

**Modeling a Dynamic and Uncertain World I:  
Symbolic and Probabilistic Reasoning about  
Change<sup>1</sup>**

Steve Hanks, Drew McDermott<sup>2</sup>

Department of Computer Science and Engineering  
University of Washington  
Seattle, WA 98195

Technical Report 93-06-10  
June 15, 1993

<sup>1</sup> To appear, *Artificial Intelligence*

<sup>2</sup> Department of Computer Science, Yale University

# Modeling a Dynamic and Uncertain World I: Symbolic and Probabilistic Reasoning about Change

Steve Hanks  
Department of Computer Science & Engr.  
University of Washington, FR-35  
Seattle, WA 98195

Drew McDermott  
Department of Computer Science  
Yale University  
Box 2158 Yale Station  
New Haven, CT 06520

June 28, 1993

To appear: *Artificial Intelligence*

## Abstract

Intelligent agency requires some ability to predict the future. An agent must ask itself what is presently its best course of action given what it now knows about what the world will be like when it intends to act. This paper presents a system that uses a probabilistic model to reason about the effects of an agent's proposed actions on a dynamic and uncertain world, computing the probability that relevant propositions will hold at a specified point in time. The model allows for incomplete information about the world, the occurrence of exogenous (unplanned) events, unreliable sensors, and the possibility of an imperfect causal theory. The system provides an application program with answers to questions of the form "is the probability that  $\varphi$  will hold in the world at time  $t$  greater than  $\tau$ ?" It is unique among algorithms for probabilistic temporal reasoning in that it tries to limit its inference according to the proposition, time, and probability threshold provided by the application. The system will also notify the application if subsequent evidence invalidates its answer to a query.

## Acknowledgements

Thanks to Denise Draper, Oren Etzioni, and Dan Weld, who commented on several drafts. David Madigan improved the presentation as well. The referees made careful and helpful comments on an earlier version. This work was supported in part by DARPA grant DAAA15-87-K-0001 (Yale University) and by NSF grant IRI-9008670 (University of Washington)

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	System Overview . . . . .	4
1.2	Issues . . . . .	5
1.3	Paper overview . . . . .	7
<b>2</b>	<b>Preliminaries: Notation and Temporal Concepts</b>	<b>7</b>
2.1	Temporal ontology . . . . .	7
2.2	Temporal database management . . . . .	8
2.3	Propositions and probabilities . . . . .	8
2.3.1	The form of propositions . . . . .	8
2.3.2	Facts and events . . . . .	9
2.3.3	Contextual propositions . . . . .	9
2.3.4	Reasoning with probabilities . . . . .	10
<b>3</b>	<b>Uncertainty, Evidence, and Probabilistic Updates</b>	<b>10</b>
3.1	The nature of uncertainty . . . . .	10
3.2	Updates on sensory observations . . . . .	11
3.2.1	The observation update equations . . . . .	13
3.2.2	An equivalent graph representation . . . . .	14
3.2.3	Comparison to the likelihood ratio . . . . .	14
3.3	Updates on causal rules . . . . .	15
3.3.1	The structure of causal rules . . . . .	16
3.3.2	Incomplete theories . . . . .	16
3.3.3	The update equation for causal rules . . . . .	17
3.3.4	Causal updates: summary . . . . .	19
3.4	Unpredicted changes . . . . .	19
3.5	The update process for fixed evidence . . . . .	21
3.5.1	Conditioning on contextual propositions . . . . .	22
3.6	An equivalent graph structure . . . . .	22
3.6.1	Projected probabilities . . . . .	24
3.7	Time-based versus state-based models . . . . .	24
<b>4</b>	<b>Identifying important evidence</b>	<b>25</b>
4.1	Bounding the reliability of evidence . . . . .	27
4.2	Bounding the search's temporal scope . . . . .	28
<b>5</b>	<b>Algorithm summary</b>	<b>28</b>
5.1	Algorithm description . . . . .	28
5.2	Maintaining the belief: monitors and signals . . . . .	29
<b>6</b>	<b>Example</b>	<b>29</b>
<b>7</b>	<b>Assessing joint events</b>	<b>32</b>
7.1	The nature of probabilistic dependencies . . . . .	33
7.2	Dependencies in the causal rules . . . . .	34
7.3	Dependencies in the background probabilities . . . . .	34
<b>8</b>	<b>Assumptions</b>	<b>35</b>

<b>9</b>	<b>Related Work</b>	<b>36</b>
9.1	Logical approaches . . . . .	36
9.2	Markov processes . . . . .	36
9.3	Expert-system approaches . . . . .	37
9.4	Rule-based approaches . . . . .	37
9.5	Belief updating by assessing impact . . . . .	38
<b>10</b>	<b>Conclusion</b>	<b>39</b>

# 1 Introduction

Acting intelligently in the world requires the ability to look ahead in time in order to consider the consequences of one’s actions. Planning a morning of errands may force us to decide ahead of time when to go to the grocery store, thinking ahead to the fact that we intend to buy ice cream which will melt if left in the car too long. We mop the floor so as to avoid mopping ourselves into the corner. Cooking a meal requires a great deal of looking ahead, hypothetical scheduling if you will, to make sure things get done on time and that any step in the preparation is finished before its results are needed by some other step. And a doctor, in the midst of prescribing some medication, must think carefully about the future course of the patient’s illness, potential allergic reactions, and about potential interactions among whatever drugs she is suggesting.

Looking ahead requires in turn that the agent maintain some sort of dynamic model of its world—tracking those aspects of the world necessary for it to make good predictions and in addition remembering those aspects of the world necessary to allow it to recover its failures and learn from its experiences.

The model will roughly consist of two parts:

- some theory of how the domain operates: rules governing expected changes, assumptions about the occurrence of events, the truth of propositions, and so on
- evidence about the current state of affairs, generally in the form of reports from sensors of some sort (whether they be mechanical, verbal, or whatever).

This work describes a “world-model manager”—a program module that stores and operates on an agent’s model of its world.<sup>1</sup> Such a program is similar in spirit to temporal database managers, e.g. in [Dean 1985], though we will introduce significant differences both in representation and functionality, the most significant of which is our explicit treatment of uncertainty.

The main contribution of this paper is an algorithm for probabilistic reasoning—estimating the projected probability<sup>2</sup> that a proposition will be true at a point in time. We define a graphical framework for computing this probability, then point out the computational difficulty: that predicting the world may require instantiating the graph at every time point between the time the first piece of evidence was collected to the time of the query.

Such a scheme may be both intractable and unnecessary, however. On the one hand we may want to reason about the world over a long period of time. On the other hand, it may be the case that the world only infrequently changes state in any “interesting” way (i.e. in any way relevant to the probability we are trying to compute). Our algorithm tries to identify the relevant state changes and instantiate the graph only for those propositions and at those times relevant to the prediction task at hand. By examining the propositional content of the query and our causal model of the domain we can ignore certain propositions and evidence as irrelevant, and by reasoning about the qualitative relationship of the probability to the user-supplied threshold (instead of trying to compute the exact probability), we can ignore evidence and abandon lines of reasoning when they become too tenuous to affect that relationship.

## 1.1 System Overview

The diagram in Figure 1 shows how the world-model manager is embedded in an architecture for planning and execution.<sup>3</sup> Note that the planner’s information about the world comes entirely from the world model, and its plan commitments are posted through that module as well.

The world model can be viewed as an interconnected network of beliefs, each of which represents a summary of the agent’s information about the state of some proposition in the world at some point in time.

---

<sup>1</sup>Note the ambiguity of the term “model,” which may refer to the agent’s underlying theory of the domain or to the beliefs generated from the theory and specific evidence, or even to the program that examines the theory and generates the beliefs. Context should make clear the word’s sense.

<sup>2</sup>“Projected” means that the probability that a proposition will be true at time  $t$  is computed on the basis only of the world’s state at time prior to  $t$ . See Section 3.6.1 for more detail.

<sup>3</sup>See [Hanks and Firby 1990] for more details.

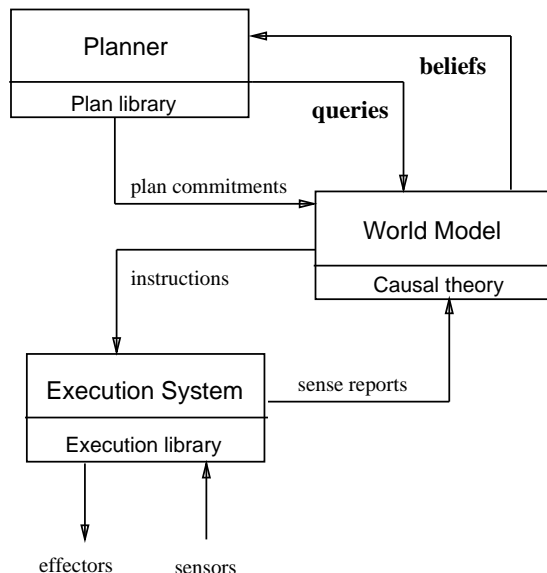


Figure 1: The world model imbedded in an agent architecture

The summary takes the form of a probability, so each belief consists of a probability distribution over the proposition’s state<sup>4</sup> along with the time point at which the proposition’s state was queried.

Beliefs are built (probabilities computed) on the basis of three sorts of information: evidence from the sensors, a causal theory that predicts changes in the proposition’s state, and plan commitments that come from the planner. Computing states of the world based on plan commitments is the problem typically known as “plan projection,” and is the topic of [Hanks and McDermott 1992]. Here we will concentrate on the process of computing probabilities based on sensory evidence and the causal theory.

Central to this work is the interaction between the planner and the world model, in which the planner posts queries to the world model and receives answers back in the form of *belief* data structures. The query indicates to the world-model manager *what* is important, *when* it is important, and *how important* the information is. The model manager tries to perform only the processing necessary to satisfy the planner’s request. The planner then uses the belief’s information to make plan choices (to select new or refine existing plans).

Information about a proposition’s state can change after the belief has been computed and returned to the planner, as the sensors provide more information about the world. A change in the proposition’s predicted state can obviously upset plan choices, so the planner needs to be notified if subsequent information changes a belief significantly. The belief structure returned to the planner carries with it such a promise: the “owner” of the belief will be notified if information subsequently invalidates its estimate; it is up to the owner then to take appropriate action.

## 1.2 Issues

The main issue confronted in this paper is how to build and reason efficiently with a model of a dynamic and uncertain world. The model will have both symbolic and numeric components. Symbolic information, like rules that predict the effects of executing a particular action, is necessary to support processes like plan selection and repair. Numeric information, on the other hand, provides a concise representation of the world’s uncertainty. Our job will be to integrate the two sorts of information.

<sup>4</sup>This paper will consider only binary propositions, so the distribution amounts to a single number—the probability that the proposition is (or was, or will be) true.

Consider, for example, the problem of determining whether one’s car is in the parking lot downstairs. One generally knows where it was put in the morning, but the car may be stolen, or towed if one parks in a lot without a permit.<sup>5</sup> One may have additional information—somebody reporting having seen the car recently, seeing a tow truck in the vicinity, parking in a high-crime area. One needs to reason about changes to the car’s state over time based on evidence, specific knowledge about the domain, and knowledge about how well specific information allows predictions about the case at hand.

The causal theory mentioned in Figure 1 consists of two parts. The first, mentioned above, is a set of “causal rules” that predict change in the world. These rules, similar to those seen in the literature on formal models of change ([McDermott 1982] for example) describe relationships of the form “for all times  $t$ ,  $P$  true at  $t$  causes  $Q$  to become true at  $t+1$ .”

