Towards a Second Generation of Computer Interpretable Guidelines

Paolo Terenziani, Alessio Bottrighi, Laura Giordano, Giuliana Franceschinis, Stefania Montani, Luigi Portinale and Daniele Theseider Dupre
DISIT - Computer Science Institute, University of Eastern Piedmont, Alessandria, Italy
{alessio.bottrighi, laura.giordano, giuliana.franschenis, stefania.montani, luigi.portinale, paolo terenziani, daniele.theseider}@unipmn.it

Keywords: Computer Interpretable Guidelines, Integration of Different Forms of Medical Knowledge, Logical, Analogical, Defeasible and Probabilistic Reasoning, Temporal Data.

Abstract: Computer Interpretable Guidelines (CIG) are an emerging area of research, to support medical decision making through evidence-based recommendations. However, new challenges in the data management field have to be faced, to integrate CIG management with a proper treatment of patient data, and of other forms of medical knowledge (e.g., causal and behavioral knowledge). In this position paper, we summarize a proposal for a research agenda that, in our opinion, can lead to a significant advancement in the field. The goal of the work is to provide suitable models and reasoning methodologies to cope with the aforementioned aspects, and to properly integrate them for medical decision support. Achieving such a goal requires advances in data management, and, in particular, in the treatment of indeterminate valid-time data in relational databases, of temporal abstraction on time series, of case retrieval on time series, of design-time and run-time model-based verification of guidelines, of case-based reasoning, of non-monotonic logics, of formal ontologies, of probabilistic graphical models (Bayesian Networks and Influence Diagrams).

1 INTRODUCTION

This position paper proposes and summarizes a research agenda we are going to present as a four-year project proposal for an ERC Advanced Grant (http://erc.europa.eu/advanced-grants).

Context. A lot of attention has been recently devoted to the analysis of clinical processes, and thus to clinical guidelines, which constitute the tools identified by physicians in order to encode the "best practice" clinical procedures. Clinical guidelines are - as defined by US Institute of Medicine - "systematically developed statements to assist practitioner and patient decisions about appropriate health care in specific clinical circumstances". In the last years, thousands of clinical guidelines have been developed by local, national, and international organizations. Despite such a large effort, it is widely recognized that clinical guidelines have not provided all the expected advantages in the clinical practice. The recent research in medical informatics has widely demonstrated that Computer Science can help to drastically improve the impact of clinical guideline (see, e.g., (Teije et al., 2008)). For such reasons, research in this sector is becoming more and more important in the area of Medical Informatics, and of Computer Science in general. In particular, in the last twenty years, several approaches have been developed to deal with Computer Interpretable Guidelines (henceforth: CIG; see, e.g., the comparisons in (Peleg et al., 2004), and the book (Teije et al. 2008)). As a matter of fact, a wide scientific literature in the last twenty years have demonstrated that, besides having an important practical impact, the computer-based treatment of clinical guidelines arise several challenging open-problems for the scientific community, and that much more theoretical and practical efforts must be spent to achieve all the potential impact on the clinical practice.

Objectives. The "visionary" idea underlying our research proposal is that, to achieve a significant step forward in the state-of-the-art, it is important to provide a homogeneous approach integrating (at least) three different aspects:

(1) Computer Interpretable Guidelines

(2) Treatment of patient data

(3) Treatment of "basic" medical knowledge (BMK)

Such aspects have been often considered in isolation by the Medical Informatics literature and a step forward in data management should be achieved in order to integrate them.

Regarding CIGs, given the dimension of clinical guidelines, their correctness and their adherence to specifications cannot be verified by-hand by physicians. Both *design-time* and *run-time verification* must be considered to cope with the different phases of guideline life-cycle. Thus, new formal methodologies must be identified. Different models, with different features like, *logical* and *Petri-net-based* models, are definitely good candidates for model-based verification.

