

The *CLINICON* Framework for Context Representation in Electronic Patient Records

Robert Müller

Department of Computer Science, University of Leipzig, Germany

To appear in: Proceedings of the 1997 Fall Symposium (formerly SCAMC) of the American Medical Informatics Association, Nashville, TN, Oct. 1997.

A well-known problem of current electronic patient records is that they usually fail to represent the semantic relationships between the involved clinical data. This has to be viewed as a problem especially in the domains characterized by a complex and long-term treatment, as the medical decision making process may not be comprehensible anymore from the data entries themselves. Context representation can overcome these limitations, enabling the record to express causality, revisions, conflicts, or individual heuristics explicitly. This article introduces CLINICON which is a formal framework for domain-independent context representation based on Sowa's conceptual graphs.

INTRODUCTION

In general, the term *context* indicates that a thing, object, event, or activity is related to its background, circumstances or neighborhood, for a better understanding or improved reasoning about it. In well-written paper-based clinical records and narratives that are in the tradition of the problem-oriented record [1], entries and statements such as laboratory values, findings, conclusions, impressions, diagnoses, and therapeutic activities seldom occur without any reference to their particular medical context. This may be of a causal or temporal nature (e.g., the justification of a diagnosis or medication; a significant trend of a laboratory value), or may express conflict situations, substitutions and revisions, intra-patient similarity patterns, or patient-specific heuristics. Although context representation is essential especially in complex domains with long-term treatment (e.g., oncology, metabolism disorders) in order to express the medical decision process, most electronic records neglect the explicit and structured representation of contexts, since they are able to store the incorporated data, events and actions, but not the relationships that group or cluster the data together into contexts. Work already done in this area (e.g., Bell et al. [2], Barrows et al. [3]) restricts context representation to special causal contexts, or concentrates on temporal aspects (e.g., Shahar [4]). Therefore, and in concordance with the postulations of Rector et al. [5], this article suggests the formal framework *CLINICON* (=Clinical Contexts) for an explicit context representation in patient records. The framework, based on Sowa's conceptual graphs [6], is a formalization and generalization of a former semi-formal context approach [7], which has been applied to clinical progress notes in the domain of pediatric oncology within the knowledge-based system THEMPO [8]. Furthermore, the approach overcomes several limitations of the former one by supporting, for instance, enhanced context granularity, hierarchical contexts, and intra-patient similarity clusters. Local

adaption of the framework to a particular subdiscipline can be achieved, for example, through inheritance from concepts of the framework, or, in general, by graph specialization.

CONTEXTS IN CLINICAL CARE

The occurrence of context information can be illustrated best by examples. The first example is a report describing an oncological case that contains several contexts (context-oriented words are in *italics*, the context type is inserted in []-brackets):

On 3 Feb. 1997, a five year old child with fever, abdominal complaints, and a palpable abdominal mass was presented. Sonography and CT showed a large abdominal tumor ($\emptyset = 9$ cm) with involvement of the left kidney. *Therefore* [-> causal context] the diagnosis of Wilms' tumor was made on 5 Feb. 1997 (the alternative diagnosis, neuroblastoma, was considered unlikely, as the catecholamine metabolite finding¹ in urine and serum was normal). Then, the pre-surgical chemotherapy of the Wilms' protocol (ACT-D, VCR) was ordered, and begun on 6 Feb. 1997². *However*, weekly sonography and CT on 3 March 1997 showed a non-regressive tumor without any tendency to Wilms'-specific morphological changes during chemotherapy. This situation was viewed as being *inconsistent* with the applied chemotherapy and the diagnosis of Wilms' tumor [-> conflict context]. *Therefore* [-> causal context] the diagnosis of Wilms' tumor *was replaced* [-> revision context] by the diagnosis of neuroblastoma, and a surgery was ordered and carried out on 5 March 1997. A final histological examination confirmed the diagnosis of a neuroblastoma (showing unusual extents and a very atypical catecholamine metabolite behavior).

