ILL: 16293525	Call Number: Q 334.C65	DateReq: 1/25/2006	Yes
	Location:	Date Rec: 1/25/2006	No
	Maxcost: $25IFM	Borrower: RBN	Conditional
		Affiliation:	
Request Type:	Source: ILLiad	LenderString: *JHE,JHE,BOS,BXM,NHM	
OCLC Number: 12073389	DueDate:	Verified: <TN:201879> OCLC 0824-7935	

Email: SCI-ILL@brown.edu

Fax: 128.148.41.3 Fax 401-863-2753

Billing Notes:

Title: Computational intelligence = Intelligence informatique.

Uniform Title:

Author:

Edition:

Imprint: Ottawa : National Research Council of Canada = Conseil national de recherche

Article: Thomas Dean and Keiji Kanazawa: A Model for Reasoning About Persistance and Causation

Vol: 5 **No.:** 3 **Pages:** 142-150 **Date:** 1989

Dissertation:

Borrowing Notes: If copy, please send by Ariel or fax. Please do not batch Ariel transmissions. Thanks.

ShipTo: ILL/Sciences Library,201 Thayer St. entrance/Brown University/Providence, RI 02912-9109

Ship Via: ARIEL for copies

ShipVia: **ARIEL for copi**

Return To:

Milton S. Eisenhower Library/JHU
Interlibrary Services Department
3400 N. Charles Street
Baltimore, MD 21218 USA

Ship To:

ILL

Sciences Library,201 Thayer St. entrance

Brown University

Providence, RI 02912-9109

NeedBy: 2/24/2006

Borrower: RBN

ILL: **16293525** **Lender:** JHE

Req Date: **1/25/2006** **OCLC #:** 12073389

Patron: **Reed, Dawn**

Author:

Title: **Computational intelligence = Intelligence inform**

Article: Thomas Dean and Keiji Kanazawa: A Model for Reasoning About Persistance and Causation

Vol.: 5 **No.:** 3

Date: 1989 **Pages:** 142-150

Verified: <TN:201879> OCLC 0824-7935

Maxcost: $25IFM **Due Date:**

Lending Notes:

Bor Notes: If copy, please send by Ariel or fax. Please do not batch Ariel transmissions. Thanks.

A model for reasoning about persistence and causation

THOMAS DEAN AND KEIJI KANAZAWA

Department of Computer Science, Brown University, Box 1910, Providence, RI 02912, U.S.A.

Received October 11, 1988

Revision accepted September 25, 1989

Reasoning about change requires predicting how long a proposition, having become true, will continue to be so. Lacking perfect knowledge, an agent may be constrained to believe that a proposition persists indefinitely simply because there is no way for the agent to infer a contravening proposition with certainty. In this paper, we describe a model of causal reasoning that accounts for knowledge concerning cause-and-effect relationships and knowledge concerning the tendency for propositions to persist or not as a function of time passing. Our model has a natural encoding in the form of a network representation for probabilistic models. We consider the computational properties of our model by reviewing recent advances in computing the consequences of models encoded in this network representation. Finally, we discuss how our probabilistic model addresses certain classical problems in temporal reasoning (e.g., the frame and qualification problems).

Key words: temporal reasoning, causal reasoning, uncertainty, probabilistic models, probabilistic inference, belief networks.

Le raisonnement à propos des modifications nécessite de prédire combien de temps une proposition demeurera vraie une fois qu'elle l'est devenue. En l'absence de connaissances parfaites, un agent peut être forcé de croire qu'une proposition persiste indéfiniment simplement parce qu'il lui est impossible d'inférer avec certitude une proposition contraire. Dans cet article, nous décrivons un modèle de raisonnement causal qui tient compte des connaissances concernant la relation entre la cause et l'effet et la tendance des propositions à persister ou non en fonction du temps. Le modèle comporte un encodage naturel sous forme de représentation en réseau pour les modèles probabilistiques. Les propriétés informatiques du modèle sont prises en considération à la lumière des récents développements dans le domaine de l'estimation des conséquences des modèles encodés dans cette représentation en réseau. Enfin, on y discute comment le modèle probabilistique traite certains problèmes classiques de raisonnement temporel (par exemple, les problèmes de cadre et de qualification).

Mots clés : raisonnement temporel, raisonnement causal, incertitude, modèles probabilistiques, inférence probabilistique, réseau de croyances.

[Traduit par la revue]

Comput. Intell. 5, 142–150 (1989)

1. Introduction

The commonsense law of inertia (McCarthy 1986) states that a proposition once made true remains so until something makes it false. Given perfect knowledge of initial conditions and a complete predictive model, the law of inertia is sufficient for accurately inferring the persistence of propositions. In most circumstances, however, our predictive models and our knowledge of initial conditions are less than perfect. The law of inertia requires that, in order to infer that a proposition ceases to be true, we must predict an event with a contravening effect. Such predictions are often difficult to make. Consider the following examples:

- a cat is sleeping on the couch in our living room
- you can leave your umbrella on the 8:15 commuter train
- a client on the telephone is asked to hold

In each case, there is some proposition initially observed to be true, and the task is to determine if it will be true at some later time. The cat may sleep undisturbed for an hour or more, but it is extremely unlikely to remain in the same spot for more than six hours. Your umbrella will probably not be sitting on the seat when you catch the train the next morning. The client will probably hold for a few minutes, but only the most determined of clients will be on the line after 15 minutes. Sometimes we can make more accurate predictions (e.g., a large barking dog runs into the living room), but, lacking specific evidence, we would like past experience to provide an estimate of how long certain propositions are likely to persist.

Events precipitate change in the world, and it is our knowledge of events that enables us to make useful predictions about the future. For any proposition, P, that can hold in a situation, there are some number of general sorts of events (referred to as *event types*) that can affect P (i.e., make P true or false). For any particular situation, there are some number of specific events (referred to as *event instances*) that occur. Let O correspond to the set of events that occur at time t, A correspond to that subset of O that affect P, $K(O)$ that subset of O known to occur at time t, and $K(A)$ that subset of A whose type is known to affect P. Figure 1 illustrates how these sets might relate to one another in a specific situation. In many cases, $K(O) \cap K(A)$ will be empty while A is not, and it may still be possible to provide a reasonable assessment of whether or not P is true at t. In this paper, we provide an account of how such assessments can be made probabilistically.

2. Prediction and persistence

In the following, we distinguish between two kinds of propositions: propositions, traditionally referred to as *fluents* (McCarthy and Hayes 1969), that, if they become true, tend to persist without additional effort, and propositions, corresponding to the occurrence of events, that, if true at a point, tend to precipitate or trigger change in the world.

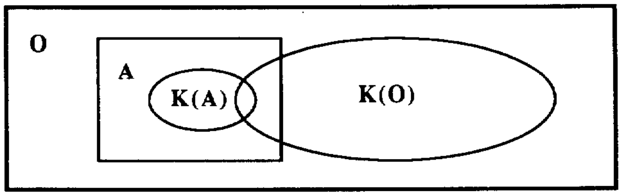

FIG. 1. Events precipitate change in the world.