Of course, CIGs must be executed considering patients' data. Supporting a proper treatment of such data is still a challenging open issue since, for instance, difficult temporal issues have to be managed (such as, e.g., *temporally indeterminate* patient data). Additionally, in many cases, raw data cannot be directly used in guidelines. *Time series* must be managed and retrieved, usually at different levels of *abstraction.* Clinical guidelines capture medical evidence, thus coping with "typical" patients, and not with patients' peculiarities. However, patient data may be such that the patient is somehow "not typical". In such cases, physician may need to resort to other kinds of BMK, and other forms of reasoning on it, including medical *ontologies*, knowledge from *experience* and *analogical* reasoning (e.g. case-based reasoning), *causal*, *behavioral* and uncertain (probabilistic) knowledge, and knowledge about patient *evolution*.

However, although coping separately with each one of the above aspects is a challenging goal of our project, our main goals lie in the definition of appropriate models and methodologies to cope with their interactions. In particular, the main challenge concerns the interactions between CIG (which are executed considering patients data) and BMK. Indeed, for specific patients, the knowledge in the guideline and the BMK may suggest different (and possibly contrasting) actions, and only domain experts seem able to choose between them, on the basis of the patient's status.

Example 1. Clinical Guideline: Patient with acute myocardial infarction presenting with acute pulmonary edema; before performing coronary angiography it is mandatory to treat the acute heart failure.

BMK: The execution of any clinical guideline may be suspended, if a problem threatening the patient's life suddenly arises.

Example 2. Clinical Guideline: In a patient affected by unstable angina and advanced predialytic renal failure, coronary angiography must be executed.

BMK: the contrast media administration may cause a further final deterioration of the renal functions, leading the patient to dialysis (see (Bottrighi et al., 2009) for more examples).

In Example 1 the execution of a CG is suspended, due to the presence of a problem threatening the patient's life. In Example 2 instead the treatment is performed even if it may involve contrast media administration, which may dangerous for the patient. This example shows that not only some guideline's prescriptions are "defeasible", since they may be overridden by BMK, but the same also holds for part of BMK. Thus, defeasible (non-monotonic or probabilistic) reasoning may be applied for the integration. Indeed, example 2 involves also a form of cost/benefit analysis, which, in turn, needs to be supported by physicians' knowledge about typical patient evolution.

Providing physicians with decision-support tools considering these different forms of knowledge and reasoning about their interactions is the ultimate goal of our research vision.

Expected Results. The project aims to enhance the state-of-the-art in several areas of research, including: the treatment of indeterminate valid-time data in relational databases, temporal abstraction on time series, case retrieval on time series, design-time and run-time model-based verification of guidelines, case-based reasoning, non-monotonic logics, formal ontologies, probabilistic graphical models (Bayesian Networks and Influence Diagrams). Prototypes will also be developed and tested by physicians, in order to demonstrate the practical feasibility and the usefulness of the proposed methodologies.

Methodology. The starting point of our proposal is GLARE, a prototypical tool for clinical **G**uide**L**ine **A**cquisition, **R**epresentation and **E**xecution (Terenziani et al., 2001). GLARE already interacts with patient data for CIG execution, and provides physicians with different advanced facilities, based on formal Temporal Database and Artificial Intelligence techniques (Terenziani et al., 2001; (Anselma et al., 2006); (Terenziani et al., 2007); (Terenziani et al., 2008); (Leonardi et al., 2012).

The idea is to deeply extend GLARE, by considering an agenda structured as shown in the following figure 1. In the rest of the paper, each one of the tasks in figure 1 will be briefly addressed.

1. **Clinical guideline verification**
2. **Modeling patient clinical data**
   2.1 **Valid-time data**
   2.2 **Temporal abstraction and time series data**
   2.3 **Querying data and data retrieval**
3. **Basic medical knowledge (BMK)**
   3.1 **Medical Ontologies**
   3.2 **Experence-based knowledge and past cases**
   3.3 **Causal and behavioral knowledge**
   3.4 **Patient evolution**
4. **Integration**
5. **Prototypes**

Figure 1: Structure of the research proposal.

