

**Knowledge Systems Laboratory
Working Paper KSL 88-30**

**July 1988
(Revised April 1989)**

A Method for Temporal Probabilistic Reasoning

by

Gregory F. Cooper, Eric J. Horvitz, and David E. Heckerman

**KNOWLEDGE SYSTEMS LABORATORY
Medical Computer Science
Stanford University
Stanford, California 94305-5479**

#210

A Method for Temporal Probabilistic Reasoning

Gregory F. Cooper, Eric J. Horvitz, and David E. Heckerman

Section on Medical Informatics
Stanford University School of Medicine
Stanford, CA 94305-5479

1. Introduction

Temporal reasoning in complex environments often involves the interpretation of evidence under uncertainty. We have been studying the application of probabilistic reasoning techniques to the problem of temporal-data interpretation. In this report, we discuss techniques for inferring the probability of current hypotheses given temporal evidence. After reviewing the background of the problem, we present our problem-solving approach and discuss a prototype implementation.

Temporal reasoning has been recognized as an important and difficult area of artificial intelligence (AI) research. During the last decade, the volume of temporal-research literature has increased markedly. A good survey of research up to 1982 is contained in [4]. Key early papers in the field include [1, 2, 16, 23]; more recent literature can be found in [7, 10, 12, 15, 19, 20, 21, 26, 30, 31, 32]. A common theme in much of this research is the view of temporal reasoning as logical entailment. Issues of concern include the logical relationships of different representations (e.g., point-based versus interval-based), the completeness of particular methods in inferring all possible temporal relationships from given data, the consistency of those inferences, and the computational complexity of temporal inference. The ability to represent and reason with uncertainty is not addressed in these research projects. However, recent work by Higgins [13], and by Dean and Kanazawa [7] has focused on temporal probabilistic reasoning. Higgins' system reasons about the dynamics of probabilistic systems. An acyclic network structure is used to specify the qualitative dependencies among variables, and probability

density functions are used to quantify those dependencies. Higgins' system is limited to a network in which there are directed arcs from evidence to hypotheses, but no directed arcs from hypotheses to evidence.

Temporal reasoning is particularly important in medical diagnosis and management. Thus, there have been numerous articles addressing research on medical systems for reasoning over time. Good reviews are [8, 29]. Previous work of particular note include [9, 11, 14, 17, 18, 24, 27, 28]. This medical research is often more applied than theoretical. Work described in [3, 8, 22, 33] addresses some temporal uncertainty issues. In particular, Weiss describes a system, called CASNET, that reasons about probabilistic causal *precedence* in performing diagnosis [33]. The system has no notion of the specific time of events and therefore is severely limited in the type of temporal reasoning it can perform. Blum has developed a system called RX that infers causal relationships from a database [3]. RX represents causal relationships using a simple, restricted form of probabilistic dependency. RX does not, however, use its knowledge for probabilistic data interpretation. Long has developed a system that uses qualitative probabilities to simulate a physiological model over time [22]. The research on IDEFIX by de Zegher-Geets deals explicitly with temporal probabilistic medical diagnosis by using temporal probability functions among variables [8]. The inference techniques in IDEFIX are not normative, however, and instead rely on numerous heuristics. In summary, there are several medical systems that perform temporal reasoning under uncertainty, but their representational and inference capabilities are quite limited.

2. The Complexity of Probabilistic Temporal Reasoning

Reformulating the probabilistic diagnostic problem from an atemporal one to that of assigning likelihood to alternative hypotheses over time adds significant complexity. The complexity is affected by (1) how finely time is divided into discrete perspectives, and (2) how far back in time a reasoner will consider evidence as potentially relevant to the current belief. The most general approach to the temporal-reasoning problem would theoretically require an infinite number of joint probabilities among a set of variables, one

for each *point* in (continuous) time. Thus, in the most general case, probabilistic temporal reasoning is clearly intractable both in computation time and in knowledge-base storage.

