Temporal Representation Design Principles: An Assessment in the Domain of Liver Transplantation

Constantin F. Aliferis, M.D., Ph.D., and Gregory F. Cooper, M.D., Ph.D.

Center for Biomedical Informatics, and the Intelligent Systems Program
University of Pittsburgh, Pittsburgh PA

ABSTRACT

Time modeling is an important aspect of medical decision-support systems engineering. At the core of effective time modeling lies the challenge of proper knowledge representation design. In this paper, we focus on two important principles for effective time-modeling languages: (a) hybrid temporal representation, and (b) dynamic temporal abstraction.

To explore the significance of these design principles, we extend a previously-defined formalism (single-granularity modifiable temporal belief networks – MTBN-SGs) to accommodate multiple temporal granularities and dynamic query and domain-specific model creation. We call the new formalism multiple-granularity MTBNs (MTBN-MGs). We develop a prototype system for modeling aspects of liver transplantation and analyze the resulting model with respect to its representation power, representational tractability, and inferential tractability.

Our experiment demonstrates that the design of formalisms is crucial for effective time modeling. In particular: (i) Hybrid temporal representation is a desirable property of time-modeling languages because it makes knowledge acquisition easier, and increases representational tractability. (ii) Dynamic temporal abstraction improves inferential and representational tractability significantly.

We discuss a high-level procedure for extending existing languages to incorporate hybrid temporal representation and dynamic temporal abstraction.

INTRODUCTION

Time modeling is an *important and challenging* aspect of engineering medical decision support systems (MDSSs) [1]. During the last two decades significant progress has been made in several medical informatics areas with respect to representation of, and reasoning about temporal concepts. The research includes (but is not exhausted by) temporal abstraction and summarization systems, temporal databases, temporal maintenance systems, temporal logics, dynamic Bayesian networks, recurrent neural networks, Markov decision processes and time-specific statistical models (time series, reliability and survival analysis) [1,2].

In previous work we introduced an extension of Bayesian belief networks called single-granularity modifiable temporal belief networks (MTBN-SGs) with the intent to incorporate formal temporal and causal semantics to probabilistic temporal reasoning [3]. In the present paper, we focus on the *implementation and empirical assessment* of two aspects of temporal language design which we had previously theoretically identified [4]: **Dynamic temporal abstraction,** which refers to the ability of a modeling language to allow a *condensed representation* of a temporal model for a domain, from which appropriate *problem-specific submodels* can be created when problem instances are encountered. Dynamic temporal abstraction can be viewed as an instance of the *knowledge-based model construction* and *object-oriented model construction* methods that is specifically tailored to time modeling [5,6].

Hybrid temporal representation which refers to the ability of a representation to allow expression of objects or relations of the problem-solving domain in several levels of temporal detail [4].

GOALS

The goals of the present research are: (a) To introduce a representation language that implements these two design principles *fully and in an integrated manner*, (b) to empirically investigate the modeling effectiveness of using hybrid temporal representation and dynamic temporal abstraction by applying the language of part (a) in a time-sensitive medical domain, and (c) to identify general principles for MDSSs time-modeling language design.

DOMAIN BACKGROUND

Liver transplantation is a medical domain with important technical challenges, as well as decision problems with profound ethical and economic repercussions. Under the current organ donor shortage, one of the most pressing issues is the optimal allocation of organs. Two related questions are under which conditions (if any) retransplantation should be attempted, and what is the optimal timing for transplantation [7,8,9].

The liver transplantation domain *is highly temporal and is driven by complex tasks such as prediction/ forecasting, intervention and therapy planning, as well as policy formulation and evaluation.* As such, it can serve as a useful application environment for testing new methods for time modeling.

In the medical literature one can find several statistical and machine learning methods that have been applied in the domain of liver transplantation to address aspects of the above issues [7,8,10].

In this modeling experiment we focused on creating a model that would be able to: (a) predict the outcomes of patients given specific interventions and patient background information, and (b) evaluate the expected utilities of competing policies for transplantation timing and retransplantation constraints.