2 CIG VERIFICATION

One of the main obstacles for a wide applicability of clinical guidelines is their complexity and, as a consequence, the difficulty in defining a semantically accurate representation, and in verifying their properties (such as consistency). This motivates the use of standard verification tools for the verification of clinical guidelines. In the project, at least two different types of verification tools will be explored, and compared: logical-based tools, such as model checkers and theorem provers, and Petri-Net based approaches.

The logical-based verification requires, on the one hand, the translation of the guideline in the specification language of the model checker and, on the other hand, the encoding of the guidelines properties as temporal formulas. This approach has been pursued, for instance, in (Giordano et al., 2006); (Bottrighi et al., 2010), by providing a translation of GLARE clinical guidelines into the language Promela of the LTL model checker SPIN. In this direction, the proposal aims at exploiting the logical language Temporal ASP (Giordano et al., 20120, for the specification and the verification of clinical guidelines. Temporal ASP allows for the definition of temporal constraints in a domain specification, as Dynamic Linear Time Temporal Logic (DLTL) formulas. Such constraint could be used, together with ASP rules, for a declarative specification of CIGs. In particular, to do so, a mapping of GLARE constructs into Temporal ASP has to be defined. The verification of runtime and design time properties of the CIGs formulated in Temporal ASP can then be performed by DLTL Bounded Model Checking. The runtime verification that the treatment of a given patient is compliant with the guideline can be modeled as a satisfiability problem. Instead, the verification of temporal properties of the guideline can be modeled as validity checks.

Petri Net (PN) models and their extensions have been proposed as a flexible and general formalism that can be used as a common semantics for different application-oriented modeling languages. In the literature the feasibility of the (semi)automatic translation of CIG models into PNs has been already illustrated (Beccuti et al., 2009); (Quaglini et al., 2001); (Peleg et al., 2005): the resulting model may be used for verification of both qualitative and quantitative properties. In this context, our research work will focus on one hand towards extending the compositional translation rules to represent the relevant steps in each CIG that produce a request for resources at given points in time; in other words the goal is to extract a workload stochastic model from CIGs represented with GLARE. Multiple instantiations of sets of similar CIG models can be composed in a compact way through High Level Stochastic PNs (in particular Well-Formed Nets, leading to efficient analysis methods): such a model can then be merged with another one representing the available resources and the scheduling constraints to perform what-if analysis on alternative ways of organizing healthcare services. In addition the most suitable (logic or automata based) language for expressing the non functional properties will be selected, and tested on case studies. We plan to define both system oriented performance properties (useful for the evaluation of costs and resources utilization) and user oriented performance properties (useful to evaluate customer satisfaction indicators). The experiments will be performed mainly through the GreatSPN tool (Baarir et al., 2009).

3 MODELING PATIENT CLINICAL DATA

This part of the project focuses on the treatment of patient clinical data, with specific emphasis on temporal issues.

3.1 Valid-time Data

In the medical domain, as well as in many other real-world domains, often the exact time of occurrence of facts is not known and temporal indeterminacy (i.e., “don’t know when”) occurs. Temporal indeterminacy is widespread within medical data (consider, e.g., patient complaint in Ex.1).

Ex.1 On January 1st 2012 Mary had headache starting between 8am and 9am and ending between 1pm and 2pm.

However, in the area of (relational) databases, the treatment of temporal indeterminacy has been quite neglected (see the survey in (Dyreson, 2009) and (Anselma et al., 2012)). We aim at overcoming such a limitation by proposing (i) a new data model, to model temporally indeterminate relational data, and (ii) a new relational algebra, to query it. To guarantee implementability on top of current technologies and interoperability, we will prove that our data model is a consistent extension of TSQL2's one (which, in turn, is a consistent extension of SQL) (Snodgrass et al., 1995), and our algebra is reducible to the standard relational one. Also, we plan to study the interactions between temporal indeterminacy and telicity, which is important to characterize patient data (Terenziani et al., 2007).