Consider the task of determining the probability of a hypothesis H being true at time t , given all the evidence E that has been observed *before* time t . We denote this probability as $P(H_i(t) \mid E(0, t))$, where time 0 is the earliest possible absolute time at which evidence is available. In general, E must represent every possible value pattern of the data variables in the system. If time is continuous, then there is potentially an infinite number of possible patterns of E that must be stored in the knowledge base.

Assume, however, that there is some minimum granularity of time called T_g ; that is, no time measurement shorter than T_g is of interest; for example, T_g might be one second. Furthermore, suppose that each evidence variable has V possible values. The value of $E(0, t)$ is some time pattern of the values of the system evidential variables from time 0 to time t , discretized into time periods of length T_g . Thus, there are $V^{|E|} t/T_g$ probabilities to store in the knowledge base for each hypothesis, where $|E|$ is the total number of evidential variables. This general case leads to an astronomical number of states to acquire and store. For example, suppose there are 20 evidential variables, T_g is one second, and V is 3 (e.g., the values might be *true*, *false*, or *unknown*). A system that reasons about a patient's condition in the intensive care unit (ICU) over a 2-hour (7200 second) period would theoretically require $3^{20 \times 7200/1} = 3^{144,000} \cong 10^{68,705}$ temporal patterns of the form $P(H_i(t) \mid E(0, t))$ for each hypothesis $H_i(t)$ about the patient's condition. Obviously, it is not possible to acquire or store this many probabilities.

3. Methods for Obtaining Tractable Inference

To make probabilistic temporal reasoning tractable, we must impose practical restrictions on the general case introduced in Section 2. In this section, we discuss three types of restrictions.

3.1 Representing Temporal Conditional Independence

The assumption of conditional independence can simplify temporal probabilistic reasoning. For instance, if all evidential variables are conditionally independent given each hypothesis, then only $|E|V^t/T_g$ probabilities are necessary for each hypothesis, rather than the $V^{|E|} t/T_g$ probabilities needed as before. All variables do not need to be conditionally independent in order to achieve a reduction. The more independence that exists, the greater the reduction.

3.2 Decreasing the Evidential Time Horizon

In the general case described in Section 2, each set of evidence E was presumed to influence each hypothesis H_i beginning at some distant time 0 through the current time t . Often, this assumption is not realistic. Usually, evidence will have a more bounded, specific span of influence. Let T_h be the maximum amount of time before t that any evidence can significantly influence H_i (i.e., T_h is a time horizon). In this case, there are T_h/T_g possible time periods during which evidence must be considered. Thus, with no assumptions of conditional independence, there are at most $V^{|E|} T_h/T_g$ possible states of evidence bearing on each hypothesis H_i . The temporal horizon imposes a time window of dependency between past evidence and current hypotheses. By shortening the time horizon, we shorten the length of the time period for which past evidence is assumed to influence current hypotheses. The evidential time horizon, T_h , may depend on both the type of evidence and the particular hypotheses under consideration. Note that more generally, each piece of evidence may have its own specific time horizon with respect to specific hypotheses. We shall return to the issues of evidential time horizons in our discussion of a particular representation and inference strategy in Section 5.

3.3 Increasing the Time Granularity

Maintaining a small time granularity T_g for each evidential variable often is unnecessary. For example, in many areas of medicine, it is sufficient to maintain a time granularity on many evidential variables that is on the order

of hours or days, rather than of seconds. Increasing the time granularity can dramatically reduce the number of states considered. Dynamic changes in the time granularity, T_g , may be useful. For instance, increasing the granularity as some function of the age of the evidence might be useful, because the importance of being precise often decreases with the age of the evidence [9]. For example, over a give 10-minute period, the value of a given piece of evidence may be important second by second, but its value 3 hours previously may be important only on a minute-to-minute basis.

4. A General Model for the Interpretation of Temporal Probabilistic Data

Our problem-solving objective is to compute the current probability of each hypothesis H_i , given all past relevant evidence. In the previous section we discussed the complexity of this task, and examined several methods for rendering the task more tractable. We now introduce a model that facilitates the application of these methods. The essence of the model is the way in which evidence and hypotheses are conditioned on each other. The model attempts to capture a natural form of conditioning that allows the intuitive expression of subjective probabilities and the efficient application of known forms of conditional independence.

