

# Interface of Inference Models with Concept and Medical Record Models

Alan L. Rector<sup>1</sup>, Peter D Johnson<sup>2</sup>, Samson Tu<sup>3</sup>, Chris Wroe<sup>1</sup> & Jeremy Rogers<sup>1</sup>

<sup>1</sup>Medical Informatics Group, Dept of Computer Science, University of Manchester, UK  
(arector jrogers wroec)@cs.man.ac.uk

<sup>2</sup>Sowerby Centre for Health Informatics in Newcastle, University of Newcastle upon Tyne, UK  
pete@mimir.demon.co.uk

<sup>3</sup>Stanford Medical Informatics, Stanford University, USA  
tu@SMI.Stanford.EDU

## Citation:

The formal version of this paper appeared as:

A L Rector PD Johnson S Tu C Wroe and J Rogers (2001) Interface of inference models with concept and medical record models. in S Quaglini, P Barahona and S Andreassen (eds) *Proc Artificial Intelligence in Medicine Europe (AIME-2001)* Springer:314-323

**Abstract.** Medical information systems and standards are increasingly based on principled models of at least three distinct sorts of information – patient data, concepts (terminology), and guidelines (decision support). Well defined interfaces are required between the three types of model to allow development to proceed independently. Two of the major issues to be dealt with in the defining of such interfaces are the interaction between ontological and inferential abstractions – how general notions such as ‘abnormal cardiovascular finding’ are abstracted from concrete data – and the management of the meaning of information in guidelines in different contexts. This paper explores these two issues and their ramifications.

## 1. Introduction

Medical Information systems and standards are increasingly based on principled models. Recent papers by two of the authors [1] [2] discussed the interface between three different types of models used by different groups to represent different aspects of clinical information:

1. *‘Information’ or ‘patient-data’ models* – the structure of the information to be stored, e.g. HL7’s Reference Information Model (RIM) [3], CEN’s Electronic Healthcare Record Architecture [4] or the models underlying GEHR’s architecture [5, 6] typically expressed in UML object models. Also referred to as the ‘Patient Data Model’ [2]
2. *‘Concept models’ or ‘ontologies’* – the meaning of what is stored (referred to in [2] as the ‘medical specialty model’). In this framework, a ‘terminology’ consists of the combination of a ‘concept model’ plus one or more ‘lexicons’ supplying natural language words to signify the concepts. For example, SNOMED-RT’s compositional representation [7] or GALEN’s Common Reference Model [8, 9] are concept models in this sense. Concept models are typically expressed in hierarchies, frame systems, or their more modern successors, description logics.
3. *‘Inference’ models* – models that encapsulate knowledge needed to derive the conclusions, decisions, and actions that follow from what is stored. In this paper, we will use models of clinical guidelines and protocols as examples of such inference models, although the issues discussed here apply to medical decision-support systems more generally e.g. Medical Logic Modules (MLMs) [10] EON [11], Prodigy [12], PROForma [13] or GLIF [14].

This tripartite functional division mirrors an analysis of the operational information requirements of guideline-based decision-support systems in which those information needs are divided into two broad categories and the second category divided again into two further categories as follows:

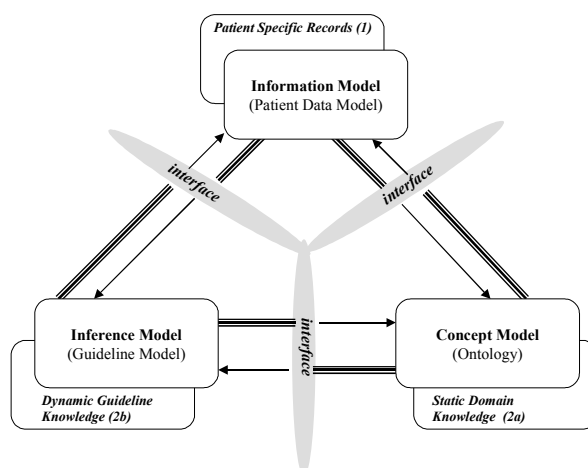
1. *Information about specific patients and clinical situations* – information gathered from healthcare workers about clinical observations and results and also, potentially, inferences made by guidelines or other software agents based on the human input.
2. *General patient-independent information about medicine and medical practice (also known as ‘domain knowledge’)*  
This information can be subdivided further:

2a. *Guideline independent static knowledge* – definitions and intensions of concepts and closely associated facts such as indications and contraindications of drugs, causes and symptoms of diseases, and mechanisms of tests and procedures which are both relatively independent of individual guidelines and directly linked to particular concepts without complex inference.