1091—8280/98/$5.00

RESULTS

1. Extending MTBNs to incorporate dynamic temporal abstraction and hybrid temporal representation.

In [3] we introduced MTBNs with a single granularity (MTBN-SGs). The main differences of those from ordinary belief networks and temporal belief networks are: (a) MTBN-SG models have well-defined temporal and causal semantics. (b) There are two additional types of variables besides ordinary ones: arc (mechanism variables) and cause-effect delay (lag) variables. All types of variables can interact (e.g., an ordinary variable may cause a lag variable, or an arc variable may cause another arc variable). (c) There is a condensed and a deployed graph. The former serves for defining and describing a model, while the later contains all variable instances over time and is used for inference.

MTBN-SG models have variables that are either indexed with a single temporal granularity, or non-indexed so that a simple form of hybrid temporal representation is supported. In the remainder of this section we introduce **multiple granularity MTBNs (MTBN-MGs)** that are enhanced so as to support sophisticated hybrid temporal representation and dynamic temporal abstraction by having the following additional enhancements over single granularity MTBNs:

(i) We introduce the notion of **temporal context** which is a collection of related variables. Variables are organized into variable types, variable types into temporal context instances, and temporal context instances into temporal context types. These elements are hierarchically related in an object-oriented ontology.

(ii) We introduce the **notions of inclusion and exclusion of a variable in a model**. An included variable is one that is relevant to solving a particular problem instance, whereas excluded ones are not.

(iii) The **model of time** is still discrete and linear, however it may consist of several temporal granularities (e.g., seconds, minutes, days).

(iv) Each temporal context instance has four additional **meta-variables** associated with it: *activation* (determines whether all context instance variables are included or not in the model), *duration*, *location* in time, and *granularity* (which can also change dynamically). Figure 1 shows graphically the example of a temporal context instance having a variable A, granularity of 1 month, duration of 2 months, starts at the 1[st] month, and its activation is not constant (but determined as a random variable).

Enhancements (i) to (iv) implement a hybrid temporal representation in MTBNs. The following additions implement dynamic temporal abstraction:

(v) **Generalized arc variables** describe connections among multiple variables (e.g., ARC[TC1(1-10).$A_{1\text{-}5}$ -> TC2(1-3).$B_{1\text{-}20}$] connects all variable instances of variable type A (belonging to context instances 1 to 10 of context type 1) and having temporal index 1 to 5, to all variable instances of variable type B (belonging to context instances 1 to 3 of context type 2) and having temporal index 1 to 20. The temporal index ranges, instance ranges and connection patterns are not necessarily constant but may be modeled as random variables themselves).

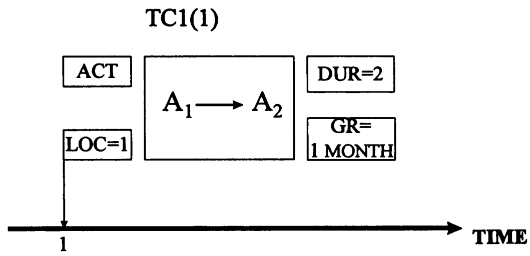

Figure 1: Temporal context instance example.

(vi) Inference is augmented with three optimizations. **Temporal context deactivation optimization** omits variables that belong to context instances that have been deactivated in a particular simulation cycle (Figure 2a). **Backward search optimization** builds only the portion of the deployed graph that needs to be instantiated to answer a specific query (Figure 2b). Finally, **dynamic duration optimization** omits variables that fall outside the dynamically-determined length of a temporal context instance (Figure 2c).