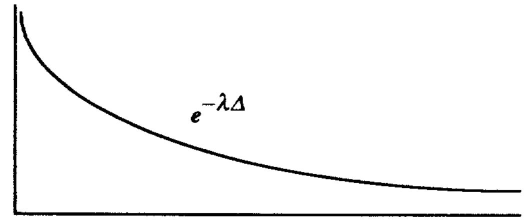

FIG. 2. A survivor function with exponential decay.

Let $\langle P, t\rangle$ indicate that the fluent P is true at time t, and $\langle E, t\rangle$ indicate that an event of type E occurs at time t. We use the notation E_P to indicate an event corresponding to the fluent P becoming true.

Given our characterization of fluents as propositions that tend to persist, whether or not P is true at some time t may depend upon whether or not it was true at $t - \Delta$, where $\Delta > 0$. We can represent this dependency as follows:[1]

$$[1] \quad p(\langle P, t\rangle) = p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle)p(\langle P, t - \Delta\rangle) + p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle)p(\neg\langle P, t - \Delta\rangle)$$

where $\neg\langle P, t\rangle \equiv \langle\neg P, t\rangle$.

The conditional probabilities $p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle)$ and $p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle)$ are related to the *survivor function* in classical queuing theory (Syski 1979). Survivor functions encode the changing expectation of a fluent remaining true over the course of time. We employ survivor functions to capture the tendency of propositions to become false as a consequence of events with contravening effects. With survivor functions, one need not be aware of a specific instance of an event with a contravening effect in order to predict that P will cease being true. As an example of a survivor function,

$$p(\langle P, t\rangle) = e^{-\lambda\Delta}\, p(\langle P, t - \Delta\rangle)$$

indicates that the probability that P persists drops off as a function of the time since P was last observed to be true at an exponential rate determined by λ (Fig. 2). The exponential decay survivor function is equivalent to the case where

$$p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle) = e^{-\lambda\Delta}$$

and

$$p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle) = 0$$

Referring back to Fig. 1, survivor functions account for that subset of A corresponding to events that make P false, assuming that $K(A) = \{\}$.

If we have evidence concerning specific events known to affect P (i.e., $K(A) \cap K(O) \neq \{\}$), [1] is inadequate. As an interesting special case of how to deal with events known to affect P, suppose that we know about all events that make P true (i.e., we know $p(\langle E_p, t\rangle)$ for any value of t), and none of the events that make P false. In particular, suppose that P corresponds to John being at the airport, and E_P corresponds to the arrival of John's flight. We are interested in whether or not John will still be waiting at the airport when we arrive to pick him up. Let $\sigma(t) = e^{-\lambda t}$ represent John's tendency to hang around airports, where λ is a measure of his impatience. If $f(t) = p(\langle E_P, t\rangle)$, then we can compute the probability of P being true at t by convolving f with the survivor function σ as in

$$[2] \quad p(\langle P, t\rangle) = \int_{-\infty}^{t} p(\langle E_p, z\rangle)\sigma(t - z)\, dz$$

A shortcoming of [2] is that it fails to account for evidence concerning specific events known to make P false. Suppose, for instance, that $E_{\neg P}$ corresponds to Fred meeting John at the airport and giving him a ride to his hotel. In certain cases,

$$[3] \quad p(\langle P, t\rangle) = \int_{-\infty}^{t} p(\langle E_P, z\rangle)\sigma(t - z) \times \left[1 - \int_{z}^{t} p(\langle E_{\neg P}, x\rangle)\, dx\right] dz$$

provides a good approximation. Figure 3 illustrates the sort of inference licensed by [3].

There are some potential problems with [3]. The survivor function, σ, was meant to account for all events that make P false, but [3] counts one such event, John leaving the airport with Fred, twice: once in the survivor function and once in $p(\langle E_{\neg P}, t\rangle)$. In certain cases, this can lead to significant errors (e.g., Fred always picks up John at the airport). To combine the available evidence correctly, it will help if we distinguish the different sorts of knowledge that might be brought to bear on estimating whether or not P is true. We will also reinterpret the event type E_P to mean an event *known* to make P true. The following formula makes the necessary distinctions and indicates how the evidence should be combined:[2]

$$\begin{aligned}
[4] \quad p(\langle P, t\rangle) = {} & \\
& p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle \wedge \neg(\langle E_P, t\rangle \vee \langle E_{\neg P}, t\rangle)) && \text{(N1)}\\
& \quad \times p(\langle P, t - \Delta\rangle \wedge \neg(\langle E_P, t\rangle \vee \langle E_{\neg P}, t\rangle)) \\
& + p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle \wedge \langle E_P, t\rangle) && \text{(N2)}\\
& \quad \times p(\langle P, t - \Delta\rangle \wedge \langle E_P, t\rangle) \\
& + p(\langle P, t\rangle \mid \langle P, t - \Delta\rangle \wedge \langle E_{\neg P}, t\rangle) && \text{(N3)}\\
& \quad \times p(\langle P, t - \Delta\rangle \wedge \langle E_{\neg P}, t\rangle) \\
& + p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle \wedge \neg(\langle E_P, t\rangle \vee \langle E_{\neg P}, t\rangle)) && \text{(N4)}\\
& \quad \times p(\neg\langle P, t - \Delta\rangle \wedge \neg(\langle E_P, t\rangle \vee \langle E_{\neg P}, t\rangle)) \\
& + p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle \wedge \langle E_P, t\rangle) && \text{(N5)}\\
& \quad \times p(\neg\langle P, t - \Delta\rangle \wedge \langle E_P, t\rangle) \\
& + p(\langle P, t\rangle \mid \neg\langle P, t - \Delta\rangle \wedge \langle E_{\neg P}, t\rangle) && \text{(N6)}\\
& \quad \times p(\neg\langle P, t - \Delta\rangle \wedge \langle E_{\neg P}, t\rangle)
\end{aligned}$$

Consider the contribution of the individual terms corresponding to the conditional probabilities labeled N1 through N6 in [4]. N1 accounts for *natural attrition*: the tendency for propositions to become false given no direct evidence of events known to affect P. N2 and N5 account for *causal accretion*: accumulating evidence for P due to events known

[1]The equality in [1] follows from the *generalized addition law*: if $A_1, \ldots, A_n$ are exclusive and exhaustive and B is any event, then

$$p(B) = \sum_{i=1}^{n} p(B \mid A_i)p(A_i).$$

[2]To justify our use of the generalized addition law in [4], we assume that $p(\langle E_P, t\rangle \wedge \langle E_{\neg P}, t\rangle) = 0$ for all t.

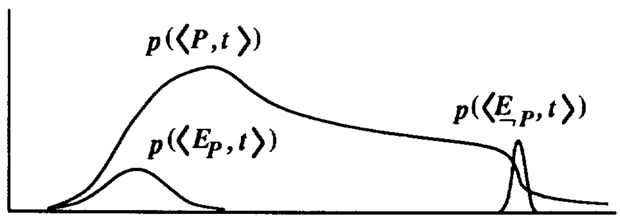

FIG. 3. Probabilistic predictions.

to make P true. N2 and N5 are generally 1. N3 and N6, on the other hand, are generally 0, since evidence of $\neg P$ becoming true does little to convince us that P is true. Finally, N4 accounts for *spontaneous causation*: the tendency for propositions to suddenly become true with no direct evidence of events known to affect P.

