Research Paper ■

Medical-concept Models and Medical Records: An Approach Based on GALEN and PEN&PAD

A. L. RECTOR, MD, PHD, A. J. GLOWINSKI, W. A. NOWLAN, ANGELO ROSSI-MORI

Abstract **Objectives**: To investigate the issues raised in applying a preliminary version of the GALEN compositional concept reference (CORE) model to a series of radiographic reports, and to demonstrate that the same underlying concept model could be used in conjunction with both a detailed, fine-grained model of medical records based on that used in the PEN&PAD project and with other more conventional medical-record models.

Design: Following analysis and representation of concepts from a set of reports, a single report was taken as a "case study." This report was analyzed in detail in its entirety and represented using each of the medical-record models.

Results: The reports were successfully represented within the limits of the study, but a number of significant issues were raised.

Conclusion: The compositional approach plus the PEN&PAD medical-record model allowed detailed information in the radiographic report to be represented, including information about the inferences and the clinical process. The resulting representation was large, and more compact representations may be necessary for some systems. Alternative encapsulations of the information as might be used in such systems were successfully prepared. The compositional approach avoided many issues that often cause controversy in the design of traditional coding and classification systems, but it raised other issues, including the handling of ambiguity and underspecification, linkage to information not explicitly present in the report, and questions concerning the focus of individual concepts. All work is preliminary and definitive conclusions await further studies and systematic evaluation.

■ **J Am Med Informatics Assoc.** 1995;2:19–35.

The GALEN project, whose name stands for Generalized Architecture for Language Encyclopædias and Nomenclatures in Medicine, is funded by the European Union under the initiative for Advanced Informatics in Medicine (AIM).*

The GALEN project is developing the architecture and prototypes for a "terminology server," which is intended as a key component of open and distributed health care information systems.[1] The goals of the terminology server are:

Affiliations of the authors: Medical Informatics Group, Department of Computer Science, University of Manchester, Manchester, United Kingdom (ALR, AJG, WAN), and Consiglio Nazionale delle Ricerche, Italy (AR-M).

Supported in part by grants from the UK Department of Health, grant number 8800091 from the Medical Research Council, and Project A2012 (GALEN) under the EC initiative for Advanced Informatics in Medicine (AIM).

Correspondence and reprints: A. L. Rector, MD, PhD, Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom. e-mail: rector@cs.man.ac.uk

Received for publication: 11/16/93; accepted for publication: 8/15/94.

*Members of the GALEN consortium: University of Manchester (Coordinator), Hôpital Cantonal Universitaire de Genève, Consiglio Nazionale delle Richerche (Italy), Conser Systemi Avanzati (Italy), The Association of Finnish Local Authorities, The Finnish Technical Research Centre, GSF-Medis Institut, (Germany), Hewlett-Packard Healthcare Products, Katholieke Universiteit Nijmegen (Holland), University of Linköping (Sweden), and University of Liverpool.

1. To provide a flexible, extensible basis for coherent exchange of clinical information among electronic medical-record systems, decision support systems, clinical interfaces, information-retrieval systems, and other systems using clinical information without imposing a single rigid coding system.
2. To serve as an accessible repository of language-independent medical-concept representations, and to map this repository to different natural languages.
3. To provide infrastructure support for the development of new clinical information systems.
4. To convert between existing coding schemes and between existing representations.
5. To provide dynamically generated local nomenclatures or "coding schemes" that are more comprehensive and thoroughly organized than are those that can be maintained as static structures or managed manually.

The heart of the project is the development of a concept reference (CORE) model of medical concepts, which serves as the "interlingua"[2] in which the concepts used in the medical records or referred to by other coding systems and nomenclatures are represented. The goal is that the CORE model is to be independent of local language and particular applications so as to provide a reusable foundation for large-scale development across Europe. The goal is both to facilitate new developments and to enhance existing systems by unpacking the information that is encapsulated in their coding systems so that it can be viewed and recombined in different ways for different purposes.

GALEN in one response to the need felt by many and expressed recently in this journal by the Canon Group for a formal "medical-concept representation language"—a logical, well-defined system for manipulating concepts.[3] Like the Canon group, GALEN sees the primary task as being the manipulation of such concepts formally on computer. It sees the need for formal concept manipulation as essential to both system integration (systems must be able to interpret each other's information) and human–computer interaction (systems dealing with clinical content need to behave in ways that approximate the use of clinical concepts by clinicians.[4,5] In achieving the required level of formal structure, GALEN aims to complement systems for direct mapping between coding schemes such as the Unified Medical Language System (UMLS) metathesaurus[6] and the semantic network[7] by making it possible to represent much more information about the concepts represented, and about the constraints on the use of those concepts, than has been possible previously.

A key objective of the GALEN project is to demonstrate that such a model can be built, scaled up to realistic size, and made an effective tool for building and integrating clinical systems. The background of the approach lies in the experience of the PEN&PAD project[8,9] and the work on predictive data entry and medical records that underpins it.[10,11] The PEN&PAD project uses a fine-grained model of medical records based on semantic networks. The GALEN CORE model is intended to support such models directly. However, it is also intended to be able to support other medical-record models representing less detail in structural form. The ability to use the same concept model to support different medical-record models is one aspect of making the concept model reusable and application-independent.

This article is based on experiments using an early version of the GALEN CORE model in conjunction with the material compiled for a workshop held in January 1993, by the Canon Group.[12] The workshop took place during the first year of GALEN and the material presented should be regarded as a case study. The outcome is a series of issues and observations rather than definitive conclusions or a systematic evaluation of a finished product.

The primary goal of these experiments was to demonstrate how the full detail of the radiographic reports might be represented using the PEN&PAD medical-record model using concepts in the style of the GALEN CORE model as it then existed. Our goal was to represent reports in their entirety, including both the purely clinical concepts (the concepts describing the investigation process) and the concepts whose existence was alluded to in the report but which were not described explicitly (e.g., "previous radiotherapy"). This article presents one such report. The choice of attempting to represent entire reports rather than simply extracting concepts from reports was made because extracting concepts makes it too easy to gloss over difficulties and to bias the results toward what is easy to do in a given representation. At this stage, we believe the task is to represent everything or to give principled reasons why some things cannot, or should not, be represented.

In addition, a brief experiment was conducted to demonstrate how the same concept model might be reused to support other medical-record models that represent less detail in structured form. The central structures of two alternative medical-record models were abstracted from existing systems: a "simple code-based model" consisting of only dates and codes and a "linked relational model" that allows tests and values plus some linking of statements to indicate causal relations and uncertainty.

Background

Goals and Antecedents of the GRAIL Kernel

The GALEN CORE model of medical concepts is based on a theory about the generation and subsumption of composite concepts—a theory of how to build models that can generate, classify, and verify entities representing "all and only" sensible medical concepts. This theory is embodied in the GALEN Representation and Integration Language (GRAIL) Kernel. It is based on previous work on the Structured Meta Knowledge (SMK) formalism used in the original PEN&PAD project.[13–15]†

Concepts are represented in GRAIL by entries in a semantic network. The GRAIL Kernel provides for five primary functions on these entities:

1. Sanctioning of the composition of entities to form more composite entities (e.g., that from "pneumonia" and "severe" may be composed "severe pneumonia");
2. Classification of those composite entities (e.g., that "severe pneumonia" is a kind of "pneumonia");
3. Transformation of those composite entities to a canonical form (e.g., that "pneumonia of the lungs" is just "pneumonia");
4. Tests to determine whether a proposed composite entity is coherent with the model (e.g., that a "mild severe pneumonia" or a "pneumonia of the feet" is incoherent);
5. Tests to determine whether the model is coherent globally (i.e., to show that there is no implied contradiction or ambiguity).