In this model, we make the following assumptions:

Assumption 1. Time is viewed as discrete and is represented by an integer that increases by 1 with each incremental increase in time.

We therefore develop models that will use summation. It is possible to generalize to continuous time and to use integration, but we do not address this issue in this paper.

Assumption 2. Evidence from before some finite t_0 is not relevant, and hypotheses are not considered before t_0 .

Note that t_0 equals the current time minus T_h .

Assumption 3. One and only one hypothesis can be true at a given time.

We can generalize Assumption 3 by allowing hypotheses to be defined as conjunctions of other hypotheses. In the remainder of this paper, we consider a hypothesis to be either some single patient disease or the absence of all diseases. Techniques for diagnosing atemporal multiple diseases in the medical domain can be adapted for temporal multiple-disease diagnosis; however, temporal, multiple-disease diagnosis is generally much more computationally expensive than is atemporal diagnosis.

The current time is represented as t_{now} . All the evidence from t_0 to t_{now} is represented as $E(t_0, t_{\text{now}})$. Therefore, the objective is to compute for each disease d_i the probability $P(d_i(t_{\text{now}}) \mid E(t_0, t_{\text{now}}))$.

Unfortunately, it is often difficult to acquire subjective probabilities directly of the form $P(d_i(t_{\text{now}}) \mid E(t_0, t_{\text{now}}))$. Furthermore, this form of probabilistic representation does not take advantage of possible conditional independencies of the evidence given diseases and in not doing so it may lead to an intractable representation, as discussed in Section 2. We have therefore chosen to decompose this complex probability into the following four probabilities:

- a. $P(d_i\{t_x\} \mid E(t_0, t_x - 1))$. The open-bracket notation in $d_i\{t_x\}$ means that disease d_i *began* at t_x . That is, d_i did not exist immediately before t_x , and does exist at t_x . Thus, $P(d_i\{t_x\} \mid E(t_0, t_x - 1))$ represents the probability of d_i beginning at time t_x , given the evidence up to (but not including) time t_x . For simplicity, we will consider that $T_g = 1$. By varying t_x , we obtain a probability distribution on the expected start times of each disease d_i given the evidence. Note that the start times for a given d_i are mutually exclusive.
- b. $P(d_i\{t_x, t_{\text{now}}\} \mid d_i\{t_x\}, E(t_0, t_x - 1))$. This term represents the probability of disease d_i extending from time t_x until the current time t_{now} , given that d_i began at time t_x and given all the past evidence. Note that this probability does not state that d_i *ended* at the current time t_{now} , but only that it exists at t_{now} . By varying t_x , we obtain a probability

distribution on the expected durations of a hypothesis through t_{now} . Here we have used the notation $d_i\{t_x, t_{\text{now}}\}$ to represent the expression $d_i\{t_x\} \wedge d_i(t_x, t_{\text{now}})$.

- c. $P(E(t_x, t_{\text{now}}) \mid d_i\{t_x, t_{\text{now}}\}, E(t_0, t_x - 1))$. This term represents the probability that the evidence that is observed from time t_x to t_{now} would be in fact be observed, given that disease d_i began at time t_x and continued until at least the current time, and given all the evidence that is observed before t_x . This probability expresses the chance of observing recent evidence in the context of an ongoing disease and more distant evidence.
- d. $P(E(t_0, t_x - 1))$. This term represents the probability of observing all the evidence that is observed before t_x . Later in this section we show how to compute this probability using probabilities a through c .

Probabilities a through c may appear to be more complex than $P(d_i(t_{\text{now}}) \mid E(t_0, t_{\text{now}}))$, which they are meant to replace as the focus of knowledge acquisition. In Section 5, however, we introduce assumptions that will render the three probabilities more intuitive to assess and more tractable to use in computing the time-varying probabilities of diseases.

