Reports from the Canon Group
JAMIA

Research Paper ■

The Canon Group's Effort: Working Toward a Merged Model

CAROL FRIEDMAN, PHD, STANLEY M. HUFF, MD, WILLIAM R. HERSH, MD, EDWARD PATTISON-GORDON, MS, JAMES J. CIMINO, MD

Abstract **Objective:** To develop a representational schema for clinical data for use in exchanging data and applications, using a collaborative approach.

Design: Representational models for clinical radiology were independently developed manually by several Canon Group members who had diverse application interests, using sample reports. These models were merged into one common model through an iterative process by means of workshops, meetings, and electronic mail.

Results: A core merged model for radiologic findings present in a set of reports that subsumed the models that were developed independently.

Conclusions: The Canon Group's modeling effort focused on a collaborative approach to developing a representational schema for clinical concepts, using chest radiography reports as the initial experiment. This effort resulted in a core model that represents a consensus. Further efforts in modeling will extend the representational coverage and will also address issues such as scalability, automation, evaluation, and support of the collaborative effort.

■ **J Am Med Informatics Assoc.** 1995;2:4–18.

One point of general agreement in the medical informatics community is that no common representation for medical information now exists that is generally accepted for use across clinical systems. A possible approach to addressing this need is through collaborative efforts. Several collaborators might effectively divide the labor of modeling a large domain. More importantly, collaborators working in a variety of application areas and representing a diversity of motivations should produce a broader, more useful representation.

Well-defined methods for collaborative development have yet to be described. Simply combining contributions may lead to misunderstandings, redundancies, or ambiguities. "Design by committee" may produce results that end up satisfying no one.[1] The collaborative development of objective, reproducible methods for representing medical information would be a valuable contribution to the field.

Affiliations of the authors: Queens College of the City University of New York, Department of Computer Science (CF), and Columbia University, Center for Medical Informatics (CF, JJC), New York, NY; University of Utah, Salt Lake City, UT (SMH); Oregon Health Sciences University, Portland, OR (WRH); and Brigham and Women's Hospital, Boston, MA (EP-G).

Correspondence and reprints: Carol Friedman, PhD, Queens College of CUNY, Computer Science Department, 65-30 Kissena Boulevard, Flushing, NY 11367. e-mail: friedma@cucis.cis.columbia.edu

Received for publication: 7/23/94; accepted for publication: 8/30/94.

The Canon Group was formed by a group of medical informatics researchers with the common interest of addressing the scientific and technical issues surrounding medical vocabularies,[2] such as the need for coherent conceptual representation across applications and subject domains. A primary tenet of this group is that terms in controlled medical vocabularies should correspond to concepts that have their meanings made explicit through a deep representational structure[3] that may have a variety of uses. Five important aspects associated with an adequate representational structure were identified by the Group: controlled vocabulary, typology, concept model, notation, and granularity. We describe an initial effort in the development of such a structure.

Another tenet of the Canon Group is that the establishment of a broad collaborative effort is necessary to achieve a sharable representational structure and controlled vocabulary. Additionally, because many collaborators are involved, it is especially necessary to delineate and establish procedures, methods, and tools to support the collaborative effort. A recent editorial, which appeared in this journal, challenged the Canon Group to demonstrate the collaborative approach to controlled vocabulary construction.[4] The construction of a sharable representational model for chest radiography report findings is the first experiment in testing the above hypotheses.

Background

The broad goal of the Canon Group, the establishment of a basis for the canonical representation of medical concepts, has been previously described in this journal.[2] The need for such a representation is based on increasing demands for computerized medical information and automated health care systems. The quality of the information and the health care systems is closely related to the quality of the medical-concept representation. Additionally, a common representational model would facilitate the sharing of applications and data among health care institutions. For example, if clinical data had the same representation, it would be possible to collect specified data from different institutions in order to obtain a large amount of data associated with a broad patient population; these data could then be used to perform medical research studies. Even if the different institutions retained their own individual representations, the data could still be shared if they were mapped into the representational form of the common model.

One of the first steps toward establishment of a canonical representation of medical concepts was a meeting of the Canon Group for the purposes of discussing requirements for a representational model and gaining a broader perspective concerning the modeling of medical language. The following tasks were agreed on to initiate the Canon Group's effort:

1. An initial domain consisting of chest radiography reports was chosen.
2. Text from this domain was obtained from four sites to serve as a source for automated text processing.
3. Eighteen reports were chosen for the purposes of information modeling and detailed comparison of each group's approach.
4. A meeting was arranged in January 1993, in Harriman, New York, where different participants presented their models and demonstrated how they would represent the medical information in the chosen reports.

The meeting did not result in agreement on methods or on details of the model. However, a better understanding of the different models and an initial attempt to develop a "merged" model by combining desirable features of the different models did emerge. Although a merged model was not achieved due to time constraints, enough groundwork was covered so that it was possible to begin development of an experimental, yet tangible, model resulting from a collaborative effort.

The Individual Models

Models simplify reality by ignoring certain details of a system in order to focus attention on aspects of the world that suit the purpose of the modeler. The structure and content of a model are driven by the purpose of the model, and models of a given domain will typically converge to the degree that the underlying purposes of the models converge. In the case of the models presented at the Harriman retreat, there are differences relating to the purposes (requirements) of the models. For instance, models such as MedSORT,[5] MOSE,[6] and Galen[7,8] emphasize knowledge-intensive models of medical concepts. MOSE and Galen both aim to define an extensible and application-independent framework that is suitable for building and integrating different terminologies. MedSORT and Galen aim to represent all (and only) valid medical concepts, and to reveal all implicit relationships associated with the concepts. The Queens/Columbia model[9] and the Utah model[10] emphasize medical-data representation geared for decision support and natural-language applications.

These models aim to represent data found in clinical events in a way that is convenient for access by database management systems and for use by clinical applications. The models from Stanford[11] and Harvard[12] emphasize the support of structured data collection. The Stanford model was based on formal logic because it is well suited for reasoning. The Carnegie Mellon/Ohio Health Sciences (CMU/OHSU) model[13] focuses on language normalization to allow translation across representations, from controlled vocabularies to natural-language text. However, the individual models presented at the Harriman retreat also share common factors, and therefore it is possible to frame a discussion of the individual models around these five highly interacting and overlapping themes: terminology (controlled vocabulary), typology, concept model, notation, and granularity.

Every model mentioned controlled vocabulary as a distinct component of the model. Cimino et al. have defined a set of requirements for a controlled vocabulary. These requirements help to define the functional role of the vocabulary in the Queens/Columbia model. The Galen group[8] separate knowledge about terminology from knowledge about contents, an approach also taken by the Queens/Columbia[9] and Harvard[12] groups.

Besides the recognition of vocabulary as a part of the model, there are similarities in the requirements and use of vocabularies among the models. In particular, several groups stated the need for definitions, the need to manage synonyms and homonyms, and the need for domain completeness and nonredundancy. The models expressed the idea that vocabulary terms are symbolic names for underlying medical concepts. Additionally, the CMU/OHSU group[13] expressed the need for molecular to atomic mappings (decomposition) and illustrated the utility of semantically typing the concepts as part of the process of creating canonical representations. Similarly, the Queens/Columbia group[15] expressed the need for compositional mappings that also specified the semantic relations of the atomic components.

A second common theme in most of the models is typology—the use of a semantic network or a hierarchy to organize the terms (representing concepts) into semantic domains that are then referenced within the models.