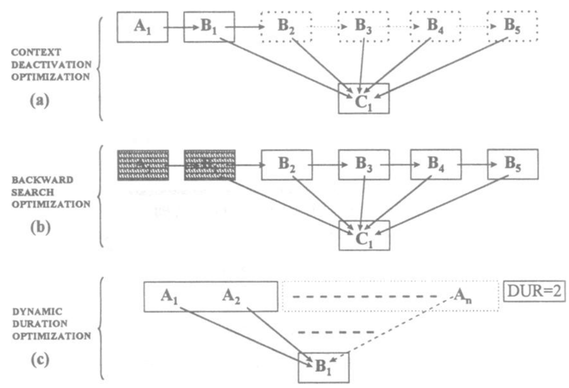

Figure 2: MTBN-MG optimizations. (a) Temporal context deactivation optimization. The context instances in dotted lines are deactivated, therefore we need not instantiate them. (b) Backward search optimization. The example query is "p(A_1=1 | B_1=2) =?". Only A_1 and B_1 (belonging to the context instances with the shaded background) need be instantiated. (c) Dynamic duration optimization. The top context instance is assigned a duration of 2, therefore variable instances A_3 to A_n need not be instantiated.

2. The liver transplantation model.

With the help of a domain expert we built a liver transplantation model capable of answering the query types described in the "Domain background" Section using a combination of:

(a) *Knowledge engineering* (for the domain structure, qualitative relationships, and some quantitative relationships expressed as conditional probabilities),

(b) *Domain theory* (in the form of textbook physiological, or surgical facts, as well as published analyses and predictive models).

(c) *Statistical analysis/machine learning* from patient data.

A total of approximately 15 two-hour knowledge-acquisition sessions with the expert were conducted, about 20 published articles were used, and 2 patient data sets from the expert's file were analyzed.

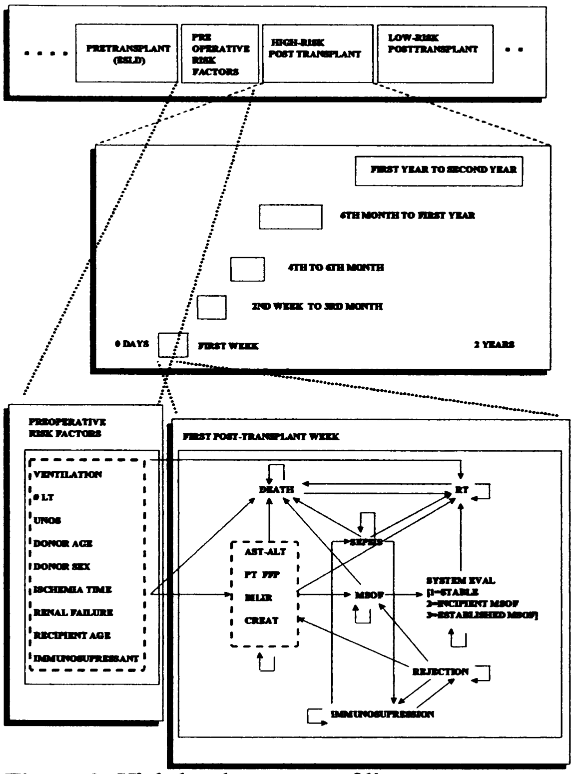

Figure 3: High-level structure of liver transplantation model, and condensed graph for two of the component temporal context instances.

Figure 3 (top) presents the high-level domain structure. The bottom of Figure 3 elaborates on two of the components of the high-level structure. We can see that the history of liver transplantation starts from the end-stage liver disease context (ESLD – top left of Figure 3), in which a patient with failing liver function is waiting for a transplantation. During this waiting period the patient may die, suffer a complication, or get transplanted. Once a transplantation occurs, a number of preoperative factors are observed (Bottom left of Figure 3) that determine the subsequent outcome of the transplantation. Some of these factors, such as renal failure or liver ischemia time are very important, and are considered *high-risk indicators*, whereas others (e.g., the donor's age) are considered *less important risk factors*. Subsequent to transplantation, the patient goes through a *high-risk post-transplant period* during which graft or donor failure has higher chances of occurring. This period lasts for approximately 2 years and can be further decomposed into a number of smaller temporal contexts corresponding to the first week, 2nd week to 3rd month, 4th month to 6th month, 6th month to 1st year, and 2nd year periods. These periods differ little in terms of the involved variables and more in terms of the quantitative strength of association among the variables. Looking at the bottom right of Figure 3, we see that during the first week after transplantation, a number of laboratory variables (hepatic enzymes – AST/ALT, prothrombin time – PT, bilirubin – BILIR, and creatinine – CREAT) are closely observed so that an assessment is being made as to whether the patient is stable, or is having incipent or established multiple system organ failure (MSOF), and also regarding the need to administer fresh frozen plasma (FFP). The type of immunosupression administered, the presence of sepsis, and a possible rejection taking place, all influence the laboratory findings, and the assessment made by the surgical team about a possible retransplantation. In turn, the decision for retransplantation along with the preoperative risk factors and the current systemic status will determine the patient's chances for survival within this time frame.