The numeric part of the theory quantifies the extent of our confidence in the rules and also in the reliability of reports received by the sensors.

Parameters of the former type measure the extent to which rules that predict changes in  $P$ ’s state (like the one above) accurately reflect the exact circumstances under which  $P$  will *actually* change state. Since the rules might be neither necessary nor sufficient, we need parameters of two types: the probability that  $P$  will be false even if a rule predicts that it should have become true, and the probability that  $P$  will become true *even if* none of the rules predicted that change.

We use probabilities to represent our confidence in evidence provided us by the execution system (via its sensors) as well: a numeric parameter attached to a symbolic sensor report measures our belief that the sensor reporting on  $P$ ’s state is reporting accurately—that is, that the report is an accurate reflection of the actual state of the world.

Apart from these problems of representation we are faced with a difficult computational problem, which is to cope with the staggering body of evidence that might be brought to bear in answering a query. Our causal rules are quantified over all time points, thus we might consider an infinite number of rule instances in the process of predicting  $P$ ’s state. Sensory evidence about  $P$ ’s state can extend arbitrarily far back into the past. Our task therefore will be to consider enough evidence to produce a good estimate of  $P$ ’s exact probability, ignoring enough evidence so that we can do so quickly. In other words, we have to build an instance of the causal theory (consisting of instances of the rules along with selected observations) appropriate to the query at hand.

The algorithm described in the paper can be described briefly as follows:

- Inputs:
  - A proposition  $\varphi$ , a time point  $t$ , and a probability threshold  $\tau$ .
  - A causal theory, consisting of
    - \* A set of causal rule schemas describing the circumstances under which various propositions (including  $\varphi$ ) change state.
    - \* A set of probability parameters characterizing the extent to which the model is considered complete, the rate at which various events occur, and so on.
  - Situation-specific evidence, consisting of
    - \* A set of observations, each consisting of a proposition (which might be  $\varphi$ ), a time point, and a probability parameter characterizing our confidence in the observation’s veracity.
- Output:
  - A *belief* data structure, consisting of
    - \* An estimate  $E$  of  $P(\varphi_t)$  with the property that if  $E < \tau$  then  $P(\varphi_t) < \tau$ , and if  $E > \tau$  then  $P(\varphi_t) > \tau$ .
    - \* Monitors on the temporal database that will notify the application program if subsequent evidence changes the estimate with respect to the threshold.

---

<sup>5</sup>This example is purely hypothetical, and does constitute an admission of guilt on either author’s part.

- Method:

- The program first computes an estimate of  $\varphi$ 's probability based on current evidence, then tries to characterize (in terms of recency and reliability) what new evidence it would have to be to change the current estimate's value with respect to  $\tau$ . It searches for this evidence, and if doesn't find any, reports the current estimate. If it does find new evidence it re-computes the estimate and repeats the process.

The algorithm is “interruptable” in that it always maintains an estimate of  $\varphi$ 's probability, and can produce that estimate at any time the application needs it (yielding a rough estimate quickly, for example). Although the algorithm will always consider more evidence in computing the probability given more time to compute, it is not convergent, in the sense that  $|E - P(\varphi_t)|$  does not necessarily decrease as the algorithm takes more time.

### 1.3 Paper overview

Describing the model manager amounts to explaining how the module takes a query as input and (1) builds and solves an appropriate instance of the causal theory, yielding an estimate of the query proposition's probability, and (2) sets up the dependency information necessary to detect and respond appropriately to subsequent relevant information. We describe the work in the following steps:

1. a discussion of representation issues—time, chance, and their interplay
2. a method for computing probabilities for a proposition at a time point, *given* an appropriate body of evidence
3. a method for choosing an appropriate body of evidence for the query
4. an example
5. a discussion of the central assumptions underlying the algorithm
6. a discussion of related and future work.

## 2 Preliminaries: Notation and Temporal Concepts

We first need to establish some notational conventions, as well as introduce the temporal model we will exploit in later sections. We will also discuss briefly the temporal database manager module, which is used by the rest of the system. This work is covered more thoroughly in [Hanks 1987] and [Hanks 1990b, Chapter 2].

### 2.1 Temporal ontology

First of all we will adopt a point-based representation of time, meaning that the temporal individuals will denote instants in time. We will use variables like  $t$  and  $t_i$  to refer to these individuals. Our analysis does not require us to commit to whether the underlying model of time is continuous or discrete, though the implementation assumes discrete time for the sake of convenience.

An interval denotes an uninterrupted stretch of time defined by its endpoints. We will write an interval as a pair  $[t_i, t_j]$  where  $t_i$  must occur temporally no later than  $t_j$ . An interval includes both of its endpoints.

The function **distance** or **d** relates the instants denoted by time points.  $d(t_i, t_j)$  indicates the amount of time elapsed between  $t_i$  and  $t_j$ . We therefore enforce the following constraints:

$$\begin{aligned} \forall i, j, k [d(t_i, t_j) + d(t_j, t_k) &= d(t_i, t_k)] \\ \forall i, j [d(t_i, t_j) &= -d(t_j, t_i)] \\ \forall i [d(t_i, t_i) &= 0], \end{aligned}$$



and introduce the notation

$$t_i \preceq t_j \equiv d(t_i, t_j) \geq 0.$$

## 2.2 Temporal database management

The temporal database manager is a module that maintains a network of time points and intervals along with constraints on their temporal distances. It provides four functions:

1. managing the addition and deletion of points and intervals,
2. maintaining the set of constraints on inter-point distances,
3. fetching information from the database,
4. monitoring the database for the addition of new information.

Application programs can add constraints of the form

$$\mathbf{low} \leq d(t_i, t_j) \leq \mathbf{high}$$

where **low** and **high** are integers with  $\mathbf{low} \leq \mathbf{high}$ , and then pose queries of the form “is  $d(t_i, t_j) \leq k$ ?” for time points  $t_i$  and  $t_j$ , and some integer  $k$ .

The other type of query is the “temporally scoped fetch,” which returns all time points (or intervals) that fall within a application-supplied time interval, and whose *tokens* satisfy an application-supplied fetch criteria (supplied in the form of a function). Each time point and interval has an associated *token*, indicating, for example, what happened at the time point or over the interval. The database attaches no particular meaning to the tokens, using them only in the process of fetching.

The application can ask the database monitor to notify it (by calling an application-supplied function) when a point or interval gets added that (1) occurs within a particular interval of time, and (2) whose token satisfies a particular criterion. An add monitor is attached to the interval of interest and checks points and intervals as they are added to the database, notifying the application if an “interesting” one arrives. The idea is that new information might invalidate a probabilistic estimate— $\varphi$ ’s probability was thought to be greater than  $\tau$  but new evidence shows it to be less—at which point the application should reconsider any plans that rely on  $\varphi$  being true.

## 2.3 Propositions and probabilities

This paper is concerned with how to compute probabilities of the form

$$P(\text{“}\varphi \text{ true at time } t\text{”}),$$

where  $\varphi$  is a proposition and  $t$  is a time point. We have already introduced the notion of a time point; now we move on to the notion of a proposition, then to a brief discussion of why we are using probabilities to represent the agent’s uncertainty.

### 2.3.1 The form of propositions

We will use Greek letters, most often  $\varphi$ , to stand for the propositions we wish to assess. “Assessing  $\varphi$ ” means computing  $\varphi$ ’s probability at some time point  $t$ . We will abbreviate this quantity as  $P(\varphi_t)$ . In cases where  $t$  is an indexed time point, e.g.  $t_i$ , we may further abbreviate  $P(\varphi_{t_i})$  as  $P(\varphi_i)$ .

These propositions will in practice look like the usual AI logical propositions: (RAINY), (LOC MY-CAT MY-LIVING-ROOM), and so on—the *atomic formulas* in a first-order language. Propositions thus do not contain logical connectives. Instead of logical negation and conjunction, for example, probability theory speaks of *complementary* and *joint* events. If  $\varphi$  and  $\psi$  are propositions of interest (thus

are events in the probability space) the events  $\bar{\varphi}$  and  $\bar{\psi}$  are the corresponding complementary events, corresponding to (NOT  $\varphi$ ) and (NOT  $\psi$ ). The axioms of probability theory dictate that  $P(\varphi) = 1 - P(\bar{\varphi})$ . The joint event  $(\varphi, \psi)$  corresponds to the conjunction  $(\varphi \wedge \psi)$ . We will discuss below the problem of computing the probabilities of joint events. We allow quantification only in special domain-dependent cases, which we will discuss in [Hanks and McDermott 1992]. When we speak of propositions such as  $\varphi$ , however, we are talking about atomic propositions—no connectives.

Our causal model, inspired by McDermott’s concept of the “pcause” rule [McDermott 1982], grants special semantic status to a class of propositions known as “events.” Events are the propositions that effect changes in the world. The agent’s actions are events. The distinction between events and “facts” (as the remaining propositions are called) is somewhat arbitrary, as Shoham points out in [1988, Chapter 2]. Is there a meaningful difference between the occurrence of a rainstorm and some period of time over which it is raining?

### 2.3.2 Facts and events

Theories of change have typically divided the set of proposition types into *facts* and *events*. The latter are usually associated with the agent’s actions. In fact most theories of change, and most planners, assume that no events *except* the agent’s actions will ever occur.

One reason for making this distinction is to separate the intuition that if a fact F is said to be true over the interval we want to say it’s true over all subintervals as well. In contrast, if an event E is said to be true over an interval we really want to say it occurs exactly once, over the whole interval, and not over any subinterval. This distinction avoids a technical problem with the causal rules: if we have a rule that says if F is true when E occurs then P becomes true, and we assert that E occurred over an interval, then we want to assert that P becomes true exactly once, and not at every time point contained within the interval.

In practice this problem rarely arises, mainly because events are almost always assumed to occur at a single point in time. We will begin by making this assumption, then relax it so that certain events (in particular the agent’s actions) are allowed to occur over an interval rather than at a point in time.

A more important distinction arises in the probabilistic assumptions we make about the occurrence of events as opposed to the persistence of facts.

The first assumption is that events occur “dispassionately” in the world. More particularly we assume that the probability that an exogenous event<sup>6</sup> occurs does not depend on the planning choices made by the agent. We are thus assuming a situation of uncertainty rather than one of risk or gaming.

A second assumption we make about events is a rather vague one; it is important for the efficiency of the assessment and projection algorithms but does not otherwise affect the analysis. The assumption is that causally relevant events occur, *unknown to the agent*, with low probability. We characterize as “chaotic” those domains in which the low-probability assumption does not hold. Predicting the future becomes very difficult in chaotic situations; our algorithm becomes very slow. We will see below how this assumption of “low probability events” can be used to guide the search for important causal inferences.

Note that these assumptions—and thus the distinction between facts and events—is essentially a practical one. In fact the distinction between events and facts will usually be unimportant to our analysis; when it matters we will use the letter  $\varepsilon$  or E to stand for the former and the letter  $\varphi$  or P to stand for the latter.

### 2.3.3 Contextual propositions

Contextual propositions are facts we have some information about but do not reason about directly using our temporal theory; they are assumed not to change in the interval over which we model the temporal facts and events. The weather, the current temperature, a patient’s age are examples.