In addition to these three context types indicated in the case example (causal, conflict, revision), some other context types occur frequently in medical records too:

Temporal Contexts

- In the interval from 12 March '97 up to 26 March '97, the patient showed a seriously decreasing creatinine clearance trend (150 ml/min -> 70 ml/min) (indicating a severe renal failure).

Intra-Patient Similarity Contexts

- The infection symptoms of the patient and the infection progress during this chemotherapy *are very similar* to his infection observed during the

1. an important tumor marker for neuroblastoma
2. pre-surgical chemotherapy is performed to make the tumor operable at all (by forcing it to become more solid)

chemotherapy from 5 May - 8 May.

Patient-Specific Heuristics (with integrated causal context):

- In April '96, a patient with insulin-dependent diabetes mellitus complained about extreme tiredness after an infection of the respiratory tracts. This was first explained by the infection and the diabetes itself. However, three weeks later, a pericard effusion was detected, and therefore, together with serodiagnostics, the diagnosis of a post-infectious, atypical pericarditis, induced by the infection of the respiratory tracts, was made. In Dec. '96, the patient again had an infection of the respiratory tracts with pericardial involvement, indicating a strong disposition of the patient to pericarditis. *Therefore*, the patient-specific heuristic was made that whenever this patient again has a respiratory tracts infection, controlling cardio-diagnostics should be carried out.

Contexts like these may have been established by the automated reasoner of a knowledge-based system or by the physician. Their explicit representation enables the record to give unambiguous information about the medical decision process (causal and revision contexts), to draw the physician's or machine agent's attention to data entries which do not exactly fit into the current view of the patient's situation (conflict contexts), and to support the decision process by providing knowledge that is highly patient-adapted, valid only for one particular patient (intra-patient similarity contexts, patient-specific heuristics).

CONCEPTUAL GRAPHS

Conceptual graphs, first introduced by Sowa [6], have been applied to several medical domains (e.g., Baud et al. [9], Bell et al. [2], Campbell et al. [10], Johnson [11], Bernauer et al. [12]). They combine the representational and inferential power of first-order predicate logic with a graphical and natural language-oriented notation. A conceptual graph is a bipartite, directed graph consisting of *conceptual* nodes (denoted as boxes), which are connected with *relation* nodes (denoted as circles). In the alternative linear notation, conceptual nodes are written within []-brackets while relation nodes are denoted within ()-brackets. A conceptual node consists of a type information, and, optionally and separated by a colon, a referent specifying individuals, instances, linguistic articles, quantifiers, plural nouns, or sets (the latter by using {}-brackets within the referent), e.g., [Diagnosis: Endocarditis Lenta], or [Symptom: {Fever, Abdominal Pain}]. The term "{*}" denotes a set of zero or more elements, additional cardinality constraints can be expressed, for example, by "{*}@5" (set of 5 elements) or "{*}@>4" (set of more than 4 elements). Relation nodes consist only of a type information. Both type classes are ordered in a type lattice showing their inheritance relationships. A concept which may have one or more graphs as its referent is called a *context* [13] (which, at first, is nothing more than a syntactic construct, and has to be distinguished from a *clinical* context). Through contexts, conceptual graphs achieve the power of first order log-

ics and can be easily extended to modal logics by linking monadic relations such as (**Believe**) to context concepts. In the following, the context construct of the conceptual graph calculus will be used to represent *clinical* contexts.

To define the basic structure and selectional constraints of the graphs representing domain knowledge and facts, so-called *canonical* conceptual graphs must be defined. Together with the common set of the inference-compliant graph operations *copy*, *simplify*, *restrict*, *join*, *unrestrict*, and *detach*, canonical graphs form a domain-specific *graph-grammar*, i.e. a graph is valid if it can be derived from a set of canonical graphs using a combination of the operations above. The framework introduced here therefore mainly consists of a set of canonical conceptual graphs (CCG) for clinical context representation. Instances of these CCGs then appear as referents in the context concepts shown in fig. 1.