The third theme discussed at the Harriman retreat involves the way in which concepts are combined to make meaningful expressions of more complex and complete medical concepts. Every model used some formalism to express relationships between vocabulary terms (the names of concepts) to form more complex terms. The Utah group[10] used a frame-oriented paradigm where the vocabulary elements were used as fillers of slots. The name of each slot corresponds to the semantic role that the vocabulary element plays in the model. For instance, the term **lung** is used to fill the slot **body part** and the term **nodule** is used to fill the slot **radiology finding**. However, the most common approach of the modelers was to treat the vocabulary items as nodes in a network and connect the nodes by links that were named relationships. Using this methodology, the radiology finding above would be expressed as **nodule-has location-lung**. In this example, the controlled vocabulary terms **nodule** and **lung** have been connected by the relationship **has location**. The second approach was used in all of the models except the model proposed by the Utah group,[10] but it was noted during discussions that the frame-oriented representation could easily be converted to the named-relationship form.

The fourth theme is related to the notation used to represent the models. The Utah group[10] used frames and slots to describe the model, while the most popular mode of expression used was the conceptual graph (CG) notation,[16] which was used by the Stanford, Columbia, and Harvard groups. The Galen group used the Semantic Meta Knowledge (SMK) notation associated with the Galen work,[7] while the MOSE group described the notation used with the MOSE project.[6] Basically, these different notations are all similar and convertible.

The final theme of the Harriman retreat relates to the granularity of terms. Specific differences in the models are attributable to differences in the granularity of concepts in the vocabulary, the granularity of the hierarchy, and the symbols used to represent the concepts. With regard to concept granularity, there were classic "lumpers" and "splitters." What lumpers expressed as a single concept, splitters would express as closely related simpler concepts. Lumpers preferred to think of **hilar adenopathy** as a single concept, whereas splitters preferred to think of **adenopathy** as the concept and **hilum** as a body part that is the location. Again, it is obvious that the two forms of expression are equivalent and interchangeable as long as the simpler terms are presented in the vocabulary and their relationship to the more complex terms is understood.

In the case of the granularity of the hierarchy, some models had a very fine network of semantic classes, whereas others had a coarser network. For example, those with a flatter semantic network preferred to classify modifier terms such as **large**, **round**, and

white as children of one term, **property**, whereas those with a deeper hierarchy would create intermediate classes such as **size**, **shape**, and **color** that would then be classified as children of **property** and parents of **large**, **round**, and **white**, respectively. Another example of this difference occurred when specifying synonymous terms. For some models *enlarged cardiac silhouette* was synonymous with *enlarged heart*, whereas for others these were similar but non synonymous concepts. Another difference was associated with the actual symbolic names assigned to the concepts. For example, one group would use the name **enlarged heart**, whereas another would use **cardiomegaly**. This difference can be resolved by straightforward substitution, if there is a one-to-one mapping and a precise understanding of the underlying concept.

Two other interesting phenomena that caused some differences in the detailed models were noted. The differences were due to the ambiguous nature of the source material and to the task that consisted of mapping the relevant text of the reports to corresponding concepts in the model. The clinical information was not always interpreted in the same way by different participants. Because the clinical information being modeled was obtained from natural-language reports, the expression of the underlying clinical information often contained ambiguities that were sometimes resolved differently by different participants. For example, the phrase *increased paramediastinal opacity* was interpreted by some participants as referring to a temporal concept denoting *an increase in opacity over time*, whereas others interpreted the phrase as referring to a degree concept denoting *an above normal opacity*. Because differences in interpretations of the actual reports were not part of the modeling exercise per se, it was frequently necessary to disambiguate differences in interpretations from differences in the model itself. It was decided that for the modeling exercise, it was more critical to understand the interpretation being modeled than to decide which interpretation was correct. Although differences in the interpretations of the reports were not particularly relevant to the modeling exercise, the test data supported the argument that free text is inherently ambiguous.

The second problem occurred when the textual expressions in the actual reports were mapped to the model. There was a tendency for the models to associate their symbolic terms with identical or similar surface forms. For example, the word *cardiomegaly* in the reports was typically associated with a concept called **cardiomegaly** in the models. This method of naming concepts was adequate when the word itself was unambiguous. However, when the word was ambiguous, the corresponding concept was likely to be ambiguous also. For example, the word *increased* frequently occurred in the reports, but had at least two different meanings. Therefore, the symbolic name **increased** would be a poor symbol to use in the model because the underlying meaning would probably be misunderstood inadvertently, even if it were described precisely. A better approach would be to use symbols that are unambiguous. For example, it would be more appropriate to use the symbols **temporal increase in** and **above normal degree** to represent the two different concepts associated with the word *increased*.

Methods

To facilitate the modeling effort, the task of representing the entire report was broken up into several well-defined subtasks. The initial subtask chosen was the modeling of individual findings from a small number of sample reports. For example, the phrase *new plate like opacities in left lower lung zone compatible with atelectasis* contains two interrelated findings—*new plate like opacities* and *atelectasis*. These two findings were modeled independently and the connective relation *compatible with* was ignored. This stepwise method of model building made possible the development of a tangible core merged model within a reasonable amount of time so that it could be critiqued by those involved in the collaborative effort and by others working in medical informatics.

Three conventions were adopted that were considered prerequisites for the merging of the models: a common notation was adopted, a common database was established, and a common convention for commenting in the notation system was adopted. A common notation, Sowa's conceptual graph formalism,[16] was chosen as the representational notation for the initial effort because it is widely used in medical informatics[11,17–19] and can be mapped into other knowledge-representation schemas and database forms. Currently, the Knowledge Interchange Format (KIF)[20] developed by the knowledge-sharing project[21] sponsored by the Advanced Research Projects Agency (ARPA) provides a means whereby a representation consisting of CGs can be translated into other knowledge-representation schemas. A database of individual findings from the set of selected reports was established. This database was necessary solely for the collaborative modeling exercise because it provided uniquely labeled findings for identification purposes and helped to disambiguate the interpretations of the clinical information in the sample reports. Because the initial task was restricted to the individual findings only, connective relations were included for completeness but were enclosed in parentheses, as were comments. In addition, conjoined phrases were expanded (when appropriate) and words added (enclosed in square brackets) to make the conjunction more explicit. For example, *new plate like opacities in left mid and lower lung zone* contains the conjoined body location phrase *in left mid and lower lung zone.* That finding was represented as *new plate like opacities in left mid [lung zone] and [left] lower lung zone* because in the original sentence *lung zone* does not immediately follow *left mid* and *left* does not immediately precede *lower lung zone.* The individual findings from a sample report called BWH22 are shown in Appendix A. A convention was also established for representing comments within the notation. A percent sign (%) indicates the start of a comment, which continues to the end of the line. Although comments did not represent information in the reports, they facilitated collaboration and were useful for documentation purposes. In addition, an Internet FTP site was established and members of the Canon Group were given access to it. The different versions of the model and the modeling exercises are maintained at the FTP site.

The current version of the merged model was developed in an iterative fashion by means of workshops, meetings, and electronic interchanges. The first version of the model was developed at a workshop by a subcommittee. It was presented to all of the Canon Group participants, who analyzed it and discussed its problems. It was subsequently modified in accordance with the discussion and placed in a directory at the FTP site. In addition, findings from the database of sample reports were also modeled in accordance with each version of the merged model and added to the server. Every participant was asked to review the latest merged model and the modeling of the findings. The current model represents a consensus that was reached after several rounds of reviews and modifications.