Our implementation assumes an atemporal probabilistic model of the contextual propositions, meaning that we can ask for the probability of these propositions, and these probabilities may be interdependent, but the probabilities remain unchanged over the interval in question.

---

<sup>6</sup>An exogenous event is one that is not explicitly planned by the agent.

Contextual propositions will be relevant only in Sections 3.4 and 7.

### 2.3.4 Reasoning with probabilities

We will quantify the agent’s uncertainty about the world using a probability number, thus will be speaking of assessing the probability of  $\varphi$  at a time point  $t$ , abbreviated to  $P(\varphi_t)$ . This approach is controversial. Some researchers of a philosophical bent contend that a probability is not a valid measure of credal commitment. Practitioners in AI argue a number of points: that probabilistic reasoning requires a tremendous amount of input data (prior probabilities and the like) that cannot reasonably be expected from an application program, that probabilistic reasoning is computationally intractable, and that a single probability number is insufficient to capture concepts, like ignorance and the value of information, that are essential components of the decision-making or planning process.

The philosophical issues begin a debate we have no wish to enter at this point; there is an excellent summary of the issues in [Pearl 1988, Chapter 1]. In any case, the analysis in this chapter does not require the reader to view the computed probabilities as states of belief<sup>7</sup>

The second group of arguments touches on empirical issues, and we hope to show in this paper that reasoning within an application-supplied probabilistic model can be accomplished efficiently. Our contention will be that although computing *exact* probability values may be impractical, one may nonetheless compute approximations that (1) can be arrived at quickly, and (2) provide the agent with the information necessary to make decisions about its future actions.<sup>8</sup> The final point, that a single probability number does not suffice to represent concepts related to the agent’s state of information, will be addressed in [Hanks and McDermott 1992].

## 3 Uncertainty, Evidence, and Probabilistic Updates

Estimating a proposition’s probability can be considered in two steps:

1. choosing an appropriate body of evidence  $E$ , and
2. computing the probability  $P(\varphi_t|E)$ .

Our implementation actually interleaves these two steps, but viewing them separately serves to isolate the issues involved in the two processes. This section treats the second step, computing a probability given a particular set of evidence. We explain the process by confronting these questions:

1. How does uncertainty manifest itself in the system?
2. What form does evidence take?
3. What is the proper update for each type of evidence?

### 3.1 The nature of uncertainty

There are two ways to arrive at a belief that some proposition, say  $\varphi$ , holds in the world at some point in time  $t$ . The first is to observe  $\varphi$  at  $t$  either directly, or by being told by someone else, or by some other indirect evidence. We then base a belief in  $\varphi$ ’s truth on

1. the observation or evidence itself, and
2. the belief that the observation at time  $t$  accurately reflected  $\varphi$ ’s state at that time.

---

<sup>7</sup> We may occasionally lapse into language betraying the fact that we are interpreting these probabilities in exactly that way—e.g. by naming the main data structure a *belief*. Nonetheless, what we are doing is computing probabilities, and the reader is free to interpret them any way he or she wants.

<sup>8</sup> [Wellman 1988] advances this position as well.

We will assume that we will get observational evidence about  $\varphi$  only in the form of direct (but possibly inaccurate) testimony.

We can also reason indirectly about  $\varphi$ 's state, provided we have some sense of the circumstances under which  $\varphi$  can become true. For example, knowing that the occurrence of an event  $\varepsilon$  causes  $\varphi$  to become true, and knowing that  $\varepsilon$  actually occurred at a time  $t'$  prior to  $t$ , we can base a belief in  $\varphi$ 's truth on

1. the belief that  $\varepsilon$  causes  $\varphi$ , and
2. the belief that  $\varepsilon$  occurred at  $t'$ , and
3. the belief that nothing happened between  $t'$  and  $t$  to cause  $\varphi$  to become false.

The two modes of reasoning are not exclusive—we may have arrived at the belief that  $\varepsilon$  occurred at  $t'$  by direct observation, or by some other chain of causal inference, or both.

Belief in the truth of  $\varphi$  at  $t$  will be uncertain to the extent that we doubt:

- whether an observation was a good reflection of  $\varphi$ 's state at  $t$
- whether some relevant  $\varepsilon$  actually occurred at  $t'$
- whether  $\varepsilon$  indeed causes  $\varphi$
- whether nothing actually changed  $\varphi$ 's state in the interim, because
  - maybe some other known relevant event occurred but its occurrence was unknown
  - maybe some other relevant event occurred but its *relevance* was unknown.

To summarize the above, one is uncertain about the state of the world to the extent that one lacks confidence in the veracity of one's observations, lacks perfect knowledge of all causally relevant events, and lacks a perfect predictive model of the world.

The discussion identifies two types of evidence: reports from the sensors and predictive information from the causal rules. Each has an associated update equation—a function that incorporates the new evidence into one's belief about  $\varphi$ 's state at  $t$ . In other words, if one's current state of information about  $\varphi$ , say  $E$ , leads one to a certain probability  $P(\varphi_t|E)$ , and a new piece of evidence  $e$  is obtained, then the update equation determines the new probability  $P(\varphi_t|E, e)$  on the basis of the old probability and an estimate of  $e$ 's veracity.

### 3.2 Updates on sensory observations

Suppose we are reasoning about the state of  $\varphi$  at a point in time  $t$ . We are thus computing the probability  $P(\varphi_t)$ , and we have so far amassed a body of evidence  $E$  that bears on  $\varphi$ , and have used it to compute a probability  $P(\varphi_t|E)$ . Then we get a sensory observation  $O$  to the effect that  $\varphi$  was observed true at time  $t$ .<sup>9</sup>

The complementary event is  $\bar{O}$ , which denotes an observation that  $\varphi$  is false. The question is how to update the probability  $P(\varphi_t|E)$  to incorporate the new piece of evidence; that is, to compute  $P(\varphi_t|E, O)$ , bearing in mind that  $O$  may or may not accurately reflect  $\varphi$ 's state at  $t$ .

Associated with each observation  $O$  we have several parameters: the proposition it claims to have observed, whether that proposition was true or false, the time point at which the observation occurred, and

---

<sup>9</sup>We should note a potential confusion here among time points. Two time points are potentially relevant in considering an observation: the time at which the observation was *made* and the time at which the observation was posted to the database. Likewise when one makes a query like “what is  $\varphi$ 's probability at  $t$ ” we can consider both the time  $t$  and whatever time the application *posed* the query. We are not considering times of the second sort: the times at which observations are made and the times at which queries are posed. Our network of beliefs represents what the application believes *now*, on the basis of what information it has received to this point. We cannot therefore reason about changes to the agent's state of information over time. Although we can represent statements like “I believe that  $\varphi$  was true five minutes ago but will not be true in one hour,” we cannot explicitly represent statements like “five minutes ago I believed that  $\varphi$  would be true in one hour, but now I do not.”

one or more parameters characterizing the relationship between  $\varphi$ 's state and the sensor's report about  $\varphi$ 's state.

The last parameters are most interesting. The sensor either was functioning properly when it made the observation, in which case we say the observation was “reliable,” or it was malfunctioning, in which case we say the observation was “unreliable.” Let  $\mathbf{R}$  denote a reliable observation. The probability space then consists of the following events:

- $\varphi$  = proposition  $\varphi$  is true in the world
- $\bar{\varphi}$  = proposition  $\varphi$  is false
- $\mathbf{O}$  = sensor reports  $\varphi$  is true
- $\bar{\mathbf{O}}$  = sensor reports  $\varphi$  is false
- $\mathbf{R}$  = observation  $\mathbf{O}$  is reliable
- $\bar{\mathbf{R}}$  = observation  $\mathbf{O}$  is unreliable

all of which are relative to time  $t$ .

We can constrain the joint distribution over these events so that it corresponds to our definitions of observations and reliability. We first assume that a *reliable* observation of  $\varphi$  is definitive—it should determine our belief in  $\varphi$  at the time the observation was made. A reliable observation of  $\bar{\varphi}$  (i.e.  $\bar{\mathbf{O}}$ ) means that we believe absolutely that  $\varphi$  is false at  $t$ . This assumption translates into the following two constraints:

$$P(\varphi|\mathbf{O},\mathbf{R},\mathbf{E}) = 1 \tag{1}$$

$$P(\varphi|\bar{\mathbf{O}},\mathbf{R},\mathbf{E}) = 0 \tag{2}$$

for any body of evidence  $\mathbf{E}$ .

It is harder to characterize the relationship between an observation and  $\varphi$  if the observation was *unreliable*: a malfunctioning sensor may still report either  $\mathbf{O}$ , or  $\bar{\mathbf{O}}$ , and hence an observation may correctly report  $\varphi$ 's state even though it was unreliable—*i.e.* the sensor may do the right thing for the wrong reason.

Furthermore, we generally won't know whether any individual observation was reliable or not, and may therefore have to base our update on statistical information about the sensor's behavior. First let us assume we have a probability  $r = P(\mathbf{R})$  that represents the likelihood that the observation was reliable. We next assume that that there is no systematic relationship between  $\varphi$ 's state and the observation's reliability. That is,  $\varphi$  being true in no way *causes* the sensor to malfunction. This assumption amounts to an assertion that  $\mathbf{R}$  is probabilistically independent of  $\varphi$ :

$$P(\mathbf{R}|\varphi) = P(\mathbf{R}|\bar{\varphi}) = P(\mathbf{R}) = r. \tag{3}$$

Next we have to characterize the relationship between the observation and  $\varphi$  for those cases when the sensor malfunctions. This information corresponds to the parameters  $P(\mathbf{O}|\bar{\mathbf{R}},\varphi)$  and  $P(\mathbf{O}|\bar{\mathbf{R}},\bar{\varphi})$ . In some cases we may have a good model of the sensor's “failure-mode” behavior—for example, the sensor may always report “true” when it malfunctions, in which case we would set

$$P(\mathbf{O}|\bar{\mathbf{R}},\varphi) = P(\mathbf{O}|\bar{\mathbf{R}},\bar{\varphi}) = 1.0.$$

This example introduces a reasonable assumption about unreliable observations: that when the sensor malfunctions its output is no longer sensitive to  $\varphi$ 's state. This assumption amounts to the assertion that  $P(\mathbf{O}|\bar{\mathbf{R}},\varphi) = P(\mathbf{O}|\bar{\mathbf{R}},\bar{\varphi})$ . If we had less precise information about the sensor's behavior when it malfunctions, we might introduce into the model a “failure-mode” parameter  $f$ : the probability that the sensor reports “true” given that it is malfunctioning:

$$P(\mathbf{O}|\bar{\mathbf{R}},\varphi) = P(\mathbf{O}|\bar{\mathbf{R}},\bar{\varphi}) = f. \tag{4}$$

In other cases we may have *no* explicit model of the sensor’s failure-mode behavior, and thus have no values for the probabilities  $P(O|\bar{R},\varphi)$  and  $P(O|\bar{R},\bar{\varphi})$ . The assumption we made above—that the sensor’s failure-mode behavior does not depend on  $\varphi$ ’s state—is still reasonable, but we may have no  $f$  parameter available. Two other reasonable assumptions might be:

1. That one’s state of belief about  $\varphi$  should not change as a result of an unreliable sensor report. Formally we might say that  $P(\varphi | O, \bar{R}) = P(\varphi | \bar{O}, \bar{R}) = P(\varphi)$ , though it is not immediately clear that this assumption determines a unique distribution or even that it is consistent with the rest of the constraints.
2. That our assumption about the failure-mode probabilities should add the least possible amount of information to the distribution. This is in effect the “maximum-entropy” solution ([Jaynes 1979] [Hunter 1989]): choose the (unique) distribution consistent with the constraints that exhibits maximum entropy (minimum information).