THE CLINICON CONCEPTUAL GRAPH FRAMEWORK FOR CONTEXTS

Causal Contexts

Principally, causal contexts represent the patient-specific rationale of the medical decision making process, i.e. the reasons why something has been done and why other actions have been rejected. The most general causal CCG is:

[Do]- (1)
(Agnt)->[Medical Agent: {*}@>0]
(Ptnt)->[Medical Activity: {*}@>0]
(Because Of)->[Medical Event: {*}]

It describes that a medical agent (e.g. a physician) is doing or has done³ (e.g., concluded, ordered, or prescribed) a number of medical activities *because of* a number of medical events (the (Ptnt)-relation points to the accusative in the sentence; local, non-context information of the conceptual nodes - such as the date or time of the event, is omitted in this CCG; see figure 1 for the local attributes of events). If a medical activity is based on events of different subtypes of **Medical Event** or several activities of different subtypes of **Medical Activity** are done, this can be represented by a graph joining several graphs of type (1).

In *CLINICON*, the most frequent specializations of these CCGs are diagnostic interpretations - see (2) - and the ordering of diagnostic or therapeutic procedures - see (3). In the following, all finding or diagnosis nodes of the canonical context graphs should be viewed as *contractions* of more complex subgraphs representing the details of a finding, diagnosis or activity in the granularity of the specific subdiscipline⁴. *CLINICON*

3. The specification of the particular tense of the represented fact is neglected here as this can be done by linking a monadic relation such as (**Past**) to the graph, or as the tense implicitly is given by the time information linked to every **Medical Event** node.
4. In particular, it depends on the subdomain whether to model a diagnosis such as Myopericarditis as a conceptual subtype of **Diagnosis**, or via the referent in a **Diagnosis** node, as done in the examples here.

