An Empirical Investigation into the Conceptual Structure of Chest Radiograph Findings

Edward Pattison-Gordon, MS
Robert A. Greenes, MD, PhD

Decision Systems Group, Harvard Medical School, Brigham and Women's Hospital
Boston, Massachusetts

A method of investigating the conceptual structure of findings is presented, in which: 1) finding statements are extracted manually from free text, 2) the main concepts in each finding are manually identified and classified, and 3) the resulting sets of classes are examined for insights into finding structure. This study applies the method to chest radiograph reports. Although the subjects of the findings studied fall into seven classes, the same conceptual structure can be used for most of them.

INTRODUCTION

Today's health-care system requires the processing and analysis of patient data not only for the purpose of providing medical care, but also for monitoring the quality of that care, seeking compensation for its provision, and performing research so that it may be improved. Increasing portions of the patient record are available electronically, but the full potential of computer processing of patient data will not be realized until the contents of the electronic medical record can be analyzed by computer. Although this is relatively easy when the data are numeric, such as laboratory results, most physician notes are free text. Two approaches have been explored for handling free text patient data [1]. One approach is to use natural language understanding techniques to try to interpret textual information [2]; the other approach uses structured data entry to capture information in a format that is already interpretable by computer [3, 4, 5, 6]. Both approaches, however, depend on a conceptual model of the information that is explicitly stated as well as underlying knowledge.

Knowledge representation has long been a focus of artificial intelligence, where such formalisms as semantic nets [7], frames [8], and conceptual graphs [9] have been introduced. These systems attempt to provide, in one form or another, ways to construct and use conceptual models. Such models identify concepts and define relationships between them.

We have previously observed that there are two complementary ways to set about modeling concepts and their relations, *top-down* and *bottom-up* [10], and in an earlier effort used the top-down approach to model findings [11]. This paper uses a bottom-up approach to test the following hypothesis:

> Hypothesis. Different types of findings occur in chest radiograph reports, depending upon a finding's subject. That is, their conceptual structures differ by at least one of the following:
> • the sets of relations used;
> • the values allowed for a relation.

This work was carried out as part of our participation in both the National Library of Medicine's Unified Medical Language System© project [12] and in the Canon group's efforts to create a multi-purpose model of chest radiograph findings [13]. It also extends our previous modeling efforts [14]. In this study, finding statements in free text chest radiograph reports were examined. Previous studies have looked at sentences [15], terms [16], and noun phrases [17] to determine conceptual structure.

METHODS

Chest radiograph reports, findings, and main concepts were entered into an object oriented database, Mainstay's Phyla™ for the Apple® Macintosh® computer, (a relational database, however, would be sufficient). The database sorting and retrieval functions supported the analysis of finding conceptual structure.

Reports and Sentences

Eighteen chest radiograph reports were used in the study. These reports had been previously selected for conceptual modeling by the Canon group [13] because of their intermediate length and because of the challenges they posed to modeling.

The reports were entered into the database with each sentence as a separate entry. To preserve context, the complete report was entered including, for example, headings (*e.g.*, "Description," "Impression"), which were entered as separate sentences. Figure 1 shows sample sentences from the reports.

Findings

In order to determine finding conceptual structure, it was necessary to collect all of the information

0195-4210/94/$5.00 © 1994 AMIA, Inc.

CPMC17.01. AP PORTABLE CHEST X-RAY
OHSU12.02. There is coarsening of the interstitial lung markings which are most marked in the lung bases bilaterally.
BWH39.04. There are a few scattered small nodular opacities that are not clearly vascular in origin, possibly representing granulomata.
LDS64.10. There is no evidence of pneumothorax.