By using a discrete approximation of time and fixing Δ, it is possible both to acquire the necessary values for the terms N1 through N6 and to use them in making useful predictions (Dean and Kanazawa 1987). If time is represented as integers, and $\Delta = 1$, we note that the law of inertia applies in those situations in which the terms N1, N2, and N5 are always 1 and the other terms are always 0. In the rest of this paper, we assume that time is discrete and linear and that the time separating any two consecutive time points is some constant δ (see Cooper *et al.* (1988) for a discussion of a related approach to probabilistic reasoning about change using a discrete model of time). Only evidence concerning events *known* to make P true is brought to bear on $p(\langle E_P, t\rangle)$. If $p(\langle E_P, t\rangle)$ were used to summarize all evidence concerning events that make P true, then N1 would be 1.

3. Reasoning about causation

Before we consider the issues involved in making predictions using knowledge concerning N1 through N6, we need to add to our theory some means of predicting additional events. We consider the case of one event causing another event. Deterministic theories of causation often use implication to model cause-and-effect relationships. For instance, to indicate that the occurrence of an event of type E_1 at time t causes the occurrence of an event of type E_2 following t by some $\delta > 0$ just in case the conjunction $P_1 \wedge P_2 \wedge \ldots \wedge P_n$ holds at t, we might write

$$(\langle P_1 \wedge P_2 \wedge \ldots \wedge P_n, t\rangle \wedge \langle E_1, t\rangle) \supset \langle E_2, t + \delta\rangle$$

If the caused event is of a type E_P, this is often referred to as *persistence causation* (McDermott 1982). In our model, the conditional probability

$$p(\langle E_2, t + \delta\rangle \mid \langle P_1 \wedge P_2 \wedge \ldots \wedge P_n, t\rangle \wedge \langle E_1, t\rangle) = \pi$$

is used to indicate that, given an event of type E_1 occurs at time t, and P_1 through P_n are true at t, an event of type E_2 will occur following t by some $\delta > 0$ with probability π.

In moving to a probabilistic model of causation, there are some complications that we have to deal with. Consider, for example, the two rules:

$$(\langle P, t\rangle \wedge \langle E, t\rangle) \supset \langle E_R, t + \delta\rangle$$

and

$$(\langle P \wedge Q, t\rangle \wedge \langle E, t\rangle) \supset \langle E_R, t + \delta\rangle$$

These two rules pose no problems for the deterministic theory of causation, since P and Q are either true or false, and the rules either apply or not. In fact, the second rule is redundant. However, in a probabilistic model, P and Q usually are not unambiguously true or false. Therefore, in the probabilistic causal theory consisting of

$$p(\langle E_R, t + \delta\rangle \mid \langle P, t\rangle \wedge \langle E, t\rangle) = \pi_1$$

and

$$p(\langle E_R, t + \delta\rangle \mid \langle P \wedge Q, t\rangle \wedge \langle E, t\rangle) = \pi_2$$

the second rule can no longer be considered redundant. Since the second rule is more specific than the first, it provides us with valuable additional information. In a complete account of the causes for E_R, we would also need

$$P(\langle E_R, t + \delta\rangle \mid \langle P \wedge \neg Q, t\rangle \wedge \langle E, t\rangle) = \pi_3$$

and other information as well. Providing a complete account of the interactions among causes and between causes and their effects is important in modeling change in a probabilistic framework. In the following two sections, we will consider this issue in more detail.

4. An example

The task in *probabilistic projection* is to assign each propositional variable of the form $\langle \varphi, t\rangle$ a certainty measure consistent with the constraints specified in a problem. In this section, we provide examples drawn from a simple factory domain that illustrate the sort of inference required in probabilistic projection. We begin by introducing some new event types:

Cl = "The mechanic on duty cleans up the shop"
As = " Fred tries to assemble Widget17 in Room101"

and fluents:

Wr = "The location of Wrench14 is Room101"
Sc = "The location of Screwdriver31 is Room101"
Wi = "Widget17 is completely assembled"

We assume that tools are occasionally displaced in a busy shop, and that Wr and Sc are both subject to an exponential persistence decay with a half life of 1 day; this determines N1 in [4]:

$$p(\langle \mathrm{Wr}, t\rangle \mid \langle \mathrm{Wr}, t - \Delta\rangle \wedge \neg(\langle E_{\mathrm{Wr}}, t\rangle \vee \langle E_{\neg\mathrm{Wr}}, t\rangle)) = e^{-\lambda\Delta}$$
$$p(\langle \mathrm{Sc}, t\rangle \mid \langle \mathrm{Sc}, t - \Delta\rangle \wedge \neg(\langle E_{\mathrm{Sc}}, t\rangle \vee \langle E_{\neg\mathrm{Sc}}, t\rangle)) = e^{-\lambda\Delta}$$

where $e^{-\lambda\Delta} = 0.5$ when Δ is 1 day.

The other terms in [4], N2, N3, N4, N5, and N6, we will assume to be, respectively, 1, 0, 0, 1, and 0. When the mechanic on duty cleans up the shop, he is supposed to put all of the tools in their appropriate places. In particular, Wrench14 and Screwdriver31 are supposed to be returned to Room101. We assume that the mechanic is very diligent:

$$p(\langle E_{\mathrm{Wr}}, t + \epsilon\rangle \mid \langle \mathrm{Cl}, t\rangle) = 1.0$$
$$p(\langle E_{\mathrm{Sc}}, t + \epsilon\rangle \mid \langle \mathrm{Cl}, t\rangle) = 1.0$$

Fred's competence is assembling widgets depends upon his tools being in the right place. In particular, if Screwdriver31 and Wrench14 are in Room101, then it is certain that Fred will successfully assemble Widget17.

$$p(\langle E_{\mathrm{Wi}}, t + \epsilon\rangle \mid \langle \mathrm{Wr}, t\rangle \wedge \langle \mathrm{Sc}, t\rangle \wedge \langle \mathrm{As}, t\rangle) = 1.0$$

Let T0 correspond to 12:00 p.m., February 29, 1988, and T1 correspond to 12:00 p.m. on the following day. Assume that ϵ is negligible given the events we are concerned with (i.e., we will add or subtract ϵ in order to simplify the analysis).

$$p(\langle \text{Cl, T0} \rangle) = 0.7$$
$$p(\langle \text{As, T1} \rangle) = 1.0$$

We are interested in assigning the propositions of the form $\langle \varphi, t \rangle$ a certainty measure consistent with the axioms of probability theory. We will work through an example showing how one might derive such a measure, noting some of the assumptions required to make the derivations follow from the problem specification and the axioms of probability. In the following, we will denote this measure of belief by BEL. What can we say about BEL($\langle$Wi, T1 + $\epsilon\rangle$)? In this particular example, we begin with

$$\begin{aligned}
&\text{BEL}(\langle \text{Wi, T1} + \epsilon \rangle) \\
&\quad = p(\langle E_{\text{Wi}}, \text{T1} + \epsilon \rangle) \\
&\quad = p(\langle E_{\text{Wi}}, \text{T1} + \epsilon \rangle \mid \langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle \wedge \langle \text{As, T1} \rangle) \\
&\qquad\qquad \times p(\langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle \wedge \langle \text{As, T1} \rangle) \\
&\quad = p(\langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle \wedge \langle \text{As, T1} \rangle) \\
&\quad = p(\langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle)
\end{aligned}$$