The first step in computing $P(d_i(t_{\text{now}}) \mid E(t_0, t_{\text{now}}))$ involves using the above four probability terms to compute $P(d_i\{t_x, t_{\text{now}}\}, E(t_0, t_{\text{now}}))$ for all d_i and all t_x , such that $t_0 \leq t_x \leq t_{\text{now}}$. The following equation is used to perform this computation:

$$\begin{aligned}
 P(d\{t_x, t_{\text{now}}\}, E(t_0, t_{\text{now}})) = & \tag{1} \\
 & P(E(t_x, t_{\text{now}}) \mid d_i\{t_x, t_{\text{now}}\}, E(t_0, t_x - 1)) \times \\
 & P(d_i\{t_x, t_{\text{now}}\} \mid d_i\{t_x\}, E(t_0, t_x - 1)) \times \\
 & P(d_i\{t_x\} \mid E(t_0, t_x - 1)) \times \\
 & P(E(t_0, t_x - 1))
 \end{aligned}$$

Note that when $t_x = t_0$, Equation 1 becomes:

$$P(E(t_0, t_{\text{now}}) | d_i\{t_0, t_{\text{now}}\}) \times P(d_i\{t_0, t_{\text{now}}\} | d_i\{t_0\}) \times P(d_i\{t_0\})$$

The joint probability of both observing all the evidence seen through time t_{now} and having hypothesis d_i occur at t_{now} can be calculated from Equation 1 as follows:

$$P(d_i(t_{\text{now}}), E(t_0, t_{\text{now}})) = \sum_{t_0 \leq t_x \leq t_{\text{now}}} P(d_i\{t_x, t_{\text{now}}\}, E(t_0, t_{\text{now}})) \quad (2)$$

Equation 2 can be used to compute the marginal probability of evidence from t_0 to t_{now} :

$$P(E(t_0, t_{\text{now}})) = \sum_{d_i} P(d_i(t_{\text{now}}), E(t_0, t_{\text{now}})) \quad (3)$$

Note that Equation 3 uses Assumption 3 — namely, that the diseases are exhaustive (i.e., every possible disease of *importance* is represented) and mutually exclusive (i.e., diseases do not occur simultaneously).

Note also that Equations 2 and 3 can be applied to compute $P(E(t_0, t_z))$ for t_z from t_0 to $t_{\text{now}} - 1$ before it is necessary to compute $P(d_i(t_{\text{now}}), E(t_0, t_{\text{now}}))$, using Equation 2 again.

Using Equations 2 and 3, the probability of any given disease d_i is computed as follows:

$$P(d_i(t_{\text{now}}) | E(t_0, t_{\text{now}})) = \frac{P(d_i(t_{\text{now}}), E(t_0, t_{\text{now}}))}{P(E(t_0, t_{\text{now}}))} \quad (4)$$

5. A Prototype System

We have implemented a small prototype system to investigate the feasibility of the reasoning methods discussed in Section 4.

5.1 Additional assumptions

In addition to the three assumptions in Section 4, we make the following assumptions in implementing the prototype:

Assumption 4. Variables are binary valued.

This assumption is by no means necessary, but it simplifies the program and focuses our attention in this paper on the most fundamental issues.

Assumption 5. The value of each evidential variable is known for each point in time from time t_0 to t_{now} .

General probabilistic inference using a belief network with uninstantiated variables (i.e., variables with unknown values at some points in time) is NP-hard [5]. We have chosen currently to concentrate our efforts on issues other than the techniques needed for efficient temporal probabilistic inference given uninstantiated variables.

Assumption 6. Discrete time is measured at a *fixed* granularity T_g .

Assumption 7. $P(d_i\{t_x\} \mid E(t_0, t_x - 1)) = P(d_i\{t_x\})$.

According to Assumption 7, the probability of a disease *starting* at time t_x is dependent only on t_x , and not on any past evidence. Typically, t_0 might be the time of birth of the patient, and thus t_x is the patient's age. We are assuming that evidence is not *predictive* of when a disease will begin, but rather is only a *consequence* of the disease once the disease has begun. This assumption is a particularly restrictive one that we made to simplify the development of the initial prototype system.

Assumption 8. $P(d_i\{t_x, t_{now}\} \mid d_i\{t_x\}, E(t_0, t_x - 1)) = P(d_i\{t_x, t_{now}\} \mid d_i\{t_x\})$.