Following the CG formalism, the merged model specifies canonical medical concepts in a form that consists of two major components: one component specifying the semantic classification and hierarchical organization of the concepts, the other containing canonical graphs. Every concept must be associated with a place in the overall hierarchy. The model supports two different versions of a hierarchical organization. The first version is called the "core" hierarchy and classifies the concepts for the purpose of supporting exchange of data using the model. The core hierarchy is a minimal hierarchy consisting of broad classes or axes. A minimal hierarchy was chosen to simplify classification, avoid inconsistencies in classification, and facilitate collaborative efforts because it simplifies the task of mapping to different hierarchies that are likely to be developed by individual sites in support of particular applications.

The second version of the hierarchy, the "specialized" hierarchy, consists of specializations on the core hierarchy. It was chosen to support particular applications and views of the concepts. In the specialized hierarchy, a concept may frequently have multiple parents in order to provide as many classificatory views of the concept as are necessary to support the functions for applications. The specialized hierarchy is application-dependent and is not shown here.

The second component of the model contains canonical graphs consisting of terminologic knowledge about the structure of the concepts and their semantic relationships with each other. Every concept in the model is associated with a unique preferred symbolic name that corresponds to a unique, well-defined con-

```
% TYPE HIERARCHY

concept2 < concept1.        %              concept1
concept3 < concept1.        %                / \
concept4 < concept2.        %               /   \
concept5 < concept2.        %        concept2   concept3
concept5 < concept3.        %           / \        / \
concept6 < concept3.        %          /   \      /   \
                            %     concept4  concept5    concept6

% CANONICAL GRAPHS
[concepti:{"expression1","expression2",...,"expressionM"}] -
    (relation1) -> [concepti1:c1]
    (relation2) -> [concepti2:c2]
              ...
    (relationN) -> [conceptiN:cN].
```

Figure 1 A schematic overview of the organization of the merged model. The first component specifies the type hierarchy. A graph-like version of this hierarchy can be seen on the right-hand side of this component. The second component consists of canonical conceptual graphs that specify the components of concepts along with associated relationships.

cept. Symbolic names were chosen carefully so that the underlying meaning of the concept would be as unambiguous as possible. For example, as mentioned in the Background Section, the symbolic name **increased** would not be appropriate because its meaning is ambiguous; in the merged model, two different symbolic names were assigned to represent the different meanings of the word *increased*: **temporal_increase_in** and **more_than_normal_degree**.

A schematic overview of the organization of the merged model is shown in Figure 1 as a CG. The first component specifies the type hierarchy. A hierarchical classification specifying that concept2 is a subclass of concept1 has the form **concept2** < **concept1**. In Figure 1, a graph-like version of the hierarchy is also shown (on the right-hand side) with the CG statements, because it is easier to visualize than the CG subtype statements. According to Figure 1, the highest concept is **concept1** because it has no parent. **Concept2** and **concept3** are subclasses of **concept1** because they appear to the left of the < symbol and **concept1** appears to the right. **Concept5** is a subclass of both **concept2** and **concept3**.

The second component of the model, as shown in Figure 1, consists of canonical CGs. A canonical CG specifies the components of a complex concept along with the associated relationships, and it may also specify surface form (i.e., textual) expressions of the concept. For example, in Figure 1, **concepti** is related to N other concepts. It has a relation called **relation1** to a concept called **concepti1**. A referent of a concept may be expressed by specifying a colon (:) after the related concept followed by a unique identifier or set, and corresponds to a specific instance of the concept, a set of instances, or a cardinality constraint.

An example of the canonical CG of a concept named **cardiomegaly** is shown in Figure 2. As illustrated in Figure 2, the concept **cardiomegaly** is associated with a set of two expressions found in the text—"cardiomegaly" and "enlarged heart"—that are represented as literal elements of a set following the name of the concept **cardiomegaly**. Although the name of the concept is unique, the mapping from the surface form strings (i.e., textual expressions) to the concept is not necessarily unique. For example, the string "enlarged heart" may be specified in a referent set of another concept, providing a mechanism whereby the ambiguous nature of natural-language expressions may be represented in the model because the possibility exists that a mapping from the text to a concept is not unique. This also serves to differentiate the linguistic level of expression from the unambiguous, well-defined conceptual level. The concept **cardiomegaly** consists of two more basic components. One component is the core observation, **heart**. The other component is the concept **enlarged** that describes a property of the heart. Both concepts must also be defined in the model. The representation of this concept illustrates a phenomenon that is likely to occur when there are collaborators from different orientations developing a model. The modelers did not agree on the representation of this concept, nor on the representation of other concepts consisting of body locations and associated properties. Basically, there were two different views of **cardiomegaly**: one view represented **cardiomegaly** as described above, the other view preferred to represent **cardiomegaly** so that the core observation is **enlarged** and the body location is **heart**. The former view was agreed on so that we could proceed with the merged model. It was realized that there would be differing viewpoints in certain instances, and that compromises would have to be made in the process. This was acceptable to the group members as long as the model was associated with a well-defined semantics that was consistent.

```
[cardiomegaly:{"cardiomegaly","enlarged heart"}] -
    (has_observation)     -> [heart]
    (has_property)        -> [enlarged].
```

Figure 2 The canonical conceptual graph of the concept **cardiomegaly** consists of two more elementary components, an observation concept, **heart**, and a property concept, **enlarged**.

Results

The merged model contains concepts closely associated with terms in the reports, such as **cardiomegaly** and **lung**, but also contains higher level, more abstract concepts that are not generally seen in actual reports because these concepts contain generic structural descriptions of the information rather than the information itself. For example, in the merged model there is a high-level concept called **rad_finding**, as shown in Figure 3, that represents the structure of a generic radiology finding that contains an observation and optional qualifiers. According to Figure 3, a **rad_finding** is a complex concept with components that are also concepts that are interrelated in predetermined ways. For example, the core component of **rad_finding** is a concept that is classified as an observation. Observation concepts are also specified elsewhere in the model, and represent the different observations that occur in radiologic examinations of the chest, such as *pleural effusion* and *coronary artery bypass graft*. Since the domain being modeled consists

```
[rad_finding] -
   (has_observation)->           [observation]
   (has_location)     ->         [body_location:{*}]
   (has_location_qualifier) -> [location_qualifier:{*}]
   (has_presence)     ->         [certainty:{*}]
   (has_degree)       ->         [degree:{*}]
   (has_temporal)     ->         [temporal:{*}]
   (has_quantity)     ->         [quantity:{*}]
   (has_property)     ->         [property:{*}].
```

Figure 3 The canonical conceptual graph associated with the concept **rad_finding**, the generic radiology finding. **Rad_finding** is considered a high-level concept because it is not actually seen in the reports but instead contains a description of the structure of report findings. A finding consists of a core observation relation, **observation**, with different qualifiers (e.g., **body_location**, **location_qualifier**, **certainty**, **degree**, **temporal**, **quantity**, and **property**), all of which are optional and may occur zero or more times.

of only radiologic examinations of the chest, **observation** actually is meant to refer specifically to a radiologic observation concept rather than to a broader observation concept encompassing other observations; in future models it will be renamed **radiological observation**. In **rad_finding**, the naming of the core concept underwent several rounds of changes. Some previous names were **focus**, **subject**, **central finding**, and **core**. In this case, the modelers agreed on the semantics of the component but found it difficult to assign an appropriate name to the concept. The preferred symbolic name **observation** was chosen because it seemed to be the most appropriate and generic name for the concept. The remaining components of **rad_finding** are all optional and may occur zero or more times because they contain the cardinality constraint denoted by {*}. They further qualify the **observation**, and contain different information, such as body location, certainty, degree, and temporal and other descriptive information.