The first step follows from our interpretation of E_{Wr}, and the fact that there is no additional evidence for or against Wr at T1 + ϵ. The second step employs the addition rule and the assumption that the assembly will fail to have the effect of $\langle E_{\text{Wr}}, \text{T1} \rangle$ if any one of $\langle$Wr, T1$\rangle$, $\langle$Sc, T1$\rangle$, or $\langle$As, T1$\rangle$ is false. The third step relies on the fact that assembly is always successful given that the attempt is made and Wrench14 and Screwdriver31 are in Room101. The last step depends on the assumption that the evidence supporting $\langle$Wr $\wedge$ Sc, T1$\rangle$ and $\langle$As, T1$\rangle$ are independent. The assumption is warranted in this case given that the particular instance of As occurring at T1 does not affect Wr $\wedge$ Sc at T1, and the evidence for As at T1 is independent of any events prior to T1. Note that if the evidence for As at T1 involved events prior to T1, then the analysis would be more involved. It is clear that p($\langle$Wr, T1$\rangle$) $\geq$ 0.35, and that p($\langle$Sc, T1$\rangle$) $\geq$ 0.35; unfortunately, we cannot simply combine this information to obtain an estimate of $p(\langle \text{Wr} \wedge \text{Sc, T1} \rangle)$, since the evidence supporting these two claims is dependent. We can, however, determine that

$$\begin{aligned}
&p(\langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle) \\
&\quad = p(\langle \text{Wr, T1} \rangle \wedge \langle \text{Sc, T1} \rangle \mid \langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \\
&\qquad\qquad \times p(\langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \\
&\quad = p(\langle \text{Wr, T1} \rangle \mid \langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \\
&\qquad \times p(\langle \text{Sc, T1} \rangle \mid \langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \\
&\qquad\qquad \times p(\langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \\
&\quad = p(\langle \text{Wr, T0} \rangle \wedge \langle \text{Sc, T0} \rangle) \times 0.5 \times 0.5 \\
&\quad = p(\langle E_{\text{Wr}}, \text{T0} \rangle \wedge \langle E_{\text{Sc}}, \text{T0} \rangle) \times 0.5 \times 0.5 \\
&\quad = p(\langle E_{\text{Wr}}, \text{T0} + \epsilon \rangle \wedge \langle E_{\text{Sc}}, \text{T0} + \epsilon \rangle \mid \langle \text{Cl, T0} \rangle) \\
&\qquad\qquad \times p(\langle \text{Cl, T0} \rangle) \times 0.5 \times 0.5 \\
&\quad = 0.7 \times 0.5 \times 0.5 \\
&\quad = 0.175
\end{aligned}$$

assuming that there is no evidence concerning events that are known to affect either Wr or Sc in the interval from T0 to T1, that Wr and Sc are independent, and that E_{Wr} and E_{Sc} are conditionally independent of one another given Cl.

Throughout our analysis, we were forced to make assumptions of independence. In many cases, such assumptions are unwarranted or introduce inconsistencies. The inference process is further complicated by the fact that probabilistic constraints tend to propagate both forward and backward in time. This bidirectional flow of evidence can render the analysis described above useless. In the next section, we consider a model that simplifies specifying independence assumptions, and that allows us to handle both forward and backward propagation of probabilistic constraints.

5. A model for reasoning about change

In this section, we take a slight modification of [4] as the basis for a model of persistence. Formula [4] predicts $\langle P, t \rangle$ on the basis of $\langle P, t - \Delta \rangle$, $\langle E_P, t \rangle$, and $\langle E_{\neg P}, t \rangle$, where Δ is allowed to vary. In the model presented in this section, we only consider pairs of consecutive time points, t and $t + \delta$, and arrange things so that the value of a fluent at time t is completely determined by the state of the world at δ in the past. In [4], we interpret events of type E_P occurring at t as providing evidence for P being true at t. In our new model, we interpret events of type E_P occurring at t as providing evidence of P being true at $t + \delta$. This reinterpretation is not strictly necessary, but we prefer it since the expressiveness of the resulting models can easily be characterized in terms of the properties of Markov processes. In our new model, we predict $A \equiv \langle P, t + \delta \rangle$ by conditioning on

$$C_1 \equiv \langle P, t \rangle$$
$$C_2 \equiv \langle E_P, t \rangle$$
$$C_3 \equiv \langle E_{\neg P}, t \rangle$$

and specify a complete model for the persistence of P as

$$p(A) = \Sigma\, p(A \mid C_1 \wedge C_2 \wedge C_3)\, p(C_1 \wedge C_2 \wedge C_3)$$

where the sum is over the eight possible truth assignments for the variables C_1, C_2, and C_3. Note that this model requires that we have probabilities of the form $p(A \mid C_1 \wedge C_2 \wedge C_3)$ and $p(C_1 \wedge C_2 \wedge C_3)$ for all possible valuations of the C_i.

In the following, we will make use of graphical representation for probabilistic models that will serve to clearly indicate the assumptions concerning dependence and independence underlying our models. The graphical representations that we will be using have been called *Bayes nets* (Pearl 1988), *belief networks* (Duda *et al.* 1981), and *influence diagrams* (Howard and Matheson 1984). We will use the generic term belief network to refer to a network that satisfies the following basic properties common to all three of the above representations. A belief network represents the variables or propositions of a probabilistic theory as nodes in a graph. The variables in our networks correspond to propositional variables of the form $\langle \varphi, t \rangle$. Dependence between two variables is indicated by a directed arc between the two nodes associated with the variables.

Because dependence is always indicated by an arc, belief networks make it easy to identify the conditional independence inherent in a model simply by inspecting the graph. Two nodes that are linked via a common neighbor, but for which there are no other connecting paths, are conditionally independent given the common node. For instance, in the models described in this section, $\langle P, t - \delta \rangle$ is independent of $\langle P, t + \delta \rangle$ given $\langle P, t \rangle$. Belief networks make it easy to construct and verify the correctness and reasonableness of

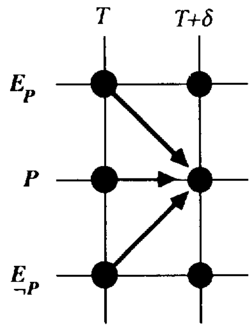

FIG. 4. The evidence for P at time $T + \delta$.