Assumption 8 states that the duration of a disease is not dependent on past evidence. Each disease, however, can have a different distribution of its duration (e.g., pneumonia generally persists longer in the elderly than in young adults).

Assumption 9. $P(E(t_x, t_{now}) \mid d_i\{t_x, t_{now}\}, E(t_0, t_x - 1)) = P(E(t_x, t_{now}) \mid d_i\{t_x, t_{now}\})$.

According to Assumption 9, the probability of observing particular evidence during a period of time is dependent on only the disease that is occurring during that period of time and not on earlier evidence before that time period. Thus, evidence is conditioned only on a possible ongoing disease. Like Assumption 7, this assumption is particularly restrictive.

Assumption 10. The sum in Equation 2 needs to be taken over only $t_{now} - horizon(d_i) \leq t_x \leq t_{now}$, where *horizon* is a function that returns the maximum amount of time that disease d_i can possibly exist.

Assumption 10 improves the efficiency of the calculation of Equation 2 by using knowledge about the maximum possible duration of each disease. We define t_h to be $t_{now} - horizon(d_i)$. Although currently *horizon* is not context-sensitive, this restriction is by no means necessary.

5.2 The Prototype Model

Applying Assumptions 7 through 10 to the terms in Equation 1 and substituting this into Equation 2, we derive the following equation, which forms the basis of our prototype system:

$$P(d_i(t_{\text{now}}), E(t_0, t_{\text{now}})) = \tag{5}$$

$$\sum_{t_h \leq t_x \leq t_{\text{now}}} P(E(t_x, t_{\text{now}}) | d_i(t_x, t_{\text{now}})) \times P(d_i(t_x, t_{\text{now}}) | d_i(t_x)) \times P(d_i(t_x)) \times P(E(t_0, t_x-1))$$

Equation 4 is then applied as before to calculate $P(d_i(t_{\text{now}}) | E(t_0, t_{\text{now}}))$ for each d_i .

The assumptions embodied in Equation 5 impose restrictions on the application of the general technique; their validity for any given domain must be evaluated. Note that there is a sliding scale from the general model represented by Equation 2 to the much more restricted implementation represented by Equation 5. In any given application, we need to assess the particular set of assumptions that seem most appropriate for that domain. The tradeoff is between the complexity of a relatively high-fidelity system with few assumptions and the simplicity of a relatively lower-fidelity system with many assumptions. Finding the appropriate tradeoff is key. The general formulation in Section 4 provides a basis for considering these tradeoffs explicitly.

5.3 Expression of Conditional Independencies

The term $P(E(t_x, t_{\text{now}}) | d_i(t_x, t_{\text{now}}))$ in Equation 5 provides the opportunity to introduce additional independences in the form of conditional independence of evidence given a disease. In particular, this term can be calculated using a belief network [25] that contains variables with a temporal dimension. To see this result, we first transform the term as follows:

$$P(E(t_x, t_{\text{now}}) | d_i(t_x, t_{\text{now}})) = \tag{6}$$

$$\prod_{1 \leq j \leq n} P(e_j(t_x, t_{\text{now}}) | e_1(t_x, t_{\text{now}}), \dots, e_{j-1}(t_x, t_{\text{now}}), d_i(t_x, t_{\text{now}}))$$

where $e_j(t_x, t_{now})$ represents the values from t_x to t_{now} of event e_j , which is one of n events in E .

Often, it is not necessary to condition the probability of each $e_j(t_x, t_{now})$ in Equation 6 on *all* of the events $e_1(t_x, t_{now}), \dots, e_{j-1}(t_x, t_{now})$, because $e_j(t_x, t_{now})$ may be conditionally independent of many of these events given $d_i\{t_x, t_{now}\}$. For the case in which each evidence variable is completely conditionally independent of the others (given $d_i\{t_x, t_{now}\}$), Equation 6 becomes

$$P(E(t_x, t_{now}) \mid d_i\{t_x, t_{now}\}) = \prod_{1 \leq j \leq n} P(e_j(t_x, t_{now}) \mid d_i\{t_x, t_{now}\}) \quad (7)$$

Although the global assumption of conditional independence in Equation 7 is extreme, for simplicity we assume it in the prototype. More typically, events will be partially dependent on some other events given a disease, but not on all other events. By selectively including only the conditioning events of relevance to the probability of each $e_j(t_x, t_{now})$ we form a belief-network architecture. Note that a more general formulation than the one we address in the prototype would allow $e_j(t_x, t_{now})$ to be conditioned on events preceding t_x .