Figure 1. Sample sentences from the eighteen chest radiograph reports. Sentence IDs are report IDs—the characters and digits to the left of the first period—followed by the number of the sentence in the report.

provided for each finding from the natural language sentences in the reports. This was done manually, for two reasons. First, because there is not a one-to-one correspondence between sentences and findings in the reports. That is, some sentences contain no findings, others several. Information about the same finding may also be presented in multiple sentences. In addition to headings, sentences without findings include sentences that describe the reason for the radiographic study (*e.g.*, "rule out sepsis and renal failure"), that provide clinical information about the patient (*e.g.*, "pre-op cataract"), that describe the imaging procedure (*e.g.*, sentence CPMC17.01 in Figure 1), that present "information concerning technique" [16] (*e.g.*, "comparison is made to prior PA and lateral chest #1, dated 7 October 1988"), and sentences with "patient management issues" [cf. 16] (*e.g.*, "follow-up film in 6-8 weeks to rule out progression is recommended, if clinically indicated").

Second, when multiple findings are contained in the same sentence, information about one finding may be intermixed with information about the other findings. Multiple findings occur in a sentence when independent findings are joined with "and" (*e.g.*, "the costophrenic angles are clear *and* the domes of the diaphragm are smooth") or when findings that have some item of information in common (*e.g.*, location) are combined so that the information is only stated once (*e.g.*, "again noted is relative elevation and eventration of the left hemidiaphragm"). Multiple findings are also placed in the same sentence when one finding is inferred from the other, linked by words expressing the radiologist's certainty of the inference (*e.g.*, sentence BWH39.04 in Figure 1).

For each finding, the words used to express it in the reports were collected as a statement of the finding. Articles and verbs that did not contribute to the information content of the finding were omitted. For example, "the" was omitted from "the left lung", but "a" was kept in "a pleural effusion" because it specifies the number of effusions (one). Figure 2 gives examples of finding statements.

Usually, identifying findings within a sentence was straightforward but, occasionally, it was difficult. For example, how many findings are in sentence OHSU12.02, shown in Figure 1? Two were chosen (see Figure 2) because the sentence makes two slightly different statements: one about the interstitial lung markings generally and the other particularly about those at the lung bases.

The finding statements were reviewed as a collection, to see if similar findings had been treated consistently, and corrections made. Ultimately, 234 findings were identified in the eighteen reports. Sixteen finding statements were found to be identical—word for word—to other finding statements. They were not removed from subsequent analysis, however, because they are simply more instances of their finding types.

Concepts: Identification and Classification
The main concepts in each finding were identified and classified. Typically, main concepts were the nouns in noun phrases. Although the modifiers in a noun phrase may also represent concepts, they often reflect the conceptual structure of the noun concept, not that of the finding as a whole. For example, consider the noun "opacities" in Finding BWH39.04.01 (Figure 2). The modifiers "a few," "small," and "scattered" specify quantity, size, and distribution, respectively, of the opacity. Thus a model of "opacity" should include these attributes.

OHSU12.02.01. coarsening of interstitial lung markings
OHSU12.02.02. [coarsening of interstitial lung markings] most marked in lung bases bilaterally
BWH39.04.01. a few scattered small nodular opacities
BWH39.04.02. not clearly vascular [opacities] (BWH39.04.01)
BWH39.04.03. possibly granulomata (BWH39.04.01)
LDS64.10.01. no evidence of pneumothorax

Figure 2. Findings from the sample sentences in Figure 1. (Note the first sentence had no findings.) Square brackets indicate text that appeared only once in the original sentence, but was copied into more than one finding. Parentheses enclose references to other findings; here, findings from which the referencing finding was inferred. Finding IDs are sentence IDs plus the number of the finding in the sentence.

Exceptions to this rule were made for modifiers belonging to one of the classes to which main concepts were assigned. In this case, the modifier was assigned to the class and not kept with the noun. "Pleural," in "pleural effusion," for example, was not left as a modifier of effusion but assigned to Body Location (see below).

Concept classification was done manually. Each main concept was examined in turn and grouped with other concepts based on the role it played in the finding (*e.g.*, location, certainty, negation). When the role of a concept was the finding's subject, the concept was further classified with similar subjects (*e.g.*, medical device, medical procedure, pathologic process or entity). New classes were introduced as needed. Classification was reviewed for consistency by examining the members of the resulting classes for concepts that seemed out of place.