The canonical CG for any concept that occurs in a radiologic observation should be representable according to the specifications set forth for **rad_finding**. For example, **hilar adenopathy** and **platelike atelectasis** are complex terms that are frequently found in radiologic examinations. They could be included in the model as new concepts using CGs similar to the one shown in Figure 3. The CGs for **hilar adenopathy** and **platelike atelectasis** are shown in Figure 4. The concept **hilar adenopathy** has two components, an observation called **adenopathy** and a body location called **hilum**, which is the location of the observation. Similarly, **platelike atelectasis** has two components, an observation, **atelectasis**, and a property qualifier, **platelike**. If the concepts **hilar adenopathy** and **platelike atelectasis** were included in the model as new concepts, they would also have to be given hierarchical assignments.

The concept **rad_finding** can also be used to represent findings from specific radiologic examinations instead of canonical concepts. In this case the CG is not a canonical CG but instead represents an instantiation of the canonical CG **rad_finding**. The instantiation of a CG is represented by an identification marker (a # symbol and an identifier) following the name of the CG. Thus, if the first two findings in a report identified as CXR123 were **possible hilar adenopathy** and **moderate platelike atelectasis**, the corresponding CGs would be as shown in Figure 5. In Figure 5, identifiers are associated with the two **rad_finding** concepts because, for comparison purposes, it is convenient to identify individual findings using a common notation. When the findings in the reports are modeled independently by different modelers, the levels of granularity tend to differ, and therefore the values of the observations may differ. For example, an equivalent CG convention could associate the identifier with instances of **hilar adenopathy** or **platelike atelectasis** instead of with **rad_finding** instances. If applications exist where this representation is necessary, a mapping could be used to transform the report findings so that the observations (i.e., **hilar adenopathy** and **platelike atelectasis**) are associated with identifiers instead.

Other CGs were developed to represent the structure of concepts found in body location information and qualifiers. Body location information is complex because it encompasses spatial information that is difficult to represent. Presently, the CG associated with body location requires more work, but a large variety of body location concepts can be represented properly. Figure 6 illustrates the CG called **body_location**, consisting of three components. One component, represented by the relation **has_location**, is needed when one body location is used to identify another, as in **lymph nodes of right hilum**. The component called **has_location_qualifier** corresponds to the

```
[hilar adenopathy] -
        (has_observation)     ->     [adenopathy]
        (has_location)        ->     [hilum].

[platelike atelectasis] -
        (has_observation)     ->     [atelectasis]
        (has_property)->      ->     [platelike].
```

Figure 4 Two canonical conceptual graphs that correspond to the concepts **hilar adenopathy** and **platelike atelectasis**. These concepts are considered lower level concepts because they are associated with actual phrases that are found in reports.

Figure 5 An instance of a canonical conceptual graph of the core hierarchy is represented by a referent—an identifier that is preceded by a # symbol. The referent is shown following the name of the concept. The representation's of two radiology findings, identified as CXR123.1 and CXR123.2 (*possible hilar adenopathy* and *moderate platelike atelectasis*), that occurred in a report identified as CXR123 are shown.

```
[rad_finding:#CXR123.1] -
        (has_observation)     ->     [hilar adenopathy]
        (has_presence)        ->     [possible].

[rad_finding:#CXR123.2] -
        (has_observation)     ->     [platelike atelectasis]
        (has_degree)          ->     [moderate degree].
```

Figure 6 The canonical conceptual graph for **body_location** contains the representation of a generic body location. A body location concept has a component that is the primary body location and an optional qualifier, **location_qualifier**, that is associated with concepts such as **left**. **Body_location** also has another optional relation, **locative**, that is associated with locative prepositions such as **under**.

```
[body_location] -
       (has_location)              ->    [body_location:{*}]
       (has_location_qualifier)    ->    [location_qualifier:{*}].
       (location_relation)         ->    [locative:{*}].
```

Figure 7 Examples of the conceptual graphs of two complex body location concepts, **left upper lobe of lung** and **medial anterior segment of left upper lobe**. The primary location of **left upper lobe of lung** is **lobe of lung**, which is qualified by two location qualifiers, **left** and **upper**. The primary location of **medial segment of left upper lobe** is **left upper lobe of lung**, which has a location qualifier, **segment**. **Segment** is qualified by **medial** and **anterior**.

```
[left upper lobe of lung] -
   (has_location)               ->    [lobe of lung]
   (has_location_qualifier)     ->    [left]
   (has_location_qualifier)     ->    [upper].

[medial anterior segment of left upper lobe of lung] -
   (has_location) ->                  [left upper lobe of lung]
   (has_location_qualifier)     ->
       [segment] -
           (has_location_qualifier)   -> [medial]
           (has_location_qualifier)   -> [anterior].
```

possible qualifiers of a body location. These could be relative locations, such as **upper** or **base**, or other body locations. For example, in the merged model, a concept **joint of left hand** consists of a body location, **joint**, with a qualifier, **hand**. Some developers may want to model **joint of left hand** using a more detailed representation. For example, **joint** may be viewed as having the relation **part-of** to the body location **hand**. Although a more detailed model of body location may be desirable at a later time, the simpler model was chosen for the current version of the merged model to shorten development time. The remaining relation, **has_location_relation**, refers to qualifiers, such as **under** or **along**, which specify the locative relation of the finding to the body location.

Examples of the CGs of two complex body location concepts, **left upper lobe of lung** and **medial anterior segment of left upper lobe**, are shown in Figure 7. The concept **left upper lobe of lung** consists of a more elementary concept called **lobe of lung** that is qualified by **left** and **upper**. The second concept in Figure 7 is more complex. It consists of the more elementary concept **left upper lobe of lung** that is qualified by the concept **segment**. Similarly, **segment** is qualified by the concepts **medial** and **anterior**.

The representations of other qualifiers, such as **temporal**, **certainty**, and **degree**, are shown in Appendix B, which contains the current version of the core model. A listing of the CGs that represent the findings in the sample report is given in Appendix C.

Discussion

The current version of the merged model was deliberately restricted to a subtask that consisted of the modeling of individual findings. However, other subtasks, such as modeling the overall structure of the report and modeling the interrelations among findings, are considered essential for the final model. Subsequent versions of the model will be extended to handle this information. Adequate representations of information containing spatial relations, uncertainty, fuzzy information, anatomic descriptions, temporal information, and causality are each very difficult and complex subjects, and there is much active research within each of these areas.[22–29] This information is presently represented in the model in very simplistic ways. To develop deeper models within these subareas, a long-range sustained effort will be needed.

Another open issue is how to represent complex concepts in the model. When there is a finding like *hilar adenopathy,* which occurs frequently in the results sections of radiology reports, there are generally two equivalent ways in which the information may be represented. One way is the method described in the Results section, which consists of adding a new definition of a canonical concept, **hilar adenopathy**, to the model. Adding complex terms can result in the proliferation of many new concepts. An alternative method does not involve adding a new concept to the model, and requires that only the more elementary concepts **hilum** and **adenopathy** be in the model. The finding *hilar adenopathy* in a report may then be represented as a **rad_finding** that consists of two related concepts—**adenopathy** where the body location is **hilum**. This method does not involve the proliferation of concepts, but retrieval of the information may be more complicated.