To recap, the probabilities that we now must obtain to construct a complete probabilistic temporal model of disease are the probabilities on the right hand sides of Assumptions 7 and 8 and Equation 7, as summarized here:

$P(d_i\{t_x\})$: the prior probability distribution of the initiation of a disease.

$P(d_i\{t_x, t_{now}\} \mid d_i\{t_x\})$: the probability distribution over the duration of a disease.

$P(e_j(t_x, t_{now}) \mid d_i\{t_x, t_{now}\})$: the probability of observing a particular pattern of evidence during a particular episode of a disease.

5.4 Implementation of the Prototype System

The temporal probabilistic-reasoning algorithm discussed in Section 5 has been implemented as an INTERLISP program on a Xerox Lisp workstation. We shall describe the knowledge base, and then shall present several sample runs.

We created a small initial knowledge base of abstract diseases and evidence. Although this knowledge base is rudimentary, it does demonstrate useful concepts. In particular, we have investigated the issues of *event precedence* and *event lag* (i.e., the degree to which events are expected to occur together temporally).

Consider two evidential events, e_1 and e_2 , that may occur in the context of three different diseases, $disease_1$, $disease_2$, and $disease_3$. In $disease_1$, e_1 typically precedes e_2 , and these two events generally do not occur close together in time. In $disease_2$, event e_1 again typically precedes event e_2 , but in contrast to $disease_2$, e_2 generally occurs soon after e_1 . In $disease_3$, e_2 typically occurs soon before e_1 .

The knowledge base uses a probability function that is based on a simple truncated Gaussian distribution. More sophisticated (and typically more realistic) distributions, could be used to implement time-precedence and time-lag relationships probabilistically. In the context of a given disease, our current distribution function contains a mean and standard deviation of the lag between two variables. For example, the mean lag between e_1 and e_2 in the context of $disease_1$ is 5 time units, whereas in the context of $disease_2$ it is only 1 time unit.

The probability distribution on the *duration* of each of the three diseases is a simple uniform distribution. Currently, all three diseases are given the same uniform distribution function.

The temporally varying prior probability of each of the diseases is different. In particular, $disease_1$ typically occurs early, $disease_3$ generally occurs late,

and disease₂ is generally intermediate between disease₁ and disease₃. A lag function is used to compute these time-varying prior probabilities.

5.5 Sample Runs of the Prototype System

Figure 1 demonstrates the probability of the diseases over time in the absence of any positive evidence. As expected, based on prior probabilities alone, disease₁ generally occurs before disease₂, which typically occurs before disease₃.

Figure 2 contains some positive evidence. This evidence was entered in the lower window using a mouse; only evidence (events) e_1 and e_2 are relevant to the examples we discuss here. In general, the user can add, modify, and delete evidence, then instruct the system to generate an interpretive plot of the probabilities of the diseases over time given the evidence. In this particular example, the pattern favors disease₁ over disease₂, because e_1 and e_2 occur far apart in time and disease₁ already has a high prior probability in this time region. Figure 3 contains evidence over an interval in which disease₂ becomes more likely (relative to the situations in Figures 1 and 2) due to the short time lag between e_1 and e_2 .

Figure 4 shows a pattern of evidence that again favors disease₂. Figure 5 demonstrates the opposite precedence pattern, in which e_2 occurs before e_1 , which favors disease₃ relative to disease₂.

These simple experiments with our prototype are encouraging, albeit only preliminary. It will be important to test the code more exhaustively and to construct a more complex knowledge base. We can use this more extensive system to evaluate further the key issues of the expressiveness of the representation and the computational efficiency of the probabilistic inference algorithm.





Figure 1. The probability distribution of the diseases over time in the absence of any evidence.