2b. *Guideline dependent dynamic knowledge* – the model of how to infer conclusions and decisions from the patient specific information and guideline independent facts.

Our assumption is that the goal is to express protocols or guidelines for care in an executable form that interacts with an electronic patient record (and broader clinical information systems). We assume that the key elements in inference models are ‘criteria’ – logical expressions which trigger inferences or actions. In order to execute a guideline for a specific patient, criteria from the inference model must be transformed into queries against the corresponding electronic patient record. We assume that in principle each model specifies a corresponding repository of information - patient records, dynamic executable guidelines, and static domain information respectively. We assume, however, that electronic patient records and clinical information systems will continue to be designed largely independently from any reusable guideline-based decision-support systems. Therefore it is essential to design a clean and well defined interface between them.

This architecture is shown graphically in Figure 1. The three models are shown as boxes shadowed by their respective repositories. For clarity, the number referring to the type the classification above is shown in brackets with each repository. The interfaces are shown as ‘lenses’. Each interface is bi-directional. In each direction, each interface filters and transforms information so that the receiving model sees only that portion of the originating model and associated repository that is relevant for its own purposes – as symbolised by the arrows emerging from each model as multiple broad lines and continuing on from the interface as single thin lines. As well as carrying information, the arrows carry obligations. Each interface places mutual obligations on the two models it connects, and all these obligations must be coherent – as symbolised by the overlap of the interfaces at the centre of the drawing.



**Fig. 1.** Interfaces, Models and corresponding repositories

The architecture in figure 1 raises many questions. This paper examines two important issues –

1. Expressing the abstractions in guideline criteria operationally in terms of the patient-specific information in the medical record
2. Relating patient independent information to individual patients in the clinical context and how to partition this information between the concept and inference models.

The ideas will be illustrated by examples drawn from the UK Department of Health sponsored PRODIGY project developing guidelines for chronic disease prescribing in primary care [12].

## 2. Abstraction and Operationalisation

The main barriers to successfully expressing guideline criteria in a form compatible with an electronic patient record (EPR) system are closely related:

1. *Differences in granularity and axes of classification.* The information in the criteria may not be expressed at the same granularity as in the medical record or the relevant abstraction may be missing in the terminology used in the record, e.g. there is no single code for “steroid” in most standard drug classifications, although there are codes for all

of the individual drugs in question and abstractions for covering various subclasses of steroids (e.g. “topically applied steroids”).

2. *Need for inferential abstractions*: Some information may be recorded as raw measurements which need to be interpreted before use - e.g. “elevated white cell count” may never be explicitly recorded by the user but must instead be inferred from a range of values. .
3. *Ambiguity of operational meaning*: The criteria in a written guideline may not correspond to any obvious information likely to be ever recorded in the medical record, e.g. “tendency to prolonged ST interval”.
4. *Differences in encapsulation and form of expression*: There are almost always many different ways in which the same information can be packaged or ‘encapsulated’ (see [1]). Information expected as a code-value pair in one (‘melanoma’ + ‘metastatic’), may be represented as a single code (‘metastatic melanoma’), or as a code plus a series of modifiers in the other (‘neoplasm’ + ‘melanocyte’ + ‘malignant’ + ‘metastatic’).

Each of the above issues will be discussed in turn.

## 2.1 Ontological and Inferential abstractions

The information recorded in medical records is usually fine grained and specific – concrete findings, diagnoses, drugs, procedures, etc. The criteria in guidelines are usually expressed as abstractions of such findings. Ontological abstraction bridges the difference in granularity between the information in the record and that being looked for by the guidelines. Roughly, the requirement is that the concept model be capable of classifying the information in the medical record under the abstractions used in the guideline.

Unfortunately, many standard terminologies cannot make this link reliably. For example, the example of the absent category “Steroid” in the current drug classifications used by the British National Formulary (BNF) is typical: although the classification provides for various types of steroids, e.g. “topical steroids,” it does not provide all needed generalisations at the correct granularity.

