology and Society” that touches on the social and ethical aspects of technology, with a special focus on the way invisible bias creeps into technology and information systems that use algorithmic decision-making.

The suggested educational policy changes for both K-12 and higher education will begin to improve students’ algorithmic literacy skills. However, there are not yet established national standards frameworks for K-12 and higher education that have proven to reliably capture algorithmic literacy competencies. Also, the correlation between algorithmic knowledge and algorithmic behavior is still unknown, and it remains unclear to what extent knowledge really changes behavior. Given that making these curriculum changes does not guarantee student behavior will change, it seems unwise to put all the responsibility on the individual student. The

entire sociotechnical infrastructure needs to change, because the current laws and regulations in the United States overseeing algorithms are inadequate and in need of updating. The country needs stronger legislation modeled after the European Union General Data Protection Regulation (EUGDPR, 2018), which requires companies to get consent to collect data and gives people the right to have their data erased, receive it in a commonly used format, have it transmitted to another data controller, and not be subject to an unfavorable decision based on application of artificial intelligence (Hautala, 2020).

The EUGDPR (2018) inspired companies even outside of Europe to start using pop-up cookie notifications asking for permission from users to track them in order to appear more transparent about privacy. Unfortunately, these are largely ineffective because users either don't read them before clicking on "accept," or users don't want to be penalized for noncompliance by not being granted access. Algorithms' complexity and obscurity have helped technology firms make the case that they are neutral platforms and avoid taking responsibility for algorithmic harms; their power has gone largely unchecked and unregulated, resulting in the proliferation of polarization and disinformation as well as the harms caused from algorithmic bias and the monetization of personal data. First Amendment protections largely succeed in shielding hate speech and discriminatory practices (Noble, 2018), while the Freedom of Information Act (FOIA; 1967) largely fails because trade secrets are a protected category, allowing third party companies to still obtain personal data (Diakopoulos, 2015). Section 230 of the Communications Decency Act (CDA) of 1996 protects online intermediaries such as social media companies from being legally responsible for what others post on their platform. As such, citizens need to put pressure on policymakers to better regulate algorithms.

Policymakers have started to introduce legislation that would better provide more uniform regulation of the use of personal data, embed better privacy protections, and offer more transparency and accountability. However, it is not yet enough to make a difference. In June of 2019, the U.S. Subcommittee on Communications, Technology, Innovation, and the Internet held a Senate hearing entitled *Optimizing for Engagement: Understanding the Use of Persuasive Technology on Internet Platforms* (*Optimizing for Engagement*, 2019). The hearing examined how algorithmic decision-making and machine learning on internet platforms might be influencing the public and whether algorithmic transparency or algorithmic explanation were policy options Congress should be considering. In April 2021, there was another hearing entitled *Algorithms and Amplification: How Social Media Platforms' Design Choices Shape our Discourse and Our Minds* (*Algorithms and Amplification*, 2021) where two proposed laws to limit some of Section 230's protections were discussed. So far, efforts to pass significant legislation have failed. For example, the Fundamentally Understanding the Usability and Realistic Evolution of Artificial Intelligence Act, or FUTURE of AI Act (2017) proposed to establish a federal advisory committee to report to the Secretary of Commerce on industry-disrupting AI issues, but it failed and then failed again after it was reintroduced as The FUTURE of AI Act in 2020.

Similarly, The Algorithmic Accountability Act (2019) failed initially, but was reintroduced in 2022—it proposed to require companies to assess the impacts of the automated systems they use and sell and create new transparency as well as report to the Federal Trade Commission (FTC) using structured guidelines. There are current (at the time of this writing) proposed bills to better monitor social media (Filter Bubble Transparency Act, 2021; Protecting

Americans from Dangerous Algorithms Act, 2021), offer better privacy protection (Mind Your Own Business Act, 2021; Consumer Data Privacy and Security Act, 2021; Deceptive Experiences to Online Users Reduction Act, 2021), and force online platforms to be more transparent and prohibit discrimination (Algorithmic Justice and Online Platform Transparency Act, 2021). The committees reviewing these bills need to stop tabling them and start acting on them. Many of the bills propose a more active regulatory role for the Federal Trade Commission (FTC), which would be a step in the right direction. The FTC could also potentially more actively enforce three laws to check the power of big technology firms. The first is Section five of the Federal Trade Commission (FTC) Act (1914), which could allow the prohibition of unfair or deceptive practices that could include racially biased algorithms (Jillson, 2021). The second is the Fair Credit Reporting Act (1971) in situations where an algorithm is used to deny people employment, housing, credit, insurance, or other benefits; and the third is the Equal Credit Opportunity Act (1974), making it illegal for a company to use a biased algorithm that results in credit discrimination on the basis of race, color, religion, national origin, sex, marital status, age, or because they receive public assistance (Jillson, 2021). The FTC needs to be better funded in order to carry out this work.

One challenge to regulation is that because algorithms are spread out across so many domains, the infrastructure is lacking to regulate all of them consistently. Other alternative approaches could include focusing on building more data transparency tools to make it clearer to users how algorithms work and putting more energy into “explainable AI” (XAI) to make algorithms more transparent, interpretable, and explainable. This would give back some agency to the user. Developers of algorithms should also collaborate more with stakeholders and act on

their input—and stakeholders could include not only end users but also domain experts, advocacy groups, and policy experts (Cech, 2021). Further, collaboration needs to happen across the entire lifecycle of a product (Madaio et al., 2020). All stakeholders, but especially the designers and deployers of algorithmic systems, need to engage in speculative futuring so they can get to a place where they can imagine other ways for algorithms to work that are less motivated by profit. Finally, advocacy organizations like A.I. For Anyone (n.d.), Electronic Frontier Foundation (n.d.), and Our Data Bodies (n.d.) can work to influence legislators, public agencies, platform service providers, and algorithm designers while simultaneously educating regular citizens and motivating them to join in the fight for algorithmic justice.

Conclusion

There is growing concern over how much of everyday life is increasingly mediated through information systems using algorithmic decision-making. Algorithms are now everywhere; they appear across many fields of study to digitally mediate everyday life and make decisions both for and with people, often without their awareness. The protection granted to algorithms as trade secrets perpetuates a “black box” society where the values embedded in the codes remain hidden. As a result, people do not realize the large amount of information being collected about them, how it is used, or what the consequences might be. Often, the public perception is that algorithmic systems can make better decisions than human beings, but many of the systems are skewed towards hegemonic (White, cis, male) perspectives. Civil rights and privacy violations occur when algorithms reflect the bias programmed into them and are used to make social and value judgements with no transparency, regulation, or accountability. One researcher made the analogy that algorithms are like racism because both make binary

predictions that all people with similar characteristics behave the same, and both are reinforced by institutional inequities and confirmation bias (O’Neil, 2016).

Research shows that college students are largely unaware of how algorithms work and the impact of algorithms on their everyday lives. Also, most university students are not being taught about algorithms as part of the regular curriculum. This study suggested knowledge components, coping behaviors, and pedagogical considerations to aid faculty in teaching algorithmic literacy to college students. Due to its interdisciplinarity, overlap with many other literacies, and ability to be easily incorporated into almost any subject area, there are many possible places to integrate algorithmic literacy more into the university curriculum. This study focused on how to embed algorithmic literacy into existing information literacy instruction by suggesting new competencies to add to the existing *ACRL Framework for Information Literacy for Higher Education* (2016). For information literacy skills to be a means of personal empowerment for lifelong learning, it is important that the curriculum be expanded to include algorithmic literacy. Doing so would empower students to not only develop greater algorithmic literacy skills but become advocates for a more democratic and open information society.

Finally, critical algorithm studies are a growing field with many opportunities for future research. The current exploratory study had several limitations, but it is the researcher’s hope that practitioners will still find its recommendations useful for teaching algorithmic literacy, while scholars will build on its findings to reach more rigorous reliability and credibility.

APPENDIX A

Email Solicitation

Dear X,

I am a doctoral student in LMU's Educational Leadership for Social Justice program, and my dissertation research is on algorithmic awareness in college students. I identified you as a key innovator and change agent in my area of study through my literature review, and therefore I would like to invite you to participate in my study.

I am conducting individual interviews and focus groups online (via Zoom) with scholars working in the area of algorithm studies. My aim is to explore the components of algorithmic awareness relevant to college students, along with the behavior and knowledge that contributes to this awareness. My study has been approved by LMU's Institutional Review Board, and I guarantee the data recordings will be stored on a secure server and permanently deleted after the project is complete. Also, your name will never be used in public dissemination of the data.

Are you available sometime within the next 2 months for a 1-hour interview or focus group? I would really appreciate and value your time and expertise, especially because there are still few scholars in this emerging field. If you are available, let me know what dates and times are convenient (I'm in PST, so let me know what time zone you are in).

Thank you very much, and I look forward to hearing from you,

Susan Archambault

APPENDIX B

Interview Protocol

Introduction: Thank you for agreeing to be part of this interview today. I really appreciate your time and expertise, and the information you provide as a content expert will be extremely valuable to my study.

As I mentioned in the initial email, I am studying algorithmic literacy in college students. I hope to solicit information from you about what the key topics are within the overall construct of algorithmic literacy and how this could best be measured in college students.

This interview will take no more than 1 hour. Your participation in this interview is completely voluntary, and you are allowed at any time to leave and terminate your involvement. Are you comfortable with video recording the session over Zoom?

Are there any questions before I send you the link for the informed consent form?

[Send link <https://bit.ly/algointerview> in the chat box for the informed consent forms]

As you read through this, please let me know if you have any questions.

[Collect informed consent form]

Are there any questions before we begin?

- What led you to become actively involved in the area of algorithm studies?
- How are algorithmic systems currently affecting people and impacting their lives?
- What are the strengths and weaknesses of algorithmic decision-making—what types of problems are a good match for algorithms versus what types are more challenging?
- What platforms or technologies using algorithmic decision-making have the greatest impact on society?
- If you were creating a list of the key elements that capture the idea of “algorithmic awareness” or “algorithmic literacy” in adult learners or college students, what content areas would you identify as most important?
 - What technical concepts related to algorithms are necessary for a basic understanding of how algorithms work?
- What teaching tools or teaching strategies can aid in learners’ understanding of algorithms?
 - Are there any examples of case studies that you use in your teaching that help students better understand algorithms?
 - Are there any additional examples of class activities that you use in your teaching that are helpful to students?
 - What can adult learners who are not in enrolled in college do to expand their understanding of how algorithms work?

- What behaviors reflect high levels of algorithmic awareness? In other words, what could someone do to demonstrate awareness of how algorithms work and their own role in the algorithmic process?
 - What dispositions do people typically have who possess high levels of algorithmic awareness?
- How can we measure or assess someone’s basic level of algorithmic awareness or algorithmic literacy?
 - How can a general assessment of algorithmic awareness be done without assuming too much technical knowledge?
 - How can a general assessment of algorithmic awareness be done without being platform-specific?
 - How might we create a rating scale to measure the idea of “algorithmic awareness”?
- How can we best embed algorithmic literacy into the mainstream curriculum for college students?
 - Where can we best communicate algorithmic concepts across the curriculum?
 - How might competencies for algorithmic literacy look different across different disciplines?