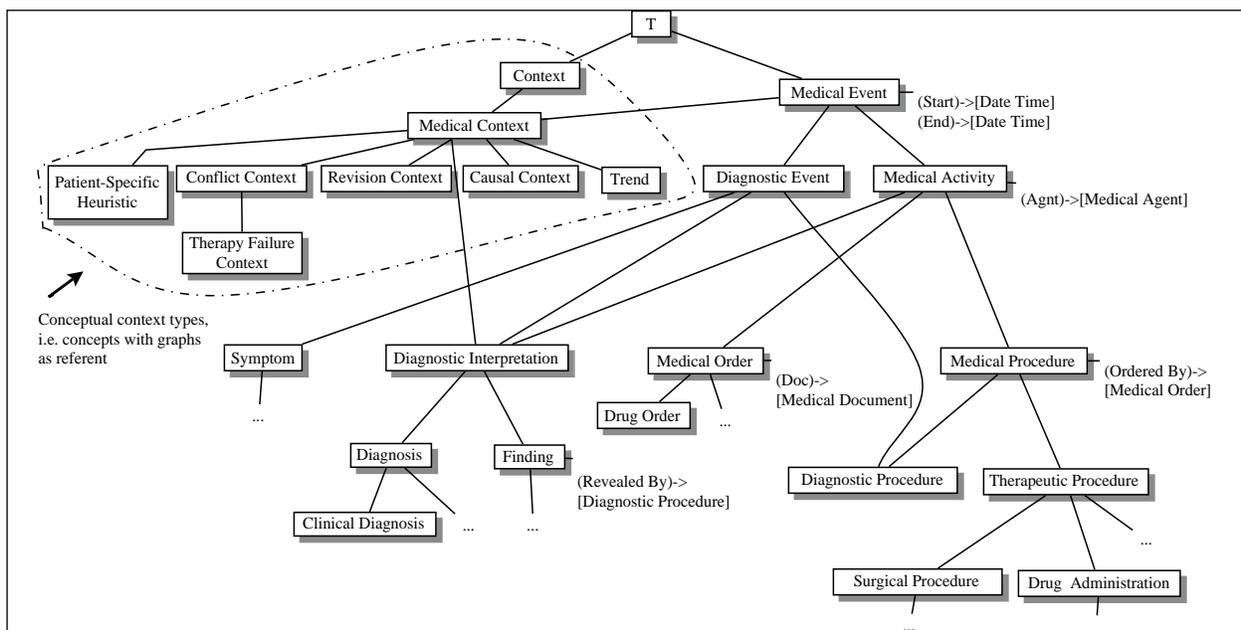


Figure 1: Conceptual type lattice. The right part is partially based on Johnson [11] and Bell [3]. The left part introduces the context types which group instances of **Medical Event** together into clinical contexts (the **Therapy Failure Context** is viewed as a subtype of a **Conflict Context**, as a therapy failure usually is a situation contradicting the current diagnosis). A concept inherits all relations of its ancestors. Dots indicate sublattices which lie beyond the scope of this approach and have to be specified by the subdiscipline. The time information of **Medical Event** may collapse to **(Ptim)->[Date Time]** or **(Ptim)->[Date]**, if the interval has length 0 (Ptim = point in time).

aims to provide conceptual structures for semantic relationships between medical events, while the modeling of the concepts of the particular discipline lies beyond the scope of this framework. Furthermore, the agent is omitted in the following graphs:

[Conclude]- (2)
(Ptnt)->[Diagnostic Interpretation: { * } @>0]
(Because Of)->[Diagnostic Event: { * }]
 e.g., as a result of joining three instances of (2):

[Conclude]- (2')
(Ptnt)->[Diagnosis: Bacterial Myopericarditis]
(Because Of)->[Symptom: Col{Tachypnoea, Retrosternal Pain, Fever}]
(Because Of)->[Hematological Finding: Increased Erythrocyte Sedimentation Rate]
(Because Of)->[Sonography Finding: Col{Pericardial Effusion, Cardiomegaly}]

[Order]- (3)
(Ptnt)->[Medical Order: { * } @>0]
(Because Of)->[Medical Event: { * }] e.g.,

[Order]- (3')
(Ptnt)->[Drug Order: Streptomycin]
(Because Of)->[Diagnosis: Endocarditis lenta]
(Because Of)->[Microbiological Finding: Streptococcus viridans positive]

or, with **Discontinue** as a subtype of **Order**:

[Discontinue]- (3'')
(Ptnt)->[Drug Order: Streptomycin]
(Because Of)->[Hematological Finding: Serious Thrombopenia]

As a simplification the following contraction for (1) can be used:

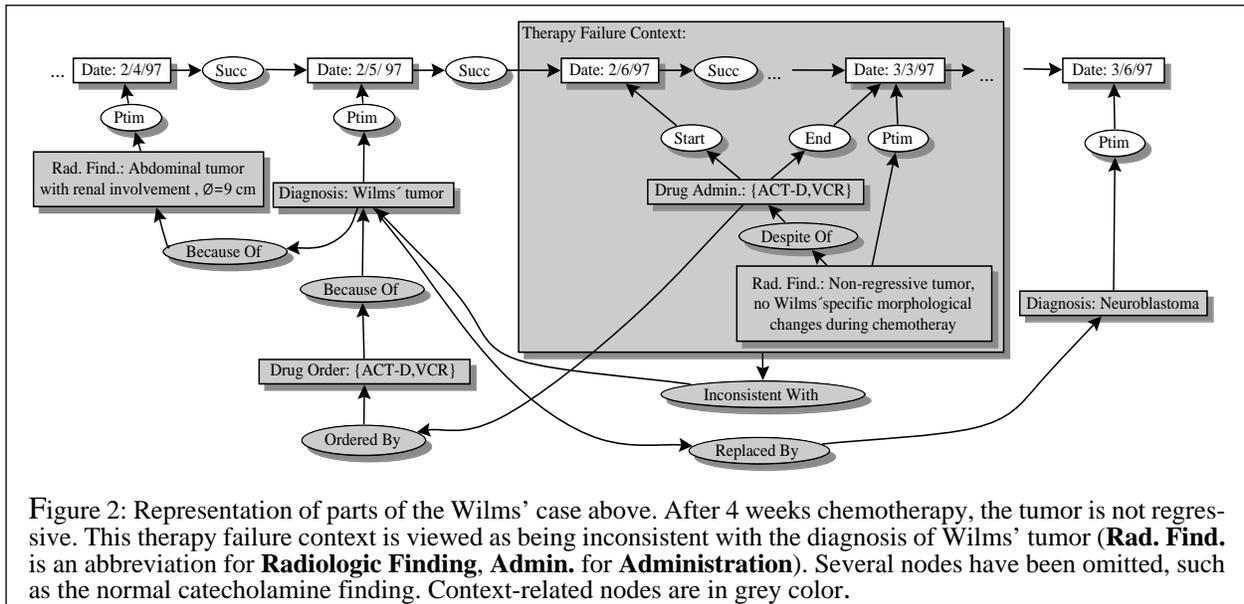
[Medical Event: { * } @>0]- (4)
(Because Of)->[Medical Event: { * }]