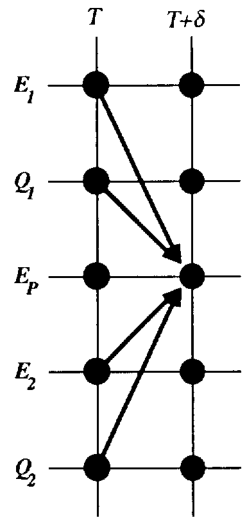

FIG. 5. The evidence for E_P at time $T + \delta$.

a model directly in terms of the corresponding graphical representation. Our model for persistence can be represented by the network shown in Fig. 4. As soon as we provide a model for causation, we will show how this simple model for persistence can be embedded in a more complex model for reasoning about change over time.

Generally, we expect that the cause-and-effect relations involving E_P will be specified in terms of constraints of the form:

$$p(\langle E_P, t + \delta\rangle \mid \langle E_1, t\rangle \wedge \langle Q_1, t\rangle) = \pi_1$$
$$p(\langle E_P, t + \delta\rangle \mid \langle E_2, t\rangle \wedge \langle Q_2, t\rangle) = \pi_2$$
$$\vdots$$
$$p(\langle E_P, t + \delta\rangle \mid \langle E_n, t\rangle \wedge \langle Q_n, t\rangle) = \pi_n$$

However, to specify a complete model, we will need some more information. To predict $A \equiv \langle E_P, t + \delta\rangle$, we condition on

$$C_1 \equiv \langle E_1, t\rangle \wedge \langle Q_1, t\rangle$$
$$C_2 \equiv \langle E_2, t\rangle \wedge \langle Q_2, t\rangle$$
$$\vdots$$
$$C_n \equiv \langle E_n, t\rangle \wedge \langle Q_n, t\rangle$$

and specify a complete model as

$$p(A) = \Sigma\, p(A \mid C_1 \wedge C_2 \wedge \ldots \wedge C_n)\, p(C_1 \wedge C_2 \wedge \ldots \wedge C_n)$$

Note that we need on the order of 2^n probabilities corresponding to the 2^n possible valuations of the propositional variables C_1 through C_n to specify this model. The

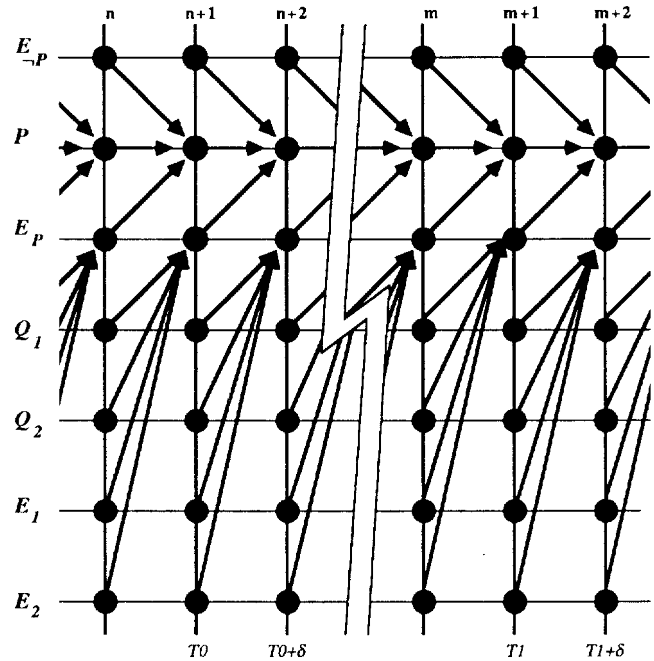

FIG 6. A temporal belief network.

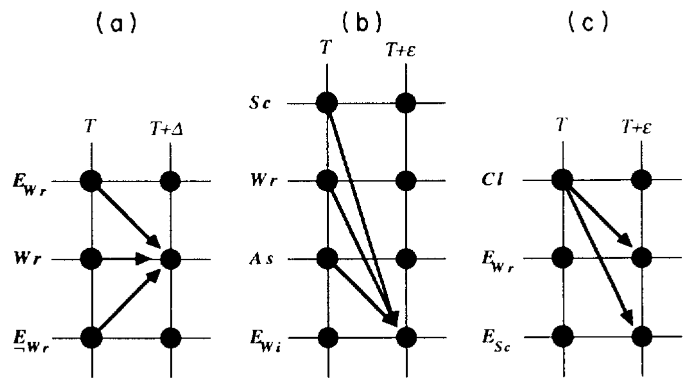

FIG. 7. Models for the factory example: (*a*) persistence of Wr; (*b*) effects of the assembling action; (*c*) effects of the cleaning action.

associated belief network is shown in Fig. 5. Similar networks would be constructed for event types other than those involving propositions becoming true or false.

Now we can construct a complete model for reasoning about change over time. Figure 6 illustrates the temporal belief network for such a complete model. For each propositional variable of the form $\langle \varphi, t\rangle$, there is a node in the belief network. The arcs are specified according to the isolated models for persistence and causation illustrated in Figs. 4 and 5. Following Pearl (1988), we can write down the unique distribution corresponding to the model shown in Fig. 6 as

$$p(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} p(x_i \mid S_i)p(S_i)$$

where x_i denotes the propositional variables in the model, and S_i is the conjunction of the propositional variables associated with those nodes for which there exists arcs to x_i in the network. How we compute this distribution is the subject of Sect. 6.

TABLE 1. Data for modeling the persistence of Wr

$p(\langle \text{Wr}, t\rangle \mid \ldots)$	$\langle \text{Wr}, t - \Delta\rangle$	$\langle E_{\text{Wr}}, t - \Delta\rangle$	$\langle E_{\neg\text{Wr}}, t - \Delta\rangle$
$e^{-\lambda\Delta}$	True	False	False
$e^{-\lambda\Delta}$	True	True	False
0.0	True	False	True
0.0	False	False	False
$e^{-\lambda\Delta}$	False	True	False
0.0	False	False	True
—	True	True	True
—	False	True	True

TABLE 2. Data for modeling the effects of the assembling action

$p(\langle E_{\text{Wi}}, t\rangle \mid \ldots)$	$\langle \text{Sc}, t - \epsilon\rangle$	$\langle \text{Wr}, t - \epsilon\rangle$	$\langle \text{As}, t - \epsilon\rangle$
0.0	False	False	False
0.0	True	False	False
0.0	False	True	False
0.0	True	True	False
0.0	False	False	True
0.0	True	False	True
0.0	False	True	True
1.0	True	True	True

TABLE 3. Data for modeling the effect of cleaning on Wr

$p(\langle E_{\text{Wr}}, t\rangle \mid \ldots)$	$\langle \text{Cl}, t - \epsilon\rangle$
0.0	False
1.0	True

TABLE 4. Data for modeling the effect of cleaning on Sc

$p(\langle E_{\text{Sc}}, t\rangle \mid \ldots)$	$\langle \text{Cl}, t - \epsilon\rangle$
0.0	False
1.0	True

As a specific instance of a temporal belief network, we reconsider the factory example of Sect. 4. We will need models for the persistence of wrenches and screwdrivers remaining in place, and models for reasoning about the consequences of cleaning and assembling actions. Figure 7*a* shows a portion of a belief network dedicated to modeling the persistence of Wr (i.e., the proposition corresponding to Wrench14 being in Room101). To completely specify the model for Wr persisting, we need the information given in Table 1. The first six entries in the table correspond to terms N1–6 in [4]. Note that the entries corresponding to N2 and N5 — assumed to be 1 in Sect. 4 — are now the same as N1 to account for our revised interpretation of events of type E_P. Figure 7*b* shows a portion of a belief network for modeling the effects of the assembly action. The complete model is specified in Table 2. Finally, Fig. 7*c* shows a portion of a belief network for modeling the effects of the cleaning action. The complete model for the effect of cleaning on the location of Wrench14 is shown in Table 3 and similarly for the effect of cleaning on the location of Screwdriver31 in Table 4.