We show in [Hanks 1990b, Section 3.4] that the three assumptions are equivalent: that the maximum-entropy distribution implies

$$P(O | \bar{R}, \varphi) = P(O | \bar{R}, \bar{\varphi}) = 0.5,$$

which in turn implies

$$P(\varphi | O, \bar{R}) = P(\varphi | \bar{O}, \bar{R}) = P(\varphi).$$

In other words, the report of a malfunctioning sensor does not depend on  $\varphi$ ’s state, and an unreliable observation conveys no information.

This model—the reliability probability parameter  $r$  coupled with a minimum-information assumption about the sensor’s failure-mode behavior—is essentially the same as that adopted for electronic components in the literature on circuit diagnosis, e.g. [de Kleer and Williams 1987]. For the remainder of the paper we will assume a model consisting of the reliability parameter  $r$  only. Adopting a model with more parameters, e.g. the  $f$  parameter from Equation (4), would require minimal changes to the algorithm.

### 3.2.1 The observation update equations

Now suppose we have available a prior probability  $p = P(\varphi | E)$  and a reliability parameter  $r = P(R)$ . We can then build the following set of constraints:

$$\begin{aligned} P(\varphi|O,R,E) &= 1 \\ P(\varphi|\bar{O},R,E) &= 0 \\ P(\varphi|E) &= p \\ P(R|E) &= P(R|\varphi) = P(R|\bar{\varphi}) = r \\ P(O|\varphi,\bar{R},E) &= P(O|\bar{\varphi},\bar{R},E) = 0.5 \end{aligned}$$

which defines a unique distribution over the probability space. We can then apply Bayes’ rule:

$$P(\varphi|E, O) = \frac{P(O | E, \varphi)P(\varphi | E)}{P(O | E)}$$

which, after expanding terms and substitution, yields:

$$P(\varphi|E,O) = \frac{.5(1+r)p}{rp + .5(1-r)} \tag{5}$$

for positive observations, and similarly for negative observations

$$P(\varphi|E, \bar{O}) = \frac{(1-r)p}{1+r-2rp}. \tag{6}$$




Figure 2: Graph representation for observation updates

Note the following properties of these equations: if the observation is certainly reliable ( $r = 1$ ) then  $P(\varphi | \mathbf{O}) = 1$  and  $P(\varphi | \overline{\mathbf{O}}) = 0$ . A reliable observation renders the agent certain about the state of the proposition in question. If the observation is certainly unreliable ( $r = 0$ ) then  $P(\varphi | \mathbf{O}) = P(\varphi | \overline{\mathbf{O}}) = p = P(\varphi)$ . An unreliable observation provides no information regarding the proposition. The positive update equation is undefined in case  $r = 1$  and  $p = 0$ ; a negative update is undefined in case  $r = 1$  and  $p = 1$ .

Now assume we have an observation  $\mathbf{o}_i$  that takes place at time  $t_i$  and has reliability parameter  $r_i$ . We can rewrite (5) and (6) as a function of the prior probability  $p$ , but which depends on (1) whether  $\mathbf{o}_i$  is positive or negative and (2) the reliability parameter  $r_i$ . We will call this function  $s_i$ , for a sensor update at time  $t_i$  and reliability parameter  $r_i$ , and note that it is monotone nondecreasing in its input parameter.<sup>10</sup>

$$s_i(p) = \begin{cases} \frac{5(1+r_i)p}{r_i p + 5(1-r_i)} & \text{if } \mathbf{o}_i \text{ is positive} \\ \frac{(1-r_i)p}{1+r-2r_i p} & \text{if } \mathbf{o}_i \text{ is negative.} \end{cases} \quad (7)$$

### 3.2.2 An equivalent graph representation

Figure 2 shows explicitly the probabilistic model implied by the update equations. Note that the state of  $\varphi$  and the observation's reliability  $R$  jointly determine whether the observation will be true or not.  $\varphi$  and  $R$  are independent, given a true or false observation. We require prior probabilities for  $\varphi$  and for  $R$  and a matrix (shown in the diagram) for all possible values of  $P(\mathbf{O}|\varphi, R)$ . The entries in this matrix are determined by the constraints and the maximum-entropy assumption introduced in the previous section.

### 3.2.3 Comparison to the likelihood ratio

The approach above does sensor updates on the basis of (1) a reliability parameter  $r$ , (2) a prior probability  $p$ , and (3) additional parameters or assumptions about the sensor's behavior when it malfunctions. A more traditional update technique (see [Pearl 1988, Chapter 2], for example) uses two likelihood ratios instead,

$$\lambda_{\mathbf{O}} = \frac{P(\mathbf{O} | \varphi)}{P(\mathbf{O} | \overline{\varphi})}$$

$$\lambda_{\overline{\mathbf{O}}} = \frac{P(\overline{\mathbf{O}} | \varphi)}{P(\overline{\mathbf{O}} | \overline{\varphi})}$$

<sup>10</sup>Note that we no longer need to condition on  $E$ . Prior information is summarized in the parameter  $p$ . The sensor update depends only on the reliability parameter, which is assumed to be independent of  $E$ .

along with the same prior probability, then applies Bayes' rule as follows (this is for positive updates):

$$P(\varphi|\mathbf{O}) = \frac{\lambda_{\mathbf{O}}\Theta}{1 + \lambda_{\mathbf{O}}\Theta} \quad (8)$$

where

$$\Theta = \frac{P(\varphi)}{1 - P(\varphi)} = \frac{p}{1 - p}.$$

As it happens, the two approaches are almost identical in their applicability. Suppose that we admit an arbitrary relationship between  $\mathbf{O}$  and  $\varphi$ —the following equations then hold:

$$\begin{aligned} P(\mathbf{O}|\varphi) &= P(\mathbf{O}|\varphi, \mathbf{R})P(\mathbf{R}|\varphi) + P(\mathbf{O}|\varphi, \overline{\mathbf{R}})P(\overline{\mathbf{R}}|\varphi) \\ P(\mathbf{O}|\overline{\varphi}) &= P(\mathbf{O}|\overline{\varphi}, \mathbf{R})P(\mathbf{R}|\overline{\varphi}) + P(\mathbf{O}|\overline{\varphi}, \overline{\mathbf{R}})P(\overline{\mathbf{R}}|\overline{\varphi}) \end{aligned}$$

But with our assumption that reliability is independent of  $\varphi$ , we can take all of the  $P(\mathbf{R}|\cdot)$  parameters to be equal to  $r$ , and then  $P(\overline{\mathbf{R}}|\cdot) = 1 - r$  as well. Furthermore,  $P(\mathbf{O}|\overline{\mathbf{R}}, \varphi) = P(\mathbf{O}|\overline{\mathbf{R}}, \overline{\varphi}) = f$ , the failure-mode parameter we introduced in Equation (4). And, of course,  $P(\mathbf{O}|\mathbf{R}, \overline{\varphi}) = 0$ . Hence, the equations above reduce to

$$\begin{aligned} P(\mathbf{O}|\varphi) &= r + (1 - r)f \\ P(\mathbf{O}|\overline{\varphi}) &= (1 - r)f. \end{aligned}$$

Subtracting, we obtain

$$r = P(\mathbf{O}|\varphi) - P(\mathbf{O}|\overline{\varphi}),$$

which establishes a relationship between  $r$  and the components of the likelihood ratio  $\lambda_{\mathbf{O}}$ . Considering  $P(\overline{\mathbf{O}})$  yields the corresponding equation

$$r = P(\overline{\mathbf{O}}|\overline{\varphi}) - P(\overline{\mathbf{O}}|\varphi).$$

We can perform a similar analysis for the  $f$  parameter:

$$\begin{aligned} f &= \frac{P(\mathbf{O}|\overline{\varphi})}{(1 - r)} \\ &= 1 - \frac{P(\overline{\mathbf{O}}|\varphi)}{(1 - r)} \end{aligned}$$

where  $f$  is undefined (but unnecessary) if  $r = 1$ .

These equations suggest certain relationships between the two update methods. They basically cover the same territory, but just parameterize it differently. Any relationship between the probabilities of  $\varphi$  and  $\mathbf{O}$  can be modelled in terms of  $r$  and  $f$  parameters, unless  $P(\mathbf{O}|\varphi) < P(\mathbf{O}|\overline{\varphi})$ , in which case it is misguided to think of  $\mathbf{O}$  as an observation of  $\varphi$  at all.

One advantage to formulating the update in terms of  $r$  and  $f$  is that it allows the case  $r = 1$ , which corresponds to a singularity in the orthodox Bayesian approach. A perfectly reliable observation fixes belief with certainty, which is impossible using Equation 8 above—the posterior probability for  $\varphi$  cannot be 1. A probabilistic purist will no doubt approve of the impossibility of forcing absolute certainty with a single observation. But in the context of computational modeling it is quite useful to be able to include perfectly reliable observations as boundary conditions, thereby cutting off expensive cogitation about slight deviations from 1.

### 3.3 Updates on causal rules

Equation (7) allows us to update a probability at a single point in time in response to a new piece of observational evidence. In this section we discuss the problem of how to update the probability between a point of time  $t$  and the next instant  $t^>$ . (If we were to assume a discrete temporal model we would set  $t^> = t + 1$ .)



### 3.3.1 The structure of causal rules

Many proposals have appeared for reasoning about the dynamics of the world, variously called theories of time [McDermott 1982], theories of action [Lifschitz 1987] or theories of change [Shoham 1988]. The idea in all cases is the same: we identify *events* as the propositions that initiate change in the world; events change the state of *facts*. The agent’s world model consists of a set of *causal rules* that dictate what event types cause what changes to what fact types under what circumstances. One such rule (a variant of which appears in [Hanks and McDermott 1987]) represents the information that shooting somebody with a loaded gun usually causes that person to become dead (actually, “not alive”):<sup>11</sup>

```
(DEFINE-RULE-SCHEMA 'GUN-SHOOTING
 '(PRECONDITIONS (LOADED ?G))
 '(TRIGGER (SHOOT ?P ?G))
 '(CONSEQUENT (ALIVE ?P))
 '(SIGN 0)
 '(EFFECTIVENESS 0.9))
```

A **SIGN** of 0 means the same thing as negating the consequent (recall that propositions cannot be negated). We will explain the **EFFECTIVENESS** parameter below.

These rules must be accompanied by axioms or algorithms that enforce the idea that the consequents indicated by the rules are realized at all appropriate time points—that is, for any point in time at which the **TRIGGER** event occurs and the **PRECONDITION** facts are all true, the **CONSEQUENT** fact is caused to be true (or false, in this case) immediately afterwards.

A first-cut interpretation of these rules is as follows: let **R** be a rule with precondition **P**, trigger event **T**, sign **S**, and consequent **C**. In the case  $S = 1$  (a “positive rule”) we can state, for all time points **t** that

$$P(C_{t^>} | P_t, T_t, \bar{C}_t, E) = 1 \tag{9}$$

and similarly for  $S = 0$ :

$$P(C_{t^>} | P_t, T_t, C_t, E) = 0 \tag{10}$$

where **E** again is our previous body of evidence.

