



## Digital Commons@

Loyola Marymount University  
LMU Loyola Law School

---

Computer Science Faculty Works

Computer Science

---

4-3-2000

## The Next Step

Benjamin Y. Dai

Lynn Thompson

John David N. Dionisio  
*Loyola Marymount University*, [dondi@lmu.edu](mailto:dondi@lmu.edu)

Hooshang Kangarloo

Ricky K. Taira

Follow this and additional works at: [https://digitalcommons.lmu.edu/cs\\_fac](https://digitalcommons.lmu.edu/cs_fac)



Part of the [Computer Sciences Commons](#), and the [Radiology Commons](#)

---

### Digital Commons @ LMU & LLS Citation

Dai, Benjamin Y.; Thompson, Lynn; Dionisio, John David N.; Kangarloo, Hooshang; and Taira, Ricky K., "The Next Step" (2000). *Computer Science Faculty Works*. 3.

[https://digitalcommons.lmu.edu/cs\\_fac/3](https://digitalcommons.lmu.edu/cs_fac/3)

This Article is brought to you for free and open access by the Computer Science at Digital Commons @ Loyola Marymount University and Loyola Law School. It has been accepted for inclusion in Computer Science Faculty Works by an authorized administrator of Digital Commons@Loyola Marymount University and Loyola Law School. For more information, please contact [digitalcommons@lmu.edu](mailto:digitalcommons@lmu.edu).

## The Next Step

In traditional radiology practice, reports are typically dictated and then transcribed. While the free-text reports represent the semantic knowledge interpreted and conveyed by a physician, the information can be hard to access. The advantages of representing medical data in a structured format using standard terminology are clearly recognized. These include the ability to implement a standardized electronic medical record, automatically invoke medical guidelines when appropriate, and conduct outcomes research. Standard structured reports facilitate intelligent indexing, searching, and retrieval of documents from clinical databases. Recent attempts have been made in the industry to enable structured data entry using preformatted templates, but these have yet to gain widespread acceptance.<sup>1,2</sup> These preformatted templates do not necessarily use standard nomenclature and tend to disturb a clinician's normal workflow. This paper presents a prototype system that incorporates the benefits of both dictated free-text reports and standard, structured reports.

### Structured Report Generator

The prototype system incorporates a visually enhanced structured relevant report generator that incorporates speech recognition, natural language processing (NLP), and mapping to a standard nomenclature, Standardized Nomenclature of Human and Veterinary Medicine (SNOMED).<sup>3</sup> The process flow is shown in Figure 1.

The report is dictated in a manner analogous to traditional report dictation. The speech recognition module transcribes the speech input to a text format. The product facilitates editing of the output of the speech recognizer. The text output of the speech module is written as a flat text file to a directory. In addition to the dictated report, the text output of the speech module has a header that includes fields for patient demographics, radiology information systems (RIS) accession number, and radiological procedure code. Note that as long as the free-text report generated is accurate, any speech recognition product can be used in this first step.

The free-text report is then processed by the NLP module that extracts relevant findings. Our module report manager polls a directory every 10 seconds for the arrival of new free-text reports. The NLP is the heart of the module and is based on a statistical model of word interactions.<sup>4</sup> The output from the NLP is a frame, with each frame containing many logical relations. A logical relation includes a head, a relation, and a value (eg, "mass," has \_ size, "equals," "5 cm"). Extensive experiments evaluating various aspects of the algorithm have been performed for chest radiology reports. The parser recall and precision performance values are reported at 89% and 90%, respectively, and semantic interpreter recall and precision values at 79% and 87%, respectively.<sup>4</sup>

The structured findings and related attributes generated from the NLP are presented as an augmented version of the original free text as well as in a structured tabular format (Figure 2). Furthermore, the location of the findings is also depicted in a graphical schematic, detailed below. The structured relevant findings can be edited using either the tabular or graphical template. The editing facilities include deletion, modification, and creation of findings and related attributes.