A need identified by the Canon Group is for tools that support collaboration. To support model building, tools are needed for parsing, browsing, and editing concept models. Our modeling effort is a collaboration of geographically separated participants using a variety of computer platforms. Therefore, our tools also need to facilitate the communication of model content and proposed changes across computer platforms and wide distances.

Researchers have noted that participants in any collaborative design effort use a "shared drawing space" to both convey and store information and to mediate interaction.[30] When the participants are gathered together in a design meeting, for example, the shared drawing space is often a white board. Our shared drawing space consists of the current state of the model, as well as the set of individual chest x-ray findings that we are trying to model. The model was represented both by diagrams (either on the white board or on paper) and by CG statements in the linear notation.

Four different collaborative interactions have been identified and classified according to whether the times and locations of the interactions are the same.[31] These interactions are 1) face-to-face (same time, same place), 2) distributed synchronous (same time, different places), 3) asynchronous (different times, same place), and 4) distributed asynchronous (different times, different places).

Obviously, when the participants are geographically separated, the same physical white board cannot serve as the shared drawing space. Instead, other tools are needed to support the different collaborative interactions. ("Groupware" is the term that has come to be applied to computer software tools designed to facilitate collaborative interaction.) Our choice of interaction was limited by the tools for collaboration available to us. Face-to-face collaboration occurred in meetings and was supported by the usual white boards, overheads, and paper handouts. E-mail supported distributed asynchronous collaboration, especially while writing papers (like this one), but we did not have, for example, tools that support distributed synchronous model building.

In fact, e-mail and access to the Internet were about the only computer tools shared by all of us. The choice of CGs as a formalism enabled collaborators to speak the same language and, using the linear notation, to exchange models via e-mail. But not every collaborator had a parser for the linear notation and CG editing and browsing tools. Parsing CGs readily reveal errors made with linear notation syntax and, occasionally, semantic errors as well. Display tools can make it easier to see a large number of concepts, and their relationships, at once. For example, an outline viewer (in which hierarchies are viewed as outlines, with descendants indented beneath ancestors) implemented by one of the Canon Group members (Bell) has helped reveal semantic errors.

As we have progressed, and as the model has grown in complexity, the need for CG parsers, browsers, and editors on multiple platforms has become more pronounced. There is also a need for a tool, such as the Standard Generalized Markup Language (SGML),[32] to provide a means of specifying and standardizing the format of the reports so that they can readily be shared by others. These tools would support distributed asynchronous collaboration, whereas tools that display CGs could also support face-to-face collaboration. If this tool set were augmented with some sort of real-time messaging system, then that would enable distributed synchronous collaboration as well.

The merged model is an experimental model, developed by merging the models of some of the participants at the Canon Group's meeting.[9–13] It is an incomplete model and in its present form accounts for a small subset of clinical information. Even though it encompasses a very small piece of the overall goal, it represents an important achievement in that it subsumes independent work at four sites associated with four applications and orientations. It provides a medium whereby participants can communicate using a common language, and thus makes it possible for those in the group to analyze and criticize the actual modeling effort more accurately. Since it is a partial and an experimental model, it will continue to change and evolve as it is extended and applied.

Future Plans

In this article, a model was described that emerged from a unified vision, as well as from continued collaborative efforts. However, it is a small part of what is truly necessary to meet the goals described in the Canon Group's position paper.[2] This section provides an overview of future work deemed necessary for the Canon Group, consisting of four major themes: 1) scalability and generalizability, 2) automation, 3) evaluation, and 4) collaboration and support.

The Canon Group's work must scale-up and generalize, both in the chest x-ray (CXR) report domain and in other domains. To this end, we must take advantage of the 10,000 remaining reports that we collected in our initial work. Creating a workable model for all of these reports will require some resources that currently do not exist. We advocated in our position paper that we needed a grammar for the formation of medical concepts, consisting of basic resources (i.e., basic lexical units, typology, and an inventory of basic concepts) and procedures (i.e., rules of composition). For the current model, we have focused more on the collaborative process than on the resources. As we scale-up and, particularly, as applications are built, we need to explicitly create these resources.

Another aspect to scalability concerns making all of our methods computationally tractable. This includes not only devising efficient algorithms, but also adopting technologies that make the product useful. One such technology is an Internet-based client-server architecture. This architecture will allow collaborators and their applications to access Canon Group resources. Once we achieve scale in the CXR report domain, we must determine which aspects of our work generalize to other domains. There is considerable interest in handling what is probably the most unstructured of all medical data, the physical examination. As with the CXR reports, we will obtain data from diverse sites and will repeat the process. Initially we will focus on one aspect of the physical examination, such as the cardiac or abdominal examination, and will build outward.

The second consideration is automation. With the large volume of medical data generated daily, there must be considerable, if not complete, automation of these processes. Modeling CXR reports by hand may enable us to collaboratively understand and analyze the underlying conceptual issues, but ultimately the modeling must be nearly fully automated. The natural-language processor developed by Friedman et al.[15] was used to automatically process and structure the CXR reports in accordance with the model proposed by their group.[9] It is possible that the same natural-language processing system could be used to automate the processing of the CXR reports in accordance with the merged model. Another part of the process that must be automated is the building of the model itself. As mentioned by Evans and Hersh[13] the CLARIT system provided automated noun-phrase extraction and first-order thesaurus construction, allowing large numbers of terms and modifiers to be discovered.

The third aspect of our future plans is the need for evaluation, both to provide us with a measure of our work and, if we are successful, to convince others to adopt our approach. There are several planned approaches to evaluation. These approaches are not mutually exclusive and include:

1. Evaluation of the model in each individual group's application, such as decision support and structured data entry. The benefit of this approach is the establishment of the operational use of the system, while the drawback is the possible inability to control for variables outside the context of the vocabulary.

2. Evaluation of the model in different sites by sharing clinical data. The benefit of this approach is the direct operational assessment of the model between sites and facilitation of sharing of data.

3. Ensuring consistent mapping back and forth between the model and the original text. The benefit of this approach is the direct assessment of mapping back and forth. The drawback is the lack of evaluation in an operational setting.

4. Presenting the model to clinicians for evaluation. The benefit of this approach is to have assessment by the people whose language we are modeling, while the drawback is its inherent subjectivity.

The fourth theme of future work concerns collaboration and support. While we have found that a small focused group has enabled us to move beyond mere ideas, we will not consider our work a success unless it is adopted for use in operational systems. Collaboration, of course, requires support in many forms. We will obviously need the support of the producers of existing vocabularies, not only to map our representations to their terminologies, but also to utilize their terminologies.

Conclusions

The development of a core merged model is a small but critical step in addressing what we identify as the central challenge of medical informatics—development of a generally accepted model for representing clinical information. Because the merged model has been developed collaboratively, it has been deemed acceptable on an experimental basis by a number of different sites involved in medical informatics. This is a substantial step in the required direction.

The effort so far enhances the level of discussion about and activity for developing a standard model for medical-concept representation. The work described here is an initial effort that provides a foundation we hope will ultimately be appropriated for use in tangible clinical applications. A widely accepted standard model for medical-concept representation would provide a mechanism whereby sharable applications and data could become a reality and true collaboration could be feasible.