In all, thirteen classes were used: X-ray Image Feature (opacity, density, silhouette), X-ray Image View (AP, PA, lateral), Body Location[1] (lung, lobe, heart), Pathologic Process or Entity (effusion, nodule, atelectasis, lung cancer), Physiologic Process (pregnancy, *i.e.*, normal processes), Medical Device (surgical clip, sternotomy wire), Medical Procedure[2] (lobectomy, cardiac transplant), Body Substance (fluid), Comparison Result[3] (new, again, persistent), Inference Certainty[4] (probably, consistent with, less likely), Relationship to Other Finding (related to, due to), Attribute-Value (enlarged, normal, distended, elevated), and Negation (no radiographic evidence).

Attribute-Value is a special class, created because attributes and values for a main concept were often not expressed as modifiers of the main concept. Instead, the attributes and values occurred in other parts of the sentence. In finding OHSU12.02.01, in Figure 2, for example, "coarsening" is not part of the prepositional phrase containing the noun. Such words were classified as Attribute-Value to see if they differed from those modifiers that had been expressed in the concepts' noun phrases. Figure 3 shows an example of concept identification and classification.

[1]Body Location corresponds to Friedman's "body part" and "body part region" modifier classes [16].
[2]Medical Device and Medical Procedure together correspond to Friedman's second "broad informational unit" [cf. 16].
[3]Comparison Result corresponds to Friedman's third broad informational unit and her "change" modifier class [cf. 16].
[4]Inference Certainty roughly corresponds to Friedman's "certainty" modifier class, although negation is treated separately in this study [cf. 16].

Concept	Class
coarsening, most marked	Attribute-Value
interstitial lung markings	X-ray Image Feature
lung bases bilaterally	Body Location

Figure 3. Classification of the main concepts in finding OHSU12.02.02. The resulting set of classes—Attribute-Value, X-ray Image Feature, Body Location—reflects the finding's conceptual structure.

ANALYSIS

To determine finding conceptual structure, the classes to which all of a finding's main concepts had been assigned were taken together, as a set, and properties of these *class sets* were examined. 48 distinct class sets were found among the 234 findings.

Co-occurrence of the 13 classes within the 48 class sets was tabulated. Table 1 shows which classes co-occurred with each other at least once. Sorting the table by the number of classes with which each class co-occurred reveals that five classes never co-occur: Body Substance, Medical Device, Medical Procedure, Pathologic Process or Entity, and X-ray Image Feature. (Physiologic Process also did not co-occur with the other five classes but, because there was only one finding with a concept in this class, it was dropped from further analysis.) The concepts in these mutually exclusive classes had filled the subject role in the findings.

Sorting the 48 class sets according to which of the mutually exclusive classes occurred in them revealed the existence of class sets that did not contain any of the mutually exclusive classes. These class sets fell into two groups: 1) a single class set containing just the class Comparison Result (*i.e.*, findings, like "no change," that summarize comparing the current study with a previous one); and 2) five class sets whose only common feature—besides the absence of the mutually exclusive classes—was the presence of the class Body Location (*i.e.*, findings that describe organs, body parts, and regions as in "continued fullness of AP window"). These class sets represent two additional finding subjects.

An examination of each of the groups of class sets with the same mutually-exclusive class, shows the mutually exclusive class in various combinations with the remaining, non-mutually exclusive, classes. Reasons for this variability may be, for example, the inappropriateness of a role in some contexts (e.g., of Comparison Result when the finding is from the only study) or the omission of details by the radiologist. For the sake of a simpler model, and because a need for this level of detail had not been established, each

	Y	M	D	S	X	P	N	V	O	C	A	R	L	
Body Location		•	•	•	•	•	•	•	•	•	•	•		L
Comparison Result			•	•	•	•			•	•	•		•	R
Attribute-Value			•	•	•			•	•	•		•	•	A
Inference Certainty		•			•	•			•		•	•	•	C
Relation Other Finding		•			•	•				•	•	•	•	O
X-ray Image View			•		•	•						•	•	V
Negation				•		•							•	N
Pathologic Process/Ent.							•	•	•	•		•	•	P
X-ray Image Feature								•	•	•	•	•	•	X
Body Substance							•				•	•	•	S
Medical Device								•				•	•	D
Medical Procedure										•	•		•	M
Physiologic Process														Y
	0	3	4	4	6	6	3	5	7	7	8	8	11	