APPENDIX C

Focus Group Protocol

Introduction: Thank you for agreeing to be part of this focus group session today. I really appreciate your time and expertise, and the information you provide as a content expert will be extremely valuable to my study.

As I mentioned in the initial email, my name is Susan Archambault, and I am a doctoral student at Loyola Marymount University studying algorithmic literacy in college students. I identified all of you from my literature review as scholars and/or teachers working in the area of critical algorithm studies. I hope to solicit information from you about what the key competencies are for algorithmic literacy and algorithmic awareness and how these could best be measured in college students. At this point, I would like to invite all of you to briefly introduce yourselves by stating your name and where you work.

Great, thank you! This focus group will take no more than 1 hour. Your participation in this focus group is completely voluntary, and you are allowed at any time to leave and terminate your involvement. Also, even though the questions I will ask are not sensitive in nature, I expect this focus group to be kept confidential and not discussed outside of this meeting.

I will serve as your moderator and time keeper. I will post each question into the chat box as well as read it out loud, and then I ask that you raise your hand when you would like to respond and I will call on you. You can either raise your hand manually or use the Reactions panel in Zoom to raise your hand. I will give everyone a 2- minute warning in the chat box right before we need to move on to the next question, and at that point I will invite anyone who has not yet responded to comment if they would like to.

Is everyone comfortable with video recording the session over Zoom?

Are there any questions before I send you the link for the informed consent forms?

[Send link <https://bit.ly/algofocus> in the chat box for the informed consent forms]

As you read through this, please let me know if you have any questions.

[Collect all informed consent forms]

Are there any questions before we begin the focus group?

- What platforms or technologies using algorithmic decision-making have the greatest impact on society?
- If you were creating a list of the key elements that capture the idea of “algorithmic awareness” or “algorithmic literacy” in college students, what content areas would you identify as most important?
 - What technical concepts related to algorithms are necessary for a basic understanding of how algorithms work?

- What behaviors reflect high levels of algorithmic awareness? In other words, what could someone do to demonstrate awareness of how algorithms work and their own role in the algorithmic process?
- How can we measure or assess someone's basic level of algorithmic awareness? How can we measure or assess someone's level of algorithmic literacy?
 - How can a general assessment of algorithmic awareness be done without assuming too much technical knowledge?
 - How can a general assessment of algorithmic awareness be done without being platform-specific?
 - How might we create a rating scale to measure the idea of "algorithmic awareness"?

Closing: Thank you all so much for your time today! I really appreciate it. Your expertise has been so valuable for my study. A summary of the results of this research study will be sent to all participants as a courtesy when my dissertation is complete.

APPENDIX D

Sample Dedoose Chart: Code Co-Occurrences

Code Co-Occurrence		Include Overlapping Excerpts																																																	
Codes	Code	Accountability	Algorithm Strengths	Algorithmic awareness (lower levels)	Algorithmic literacy (higher levels)	Algorithmic reasoning vs. human	Bias causes	Black box	Broader advocacy	Challenges of assessment	Commodification	Common misconceptions	Contextualized stories	Critical Evaluation	Data transparency tools	Education	Everyday life	Expectancy violations	Finance	Folk theories	Gaming the system	Health care	Human resources	Human-algorithm interface	Infrastructure	Key algorithmic operations	Machine learning vs	Mass media	Oppression, privilege	Other personalization	Perspective taking	Placement in the curriculum	Information literacy	Privacy of personal information	Search	Steganogram anarchy	Social media	Distortion of information	Filter bubble	Social services	Speculative future										
	Accountability		1																																																
Algorithm Strengths			1																																																
Algorithmic awareness (lower levels)				1																																															
Algorithmic literacy (higher levels)					1																																														
Algorithmic reasoning vs. human						1																																													
Bias causes							1																																												
Black box								1																																											
Broader advocacy									1																																										
Challenges of assessment										1																																									
Commodification											1																																								
Common misconceptions												1																																							
Contextualized stories													1																																						
Critical Evaluation														1																																					
Data transparency tools															1																																				
Education																1																																			
Everyday life																	1																																		
Expectancy violations																		1																																	

REFERENCES

- American Civil Liberties Union of Washington (2020, March 24) *Algorithmic equity toolkit*.
<https://www.aclu-wa.org/AEKit>
- Ada Lovelace Institute. (n.d.). *About*. <https://www.adalovelaceinstitute.org/about/>
- Ada Lovelace Institute, AI Now Institute, and Open Government Partnership. (2021, August).
Algorithmic Accountability for the Public Sector. <https://www.opengovpartnership.org/documents/algorithmic-accountability-public-sector/>
- Affelt, A. (2019). *All that's not fit to print: Fake news and the call to action for librarians and information professionals*. Emerald Publishing Limited.
- Ahlin, E. M. (2019). Semi-structured interviews with expert practitioners: Their validity and significant contribution to translational research. In *SAGE Research Methods Cases Part 2*. <https://dx.doi.org/10.4135/9781526466037>
- A.I. For Anyone. (n.d.). *About us*. <https://www.aiforanyone.org/about>
- AI Now Institute. (n.d.). *About*. <https://ainowinstitute.org/about>
- Al-Zahrani, A. (2015). Toward digital citizenship: Examining factors affecting participation and involvement in the internet society among higher education students. *International Education Studies*, 8(12), 203–217. <http://dx.doi.org/10.5539/ies.v8n12p203>
- Albitz, R. S. (2007). The what and who of information literacy and critical thinking in higher education. *portal: Libraries and the Academy*, 7(1), 97–109.
<https://doi.org/10.1353/pla.2007.0000>
- Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. (2019).
<https://www.congress.gov/bill/116th-congress/house-bill/2231/text>
- Algorithmic Accountability Act of 2022, S.3572, 117th Cong. (2022).
<https://www.congress.gov/bill/117th-congress/senate-bill/3572/committees>
- Algorithmic Justice and Online Platform Transparency Act, H.R. 3611, 117th Cong. (2021).
<https://www.congress.gov/bill/117th-congress/house-bill/3611>
- Algorithms and amplification: How social media platforms' design choices shape our discourse and our minds: Hearing before the U.S. Senate Judiciary Subcommittee on Privacy, Technology and the Law*, 117th Cong. (2021).
<https://www.judiciary.senate.gov/meetings/algorithms-and-amplification-how-social-media-platforms-design-choices-shape-our-discourse-and-our-minds>

- Alvarado, O., & Waern, A. (2018). Towards algorithmic experience: Initial efforts for social media contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (pp. 1–12). Association for Computing Machinery. <https://doi.org/10.1145/3173574.3173860>
- American Civil Liberties Union. (2022). *FAQs*. <https://www.aclu.org/faqs>
- American Library Association. (1989). *Presidential committee on information literacy: Final report*. <http://www.ala.org/acrl/publications/whitepapers/presidential>
- American Library Association, Association of College and Research Libraries. (2000). *Information literacy competency standards for higher education*. <https://alair.ala.org/handle/11213/7668>
- Anderson, J., Rainie, L., & Luchsinger, A. (2018, December 10). *Artificial intelligence and the future of humans*. Pew Research Center. https://www.elon.edu/u/imagining/wp-content/uploads/sites/964/2019/07/AI_and_the_Future_of_Humans_12_10_18.pdf
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine bias. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing?token=aGn0cAzF1vcRrdeL0ElVR2hgvgtEPpJo>
- Annenberg Public Policy Center of the University of Pennsylvania. (2022). *Factcheck.org*. <https://www.factcheck.org/>
- Araujo, T., Helberger, N., Kruikemeier, S., & de Vreese, C. H. (2020). In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & Society*, 35(3), 611–623. <https://doi.org/10.1007/s00146-019-00931-w>
- Associated Press. (2006, 9 October). *Google buys YouTube for \$1.65 billion*. NBC News. <https://www.nbcnews.com/id/wbna15196982>
- Association of College & Research Libraries. (2016, January 11). *Framework for information literacy for higher education*. <http://www.ala.org/acrl/standards/ilframework>
- Association of College & Research Libraries Research Planning & Review Committee. (2020). Top trends in academic libraries: A review of the trends and issues affecting academic libraries in higher education. *College & Research Libraries News*, 81(6), 270–278. <https://doi.org/10.5860/crln.81.6.270>
- Association for Computing Machinery US Public Policy Council. (2017, January 2017). *Statement on algorithmic transparency and accountability*. https://www.acm.org/binaries/content/assets/publicpolicy/2017_usacm_statement_algorithmics.pdf

- Astrachan, O., & Briggs, A. (2012). The CS principles project. *ACM Inroads*, 3(2), 38–42. <https://doi.org/10.1145/2189835.2189849>
- Ayres, L. (2008). Semi-structured interview. In L. M. Given (Ed.), *The SAGE encyclopedia of qualitative research methods*. SAGE Publications, Inc. <https://dx.doi.org/10.4135/9781412963909>
- Baker, P., & Potts, A. (2013). “Why do white people have thin lips?” Google and the perpetuation of stereotypes via auto-complete search forms. *Critical Discourse Studies*, 10(2), 187–204. <https://doi.org/10.1080/17405904.2012.744320>
- Bakke, A. (2020). Everyday googling: Results of an observational study and applications for teaching algorithmic literacy. *Computers and Composition*, 57, 102577. <https://doi.org/10.1016/j.compcom.2020.102577>
- Barnhart, B. (2017, November 28). *How to set up a successful twitter advertising campaign*. Sprout Social. <https://sproutsocial.com/insights/twitter-advertising/>
- Barshaba, C., Lodhi, A., Brooks, P., & Brodsky, J. (2020). *Dissociations between college students’ media literacy knowledge and fact-checking skills*. Annual Meeting of the American Educational Research Association.
- Bathurst, O., & Kinnear, R. (2018). *Black mirror season 1–3*. Just Entertainment.
- Battista, A., Ellenwood, D., Gregory, L., Higgins, S., Lilburn, J., Harker, Y., & Sweet, C. (2015). Seeking social justice in the ACRL framework. *Communications in Information Literacy*, 9(2). <https://doi.org/10.15760/comminfolit.2015.9.2.188>
- Beer, D. (2017). The social power of algorithms. *Information, Communication & Society*, 20(1), 1–13. <https://doi.org/10.1080/1369118X.2016.1216147>
- Bellingcat Investigation Team. (2020, March 25). *The coronavirus disinformation system: How it works*. Bellingcat. <https://www.bellingcat.com/news/2020/03/25/the-coronavirus-disinformation-system-how-it-works/>
- Benjamin, R. (2016, October). *Designer and discarded genomes*. <https://www.e-flux.com/architecture/superhumanity/66875/designer-and-discarded-genomes/>
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity.
- Bennett, C. L., & Keyes, O. (2019, October). What is the point of fairness? Disability, AI and the complexity of justice. *ACM SIGACCESS Accessibility and Computing*, 125, <https://doi.org/10.1145/3386296.3386301>