After the 2-year mark the patients enter a low-risk post-transplant period with chances of death roughly the same as the general population of the same age and gender (Figure 3 – top right). At any time during the high or low-risk periods the patients may need a retransplantation, in which case they will enter a new ESLD context and start the cycle anew. Note that the duration of the pre and post-2-year transplantation contexts is not constant. Also, the temporal contexts of Figure 3 can be replicated several times if multiple transplantations take place. Given the applicable retransplantation policies and the number of previous transplants, the multiple retransplantation option may not be available to some patients.

Assuming a granularity of 1 month, a maximum time span of 10 years for the ESLD and low-risk contexts, and a maximum of 10 transplantations, the deployed liver transplantation model consists of about 10,000 nodes, whereas the undeployed model consists of approximately 400 nodes, including context types, instances, and meta-variables. This difference will be explained in detail in Section 4.

3. Representational power and hybrid temporal representation.

Representational power refers to the ability to express naturally and easily the desired aspects of the problem-solving domain (expressivity) and to the corresponding knowledge acquisition tractability.

In the liver transplantation domain, we found that the **simultaneous representation of different temporal units** was necessary. Some questions and pieces of evidence are expressed at a small temporal granularity (for instance, the granularity of the first post-transplantation week variables is 1 day) while others are expressed at larger granularity levels (e.g., at the 2nd post-LT year the time granularity is 1 month).

Through our interaction with the expert, we found that another form of hybrid temporal (-abstraction) representation was also very helpful to him for parameterizing the model. In all instances of subjective assessment of quantitative relationships, the expert felt that he could not reliably assess the probabilistic conditional distributions of a node given all possible parents instantiations, but he could provide reliable **rules that abstract together many of the parent nodes' values** (e.g., "if no positive preoperative risk factors and no MSOF were present then the chances of death within the first post-transplantation week are <0.05").

Similarly, the **expert often felt uncomfortable giving assessments in one granularity but comfortable in another.** This led to an adjustment of the initial granularities for the various high-risk post-transplant period contexts to the granularities the expert was most comfortable with.

4. Representational tractability and hybrid temporal representation/dynamic temporal abstraction.

Representational tractability refers to the complexity of storing a model of the problem-solving domain.

Since the deployed liver transplantation model is very large, to facilitate runtime analyses we implemented a simplified version, which we call the *small liver transplantation model.* The simplification consists of allowing for up to two explicit transplantation periods (instead of 10) and of omitting some variables in the high-risk post-transplant contexts.

Table 1: Number of nodes and number of conditional probability distributions

MTBN-MG	110
BN WITHOUT HYBRID TEMPORAL ABSTRACTION	$3*10^5$
BN WITH HYBRID TEMPORAL ABSTRACTION	400

The MTBN small model is declared using 250 elements (110 nodes including context types, instances and meta-variables, and 140 arcs and abstraction variables) and 100 conditional probability distributions. Since the smallest granularity is one day and the model time horizon is a maximum of 54 years, assuming we have 15 distinct variables in the model only, if we wanted to express the small model in standard temporal belief network form (i.e., using temporal uniformity), we would need approximately: $54*365*15= 3*10^5$ nodes and an equal number of probability distributions. This is clearly intractable without even considering the necessary arcs.