Table 1. Examining class co-occurrence in class sets reveals a group of mutually exclusive classes; that is, classes that tend to co-occur with the other classes within the class sets, but never with each other (outlined area). Dots indicate two classes that co-occur in at least one class set. The columns were originally ordered by the total number of classes with which classes co-occur (shown across the bottom). Negation and X-ray Image View have been moved in order to group the mutually exclusive classes together. Information below the diagonal is redundant, but is shown in order to make the table easier to read.

group of class sets with the same subject was replaced with a single class created from their union. Table 2 shows the resulting seven class sets.

Attribute usage was analyzed in two ways: 1) within the same finding type, the modifiers in noun phrases were compared with those that occurred elsewhere in the finding statement and, consequently, had been assigned to the class Attribute-Value; and 2) between finding types, concepts that had been classified as Attribute-Value in one finding type were compared against those classified as Attribute-Value in the other finding types.

RESULTS

Relations. A conceptual model of a finding type can be constructed from its class set by representing each class in the set and the finding itself with a concept, then creating relations from the finding concept to each of the class concepts. Because the seven class sets each have a different set of classes, each finding type will have a different set of relations, which is one of the criteria for distinguishing between types. These results are not sufficient, however, to require that a model of chest radiograph findings contain seven finding types. Because of the variability in which classes co-occur with a mutually exclusive class among the findings in the reports and the relatively small number of findings studied, it cannot be concluded that all such co-occurrences have been observed. Examination of additional reports might show, for example, findings in which Medical Device co-occurs with Inference Certainty, or in which Body Location—in the absence of the mutually exclusive classes—co-occurs with X-ray Image View. The more such co-occurrences there are, the more the seven finding types will tend to have the same set of relations. This trend can continue potentially until the only differences between the finding types are the relations to the concepts that represent the mutually exclusive classes themselves. This may not be sufficient grounds to maintain distinct finding types in some applications, especially because a single relation, has_subject, say, could be used to relate a more generic finding type to a concept composed of all of the mutually exclusive classes. A compelling reason for preserving distinct finding types, even if they all had the same set of relations, would be if relation values were different for some finding types.

Relation Values. Of the potential relations—those to Body Location, Comparison Result, Attribute-Value, Inference Certainty, Relationship to Other Finding,

	M	D	S	X	P	R	L	
Body Location	•	•	•	•	•		•	L
Comparison Result		•	•	•	•	•	•	R
Attribute-Value		•	•	•			•	
Inference Certainty	•			•	•		•	
Relation Other Finding	•			•	•		•	
X-ray Image View		•		•	•			
Negation			•		•			
Pathologic Process/Ent.					•			P
X-ray Image Feature				•				X
Body Substance			•					S
Medical Device		•						D
Medical Procedure	•							M
	M	D	S	X	P	R	L	

Table 2. The class sets created from the union of class sets with the same mutually exclusive class. Also shown is the class set containing only Comparison Result and the class set created from the union of those sets that contain the class Body Location but don't contain any of the mutually exclusive classes. Dots show membership of a class (rows) in a set (columns). Sets are labeled with the class that characterizes them: a mutually exclusive class, Comparison Result, or Body Location.

X-ray Image View, and Negation—the only one whose values clearly vary depending on finding type is Attribute-Value. Many attribute/values are particular to specific concepts (*e.g.*, blunting – costophrenic angle, coarsening – interstitial lung markings), but this has implications for modeling those concepts, not findings. For the class Body Location, however, it was observed that modifiers in concepts' noun phrases were used to modify the location (*e.g.*, bilateral, basilar, superior, area in), whereas co-occurring concepts classified as Attribute-Value were used to describe condition or state (*e.g.*, blunted, clear, distended, tortuous). This suggests that a Body Location finding should be modeled with a relation for those attribute/values describing condition and the body locations should be modeled with relations for detailing location.