The structured data are then mapped to a standard nomenclature, SNOMED. The Topography and Morphology axes of the SNOMED 3.5 were loaded in a database (Java-based GemStone/J object-oriented database management system).<sup>5</sup> All six fields associated with each SNOMED code in the Topography and Morphology axes were included in the database, without any modifications. A modified version of the name field was included with each SNOMED code to support string searches and indexed by each word to facilitate key-word searches ignoring word order. The NLP also returns other pointers that facilitate the search: (i) a flag specifying the semantics of the term, eg, "finding type" or "location type," which sets the SNOMED axis of the search and (ii) whether a word is "singular" or "plural." The query is constructed from the NLP term that is normally made of two to four words. The first query consists of all

the words. Failure to locate this combination results in a search with a new query. The new query is formed by dropping modifiers selectively from the combination. This process is iterative until a specific combination results in a successful retrieval. In forming these queries, the key word is never dropped since that may result in retrievals that are not related to the primary concept. The user selects from a list of the best five matches found by the automated search (Figure 3). The module caches verified SNOMED mappings to structured terms in a local table. Preliminary evaluation of this method shows that this type of caching serves to reduce user interventions to verify mapping by a factor of 8 within the first 60 reports.

After completion of the mapping to standard terminology, anatomical locations with verified SNOMED codes are highlighted in the chest schematic (Figure 2, top right). This graphical view is composed of three overlays with anatomically correct representations and parts labeled with SNOMED terms. An overlay is a collection of objects labeled with SNOMED codes from the Topography axes. Each overlay is depicted at one level of the chest anatomy hierarchy in the SNOMED representation. This permits anatomy to be presented to the user at the same detail as specified in the free text. The graphic shown (Figure 2, top right) is the top-level representation of the chest anatomy. A separate graphic is used to show the lymph nodes (Figure 2, bottom right). These symbolic anatomical representations are designed to provide the user with efficient navigation to enable fast review and editing of identified locations.

## Implementation

The overall process functions in a Common Object Request Broker Architecture (CORBA) client/server environment. The stand-alone speech recognition module runs on a Microsoft Windows NT server. The NLP and the SNOMED mapper run on Sun Solaris servers. The output of the speech recognizer is transferred over the network via CORBA method calls to the NLP server for structuring. The NLP server then invokes the appropriate CORBA calls to the SNOMED terminology server for mapping. Finally, the structured, standard output is transferred back to the NT machine for visualization and editing. In this architecture, the speech recognition module and NLP module function as both CORBA clients and servers, while the SNOMED server functions solely as a CORBA server. Though CORBA does provide cross-language support, all development was conducted completely on the Java platform to allow multi-platform and multi-operating system support. Thus, the reporting modules can run on any platform that supports the Java virtual machine. Evaluations have been performed on the visualization module implemented as a stand-alone system.

The cost of the system including hardware and commercial software is listed below: PC, \$2,000; 168MHz server with 128MB of RAM, \$20,000; speech recognition software, stand-alone mode, \$5,948; SNOMED3.5 license (for 2 years), \$2,742; GemStone Object Oriented Database (for 1 year), \$5,000. The implementation described here as well as the costs quoted above are valid for a stand-alone application. No cost is attached to the natural language processing software, the SNOMED term mapper, the CORBA wrapper, and the interface for visualizing and editing since all these components were developed by the authors.

## Benefits of Report generator

Minimal impact on work flow. Compared to systems that incorporate direct structured data entry, the proposed system allows radiologists to retain the traditional method of reporting: dictation. The incorporation of commercial speech recognition software permits the radiologist to view the transcribed report in a matter of seconds as opposed to turnaround times of at least 4-5 hours in conventional transcription. The real innovation in the prototype system occurs after the free-text capture. The next step automatically generates a standard, structured representation of the relevant concepts from the free-text report. Tools are provided to review and edit data. User interventions involve verifying and editing of the outputs of the speech transcriber, natural language processor, and the standard term mapper server. A preliminary clinical evaluation of the system indicates that all three automated processes are sufficiently accurate and require only a few manual edits. Thus, a radiologist has to spend very little extra time for the entire process from dictation to generation of a standard, structured representation. The minimal impact on

radiologist work flow coupled with the traditional dictation report input method will serve to increase clinical acceptability and usage of the module.