3.2 Temporal Abstractions and Time Series Data

In medical applications, most patient data are naturally collected in the form of time series. Within the project, we aim at developing a user-friendly, interactive and flexible support for querying and retrieving time series data based on Temporal Abstraction (TA) (Shahar, 1997); (Bellazzi et al., 1998) for dimensionality reduction. Specifically, we aim at supporting multi-level abstractions of the original data (Montani et al., 2013). We foresee to abstract (and query) time series data at finer or coarser detail levels, according to two dimensions: (1) a taxonomy of symbols, and (2) a taxonomy of time granularities. Dimension (1) refers to the qualitative detail level we are interested in. As an example, referring to trend TA, we might want to identify intervals of increasing trend, or be more specific, by distinguishing between intervals of strongly increasing trend, and weakly increasing trend. As regards dimension (2), we could be very specific, and abstract intervals expressed at a "small" time granularity (e.g. minutes – depending on the application), or be coarser, and just look for the abstract behavior on a "larger" time scale (e.g. hours). We plan to convert our raw data into abstracted sequences of symbols, calculated at the finest level of detail according to both dimensions. However, to support flexible retrieval, users will be allowed to query the database according to all levels, in both dimensions.

3.3 Querying Data and Data Retrieval

The data structures (i.e. the taxonomies) and the abstraction and distance functions studied in 3.2 will be exploited to implement a flexible and efficient retrieval strategy. Namely, we plan to resort to proper index structures, able to support early pruning, defined on the basis of the symbol and time granularity taxonomies. The goal is to allow queries to be issued at any level of detail, according to all defined dimensions. A second goal will be to answer such queries efficiently both in terms of time and space. User-flexibility will also be an important issue, since we aim at: (1) allowing users to express queries in the form of regular expressions; (2) allowing users to query for sub-series.

4 BASIC MEDICAL KNOWLEDGE (BMK)

Different forms of BMK have to be investigated.

4.1 Medical Ontologies

Ontological knowledge is a relevant part of BMK. In particular, formal ontologies are important to provide a standard terminology with a clear formal semantics, to be used in the representation of all the forms of knowledge (including GIG content). A specific challenge here is that of defining tractable non-monotonic extensions of such DLs, to capture prototypical properties of concepts and inheritance with exceptions. To capture the degree of "fuzziness" of medical knowledge, we will make innovative proposals related to fuzzy and probabilistic ontologies (Klinov and Parsia, 2008) and ontologies based on different types of logics, including, e.g., preferential logics (Giordano et al., 2009).

4.2 Experience-based Knowledge and Past Cases

Maintaining patient data as cases allows the decision maker to access at least in implicit form the experience-based knowledge contained in such cases. Case-Based Reasoning (CBR) (Aamodt and Plaza, 1994) transforms such implicit chunks of information into more effective and explicit knowledge for decision making (Schmidt et al., 2001). A specific goal of our proposal will then be to integrate our treatment of time series (based on temporal abstraction and multiple abstraction levels), into the CBR loop, in such a way that both the relevant static knowledge about patients, as well as dynamic temporal information can be integrated with the objective of a final optimal decision. Furthermore, CBR has proved to be well suited for managing exceptional situations which can be neither foreseen nor preplanned. In the literature cases have often been resorted to describe exceptions, in various domains, including medical domains (Bellazzi et al., 1999). In the management and the interaction of patient data and guidelines, we aim at exploiting a CBR approach to support guideline exception management and adaptation at the local environment characteristics.