The MTBN-MG dynamic abstraction features enhance even further representational tractability. A belief network for the small model that would not use uniform temporal representation requires a total of approximately 1,000 arcs, 400 nodes, and 400 conditional probability distributions. Therefore, the MTBN small model uses 18% of total arcs and nodes needed for defining the same model in BN form, and 25% of total probability distributions (Table 1).

Once defined, the MTBN model can be *modified very efficiently* when the overall temporal context and instance structure remains the same. For example, increasing the maximum horizon of the first ESLD context from 10 years to 15 years requires changing just one number (the duration of the ESLD context type). In the belief network case, that change involves declaring 40 new death and transplantation nodes for the first ESLD context instance. To connect these nodes to the rest of the belief network model, we would need to declare an additional 120 arcs (query abstractions excluded) to capture the relations among the new nodes. In the MTBN model this can be done by changing 3 parameters. Furthermore, if we would like to *expand the basic structure of the model,* for example by adding a third explicit transplantation period in the belief network model, we would need to increase the model declaration size by 50% (700 nodes and arcs), since currently each transplant period is approximately 50% of the current model. In the MTBN case, we need to declare 28 new meta-variables, 30 new arcs, and change one parameter in each one of existing 25 arcs. The total size of the MTBN increases by 25% only. Similar expansion for a fourth transplantation requires an increase of 33% in the size of the BN but only a 17% increase in the size of the MTBN, and so on (Figure 4).

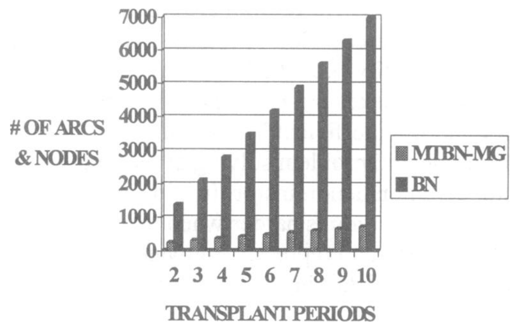

Figure 4: Rate of model size increase when adding new transplant periods.

5. Inferential tractability and dynamic temporal abstraction.

Inferential tractability refers to the complexity of reasoning with a model of the problem-solving domain to solve a set of problem instances.

To examine the effects of the 3 MTBN-MG optimizations on inferential tractability we implemented code that performs inference with the optimizations turned on and off. We used a sample of 17 queries generated from a set of 6 query types that the expert indicated were pertinent in this domain [2].

Table 2 shows the average, minimum and maximum runtime ratios of no optimization versus full optimization for a constant number of simulation cycles. The ratios are to be interpreted as the factor by which optimization allows reaching an answer (with fixed accuracy) faster than with no optimization. We see that averaged over all examined queries, the 3 MTBN-MG optimizations allow on the average for 172 times faster inference relative to no optimization. Table 2 also shows the size ratios (expressed in number of instantiated nodes) of the problem-specific model produced by the optimizations. On the average, the problem-specific model was 10.9 times smaller than the full (non-optimized) model.

Table 2: Inference time and model-size improvement ratios of dynamic temporal abstraction (all optimizations) versus inference with no optimizations in small MTBN-MG model.

	AVERAGE	*MIN*	*MAX*
TIME RATIO	172	107	340
SIZE RATIO	10.9	6.4	32

CONCLUSIONS

The above results provide support in favor of the following conclusions:

Conclusion #1: Dynamic temporal abstraction is a desirable representation language property because it *enhances inferential tractability* by focusing computational inference resources on the relevant temporal aspects of a problem instance. Furthermore, dynamic temporal abstraction *enhances representational tractability* by enabling efficient representation of a model rather than using a potentially intractable explicit model.

Conclusion #2: Hybrid temporal representation is a desirable representation language property since temporal uniformity (i.e., the lack of hybrid temporal representation) makes *inference and storage* of temporal models inefficient by focusing representation and problem solving at the smallest temporal common denominator. In addition, *knowledge acquisition* is harder when there is no flexibility to seek and represent knowledge at the level of temporal detail in which this knowledge is normally available.