- Bhatt, I., & MacKenzie, A. (2019). Just google it! Digital literacy and the epistemology of ignorance. *Teaching in Higher Education*, 24(3), 302–317. <https://doi.org/10.1080/13562517.2018.1547276>
- Big Brother Watch. (2021). *About—Big brother watch*. <https://bigbrotherwatch.org.uk/about/>
- Bobkowski, P., & Younger, K. (2018). *Be credible: information literacy for journalism, public relations, advertising and marketing students*. KU Libraries, University of Kansas. <http://hdl.handle.net/1808/27350>
- Bowles-Terry, M. (2012). Library instruction and academic success: A mixed-methods assessment of a library instruction program. *Evidence Based Library and Information Practice*, 7(1), 82–95. <https://doi.org/10.18438/B8PS4D>
- Boyd, D. & Crawford, K. (2012). Critical questions for big data. *Information, Communication & Society*, 15(5), 662–679. <https://doi.org/10.1080/1369118X.2012.678878>
- Bozdag, E. (2013). Bias in algorithmic filtering and personalization. *Ethics and Information Technology*, 15(3), 209–227. <https://doi.org/10.1007/s10676-013-9321-6>
- Brodsky, J. E., Zomberg, D., Powers, K. L., & Brooks, P. J. (2020). Assessing and fostering college students' algorithm awareness across online contexts. *Journal of Media Literacy Education*, 12(3), 43–57. <https://doi.org/10.23860/JMLE-2020-12-3-5>
- Brogan, J. (2016, February 2). Defining algorithms—A conversational explainer. *Slate Magazine*. <https://slate.com/technology/2016/02/whats-the-deal-with-algorithms.html>
- Broussard, M. (2019). *Artificial unintelligence: How computers misunderstand the world*. MIT Press.
- Brown, S., Davidovic, J., & Hasan, A. (2021). The algorithm audit: Scoring the algorithms that score us. *Big Data & Society*, 8(1), 1–8. <https://doi.org/10.1177/2053951720983865>
- Bucher, T. (2017). The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society*, 20(1), 30–44. <https://doi.org/10.1080/1369118X.2016.1154086e>
- Bucher, T. (2018). *If . . . then: Algorithmic power and politics*. Oxford University Press.
- Büchi, M., Just, N., & Latzer, M. (2017). Caring is not enough: The importance of Internet skills for online privacy protection. *Information, Communication & Society*, 20(8), 1261–1278. <https://doi.org/10.1080/1369118X.2016.1229001>

- Buolamwini, J. (2016, November). *How I'm fighting bias in algorithms*. [Video]. TED Conferences.
https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms
- Buolamwini, J., & Gebru, T. (2018). Gender shades: intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77–91. <http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>
- Burrell, J., Kahn, Z., Jonas, A., & Griffin, D. (2019). When users control the algorithms: Values expressed in practices on the twitter platform. In *Proceedings of the ACM on Human-Computer Interaction* (pp. 1–20). Association for Computing Machinery.
<https://doi.org/10.1145/3359240>
- Burton, E., Goldsmith, J., & Mattei, N. (2018). How to teach computer ethics through science fiction. *Communications of the ACM*, 61(8), 54–64. <https://doi.org/10.1145/3154485>
- California State University, Dominguez Hills. (2019). *Search engines libguide*.
<https://libguides.csudh.edu/searchengines>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
<https://doi.org/10.1126/science.aal4230>
- Carceral Tech Resistance Network. (2020, March 30). *Why*. <http://carceral.tech/why>
- Carmi, E., & Yates, S. J. (2020). What do digital inclusion and data literacy mean today? *Internet Policy Review*, 9(2). <https://doi.org/10.14763/2020.2.1474>
- Carrick Enterprises. (2018). The threshold achievement test for information literacy. [Measurement instrument]. <https://thresholdachievement.com/the-test>
- Cech, F. (2021). The agency of the forum: Mechanisms for algorithmic accountability through the lens of agency. *Journal of Responsible Technology*, 100015.
<https://doi.org/10.1016/j.jrt.2021.100015>
- Chen, L., Ma, R., Hannák, A., & Wilson, C. (2018, April). Investigating the impact of gender on rank in resume search engines. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI'18)*. (pp. 1–14). Association for Computing Machinery. <https://doi.org/10.1145/3173574.3174225>
- Cheney-Lippold, J. (2011). A new algorithmic identity: Soft biopolitics and the modulation of control. *Theory, Culture & Society*, 28(6), 164–181.
<https://doi.org/10.1177/0263276411424420>

- Chiang, T. (2019). *Exhalation*. Alfred A. Knopf.
- Chin, C. Y., Lu, H. P., & Wu, C. M. (2015). Facebook users' motivation for clicking the "like" button. *Social Behavior and Personality*, 43(4), 579–592.
<http://search.proquest.com/docview/1721707537/abstract/7276138F6AC34801PQ/1>
- Clark, H. (2019, August 25). Ways to spot fake Amazon reviews. *Atlanta Journal-Constitution*. S1.
- Clark, J. A. (2018). *jasonclark/algorithmic-awareness* [Rich Text Format]. GitHub.
<https://github.com/jasonclark/algorithmic-awareness>
- Clark, J. A., Hinchliffe, L. J., & Young, S. (2017). *Re:search—unpacking the algorithms that shape our UX*. (Research Grant Proposal, Institute of Museum and Library Services).
<https://www.imls.gov/grants/awarded/re-72-17-0103-17>
- Clarke, J., Kaptanian, J., & Bass, T. (2018, October 14). *Unpacking the algorithms that shape our UX: Algorithmic awareness as a form of information literacy* [Conference presentation]. Digital Library Federation, Las Vegas, NV, United States.
<https://dlfforum2018.sched.com/>
- Coding Rights. (2021). *The oracle for transfeminist technologies*.
<https://www.transfeministech.codingrights.org>
- Cohen, J. N. (2018). Exploring echo-systems: How algorithms shape immersive media environments. *Journal of Media Literacy Education*, 10(2), 139–151.
<https://doi.org/10.23860/JMLE-2018-10-2-8>
- Communications Decency Act of 2016, 47 U.S.C. § 230 (1996).
<https://www.law.cornell.edu/uscode/text/47/230>
- Constine, J. (2016, September 6). How Facebook news feed works. *TechCrunch*.
<https://social.techcrunch.com/2016/09/06/ultimate-guide-to-the-news-feed/>
- Constine, J. (2018, June 1). How Instagram's algorithm works. *TechCrunch*.
<https://techcrunch.com/2018/06/01/how-instagram-feed-works/>
- Consumer Data Privacy and Security Act, S. 1494, 117th Cong. (2021).
<https://www.congress.gov/bill/117th-congress/senate-bill/1494>
- Cooper, P. (2020, January 27). How the Facebook algorithm works in 2020 and how to work with it. *Social Media Marketing & Management Dashboard*. Retrieved December 20, 2021, from <https://blog.hootsuite.com/facebook-algorithm/>

- Cormen, T. H. (2013). *Algorithms unlocked*. MIT Press.
- Costello, K. L., & Floegel, D. (2020). “Predictive ads are not doctors”: Mental health tracking and technology companies. *Proceedings of the Association for Information Science and Technology*, 57(1), 1–11. <https://doi.org/10.1002/pra2.250>
- Cotter, K. (2019). Playing the visibility game: How digital influencers and algorithms negotiate influence on Instagram. *New Media & Society*, 21(4), 895–913. <https://doi.org/10.1177/1461444818815684>
- Cotter, K. M. (2020). *Critical algorithmic literacy: Power, epistemology, and platforms* (Publication No. 2435197746) [Doctoral dissertation, Michigan State University]. ProQuest Dissertations and Theses Global.
- Cotter, K. M., & Reisdorf, B. C. (2020). Algorithmic knowledge gaps: A new horizon of (digital) inequality. *International Journal of Communication*, 14(0), 21. <https://ijoc.org/index.php/ijoc/article/view/12450>
- Covington, P., Adams, J., & Sargin, E. (2016, September). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). Association for Computing Machinery. <https://doi.org/10.1145/2959100.2959190>
- Crano, W. D., Brewer, M. B., & Lac, A. (2014). *Principles and methods of social research*. Taylor & Francis Group.
- Creswell, J. W. (2009). *Research design: Qualitative, quantitative, and mixed methods approaches* (3rd ed.). Sage.
- Csapo, G., Kim, J., Klasinc, M., & Elkattan, A. (2019). *Survival of the best fit*. <https://www.survivalofthebestfit.com>
- Cunningham, A., & Dunaway, M. (2017). *Analyzing and assessing IL in the framework era: A rhetorical analysis followed by exploratory factor analysis*. [Paper presentation]. 9th Evidence Based Library and Information Practice Conference, Philadelphia, PA, United States.
- Cyr, J. (2019). *Focus groups for the social science researcher*. Cambridge University Press.
- Data & Society. (n.d.) *About*. <https://datasociety.net/about/>

- Davidson, J., Liebald, B., Liu, J., Palash, V. N., Taylor, V. V., Ullas, G., Sujoy, G., Yu, H., Lambert, M., Livingston, B., & Dasarathi, S. (2010). The YouTube video recommendation system. In *Proceedings of the Fourth ACM Conference on Recommender Systems* (pp. 293–296). Association for Computing Machinery. <https://doi.org/10.1145/1864708.1864770>
- Davies, S. (2019, February 19). *Decoding the social media algorithms in 2020: The ultimate guide*. <https://www.stedavies.com/social-media-algorithms-guide>
- Dean, B. (2021, October 10). Google’s 200 ranking factors: The complete list (2022). *Backlinko*. <https://backlinko.com/google-ranking-factors>
- Deceptive Experiences to Online Users Reduction Act, H.R. 6083, 117th Cong. (2021). <https://www.congress.gov/bill/117th-congress/house-bill/6083>
- Dedoose Version 9.0.46, web application for managing, analyzing, and presenting qualitative and mixed method research data (2021). Los Angeles, CA: SocioCultural Research Consultants, LLC. www.dedoose.com.
- DeVito, M. A. (2017). From editors to algorithms. *Digital Journalism*, 5(6), 753–773. <https://doi.org/10.1080/21670811.2016.1178592>
- DeVito, M. A. (2021). Adaptive folk theorization as a path to algorithmic literacy on changing platforms. In *Proceedings of the ACM on Human-Computer Interaction*, (pp. 1–38). Association for Computing Machinery. <https://doi.org/10.1145/3476080>
- DeVito, M. A., Gergle, D., & Birnholtz, J. (2017, May). “Algorithms ruin everything”: #RIPTwitter, folk theories, and resistance to algorithmic change in social media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 3163–3174). Association for Computing Machinery. <https://doi.org/10.1145/3025453.3025659>
- Dezuanni, M. (2021). Re-visiting the Australian media arts curriculum for digital media literacy Education. *Australian Educational Researcher*, 48(5), 873–887. <https://doi.org/10.1007/s13384-021-00472-6>
- Diakopoulos, N. (2014). *Algorithmic accountability reporting: On the investigation of black boxes*. Tow Center for Digital Journalism, Columbia University. <https://towcenter.columbia.edu/news/algorithmic-accountability-reporting-investigation-black-boxes>