There are at least two reasons for these failures or omissions:

1. The axis of classification required may not exist in the concept system or may not exist consistently throughout the system – e.g. ‘indication’ may not be used consistently as an axis and ‘contraindication’ may not exist as an axis of classification at all.
2. The concept system may be structured primarily for navigation using mixed axes, *i.e.* they are structured as ‘thesauri’ rather than as formal classifications [9]. Abstraction works reliably only along strict subsumption (“is-a”) hierarchies in formal classifications. ICD, the Read Codes v1 and v2 and most drug classifications are examples of thesauri in this sense.

Inferential abstraction, by contrast, is required when a criterion used in a protocol cannot be determined by simple logical operations directly from the meaning of individual terms in the medical record. This occurs in several situations:

1. When a single quantitative item of data or very specific observation needs to be converted to a qualitative abstraction using mathematical or logical inferences beyond the restricted set provided by the concept model – e.g. to convert a blood pressure of 190/110 to “elevated blood pressure”.
2. When a quantitative summary value must be calculated from several other values in the medical record – e.g. body mass index, predicted peak expiratory flow rate, or cardiovascular risk probability.
3. When a temporal abstraction must infer a state or trend from several items of information in the record – e.g. to obtain an abstraction such as “rising blood pressure” or “on steroids for more than three months”. Such inferences often depend on the clinical context and locality as well as the raw data (see [15]).
4. When the required abstraction cannot reliably be found in the medical record and must be inferred from proxies.

It is a common feature of inferential abstraction that it is defeasible, *i.e.* that subsequently acquired information may invalidate an earlier inference. This is particularly common in dealing with temporal abstraction or proxies in the medical record.

The details of inferential abstraction depend on three sets of issues:

1. *Operational issues* – what information is likely, in practice, to be in the patient record. Presented with a paper guideline, human readers will make the necessary inferential leaps. Computer assisted guidelines must specify them explicitly, which involves formalising the intuitions by which by which much medical data is actively reconstructed as part of the decision making process rather than predefined [16]. This is made more difficult by the fact that the medical record is rarely either complete or consistent - e.g.a recent British study found that “Ischaemic Heart Disease” was recorded in less than half the relevant cases in general practice [17].
2. *Issues of clinical context* – what are the criteria for a given inference and in what contexts is it permissible, safe, and/or necessary to make it. These often involve subtle distinctions in usage – the entry criteria for an asthma treatment guideline may be different from the criteria used to trigger warnings concerning drugs contraindicated in asthma – or local conventions – e.g. Two different guidelines may have different entry criteria, or two different authorities may have different criteria for when the risk of bronchospasm is sufficient to trigger a warning

3. *Formal issues* – how should the criteria be expressed formally internally and with respect to the information and terminology models.

As a practical example, in the PRODIGY guidelines it has been determined *operationally* by experience that a diagnosis of “asthma” is not always recorded in patients who should be recognised as matching the criterion “contraindicated in asthma”. *Clinically* it has been decided that more than one episode of “wheezing” in the past year is but one of the proxy criteria sufficient to count as “asthma” for these purposes. *Formally* the criteria are represented as logical expressions in which the terms ultimately map to sets of Clinical Terms (Read Codes).

## 2.2 Ambiguity of operational meaning

Written guidelines are written for human interpretation. It is not sufficient to represent a notion found in a text guideline or warning ‘literally’; it must also be operationally and clinically useful and effective to do so.

For example the British National Formulary contains references to drugs being contraindicated in the presence of a “tendency to prolonged ST intervals”. Even an experienced physician might have to think twice as to what this meant in terms of the concepts represented in an electronic patient record. However, since it is a recurring phrase, it is worth attempting to capture it in a general way, *e.g.* as any suggestion of heart block. This in turn suggests defining an inferential abstraction for “suggestion of heart block”.

## 2.3 Differences in encapsulation and structure amongst models

The structure of encapsulation used in various inference and information models differs widely. Encapsulation is a key issue for the interface between the information and concept models [1], and analogous issues exist with the inference model.

For example, the Arden Syntax assumes that patient data can be represented as lists of time-stamped values [18]. It defines a criteria language that includes operators to manipulate such lists (*e.g.*, *latest* or *merge*) and make comparisons of values and of timestamps (*e.g.*, *greater than* or *is before*). Prodigy, by contrast, assumes that information about patients can be expressed as instances of structured objects, such as *Investigation\_Result* and *Note\_Entry*, having time stamps or time intervals denoting periods of validity. A *Note\_Entry*, for example, specifies a domain concept and a time interval during which the condition represented by the concept is present. Eon, meanwhile, represents information as part of a temporal database that permits complex temporal queries and abstractions. None of these three inference models correspond directly to any commonly used medical record model.