The GRAIL Kernel follows a tradition of work on semantic networks and "description logics" first made fully explicit by Brachman et al.[17] Recent work in this tradition and the problems of the computational tractability of logic and semantic networks have recently been reviewed by MacGregor.[18] As do most such systems, GRAIL seeks to find an appropriate trade-off between logic, which is highly expressive but computationally intractable, and restricted terminologic languages, which are computationally tractable but inexpressive. Previous studies had indicated that one such restricted description logic was not sufficiently expressive to serve as a basis for medical-concept modeling.[19,20] The GRAIL Kernel's facilities were selected to circumvent the difficulties found in these studies and also to provide a principled approach to problems of particular importance to medicine, e.g., dealing with part–whole relationships and their interactions with the kind-of taxonomy.

More recently, several authors have advocated the use of Sowa's conceptual graphs[21] as a notation for concept representation.[22–24] Conceptual graphs are an alternative notation for logic that is particularly convenient for work on concept structures. In principle, they can express anything that can be expressed in any other form of logic, including the theory underlying GRAIL. Superficially, the notation of the GRAIL Kernel translates easily to that of conceptual graphs. To facilitate collaborations, work is in progress to extend such translations, and Bernauer and Goldberg[22] have demonstrated a theory of the relationship between part–whole and kind-of relations expressed in conceptual graphs that is closely related to that of GRAIL's. However, the GRAIL Kernel embodies a particular theory of terminology and terminologic reasoning—of how to describe what can sensibly be said and of how to classify the concepts expressed. Any translation will therefore be to a restricted subset of conceptual graphs, corresponding to the restrictions in the GRAIL theory, plus a set of definitions of GRAIL's operations in terms of standard operations on conceptual graphs.

A compiler for the GRAIL Kernel, an engine to manage the operations on GRAIL models, and associated knowledge-engineering tools have been implemented and were used as the basis for the work reported here.

Brief Overview of the GRAIL Kernel

A model in the GRAIL Kernel consists of a network of nodes, called "entities," and arcs, called "statements." Each statement consists of two entities linked by an "attribute." The GRAIL Kernel provides the rules for combining existing entities into new entities based on the sanctions expressed by the statements in the models. For example, from the entities *Effusion* and *PleuralSpace* and the attribute *hasLocation* plus the appropriate sanctions in the model, we can form the composite entity

Effusion which *hasLocation-PleuralSpace*.

For convenience, we can name composite entities to define new entities. For example, we might say

†Further details concerning the GRAIL Kernel can be found in references 4, 5, and 16.

(*Effusion* which *hasLocation-PleuralSpace*) name
'PleuralEffusion,'

which is roughly equivalent to the English sentence "Pleural effusion is defined as an effusion which is located in the pleural space." GRAIL automatically classifies composite entities in terms of all of the other entities that exist in the model (e.g., "effusions" and "things located in the pleural space") and checks the entities for self-consistency and conformance with the sanctions in the model.

In the above examples, "which," "name," and "-" are keywords and tokens from the GRAIL language;

Entity *e.g.* Lung IntrinsicallyPathologicalStructure	An elementary entity, *e.g. Lung, Lobe,* etc. By convention, primary concepts are represented by words beginning with upper case letters. Names made up of several follow the Smalltalk convention in which each word is capitalised but no spaces or underline characters ("_") are used,
\Entity *e.g.* \JohnSmith	An individual entity defined outside the primary knowledge base, *e.g.* specific patients, doctors, dates, clinics, etc.
attribute/inverseAttribute *e.g.* hasLocation/isLocationOf	An attribute and its inverse. By convention attributes begin with a lower case letter and are always verbs, *e.g. hasSeverity, hasSize, hasLocation,* etc.
attribute-Value *e.g.* hasSeverity-severe	A 'criterion'
Entity *which* attribute-Value *e.g.* Tumour *which* hasLocation-Lung.	A 'particularisation' by a single criterion,
Entity *which* < attribute-Value1 ... attribute-ValueN> *e.g.* Opacity *which* < hasForm-plateLike, hasProgress-new>	A 'particularisation' by a set of criteria. Order of the criteria is unimportant. Angle brackets <...> always indicate a criteria set.
Expression *name* 'Name'. *e.g.* (Effusion *which* hasLocation PleuralSpace) *name* 'PleuralEffusion'.	An expression which names a particularisation - *i.e.* a definition. Such expressions make the model easier to read and manage but do not change its overall behaviour. The keyword *name* may also be inserted at any point in a medical record statement to name the preceding entity without otherwise affecting the meaning of the statement.
'ACompositeEntity' *e.g.* 'RightLung', 'LeftPleuralSpace'.	A 'named particularisation', i.e. a named composite entity. Named particularisations are enclosed in single quotes.
Entity-Attribute-Entity *e.g.* RadiographicObservation22 - shows - (presence *which* isStateOf-'SurgicalClip')	A statement in the medical record.
[... , ... , ...] *e.g.* RadiographicObservation22' - hasDiagnosis -[presence *which* isStateOf- 'LeftPleuralEffusion', presence *which* isStateOf- 'LeftLowerLobeAtelectasis', presence *which* isStateOf- 'PostOperativeChange']	An expandable list. *i.e.* each item in a list in square brackets is applied to the rest of the expression, *e.g.* *Object-Attribute-[Value1 Value2 Value3] i* is equivalent to: *Object-Attribute-Value1.* *Object-Attribute-Value2.* *Object-Attribute-Value3.* Or equivalently to a branching network , *e.g.* *Object- Attribute - Value1* | *-Value2* | *-Value3*

Figure 1 Informal presentation of the notation used by the GRAIL Kernel.

all other words are entities and attributes from the CORE model. Much of the style of the CORE model is set by a relatively small number of abstract entities and attributes that together are often referred to as the "high-level schemata" or "ontology." The high-level schemata are specific to a domain or a group of related domains and are not fixed in the GRAIL language nor even dependent on it. Models using similar high-level schemata might be built using other representations such as CLASSIC or conceptual graphs. Conversely, different high-level schemata could be used in GRAIL models.

Key issues are discussed below. Figure 1 summarizes the notation used. In this article, GRAIL entities are shown in italic, natural-language equivalents of GRAIL entities are in double quotation marks, and newly introduced technical terms from GRAIL or other formalisms are in single quotation marks. Single quotation marks ('...') are used to enclose the defined names of composite GRAIL entities.

Using Composition to Coordinate Taxonomies

GALEN aims to produce models that are reusable and application-independent, i.e., models that can be developed in the course of building one application but reused in new applications or that can be developed for a group of loosely specified applications. Eventually the goal is to develop common models used by a wide range of applications.

A major means to achieving reusability and application independence is to keep different taxonomies cleanly separated. Most systems mix different taxonomies in a single hierarchy—in some places breaking things down by physiologic function, in others by means of observation or measurement (e.g., signs or symptoms), and in still others by role in decision making (e.g., indications, problems, and prognoses). Any argument of the form "Is that a sign or a symptom?" or "Is that a complaint or a diagnosis?" is taken as a clear indication of conflation of taxonomies.

For example, GALEN's CORE model defines "pleural effusion" as "fluid in the pleural space." The "presence" of a "pleural effusion" is pathologic and can be observed by either radiography or percussion. The fluid itself can also be removed, but it then becomes a separate entity to be examined by other means. Therefore, any assignment for one application of "pleural effusion" to be a sign, symptom, or potential substance for laboratory analysis will be inadequate for other applications. GALEN approaches the problem by separating the taxonomy of observation tasks from the taxonomy of pathophysiology, which includes pleural effusion. Composition can then be used to form any "sensible" expression incorporating both the pathophysiology and the mode of observation.