CONCLUSION

An empirical method of determining the conceptual structure of clinical findings, based on the analysis of finding statements expressed in free text, has been presented. Focusing on the entire statement of individual findings makes it possible to determine concept class co-occurrence within findings and, thus, finding structure.

Although the findings studied can be classified into seven groups based on each finding's subject, it was discovered that it may be possible to represent most of the finding classes using the same set of relations. These results needs to be verified with additional chest x-ray findings and studies need be made to determine the applicability of the model to other domains (e.g. mammogram and ultrasound findings).

ACKNOWLEDGMENTS

This work was supported in part by contract LM04572 from the National Library of Medicine. David Wormuth, MD, Charles Kahn, MD, and Douglas Bell , MD, provided additional medical expertise.

Reference

[1] Sager N, Lyman M, Bucknall C, Nhan N, Tick LJ. Natural language process and the representation of clinical data. J Am Med Informatics Assoc. 1994;1(2):142-160.

[2] Friedman C, Alderson PO, Austin JHM, Cimino JJ, Johnson SB. A general natural-language text Roesner D. Structured data collection and processor for clinical radiology. J Am Med Informatics Assoc. 1994;1(2):161-174.

[3] Bell DS, Greenes RA, Doubilet P. Form-based clinical input from a structured vocabulary: Initial application in ultrasound reporting. In: *Proceedings - 16th SCAMC*. New York: McGraw-Hill, 1992.

[4] Kuhn K, Zemmler T, Reichert M, Heinlein C, knowledge-based user guidance for abdominal ultrasound reporting. In: *Proceedings - 17th SCAMC*. New York: McGraw-Hill, 1993.

[5] Fischer PJ, Stratmann WC, Lundsgarrde HP. User reaction to PROMIS: Issues related to acceptability of medical innovations. In: *Proceedings - 4th SCAMC*. Los Angeles: IEEE Computer Society Press, 1980.

[6] Rector AL, Nowlan WA, Kay S. Foundations for an electronic medical record. Methods Inf Med 1991;30:179-86.

[7] Mac Randal D. Semantic networks. In: Ringland GA and Duce DA, eds. *Approaches to Knowledge Representation: An Introduction*. Letchworth, England: Research Studies Press Ltd, 1988.

[8] Minsky M. A framework for representing knowledge. In: Brachman RJ and Levesque HJ, eds. *Readings in Knowledge Representation*. Los Altos, CA: Morgan Kaufmann, 1985.

[9] Sowa JF. *Conceptual Structures: Information Processing in Mind and Machine*. Reading, MA: Addison-Wesley, 1984.

[10] Barr CE, Komorowski HJ, Pattison-Gordon E, Greenes RA. Conceptual modeling for the Unified Medical Language System. In: *Proceedings - 12th SCAMC*. IEEE Computer Society Press, 1988.

[11] Greenes RA, McClure RC, Pattison-Gordon E, Sato L. The Findings-Diagnosis Continuum: Implications for image descriptions and clinical databases. In: *Proceedings - 16th SCAMC*. New York: McGraw-Hill, 1992.

[12] Lindberg DAB, Humphreys BL, McCray AT. The Unified Medical Language System. In: *1993 Yearbook of Medical Informatics*. Internat. Medical Informatics Assoc., the Netherlands, 1993:41-51.

[13] Evans DA, Cimino JJ, *et. al.* Position statement: Toward a medical concept representation language. J. Am Med Informatics Assoc. 1994;1(3):207-217.

[14] Bell DS, Pattison-Gordon E, Greenes RA. Experiments in Concept Modeling for Radiographic Image Reports. J Am Med Informatics Assoc. 1994;1(3):249-262.

[15] Johnson SB, Gottfried M. Sublanguage analysis as a basis for a controlled medical vocabulary. In: *Proceedings - 13th SCAMC*. 1989.

[16] Friedman C. The UMLS coverage of clinical radiology. In: *Proceedings - 16th SCAMC*. New York: McGraw-Hill, 1992.

[17] Evans DA, Hersh WR. *CXR Reports: Model and Analysis*. Technical Report, Laboratory for Computational Linguistics, CMU, 1994.