In the discussion of the general model, the amount of time separating time points was assumed to be the same for all pairs of consecutive time points. In reasoning about the factory example, it will be useful to have the time separating pairs of consecutive time points differ, and to have different models for handling different separations. We will need time points close together for propagating the (almost immediate) consequences of actions, and time points separated by several hours so as not to incur the computational expense of reasoning about intervals of time during which little of interest happens. To reduce the complexity of the network for the factory example, we assume that evidence concerning the occurrence of actions such as cleaning and assembling is always with regard to the end points of 24 h intervals. Figure 8 shows the complete network for the factory example. Note that, since the evidence for actions appears only at 24 h intervals, we encode the models for action only at the time points T0 and T1; similarly, since additional evidence for events of type E_P is only available at T0 + ϵ and T1 + ϵ, we use a simpler model for persistence at T0 and T1 in which, for example, $\langle \text{Wr}, \text{T0} + \epsilon\rangle$ is completely determined by $\langle \text{Wr}, \text{T0}\rangle$. If we assume a prior probability of 0 for all nodes without predecessors in Fig. 8 excepting $\langle \text{Cl}, \text{T0}\rangle$ and $\langle \text{As}, \text{T1}\rangle$ which are, respectively, 0.7 and 1.0, then $p(\langle E_{\text{Wi}}, \text{T1} + \epsilon\rangle)$ is 0.175 in the unique posterior distribution determined by the network. This is the same as that established by the analysis of Sect. 4, but, in this case, we have made all of our assumptions of independence explicit in the structure of the temporal belief network.

It is straightforward to extend the model described above to account for new observations and updating beliefs. Suppose we have the observations $o_1, o_2, \ldots, o_n$, where each observation is of the form $\langle O, t\rangle$ and O is an event type corresponding to a particular type of observation. We assume some prior distribution specified in terms of constraints of the form:

$$p(\langle O, t\rangle) = 0.001$$

There are also constraints indicating prior belief regarding the occurrence of events other than observations. For instance, we might have

$$p(\langle E, t\rangle) = 0.001$$

Observations are related to events by constraints such as

$$p(\langle E, t\rangle \mid \langle O, t\rangle) = 0.70$$

and

$$p(\langle E, t\rangle \mid \neg\langle O, t\rangle) = 0.025$$

To update an agent's beliefs you can either change the priors,

$$p(\langle O, t\rangle) = 1.0$$

or compute the posterior distribution,

$$\text{BEL}(A) = p(A \mid o_1, o_2, \ldots, o_n)$$

Most of the standard techniques for representing and reasoning about evidence in belief networks apply directly to our model.

This paper is primarily concerned with presenting a particular model for reasoning about change. Our current research involves applying this model to problems in robotics and exploring the expressive limitations of our model. While we have only begun the latter research, a word about expressive limitations is probably in order to give the reader some idea of the potential power of our model and to encourage analysis by others. Most of the initial work in analyzing the expressive power of our model has been concerned with relating the predictive power of our model to that of Markov processes and in particular Markov chains.

Suppose that the instantaneous state of the world can be completely specified in terms of a vector of values assigned to a finite set of Boolean variables $\mathcal{P} = \{P_1, P_2, \ldots, P_n\}$, and suppose further that the environment can be accurately modeled as a Markov process in which time is discrete and the state space, Ω, corresponds to all possible valuations of the variables in $\mathcal{P}$. Given such a model including a transition matrix defined on Ω, we can generate a temporal belief network to compute the probability of any proposition in $\mathcal{P}$ being true at any time t based upon evidence concerning the values of variables in $\mathcal{P}$ at various times, and do so in accord with the transition probabilities specified in the Markov model. Conversely, given a temporal belief network such that, for all t and $P \in \mathcal{P}$, all of the predecessors of $\langle P, t\rangle$ are in the set $\{\langle P_i, t - \delta\rangle\}$, the network is said to satisfy the Markov property for temporal belief networks, and, from this network, one can construct an equivalent Markov chain. Nunez (1989) provides a proof of the equivalence of networks satisfying the Markov property for temporal belief networks and Markov chains, and provides an algorithm for translating between the two representations.

The reason that one might use a fluent-and-event-based temporal belief network model rather than an equivalent state-based Markov model is because the belief network representation facilitates reasoning of the sort required for applications in planning and decision support (e.g., computing answers to questions of the form, "What is the probability of P at t given everything else we know about the situation?"). These same answers can be computed using the Markov model, but the process is considerably less direct.

Satisfying the Markov property for temporal belief networks allows us to establish the connection between temporal belief networks and Markov chains, but it sometimes results in unintuitive network structures. Introducing a delay between an action and its consequences may appear reasonable given the intuition that causes precede effects. However, introducing a delay betwen E_P and P simply to ensure the Markov property may seem a little extreme. We can eliminate the delay between E_P and P by returning to the model for persistence in [4]. The resulting networks do not satisfy the Markov property described in this section, but they are perfectly legitimate temporal belief nets and provide a somewhat more intuitive model for representing change than networks that do satisfy the Markov property.

6. Computing probabilistic predictions

In this section, we consider the computational issues concerned with our model of temporal reasoning. We will chiefly be concerned with computing the joint conditional probability distribution for a fully specified belief network, given the available information. The problem of computing the distribution for a partial, incomplete model will be considered briefly at the end of this section.

For networks with restricted topologies, there are efficient methods for computing the joint distribution. A belief network is said to be *singly connected* if there is at most one path between any two nodes in the graph. The class of singly connected networks are basically those which form treelike topologies; both proper trees and trees in which a node may have more than one parent (called a *generalized Chow tree* (Chow and Liu 1968)). Kim and Pearl (1983) have described an efficient mechanism for computing the joint distribution for singly connected networks by local propagation of Bayes factors (i.e., the likelihood ratios $p(x)/p(\neg x)$).