For example, as expressed in the GRAIL Kernel, we can define both

presence which< *isStateOf-'PleuralEffusion'*
isObservedOn-'ChestAPRadiograph'>
presence which< *isStateOf-'PleuralEffusion'*
isObservedOn-Percussion>.

. . .

The task of the GRAIL engine is to classify expressions such as the above automatically. Consider:

(*State* which *isObservedOn-Radiograph*) name *'RadiographicSign'*
(*State* which *isObservedOn-PhysicalExamination*) name *'PhysicalSign.'*

Given these definitions, the GRAIL engine will classify each of the examples appropriately—the first as a *'RadiographicSign'* and the second as a *'PhysicalSign.'* Put another way, the issue in GRAIL is not whether "pleural effusion" is a kind of "radiographic sign" or a kind of "physical sign," but whether a radiographic or physical sign can be sensibly composed based on "pleural effusion."‡

Controlling Composition: Constraints and Sanctions

As stated earlier in this article, a GRAIL model consists of a primary hierarchy of elementary entities and a set of sanctioning statements connecting these entities. The sanctioning statements express the constraints on what composite entities can be formed. Composite entities can themselves be the topic of further sanctioning statements that, in turn, constrain the formation of more complex composite entities.

Sanctioning statements are made at three levels: the grammar level, the sense level, and the necessity level. Statements at each level must be sanctioned by statements at the next higher level. Roughly speaking, the grammar-level statements sanction queries and operations by the knowledge engineer. They correspond most closely to the type constraint found in other languages. Sense-level statements sanction the generation of new particularizations representing sensible medical concepts. Necessity-level statements prevent the generation of entities representing redundant concepts (e.g., "the hand which is a part of the arm").

‡Further discussion of the problems of separating taxonomies using examples from SNOMED-II can be found in reference 14.

The use of sanctions and constraints is illustrated by the model for lung segments. At the top level is the grammar-level statement, which indicates that it is reasonable to describe lobes of the lung as having positions "upper," "middle," and "lower."

(*Lobe* which *isSolidDivisionOf-Lung*)
grammatically *hasUpperLowerPosition* -[*upper middle lower*].

This grammar-level statement sanctions the sense-level statements, which describe which positions for left and right lobes. Consider first the upper and lower lobes that occur normally in both lungs:

(*Lobe* which *isSolidDivisionOf-Lung*)
sensibly *hasUpperLowerPosition*-[*upper lower*].

These sense-level statements sanction the generation of particularizations such as that representing "the right upper lobe of the lung":

(*Lobe* which <*isSolidDivisionOf-'RightLung'*
hasUpperLowerPosition-upper>) name *'RightUpperLobe'*.

Sense-level statements state that to generate the corresponding *particularization* is sensible. In addition, the sense-level statements sanction the necessity-level statements, which say that, conceptually, all lungs have upper and lower lobes.

'Lung' necessarily *hasSolidDivision-*
(*Lobe* which *hasUpperLowerPosition-*[*upper lower*])

These statements also indicate that, for example, the "right lung which has an upper lobe" is no different from the "right lung," i.e., that the "canonical form" of

(*Lung* which *hasLaterality-right*) which
hasSolidDivision-(*Lobe* which *hasUpperLowerPosition-upper*)

is just

Lung which *hasLaterality-right*.

Note that the sense of the attribute *hasSolidDivision* here is taken as meaning that such divisions are an essential part of our concept of lung. GRAIL models are concerned with what it is sensible to say. The surgical removal of a patient's right upper lobe does not make it any less sensible for us to discuss that patient's right upper lobe (e.g., to say that it is absent).

The middle lobe presents difficulties, since it is normally present only in the right lung but can be present in the left lung in certain congenital abnormalities. Whether this degree of detail is required depends on the purpose of the model. If this detail is required, then current practice is to construct a model equivalent to the statements: "Lungs sensibly have middle lobes"; "Normal right lungs necessarily have a middle lobe"; and "Lungs with a left middle lobe are necessarily non-normal." Applications that want only normal anatomy can request the anatomy for the "normal lung," which will exclude the "left middle lobe." (In the CORE model, a distinction is made between *nonNormal* and *pathological*. The value *nonNormal* indicates merely that something is not normally present; the value *pathological* indicates a disease process.)

Medical Records and Medical Concepts—Alternative Encapsulations

GALEN's view of the interrelationship between medical-concept models and typical information models is that the medical concepts are the chunks that fit in the boxes in the more general information models. If we assume that there is a deep representation in the form of something like a semantic network or conceptual graph, then there are clearly many ways of breaking up the underlying representation into chunks—many possible encapsulations. We can say that different potential combinations of medical-record model and concept encapsulation system are coherent if they expand to the same underlying representation, allowing for possible omissions of information or elisions.

Furthermore, any one underlying record may be encapsulated in different ways as different concepts. Consider the statement: "A pleural effusion was present on the AP View" or "The AP View showed (the presence of) a pleural effusion." The underlying information is most easily visualized as a graph or semantic network.

'ChestAPRadiograph' -shows/isObservedOn - presence -
isStateOf/hasMedicalState - 'PleuralEffusion.'

The network can be read either left to right using the primary attributes, e.g., *shows*, or right to left using their inverses, e.g., *isObservedOn*. This network implies the existence of a composite concept corresponding to each of its three component concepts, that is:

'ChestAPRadiograph' which
shows-(presence which isStateOf -'PleuralEffusion')

presence which <
isStateOf-'PleuralEffusion'
isObservedOn - 'ChestAPRadiograph'>

'PleuralEffusion' which
hasMedicalState-(presence which
isObservedOn-'ChestAPRadiograph').

Informally, we may think of picking the network up by any of its concept nodes and dangling the rest of the network from that point. Conceptual graphs[21] provide an analogous facility through the notion of the alternative "heads" of a network (alternative lambda abstractions from a given graph).

PEN&PAD Framework for Models of the Medical Records

The medical-record representations are based on the PEN&PAD approach to medical records.[10,11] Two of the important principles are:

1. That the record consists of two layers:

   Layer 1: Statements of observations and deductions pertaining directly to the patient or object observed.

   Layer 2: Statements linking the statements in layer 1 to each other and to other statements in layer 2 (e.g., statements linking diagnoses to evidence, statements ascribing clinical significance to events, and statements linking statements to problems).

2. That the record is a deletionless logbook of statements made at a particular time and place by an "agent"—usually, but not necessarily, a human clinician.

The PEN&PAD model was one of the major influences on the development of the GRAIL Kernel and was expressed entirely in the SMK notation from which GRAIL evolved.

In the PEN&PAD model there is an intimate relationship between concepts and the medical record. The record is seen as a representation of the "occurrences" of the "individuations" of the "categories" in the concept model (e.g., *'ObservationOfJohnSmithAtClinicXByDrJones'* is an occurrence of the individual *JohnSmith*, which is an individuation of the category *Patient*).

Other Medical-record Models Used

Most other medical-record models have a less intimate relation to the concept model. The models of the concept model and the medical record are separate. Concepts from the concept model (or coding system) are data in the medical-record model. Portions of two such medical-record models were used to demonstrate potential alternative ways of encapsulating the concepts represented and using them in medical records. The goal was to show alternative uses of the concept model rather than to compare alternative models for the records. Hence, simple models were chosen and, in both cases, only the central tables concerning diagnoses and findings were considered.

The simple code-based model (Fig. 5) was abstracted from widely used systems that record one or more International Classification of Diseases (ICD), Read, or similar codes against patients and dates. Some versions of such schemes also include the doctor and other details in the same table. Others normalize the information into separate tables keyed on the patient identifier and date. Some versions also include indicators for the role as in the linked relational model.