There are three approaches to this problem:

1. *Ad hoc* mapping
2. *Use a standards-based information model* such the HL7 RIM, CEN’s EHRC or GEHR.
3. *Creating a virtual medical record* within the inference model and define a mapping from that model to each individual information model

In the *ad hoc* approach, exemplified by Arden Syntax for Medical Logic Modules, differences in encapsulation are dealt with by *ad hoc* programmes. The alternative is to base the mapping on a standard information model and associated interchange format such as the HL7 RIM or CEN’s EHRC.

There are two difficulties with this approach: a) Achieving a standard view to meet the requirements of both the interface and information models may be difficult. b) No currently existing medical record systems are based firmly on the RIM or any similar standard, although several are in construction. To circumvent these problems both the PRODIGY and EON systems have defined ‘virtual medical record systems’ – in effect medical record models tailored to the needs of the guideline system.

## 3. Static Domain Information and Context

One of the major benefits of using modern ontological methods is that they are an effective means of organising complex sets of special cases along multiple axes. For example systematically creating notions such as “normal range for haemoglobin for a patient in renal failure”. In principle, this can even be extended to capture local policies “maximum dose of inhaled steroid for a patient on local research on research protocols”.

This approach provides two important benefits:

1. *Safety and consistency* – General policies can be implemented that can only be overridden explicitly – *e.g.* to express a generalisation that “beta blockers are contraindicated by the risk of bronchospasm” but include specific exceptions for certain beta blockers in specific clinical situations. In this way any new beta blocker will ‘inherit’ the contraindication unless it is explicitly overridden.

2. *Scalability* – The total number of such special cases is potentially combinatorially explosive. As the number of guidelines and variants increases, managing the underlying information consistently becomes increasingly difficult. This is most obvious currently in fields such as drug information, but as guideline libraries grow in size, it can be expected to become an increasingly important issue.

Given such a broader ontological knowledge base, there may be tight integration between ontological and inferential abstraction and between the concept and inference models. For inferential abstraction, the inference itself must be represented in the inference model, but the information used by the inference and further abstractions from that initial inference may be represented in the ontology – e.g. both the concept of the lower bound for a normal white-blood cell count (WBC) and the concept represented by the qualitative abstraction – e.g. “leukopenia” – might be represented in the ontology. The inference from a specific measured value to the corresponding qualitative abstraction – i.e., the criteria that relates the lower bound to the abstraction – resides in the inference model.

#### 4. Example

Consider the criteria in the example below of a condition action rule from a decision support system:

IF

CRITERIA:

a patient has a diagnosis of *confirmed hypertension* recorded AND the patient has either  
*diabetes mellitus* OR

*persistent highly elevated cholesterol* OR

*a family history of heart disease*

AND the patient has *no renal disease*

AND the patient is on no *drug* with a *serious interaction* to *chronic therapy* with *ACE inhibitors*

THEN

INFERENCE: an *ACE inhibitor* is a candidate treatment for the *hypertension*”

Taking the concepts in turn, the construct “a diagnosis of ... recorded” must be mapped to a corresponding encapsulation in the information model.

“*Confirmed hypertension*” may be specific to this protocol “*confirmed hypertension as per guideline\_h\_23*”. Alternatively, it may be a locally tailorable notion in the ontology used by several different guidelines, or a widely agreed concept within a clinical community. Whichever, the implication is that it is a predefined notion either entered directly by a clinician or, alternatively, was inferred earlier from a machine analysis of a series of documented blood pressure readings. It may be that the ontology allows for specialisations of such concepts, so that it includes, for example, “confirmed surgical hypertension”, “confirmed essential hypertension”, “confirmed renin dependent hypertension”, etc. An ontology might provide a range of such notions, possibly with criteria for each. Alternatively, a post-co-ordinated concept system might provide a general notion of ‘confirmed’ which could be either entered or inferred.

The concept “*diabetes mellitus*” may be found literally or, more likely, will require simple ontological abstraction from either “insulin dependent diabetes mellitus” or “non-insulin dependent diabetes mellitus”, being those terms at the minimum level of disambiguation that is clinically useful or necessary and, therefore, what is likely to be in the record. Additionally, it might be inferred by proxy for patients taking injectable insulin even in the absence of any diabetes-related diagnosis term.