Members of the Canon Group who collaborated in the development of the merged model include (alphabetically):

Douglas S. Bell, MD
Keith E. Campbell, MD
Christopher G. Chute, MD, DrPH
James J. Cimino, MD
David A. Evans, PhD
Carol Friedman, PhD
Robert A. Greenes, MD, PhD
William R. Hersh, MD
Stanley M. Huff, MD
Stephen B. Johnson, MD
Robert C. McClure, MD
Mark A. Musen, MD, PhD
Edward Pattison-Gordon, MS
Alan Rector, MD, PhD
Roberto Rocha, MD

References ■

1. Feinstein AR. ICD, POR, and DRG: unsolved scientific problems in the nosology of clinical medicine. Arch Intern Med. 1988;148(10):2269–74.
2. Evans DA, Cimino JJ, Hersh WR, Huff SM, Bell DS. Toward a medical-concept representation language. J Am Med Informatics Assoc. 1994;1:207–17.
3. Musen MA. Dimensions of knowledge sharing and reuse. Comput Biomed Res. 1992;25(5):435–67.
4. Tuttle MS. The position of the Canon Group: a reality check. J Am Med Informatics Assoc. 1994;1:298–9.
5. Evans D. Final Report on the MedSORT-II Project: Developing and Managing Medical Thesauri. Technical report. Pittsburgh, PA: Laboratory for Computational Linguistics, Carnegie Mellon University, 1987.
6. Rossi-Mori A, Bernauer J, Pakarinen V, et al. CEN/TC251/PT003 models for representation of terminologies and coding systems in medicine. In: Proceedings of the Seminar Opportunities for European and U.S. Cooperation in Standardization in Health Care Informatics. Geneva, Switzerland, September 1992.
7. Rector AL, Nowlan WA, Kay S. Conceptual knowledge: the core of medical information systems. In: Lun KC, Degoulet P, Plemme TE, Rienhoff O, eds. Proceedings of MEDINFO 92. Amsterdam: North-Holland, 1992:1420–6.
8. Rector AL, Nowlan WA, Glowinski A. Goals for concept representation in the GALEN project. In: Safran C, ed. Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1994:414–8.
9. Friedman C, Cimino JJ, Johnson SB. A schema for representing medical language applied to clinical radiology. J Am Med Informatics Assoc. 1994;1:233–48.
10. Huff SM, Rocha RA, Haug PJ, Bray BE, Warner HR. An Event Model of Medical Information Representation. Technical report. Salt Lake City, UT: University of Utah, 1994.
11. Campbell KE, Das AK, Musen MA. A logical foundation for representation of clinical data. J Am Med Informatics Assoc. 1994;1:218–32.
12. Bell DS, Pattison-Gordon E, Greenes RA. Experiments in concept modeling for radiographic image reports. J Am Med Informatics Assoc. 1994;1:249–62.
13. Evans D, Hersh B. The CXR Reports: Model and Analysis: Background paper for CANON Group workshop. Technical report. Pittsburgh, PA: Laboratory for Computational Linguistics, Carnegie Mellon University, 1993.
14. Cimino JJ, Clayton PD, Hripcsak G, Johnson SB. Knowledge-based approaches to the maintenance of a large controlled medical terminology. J Am Med Informatics Assoc. 1994;1:35–50.
15. Friedman C, Alderson PO, Austin JHM, Cimino JJ, Johnson SB. A general natural-language test processor for clinical radiology. J Am Med Informatics Assoc. 1994;1:161–74.
16. Sowa JF. Conceptual Structures. Reading, MA: Addison-Wesley, 1984.
17. Bernauer J. Conceptual graphs as an operational model for descriptive findings. In: Clayton PD, ed. Fifteenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1991:214–8.
18. Schroder M. Knowledge based analysis of radiological reports using conceptual graphs. In: Pfeiffer HD, ed. Proceedings of the Seventh Annual Workshop on Conceptual Graphs, Las Cruces, New Mexico. Berlin: Springer-Verlag, 1992:213–22.
19. Rassinoux AM, Baud RH, Scherrer JR. Conceptual graphs model extension for knowledge representation of medical texts. In: Lun KC et al., ed. Proceedings of MEDINFO92. Amsterdam: North-Holland, 1992:1368–74.
20. Genesereth MR, Fikes RE. Knowledge Interchange Format, Version 3.0 Reference Manual. Technical report logic-92-1. Stanford, CA: Stanford University, 1992.
21. Neches R, Fikes RE, Finin T, et al. Enabling technology for knowledge sharing. AI. 1991;12(3):16–36.
22. Kahn DC. Modeling time in medical decision support programs. Med Decis Making. 1991;11:249–64.
23. Console L, Furno A, Torasso P. Dealing with time in diagnostic reasoning based on causal models. In: Methodologies for Intelligent Systems, vol. 3. Amsterdam: North-Holland, 1988.
24. Allen JF. Towards a general theory of action and time. Artif Intell. 1984;23(2):123–54.
25. Zadeh LA. Commonsense knowledge representation based on fuzzy logic. IEEE Comput. 1983;16(10):61–6.
26. Maida AS, Shapiro SC. Intensional concepts in propositional semantic networks. In: Readings in Knowledge Representation. San Mateo, CA: Morgan Kaufmann, 1985:169–89.
27. Davis E. Representations of commonsense knowledge. In: Morgan Kaufmann Series in Representation and Reasoning. San Mateo, CA: Morgan Kaufmann, 1990.
28. Han R. Knowledge Representation: An AI Perspective. Tutorial monographs in cognition science. Norword, NJ: Ablex

Publishing, 1991.

29. Chen S, ed. Advances in Spatial Reasoning. Norword, NJ: Ablex Publishing, 1991.
30. Ishii H, Miyake N. Toward an open shared workspace: computer and video fusion approach of TeamWorkstation. Communications ACM. 1991;34(12):37–50.
31. Hsu J, Lockwood T. Collaborative computing. Byte. 1993; 18(3):113–20.
32. Goldfarb CF. The SGML Handbook. Oxford, UK: Clarendon Press, 1990.

APPENDIX A

Individual Findings in a Sample X-ray Report

```
BWH22.07|compatible with atelectasis (based on finding BWH22.06)
BWH22.12|consistent with coronary artery bypass graft (based on finding
+BWH22.11)
BWH22.14|[consistent with] previous lobectomy on the right (based on
finding +BWH22.13)
BWH22.10|left lower lobe atelectasis
BWH22.03|[new] 4 intact sternotomy wires
BWH22.06|new plate like opacities in left [mid lung zone and left] lower lung
+zone
BWH22.02|new surgical clips in distribution of circumflex artery
BWH22.04|persistent increased right paramediastinal opacity
BWH22.05|possibly related to previous radiation therapy (based on finding
+BWH22.04)
BWH22.11|post-operative changes
BWH22.13|[post-operative changes]
BWH22.09|slight interval decrease in left pleural effusion
BWH22.08|some interval improvement in left pleural effusion
BWH22.01|surgical clips again along right mediastinum and [along] right hilar
+region
```