4.3 Causal and Behavioural Knowledge

Modeling BMK definitely implies to have the capabilities of representing, at some level of details, basic medical, clinical processes and cause-effect associations. In the project, we will explore different kinds of formalisms to properly model BMK. Modeling BMK is a challenging goal, since, e.g., it requires the ability to reason about incomplete and defeasible knowledge. Furthermore, the properties of a given patient can be incompletely specified and, nevertheless, we may want to derive conclusions concerning that patient. One of the current trends of modeling defeasible knowledge, that will be explored in the project, is represented by Answer Set Programming (ASP) (Gelfond, 2007); (Bonatti, 2010), a logic programming language which allows for the specification of defeasible rules under the stable model semantics. ASP allows for a declarative representation of knowledge and the presence of default negation allows for the formalization of rules with exceptions. Preferences among rules can be modeled in ASP (see, for instance the approaches in (Brewka, 2004)). Also, several state of the art ASP solvers are nowadays available for computing answer sets, among which DLV, Smodels, and Clingo.

Among the other non-monotonic approaches, the project will also explore the adoption of the Event Calculus. In particular, REC (Chesani et al., 2009), a reactive axiomatization of Event Calculus, seems to be promising not only in order to capture BMK, but also to study their interaction with guidelines knowledge (Bottrighi et al., 2009).

On the other hand, when uncertainty is a primary issue, Probabilistic Graphical Models (PGMs) (Jensen and Nielsen, 2007) are a reference framework for causal knowledge. There are several points that can be of great relevance in medical decision making. For instance, the level of representation is currently assumed propositional, while in medical setting knowledge about evolutions could require more powerful representation schemes. Also, we aim at investigating both the role of probabilistic logics with first-order semantics (like Markov logic (Domingos and Lowd, 2009), probabilistic Horn abduction (Poole, 1993) or independent choice logic (Poole, 1997)) and that of Relational PGMs (Getoor et al., 2007) in the clinical decision making setting. A further challenge is the possibility of representing and reasoning with continuous variables.

4.4 Patient Evolution

In order to build flexible and useful decision support systems for therapeutic purposes, having knowledge about the patient evolution (i.e. prognostic information) is very relevant. To this extent, many different Probabilistic Graphical Models (PGM) (Koller and Friedman, 2009) have been already developed in the literature. PGMs can be suitably applied to this task. However, there are several points that deserve more attention:

(1) modeling temporal processes is one of the greatest challenges of decision networks. We aim at studying the applicability of non markovian PGMs to the clinical decision making setting. Moreover, we aim at investigating the potentiality of anytime algorithms providing approximate decision-making strategies incrementally refined in order to deal with time- or modeling-pressured situations;

(2) the majority of dynamic models rely on a discrete-time model; A further goal of the present research will be the investigation and the study of suitable inference algorithms for continuous-time PGMs, by putting emphasis also on representational issues relevant for the medical applications.

5 INTEGRATION

Different forms of reasoning will be investigated for integrating the different knowledge sources (and, in particular, CIGs and BMK).

Given the defeasible character of medical knowledge, non-monotonic logical reasoning seems to be a good candidate for the integration. However, the complexity of the problem demands for the exploration of different approaches, and for the definition of new logics. For instance, the project will explore whether a language combining ASP with some form of temporal reasoning (or reasoning about actions) can be suited for the specification of both CIGs and the BMK, and for the integrated reasoning about them. An example of such a language is given by REC (Chesani et al., 2009), a reactive axiomatization of Event Calculus. Another example is given by Temporal ASP (Giordano et al., 2012), a temporal extension of Answer Set Programming (ASP) (Gelfond, 2007), which combines ASP with temporal constraints in LTL.

Besides non-monotonic and defeasible reasoning, also utility-based reasoning will play an important role in the integration since, in many practical situations, decisions concerning patient diagnosis or therapy must be grounded on quantitative cost/benefit information. We aim at resorting to reasoning based on decision theory, by explicitly taking into account uncertainty via probabilistic modeling, combined with utilities of outcomes. Since PGMs will be introduced in the BMK setting, we rely on the features of (dynamic) decision networks to allow clinical guidelines to locally exploit optimal decisions from the underlying network model, when different alternative actions are possible and have to be evaluated by considering both uncertainty in the data or in the evolution and cost/utilities of the outcomes.