DISCUSSION

We now discuss how to incorporate dynamic temporal abstraction and hybrid temporal representation in existing representations that lack them, based on our extensions of belief networks and single-granularity MTBNs.

To implement hybrid temporal representation, we need an **explicit representation of temporal units** (or relevant components of time) of different sizes, as well as semantic rules that allow association of those with the various objects of the language.

To implement dynamic temporal abstraction we: (a) introduce (in the new language) structures that are defined as **collections of existing structures** (of the old language), based on problem-specific criteria. For example, sets of variables in temporal belief networks are declared as temporal context instances in MTBNs. (b) Once *collections* of objects in the original representation have been described as *individual* objects of the new representation, these objects can be **stored efficiently** by the new representation. For example, in MTBN-MGs, arc variables represent arbitrarily large collections of individual arcs that connect related variables. (c) Finally, we need to **extend the inference procedures** so that they take advantage of the efficient representation of objects in the enhanced representation. In MTBN-MGs this is accomplished by the backward search, variable/context-deactivation, and dynamic context length optimizations of the inference algorithm.

Since the above steps are not inherently language-dependent we believe that they can serve as a high-level guide for incorporating dynamic temporal abstraction and hybrid temporal representation in a wide variety of existing representations.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the enthusiastic and knowledgeable contribution of Dr. H. Doyle who served as the liver transplantation expert. This research has benefited greatly by the comments of Professors B.G. Buchanan, M. Pollack, and M. Wagner. This research was supported in part by grant NSF BES-9315428 from the National Science Foundation, and grant LM-05291-04 from the National Library of Medicine.

References

1. Combi C, Shahar Y. Temporal Reasoning and Temporal Data Maintenance in Medicine: Issues and Challenges. Comput. Biol. Med. 1997;27:353-368.
2. Aliferis C. A Temporal Representation and Reasoning Model for Medical Decision-Support Systems. Doctoral Thesis, University of Pittsburgh, 1998.
3. Aliferis C, Cooper GF. A new Formalism for Temporal Modeling in Medical Decision-Support Systems. 19th Annual Symposium on Computer Applications in Medical Care, 1995.
4. Aliferis CF, Cooper GF, Pollack ME, Buchanan BG, Wagner MM. Representing and Developing Temporally Abstracted Knowledge as a means Towards Facilitating Time Modeling in Medical Decision-Support Systems. Comput. Biol. Med., 1997;27:411-434.
5. Laskey KB, Mahoney SM. Network Fragments: Representing Knowledge for Constructing Probabilistic Models. Proceedings of Uncertainty in Artificial Intelligence: Morgan Kaufmann, 1997.
6. Koller D, Pfeffer A. Object-Oriented Bayesian Networks. Proceedings of Uncertainty in Artificial Intelligence: Morgan Kaufmann, 1997.
7. Roberts M. Markov Process-Based Monte Carlo Simulation: A Tool for Modeling Complex Disease and its Application to the Timing of Liver Transplantation. Winter Simulation Conference, 1991.
8. Doyle HR, Morelli F, McMichael J, et al. Hepatic Retransplantation - An Analysis of Risk Factors Associated With Outcome. Transplantation 1996;61:1-7.
9. D'Amico DF, Bassi N, Tedeschi U, Cillo U. Liver Transplantation: Procedures and Management. Milano: Masson, 1994.
10. Strasberg SM, Howard TK, Molmenti EP, Hertl M. Selecting the Donor Liver: Risk Factors for Poor Function After Orthotopic Liver Transplantation. Hepatology 1994;20:829-838.
11. Ngo L, Haddaway P, Helwig J. A Theoretical Framework for Context-Sensitive Temporal Probability Model Construction with Application to Plan Projection. Proceedings of Uncertainty in Artificial Intelligence, 1995:419-426.