- Diakopoulos, N. (2015). Accountability in algorithmic decision-making: A view from computational journalism. *ACM Queue*, *13*(9), <https://queue.acm.org/detail.cfm?id=2886105>
- Diakopoulos, N., & Koliska, M. (2017). Algorithmic transparency in the news media. *Digital Journalism*, *5*(7), 809–828. <https://doi.org/10.1080/21670811.2016.1208053>
- Digital Citizenship and Media Literacy Bill, S.B. 203, California Legislature. (2017-2018). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB203
- Digital redlining. (2020, December 23). In *Wikipedia*. https://en.wikipedia.org/wiki/Digital_redlining
- Dinger, E. (2020). Instagram, Inc. In S. Long, D. Jacques, & P. Kepos (Eds.), *International directory of company histories* (Vol. 215, pp. 152–155). St. James Press.
- Dogruel, L. (2019). Too much information!? Examining the impact of different levels of transparency on consumers’ evaluations of targeted advertising. *Communication Research Reports*, *36*(5), 383–392. <https://doi.org/10.1080/08824096.2019.1684253>
- Dogruel, L. (2021a). Folk theories of algorithmic operations during Internet use: A mixed methods study. *The Information Society*, *37*(5), 287–298. <https://doi.org/10.1080/01972243.2021.1949768>
- Dogruel, L. (2021b). What is algorithm literacy? A conceptualization and challenges regarding its empirical measurement. In M. Taddicken & C. Schumann (Eds.), *Algorithms and communication* (Digital Communication Research, 9, pp. 67–93). <https://doi.org/10.48541/dcr.v9.3>
- Dogruel, L., Facciorusso, D., & Stark, B. (2020). ‘I’m still the master of the machine.’ Internet users’ awareness of algorithmic decision-making and their perception of its effect on their autonomy. *Information, Communication & Society*. Advance online publication. <https://doi.org/10.1080/1369118X.2020.1863999>
- Dogruel, L., Masur, P., & Joeckel, S. (2021). Development and validation of an algorithm literacy scale for internet users. *Communication Methods and Measures*. Advance online publication. <https://doi.org/10.1080/19312458.2021.1968361>
- Donnellan, M. B., Oswald, F. L., Baird, B. M., & Lucas, R. E. (2006). The mini-IPIP scales: Tiny-yet-effective measures of the big five factors of personality. *Psychological Assessment*, *18*(2), 192–203. <https://doi.org/10.1037/1040-3590.18.2.192>

- Doyle, M., Foster, B., & Yukhymenko-Lescroart, M. (2019). Initial development of the Perception of information Literacy Scale (PILS). *Communications in Information Literacy*, 13(2). <https://doi.org/10.15760/comminfolit.2019.13.2.5>
- Druga, S., Yip, J., Preston, M., & Dillon, D. (2021). The 4As: Ask, adapt, author, analyze - AI literacy framework for families. In M. Ito, R. Cross, K. Dinakar, & C. Odgers (Eds.), *Algorithmic rights and protections for children: A new collection of interdisciplinary essays for open comment*. MIT Press Works in Progress. <https://doi.org/10.1162/ba67f642.646d0673>
- Edelman, B. G., & Luca, M. (2014, January 10)., Digital discrimination: The case of Airbnb.com. *Harvard Business School NOM Unit Working Paper No. 14-054*. <http://dx.doi.org/10.2139/ssrn.2377353>
- Electronic Frontier Foundation. (n.d.). *About EFF*. <https://www.eff.org/about>
- Elmborg, J. (2006). Critical information literacy: Implications for instructional practice. *The Journal of Academic Librarianship*, 32(2), 192–199. <https://doi.org/10.1016/j.acalib.2005.12.004>
- Enablex. (2022). *Enablex FaceAI demo*. <https://www.enablex.io/cpaas/faceai/>
- Equal Credit Opportunity Act, 15 U.S.C. § 1691 et seq. (1974). <https://www.ftc.gov/legal-library/browse/statutes/equal-credit-opportunity-act>
- Eslami, M., Karahalios, K., Sandvig, C., Vaccaro, K., Rickman, A., Hamilton, K., & Kirlik, A. (2016, May). First I like it, then I hide it: Folk theories of social feeds. In J. Kaye (Ed.), *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 2371–2382). Association for Computing Machinery. <https://doi.org/10.1145/2858036.2858494>
- Eslami, M., Krishna Kumaran, S. R., Sandvig, C., & Karahalios, K. (2018). Communicating algorithmic process in online behavioral advertising. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. <https://doi.org/10.1145/3173574.3174006>
- Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., & Sandvig, C. (2015, April). “I always assumed that I wasn’t really that close to [her]”: Reasoning about invisible algorithms in news feeds. In B. Begole (Ed.), *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 153–162). Association for Computing Machinery. <https://doi.org/10.1145/2702123.2702556>

- Eslami, M., Vaccaro, K., Lee, M. K., Elazari Bar On, A., Gilbert, E., & Karahalios, K. (2019). User attitudes towards algorithmic opacity and transparency in online reviewing platforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–14). Association for Computing Machinery. <https://doi.org/10.1145/3290605.3300724>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press
- European Union General Data Protection Regulation of 2018. (2018). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016. <https://gdpr-info.eu/>
- Fair Credit Reporting Act, 15 U.S.C. §§1681, et seq (1971). <https://www.ftc.gov/legal-library/browse/statutes/fair-credit-reporting-act>
- Family Educational Rights and Privacy Act of 1974, 20 U.S.C. § 1232g (1974). <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>
- Fedders, B. (2019). The constant and expanding classroom: Surveillance in K-12 public schools. *97 North Carolina Law Review* 1673. <https://scholarship.law.unc.edu/nclr/vol97/iss6/4>
- Federal Trade Commission Act, 15 U.S.C. §41 et seq. (1914). <http://www.ftc.gov/legal-library/browse/statutes/federal-trade-commission-act>
- Feng, A., & Wu, S. (2019, May 1). The myth of the impartial machine. *Parametric Press, 1*. <https://parametric.press/issue-01/the-myth-of-the-impartial-machine/>
- Fiesler, C. (2018, October 15). *Black mirror, light mirror: Teaching technology ethics through speculation*. <https://cfiesler.medium.com/the-black-mirror-writers-room-teaching-technology-ethics-through-speculation-fla9e2deccf4>
- Fiesler, C., Garrett, N., & Beard, N. (2020). What do we teach when we teach tech ethics?: A syllabi analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (pp. 289–295). Association for Computing Machinery. <https://doi.org/10.1145/3328778.3366825>
- Filter Bubble Transparency Act, S. 2024, 117th Cong. (2021). <https://www.congress.gov/bill/117th-congress/senate-bill/2024>
- Fisher, Z. (2017). Facing the frames: Using the Framework as a guide for a credit-bearing information literacy course. *College & Research Libraries News*, 78(7). <https://doi.org/10.5860/crln.78.7.354>

- Fletcher, R., & Nielsen, R. K. (2018). Generalised scepticism: How people navigate news on social media. *Information Communication and Society*, 22 (12), 1751–1769. <https://doi.org/10.1080/1369118x.2018.1450887>
- Freedom of Information Act, 5 U.S.C. § 552 (1967). <https://www.law.cornell.edu/uscode/text/5/552>
- Friggeri, A., Adamic, L. A., Eckles, D., & Cheng, J. (2014, May). Rumor cascades. In *Proceedings of the International Conference on Weblogs and Social Media* 8(1), 101–110. <https://ojs.aaai.org/index.php/ICWSM/article/view/14559>
- Fouquaert, T., & Mechant, P. (2021). Making curation algorithms apparent: A case study of “instawareness” as a means to heighten awareness and understanding of Instagram’s algorithm. *Information, Communication & Society*, 1–21. <https://doi.org/10.1080/1369118X.2021.1883707>
- Fulkerson, D., Ariew, S., & Jacobson, T. (2017). Revisiting metacognition and metaliteracy in the ACRL framework. *Communications in Information Literacy*, 11(1). <https://doi.org/10.15760/comminfolit.2017.11.1.45>
- FUTURE of Artificial Intelligence Act of 2017, H.R. 4625, 115th Cong. (2017). <https://www.billtrack50.com/BillDetail/900864>
- FUTURE of Artificial Intelligence Act of 2020, H.R. 7559, 116th Cong. (2020). <https://www.govinfo.gov/app/details/BILLS-116hr7559ih/related>
- Gallagher, J. R. (2017). Writing for algorithmic audiences. *Computers and Composition*, 45, 25–35. <https://doi.org/10.1016/j.compcom.2017.06.002>
- Gandolfi, E., & Ferdig, R. E. (2018). Scratching the coding surface: Tackling algorithms for inclusion and learning. *The International Journal of Information and Learning Technology*, 35(5), 368–378. <https://doi.org/10.1108/IJILT-06-2018-0071>
- Gardner, C. C. (2019). Teaching algorithmic bias in a credit-bearing course. *International Information & Library Review*, 51(4), 321–327. <https://doi.org/10.1080/10572317.2019.1669937>

- Garrett, N., Beard, N., & Fiesler, C. (2020). More than “if time allows”: The role of ethics in AI education. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 272–278). Association for Computing Machinery. <https://doi.org/10.1145/3375627.3375868>
- Gibson, C. & Jacobson, T. (2014, May). Informing and extending the draft ACRL information literacy framework for higher education: An overview and avenues for research. *College & Research Libraries*, 75(3), 250-254. <https://doi.org/10.5860/0750250>
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K.A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–93). MIT Press.
- Gillespie, T. (2016a). Algorithm. In B. Peters (Ed.), *Digital Keywords: A vocabulary of information* (pp. 18–30). Princeton University Press.
- Gillespie, T. (2016b). #Trendingtrending: When algorithms become culture. In R. Seyfert & J. Roberge (Eds.), *Algorithmic cultures: essays on meaning, performance and new technologies* (pp. 52–75). Taylor & Francis. <https://doi.org/10.4324/9781315658698>
- Google. (n.d.). *Google ad settings*. <https://adssettings.google.com/>
- Government of Canada. (2021) *Views of Canadians on artificial intelligence: Final report*. Innovation, Science and Economic Development Canada. <https://www.ic.gc.ca/eic/site/112.nsf/eng/07662.html>
- Graham, M., Kitchin, R., Mattern, S., & Shaw, J. (Eds.). (2019). *How to run a city like Amazon, and other fables*. Meatspace Press.
- Graham, R. (2017). Google and advertising: Digital capitalism in the context of post-Fordism, the reification of language, and the rise of fake news. *Palgrave Communications*, 3(1), 1–19. <https://doi.org/10.1057/s41599-017-0021-4>
- Gran, A. B., Booth, P., & Bucher, T. (2021). To be or not to be algorithm aware: A question of a new digital divide? *Information, Communication & Society*, 24(12), 1779–1796. <https://doi.org/10.1080/1369118X.2020.1736124>
- Granka, L. A. (2010). The politics of search: A decade retrospective. *The Information Society*, 26(5), 364–374. <https://doi.org/10.1080/01972243.2010.511560>
- Gregory, L., & Higgins, S. (Eds.). (2013). *Information literacy and social justice: Radical professional praxis*. Library Juice Press.