Enables data mining. The reports in a conventional RIS are not indexed by the content of the report, so that a large amount of data is potentially lost. The dictated reports represent knowledge conveyed by a physician in a reading of an image. Unfortunately, the free text forms a barrier to the extraction of that information. Traditional database queries that use demographic information from the header of reports do not take advantage of this knowledge. Current approaches of raw keyword database queries lack the necessary specificity to incorporate the semantic relationships of the content in the reports. The problem is exacerbated in very large RIS databases. For example, querying a thoracic database for all reports with a finding of "mass in the right upper lobe" would be conveyed as a query for all reports with the following key words: "mass," "right," "upper," and "lobe." The resulting set of reports would contain all thoracic reports that contain those four key words, which do not necessarily contain the specific finding. The proposed system allows for processing legacy data in RIS databases, permitting queries related to the knowledge in the reports through the indexing of generated standard structured data. For the thoracic database example, the query for the reports would be to find all reports with the finding "mass right upper lobe." The resulting set would contain only those reports with the described finding. The potential of this type of data mining is enormous since it enables data collection for large-scale epidemiological studies.

Relevant retrieval from online medical literature sources. The standard representation in SNOMED codes will allow for mappings to other coding schemes such as Medical Subjects Heading (MeSH)<sup>6</sup> to retrieve relevant medical documents from databases such as MEDLINE. MeSH codes can then be used to retrieve pertinent documents from on-line medical literature and also from on-line indexed radiology teaching sites, providing the user with decision support at the time of image interpretation. The transparent access to online support will benefit both the radiology resident as a teaching tool and the experienced radiologist in the event of a difficult or rare case presentation.

Potential teaching tool. Traditionally, radiologists create teaching files by manually collecting and indexing relevant studies. Most teaching files are static, since a lot of effort is involved in editing these files. A standard, structured representation of concepts in a report offers the exciting potential of direct creation of a teaching file that is dynamic. Further, various modifications to traditional teaching files are possible, since now reports can be created with indexes on findings, diagnoses, and/or location of findings, since all these fields are part of the standard, structured output.

Another aspect of the module that could be used as a teaching tool for both residents as well as referring physicians is the graphical representation of chest anatomy. As several layers of detail are represented in this schematic, it provides an automated method of teaching chest anatomy to nonspecialists and residents. Future implementations will also model findings to expand

the scope of the graphics to include radiology findings.

## **Conclusion**

A radiology reporting workstation incorporates speech recognition, natural language processing, and mapping to a standard representation. It describes the process from a free-text report to a standard, structured one. The potential advantages of such a system include significant applications to data mining, teaching, and decision support. The clinical feasibility has been evaluated and preliminary results are promising in terms of user acceptability, minimal time overhead, minimal user interaction, and accuracy of the standard, structured representation.

*Benjamin Y. Dai, MS, is the associate director of technology, Telemedicine Division*

*Ricky K. Taira, PhD, is associate professor, Department of Radiological Sciences, Children's Hospital & Medical Center, Seattle.*

*Lynn Thompson, MS, senior program analyst, John David N. Dionisio, PhD, is assistant professor, and Hooshang*

*Kangaroo, MD, is professor, Department of Radiological Sciences, University of California, Los Angeles. The authors were awarded a Certificate of Merit Award for the work described in this paper at InfoRad '99, at the annual meeting of the Radiological Society of North America, Chicago.*

**References:**

1. Kahn CE Jr. A generalized language for platform-independent structured reporting. *Methods Inf Med.* 1997;36:163-171.
2. Kahn CE Jr, Wang K, Bell DS. Structured entry of radiology reports using World Wide Web technology. *Radiographics.* 1996;16:683-691
3. Spackman KA, Campbell KE, Cote RA. SNOMED RT: a reference terminology for health care. *Proc AMIA Fall Symposium.* 1997:640-644.
4. Taira RK, Soderland SG. A statistical natural language processor for medical reports. *Proc AMIA Fall Symposium.* 1999: 970-974.
5. Butterworth P, Otis A, Stein J. The GemStone object database management system. *CACM.* 1991;34:64-77.
6. Fowler J, Maram S, Kourmajian V, Devadhar V. Automated MeSH indexing of the World Wide Web. *Proc Annu Symp Comput Appl Med Care.* 1995;19:893-897.