“*Persistent highly elevated cholesterol*” by contrast will usually require two steps of inferential abstraction. Step one is to abstract from each numerical measurement of cholesterol, to determine whether or not it meets the definition of “highly elevated cholesterol”. The second, to abstract from a temporal series of such abstractions, determining whether the elevation is persistent. Alternatively, a more sophisticated temporal abstraction mechanism might perform the two operations together.

“*Family history of heart disease*” and “*no renal disease*” present superficially analogous problems. Firstly the encapsulation and representation of “family history of ...” and “no ...” must be determined from the interface to the information model. Secondly, there is the question of what constitutes “Heart disease” and “Renal disease”. In principle, the second step ought to be a matter of ontological abstraction. In the case of “family history of heart disease” this is probably sufficient or at least all that is possible. However, in context, this may not be sufficient for “no renal disease”. Since this is in the context of a contraindication, it should be reformulated as a more sensitive criterion such as “no evidence suggesting renal disease” and inferred from the medical record.

“*No serious contraindication to chronic therapy with ACE inhibitors*” requires coping with the issues of negation and the interface to the information model for drugs, and a rich ontological knowledge base able to perform the ontological abstractions indicated by “...*serious contraindication to chronic therapy with...*”

## 5. Consequences for the interface

The above analysis suggests a number of requirements for the interface from the inference model to the information and concept models and amongst the three models.

1. Amongst the three models:
  - *Consistency of meaning*: The meaning of all concepts used, wherever they appear, must be the same. In principle, all should be derived from the definitions in the concept model.
  - *Unambiguous encapsulation of information between information and concept models*: In order for “family history of heart disease” and “no renal disease” to be represented unambiguously, the interface between the concept model and the information model must specify responsibility and structures for status concepts such as “family history of ...”
2. From the inference model to the patient information model
  - *Adequacy of structure*: The information model must hold the patient information necessary to determine if criteria are satisfied, in the example above this implies at least “observation”, “diagnosis”, “laboratory test result”, “drug sensitivity” and “medication”, but a more complete analysis would extend this list.
3. From the inference model to the concept model
  - *Axes of the concept model and ontological knowledge base*: The interaction check requires that interactions be held in the ontological KB, in the example that it include an appropriate notion of the seriousness, and that the use of the drug and whether it is acute or chronic be axes along which interactions are classified.
  - *Level of abstraction in the concept model*: The inference model requires abstractions at the level of “diabetes mellitus”, “renal disease” and “heart disease” to be provided by the concept model. [swt1]
  - *Content of the ontological knowledge base*: The inferential abstraction from laboratory values to “persistently highly elevated cholesterol” will require relevant thresholds to be stored in the ontological knowledge base. If “no renal disease” is to be obtained by inferential abstraction, it too will require relevant thresholds in the ontological knowledge base.

## 6. Discussion

Parts of the architecture in Figure 1 have been implemented in by various groups, but a complete implementation in which the interfaces are well enough specified to allow easy integration and scaling has yet to be achieved.

For example, the ‘virtual medical record’ in PRODIGY and EON can be regarded as a ‘thick’ version of the interface between the inference and information models. The inference model side of the interface is well defined, but the information model and its mapping to site-specific electronic medical-record system consists still of *ad hoc* programming[2]. HL7’s Message Development Framework (MDF) [19] with its progression from a Reference Information Model (RIM) through a Message Information Model (MIM) to a Refined Message Information Model (R-MIM) represents one part of an interface from the Information Model to the Inference Model, which potentially deals with many issues of encapsulation and transformation, but within a framework in which each message must be worked out (almost) individually. GALEN’s perspective mechanism [20] and the proposed SNOMED-CT subset mechanism [21], each provide parts of the interface between the ontological knowledge base and the inference and information models, but neither is fully formalised and neither deals adequately with the mutual obligations required to represent unambiguously negation and notions such as “family history of” (‘status concepts’). A more systematic approach to these issues is sketched in a previous paper [1] but is far from being implemented.

Standard frame systems such as those derived from PROTÉGÉ used in EON provide a knowledge base linked to a relatively simple ontology. The UK Drug Ontology is moving towards representation of context within the ontological knowledge base as did the PEN&PAD user interface [22], but the mechanisms used need to be generalised and formalised.