Moreover, a specific goal of the proposal is the study of the integration of analogical (Case-Based) reasoning with the above forms of reasoning. Usually, CBR tools are able to extract relevant knowledge, but that leave to the user the responsibility of providing its interpretation and of formulating the final decision. This is because a strict interaction with the BMK has to be established. Another goal of the proposal will be to study such an interaction, both at the most general level, as well as at the level of specific medical applications.

To demonstrate the practical feasibility of our approach, prototypes will be developed, for each specific task in the agenda. We propose to identify a set of case studies, in strict cooperation with the Hospitals and Health Agencies that will take part to the work. Such case studies will constitute the glue to relate the different prototypes, and, thus the different approaches and methodologies developed within the research work.

ACKNOWLEDGEMENTS

The work described in paper was partially supported by Compagnia di San Paolo, in the Ginseng project.

REFERENCES

Aamodt, A., Plaza, E., 1994. Case-based reasoning: foundational issues, methodological variations and systems approaches. *AI Communications* 7, 39–59

Montani, S., Bottrighi, A., 2006. Towards a Comprehensive Treatment of Repetitions, Periodicity and Temporal Constraints in Clinical Guidelines. *AIIM 38*, Elsevier, 171-195,

Anselma, L., Terenziani, P., Snodgrass, R. T., 2012. Valid-Time Indeterminacy in Temporal Relational Databases: Semantics and Representations, *IEEE TKDE*, vol. 99, PrePrints.

Baarir, S., Beccuti, M., Cerotti, D., De Pierro, M., Donatelli, S., Franceschinis, G., 2009. The GreatSPN tool: recent enhancements. *ACM SIGMETRICS Performance Evaluation Review*, Vol.36 Issue 4, 4-9

Beccuti, M., Bottrighi, A., Franceschinis, G., Montani, S. , Terenziani, P. 2009. Modeling and verifying clinical guidelines through Petri Nets*, Proc. AIME'09*, LNAI 5651.

Bellazzi, R., Larizza, C., Riva, A., 1998. Temporal abstractions for interpreting diabetic patients monitoring data. *Intelligent Data Analysis* 2, 97–122.

Bellazzi, R., Montani, S., Portinale, L., Riva, A. 1999. Integrating Rule-Based and Case-Based Decision Making in Diabetic Patient Management, *Proc. ICCBR'99, Munich, LNCS 1650*, 386-400.

Bonatti, P., Calimeri, F., Leone, N., Ricca, F. 2010. Answer Set Programming, chapter 5. *LNCS* Volume 6000 of Dovier and Pontelli,.

Bottrighi, A., Chesani, F., Mello, P., Molino, G., Montali, M., Montani, M., Storari, S., Terenziani, P., Torchio, M., 2009. An Hybrid Approach to Clinical Guideline Conformance. *Proc. AIME'09*, LNAI 5651.

Bottrighi, A., Giordano, L., Molino, G., Montani, S.. P.Terenziani, M.Torchio, 2010. Adopting model checking techniques for clinical guidelines verification, *AIIM 48*(1), 1-19.

Brewka, G., 2004.: Complex Preferences for Answer Set Optimization, *KR 2004*, 213-223.

Chesani, F., Mello, P., Montali, M., Torroni, P., 2009. Commitment tracking via the reactive event calculus. In *Proc. IJCAI 2009*, 91-96.

Domingos P., Lowd, D. 2009. Markov Logic: an Interface Layer for Artificial Intelligence. Morgan & Claypool Publ.

Dyreson, C.. 2009.Temporal Indeterminacy. In L. Liu and M. Tamer Özsu. *Encyclopedia of Database Systems*, Springer.

Gelfond, M., 2007. Handbook of Knowledge Representation, chapter 7, *Answer Sets*. Elsevier.

Getoor, L., Friedman, N., Koller, D., Pfeffer, A., Taskar, B., 2007. Probabilistic Relational Models. In L. Getoor and B. Taskar, editors, *Introduction to Statistical Relational Learning*.