- Grgić-Hlača, N., Weller, A., & Redmiles, E. M. (2020). Dimensions of diversity in human perceptions of algorithmic fairness. *ArXiv:2005.00808* <http://arxiv.org/abs/2005.00808>
- Guarino, B. (2016, June 10). Google faulted for racial bias in image search results for black teenagers. *Washington Post*. <https://www.washingtonpost.com/news/morning-mix/wp/2016/06/10/google-faulted-for-racial-bias-in-image-search-results-for-black-teenagers/>
- Hager, G. D., Drobnis, A., Fang, F., Ghani, R., Greenwald, A., Lyons, T., Parkes, D. C., Schultz, J., Saria, S., Smith, S. F., & Tambe, M. (2019). Artificial intelligence for social good. *ArXiv:1901.05406*. <https://doi.org/10.48550/arXiv.1901.05406>
- Haider, J., & Sundin, O. (2020). Information literacy challenges in digital culture: Conflicting engagements of trust and doubt. *Information, Communication & Society*. Advance online publication. <https://doi.org/10.1080/1369118X.2020.1851389>
- Haider, J., & Sundin, O. (2021). Information literacy as a site for anticipation: Temporal tactics for infrastructural meaning-making and algo-rhythm awareness. *Journal of Documentation*. Advance online publication. <https://doi.org/10.1108/JD-11-2020-0204>
- Hamilton, K., Karahalios, K., Sandvig, C., & Eslami, M. (2014). A path to understanding the effects of algorithm awareness. In M. Jones (Ed.), *CHI'14 Extended Abstracts on Human Factors in Computing Systems* (pp. 631–642). Association for Computing Machinery. <https://doi.org/10.1145/2559206.2578883>
- Hargittai, E., Gruber, J., Djukaric, T., Fuchs, J., & Brombach, L. (2020). Black box measures? How to study people's algorithm skills. *Information, Communication & Society*, 23(5), 764–775. <https://doi.org/10.1080/1369118X.2020.1713846>
- Hargittai, E., & Micheli, M. (2019). Internet skills and why they matter. In M. Graham & W. H. Dutton (Eds.), *Society and the Internet: How networks of information and communication are changing our lives* (2nd ed., pp. 109–126). Oxford University Press.
- Harris, B. (2010). Encountering values: The place of critical consciousness in the competency standards. In M. T. Accardi, E. Drabinski, & A. Kumbier (Eds.), *Critical library instruction: Theories and methods* (pp. 279–291). Library Juice Press.
- Hartman-Caverly, S., & Chisholm, A. (2020). Privacy literacy instruction practices in academic libraries: Past, present, and possibilities: *IFLA Journal*, 46(4), 305-327. <https://doi.org/10.1177/0340035220956804>

- Hautala, L. (2020, January 3). CCPA is here: California's privacy law gives you new rights. *CNET*. <https://www.msn.com/en-us/news/technology/ccpa-is-here-californias-privacy-law-gives-you-new-rights/ar-BBYAH3z>
- Head, A. J. (2012). *Learning curve: How college students solve information problems once they join the workplace*. Project Information Literacy Research Institute. https://www.projectinfolit.org/uploads/2/7/5/4/27541717/pil_fall2012_workplacestudy_fullreport_revised.pdf
- Head, A. J. (2016). *Staying smart: How today's graduates continue to learn once they complete college*. Project Information Literacy Research Institute. https://www.projectinfolit.org/uploads/2/7/5/4/27541717/staying_smart_pil_1_5_2016b_fullreport.pdf
- Head, A. J., & Eisenberg, M. B. (2009). *Lessons learned: How college students seek information in the digital age*. Project Information Literacy. (ERIC Document Reproduction Service No. ED535167). <https://eric.ed.gov/?id=ED535167>
- Head, A. J., & Eisenberg, M. B. (2010). *Assigning inquiry: How handouts for research assignments guide today's college students*. Project Information Literacy Research Institute. https://www.projectinfolit.org/uploads/2/7/5/4/27541717/pil_handout_study_finalv_july_2010.pdf
- Head, A. J., DeFrain, E., Fister, B., & MacMillan, M. (2019). Across the great divide: How today's college students engage with news. *First Monday*, 24(8). <https://doi.org/10.5210/fm.v24i8.10166>
- Head, A. J., Fister, B., & MacMillan, M. (2020). *Information literacy in the age of algorithms: Student experiences with news and information, and the need for change*. Project Information Literacy Research Institute. <https://www.projectinfolit.org/uploads/2/7/5/4/27541717/algoreport.pdf>
- Head, A., Wihbey, J., Metaxas, P. T., MacMillan, M., & Cohen, D. (2018). *How students engage with news: Five takeaways for educators, journalists, and librarians*. Project Information Literacy Research Institute. <https://doi.org/10.2139/ssrn.3503183>
- Heffernan, K. (2020). Loaded questions: The framework for information literacy through a DEI lens. *College & Research Libraries News*, 81(8), 382-386. <https://doi.org/10.5860/crln.81.8.382>
- Herman, E. S. and Chomsky, N. (1988). *Manufacturing consent: The political economy of the mass media*. Pantheon.

- Hersher, R. (2017, January 10). What happened when Dylann Roof asked Google for information about race? *NPR*. <https://www.npr.org/sections/thetwo-way/2017/01/10/508363607/what-happened-when-dylann-roof-asked-google-for-information-about-race>
- Heuer, H., & Breiter, A. (2020). How fake news affect trust in the output of a machine learning system for news curation. In M. van Duijn, M. Preuss, V. Spaiser, F. Takes, & S. Verberne (Eds.), *Disinformation in open online media* (pp. 18–36). Springer International Publishing. https://doi.org/10.1007/978-3-030-61841-4_2
- Hinchliffe, L. J., Rand, A., & Collier, J. (2018). Predictable information literacy misconceptions of first-year college students. *Communications in Information Literacy*, 12(1), 4–18. <https://doi.org/10.15760/comminfolit.2018.12.1.2>
- Hobbs, R. (2020). Propaganda in an age of algorithmic personalization: Expanding literacy research and practice. *Reading Research Quarterly*, 55(3), 521–533. <https://doi.org/10.1002/rrq.301>
- Huddleston, T., Jr. (2017, November 7). Twitter expands to 280 character limits for most of the world. *Fortune*. <http://fortune.com/2017/11/07/twitter-280-characters-worldwide/>
- Human, S., & Cech, F. (2021). A human-centric perspective on digital consenting: The case of GAFAM. In A. Zimmermann, R. J. Howlett, & L. C. Jain (Eds.), *Human centred intelligent systems* (pp. 139–159). Springer. https://doi.org/10.1007/978-981-15-5784-2_12
- International Business Machines. (n.d.). *AI fairness 360*. <https://aif360.mybluemix.net/>
- Institute of Electrical and Electronics Engineers Global Initiative on Ethics of Autonomous and Intelligent Systems. (2019). *Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems*. Institute of Electrical and Electronics Engineers. <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/autonomous-systems.html>
- Independent High-Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI*. European Commission. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419
- Instagram. (2021). *Instagram help center*. <https://help.instagram.com/>
- International Federation of Library Associations and Institutions. (2005). *Beacons of the information society: The Alexandria proclamation on information literacy and lifelong learning*. <https://www.ifla.org/publications/beacons-of-the-information-society-the-alexandria-proclamation-on-information-literacy>

- International Society for Technology in Education. (2016, June). *ISTE Standards for Students*.
<https://www.iste.org/standards/iste-standards-for-students>
- Ip, C. (2018, September 4). Who controls your data? *Engadget*.
<https://www.engadget.com/2018-09-04-who-controls-your-data.html>
- Jillson, E. (2021, April 19). Aiming for truth, fairness, and equity in your company's use of AI. *Federal Trade Commission*. <http://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>
- Johnson, K. (2019, July 17). How AI companies can avoid ethics washing. *VentureBeat*.
<https://venturebeat.com/2019/07/17/how-ai-companies-can-avoid-ethics-washing/>
- Jongejan, J., Rowley, H., Kawashima, T., Kim, J., & Fox-Gieg, N. (n.d.). *Quick, draw!* Google Creative Lab. <https://quickdraw.withgoogle.com/>
- Kamberelis, G., & Dimitriadis, G. (2013). *Focus groups: From structured interviews to collective conversations*. Routledge.
- Kantayya, S., Hoffman, S., Acevedo, S., Gilwit, A., Ludescher, Z., & Mihailova, K. (2020). *Coded bias*. Women Make Movies.
- Kavanagh, J., & Rich, M. D. (2018). *Truth decay: An initial exploration of the diminishing role of facts and analysis in American public life*. Rand Corporation.
https://www.rand.org/pubs/research_reports/RR2314.html
- Keegan, J. (2016). *Blue feed, red feed*. Wall Street Journal. <https://graphics.wsj.com/blue-feed-red-feed/>
- Keegan, J. (2021, March 11). Introducing “split screen.” *The Markup*.
<https://themarkup.org/citizen-browser/2021/03/11/introducing-split-screen>
- Keyes, O. (2018). The misgendering machines: Trans/HCI implications of automatic gender recognition. In *Proceedings of the ACM on Human-Computer Interaction* (pp. 88:1–88:22). Association for Computing Machinery. <https://doi.org/10.1145/3274357>
- Keyes, O., Hitzig, Z., & Blell, M. (2021a). Truth from the machine: Artificial intelligence and the materialization of identity. *Interdisciplinary Science Reviews*, 46(1-2), 158–175.
<https://doi.org/10.1080/03080188.2020.1840224>

- Keyes, O., Hoy, J., & Drouhard, M. (2019). Human-computer insurrection: Notes on an anarchist HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. <https://doi.org/10.1145/3290605.3300569>
- Keyes, O., May, C., & Carrell, A. (2021b). You keep using that word: Ways of thinking about gender in computing research. In *Proceedings of the ACM on Human-Computer Interaction* (pp. 1-23). Association for Computing Machinery. <https://doi.org/10.1145/3449113>
- Kids Code Jeunesse. (2022). *The algorithm and data literacy project*. <https://algorithmliteracy.org/>
- Kim, S., Jang, Y., Kim, W., Choi, S., Jung, H., Kim, S., & Kim, H. (2021, May). Why and what to teach: AI curriculum for elementary school. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 15569-15576. <https://www.aaai.org/AAAI21Papers/EAAI-84.KimS.pdf>
- Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society*, 20(1), 14–29. <https://doi.org/10.1080/1369118X.2016.1154087>
- Kitchin, R., & Lauriault, T. P. (2015). Small data in the era of big data. *GeoJournal*, 80(4), 463–475. <https://doi.org/10.1007/s10708-014-9601-7>
- Knuth, D. E. (1998). *The art of computer programming: Sorting and searching* (Vol. 3). Pearson Education.
- Koenig, A. (2019). *Algorithmic literacy: Developing a heuristic for algorithmic literacy in the technical communication classroom and beyond* [Doctoral thesis, Texas Tech University]. Texas Tech University Electronic Theses and Dissertations (ETDs) Digital Archive. <https://hdl.handle.net/2346/85989>
- Koenig, A. (2020). The algorithms know me and I know them: Using student journals to uncover algorithmic literacy awareness. *Computers and Composition*, 58, 102611. <https://doi.org/10.1016/j.compcom.2020.102611>
- Kong, S. C., Man-Yin Cheung, W., & Zhang, G. (2021). Evaluation of an artificial intelligence literacy course for university students with diverse study backgrounds. *Computers and Education: Artificial Intelligence*, 2, 100026. <https://doi.org/10.1016/j.caeai.2021.100026>
- Koul, A. (2019, April 2). *Understanding YouTube’s algorithm in 2019*. Colorado State University. <http://social.colostate.edu/2019/04/02/youtube-algorithm/>