The linked relational model (Fig. 6) was abstracted from formats commonly found in systems derived from laboratory information systems based on tests codes and results. It differs from the simple code-based model in the inclusion of a separate column for "results" and the availability of an identifier for each row or "tuple." The identifier for each tuple allow statements to be made concerning the fact expressed in that tuple. Since tuple identifiers may appear in the "subject" column, this model allows statements about statements (i.e., meta-statements) to be represented, although the multiple self-joins or nested queries required to retrieve such chains of statements from relational databases are computationally inefficient. As in many such systems, an additional column for a "role" is included, which is used to divide the information into findings, diagnosis, and physical examination. This specific model was abstracted from that used in the Nottingham General Practice System.[25]

Methodology

The experiments represented here were based on an early version of the GALEN CORE model (version 1.5). At the beginning of the experiment the CORE model contained a model of the anatomy and basic pathology of the lung, but no concepts specific to radiography or radiographic findings. Although closely based on the PEN&PAD work, the GALEN CORE model has not yet taken on the medical-record constructs. The medical-record constructs were therefore taken from the PEN&PAD medical-record model itself, with minor adaptations for consistency with

PA view is compared to the previous examination dated 10-22-91. Surgical clips are again seen along the right mediastinum and right hilar region. There are new surgical clips in the distribution of the circumflex artery as well as 4 intact sternotomy wires. There is persistent increased right paramediastinal opacity, possibly related to previous radiation therapy. New plate like opacities are seen in the left and mid lower lung zones, compatible with atelectasis. There has been some interval improvement in the left pleural effusion.
1. Slight interval decrease in left pleural effusion.
2. Left lower lobe atelectasis.
3. Post-operative changes consistent with coronary artery bypass graft as well as previous lobectomy on the right.

Figure 2 The radiographic report used for all results reported in this article.

GALEN's current usage. The radiographic report that was used for all results reported here is shown in Figure 2.

Concepts to be modeled were identified from two sources. The first was the first two layers of the automatic analysis using CLARIT[26] performed on a large corpus of reports pooled from many centers. The CLARIT analysis produces lists of candidate concepts in order of their "importance" to the corpus analyzed. Typical concepts that appeared were 'right_pleural_effusion' or 'blunted_costophrenic_angle.' The second source was manual extraction of any remaining concepts not already found in the CLARIT analysis from seven of the 18 reports used by members of the CANON Group for model development. This article focuses in detail on a single report, which is presented in full. At this early stage in the development of the field, we believe this to be the most appropriate way to present results. However, we have noted in detail the origin of each of the concepts used that did not come from the GALEN CORE model as it then existed.

The concepts identified were compared manually with the GALEN CORE model. The concepts identified were divided into three categories: concepts already representable within the GALEN CORE model, new candidate primitive concepts, and new candidate composite concepts. The decision as to whether a given concept should be represented as a primitive or composite was made subjectively using the same policies as in other parts of the GALEN model. Similarly, additional attributes and value types were defined where required to form the candidate composite concepts along with the sanctions and constraints required for a plausible account of the data available. The existing GALEN CORE model was then extended manually to include entities representing these constructs.

To demonstrate that the model was complete with respect to the concepts identified, the composite entity for each concept identified from the sources was defined and named. The resulting file was compiled using the GRAIL engine to demonstrate that the model was formally coherent and self-consistent, at least insofar as the inbuilt checks in the engine can determine. Defining named entities representing every concept encountered was a methodologic tool for the Canon workshop but was not standard practice for GALEN where many entities such as *'LeftPleuralSpace'* would be generated as needed rather than stored explicitly.

For the experiments in using the concept model to represent entire radiographic reports as they would appear in the medical record, the reports were examined for the external "individuals" implied—the doctors, the patients, and the radiographic films. The attributes in the original PEN&PAD model were adapted to link these expressions with the concepts identified. In addition, the model was extended with additional new attributes to deal with uncertainty and time, which were not yet included in the GALEN CORE model and had been dealt with differently in PEN &PAD.

To represent the linkages to previous examinations and operations implied but not explicitly present in the report, special constants for all such observations were created. These constants were created as primitives in GRAIL but have been shown in the text using the pseudo-operator *a/an*.

All of the entities required for the medical record were again compiled using the current GRAIL compiler to check their accuracy. The PEN&PAD model was slightly adapted to account for the differences between the concept models underlying PEN&PAD and the GALEN CORE model and for the changes during the development from SMK formalism to GRAIL. The record was then formulated as shown in Figure 4 using this model.

Given the basic structures of the two other medical-record models, the additional composite concepts necessary for the simple code-based and the linked relational models were defined and named and the specific examples constructed manually.

```
MedicalConcept
  AnatomicalConcept
    Structure
      BodyPart
        SurfaceBodyStructure
          Chest[1]
        Organ
          Lung
          Hilum

  RightCircumflexArtery*
        GenericBodyPart
          Lobe
          Region
          Zone*
          Membrane
            Pleura
      Cavity
        AnatomicalSpace
          Mediastinum
          PotentialSpace
      AnatomicalFeature
        CostophrenicAngle*
      Accumulation
        Effusion*

  IntrinsicallyPathologicalStructure
        Atelectasis*
    ExtrinsicStructure
      SurgicalImplant
        SurgicalClip*
        SternotomyWire*
```

```
Substance
    Fluid
  Procedure
    Excision
    RadiationTherapy*
    RadiographicProcedure*
      PlainXRay*
  State
    PresenceOrAbsence
      [presence absence]
    IncreaseOrDecreaseInSize
      [increaseInSize noChangeInSize ...]
  Status
    PathologicalOrPhysiological
      [pathological physiological]
    NormalOrNonNormal
      [normal nonNormal]
IndependentObject
  Agent§
    Person§
      Clinician§
      Patient§
  ObservableObject§
    Person^
  DiagnosticSampleOrObject§
    RadiographicFilm*
```

Composite Entity	Definition
Lungs and Pleura	
'LeftLung'	Lung *which* hasLaterality-left
'PleuralSpace'	Space *which* isSpaceFormedBy-Pleura
'LeftPleuralSpace'	Space *which* isSpaceFormedBy- (Pleura *which* isLayerOf -'LeftLung')
Pleural Effusion	
'Effusion'	Accumulation *which* isMadeOf-Fluid .
'LeftPleuralEffusion'	'Effusion' *which* isContainedIn-'LeftPleuralSpace'
Lung Zones	
'LungZone'	Zone *which* < isSurfaceAnatomicalDivisionOf-Chest isAnatomicallyRelatedTo-Lung>
'RightParaMediastinal-Region'	Region *which* < isSolidDivisionOf-Chest isNear-Mediastinum hasLaterality-right>
'LowerLungZone'.	'LungZone' *which* hasUpperLowerPosition-lower
Opacities and Atelectasis	
'PlateLikeOpacity'	Opacity *which* hasForm-plateLike
'LeftLowerLobe-Atelectasis'.	Atelectasis *which* hasLocation -(Lobe *which* < isSolidDivisionOf-'LeftLung' hasUpperLowerPosition-lower>)
Radiographs	
'ChestRadiograph'	PlainXRay *which* isPerformedOn-Chest
'PAViewChest'	'ChestRadiograph' *which* hasRadiographicPosition-posteriorAnterior

Figure 3 (**Top**) Major elementary entities in the model. The * indicates that the entity was created for the Canon workshop exercise and was not previously in the CORE model as of July 1993. The § indicates concepts borrowed from the PEN&PAD model. Lists in square brackets [. . .] are terminal leaves printed linearly to save space. A number of new concepts such as effusion have been treated as elementary for this exercise, although they may be defined as composite terms eventually, e.g., as something roughly equivalent to "an accumulation of fluid." The hierarchy allows multiple parents, although only the most important duplications are shown. Second appearances of a concept are indicated by the symbol ^ and the subtaxonomy underneath the concept is omitted on the second and subsequent appearances. In the GALEN CORE model, the term "Thorax" is used rather than 'Chest' to conform with SNOMED's usage. (**Bottom**) Sample named composite entities defined from the model for the radiographic report for the Canon workshop exercise.