Giordano, L., Terenziani, P., Bottrighi, A., Montani, S., Donzella, L., 2006. Model Checking for Clinical Guidelines: an Agent-based Approach. in *Proc. AMIA'06*, Washington.

Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G. L., 2009. ALC + T: a Preferential Extension of Description Logics. *Fundam. Inform. 96*(3): 341-372.

Giordano, L., Martelli, A., Theseider Dupré, D., 2012. Reasoning about actions with Temporal Answer Sets. In *Theory and Practice of Logic Programming*. 2012, to appear. Available at http://arxiv.org/abs/1110.3672

Jensen, F.V., Nielsen, T.D., 2007. Bayesian Networks and Decision Graphs. Springer.

Koller, D., Friedman, N., 2009. Probabilistic Graphical Models, *MIT Press*.

Klinov, P., Parsia, B., 2008. Optimization and evaluation of reasoning in probabilistic description logic: Towards a systematic approach, in *International Semantic Web Conference*, LNCS 5318, 213-228. Springer.

Leonardi, G. , Bottrighi, A., Galliani, G., Terenziani, P. Messina, A., Della Corte, F., 2012. Exceptions Handling within GLARE Clinical Guideline Framework, *Proc. AMIA 2012, Distinguished Paper Award*.

Montani, S., Leonardi, G., Bottrighi, A., Portinale, L., Terenziani, P., 2013. Supporting Flexible, Efficient and User-Interpretable Retrieval of Similar Time Series. *IEEE TKDE*, ISSN: 1041-4347, doi: 10.1109/TKDE.2011.264

Peleg, M., Tu, S., Bury, J., Ciccarese, P., Fox, J., Greenes, R.A., Hall, R., Johnson, P.D., Jones, N., Kumar, A., Miksch, S., Quaglini, S., Seyfang, A., Shortliffe, E.H., Stefanelli, M., 2004. Comparing computer interpretable GL models: a case-study approach, *JAMIA 10*, 52–68.

Peleg, M., Rubin, D., Altman, R.B., 2005. Using petri nets tools to study properties and dynamics of biological systems. *JAMIA 12*:181–199.

Poole, D., 1993. Probabilistic Horn Abduction and Bayesian Networks. *Artificial Intelligence*, 64(1): 81-129.

Poole, D., 1997. The Independent Choice Logic for modelling multiple agents under uncertainty, *Artificial Intelligence*, 94(1-2): 7-56.

Quaglini, S., Stefanelli, M., Lanzola, G., Caporusso, V., Panzarasa, S., 2001. Flexible guideline-based patient careflow systems. *AIIM 22*:65-80.

Schmidt, R., Montani, S., Bellazzi, R., Portinale, L., Gierl, L. 2001. Case-Based Reasoning for Medical Knowledge-based Systems, *International Journal of Medical Informatics 64*(2), 355-367.

Shahar, Y., 1997. A framework for knowledge-based temporal abstractions. *Artificial Intelligence 90*, 79–133.

Snodgrass, R. T. (Editor), 1995: The TSQL2 Temporal Query Language, Kluwer Academic Publishers.

Ten Teije, A., Miksch, S., Lucas, P., 2008Computer-based Medical Guidelines and Protocols: A Primer and Current Trends, Volume 139 *Studies in Health Technology and Informatics*, ISBN: 978-1-58603-873-1.

Terenziani, P., Molino, G., Torchio, M, 2001. A Modular Approach for Representing and Executing Clinical Guidelines. *AIIM 23*, 249-276.

Terenziani, P. , Snodgrass, R.T., Bottrighi, A., Torchio, M., Molino, G., 2007. Extending temporal database to deal with telic/atelic medical data, *AIIM 39*(2) 113-126.

Terenziani, P., Montani, S., Bottrighi, A., Molino, G., Torchio, M., 2008. Applying Artificial Intelligence to Clinical Guidelines: the GLARE Approach, in (Tejie et al., 2008).