- Koumchatzky, N., & Andryeyev, A. (2017, May 9). Using deep learning at scale in Twitter's timelines. *Twitter Engineering Blog*.
https://blog.twitter.com/engineering/en_us/topics/insights/2017/using-deep-learning-at-scale-in-twitters-timelines.html
- Laestadius, L. (2017). Instagram. In L. Sloan & A. Quan-Haase (Ed.), *The SAGE handbook of social media research methods*. (pp. 573–592). SAGE Publications Ltd.
<https://us.sagepub.com/en-us/nam/the-sage-handbook-of-social-media-research-methods/book245370>
- Lao, N. (2020). *Reorienting machine learning education towards tinkerers and ML-engaged citizens* (Publication No. 28791133) [Doctoral dissertation, Massachusetts Institute of Technology]. Proquest Dissertations and Theses Global.
- Latzer, M., & Festic, N. (2019). A guideline for understanding and measuring algorithmic governance in everyday life. *Internet Policy Review*, 8(2), 1–19.
<https://doi.org/10.14763/2019.2.1415>
- Latzer, M., Festic, N., & Kappeler, K. (2020). Awareness of algorithmic selection and attitudes in Switzerland. *Report 2 from the project: The significance of algorithmic selection for everyday life: The case of Switzerland*. University of Zurich.
<http://mediachange.ch/research/algosig>
- Lee, I., Ali, S., Zhang, H., DiPaola, D., & Breazeal, C. (2021). Developing middle school students' AI literacy. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (pp. 191–197). Association for Computing Machinery.
<https://doi.org/10.1145/3408877.3432513>
- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131.
<https://doi.org/10.1177/1529100612451018>
- Lewandowski, D. (2017). Is Google responsible for providing fair and unbiased results? In M. Taddeo & L. Floridi (Eds.), *The responsibilities of online service providers* (pp. 61–77). Springer International Publishing. https://doi.org/10.1007/978-3-319-47852-4_4
- Lewis, P., & McCormick, E. (2018). *How an ex-YouTube insider investigated its secret algorithm*. <https://www.theguardian.com/technology/2018/feb/02/youtube-algorithm-election-clinton-trump-guillaume-chaslot>
- Library Freedom Project. (n.d.). *Resources*. <https://libraryfreedom.org/resources/>

- Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76-80. <https://doi.org/10.1109/mic.2003.1167344>
- Lindh, M., & Nolin, J. (2016). Information we collect: Surveillance and privacy in the implementation of Google apps for education. *European Educational Research Journal*, 15(6), 644-663. <https://doi.org/10.1177/1474904116654917>
- Lindner, A., Seegerer, S., & Romeike, R. (2019). Unplugged activities in the context of AI. In S. N. Pozdniakov & V. Dagienė (Eds.), *Informatics in schools: New ideas in school informatics* (pp. 123–135). Springer International Publishing. https://doi.org/10.1007/978-3-030-33759-9_10
- Lloyd, A. (2019). Chasing Frankenstein’s monster: Information literacy in the black box society. *Journal of Documentation*, 75(6), 1475-1485. <https://doi.org/10.1108/jd-02-2019-0035>
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Long, D., Jacob, M., Davis, N., & Magerko, B. (2017). Designing for socially interactive systems. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition* (pp. 39–50). Association for Computing Machinery. <https://doi.org/10.1145/3059454.3059479>
- Long, D., & Magerko, B. (2020, April). What Is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- Loyola Marymount University. (2022). *Community of online research assignments*. <https://projectcora.org/>
- Lu, S. (2020). Taming the news feed on Facebook: Understanding consumptive news feed curation through a social cognitive perspective. *Digital Journalism*, 8(9), 1163–1180. <https://doi.org/10.1080/21670811.2020.1837639>
- Luke, A., & Kapitzke, C. (1999). Literacies and libraries: Archives and cybraries. *Pedagogy, Culture & Society*, 7(3), 467–491. <https://doi.org/10.1080/14681369900200066>

- MacCormick, J., & Bishop, C. (2013). *Nine algorithms that changed the future: The ingenious ideas that drive today's computers*. Princeton University Press.
- Machine bias: Investigating algorithmic Injustice [series]. (2015-2022). *ProPublica*.
<https://www.propublica.org/series/machine-bias>
- Mackey, T. P., & Jacobson, T. E. (2011). Reframing information literacy as a metaliteracy. *College & Research Libraries*, 72(1), 62–78. <https://doi.org/10.5860/crl-76r1>
- Madaio, M. A., Stark, L., Wortman Vaughan, J., & Wallach, H. (2020, April). Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–14). Association for Computing Machinery.
<https://doi.org/10.1145/3313831.3376445>
- Marques, L. S., Gresse von Wangenheim, C., & Hauck, J. C. R. (2020). Teaching machine learning in school: A systematic mapping of the state of the art. *Informatics in Education*, 19(2), 283–321. <https://doi.org/10.15388/infedu.2020.14>
- Marshall, P. (2018, July 6) Algorithms and artificial intelligence. *CQ Researcher*, 28(24), 561–584. <http://library.cqpress.com/>
- Massachusetts Institute of Technology Media Lab. (2020). *Project overview: Moral machine*.
<https://www.media.mit.edu/projects/moral-machine/overview/>
- McGrew, S., Breakstone, J., Ortega, T., Smith, M., & Wineburg, S. (2018). Can students evaluate online sources? Learning from assessments of civic online reasoning. *Theory & Research in Social Education*, 46(2), 165-193.
<https://doi.org/10.1080/00933104.2017.1416320>
- Melsión, G. I., Torre, I., Vidal, E., & Leite, I. (2021, June 24-30). Using explainability to help children understand gender bias in AI. [Paper presentation]. *Interaction Design and Children 2021*, Athens, Greece. <https://doi.org/10.1145/3459990.3460719>
- Meyer, J., & Land, R. (Eds.). (2006). *Overcoming barriers to student understanding: Threshold concepts and troublesome knowledge*. Routledge.
- Meyer, J., Land, R., & Baillie, C. (Eds.). (2010). *Threshold concepts and transformational learning*. Sense Publishers.

- Meyer, J. H. F., & Land, R. (2003, May). Threshold concepts and troublesome knowledge: Linkages to ways of thinking and practicing within the disciplines. *Enhancing teaching and learning environments in undergraduate courses, Occasional Report 4*. <http://www.etl.tla.ed.ac.uk/docs/ETLreport4.pdf>
- Microsoft. (2021). *FATE: Fairness, accountability, transparency, and ethics in AI*. <https://www.microsoft.com/en-us/research/theme/fate/>
- Mills, G. E., & Gay, L. R. (2019). *Educational research: Competencies for analysis and applications*. Pearson.
- Min, S. J. (2019). From algorithmic disengagement to algorithmic activism: Charting social media users' responses to news filtering algorithms. *Telematics and Informatics 43*: 101251. <https://doi.org/10.1016/j.tele.2019.101251>
- Mind Your Own Business Act, S. 2829, 117th Cong. (2021) <https://www.congress.gov/bill/117th-congress/senate-bill/2829>
- Mitchell, S., Potash, E., Barocas, S., D'Amour, A., & Lum, K. (2021). Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application*, 8(1), 141–163. <https://doi.org/10.1146/annurev-statistics-042720-125902>
- Molen, G. R., Frank, S., Cohen, J., Spielberg, S., Shusett, R., Goldman, G., Cruise, T., Farrell, C., Morton, S., Sydow, M. von, Smith, L., Stormare, P., Nelson, T. B., Harris, S., Morris, K., Binder, M., London, D., McDonough, N., Harper, J., . . . Dick, P. K. (2002). *Minority report*. DreamWorks Home Entertainment.
- Monzer, C., Moeller, J., Helberger, N., & Eskens, S. (2020). User perspectives on the news personalisation process: Agency, trust and utility as building blocks. *Digital Journalism*, 8(9), 1142–1162. <https://doi.org/10.1080/21670811.2020.1773291>
- Mooney, H. (2019). Sociology of information disorder: An annotated syllabus for informed citizens. In A. Baer, E. S. Cahoy, & R. Schroöder (Eds.), *Libraries promoting reflective dialogue in a time of political polarization* (pp. 113–148). Association of College & Research Libraries. <https://hdl.handle.net/2027.42/151764>
- Morgan, D. (2019). *Basic and advanced focus groups*. SAGE Publications, Inc.
- Mosseri, A. (2018, May 22). *News feed ranking in three minutes flat*. <https://about.fb.com/news/2018/05/inside-feed-news-feed-ranking/>
- National Leadership Council for Liberal Education and America's Promise. (2007). *College learning for the new global century*. Association of American Colleges and Universities.

- Nemeth, C. (2020, February 11). *How the Twitter algorithm works in 2020*. Sprout Social. <https://sproutsocial.com/insights/twitter-algorithm/>
- Newman, N., Fletcher, R., Schulz, A., Andi, S., & Nielsen, R. K. (2020). *Reuters institute digital news report 2020*. Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2020-06/DNR_2020_FINAL.pdf
- News Literacy Project. (2022). The News Literacy Project. <https://newslit.org/>
- Nichols, T. P., & Stornaiuolo, A. (2019). Assembling “digital literacies”: Contingent pasts, possible futures. *Media and Communication*, 7(2), 14–24. <https://doi.org/10.17645/mac.v7i2.1946>
- Noble, S. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
- Nopper, T. K. (2019). Digital character in “the scored society”: FICO, social networks, and competing measurements of creditworthiness. In R. Benjamin (Ed.), *Captivating technology: Race, carceral technoscience, and liberatory imagination in everyday life* (pp. 170–187). Duke University Press. <https://doi.org/10.1515/9781478004493>
- Oakleaf, M. (2011). All together now: Getting faculty, administrators, and staff engaged in information literacy assessment. *portal: Libraries and the Academy*, 11(3), 831–852. <https://doi.org/10.1353/pla.2011.0035>
- Oakleaf, M. (2014). A roadmap for assessing student learning using the new framework for information literacy for higher education. *The Journal of Academic Librarianship*, 40(5), 510–514. <http://dx.doi.org/10.1016/j.acalib.2014.08.001>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- Ochigame, R., & Ye, K. (2021). Search atlas: Visualizing divergent search results across geopolitical borders. In W. Ju, L. Oehlberg, & S. Follmer (Eds.), *Designing Interactive Systems Conference 2021* (pp. 1970–1983). Association for Computing Machinery. <https://doi.org/10.1145/3461778.3462032>
- O’Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