This simplification is useful, when the agent of a medical activity can be derived from other data, or if the verb **DO** is superfluous as the special subtype of the verb **DO** can be derived from the particular subtypes of the involved **Medical Event** nodes. Furthermore, a modification of (4) can be used when there is no active agent involved in the context, i.e. when a medical event happens because of (is caused by⁵) other medical events without any intervention of the medical staff:

[Conclude]- (5)
(Ptnt)->[Diagnostic Interpretation: [Medical Event: Anaphylactic Shock]- (Start)->[Date Time: 7/24/96 09:50] (Caused By)->[Drug Administration: L-Asparaginase]] // end of interpretation node

which represents that the administration of L-

5. for this special type, **(Caused By)** is used to express the slightly different semantics of causality



Asparaginase is interpreted as being the causal reason for the patient's shock event (as shown in fig. 1, **Diagnostic Interpretation** inherits from the type **Medical Context**, and therefore may have a graph as its referent).

Revision Contexts

Revision contexts cover the revision of diagnostic interpretations or the replacing of therapeutic procedures. The basic CCG is:

- [**Replace**]- (6)
 (Ptnt)->[**Medical Activity: {*}@>0**]
 (By)->[**Medical Activity: {*}@>0**] with
- [**Medical Activity: {*}@>0**]- (6')
 (**Replaced By**)->[**Medical Activity: {*}@>0**]

as a shortcut. Revision contexts often are combined with causal contexts to justify the revision, e.g. (as a combination of (4) and (6')),

- [**Revision Context:** (6'')
 [**Drug Order: Doxycyclin**]-
 (**Replaced By**)->[**Drug Order: Amphotericin B**]
]-(**Because Of**)->[**Diagnosis: Systemic Mycosis**]

to represent that an antibiotic drug was replaced by a antimycotic drug because of the (new) diagnosis of a serious fungal infection (which itself may have replaced a former diagnosis).

Conflict Contexts

Conflict contexts express that data are viewed to be inconsistent or that medical activities may raise conflicts. Although conflicts usually are resolved after some time (by having more information about a patient), they should be represented for a better understanding of the decision process. The CCG is:

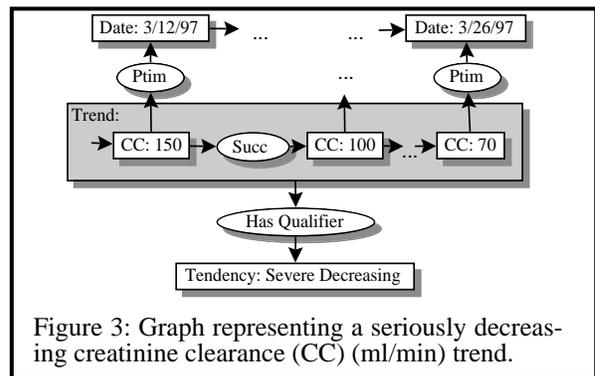
- [**View As Inconsistent**]- (7)
 (Ptnt)->[**Medical Event: {*}@>0**]
 (**With**)->[**Medical Event: {*}@>0**]

with the shortcut (see fig. 2 for an example):