APPENDIX B

Core Merged Model

```
/**********************************************************************/
/*                     THE CORE MERGED MODEL                          */
/**********************************************************************/

% Statements bracketed by "/*  */" and statements following "%" are comments.

% The concepts specified in the type assignment statements
% specify the controlled names that are present in the model.

% TYPE HIERARCHY

% Rad Findings
% higher level findings
 rad_finding < finding.

% lower level findings
 air_bronchogram < rad_finding.
 atelectasis < rad_finding.
 blunting_of_costophrenic_angle < rad_finding.
 bony_changes < rad_finding.
 calcification_of_lymph_node < rad_finding.
 cardiac_silhouette_upper_limits_of_normal < rad_finding.
 cardiac_silhouette_within_normal_limits < rad_finding.
 cardiomegaly < rad_finding.
 clear_lung < rad_finding.
 coronary_artery_bypass_graft < rad_finding.
 degeneration_of_thoracic_spine < rad_finding.
 degenerative_joint_disease < rad_finding.
 discoid_atelectasis < atelectasis.
 elevation_of_hemidiaphragm < rad_finding.
 granulomatous_disease < rad_finding.
 interstitial_markings_in_lung < rad_finding.
 kyphosis < rad_finding.
 kyphosis < rad_finding.
 nasogastric_tube < rad_finding.
 peribronchial_cuffing < rad_finding.
 pneumothorax < rad_finding.
 pleural_effusion < rad_finding.
 pleural_thickening < rad_finding.
 scoliosis < rad_finding.
 s_shaped_scoliosis < scoliosis.
 sternotomy < rad_finding.
 sternotomy_wire < rad_finding.
 subsegmental_atelectasis < atelectasis.
 tortuous_aorta < rad_finding.
 widening_of_mediastinum < rad_finding.

% Observations
rad_finding < observation.
body_location < observation.
active_disease < observation.
acute_disease < observation.
air_accumulation < observation.
calcified_granuloma < granuloma.
cancer < observation.   %disease
circumscribed_density < observation.
collapse < observation.
compression_fracture < fracture.
consolidation < observation.  % is "lung" implied
curvilinear_density < density.
edema < observation.
effusion < observation.
fibrosis < observation.
fluid < observation.
fluid_overload < observation.
focal_opacity < circumscribed_density.
fracture < observation.
granuloma < observation.
granulomatous_disease < observation.
infiltrate < observation.  %is "lung" implied
linear_fibrosis < fibrosis.
linear_opacity < circumscribed_density.
multiple_sclerosis < observation.  % disease
nodular_opacity < circumscribed_density.
nonunionized_fracture < fracture.
pleural_effusion < effusion.
rounded_density < circumscribed_density.
scarring < observation.
trauma < observation.

% Medical Devices

nasogastric_tube < observation.
prosthetic_valve_ring < observation.
sternotomy_wire < observation.
surgical_clip < observation.
wire < observation.

% Surgical procedures
therapeutic_procedure < observation.
surgical_procedure < observation.
coronary_artery_bypass_graft < observation.
lobectomy < observation.
sternotomy < observation.

% Body Location Concepts

7th_rib < rib.
aorta < body_location.
aortic_valve < body_location.
aorto_pulmonary_window < body_location.
blood_vessels < body_location.
cardiac_silhouette < body_location.
cardiopulmonary < body_location.
chest_wall < body_location.   %partof chest
costophrenic_angles < body_location.
diaphragm < body_location.
distribution_of_circumflex_artery < body_location.
extrathoracic < body_location.
heart < body_location.
hemidiaphragm < body_location.
hilum < body_location.
left_lower_lobe_of_lung < body_location.    %part_of lung
```

```
left_lower_lung_zone < body_location.        %part_of lung
left_upper_lobe_of_lung < body_location.     %part_of lung
lobe_of_lung < body_location.
lung < body_location.
lymph_node < body_location.
major_fissure < body_location.
mediastinum < body_location.
paramediastinum < body_location.
pleura < body_location.
pleural_space < body_location.
pulmonary_blood_vessels < blood_vessels.
rib < body_location.
right_lower_lobe_of_lung < body_location.   %part_of lung
soft_tissue < body_location.
spine < body_location.
subpulmonic < body_location.
thoracic_spine < body_location.  %is this same as thoracic vertebral body?
thoracic_vertebral_body < body_location.
vertebral_body < body_location.

% Location Qualifiers
body_location < location_qualifier.
body_location_part < location_qualifier.
laterality < location_qualifier.
locative < location_qualifier.
orientation < location_qualifier.
quantity < location_qualifier.
relative_location < location_qualifier.

% Laterality

bilateral < laterality.
right < laterality.
left < laterality.

% Body_location_part

area_of < body_location_part.
bibasilar < body_location_part.
border < body_location_part.
field < body_location_part.
lobe < body_location_part.
region < body_location_part.
segment < body_location_part.
wall < body_location_part.
zone < body_location_part.

% Relative Locations

anterior < relative_location.
base < relative_location.
inferior < relative_location.
lateral < relative_location.
lower < relative_location.
median < relative_location.
mid < relative_location.
posterior < relative_location.
upper < relative_location.

% Orientation

anterior_posterior < orientation.
horizonal < orientation.
lateral < orientation.
transverse < orientation.

% Qualifiers
% types of qualifiers
degree < qualifier.
orientation < qualifier.
position < qualifier.
quantity < qualifier.
temporal < qualifier.
property < qualifier.

% lower level qualifiers
calcification < property.
coarse < property.
density < property.
elevated < position.
focal < property.
hazy < property.
intact < property.
scattered < property.
smooth < property.
```

```
% Density Qualifiers

clear < property.
curvilinear density < property.
opaqueness < property.

% Shape Qualifiers

shape_qualifier < property.
lateral_deviation < shape_qualifier.
platelike < shape_qualifier.
round < shape_qualifier.
s_shaped_deviation < lateral_deviation.
tortuous < shape_qualifier.

% Degree Qualifiers
extensive < high_degree.
high_degree < degree.
large_amount < degree.
mild_degree < degree.
minimal < mild_degree.
moderate_degree < degree.
more_than_normal_degree < degree.
severe_degree < degree.
slight < mild_degree.
some < mild_degree.

% Size Qualifiers

qualitative_size < property.
quantitative_size < property.
enlargement < qualitative_size.
large < qualitative_size.
normal_size < qualitative_size.
prominent < qualitative_size.
size_within_normal_limits < qualitative_size.
small < qualitative_size.
thickening < qualitative_size.
widening < qualitative_size.

% Temporal Qualifiers

change < temporal.
again < temporal.
chronic < temporal.
decrease_in < change.
decrease_in_size < decrease_in.
healed < change.
improved < change.
temporal_increase_in < change.
temporal_increase_in_intensity < temporal_increase_in.
temporal_increase_in_number < temporal_increase_in.
temporal_increase_in_size < temporal_increase_in.
interval < temporal.
interval_development < temporal.
new < temporal.
no_change < change.
no_change_from_previous_exam < no_change.
no_change_in_intensity < no_change.
no_change_in_number < no_change.
no_change_in_position < no_change.
no_change_in_size < no_change.
persistent < temporal.
post_operative change < change.
previous < temporal.
remain < change.
remain_in_place < remain.
resolved < change.
status post < temporal.

% Certainty Qualifiers

absent < certainty.
cannot_rule_out < low_certainty.
connective < certainty.
evidence_of < moderate_certainty.
high_certainty < certainty.
history_of < high_certainty.
likely < high_certainty.
low_certainty < certainty.
moderate_certainty < certainty.
possible < moderate_certainty.
present < high_certainty.
probable < moderate_certainty.
```

```
unlikely < low_certainty.
undetermined < certainty.

% Quantities

'>1' < fuzzy_quantity.
a_few < fuzzy_quantity.
fuzzy_quantity < quantity.
many < fuzzy_quantity.
multiple < fuzzy_quantity.
number < quantity.

% Connective Relations

compatible_with < connective.
consistent_with < connective.
may_represent < connective.
most_likely_represent < connective.
related_to < connective.

% Dimensions

diameter < dimension.
length < dimension.
volume < dimension.
width < dimension.

% Units

cm < unit.
mm < unit.

% Locative

along < locative.
under < locative.
adjoining < locative.
in < locative.


% CANONICAL GRAPHS
% Canonical graphs are defined only for concepts which are complex -
% i.e. they are composed of other concepts that are more elementary.

[rad_finding] -
      (has_observation)->          [observation]
      (has_location)   ->          [body_location]
      (has_location_qualifier) -> [location_qualifier:{*}]
      (has_certainty)   ->         [certainty:{*}]
      (has_degree)      ->         [degree:{*}]
      (has_temporal)    ->         [temporal:{*}]
      (has_quantity)    ->         [quantity:{*}]
      (has_property)    ->         [property:{*}].

[body_location] -
       (has_location)           -> [body_location:{*}]
       (has_location_qualifier) -> [location_qualifier:{*}]
       (location_relation)      -> [locative:{*}].

[location_qualifier] -
       (has_location_qualifier) -> [location_qualifier: @>0].

[quantitative_size] -
       (has_dimension)   -> [dimension]
       (has_measurement) -> [measurement]
       (has_orientation) -> [orientation].

[measurement] -
       (has_quantity) -> [number]
       (has_unit)     -> [unit].

[temporal] -
       (has_degree) -> [degree:{*}]
       (has_certainty) -> [certainty:{*}].

[certainty] -
       (has_degree) -> [degree:{*}].

% Rad Findings

[calcification_of_lymph_node] -
      (has_observation)  -> [calcification]
      (has_location) -> [lymph_node].

[pneumothorax:{"pneumothorax","air in pleural space"}] -
       (has_observation)  -> [air_accumulation]
       (has_location) -> [pleural_space].

[pleural_effusion] -
       (has_observation)  -> [effusion]
       (has_location) -> [pleural_space].

[cardiac_silhouette_within_normal_limits:
    {"cardiac silhouette within normal limits","normal size heart"}] -
        (has_observation)-> [cardiac_silhouette]
        (has_property) -> [size_within_normal_limits].

[cardiomegaly:
    {"cardiomegaly","cardiac enlargement","enlargement of heart",
          "enlargement_of_cardiac_silhouette"}] -
        (has_observation)    -> [heart]
        (has_property) -> [enlargement].

[clear_lung] -
        (has_observation)    -> [lung]
        (has_property) -> [clear].

[elevation_of_hemidiaphragm] -
        (has_observation)    -> [hemidiaphragm]
        (has_property) -> [elevated].

[pleural_thickening] -
        (has_observation)   -> [pleura]
        (has_property) -> [thickening].

[scoliosis:{"scoliosis","lateral deviation of spine"}] -
        (has_observation)    -> [spine]
        (has_property) -> [lateral_deviation].

[s_shaped_scoliosis] -
        (has_observation)    -> [spine]
        (has_property) -> [s_shaped_deviation].

[tortuous_aorta:{"tortuous aorta","uncoiled aorta","unrolled aorta"] -
         (has_observation)    -> [aorta]
         (has_property) -> [tortuous].

[widening_of_mediastinum] -
         (has_observation)    -> [mediastinum]
         (has_property) -> [widening].

[circumscribed_density:{"opacity","density"}].

% Body Locations
[costophrenic_angles:
    {"costophrenic angle","costophrenic angles","
      costophrenic sulci","costophrenic sulcus"}].

[left_upper_lobe_of_lung:{"left upper lobe of lung","left upper lobe"]-
          (has_location)            -> [lobe_of_lung]
              (has_location_qualifier) -> [left]
              (has_location_qualifier) -> [upper].

% Density Qualifiers
[opaqueness:{"opacity","density","opaque","opaqueness"}].
```