- Optimizing for engagement: Understanding the use of persuasive technology on internet platforms: Hearing before the U.S. Senate Committee on Commerce, Science, and Transportation*, 116th Cong. (2019).
<https://www.commerce.senate.gov/2019/6/optimizing-for-engagement-understanding-the-use-of-persuasive-technology-on-internet-platforms>
- Our Data Bodies. (n.d.). *About us*. <https://www.odbproject.org/about-us-2/>
- Parikh, R. B., & Navathe, A. S. (2022, May 3). How the pandemic made algorithms go haywire. *Slate*. <https://slate.com/technology/2022/05/algorithms-pandemic-health-care-patients-finance.html>
- Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin Press.
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Payne, B. (2019, August). AI + ethics curriculum for middle school. Massachusetts Institute of Technology Media Lab. <https://www.media.mit.edu/projects/ai-ethics-for-middle-school/overview/>
- Pegasystems Inc. (2019, January). *What consumers really think about AI: A global study*. <https://www.pega.com/system/files/resources/2019-01/what-consumers-really-think-about-ai-study.pdf>
- Pessach, D., & Shmueli, E. (2020). Algorithmic fairness. *Arxiv:2001.09784*. 1-31.
<https://doi.org/10.48550/arXiv.2001.09784>
- Pierson, E. (2018). Demographics and discussion influence views on algorithmic fairness. *ArXiv:1712.09124 [Cs]*. <http://arxiv.org/abs/1712.09124>
- Plane, A. C., Redmiles, E. M., Mazurek, M. L., & Tschantz, M. C. (2017). Exploring user perceptions of discrimination in online targeted advertising. In *26th USENIX Security Symposium (USENIX)* (pp. 935–951). The Advanced Computing Systems Association. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/plane>
- Post-truth. (2021, March 25). In *Wiktionary*. <https://en.wiktionary.org/wiki/post-truth>
- Powers, E. (2017). My news feed is filtered? *Digital Journalism*, 5(10), 1315–1335.
<https://doi.org/10.1080/21670811.2017.1286943>
- Poynter Institute. (n.d.). *Politifact*. <https://www.politifact.com/>

- Prado, J. C., & Marzal, M. Á. (2013). Incorporating data literacy into information literacy programs: Core competencies and contents. *Libri*, 63(2). <https://doi.org/10.1515/libri-2013-0010>
- Proferes, N. (2017). Information flow solipsism in an exploratory study of beliefs about Twitter. *Social Media + Society*, 3(1). <https://doi.org/10.1177/2056305117698493>
- Project Information Literacy. (2022). *What is PIL?* <https://www.projectinfolit.org/>
- Protecting Americans from Dangerous Algorithms Act, H.R. 2154, 117th Cong. (2021). <https://www.congress.gov/bill/117th-congress/house-bill/2154>
- Qualtrics. (2021). Qualtrics XM—experience management software. <https://www.qualtrics.com/>
- Rader, E. (2017). Examining user surprise as a symptom of algorithmic filtering. *International Journal of Human-Computer Studies*, 98, 72–88. <https://doi.org/10.1016/j.ijhcs.2016.10.005>
- Rader, E., Hautea, S., & Munasinghe, A. (2020). “I have a narrow thought process”: Constraints on explanations connecting inferences and self-perceptions. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. (pp. 457-488). The Advanced Computing Systems Association. <https://www.usenix.org/conference/soups2020/presentation/rader>
- Rader, E., & Slaker, J. (2017). The importance of visibility for folk theories of sensor data. In *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)* (pp. 257–70). The Advanced Computing Systems Association. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/rader>
- Rainie, L. & Anderson, J. (2017, February). *Code-dependent: Pros and cons of the algorithm age*. Pew Research Center. <https://www.pewresearch.org/internet/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/>
- Rao, A., Schaub, F., Sadeh, N., Acquisti, A., & Kang, R. (2016). Expecting the unexpected: Understanding mismatched privacy expectations online. In *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)* (pp. 77–96). The Advanced Computing Systems Association. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/rao>
- Reed, K. L. (2015). Square peg in a round hole? The framework for information literacy in the community college environment. *Journal of Library Administration*, 55(3), 235–248. <https://doi.org/10.1080/01930826.2015.1034052>

- Register, Y., & Ko, A. J. (2020). Learning machine learning with personal data helps stakeholders ground advocacy arguments in model mechanics. In *Proceedings of the 2020 ACM Conference on International Computing Education Research* (pp. 67–78). Association for Computing Machinery. <https://doi.org/10.1145/3372782.3406252>
- Reidsma, M. (2019). *Masked by trust: Bias in library discovery*. Litwin Books.
- Reviglio, U., & Agosti, C. (2020). Thinking outside the black-box: The case for “algorithmic sovereignty” in social media. *Social Media + Society*, 6(2), 2056305120915613. <https://doi.org/10.1177/2056305120915613>
- Rhoades, A. (2020). Comment: Big tech makes big data out of your child: The FERPA loophole EdTech exploits to monetize student data. *9 American University Business Law Review* 445.
- Rhodes, L. (Producer). (2020). *The Social Dilemma* [Video]. Netflix. <http://www.netflix.com>
- Ridley, M., & Pawlick-Potts, D. (2021). Algorithmic literacy and the role for libraries. *Information Technology and Libraries*, 40(2). <https://doi.org/10.6017/ital.v40i2.12963>
- Rieder, B. (2017). Scrutinizing an algorithmic technique: The Bayes classifier as interested reading of reality. *Information, Communication & Society*, 20(1), 100–117. <https://doi.org/10.1080/1369118X.2016.1181195>
- Rieder, B., Matamoros-Fernández, A., & Coromina, Ò. (2018). From ranking algorithms to “ranking cultures”: Investigating the modulation of visibility in YouTube search results. *Convergence*, 24(1), 50–68. <https://doi.org/10.1177/1354856517736982>
- Rodríguez-García, J. D., Moreno-León, J., Román-González, M., & Robles, G. (2020). Introducing artificial intelligence fundamentals with LearningML: Artificial intelligence made easy. In F. J. Garcia-Penalvo & A. Garcia-Holgado (Eds.), *Eighth International Conference on Technological Ecosystems for Enhancing Multiculturality* (pp. 18–20). Association for Computing Machinery. <https://doi.org/10.1145/3434780.3436705>
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Slavin, A. R., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C., Ng, A. Y., Hassabis, D., Platt, J. C., Creutzig, F., . . . Bengio, Y. (2019). Tackling climate change with machine learning. *arXiv:1906.05433*. <https://doi.org/10.48550/arXiv.1906.05433>
- Rosenberg, M., Confessore, N., & Cadwalladr, C. (2018, March 17). How Trump consultants exploited the Facebook data of millions. *New York Times*. <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html>

- Saracevic, T. (2014, October 20–23). Information literacy in the United States: Contemporary transformations and controversies [Conference session]. *European Conference on Information Literacy*, Dubrovnic, Croatia.
<http://tefkos.comminfo.rutgers.edu/articles.htm>
- Saunders, L. (2007). Regional accreditation organizations' treatment of information literacy: Definitions, collaboration, and assessment. *The Journal of Academic Librarianship*, 33(3), 317–326. <https://doi.org/10.1016/j.acalib.2007.01.009>
- Saunders, L. (2017). Connecting information literacy and social justice: Why and how. *Communications in Information Literacy*, 11(1), 55-75.
<https://doi.org/10.15760/comminfolit.2017.11.1.47> [10.15760/comminfolit.2017.11.1.47](https://doi.org/10.15760/comminfolit.2017.11.1.47)
- Schep, T. (2020). *How normal am I*. Sherpa Pieces. <https://www.sherpapieces.eu/overview/how-normal-am-i>
- Schmidt, J. H., Merten, L., Hasebrink, U., Petrich, I., & Rolfs, A. (2019). How do intermediaries shape news-related media repertoires and practices? Findings from a qualitative study. *International Journal of Communication*, 13, 853–873.
<https://ijoc.org/index.php/ijoc/article/view/9080>
- Schneier, B. (2015). *Data and goliath: The hidden battles to collect your data and control your world*. W.W. Norton & Company.
- Schultheiß, S., Sünkler, S., & Lewandowski, D. (2018). We still trust in Google, but less than 10 years ago: An eye-tracking study. *Information Research: An International Electronic Journal*, 23(3). (ERIC Document Reproduction Service No. EJ1196314).
<https://eric.ed.gov/?id=EJ1196314>
- Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*, 4(2), 2053951717738104.
<https://doi.org/10.1177/2053951717738104>
- Seaver, N. (2019). Knowing algorithms. In J. Vertesi, D. Ribes, C. DiSalvo, Y. Loukissas, L. Forlano, D. K. Rosner, S. J. Jackson, & H. R. Shell (Eds.), *Digital STS* (pp. 412–422). Princeton University Press. <https://doi.org/10.2307/j.ctvc77mp9.30>
- Seidman, I. (2013). *Interviewing as qualitative research: A guide for researchers in education and the social sciences*. Teachers College Press.
- 7th Empire Media. (2020). *Coded bias activist toolkit*. <https://www.codedbias.com/takeaction>

- Shapiro, B. R., Meng, A., O'Donnell, C., Lou, C., Zhao, E., Dankwa, B., & Hostetler, A. (2020). Re-shape: a method to teach data ethics for data science education. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376251>
- Shi, Z. R., Wang, C., & Fang, F. (2020). Artificial intelligence for social good: A survey. *ArXiv:2001.01818* <http://arxiv.org/abs/2001.01818>
- Shin, D. (2020). How do users interact with algorithm recommender systems? The interaction of users, algorithms, and performance. *Computers in Human Behavior, 109*, 106344. <https://doi.org/10.1016/j.chb.2020.106344>
- Shin, D. (2021). How do people judge the credibility of algorithmic sources? *AI & Society*. Advance online publication. <https://doi.org/10.1007/s00146-021-01158-4>
- Shin, D., Kerk, K. F., & Shin, E. Y. (2022). Algorithm awareness: Why user awareness is critical for personal privacy in the adoption of algorithmic platforms? *International Journal of Information Management, 65*, 102494. <https://doi.org/10.1016/j.ijinfomgt.2022.102494>
- Shin, D., & Park, Y. J. (2019). Role of fairness, accountability, and transparency in algorithmic affordance. *Computers in Human Behavior, 98*, 277–284. <https://doi.org/10.1016/j.chb.2019.04.019>
- Shin, D., Rasul, A., & Fotiadis, A. (2021). Why am I seeing this? Deconstructing algorithm literacy through the lens of users. *Internet Research*. Advance online publication.
- Shin, D., Zhong, B., & Biocca, F. A. (2020). Beyond user experience: What constitutes algorithmic experiences? *International Journal of Information Management, 52*, 102061. <https://doi.org/10.1016/j.ijinfomgt.2019.102061>
- Siles, I., Segura-Castillo, A., Solís, R., & Sancho, M. (2020). Folk theories of algorithmic recommendations on Spotify: Enacting data assemblages in the global south. *Big Data & Society, 7*(1), 205395172092337. <https://doi.org/10.1177/2053951720923377>
- Simon, T., & Swartz, A. (2012). *Image atlas*. Net Art Anthology. <https://anthology.rhizome.org/image-atlas>
- Sloan, R., & Thompson, M. (2004). *Epic 2015*. Albino Blacksheep. <https://www.albinoblacksheep.com/flash/epic2015>
- Smith, A. (2018). *Public attitudes toward computer algorithms*. Pew Research Center. <https://www.pewresearch.org/internet/2018/11/16/public-attitudes-toward-computer-algorithms/>