Figure 2. A pattern of evidence characteristic of disease₁, due to the high prior probability of disease₁ in this time region, and to the long time lag from e₁ to e₂.





Figure 3. A pattern of evidence characteristic of disease₂, in which e₁ is quickly followed by the occurrence of e₂.


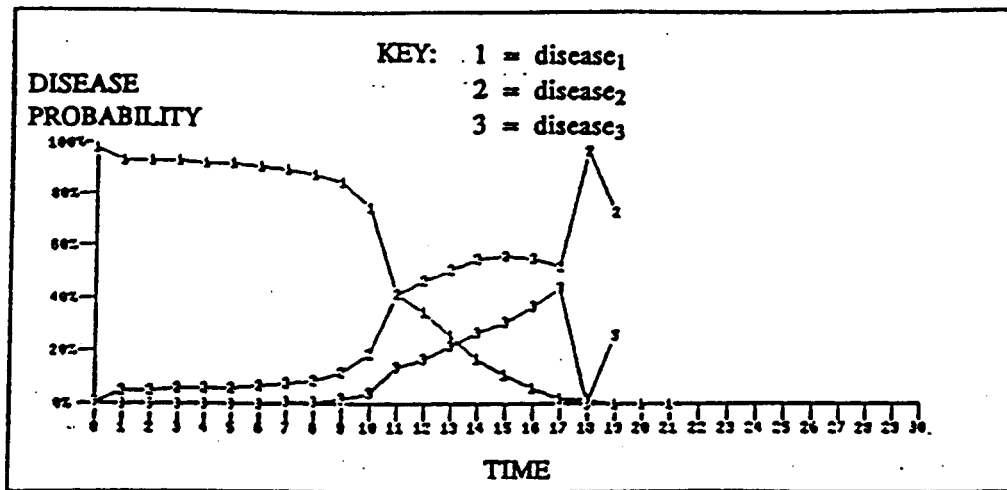


Figure 4. A pattern of evidence favoring disease₂, because e₁ occurs before e₂, and the two events occur close together.





Figure 5. A pattern of evidence favoring disease₃, because e₂ occurs before e₁.

Acknowledgements

We thank Ren Curry for his input and feedback on this work. We thank Joan Walton for her assistance in developing the graphical interface of the prototype. Lyn Dupre provided helpful comments on earlier versions of this manuscript.

References

1. Allen, J.F., Towards a general theory of action and time, *Artificial Intelligence*, 23 (1984) 123 - 154.
2. Allen, J.F and Hayes, P.J., A common-sense theory of time, In: *Proceedings of IJCAI-85* (Morgan Kaufmann, Los Altos, CA, 1985).
3. Blum, R.L., Discovery and representation of causal relationships from a large time-oriented clinical database: The RX project, In: D.A.B. Lindberg and P.L. Reichertz (Eds.), *Lecture Notes in Medical Informatics* (Springer-Verlag, New York, 1982).
4. Bolour, A., Anderson, T.L., Dekeyser, L.J., and Wong, H.K.T., The role of time in information processing: A survey, *ACM SIGART Newsletter*, 80 (1982) 28 - 48.
5. Cooper, G.F., Probabilistic inference using belief networks is NP-hard, Report KSL-87-27, Medical Computer Science Group, Stanford University, Stanford, CA, 1987.
6. Dean, T., Large-scale temporal databases for planning in complex domains, In: *Proceedings of IJCAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
7. Dean, T. and Kanazawa, K., A model for reasoning about persistence and causation, Research Report, Department of Computer Science, Brown University, Providence, RI, 1988.