The equation says, in effect that the probability that the (positive) rule will effect a change in **C**’s state from false to true between time **t** and time **t**<sup>></sup> is 1. The equation is incorrect in that it assumes that the rule is a perfect predictor of state changes in **C**—that is, the rule’s preconditions being true are sufficient to predict a change in **C**’s state. The next section addresses this issue.

We can characterize a rule fully by its precondition, trigger, consequent, and sign, so let  $R = (P^R, T^R, C^R, S^R)$ . We will be immediately interested only in rules for which  $C^R = \varphi$ , for these are the only known conditions under which  $\varphi$ ’s state can change.

Let us say that a rule “fires” at a point in time **t** if its precondition facts are all true at **t**, its trigger event occurs at **t**, and  $\varphi$  is false at **t** for positive rules or true for negative rules. For positive rules **R** we will define

$$F_t^R \equiv (P_t^R, T_t^R, \bar{\varphi}_t)$$

and likewise for negative rules

$$F_t^R \equiv (P_t^R, T_t^R, \varphi_t)$$

### 3.3.2 Incomplete theories

Equations 9 and 10 imply that we know exactly the circumstances under which a rule’s effects will be realized—in particular we assume that the rule’s firing is sufficient to change  $\varphi$ ’s state. This may not always

---

<sup>11</sup> This is the actual LISP form evaluated by the implementation

be the case: referring to the gun-firing rule, for example, it may be the case that shooting somebody with a loaded gun will usually, but not always, cause their death. We can express this notion—that the rule’s stated preconditions are necessary but not sufficient for the rule to effect change—by associating with each rule a parameter called “effectiveness.” Effectiveness represents the proportion of situations in which the rule’s consequent will be realized, given that it fires. Giving a rule an effectiveness parameter of  $p$  is equivalent to defining an additional precondition of “(CHANCE  $p$ )” which always has a probability of  $p$  and is known to be independent of all other propositions. The independence assumption is justified here because we are quantifying the extent to which our model *fails* to capture certain relationships; presumably any systematic relationships would be captured in the model. The advantage of using an explicit **EFFECTIVENESS** parameter instead of a (CHANCE  $p$ ) precondition is practical rather than semantically meaningful, so the present discussion will ignore the rule’s effectiveness (assuming it is 1.0). (But see below, Section 4.1.)

Just as effectiveness represents an admission that the rule’s precondition might not be restrictive enough (but we don’t know exactly how to restrict it further), we might also want to represent the possibility that the precondition might be *too* restrictive (but we don’t know exactly how to weaken it). We have no explicit parameter to represent this case—such information is coded in the application-supplied parameter  $\Delta^\varphi$ .  $\Delta^\varphi$  is the probability that  $\varphi$  changed state, but the cause of that change is unknown. Section 3.4 discusses this parameter in more detail. The  $\Delta^\varphi$  parameter does not distinguish between the case in which the unknown cause of  $\varphi$ ’s change is a known rule that should have fired and the case in which the change’s cause was not represented at all.

### 3.3.3 The update equation for causal rules

Now we return to the problem of tracking how belief changes from time  $t$  to time  $t^>$  where our theory of change consists of a set of rules like the one above. Suppose that  $R^1, R^2, \dots, R^n$  are all the rules whose consequent is  $\varphi$ , and the corresponding “firing” events are  $F^1, F^2, \dots, F^n$ . Assume for the sake of notational brevity that  $n = 2$ . The following causal update equation indicates that  $\varphi$ ’s state at  $t^>$  can be determined by considering its state at  $t$  and considering what happens under the (mutually exclusive and exhaustive) circumstances that (1) neither rule fires, (2) one rule fires but the other does not, and (3) both rules fire.

$$P(\varphi_{t^>} | E) = P(\varphi_{t^>} | F_t^1, F_t^2, \varphi_t) P(F_t^1, F_t^2, \varphi_t | E) + \tag{11}$$

$$P(\varphi_{t^>} | \overline{F}_t^1, F_t^2, \varphi_t) P(\overline{F}_t^1, F_t^2, \varphi_t | E) + \tag{12}$$

$$P(\varphi_{t^>} | F_t^1, \overline{F}_t^2, \varphi_t) P(F_t^1, \overline{F}_t^2, \varphi_t | E) + \tag{13}$$

$$P(\varphi_{t^>} | \overline{F}_t^1, \overline{F}_t^2, \varphi_t) P(\overline{F}_t^1, \overline{F}_t^2, \varphi_t | E) + \tag{14}$$

$$P(\varphi_{t^>} | F_t^1, F_t^2, \overline{\varphi}_t) P(F_t^1, F_t^2, \overline{\varphi}_t | E) + \tag{15}$$

$$P(\varphi_{t^>} | \overline{F}_t^1, F_t^2, \overline{\varphi}_t) P(\overline{F}_t^1, F_t^2, \overline{\varphi}_t | E) + \tag{16}$$

$$P(\varphi_{t^>} | F_t^1, \overline{F}_t^2, \overline{\varphi}_t) P(F_t^1, \overline{F}_t^2, \overline{\varphi}_t | E) + \tag{17}$$

$$P(\varphi_{t^>} | \overline{F}_t^1, \overline{F}_t^2, \overline{\varphi}_t) P(\overline{F}_t^1, \overline{F}_t^2, \overline{\varphi}_t | E). \tag{18}$$

First consider the left-hand side of each addend— $P(\varphi_{t^>} | F_t^1, F_t^2, \varphi_t)$ ,  $P(\varphi_{t^>} | \overline{F}_t^1, F_t^2, \varphi_t)$ , and so on. These quantities divide into three groups: those in which none of the causal rules fire: (14), (18), those in which exactly one of the rules fire: (12), (13), (16), (17), and those in which more than one rule fires: (11), (15).

The first group represents the case in which nothing was *known* to cause  $\varphi$  to change state between  $t$  and  $t^>$ . While we admit the possibility of *unpredicted* state changes over intervals (Section 3.4), we assume that the probability that such a change happens in the instant  $[t, t^>]$  is negligible. We can therefore rewrite the left-hand-sides of terms (14) and (18) as follows:

$$P(\varphi_{t^>} | \overline{F}_t^1, \overline{F}_t^2, \varphi_t) \approx 1$$

$$P(\varphi_{t^>} | \overline{F}_t^1, F_t^2, \overline{\varphi}_t) \approx 0.$$

The second group of terms—(12), (13), (16), (17)—represents the case where exactly one rule fires. Consider (12) and (16), for example:

$$\begin{aligned} & P(\varphi_{t>} | \overline{F}_t^1, F_t^2, \varphi_t) P(\overline{F}_t^1, F_t^2, \varphi_t | E) \\ & P(\varphi_{t>} | \overline{F}_t^1, F_t^2, \overline{\varphi}_t) P(\overline{F}_t^1, F_t^2, \overline{\varphi}_t | E). \end{aligned}$$

The rule that does fire,  $R^2$ , must either be a positive rule or a negative rule. If  $R^2$  is positive we have directly from Equation (9) that

$$P(\varphi_{t>} | \overline{F}_t^1, F_t^2, \overline{\varphi}_t) = P(\varphi_{t>} | F_t^2) = 1.$$

Further, if  $R^2$  is positive and *did* fire then  $\varphi$  must have been false at  $t$ , so

$$P(F_t^2 | \varphi_t) = 0$$

and

$$P(\overline{F}_t^1, F_t^2, \varphi_t | E) = 0.$$

Likewise, if  $R^2$  is a negative rule,

$$\begin{aligned} P(\varphi_{t>} | \overline{F}_t^1, F_t^2, \overline{\varphi}_t) &= P(\varphi_{t>} | F_t^2) = 0 \\ P(\overline{F}_t^1, F_t^2, \overline{\varphi}_t | E) &= 0. \end{aligned}$$

The case of two rules firing simultaneously is trivial:  $R^1$  and  $R^2$  must either both be positive or both be negative, otherwise they could not both fire at the same instant. If they are both positive then  $P(\varphi_{t>}) = 1$  since by assumption either is sufficient to make  $\varphi$  true; if both are negative then  $P(\varphi_{t>}) = 0$ .

Next we turn attention to computing the firing probabilities—the right sides of the causal update equation, for example

$$\begin{aligned} & P(\overline{F}_t^1, F_t^2, \varphi_t | E) \\ & P(\overline{F}_t^1, F_t^2, \overline{\varphi}_t | E). \end{aligned}$$

The questions to address are to what extent the rule firings interact, and to what extent rule firings depend on the previous evidence  $E$ .

First of all, as we hinted above, we will not worry about multiple rule firings: interacting positive and negative rules are not the issue here, so the danger is that we overestimate the probability of a state change by double counting evidence common to the firing of two positive (or two negative) rules.

The assumption that tends to mitigate this problem involves independence of event occurrences. In Section 2.3.2 we assumed events occurred “dispassionately” in the world: that the probability of an event occurring did not depend on that event’s implications. We can formalize that assumption by assuming that an event occurs independent of the preconditions of the causal rules in which it participates. So, considering a positive rule  $R^1$ , we can apply these assumptions and write its firing probability as follows:

$$\begin{aligned} P(F_t^1 | E) &\equiv P(P_t^1, T_t^1, \overline{\varphi}_t | E) \\ &= P(T_t^1 | P_t^1, \overline{\varphi}_t, E) P(P_t^1 | \overline{\varphi}_t, E) P(\overline{\varphi}_t | E) \\ &= P(T_t^1) P(P_t^1 | \overline{\varphi}_t, E) (1 - p) \end{aligned}$$

where  $p$  is the estimate of  $\varphi$  we get from prior updates. The independence assumption on events allows us to write  $P(T_t^1 | P_t^1, \overline{\varphi}_t, E) = P(T_t^1)$ : we can assess the probability of events occurring separate from the rules in which they participate.

We will further make a Markov assumption with regards to the rule’s precondition: that  $\varphi$ ’s state at  $t$  renders all prior information irrelevant—that  $P(P_t^1 | \overline{\varphi}_t, E) = P(P_t^1 | \overline{\varphi}_t)$ . This assumption is questionable—later in the paper we will examine it in more detail.

### 3.3.4 Causal updates: summary

The assumptions above allow us to consider rule firings (relevant to  $\varphi$ ) at a time point  $t$  to be mutually exclusive and exhaustive. That is, either exactly one of the rules fire, or none of the rules fire. Under these assumptions the causal update equation can be rewritten as follows: let

- $R^+$  be the set of all positive rules whose consequent is  $\varphi$ ,
- $R^-$  be the set of all negative rules whose consequent is  $\varphi$ ,
- $F_t^+ \equiv \sum_{r \in R^+} P(F_t^r | E)$
- $F_t^- \equiv \sum_{r \in R^-} P(F_t^r | E)$
- $p = P(\varphi_t | E)$ ,

then  $\varphi$ 's probability at  $t^>$  can be computed by realizing that  $\varphi$  is true at  $t^>$  just in case it was true at  $t$  and no negative rules fired at  $t$ , or if  $\varphi$  was false at  $t$  and at least one positive rule fired at  $t$ :

$$P(\varphi_{t^>} | E) = p(1 - F_t^-) + (1 - p)F_t^+. \quad (19)$$

We will once again rewrite this update equation as a function of the prior probability  $p$  and relative to a time point  $t_i$ , giving:

$$c_i(p) = p(1 - F_i^-) + (1 - p)F_i^+ \quad (20)$$

where  $F_i^+$  and  $F_i^-$  are the firing probabilities for all positive and negative rules respectively, evaluated at  $t_i$ .