APPENDIX C

Structured Findings in X-ray Report BWH22

```
/* BWH22.01. surgical clips again along right mediastinum and [along] right
hilar region */

[rad_finding:#BWH22.01] -
        (has_observation) -> [surgical_clip]
        (has_location) -> [mediastinum] -
                (has_location_qualifier) -> [right]
                (location_relation) -> [along],
        (has_location) -> [hilum] -
                (has_location_qualifier) -> [right]
```

```
		(has_location_qualifier) -> [region]
		(location_relation) -> [along],
	(has_temporal) -> [again]
	(has_quantity) -> [">1"].

/* BWH22.02. new surgical clips in distribution of circumflex artery */

[rad_finding:#BWH22.02] -
	(has_observation) -> [surgical_clip]
	(has_location) -> [distribution_of_circumflex_artery]
	(has_temporal) -> [new]
	(has_quantity) -> [">1"].

/* BWH22.03. [new] 4 intact sternotomy wires */

[rad_finding:#BWH22.03] -
	(has_observation) -> [sternotomy_wire]
	(has_temporal) -> [new]
	(has_quantity) -> [4]
	(has_property) -> [intact].

/* BWH22.04. persistent increased right paramediastinal opacity */

[rad_finding:#BWH22.04] -
	(has_observation) -> [circumscribed_density]
	(has_location) -> [paramediastinum] -
		(has_location_qualifier) -> [right],
	(has_temporal) -> [persistent]
	(has_degree) -> [more_than_normal_degree].

/* BWH22.05.  (possibly related to) previous radiation therapy */

[rad_finding:#BWH22.05] -
		(has_observation)  -> [radiation_therapy]
		(has_temporal) -> [previous].

/* BWH22.06. new plate like opacities in left mid lung zone and left lower
lung zone */

[rad_finding:#BWH22.06] -
	(has_observation) -> [opacity]
	(has_location) -> [lung] -
		(has_location_qualifier) -> [zone] -
			(has_location_qualifier) -> [left]
			(has_location_qualifier) -> [mid],
		(has_location_qualifier) -> [zone] -
			(has_location_qualifier) -> [left]
			(has_location_qualifier) -> [lower],
	(has_temporal) -> [new]
	(has_property) -> [platelike].

/* BWH22.07. (compatible with) atelectasis (based on finding BWH22.06) */

[rad_finding:#BWH22.07] -
	(has_observation) -> [atelectasis].

/* BWH22.08. some interval improvement in left pleural effusion */

[rad_finding:#BWH22.08] -
	(has_observation) -> [pleural_effusion]
	(has_location_qualifier) -> [left]
	(has_temporal) ->
		[improved] -
		  (has_degree) -> [some]
		  (has_temporal) -> [interval].

/* BWH22.09. slight interval decrease in left pleural effusion. */

[rad_finding:#BWH22.09] -
	(has_observation) -> [pleural_effusion]
	(has_location_qualifier) -> [left]
	(has_temporal) ->
	    [decrease_in] -
		(has_degree)   -> [slight]
		(has_temporal) -> [interval].

/* BWH22.10. left lower lobe atelectasis */

[rad_finding:#BWH22.10] -
	(has_observation) -> [atelectasis]
	(has_location) -> [left_lower_lobe_lung].

/* BWH22.11 (& BWH22.13). post-operative changes */
[rad_finding:#BWH22.11] -
      (has_observation) -> [observation]
      (has_temporal) ->  [post_operative_changes].

/* BWH22.12. (consistent with) coronary artery bypass graft (based on finding
BWH22.11) */

[rad_finding:#BWH22.12] -
	(has_observation) -> [coronary_artery_bypass_graft].

/* BWH22.14. (consistent with) previous lobectomy on the right (based on
finding BWH22.13) */

[rad_finding:#BWH22.14] -
	(has_observation) -> [lobectomy]
	(has_location_qualifier) -> [right]
	(has_temporal) -> [previous].
```