Results

Initial Analysis of Concepts Used

The concepts and model are shown in Figure 3. Figure 3 (top) shows the primitive concepts used. Figure 3 (bottom) shows a sample of the definitions of composite concepts used in the radiographic report represented.§

Figures concerning the number of entities in GRAIL models are often misleading, since most of the concepts implied are not defined explicitly. A methodology is being developed to describe GRAIL models in terms of coverage rather than size. However, to provide a rough indication, at the time of the experiment the GALEN CORE model contained approximately 1,200 primitive concepts and a further 2,500 explicitly named concepts. The PEN&PAD model contained approximately 2,500 primitive concepts of which 1,200 were drugs or drug-related. Each model used on the order of 50 attributes. In the course of the exercise, 65 new entities including attributes and values were added plus 16 new individuals for the specific medical record.

Examination of the models for the attributes added is more informative than raw numbers. Asterisks indicate the new elementary categories and attributes added; the § sign indicates an attribute adapted from PEN&PAD but not present in the GALEN CORE model. All of the anatomic concepts and connectives were present in the CORE model, but, as expected, the specific radiographic findings were not. The concepts concerning clinical care, such as *Agent* and *Clinician*, were borrowed from the PEN&PAD model. Terms such as *Zone* have been treated as new terms distinct from the CORE model's *AnatomicalRegion*, since the impression is that "Zone of the Lung on Radiograph" is used differently from "Anatomical region of the lung." If further experience fails to confirm this difference in usage, the two terms will be coalesced. The constructs for uncertainty were developed for this exercise. They have any direct counterpart in the original PEN&PAD schema, which confined itself to the inbuilt qualifiers *true*, *query*, and *false*.

Three Representations of the Radiographic Report

Adapted PEN&PAD network model. The representation of the radiographic report using the adaptation of the PEN&PAD network model is shown in Figure 4. The representation follows the framework previously outlined in the literature using the more recent syntax used in the GRAIL Kernel.[10,11] The first layer, as described earlier in this article, uses the attribute *shows*. The second layer uses three further attributes as shown below:

Layer 1:	Observations and deductions pertaining directly to the patient or object observed.	
	a) The direct observations of the radiographic film	*shows* (applied to radiograph)
	b) Diagnoses/findings accepted as pertaining to the patient and acknowledged by the patient's own physician	*shows* (applied to patient)
Layer 2:	Statements linking the statements in layer 1 to each other and to other statements in layer 2.	
	c) Implications needed to formulate the remainder but not directly mentioned as observations	*implies*
	d) Interpretations made by the radiologist as to the underlying phenomena causing the direct observations	*hasInterpretation*
	e) Diagnoses assigned by the radiologist to the patient on the basis of the above	*hasDiagnosis*

Additional externally defined individuals, which are taken as primitive for the purposes of this record fragment are written in the form *EntityName*. (Some of the times, places, or agents that are represented here as distinct might be found to be identical when referred back to the record as a whole.)

Every statement has as its topic an "observation," which encapsulates the time, the place, and the person making the observation. For example, *'PatientXObs3'* is defined by:

```
\ PtXsOwnPhysician which <
  isAttendingAt-\ Place5
  isAt\ Time5 >
    name 'AgentSession3'

(\ PtX which isObservedByAt 'AgentSession3')
  name 'PatientXObs3'.
```

§The schema used in describing the anatomy is discussed in more detail in reference 27.

```
RadiographicObservation22' -
  [
          shows-[   presence which isStateOf-'SurgicalClip' which<
                        hasLocation-'RightMediastinalRegion'
                        hasProgress-persistent
                        hasNumber-multiple>,
                    presence which isStateOf-'SurgicalClip' which <
                        hasLocation-'RightHilarRegion'
                        hasProgress-persistent>,
                    presence which isStateOf-'SurgicalClip' which <
                        hasLocation - (Region which isServedBy-'RightCircumflexArtery')
                        hasProgress-new>,
                    presence which isStateOf-SternotomyWire which <
                        hasDamage-intact
                        hasNumber-4>,
                    (presence which isStateOf-Opacity which
                        hasLocation-'RightParaMediastinum') name 'Opacity1',
                    (presence which isStateOf-Opacity which <
                        hasForm-plateLike
                        hasLocation-'LeftLowerLungZone'
                        hasProgress-new>) name 'Opacity2'>,
                    (presence which isStateOf-Opacity which <
                        hasForm-plateLike
                        hasLocation -'MidLowerLungZone'
                        hasProgress-new>) name 'Opacity3'>,
                    presence which isStateOf-'LeftPleuralEffusion' which
                        hasProgress-(improved which
                                isRelativeTo-'PreviousRadiographicFilm')
                 ],
        implies - [ (a performance which isStateOf-
                        (RadiationTherapy which <
                           isTreatmentOf-PtX
                           isAtTime-(a Time which isPreviousTo-\Time1 )>))
                                                   name 'RadiationTreatment1',
                    (a performance which
                        isStateOf- ('CoronaryArteryBypassGraft' which <
                           isTreatmentOf-PtX
                           isAtTime-a Time which isPreviousTo-\Time1 )>)) name 'Op1',
                    (a performance which isStateOf-
                        (Excision which <
                           isPerformedOn-(Lobe which isSolidDivisionOf-'RightLung')
                           isTreatmentOf-PtX
                 ],        isAtTime-(a Time which isPreviousTo-\Time1)>)) name 'Op2'

  hasInterpretation-[ ('Opacity1' - isConsequenceOf - 'RadiationTreatment1'),
                    ('Opacity2' - isConsequenceOf - Atelectasis) which hasCertainty-compatibleWith,
                    ('Opacity3' - relatedTo - Atelectasis) which hasCertainty-compatibleWith
                 ],
     hasDiagnoses[  presence which isStateOf- (LeftPleuralEffusion' which
                        hasProgress-(decreased which isRelativeTo-'PvsXRay')),
                    presence which isStateOf-'LeftLowerLobeAtelectasis',
                    presence which isStateOf- ('PostOperativeChange' which <
                        (isConsequenceOf-'Op1') which  hasCertainty-consistentWith
                        (isConsequenceOf-'Op2') which  hasCertainty-consistentWith>)
                 ]
 ].
'PatientXObs3'-
          shows-[   presence which isStateOf-'LeftPleuralEffusion' which <
                        hasProgress-(decreased which isRelativeTo-'PvsXRay')
                        isReportedOn-'RadiographicObservation22'>,
                    presence which isStateOf-('LeftLowerLobeAtelectasis' which
                        isReportedOn-'RadiographicObservation22')
                 ]
```

Figure 4 Representation of the primary observations of the radiographic report using the adaptation of the PEN&PAD network model. The interpretation of the attributes is described in the Results section.