There remains a question as to whether ‘abstraction’ in the sense used in this paper should remain an emergent property of cooperation between the three models or requires a separate mechanism such as the Tzolkin module used in the EON architecture [2].

Putting all of the pieces together remains a challenge. New technical frameworks such as those provided by Franconi [23] for schema evolution and the development of standards for computable ontologies – the Ontology Inference Layer (OIL) [24] may provide further tools. However, what is urgently needed are practical tools for defining each interface.

## References

1. Rector, A.L. The Interface between Information, Terminology, and Inference Models. in Medinfo-2001 (2001) (in press).
2. Tu, S.W. and M.A. Musen. Modeling data and knowledge in the EON guideline architecture. in Medinfo 2001 (2001) (in press).
3. HL7, HL7 Data Model Development, (2000), <http://www.hl7.org/library/data-model/>
4. CEN/WG1, ENV13606: Electronic Healthcare Record Architecture, . 1999, CEN.
5. Ingram, D., GEHR: The Good European Health Record, in Health in the New Communications Age, M. Laires, M. Ladeira, and J. Christensen, Editors. IOS Press: Amsterdam. (1995) 66-74.
6. Education, C.f.H.I.a.M., The GEHR Homepage, (1997), <http://www.cww.chime.ucl.ac.uk/HealthI/GEHR/>
7. Spackman, K.A., K.E. Campbell, and R.A. Côté, SNOMED-RT: A reference Terminology for Health Care. Journal of the American Medical Informatics Association (JAMIA), ((Symposium special issue)): .(1997) 640-644.
8. OpenGALEN, *OpenGALEN Home Page*, [www.opengalen.org](http://www.opengalen.org)
9. Rector, A., Thesauri and formal classifications: Terminologies for people and machines. *Methods of Information in Medicine*, . 37(4-5): .(1998) 501-509.
10. Pryor, T. and G. Hripscak, The Arden syntax for medical logic modules. *Int J Clin Monit Comput*, . 10(4): .(1993) 214-224.
11. Tu, S.W. and M.A. Musen. A flexible approach to guideline modelling. in AMIA Fall Symposium. Hanley and Belfus (1999) 420-424.
12. Johnson, P.D., *et al.*, Using scenarios in chronic disease management guidelines for primary care. *JAMIA*, (Symposium Special Issue): .(2000) 389-393.
13. Fox, J. and S. Das, *Safe and Sound*, Cambridge MA: MIT Press. (2000) .
14. Peleg, M., *et al.* GLIF3: The evolution of a guideline representation format. in AMIA Fall Symposium (2000) 645-649.
15. Shahr, Y. Timing is everything: temporal reasoning and temporal data maintenance in medicine. in *Seventh Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making*. Springer Verlag (1999)
16. Berg, M., Medical work and the computer-based patient record: A sociological perspective. *Methods of Information in Medicine*, . 37: .(1998) 294-301.
17. Bray, J., *et al.*, Identifying patients with ischaemic heart disease in general practice: cross sectional study of paper and computerised medical records. *British Medical Journal*, . 321: .(2000) 548-050.
18. Hripscak, G., *et al.* The Arden Syntax for Medical Logic Modules. in *Fourteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-90)*. McGraw Hill (1990)
19. HL7, HL7 Home Page, [www.hl7.org](http://www.hl7.org)
20. Solomon, W., *et al.*, Having our cake and eating it too: How the GALEN Intermediate Representation reconciles internal complexity with users' requirements for appropriateness and simplicity. *Journal of the American Medical Informatics Association*, (Fall Symposium Special Issue): .(2000) 819-823.
21. College of American Pathologists, SNOMED Home Page, [www.snomed.org](http://www.snomed.org)
22. Nowlan, W.A., Clinical workstation: Identifying clinical requirements and understanding clinical information. *International Journal of Bio-Medical Computing*, . 34: .(1994) 85-94.
23. Franconi, E. A Semantic approach for schema evolution and versioning in object-oriented databases. in *6th International Conference on Rules and Objects in Databases (DOOD'2000)* (2000)
24. The OIL Home Page, (2000), [www.ontoknowledge.org/oil/](http://www.ontoknowledge.org/oil/)

[swt1] Indeed there are multiple possible subsets. In HL7 RIM, diagnosis and laboratory-test result are subsumed under “observation”. How about “to record at least basic demographic information, “observations” (including diagnosis and laboratory-test results, “drug sensitivity,” “medication,” and other medical procedures.”