8. de-Zegher-Geets, I., IDEFIX: Intelligent summarization of a time-oriented medical database, Report KSL-88-34, Medical Computer Science Group, Stanford University, Stanford, CA, 1987.
9. Fagan, L.M., Kunz, J.C., Feigenbaum, E.A., and Osborn, J.J., In: *Extensions to the Rule-Based Formalism for a Monitoring Task* (Addison-Wesley, Reading, MA, 1984).
10. Hall, R.J., Lathrop, R.H., Kirk, R.S., A multiple representation approach to understanding the time behavior of digital circuits, In: *Proceedings of AAAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
11. Hamlet, I. and Hunter, J., A representation of time for medical expert systems, Report AIMG-13, Artificial Intelligence in Medicine Group, University of Sussex, Brighton, 1987.
12. Haugh, B.A., Non-standard semantics for the method of temporal arguments, In: *Proceedings of IJCAI-87*, (Morgan Kaufmann, Los Altos, CA, 1987).
13. Higgins, M.C., NAIVE: A method for representing uncertainty and temporal relationships in an automated reasoner, In: *Proceedings of AAAI Workshop on Uncertainty in Artificial Intelligence*, 1987.
14. Hunter, J.R. and Sinnhuber, R.K., Representation of disease development, *Expert Systems* (1983) 174 - 183.
15. Kaemmerer, W.F. and Allard, J.R., An automated reasoning technique for providing moment-by-moment advice concerning the operation of a process, In: *Proceedings of AAAI-87*, (Morgan Kaufmann, Los Altos, CA, 1987).
16. Kahn, K. and Gorry, G.A., Mechanizing temporal knowledge, *Artificial Intelligence* 9 (1977) 87 - 108.

17. Kaihara, S. and Koyama, T., Medical consultation system with practical requirements — Development of MECS-AI, *Japan Annual Review of Electronics, Computers, and Telecommunications* 4 (1982) 322 - 333.
18. Kastner, J.K., Weiss, S.M., and Kulikowski, C.A., An efficient scheme for time-dependent consultation systems, In: J.H. van Bemmelen, M.J. Ball, and O. Wigertz (Eds.), *Proceedings of Medinfo*, Amsterdam, 1983.
19. Kumar, K., Temporal event conceptualization, In: *Proceedings of IJCAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
20. Ladkin, P., Models of axioms of time intervals, In: *Proceedings of AAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
21. Ladkin, P., The completeness of a natural system for reasoning with time intervals, In: *Proceedings of IJCAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
22. Long, W.J., Reasoning about state from causation and time in a medical domain, In: *Proceedings of AAI-83* (Morgan Kaufmann, Los Altos, CA, 1983).
23. McDermott, D., A temporal logic for reasoning about processes and plans, *Cognitive Science* 6 (1982) 101 - 155.
24. Mittal, S., Event-based organization of temporal databases, In: *Proceedings of the Fourth National Conference of the Canadian Society for Computational Studies of Intelligence*, 1982.
25. Pearl, J., Fusion, propagation, and structuring in belief networks, *Artificial Intelligence* 29 (1986) 241 - 288.
26. Pelavin, R.N. and Allen, J.F., A model for concurrent actions having temporal extent, In: *Proceedings of AAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).

27. Shibahara, T., Tsotos, J.K., Mylopoulos, J., and Covvey, H.D., CAA: A knowledge-based system using causal knowledge to diagnosis cardiac rhythm disorders, In: *Proceedings of IJCAI-83* (Morgan Kaufmann, Los Altos, CA, 1983).
28. Shortliffe, E.H., Scott, A.C., Bischoff, M., Campbell, A.B., van Melle, W., and Jacobs, C., ONCOCIN: An expert system for oncology protocol management, In: *Proceedings of IJCAI-84* (Morgan Kaufmann, Los Altos, CA, 1984).
29. Sinnhuber, R. and Hunter, J., Temporal information processing for medical expert systems: A review of recent work, Report AIMG-6, Artificial Intelligence in Medicine Group, University of Sussex, Brighton, 1984.
30. Tsang, E.P.K., The consistent labelling problem in temporal reasoning, In: *Proceedings of AAAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
31. Tsang, E.P.K., Time structures in AI, In: *Proceedings of IJCAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
32. Valdes-Perez, R.E., The satisfiability of temporal constraints, In: *Proceedings of AAAI-87* (Morgan Kaufmann, Los Altos, CA, 1987).
33. Weiss, J.M., Kulikowski, C.A., Amarel, S., and Safir, A., A model-based method for computer-aided medical decision making, *Artificial Intelligence* 11 (1978) 145 - 172.