Code	Expression for Code
code_1110	decrease *which* isStateOf ('PleuralEffusion' *which* hasLaterality-left)
code_1120	presence *which* isStateOf (atelectasis *which* hasLocation 'LeftLowerLobe')
code_1140	presence *which* isStateOf-(changes *which* isConsequenceOf-SurgicalProcedure)
code_1141	presence *which* isStateOf- (changes *which* isConsequenceOf-'CoronaryArterBypassGraft'
code_1142	presence *which* isStateOf (changes *which* isConsequenceOf-Excision of (Lobe *which* isSolidDivisionOf-(Lung *which* hasLaterality-right))
code_1142.1	presence *which* isStateOf (changes *which* <isConsequenceOf 'CoronaryArteryBypassGraft' isConsequenceOf (Excision of (Lobe *which* isSolidDivisionOf (Lung *which* hasLaterality-right))>

patient	Date	Code
pid_12345	10-28-91	code_1110
pid_12345	10-28-91	code_1120
pid_12345	10-28-91	code_1142.1

Figure 5 (Top) Codes and expressions for a simple code-based medical record. (**Bottom**) Simple code-based model of a medical record consisting of patient identifier, date, and code. As in many simple medical record or diagnostic register systems, only the central table is shown.

The representation distinguishes clearly between the "radiographic film," the "act of radiographic examination" that produces that film, and the "interpretation" of that film by a radiologist.

Simple code-based model. Figure 5 (top) shows the encapsulation of GRAIL entities to form "codes."¶ Figure 5 (bottom) shows the use of these codes in a simple three-column model of the medical record as a series of patient identifiers, dates, and codes—fully encapsulated nominalizations of the underlying network. Only the conclusions have been encoded, since few such systems attempt to record the findings separately. Only the central table concerning diagnoses and findings is shown.

Linked relational model. The central table from the linked relational model for the medical record is shown in Figure 6. Each entity in Figure 6 represents a single code—a fixed-length identifier corresponding to the named entity. (The table could have been presented split into two as in Figure 5, but the result would have been difficult to read.)

¶These codes are natural encapsulations of the model developed here used to illustrate the basic principles. The full complexity of coping with the idiosyncracies of long-established coding systems such as ICD9 is beyond the scope of this article.

Issues Arising in the Medical-record Experiment

Quantity of Information

The most striking thing about the complete representation of the radiographic report in Figure 4 is the amount of detail required to capture the complete meaning of even a brief report. Medical records are succinct, information-rich documents. There may be arguments over details, but we believe the order of magnitude of information in Figure 4 is what will be needed.

Much of the information concerns the process of investigation and care—"clinic visits," "radiographic films," "interpretation," and "uncertainty." In fact, only about half of the statements in the representation of the report using the PEN&PAD model represent the basic observations themselves. The rest are "second layer" statements, or meta-statements, that describe the use of those observations and the inferences made from them. This is consistent with experience with the PEN&PAD system itself where the direct patient observations constituted considerably less than half of the total record. The ability to capture this detail is one of the great strengths of the

PEN&PAD model. As van der Lei[28] has pointed out, information about how and why information was obtained and used is essential to the effective use of the medical record as a basis for decision support or detailed clinical audit.

However, this degree of detail may be inappropriate, uneconomical, or technically unfeasible for some applications. Successful dissemination of detailed concept models will depend on their being usable in relatively simple systems as well (hence, the experiments reported here with two alternative medical-record formats).

Presence and Absence and the Question of Focus

Most pathology is the presence of an abnormal structure, but it is also necessary to be able to represent the absence of normal structures. The report given did not explicitly report the "absence of a lobe of the right lung" although it can be inferred from the phrase "previous lobectomy on the right." A report might have included an observation such as

absence which *isStateOf'LeftLungLobe'*.

The schema presented here using *presence/absence* has the advantage that all conditions of the lungs are united under

PathologicalState which *isStateOf-Lung*.

The result is subtly different if we use expressions of the form

'LeftLungLobe' which *hasState-absent*,

since this entity will be classified as a kind of

'LeftLungLobe'

For most purposes, it seems better to subsume the "absence of a lobe of the left lung" under a "pathological state of the left lung" than to treat it as a peculiar kind of anatomic structure (i.e., one that is absent). This usage conforms to that of the GALEN CORE model (version 1.5), but it remains a matter of experiment whether it proves suitable in the long run.

Encapsulation and Transformation of Alternative Nominalizations

Different decisions were taken for each of the three medical-record structures described about what information to encapsulate and which of the nominalizations to include explicitly. In the simple code-based model, the codes represent concepts at varying levels of detail. Any fixed coding system must decide what level of detail is practical. The first two concepts, "left lower lobe atelectasis" and "decreasing left pleural effusion," are almost certainly concepts that would be found in any radiologic coding scheme. However, the third concept, represented by *code_1142.1*—"Post-operative changes consistent with coronary artery bypass graft as well as previous lobectomy on the right," is too detailed to be plausible for any static, precompiled coding scheme, even one including modifiers. An identifier for such a complex code might be generated dynamically by a terminology server as proposed by GALEN. Otherwise, the most likely solution would be simply to make *code_1140* equivalent to "post operative changes." Alternatively, some schemes might allow two codes to be applied. Some coding systems propose adding modifiers for concepts, such as "left" and "right" or "severity," in order to reduce the combinatorial proliferation of codes.

In the linked relational model, the encapsulation is reduced and the information in each previous code

date	subject	role	indicant	result	identifier
10-28-91	pid_12345	xrf	'SurgicalClip'	present	r1
10-28-91	pid_12345	xrf	Opacity	'RightParaMediastinal Region'	r2
10-28-91	pid_12345	xrf	'PlateLikeOpacity'	'LeftLowerLungZone'	r3
10-28-91	pid_12345	xrf	'PlateLikeOpacity'	'MidLowerLungZone'	r4
10-28-91	pid_12345	dx	'LeftPleuralEffusion'	Improved	r5
10-28-91	pid_12345	dx	'LeftLowerLobe Atelectasis'	present	r6
10-28-91	pid_12345	dx	'PostOp Changes'	present	r7
10-28-91	r6	int	based on	r3	r8
10-28-91	r65	int	based on	r4	r9

Figure 6 Linked relational model. Each entry in the "indicant" and "result" columns as a fixed code, but to make the table easier to read, the codes are represented by the corresponding GRAIL expressions rather than presented in a separate figure.

is split across two columns, one for the "indicant" and the other for the "result." There is a significant problem for reusability and application independence in that the implementor must decide which information should appear in the result column. In this version, we have chosen the "location" in the case of findings and the "presence or absence" or "improvement" in the case of examination results. This choice is arbitrary and almost certainly application-dependent. For example, location might have been encapsulated with the indicant for findings. Currently, there is no widely accepted method for articulating and documenting these choices or for ensuring that they are made consistently across all parts of an application. One role for GALEN's concept modeling will be to contribute to a standard means for articulating such choices and the mutual obligations that they imply between developers of concept systems and developers of the information systems that will use those concepts. A second role will be to provide the basis for mediation between systems based on different choices.

Transformations between models may well go beyond such structural transformations. For example, the original PEN&PAD model did not include "presence" explicitly in the form shown in the current GALEN CORE model but did provide an additional mechanism for dealing with negation. The mechanisms for dealing with presence, absence, and negation were part of the procedural rules in the program code. Transformations involving such procedural operations are outside the scope of the purely structural transformations supported by GRAIL and require additional mechanisms.

Clinical Ambiguity, Imprecision, and Usage and the "Faithfulness" of the Record

The prime principle of the PEN&PAD framework for the medical record is that it can be "faithful" to clinicians' intentions. Clinical usage is often imprecise or ambiguous and often blurs formal distinctions. Knowledge engineers tend to elaborate schemas to produce ontologically clean models. The conflict between the natural behaviors of clinicians and knowledge engineers leads to the danger of overinterpretation with consequent loss of faithfulness.