- Smith, A. (2019). *7 things we've learned about computer algorithms*. Pew Research Center. <https://www.pewresearch.org/fact-tank/2019/02/13/7-things-weve-learned-about-computer-algorithms/>
- Smith, B., & Linden, G. (2017). Two decades of recommender systems at amazon.com. *IEEE Internet Computing*, 21(3), 12–18. <https://doi.org/10.1109/MIC.2017.72>
- Smith, C. L., & Matteson, M. L. (2018). Information literacy in the age of machines that learn: Desiderata for machines that teach. *Libri*, 68(2), 71–84. <https://doi.org/10.1515/libri-2017-0025>
- Snopes Media Group Inc. (2022). *Snopes*. <https://www.snopes.com/>
- Souto-Otero, M., & Beneito-Montagut, R. (2013). “Power on”: Googlecracy, privatisation and the standardisation of sources. *Journal of Education Policy*, 28(4), 481-500. <https://doi.org/10.1080/02680939.2013.775347>
- Starke, C., Baleis, J., Keller, B., & Marcinkowski, F. (2021). Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *arXiv:2103.12016*. <https://doi.org/10.48550/arXiv.2103.12016>
- Striphas, T. (2015). Algorithmic culture. *European Journal of Cultural Studies*, 18(4–5), 395–412. <https://doi.org/10.1177/1367549415577392>
- Sulmont, E. (2019). *Machine learning for end-users: Exploring learning goals and pedagogical content knowledge*. (Publication No. 28254426) [Master's thesis, McGill University]. Proquest Dissertations and Theses Global.
- Surveillance Technology Oversight Project. (2022). *Our vision*. <https://www.stopspying.org/our-vision>
- Swanson, B. (2010). Information is personal: Critical information literacy and personal epistemology. In M.T. Accardi, E. Drabinski, & A. Kumbier (Eds.), *Critical library instruction: Theories and methods* (pp. 265–277). Library Juice Press.
- Swart, J. (2021). Experiencing algorithms: How young people understand, feel about, and engage with algorithmic news selection on social media. *Social Media + Society*, 7(2). <https://doi.org/10.1177/20563051211008828>
- Sweeney, L. (2013). Discrimination in online ad delivery. *Communications of the ACM*, 56(5), 44-54. <https://doi.org/10.1145/2447976.2447990>

- Takavarasha, S., Cilliers, L., & Chinyamurindi, W. (2018). Navigating the unbeaten track from digital literacy to digital citizenship: A case of university students in South Africa's Eastern Cape province. *Reading & Writing: Journal of the Reading Association of South Africa*, 9(1), 1-15. <https://doi.org/10.4102/rw.v9i1.187>
- Tewell, E. (2015). A decade of critical information literacy: A review of the literature. *Communications in Information Literacy*, 9(1), 24–43. <https://doi.org/10.15760/comminfolit.2015.9.1.174>
- Thomas, P., Billerbeck, B., Craswell, N., & White, R. W. (2019). Investigating searchers' mental models to inform search explanations. *ACM Transactions on Information Systems*, 38(1), 10:1-10:25. <https://doi.org/10.1145/3371390>
- Thurman, N., Moeller, J., Helberger, N., & Trilling, D. (2019). My friends, editors, algorithms, and I. *Digital Journalism*, 7(4), 447–469. <https://doi.org/10.1080/21670811.2018.1493936>
- TinEye. (2022). *TinEye*. <https://tineye.com/>
- Toner, A., Eckersley, P., & Budington, B. (2017, November 28). *Panopticklick 3.0*. <https://www.eff.org/deeplinks/2017/11/panopticklick-30>
- Touretzky, D. S., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019, June 23–26). *Guidelines for artificial intelligence: What students should know*. [Paper presentation]. ISTE Conference, Philadelphia, PA, USA. https://uocswb03.uocslive.com/ISTE/ISTE2019/PROGRAM_SESSION_MODEL/HANDOUTS/112142285/ISTE2019Presentation_final.pdf
- Townsend, L., Brunetti, K., & Hofer, A. R. (2011). Threshold concepts and information literacy. *portal: Libraries and the Academy*, 11(3), 853–69. <https://doi.org/10.1353/pla.2011.0030>
- Townsend, L., Hofer, A., Hanick, S. L., & Brunetti, K. (2016). Identifying threshold concepts for information literacy: A delphi study. *Communications in Information Literacy*, 10(1). <https://doi.org/10.15760/comminfolit.2016.10.1.13>
- Trepte, S., Teutsch, D., Masur, P. K., Eicher, C., Fischer, M., Hennhöfer, A., & Lind, F. (2015). Do people know about privacy and data protection strategies? Towards the “online privacy literacy scale” (OPLIS). In S. Gutwirth, R. Leenes, & P. de Hert (Eds.), *Reforming European Data Protection Law* (pp. 333–365). Springer Netherlands. https://doi.org/10.1007/978-94-017-9385-8_14
- Turow, J. (2011). *The daily you: How the new advertising industry is defining your identity and your worth*. Yale University Press.

- University of Helsinki. (2019). *Elements of AI*. <https://www.elementsofai.com/>
- Upbin, B. (2012, April 9). Facebook buys Instagram for \$1 billion. Smart arbitrage. Forbes. <https://www.forbes.com/sites/bruceupbin/2012/04/09/facebook-buys-instagram-for-1-billion-wheres-the-revenue/>
- Vaidhyanathan, S. (2011). *The Googlization of everything: (And why we should worry)*. University of California Press.
- Vaidhyanathan, S. (2018). *Antisocial media: How Facebook disconnects us and undermines democracy*. Oxford University Press.
- Valtonen, T., Tedre, M., Mäkitalo, K., & Vartiainen, H. (2019). Media literacy education in the age of machine learning. *Journal of Media Literacy Education*, 11(2), 20–36. <https://doi.org/10.23860/JMLE-2019-11-2-2>
- Van Brummelen, J., Heng, T., & Tabunshchyk, V. (2021, May). Teaching tech to talk: K-12 conversational artificial intelligence literacy curriculum and development tools. In *Proceedings of the 2021 AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*. (pp. 15655–15663). Association for the Advancement of Artificial Intelligence. <https://www.aaai.org/AAAI21Papers/EAAI-20.VanBrummelenJ.pdf>
- van Dijck, J. (2013). *The culture of connectivity: A critical history of social media*. Oxford University Press.
- von Wangenheim, C. G., Marques, L. S., & Hauck, J. C. (2020). Machine learning for all—Introducing machine learning in K-12. *SocArXiv*, 1-10. <https://doi.org/10.31235/osf.io/wj5ne>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Wachter-Boettcher, S. (2017). *Technically wrong: Sexist apps, biased algorithms, and other threats of toxic tech*. W. W. Norton & Company.
- Wang, L., & Fussell, S. R. (2020). More than a click: exploring college students' decision-making processes in online news sharing. In J. Nichols (Ed.), *Proceedings of the ACM on Human-Computer Interaction* (pp. 1–20). Association for Computing Machinery. <https://doi.org/10.1145/3375189>
- Wardle, C. (2018, July 9). Information disorder, part 1: The essential glossary. *First Draft*. <https://medium.com/1st-draft/information-disorder-part-1-the-essential-glossary-19953c544fe3>

- Warshaw, J., Taft, N., & Woodruff, A. (2016, June 22–24). Intuitions, analytics, and killing ants: Inference literacy of high school-educated adults. In *Proceedings of SOUPS 2016: Twelfth Symposium on Usable Privacy and Security (SOUPS)* (pp. 271-285). USENIX: The Advanced Computing Systems Association.
<https://www.usenix.org/conference/soups2016/technical-sessions/presentation/warshaw>
- Weber, M. S., & Kosterich, A. (2018). Coding the news. *Digital Journalism*, 6(3), 310–329.
<https://doi.org/10.1080/21670811.2017.1366865>
- Williams, R., Kaputsos, S. P., & Breazeal, C. (2021, May). Teacher perspectives on how to train your robot: A middle school AI and ethics curriculum. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 15678-15686.
<https://www.aaai.org/AAAI21Papers/EAAI-43.WilliamsR.pdf>
- Willson, M. (2017). Algorithms (and the) everyday. *Information, Communication & Society*, 20(1), 137–150. <https://doi.org/10.1080/1369118X.2016.1200645>
- Wineburg, S., Breakstone, J., Ziv, N., & Smith, M. (2020). *Educating for misunderstanding: How approaches to teaching digital literacy make students susceptible to scammers, rogues, bad actors, and hate mongers* (Stanford History Education Group Working Paper A-21322). Stanford University. <https://purl.stanford.edu/mf412bt5333>
- Wineburg, S., & McGrew, S. (2017). *Lateral reading: Reading less and learning more when evaluating digital information* (Stanford History Education Group Working Paper No. 2107-A1). Stanford University. <http://dx.doi.org/10.2139/ssrn.3048994>
- Wineburg, S., McGrew, S., Breakstone, J., & Ortega, T. (2016). *Evaluating information: The cornerstone of civic online reasoning*. Stanford Digital Repository.
<http://purl.stanford.edu/fv751yt5934>
- Woodruff, A., Fox, S. E., Rousso-Schindler, S., & Warshaw, J. (2018). A qualitative exploration of perceptions of algorithmic fairness. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (pp.1–14). Association for Computing Machinery.
<https://doi.org/10.1145/3173574.3174230>
- Yang, K. C., Varol, O., Davis, C. A., Ferrara, E., Flammini, A., & Menczer, F. (2019). Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, 1(1), 48-61. <https://doi.org/10.1002/hbe2.115>
- Zarouali, B., Boerman, S. C., & de Vreese, C. H. (2021). Is this recommended by an algorithm? The development and validation of the algorithmic media content awareness scale (AMCA-scale). *Telematics and Informatics*, 62, 101607.
<https://doi.org/10.1016/j.tele.2021.101607>

- Zarouali, B., Boerman, S. C., Voorveld, H. A. M., & van Noort, G. (2022). The algorithmic persuasion framework in online communication: Conceptualization and a future research agenda. *Internet Research*. Advance online publication. <https://doi.org/10.1108/intr-01-2021-0049>
- Zhang, B., & Dafoe, A. (2019). *Artificial intelligence: American attitudes and trends*. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3312874
- Zhou, X., Van Brummelen, J., & Lin, P. (2020). Designing AI learning experiences for K-12: Emerging works, future opportunities and a design framework. *ArXiv:2009.10228 [Cs]*. <http://arxiv.org/abs/2009.10228>
- Zhu, S., Hao Yang, H., Xu, S., & MacLeod, J. (2020). Understanding social media competence in higher education: Development and validation of an instrument. *Journal of Educational Computing Research*, 57(8), 1935–1955. <https://doi.org/10.1177/0735633118820631>
- Zoom. (2021). *Zoom meetings*. <https://zoom.us/meetings>
- Zuboff, S. (2019). *The Age of surveillance capitalism: The fight for a human future at the new frontier of power*. Public Affairs.
- Zurkowski, P. B. (1974). The information service environment relationships and priorities. Related paper No.5., National Commission on Libraries and Information Science. (ERIC Document Reproduction Service No. ED100391). <https://files.eric.ed.gov/fulltext/ED100391.pdf>