This equation, like Equation (7) dictates how to update a probability given a new piece of evidence (in this case a causal rule instance). Note the difference in the forms of the equations, however: Equation (7) takes a prior at  $t$  and produces a posterior at  $t$  as well whereas Equation (20) takes a prior at  $t$  and produces a posterior at  $t+1$ .

## 3.4 Unpredicted changes

We discussed in the previous section the case of a causal theory in which a rule was not sufficient to predict a change in its consequent's state. The problem is the case in which all the rules are not necessary to predict a state change—that is, we must admit the possibility that a proposition might change even though no rule actually predicted that it would.

In the previous section we made the probably innocuous assumption that unpredicted change can't happen in the instant that passes between  $t$  and  $t^>$ , but it is unrealistic to expect the assumption to hold over longer periods of time, e.g. between times at which our sensors provide us with observations of the proposition. For example, parking one's car in a parking lot and returning three weeks later to pick it up requires that one admit the possibility that it may not still be there, even without being able to say with any precision how it might have disappeared. Leaving one's wallet on the New York subway even for a short period of time carries with it the strong possibility that its location has changed.

The point is that we may not have control over the times at which we get evidence about a proposition of interest, and therefore have to be able to reason about the possibility of change even in the absence of explicit evidence.

Assume now that we have a set of time points  $\{t_1, t_2, \dots, t_n\}$  such that  $t_i < t_{i+1}$ . Each  $t_i$  will correspond to a point in time at which a sensor observation was received or at which a relevant causal rule was likely to have fired—in other words, a time point at which we can invoke one of the update rules above. For the sake of notational simplicity we will assume that the interval of time between  $t_i$  and  $t_i^>$  is negligible, so we will speak of updating over the interval  $[t_i, t_{i+1}]$  when we really mean  $[t_i^>, t_{i+1}]$ . Recalling that  $\varphi_i$  abbreviates  $\varphi_{t_i}$ , we can write this probability

$$P(\varphi_{t_{i+1}}) = P(\varphi_{t_{i+1}} | \varphi_i)P(\varphi_i) + P(\varphi_{t_{i+1}} | \bar{\varphi}_i)P(\bar{\varphi}_i). \quad (21)$$

$P(\varphi_i)$  is the probability resulting from the previous observational or causal update. A state change between  $t_i$ <sup>></sup> and  $t_{i+1}$  is an *unpredicted* state change.

The usual approach to reasoning about probabilistic dynamic systems, e.g. in [Howard 1971], is to assume the probabilities  $P(\varphi_{i+1}|\varphi_i)$  and  $P(\varphi_{i+1}|\bar{\varphi}_i)$  (called “transition probabilities”) are available directly as inputs to the system. Usually this assumption is justified by controlling the amount of time that passes between any  $t_i$  and the next  $t_{i+1}$ . By making this interval short enough compared to the expected frequency at which  $\varphi$  changes state, one can assume that  $\varphi$  can change state *at most once* in any interval  $[t_i, t_{i+1}]$ . Believing that  $\varphi$  changed state in an interval therefore implies that its state at  $t_{i+1}$  is the opposite of what it was at  $t_i$ , and of course believing that  $\varphi$  did *not* change state in the interval implies the belief that its state at  $t_{i+1}$  is the *same* as it was at  $t_i$ . Predicting  $\varphi$ ’s new state is therefore the same as predicting whether it changed state at all.

When our view of assessment involves reasoning about the persistence of  $\varphi$  over longer intervals of time, however,<sup>12</sup> we invalidate the assumption that a state change amounts to a known transition. We may be reasonably sure that  $\varphi$  changed state at least once during a (long) interval of time, but still be unsure of  $\varphi$ ’s state at the end of that interval.

Consider, for example, the problem of predicting whether a certain cat is currently in the living room. Suppose that it’s now 4PM, and that it is known (by observation) that she was there at noon, four hours ago. We have no *specific* information that she has left (mainly because we don’t have a very good causal model of her behavior), but we can be fairly certain at least that she has not stayed put for the whole time. That is, the probability of a state change over that four-hour period is very high. Does that mean she is not in the room? No. She might have left, then reentered, left, entered and left again, and so on. In fact, to the extent that it is likely she has left the room, our causal model loses its predictive power—an observation of her location at noon loses its predictive impact. At that point predictions must start to depend only on her “typical” behavior.

To represent this method of reasoning we introduce a new event into the probability space:

$$\Delta_{i,j}^\varphi \equiv \text{“}\varphi \text{ experienced at least one unpredicted state change between } t_i \text{ and } t_j\text{,”}$$

where  $t_i$  and  $t_j$  are two arbitrary time points such that  $t_i \preceq t_j$  and “unpredicted” means that no known causal rule firing explains the change. Then we rewrite the transition probabilities in (21) to reflect the fact that  $\varphi$  can be true at  $t_{i+1}$  just in case (1) it was true at  $t_i$  and nothing changed it between  $t_i$  and  $t_{i+1}$ , or (2) it was true at  $t_i$  and changed, but ended up true, or (3) it was false at  $t_i$  and changed and ended up true:

$$\begin{aligned} P(\varphi_{i+1}|\varphi_i) &= P(\varphi_{i+1} | \varphi_i, \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi | \varphi_i) + \\ &\quad P(\varphi_{i+1} | \varphi_i, \bar{\Delta}_{i,i+1}^\varphi)P(\bar{\Delta}_{i,i+1}^\varphi | \varphi_i) \\ &= P(\varphi_{i+1} | \varphi_i, \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi | \varphi_i) + \\ &\quad 1 - P(\Delta_{i,i+1}^\varphi | \varphi_i) \\ P(\varphi_{i+1} | \bar{\varphi}_i) &= P(\varphi_{i+1} | \bar{\varphi}_i, \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi | \bar{\varphi}_i) + \\ &\quad P(\varphi_{i+1} | \bar{\varphi}_i, \bar{\Delta}_{i,i+1}^\varphi)P(\bar{\Delta}_{i,i+1}^\varphi | \bar{\varphi}_i) \\ &= P(\varphi_{i+1} | \bar{\varphi}_i, \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi | \bar{\varphi}_i) \end{aligned}$$

The parameters  $P(\Delta_{i,i+1}^\varphi|\varphi_i)$  and  $P(\Delta_{i,i+1}^\varphi|\bar{\varphi}_i)$  are called the *state-change* probabilities for  $\varphi$ . If we assume that  $\varphi$ ’s state at  $t_i$  is irrelevant to its state at  $t_{i+1}$  *assuming a state change took place in the interim*, we can rewrite

$$P(\varphi_{i+1} | \varphi_i, \Delta_{i,i+1}^\varphi) = P(\varphi_{i+1} | \bar{\varphi}_i, \Delta_{i,i+1}^\varphi) = P(\varphi_{i+1} | \Delta_{i,i+1}^\varphi)$$

which we call the *post-change* probability for  $\varphi$ . The state-change and post-change probabilities together we call the *background* probabilities for  $\varphi$ .

---

<sup>12</sup>Recall that we have no *a priori* bound on the amount of time that might elapse between a point  $t_i$  and the next distinguished time point  $t_{i+1}$ .

Now note that  $\overline{\Delta}_{i,i+1}^\varphi$  represents a situation in which *no* state change, predicted or unpredicted, occurs between  $t_i$  and  $t_{i+1}$  and therefore by definition

$$\begin{aligned} P(\varphi_{i+1} \mid \varphi_i, \overline{\Delta}_{i,i+1}^\varphi) &= 1 \\ P(\varphi_{i+1} \mid \overline{\varphi}_i, \overline{\Delta}_{i,i+1}^\varphi) &= 0. \end{aligned}$$

The new equation for updating probabilities over intervals then becomes:

$$\begin{aligned} P(\varphi_{i+1}) &= P(\varphi_{i+1} \mid \varphi_i)P(\varphi_i) + P(\varphi_{i+1} \mid \overline{\varphi}_i)(1 - P(\varphi_i)) \\ &= (P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi, \varphi_i)P(\Delta_{i,i+1}^\varphi \mid \varphi_i) + P(\varphi_{i+1} \mid \overline{\Delta}_{i,i+1}^\varphi, \varphi_i)P(\overline{\Delta}_{i,i+1}^\varphi \mid \varphi_i))P(\varphi_i) + \\ &\quad (P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi, \overline{\varphi}_i)P(\Delta_{i,i+1}^\varphi \mid \overline{\varphi}_i) + P(\varphi_{i+1} \mid \overline{\Delta}_{i,i+1}^\varphi, \overline{\varphi}_i)P(\overline{\Delta}_{i,i+1}^\varphi \mid \overline{\varphi}_i))P(\overline{\varphi}_i) \\ &= (P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi \mid \varphi_i) + (1)(1 - P(\Delta_{i,i+1}^\varphi \mid \varphi_i)))P(\varphi_i) + \\ &\quad (P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi \mid \overline{\varphi}_i) + (0)(1 - P(\Delta_{i,i+1}^\varphi \mid \overline{\varphi}_i)))(1 - P(\varphi_i)) \\ &= P(\varphi_i)(P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi \mid \varphi_i) + 1 - P(\Delta_{i,i+1}^\varphi \mid \varphi_i)) + \\ &\quad (1 - P(\varphi_i))P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi)P(\Delta_{i,i+1}^\varphi \mid \overline{\varphi}_i) \end{aligned} \tag{22}$$

which depends on the two state-change parameters, one post-change parameter, and a prior on  $\varphi_i$ . The equation says essentially that  $\varphi$  will be true at  $t_{i+1}$  if and only if

- it was true at  $t_i$  and there was no state change in the interim, or
- it was true at  $t_i$  and there was a state change in the interim, but it ended up true anyway, or
- it was false at  $t_i$  and there was a state change, and it ended up true.

Now for some  $t_i$  we will summarize the necessary probability parameters as follows:

$$\begin{aligned} d_i^+ &= P(\Delta_{i,i+1}^\varphi \mid \varphi_i) \\ d_i^- &= P(\Delta_{i,i+1}^\varphi \mid \overline{\varphi}_i) \\ b_i &= P(\varphi_{i+1} \mid \Delta_{i,i+1}^\varphi), \end{aligned}$$

which allows us to write the persistence update equation once again as a function of an input parameter  $p = P(\varphi_i)$ :

$$w_i(p) = p(1 - d_i^+) + pd_i^+ b_i + (1 - p)d_i^- b_i. \tag{23}$$

It is often reasonable to assume that the two state-change probabilities are exponential and that the background probability is constant, in which case this equation can be reduced to an exponential with a single parameter (see Section 6). Our implementation makes the less restrictive assumption that background probabilities can depend on the contextual propositions but not on the temporal propositions themselves. It would therefore be more accurate to write the background probabilities explicitly in terms of the contextual propositions  $C$ , e.g.  $d_i^+ = P(\Delta_{i,i+1}^\varphi \mid \varphi_i, C)$ .