For example, in the schema presented in this article, every modifier has been attached to the thing itself rather than to its "presence." However, there is a strong argument that certain modifiers that concern the time course should be applied to the "presence" or "absence" rather than to the thing itself [e.g., "persistent presence or "acute (onset of) presence"]. Using such an alternative schema, the first part of the record might be

presence which <
  hasProgress-persistent
  isStateOf-'SurgicalClip'' which <
    hasLocation-'RightMediastinalRegion
    hasNumber-multiple>
  >

This formulation has obvious attractions but raises problems. Clinicians tend to use "persistent" as part of a continuum including "worsening" and "improving." Separating the continuum requires sacrificing clinical intuition. Recovering some of that intuition requires introducing an explicit notion of "unchanged" to go with "worsened" and "improving."

However, the combination of "persistent presence" and "progress unchanged" presents a potential temptation with regards to faithfulness. To maintain a uniform representation, we need to insert "persistent" as an implied modifier for "presence" whenever we meet "worsened" or "improved"; otherwise, attempts to retrieve all persistent pleural effusions will fail. The effusion must be "persistently present" for its "progress" to be described as "improved," "unchanged," or "worsened." However, the converse—adding "progress unchanged" as part of the interpretation of "persistent pleural effusion"—goes beyond the information explicitly included in the report and requires the system to make inferences that are not justified.

Problems of Reference

Expressions such as "opacity" and "left and mid lower lung zone" raise problems of reference, as well as of ambiguity. Opacities do not occur in the lung but rather on radiographs of the lung. Likewise, the natural model of "the zone of the left lung" fails entirely when applied by analogy to the "zone of the mid lung" since there is no such anatomic structure. To make sense, the "zone" must apply to the "chest," and the model here represents the expression as the "zone of the chest which is related to lung." Even this is inadequate since zones, like opacities, occur on radiographs rather than in the body. In the model presented, the fact that *Radiograph22-shows-Opacity. . .* indicates that the *Opacity* is an *Opacity which isShownOn-Radiograph. . . .* Formal transformations so that the "lung" in which the "opacity" was located was likewise seen on "radiograph" are under development but are not yet fully defined or implemented. In the current version of the GRAIL Kernel, "opacities" and "zones" are simply modeled as located in "lungs."

Constraints, Verification, and Generation

The focus of this article has been primarily on the expressiveness of the CORE model and issues of encapsulation of composite concepts in medical records. A third, equally important issue is that of constraining what can be said to that which is sensible. The GRAIL Kernel differs from much other work,[29] in the detail of its constraints and sanctions on encompassing concepts and in its emphasis on generation of new composite entities in response to queries from users and applications. The purpose of the detailed constraints is to limit the concepts generated to those that are sensible.

Constraints and sanctions on the creation of composite entities are important in at least four other areas:

1. In verifying that input is correct—Is the proposed composite entity sanctioned?
2. In using the model as a source of implied concepts for generating data-entry forms for predictive data entry[13]—What are the sanctioned descriptors for this entity that might appear on a form?
3. For enforcing consistent usage when the model is used for knowledge engineering—How can this entity be used?
4. As semantic constraints to disambiguate expressions in natural-language processing—How can a set of entities be fit together to make a sensible sentence?[30]

The basic approach to constraints and sanctions was mentioned earlier in this article. Methods for evaluating the adequacy of the constraints are still under development but are based around methods for testing whether all of the concepts generated on the basis a model are clinically sensible. We would maintain that to form a sound basis for clinical information management, any approach to concept modeling must be evaluated both for its ability to express what needs to be said and for its ability to reject nonsense.

Incomplete Information and Linkage to External Information

The PEN&PAD model of the medical record is of a complete logbook. All information is complete, and all references to previous or subsequent observations are explicit. In dealing with a fragment such as a single radiographic report or in providing any other data-interchange format, provision for linkage to omitted or unknown data is required (e.g., to a previous radiographic film). An example of this would be

```
(RadiographicFilm which isGeneratedBy-
  ('PAViewLungs' which <
    isActionPerformedOn-\PtX
    isPerformedAt-'AgentAtSession1'>)
) name 'PreviousRadiographicFilm'.
```

Since the radiologist explicitly compared the current film with the previous one, presumably a record of the previous radiograph would be in the patient record someplace, but the system must make provisions for linking to it if the report is to be presented or stored in isolation.

More seriously, in references such as "previous radiation therapy" there is the implication of an event for which no record may be available. The GRAIL kernel can only deal with external or undefined references by creating new individuations of corresponding category, such as *RadiationTherapy*. The keyword *a/an* was introduced for this exercise to indicate such new individuations and is used in the example below.

```
(a performance which isStateOf-
  (RadiationTherapy which <
    isTreatmentOf-PtX
    isAtTime-(a Time which isPreviousTo-\Time1)>)
) name 'RadiationTreatment1',
```

The extent to which such references can be used to infer the actual existence of the event referred to is a serious issue—under what circumstances is it appropriate to retrieve this patient in response to a request for "all patients who have had radiotherapy" in the absence of any explicit record of that therapy? After all, the radiologist might have misinterpreted the film. This is not an issue that can be solved by transformations of concept models alone.

Time and Temporal Attributes

Temporal attributes and descriptors can be modeled in the GRAIL Kernel, but temporal interpretation is considered to be outside the range of transformations on the concept model or terminologic reasoning. The most a GRAIL model can represent are the constraints on what is sensible, e.g., that it is not sensible for one event to be both before and after another event. GALEN's approach to uncertainty is similar—concept models may be built to represent how uncertainty may be expressed sensibly, but the GRAIL engine does not perform inferences based on that uncertainty. In both cases, such inferences are the province of external-reasoning systems.

Numbers and Quantities

The current representation in the GRAIL Kernel supports only a minimal notion of number, although extensions have been considered. Expressions such as *SternotomyWire which hasNumber-4* have clear limitations, although they have served well in practice.

Conclusion

This article presents the results of a case study performed using an early version of GALEN's compositional approach to representing medical concepts. It reports preliminary work from early in the life of a long-term project. In outline, the compositional approach was deceptively easy to apply, although some may find surprising the volume of information in even a brief report. Whatever scheme is used, we doubt that this result can be avoided—complete detailed electronic patient records will be large.

The outline of how to apply the compositional approach may be clear, but the details raise a series of thorny issues. The ability to form compositional models requires choices to be made concerning questions that previously could not be expressed (e.g., whether the disease is the "ulcer" or the "presence of the ulcer"). To understand fully the consequences of the various potential choices takes time and experimentation. Many of the constructs used in this article must remain tentative. Experience within the project suggests that *a priori* dogmatism as to the theoretically "correct" solution is unwise. The goal is to produce systems that work with real applications—indeed, with more than one real application. Informed choices require practical experience.

While the use of a compositional model raises certain issues, it avoids many issues that often cause controversy in the design of traditional coding and classification systems. If an entity can be subsumed by several categories, then arguments about which way to classify the entity are unnecessary. If the information about task and method of observation can be cleanly separated from pathophysiologic concepts, then the question as to whether an entity is a symptom, sign, or diagnosis becomes transformed into a question as to whether it is sensible to compose a symptom, sign, or diagnosis from that entity (i.e., whether it is sensible for the entity in question to be reported by the patient, observed by the doctor, or inferred by the doctor in the role of diagnosis). No argument is necessary; the model may sanction any, all, or none of these compositions. Similarly, if different formulations can be seen as different nominalizations of the same underlying network, discussions of whether the information system should be concerned with the "opacity which is interpreted as atelectasis" or the "atelectasis which is inferred from the opacity" become pointless as decisions can be taken locally as convenient and the results transformed as required. .