- [**Medical Event: {*}@>0**] (8)
 (**Inconsistent With**)->[**Medical Event: {*}@>0**]

Temporal Contexts

The basic structure of a temporal context is shown in fig. 3. Temporal contexts are frequently embedded into causal contexts, e.g., the creatinine clearance trend in fig. 3 could be used as a causal justification of a finding such as "Serious renal toxicity of current chemotherapy".



Intra-Patient Similarity Contexts

In long-term records, event patterns and clinical contexts of a patient may recur several times in a similar way, expressing for example patient-specific dispositions to infections or other complications. The CCG is:

- [**Medical Context**]- (9)
 (**Similar To**)->[**Medical Context**]

Similarity contexts are usually embedded within larger contexts which establish patient-specific heuristics (see

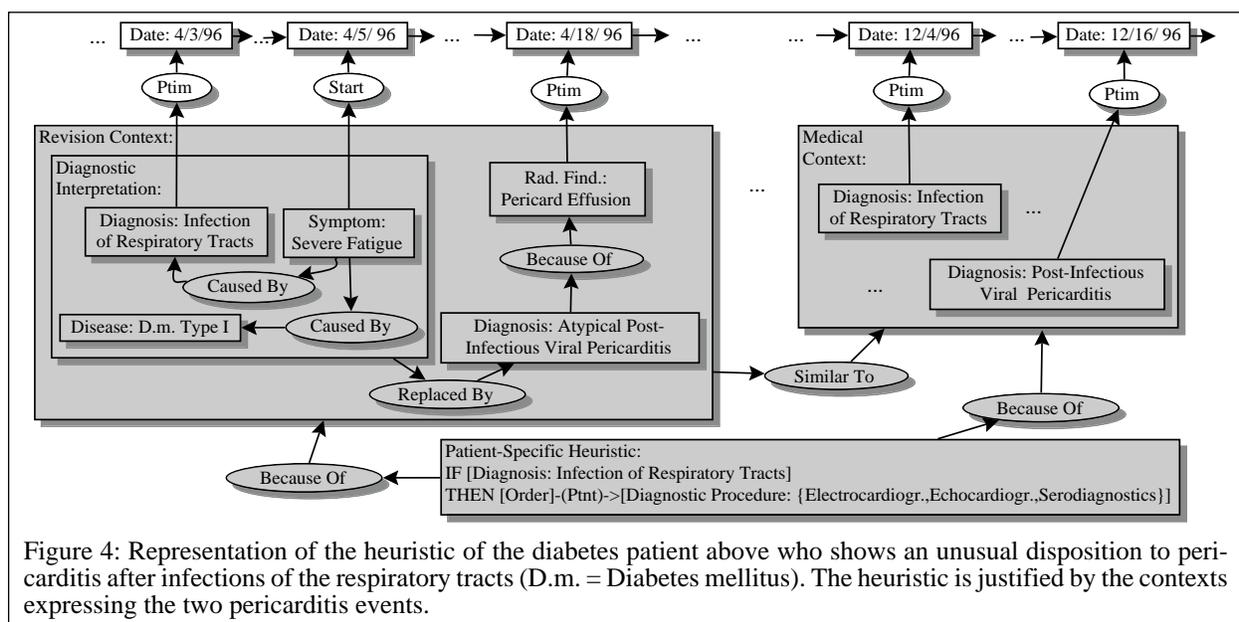


Figure 4: Representation of the heuristic of the diabetes patient above who shows an unusual disposition to pericarditis after infections of the respiratory tracts (D.m. = Diabetes mellitus). The heuristic is justified by the contexts expressing the two pericarditis events.

fig. 4 for an example).