### 3.5 The update process for fixed evidence

Recall the three update equations, corresponding to the evidential effects of sensor observations, causal rules, and “uncaused persistence” respectively:

$$s_i(p) = \begin{cases} \frac{5(1+r_i)p}{r_i p + 5(1-r_i)} & \text{if } o_i \text{ is positive} \\ \frac{(1-r_i)p}{1+r-2r_i p} & \text{if } o_i \text{ is negative.} \end{cases} \tag{24}$$

$$c_i(p) = p(1 - F_i^-) + (1 - p)F_i^+ \tag{25}$$

$$w_i(p) = p(1 - d_i^+) + pd_i^+ b_i + (1 - p)d_i^- b_i \tag{26}$$

which suggests the following algorithm for computing the probability  $P(\varphi_{n+1})$  based on a set of distinguished time points  $\{t_i\}$  representing a set of causally relevant events that occurred at those times:

1. Compute a prior on  $\varphi$  at  $t_1$ :  $P \leftarrow P(\varphi_1)$ .
2. For  $i = 1, 2, \dots, n$ 
  - (a) compute and apply the appropriate update function  $(s_i, c_i)$  depending on what sort of event occurred at  $t_i$ , e.g.,  $P \leftarrow s_i(P)$
  - (b) “persist  $P$  forward” from  $t_i$  to  $t_{i+1}$ : compute  $w_i$  for the interval  $[t_i, t_{i+1}]$  and let  $P \leftarrow w_i(P)$ .

One can alternatively view this algorithm as a function  $A_{\mathbf{E}}(\mathbf{p})$ , where  $\mathbf{E}$  is the set of events that occurred at the  $\{t_i\}$ .  $A_{\mathbf{E}}$  is the composition of appropriate instances of  $s$ ,  $c$ , and  $w$ . The input parameter  $\mathbf{p}$  is the prior  $P(\varphi_{t_1})$ . Note that  $A(\mathbf{p})$  is monotone nondecreasing in  $\mathbf{p}$  since all of its component functions are nondecreasing in their input parameters.

### 3.5.1 Conditioning on contextual propositions

We have left out one aspect of the update algorithm: the role of the contextual propositions. Recall that the equations that define the update algorithm depend on background probabilities (priors on the facts and events and probabilities for unpredicted state changes over intervals), and these background probabilities in turn depend on the contextual propositions. The same contextual proposition can figure in more than one background probability used in an update, so to avoid double counting we must explicitly condition on these propositions. Strictly speaking, then, we run the update algorithm once for each possible assignment of truth values to the contextual propositions that appear in the equations, weighting the result by the probability associated with that assignment. While this represents a potential explosion of computation time, in practice it has not been an issue: for all applications to this point (see Section 6 for example) contextual variables have not figured in the background probabilities at all. Generally a state-independent number suffices for a prior, and for the state-change probabilities we have been using an exponential function that depends only on the duration of the interval. Also note that since the states of contextual propositions do not depend on the temporal facts and events, we need calculate a probability distribution over those propositions only once.

## 3.6 An equivalent graph structure

So far we have defined a model for computing probabilities for propositions that vary over time, stating this model implicitly in functional or algorithmic terms. We can also describe the model in terms of a graph [Pearl 1988]: each node in the graph represents a quantity that can be assigned a probability and the arcs connecting nodes indicate a probabilistic dependency. Figure 3 shows the structure for a simple example:

- There are two fact types,  $P$  and  $Q$ .
- There are two event types,  $E$  and  $F$ .
- There are two rules:
  - R1: If  $E$  occurs when  $P$  is true then  $Q$  becomes true (perhaps with some probability).
  - R2: If  $F$  occurs then  $P$  becomes true (perhaps with some probability).

The graph shows the model’s structure for a single update, between two time points,  $t_i$  and  $t_{i+1}$ . Recall that an arbitrary amount of time might separate  $t_i$  and  $t_{i+1}$ , but the expectation is that the system did not make any “interesting” state changes in the interim.

Each row on the grid represents a proposition, each column represents a point in time. In addition to the facts, events, and rules we mentioned above we model the possibility that a fact is observed to be true




Figure 3: Graph structure for a single update



(or false) and that an event is observed to have occurred (or not to have occurred). Each observation has an associated reliability parameter. The nodes in the rows labelled R1 and R2 represent the possibility that the rules fired at the corresponding time point.

The graph does not show the probabilistic parameters—the numbers that determine how a node’s probability is computed based on the states of those nodes that point to it. This application supplies this information in the form of prior probabilities for those nodes with no incoming arcs. The application communicates whether an observation was made by providing a value of `true` or `false` or `none` for the appropriate `obs(x)` node, and also provides the reliability parameter  $r$  as a prior for the associated `rel(x)` node. The application also supplies priors for each fact and event proposition and for the state-change parameters  $\Delta_x$ . These priors can all depend on the contextual propositions, so there should be an arrow pointing from the contextual propositions to all nodes with no other incoming arcs. (These arrows were omitted to unclutter the graph’s diagram.)

The other relationships are determined by the equations developed earlier in the paper. Considering the graph from left to right, we see that updating occurs in four stages, each corresponding to one of the equations presented in this section:

- Equation (7) determines a fact’s or event’s state based on its prior state, whether it was observed true or false, and how reliable that observation was.
- Equation (19) determines whether a rule fired based on the states of its precondition facts and trigger event and on the rule’s effectiveness parameter.
- Equation (20) determines whether a fact will be true based on whether it was previously true and whether any of its rules fired.
- Equation (23) determines whether a fact will be true based on whether it was true before and whether an unpredicted state change took place over the interval.

### 3.6.1 Projected probabilities

We should note that although Figure 3 illustrates the probability space and how elements of the space are related, our update algorithm is *not* the same as the Bayesian propagation algorithm (as described in [Pearl 1988] for example). The difference is that our algorithm computes probabilities only from earlier to later time points, so a proposition’s state at time  $t$  can only have an effect on the system’s state at times subsequent to  $t$ . A full propagation algorithm, on the other hand, also considers the effect of later evidence on earlier propositions.

Suppose, for example, that  $P$  were observed to be false (fairly reliably) at time  $t_i$  but was observed to be true (again fairly reliably) at a subsequent time  $t_{i+1}$ . The full propagation algorithm would reason that  $P$  is probably true at  $t_{i+1}$  and therefore that it must either have been true at time  $t_i$  or it had to have been made true between  $t_i$  and  $t_{i+1}$ . The algorithm would therefore increase the probability that  $P$  was true at  $t_i$  and also that the observation made at  $t_i$  was unreliable. It would also increase the probability that one of  $P$ ’s causal rules fired at  $t_i$  and that  $P$  changed state from false to true between  $t_i$  and  $t_{i+1}$ . These changes in turn would cause the probabilities of propositions at times prior to  $t_i$  to change. In other words, the observation at  $t_{i+1}$  affects probabilities of propositions at  $t_i$  and before.

Our algorithm determines  $P$ ’s state at  $t_i$ , then at  $t_i^>$ , then at  $t_{i+1}$ .  $P$  would therefore be given low probability at  $t_i^>$  based on the negative observation. That would provide a low prior for the observation-update stage at  $t_{i+1}$ . That prior would be combined with the observational evidence at  $t_{i+1}$ , tending to lower the posterior. But the algorithm does *not* then go back and revise its beliefs about  $P$  at  $t_i$  based on the observation at  $t_{i+1}$ —subsequent evidence has no effects on prior beliefs.

## 3.7 Time-based versus state-based models

Our equations so far tell us how to update probabilities given a set of “important” time points—those time points at which the system is likely to change state—but we have no way of reasoning about the system over

some interval of time over which several changes might have occurred. In other words we have no way of identifying the important time points.

One easy way to model the system completely is to consider *every* time point to be potentially important. Suppose we want to project the course of the system over an interval of time  $[t_0, t_n]$ . One way to do so is choose some short interval of time, say  $\delta$ , short enough so that no event is likely to occur within any interval  $[t, t + \delta]$ . This interval is usually called a *time mesh*. We can then instantiate the model by letting the first important time point be  $t_0$ , the second be  $t_0 + \delta$ , the third be  $t_0 + 2\delta$ , and so on to  $t_n$ , duplicating the graph in Figure 3 at every time point. We then provide probabilities for the facts at the initial time point  $t_0$  and propagate that information through the graph. This method is equivalent to the usual method for reasoning about Markov random processes, and is similar to the approach suggested by [Dean and Kanazawa 1989].<sup>13</sup> We call this approach a “time-based” approach to solving the system.

While this is an adequate formal model for solving the system it can be quite inefficient, for several reasons. First it requires reasoning about every time point whether or not anything changes between one time and the next. We would prefer to focus attention only on those times at which a change occurs and ignore those intervals (those parts of the graphs) over which the system is stable. A time-based system forces us to reason about every time point, whether or not a change is likely. We therefore risk spending large amounts of processing resources propagating information through the model over intervals in which nothing happens.

Second, it requires reasoning about the state of every proposition whether or not the application program is interested in that proposition’s state. We would prefer to focus attention only on those propositions that are of direct interest to the application (e.g. its goals) and to those that have a causal impact on those propositions. Third, it requires reasoning about the entire system using a level of granularity (time mesh) short enough to be appropriate for the fastest-changing proposition in the system. Reasoning about hummingbirds and boulders simultaneously requires thinking about whether the boulder has moved on a second-by-second basis.

The alternative to a time-based approach is to try to identify the propositions of interest and the time points at which they are likely to change, and instantiate only those parts of the model relevant to the specific reasoning task at hand. We call this approach a “state-based” approach because it focuses on significant changes in the system’s state. To implement this approach we must identify what changes qualify as significant, and how likely they have to be to demand our attention.<sup>14</sup>

The next section attacks the problem of deciding which parts of the model to instantiate; in the language of the previous sections, this process can also be viewed as the process of identifying an appropriate set of evidence (observations, events, and rules) to reason about explicitly.

## 4 Identifying important evidence

Analysis to this point has relied on the existence of a finite set of relevant events—sensory observations or causal rule instances—occurring at the  $t_i$  points that together should determine the probability  $P(\varphi_{n+1})$ . The problem remains, of course, how we collect these events in the first place. Virtually any piece of evidence—sensory observation or causal rule instance mentioning  $\varphi$  and occurring prior to  $t_n$ —will have *some* impact on  $P(\varphi_{n+1})$ . The number of *actual* pieces of evidence will certainly be enormous: we may have collected many sensory observations of  $\varphi$  over a long period of time, and we have to consider potential causal rule instances for every rule whose consequent is  $\varphi$  *at every point in time prior to  $t_{n+1}$* . Of course if the firing probabilities of these rule instances are close to zero, we see from Equation (25) that they will have little effect on  $P(\varphi)$ .

The intuition that we can get away with examining only a few important events rests not only on the assumption that the important events are fairly decisive (have reliabilities or firing probabilities near 1.0,

---

<sup>13</sup> Also see Section 9.4

<sup>14</sup> The statistics literature sometimes refers to the time-based approach as a “discrete-time” model and the state-based approach as a “continuous-time” model. We avoid these terms to head off confusion with the literature on temporal logics, in which discrete and continuous refer to whether time is represented by the integers or the reals.

thus pretty much determine belief at that point in time) but also on the realization that we generally need only an *approximation* of the exact probability.

My decision whether or not to ride my bicycle to work, for example, may depend on whether or not I think it will be raining the next morning. Weighing the discomfort of getting wet against the inconvenience of having to find a parking place, I arrive at some threshold for the probability of rain, say 0.6, above which I will decide to drive. The threshold presumably comes from an analysis of the planner’s current goals—see [Haddawy and Hanks 1990]. At that point the *exact* probability of rain is of no interest to me; all that matters is what side of the threshold the true probability lies. And by implication, the only evidence of interest to me is some subset of the evidence that will convince me that the probability of rain is above or below that threshold. If, for example, I have a current estimate of the probability that is above the threshold, I needn’t consider *confirming* evidence at all, since it can’t possibly change my judgement relative to the threshold and thus will not affect the decision at hand.