Avoiding pointless arguments has become one of the explicit goals for the further development of the GALEN modeling framework and the GRAIL Kernel. It remains to be seen to what extent this goal can be achieved and which of the new issues raised by the compositional models themselves can be similarly reduced to matters of formal manipulations and therefore be made formally pointless and which are substantive.

Even if certain arguments are rendered formally pointless, the computational tasks involved in the manipulation of the models are nontrivial. Hence, the goal of a terminology server as an active engine for manipulating concept representations. A major function of the terminology server is to relieve applications of the need to perform these operations for themselves or even to know the details of how they are performed.

Is the goal of a reusable, application-independent model of medical terminology feasible on a broader scale? History demands that we maintain a healthy skepticism, particularly with regards to problems of scaling. GALEN's strategy is twofold. On the one hand, it restricts the range of what is meant by a concept model, resisting the temptation to incorporate evermore forms of inference within the GRAIL Kernel. On the other hand, it provides facilities to allow clean separation of different taxonomies and to coordinate subsumption with other transitive relations.

If such a model could be constructed, would it be useful? Again, we must remain skeptical pending further results. GALEN seeks to manage the technical complexity by encapsulating it in a terminology server. It seeks to test the usefulness of the CORE model and terminology server through a series of experiments using them, even in their current preliminary states, to drive practical applications. The results reported here are from a preliminary "experiment of opportunity" that was one part of that effort. Within the limits of a case study, it was possible to represent the report in great detail and to reuse the same concept model for two alternative encapsulations of the information, but numerous problems were raised that remain to be solved. Other more extended and more systematic experiments are under way. Better estimates of both the feasibility and the usefulness of GALEN's models and terminology server will become available as the project matures.

References ■

1. Nowlan W, Rector A, Rush T, Solomon W. From terminology to terminology services. In: Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care. Washington, DC, 1994 (in press).
2. Masarie F Jr, Miller R, Bouhaddou O, Giuse N, Warner H. An interlingua for electronic interchange of medical information: using frames to map between clinical vocabularies. Comput Biomed Res. 1991;24(4):379–400.
3. Evans DA, Cimino JJ, Hersh WR, Huff SM, Bell DS. Toward a medical-concept representation language. J Am Informatics Assoc. 1994;1:207–17.
4. Rector A, Nowlan W, Glowinski A. Goals for concept representation in the GALEN project. In: Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1993:414–8.
5. Rector A, Nowlan W, Kay S. Conceptual knowledge: the core of medical information systems. In: Lun K, Degoulet P, Pierre T, Rienhoff O, eds. Seventh World Congress on Medical Informatics, MEDINFO-92. Geneva: North-Holland, 1991:1420–6.
6. Lindberg D, Humphreys B, McCray A. The Unified Medical Language System. In: van Bemmel J, ed. 1993 Yearbook of Medical Informatics. Amsterdam: International Medical Informatics Association, 1993:41–53.
7. McCray A, Hole W. The scope and structure of the first version of the UMLS semantic net. In: Miller R, ed. Fourteenth Symposium on Computer Applications in Primary Care. Washington, DC: IEEE Computer Society, 1990:126–30.
8. Nowlan W, Rector A, Kay S, Horan B, Wilson A. A patient care workstation based on a user centered design and a formal theory of medical terminology: PEN&PAD and the SMK formalism. In: Clayton P, ed. Fifteenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1991:855–7.
9. Horan B, Rector A, Sneath E, et al. Supporting a humanly impossible task: the clinical human-computer environment. In: Diaper D, ed. Interact 90. Geneva: North-Holland, 1990:247–52.
10. Rector A, Nowlan W, Kay S. Foundations for an electronic medical record. Methods Inf Med. 1991;30:179–86.
11. Rector A, Nowlan W, Kay S, Goble C, Howkins T. A framework for modelling the electronic medical record. Methods Inf Med. 1992;32(2):109–19.
12. Cimino JJ. Controlled medical vocabulary construction: methods from the Canon Group. J Am Med Informatics Assoc. 1994;1:296–7.
13. Nowlan W, Rector A. Medical knowledge representation and predictive data entry. In: Stefanelli M, Hasman A, Fiesch M, Talmon J, eds. Artificial Intelligence in Medicine Europe (AIME-91). Maastricht: Springer-Verlag, 1991:105–16.
14. Rector A, Nowlan W, Kay S. Unifying medical information using an architecture based on descriptions. In: Miller R, ed. Fourth Annual Symposium on Computer Applications in Medical Care. Washington, DC: IEEE Computer Society Press, 1990:190–4.
15. Nowlan W. Representation of Medical Terminology (dissertation). Manchester, UK: University of Manchester, 1993.
16. Rector A, Nowlan W. The GALEN Representation and Integration Language (GRAIL) Kernel, version 1. In: The GALEN Consortium for the EC AIM Programme. 1993. (Available from the Medical Informatics Group, University of Manchester).
17. Brachman R, Fikes R, Levesque H. An essential hybrid reasoning system: knowledge and symbol level accounts of KRYPTON. In: International Joint Conference on Artificial Intelligence. San Mateo, CA: Morgan Kaufmann, 1985:532–9.
18. MacGregor R. The evolving technology of classification-based knowledge representation systems. In: Sowa J, ed. Principles of Semantic Networks: Explorations in the Representation of Knowledge. San Mateo, CA: Morgan Kaufmann, 1991:385–400.
19. Haimowitz I, Patil S, Szolovits P. Representing medical knowledge in a terminological language is difficult. In: Twelfth Annual Symposium on Computer Applications in Medical Care. Washington, DC: Computer Society Press, 1988:101–5.
20. Doyle J, Patil R. Two theses of knowledge representation: language restrictions, taxonomic classification and the utility of representation services. Artif Intell. 1991;48:261–7.
21. Sowa J. Conceptual Structures: Knowledge Representation in Mind and Machine. New York: John Wiley & Sons, 1985.
22. Bernauer J, Goldberg H. Compositional classification based on conceptual graphs. In: Artificial Intelligence in Medicine Europe. Munich, 1993.
23. Campbell KE, Das AK, Musen MA. A logcial foundation for representation of clinical data. J Am Med Informatics Assoc. 1994;1:218–32.
24. Baud R, Rassinoux A-M, Scherrer J-R. Knowledge representation of discharge summaries. In: Stefanelli M, Hasman A, Fieschi M, eds. Third European Conference on Artificial Intelligence in Medicine. Marseilles: Springer-Verlag, 1991:173–84.
25. Rector A. Data decay, significance and confidentiality: a time-oriented data model for comprehensive care. Med Infr (London). 1981;6:187–93.
26. Evans D, Ginther-Webster K, Hart M, Lefferts R, Monarch I. Automatic indexing using selective NLP and first-order thesauri. In: Lichnerowicz A, ed. Intelligent Text & Image Handling. Proceedings of RIAO '91. Amsterdam: Elsevier, 1991:624–44.
27. Rector A, Gangemi A, Galeazzi E, Glowinski A, Rossi-Mori A. The GALEN CORE model schemata for anatomy: towards a re-usable application-independent model of medical concepts. In: Barahona P, Veloso M, Bryant J, eds. Twelfth International Congress of the European Federation for Medical Informatics. Lisbon, Portugal: 1994:229–33.
28. van der Lei J. Experience from computer-based patient records for computer-assisted decision making (editorial). Methods of Inf Med. 1993;32(1):14–5.
29. 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.
30. Morel-Rassinous A-M, Baud R, Scherrer J-R. Proximity processing of medical texts. In: Medical Informatics Europe. Berlin: Springer-Verlag, 1990:625–30.