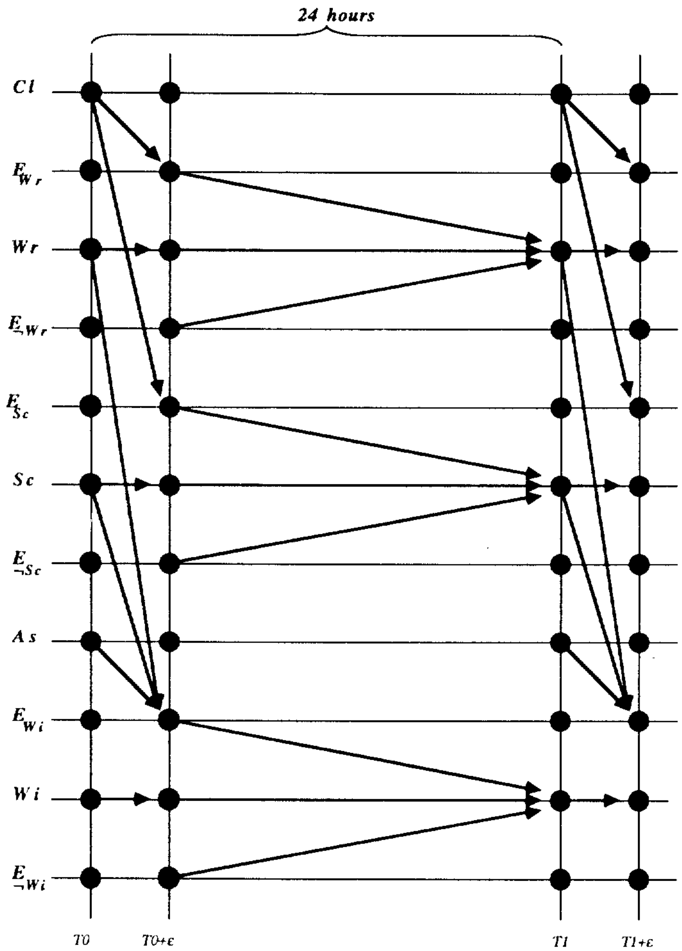

FIG. 8. A belief network for the factory example.

The Kim–Pearl algorithm is linear in the diameter of the network. However, it is restricted to singly connected networks. Unfortunately, most interesting causal theories, such as those arising in medicine, and our own temporal models involve *multiple connected* networks; networks that contain cycles. The general problem of probabilistic inference in belief networks, including multiple connected networks, has been shown to be NP-hard (Cooper 1987). Although this points towards some possible problems in terms of computational efficiency, it should not discourage us from using belief networks for probabilistic inference. The NP-hardness result tells us that a search for an efficient exact algorithm for all cases is probably misguided, but it says little about what to expect in practice. In general, asymptotic complexity results should serve as a guide in algorithm design, not as a justification for abandoning a research area. Such results tell us to look for algorithms that have good expected-case performance and to consider possible tradeoffs in terms of efficiency, soundness, or completeness.

There have been three principal types of algorithms proposed for probabilistic inference with general belief networks: *exact*, *bounding*, and *approximation* algorithms. The exact algorithms attempt to exploit the topology of the network; they include the method of *conditioning* (Pearl 1988) and the Lauritzen–Spiegelhalter (L–S) algorithm (Lauritzen and Spiegelhalter 1988). In a multiply connected network, if a node which has multiple children is instantiated to a given value, then the values of the children become independent of one another. The path(s) connecting the children through the parent node become "blocked," and the connectivity of the graph changes. If enough nodes are instantiated, then the resulting network becomes singly connected, so that the Bayes propagation algorithm can be used. A set of nodes which, when instantiated, render a multiply connected network singly connected corresponds to a cutset of the underlying graph. Conditioning computes a weighted average of the joint conditional probability distributions for all possible instantiations of such a cutset (for further details, see Pearl (1988)).

In the L–S algorithm, a set of subsidiary arcs is added to the graph to *triangulate* the graph, and the joint distribution can be efficiently computed in terms of the cliques of the original graph. A graph is triangulated if there are no cycles of length 4 or more without a chord or shortcut. L–S represents the network in terms of a clique tree, a tree of the cliques in the triangulated graph. Computing in terms of the tree is very efficient, as with the Kim–Pearl algorithm; the L–S algorithm alters the connectivity of a belief network much like conditioning does, rendering a multiply connected network singly connected so that an efficient Kim–Pearl type propagation algorithm may be employed. The joint distribution for the network shown in Fig. 8 can be computed in just a few seconds using the L–S algorithm running on a workstation. An important characteristic of the L–S algorithm is that after a relatively time-consuming initialization step for a given network topology, it is capable of performing Bayesian updating in the network very efficiently.

Since conditioning and L–S compute exact results, they are directly affected by the NP-hard results; both exhibit exponential behavior in the worst case. Conditioning involves the computation of the minimum cutset which is an NP-hard problem, and L–S is exponential in the size of the largest clique in the graph. In the case of the network shown in Fig. 8, the largest clique consists of only four nodes, but it is quite easy to generate temporal belief networks for which the largest clique makes computing the joint distribution using the L–S algorithm impractical.

In light of the worst-case exponential behavior of the exact algorithms, it is worthwhile to consider how the inference task may be relaxed in order to gain efficiency. Both bounding algorithms and approximation algorithms relax the exactness criterion, trading precise answers for speed. By a bounding algorithm, we mean an algorithm that exploits the structure of a network and the constraints to assign some rough bounds on each node. By an approximation algorithm, we mean a Monte Carlo simulation algorithm. The distinction between the two types of algorithms is somewhat arbitrary, in that it is possible to have an approximation algorithm which in a sense bounds its answer by an error margin. Both types of algorithms are able to supply a rough answer in a very short time, and the reliability of the answers increases with the amount of time spent computing the answer. The main difference of the two types of algorithms is that a bounding algorithm always makes sound, albeit initially very weak, inferences, whereas an approximation algorithm may take some time before the numbers it generates make some sense. Although neither type of algorithm supplies a good answer right away, they can provide some rough answers fairly quickly. Because of this, using these algorithms can have significant advantage in situations where the time spent in computing is important (Dean and Boddy 1988; Horvitz 1988).

Bounding algorithms look at the constraints and bound the distribution, typically by supplying an upper and lower bound (Horvitz *et al.* 1989; Henrion 1988*b*). The bounds are successively refined such that they approach the correct distribution in the limit. These algorithms appear promising, but more research is needed to determine their properties.

Approximation algorithms simulate the states that a belief network is likely to go through, given a set of constraints. By a state of the network, we mean an assignment of a value to each of the variables in the network. For example, given a prior probability constraint that $p(A = 1) = 0.8$ for some variable A, we would, with the use of an oracle, assign $A = 1$ in 80% of the simulated states of the network. The aim in approximation is to develop an algorithm which comes ϵ close to the correct answer almost all the time (i.e., with high probability). As noted earlier, most often, the algorithms perform iterative refinement of the approximations, such that the error is some inversely monotonic function of the time spent in computing the answer. Monte Carlo algorithms have a solid basis in the physical sciences (Metropolis *et al.* 1953). Unfortunately, the Monte Carlo algorithms developed to date for belief networks (Henrion 1988*a*; Pearl 1987) have been somewhat less than ideal; there are no proofs with regard to convergence (we cannot say anything about when the computed distribution is within ϵ of the correct distribution, or how confident we can be that it is within that range), and in practice convergence is often slow (Pearl 1988).