Patient-Specific Heuristics

Patient-specific heuristics are rules expressing knowledge which is valid only for a particular patient, and which usually evolved over a longer time period reflecting very individual experiences with a patient. According to the rule notation in conceptual graphs, the CCG is (see fig 4 for an example):

[Patient-Specific Heuristic: (10)
IF [Medical Event: {*}@>0]
THEN [Medical Activity: {*}@>0]]

During treatment of the patient, his or her heuristics can be linked to the knowledge base containing rules valid for all patients, and can then be processed by the rule interpreter.

IMPLEMENTATION ISSUES

The framework core has been prototypically implemented on the basis of the CONCEPTUAL GRAPH KNOWLEDGE ENGINEERING ENVIRONMENT (CGKEE). (<http://turing.unne.edu.au/~cgtools/CGKEE.html>) of the University of New England, Australia. Current efforts concentrate on integrating the context core into a hematological patient record of the University Hospital Leipzig.

DISCUSSION

This article introduced *CLINICON*, a framework for representing clinical contexts in patient records. It has been shown that contexts occur frequently during treatment and that especially in long-term domains (such as oncology and metabolism disorders) context representation is necessary for the medical decision process to be made transparent. The overall goal has been to provide a domain-independent, formal, and computable framework which can be adapted to a particular subdiscipline by graph specialization. Future efforts will concentrate on report generation from

context graphs (which is directly supported by conceptual graphs, as they are a representation “close” to natural language), automated context acquisition by mapping contexts from problem-solvers to the patient graph, and data mining aspects (e.g., can knowledge about diagnostic failures be derived from revision contexts such as the one in the Wilms’ case?).

Acknowledgments

The author thanks Dr. med. H.-M. Dittrich for giving valuable advice concerning the medical examples.

References

1. Weed L. Medical Records, Medical Education, and Patient Care: The Problem-Oriented Record as a Basic Tool. Chicago: Year Book Medical Publishers, 1971.
2. Barrows RC, Johnson SB. A data model that captures clinical reasoning about patient problems. Proc Symp Comp Appl Med Care, 1995: 402-405.
3. Bell DS, Pattison-Gordon E, Greenes, RA. Experiments in concept modeling for radiographic image reports. JAMIA, 1994:1(3):249-262.
4. Shahar Y, Musen MA. Knowledge-based temporal abstraction in clinical domains. Artificial Intelligence in Medicine, 1996:8:267-298.
5. Rector AL, Nowlan W, Kay S, Goble C, Howkins IJ. A framework for modelling the electronic medical record. Meth Inform Med, 1993: 32: 109-119.
6. Sowa JF. Conceptual Structures. Reading, MA: Addison Wesley, 1984.
7. Müller R, Thews O, Rohrbach C et al. A graph-grammar approach to represent causal, temporal and other contexts in an oncological patient record. Meth Inf Med, 1996:35(2):127-141.
8. Müller R, Serfl M, Nauerth U et al. TheMPO: A knowledge-based system for therapy planning in pediatric oncology. Comp Bio Med, 1997:27(3): 177-200.

9. Baud RH, Rassinoux AM, Wagner JC et al.: Representing clinical narratives using conceptual graphs. *Meth Inf Med*, 1995;34(1/2):176-186.
10. Campbell KE, Das AK, Musen MA. A logical foundation for representation of clinical data. *JAMIA*, 1994;1:218-232.
11. Johnson, SB. Generic data modeling for clinical repositories. *JAMIA*, 1996;3:328-339.
12. Bernauer J, Taphorn FJ, Pretschner DP. A model for composing medical descriptions from fragments. *Proc Medinfo*, 1992:1395-1401.
13. Esch J, Levinson R. An implementation model for contexts and negation in conceptual graphs. *Proc 3th Int Conf Conceptual Structures*, 1995: 247-263.