7. Fundamental problems in temporal reasoning

Given that our model addresses many of the same problems that concern logicians working on temporal logic, we will briefly mention how our model deals with certain classic problems in temporal reasoning: the frame, ramification, and qualification problems. We will begin by considering the frame problem stated in probabilistic terms: "Does our model accurately capture our expectations regarding fluents that are considered *not* likely to change as a consequence of a particular event occurring?" The answer is yes insofar as frame axioms can be said to solve the frame problem in temporal logic; persistence constraints are the probabilistic equivalent of frame axioms. In considering the ramification problem, we will consider two possible interpretations. First, "Does our model enable us to compute appropriate expectations regarding the value of a particular fluent at a particular point in time without bothering with a myriad of seemingly unimportant consequences?" The answer to this is a resounding no; our model commits us to predicting every possible consequence of every possible action no matter how implausible. A second interpretation (or perhaps facet is a better word) of the ramification problem is "Does our model enable us to handle additional consequences that follow from a set of causal predictions?" For instance, if A is in box B and I move B to a new location, I should be able to predict that A will be in the new location along with B. Our model provides no provision at all for this sort of reasoning. The basic idea of Bayesian inference can be extended to handle this sort of reasoning, but we have not investigated this to date. The last problem we consider concerns reasoning about exceptions involving the rules governing cause-and-effect relationships. Does our model solve the qualification problem? That is to say, "Does our model accurately capture our expectations regarding the possible exceptions to knowledge about cause-and-effect relationships?" The answer is yes; conditional probabilities would seem to be exactly suited for this sort of reasoning. It should be noted, however, that our model imposes a considerable burden on the person setting up the model. The model described in this paper requires specifying all possible causes for each possible effect and the probability of each effect for every possible combination of possible causes. It is not clear, however, that one can get away with less (Dean and Kanazawa 1988; Hanks 1988). Given the problems inherent in eliciting such information from experts, it would appear that we will have to automate the process of setting up our probabilistic models.

8. Conclusions

This paper presents a model of temporal reasoning suited for situations involving incomplete knowledge. By expressing knowledge of cause-and-effect relations in terms of conditional probabilities, we are able to make appropriate judgements concerning the persistence of propositions. By directly encoding expectations regarding how long certain propositions are likely to persist, we are able to reason about phenomena for which we do not possess an accurate cause-and-effect model as well as phenomena for which we do possess an accurate model but about which it is either difficult or impossible to acquire evidence due to sensing and real-time processing constraints. The primary contribution of this paper is ontological. This paper builds on the temporal representations of Allen (1984), McDermott (1982), and others to provide an account of time and change in terms of well-understood probabilistic models. Our model can be described in terms of a network representation for probabilistic theories that has recently received a great deal of attention in the literature, and for which there now exist a number of computational techniques that appear to be quite promising. We expect that the development of approximate methods for computing the expectations engendered by a given belief network will enable us to apply our methods to solve a number of problems in robotic control and decision support that require reasoning about time and change.

Acknowledgments

The authors would like to thank Peter Cheeseman, Steve Hanks, and Mike Wellman for helpful comments on an earlier draft. This work was supported in part by the National Science Foundation under grant IRI-8612644 and by DARPA/AFOSR under contract F49620-88-C-0132DEF.

ALLEN, J. 1984. Towards a general theory of action and time. Artificial Intelligence, **23**: 123–154.

CHOW, C., and LIU, C. 1968. Approximating discrete probability distributions with dependence trees. IEEE Transactions on Information Theory, **114**: 462–467.

COOPER, G.F. 1987. Probabilistic inference using relief networks is NP-hard. Technical Report KSL-87-27, Knowledge Systems Laboratory, Stanford University, Stanford, CA.

COOPER, G.F., HORVITZ, E.J., and HECKERMAN, D.E. 1988. A method for temporal probabilistic reasoning. Technical Report KSL-88-30, Knowledge Systems Laboratory, Stanford University, Stanford, CA.

DEAN T., and BODDY, M. 1988. An analysis of time-dependent planning. Proceedings of the Seventh National Conference on Artificial Intelligence, St. Paul, MN, pp. 49–54.

DEAN, T., and KANAZAWA, K. 1987. Persistence and probabilistic inference. Technical Report CS-87-23, Department of Computer Science, Brown University, Providence, RI.

——— 1988. Probabilistic temporal reasoning. Proceedings of the Seventh National Conference on Artificial Intelligence, St. Paul, MN, pp. 524-528.

DUDA, R.O., HART, P.E., and NILSSON, N.J. 1981. Subjective Bayesian methods for rule-based inference systems. *In* Readings in artificial intelligence. *Edited by* B.W. Webber and N.J. Nilsson. Tioga Publishing Co., Palo Alto, CA.

HANKS, S. 1988. Representing and computing temporally scoped beliefs. Proceedings of the Seventh National Conference on Artificial Intelligence, St. Paul, MN, pp. 501–505.

HENRION, M. 1988*a*. Propagating uncertainty by logic sampling in Bayes' networks. *In* Uncertainty in artificial intelligence 2. *Edited by* John F. Lemmer and Laveen F. Kanal. North-Holland Publishing Co., Amsterdam, The Netherlands. pp. 149–163.

——— 1988*b*. Towards efficient probabilistic diagnosis in multiply connected belief networks. Proceedings of the Conference on Influence Diagrams, Berkeley, CA.

HORVITZ, E.J. 1988. Reasoning under varying and uncertain resource constraints. Proceedings of the Seventh National Conference on Artificial Intelligence, St. Paul, MN, pp. 111–116.

HORVITZ, E.J., SUERMONDT, H.J., and COOPER, G.F. 1989. Bounded conditioning: flexible inference for decisions under scarce resources. Technical Report KSL-89-42, Knowledge Systems Laboratory, Stanford University, Stanford, CA.

HOWARD, R.A., and MATHESON, J.E. 1984. Influence diagrams. *In* The principles and applications of decision analysis. *Edited by* R.A. Howard and James E. Matheson. Strategic Decisions Group, Menlo Park, CA 94025.

KIM, J.H., and PEARL, J. 1983. A computational model for causal and diagnostic reasoning in inference systems. Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp. 190–193.

LAURITZEN, S.L., and SPIEGELHALTER, D.J. 1988. Local computations with probabilities on graphical structures and their application to expert systems. Journal of the Royal Statistical Society, **50**: 157–194.

MCCARTHY, J. 1986. Applications of circumscription to formalizing commonsense knowledge. Artificial Intelligence, **28**: 89–116.

MCCARTHY, J., and HAYES, P.J. 1969. Some philosophical problems from the standpoint of artificial intelligence. Machine Intelligence, **4**: 463–502.

MCDERMOTT, D.V. 1982. A temporal logic for reasoning about processes and plans. Cognitive Science, **6**: 101–155.

METROPOLIS, N., ROSENBLUTH, A.W., ROSENBLUTH, M.N., TELLER, A.H., and TELLER, E. 1953. Equation of state calculations by fast computing machines. The Journal of Chemical Physics, **21**: 1087–1092.

NUNEZ, L.M. 1989. On the relationship between temporal Bayes networks and Markov chains. M.Sc. thesis, Brown University, Providence, RI.

PEARL, J. 1987. Evidential reasoning using stochastic simulation of causal models. Artificial Intelligence, **32**: 245–257.

——— 1988. Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan-Kaufman, Los Altos, CA.

SYSKI, R. 1979. Random processes. Marcel Dekker, New York, NY.