Basically a piece of evidence can be ignored for one of three reasons: first, as above, it can be the wrong “sign”—it will only tend to confirm the current hypothesis. Second, it can be too unreliable. Recall from the update equations that sensory and causal evidence affect the estimate only to the extent that they are reliable, or likely to fire. An unreliable observation, for example, may not influence my belief enough to change it with respect to the threshold. Finally, a piece of evidence may be too remote: the farther in the past a piece of evidence occurs the more likely it is that an intervening state change will occur, which renders the evidence irrelevant. In trying to figure out where my car is, for example, I don’t want to dwell on the fact that I saw it in my driveway two years ago.

All three criteria for ignoring evidence can be applied to the causal rules as well. Suppose that the query proposition is  $\varphi$ , and we have a rule that says event  $\mathbf{E}$  and proposition  $\mathbf{P}$  together cause  $\mathbf{P}$ . Then we can compute how likely it must be for this rule to “fire” to affect our estimate of  $\varphi$ , which in turn gives us bounds on how likely  $\mathbf{E}$ ’s occurrence and  $\mathbf{P}$ ’s truth must be to warrant our attention.

These limits suggest an algorithm for computing the probability of  $\varphi$  at  $t_{n+1}$ , relative to some threshold  $\tau$ :

1. Start with no evidence,  $\mathbf{E} = \phi$ . ( $\mathbf{E}$  will be the set of relevant events.)
2. Estimate  $P(\varphi_{n+1})$  using the algorithm from the previous section, that is, letting  $E = P(\varphi_{n+1} \mid \mathbf{E}) = A_{\mathbf{E}}(\cdot)$ .
3. If  $E = \tau$ , perturb  $\tau$  slightly, letting  $(\tau \leftarrow \tau + \varepsilon)$  or  $(\tau \leftarrow \tau - \varepsilon)$ .<sup>15</sup>
4. If  $E > \tau$ , search for new evidence  $\mathbf{O}$  (observations and rule instances) such that, if it were to be found, would change the estimate such that  $E' = P(\varphi_{t_{n+1}} \mid \mathbf{E} \cup \mathbf{O}) < \tau$ .
5. Otherwise, if  $E < \tau$ , search for new evidence  $\mathbf{O}$  (observations and rule instances) such that, if it were to be found, would change the estimate such that  $E' = P(\varphi_{t_{n+1}} \mid \mathbf{E} \cup \mathbf{O}) > \tau$ .
6. If no such evidence is found, report that  $P(\varphi_{t_{n+1}})$  stands in relation to  $\tau$  as  $E$  does. Otherwise let  $\mathbf{E} \leftarrow \mathbf{E} \cup \mathbf{O}$  and go back to Step 2.

We stated the retrieval problem as that of finding a body of evidence  $\mathbf{O}$  such that adding  $\mathbf{O}$  to our current body of evidence  $\mathbf{E}$  would change our estimate of  $\varphi$ ’s probability with respect to the input threshold. In the process we might find that there *is* no such  $\mathbf{O}$ , in which case we terminate. Establishing that there is no such set requires, in the worst case, examining all possible evidence, which is exactly what we want to avoid: we would have to verify that there is no set  $\mathbf{O}$  of size 1, of size 2, and so on, that invalidates our current estimate.

We therefore have to adopt a heuristic criterion for deciding that there is no such body of evidence (implying that our current estimate is correct). Subsequent sections adopt the following rule: we will

---

<sup>15</sup>The rules below for bounding the search for evidence are based on  $\tau$ , and are undefined when  $E = \tau$ . In that case one might search arbitrarily carefully, since *any* piece of evidence, positive or negative would push  $E$  over or under the threshold  $\tau$ .

terminate the assessment if we find no *single* piece of evidence that upsets the current estimate. We base this policy on the observation that in many planning applications a single event, observation, or action often determines a proposition’s state—rarely do several unlikely events conspire to make a proposition likely. The analysis in subsequent sections could instead use heuristics that consider evidence subsets of size 2 or greater; these heuristics tend to be more expensive to calculate and require searching more of the database, however, so more accurate policies carry with them a corresponding computational burden.

## 4.1 Bounding the reliability of evidence

We now confront the problem of how to identify the new evidence. In addition to considering evidence of the correct sign, we also want to limit consideration to evidence that is reliable enough, and recent enough, to change the current estimate.

First consider the reliability problem, and assume that  $E < \tau$ , so we are looking only for positive evidence. The impact of a (single) piece of evidence—the extent to which it changes the probability estimate—can be measured by the appropriate  $s$  or  $c$  equation, but in either case the impact depends on the time at which that evidence is noted. It is generally the case that the more recent the evidence, the greater its impact. This will strictly be the case if the firing probabilities for the causal rules do not depend on  $\varphi$ . So we can conservatively ask: how reliable must a piece of evidence, *noted at*  $t_{n+1}$ , be in order that the new estimate  $E'$  would exceed  $\tau$ .

Consider the positive case for  $s$ , which measures the evidential impact of a positive sensory observation. We know that the current estimate of  $\varphi$  at  $t_{n+1}$  is  $E$ . Setting the reliability parameter in equation (7) to the desired level  $r^*$ :

$$\frac{.5(1+r^*)E}{r^*E + .5(1-r^*)} > \tau$$

or, in order for a positive observation to be interesting it must have a reliability parameter that exceeds

$$r^* > \frac{\tau - E}{E - 2\tau E + \tau}. \quad (27)$$

If we were instead searching for negative observations we would look for

$$r^* > \frac{E - \tau}{E - 2\tau E + \tau}. \quad (28)$$

Similarly for causal rules, we use equation (20) to get

$$f^* > \frac{\tau - E}{1 - E} \quad (29)$$

or for negative rules we use

$$f^* > 1 - \frac{\tau}{E}. \quad (30)$$

Recall that  $f$  is the probability that a particular rule fires at some time  $t$ , assuming that  $\varphi$  is in the correct state. Note one detail here: we have not required that  $f$  be independent of the state of  $\varphi$  at that time, so strictly speaking we cannot use this formula for bounding  $f$ . In the case where a rule depends on a precondition  $\psi$ , and  $\psi$  is true just in case  $\varphi$  is false, then  $f$  and  $\varphi$  are highly correlated and we will tend to underestimate the rule’s evidential impact. We *did* assume, however, that the probability of trigger events occurring, and also the “effectiveness” of the rule, could both be assessed independent of  $\varphi$ ’s state. So a conservative policy would be to consider all those rule instances for which

$$P(T)\chi > \frac{\tau - E}{1 - E}$$

where  $T$  is the rule’s trigger event and  $\chi$  is the rule’s effectiveness.

## 4.2 Bounding the search’s temporal scope

The final problem is to limit how far back in time we search for evidence. We ask the question: how far back in time would a perfectly reliable piece of evidence have to occur in order that if we were to find it (and it alone), it would change our estimate  $E$  with respect to  $\tau$ ?

Recall that the current estimate  $E$  is based on the estimate  $A_{\mathbf{E}}(\mathbf{p})$ , which we will abbreviate  $A(\mathbf{p})$ . If currently  $E(\mathbf{p}) < \tau$ , we can ask two questions:

- what value  $\mathbf{p}^*$  will cause  $A(\mathbf{p}^*) > \tau$ ,
- what is the latest possible time point  $t^* \prec t_1$  such that a perfectly reliable *positive* observation taken at  $t^*$  would induce a probability at least as great as  $\mathbf{p}^*$ ?

And of course the same holds if  $E > \tau$ , except substitute “ $<$ ” for “ $>$ ” and “*negative*” for “*positive*.”

The first question obviously depends on the form of the function  $A(\mathbf{p})$ . We can most easily arrive at a  $\mathbf{p}^*$  value by using binary search and repeatedly evaluating  $A$ , since we noted that  $A$  is nondecreasing in  $\mathbf{p}$ . The second quantity depends on the state change and background probabilities for  $\varphi$ :  $\mathbf{d}^+$ ,  $\mathbf{d}^-$ , and  $\mathbf{b}$ . For certain forms of these functions we can compute exactly the time  $t^*$  (see Section 6), otherwise we can compute  $t^*$  by binary search as well, since  $\Delta_{i,j}^\varphi$  must increase as  $\mathbf{d}(t_i, t_j)$  increases.

## 5 Algorithm summary

### 5.1 Algorithm description

We can now restate the estimation algorithm from Section 4 as follows:

- Inputs:
  - A proposition  $\varphi$ , a time point  $t$ , a probability threshold  $\tau$ .
  - A set of causal rule schemas such as the one described in Section 3.3.1.
  - Probabilities concerning the occurrence of trigger events and on the background probabilities of facts at points in time.
  - Sensory observations of various facts and events, each with an associated proposition, reliability parameter and time point.
- Output:
  - An estimate  $E$  of  $P(\varphi_t)$  with the property that if  $E < \tau$  then  $P(\varphi_t) < \tau$ , and if  $E > \tau$  then  $P(\varphi_t) > \tau$ .
  - Monitors in the temporal database ensuring that the estimate will be recomputed if subsequent evidence calls this relationship into question (see next section).
- Method:
  1. Initialize  $\mathbf{K} = \phi$ .
  2. Compute an estimate  $E = A_{\mathbf{K}}(\cdot)$ , using the algorithm from Section 3.5.
  3. If  $E = \tau$  then randomly set  $\tau \leftarrow \tau \pm \varepsilon$ .
  4. Determine reliability, firing, and recency bounds on evidence that would change  $E$  with respect to  $\tau$ , using, respectively, equation (27) or (28), equation (29) or (30), and the method of Section 4.2.
  5. Search through the temporal database for these new events.
  6. If none are found, report  $E$  as the estimate for  $P(\varphi_t)$ .

7. Otherwise add the newly found evidence to  $\mathbf{K}$ , and go back to step 2.

The algorithm is heuristic in the sense that it considers only a *single* piece of evidence that would change its estimate in computing the reliability and recency bounds. The estimate may be incorrect only if several pieces of evidence considered jointly would change the estimate with respect to the threshold even though any piece of evidence taken alone would be too unreliable to do so.

## 5.2 Maintaining the belief: monitors and signals

We mentioned in the introduction that a belief, as computed by the algorithm above, also represents a promise to notify the calling program in case subsequent evidence invalidates the relationship between the current estimate and the belief’s threshold.

Providing this functionality is quite easy at this point in that the last search iteration (which yielded no evidence) also tells us what new evidence (in terms of sign, recency, and reliability) we should look for. Once the estimate is established we therefore establish *add monitors* in the temporal database (Section 2.2) telling it to notify the belief-computation algorithm if any such new evidence should appear. When new evidence arrives we recompute the estimate including that new evidence. If the estimate changes with respect to the threshold we call the application-supplied signal function associated with the belief; if the new estimate stands in the same relationship to the threshold as before, we re-compute the recency and reliability bounds and post new add monitors, but do not call the signal function.

We will devote the remainder of the paper to an example illustrating the algorithm, then to a discussion of the algorithm’s implementation, the assumptions we have made, and related and future work.

## 6 Example

To illustrate the algorithm we will work through a simple example. The proposition of interest is

CAT = “the cat is now in the living room,”

and it is now 2:00 in the afternoon. There are two relevant sensory observations: at 9:00 this morning the cat was seen in the room. But at 1:00 this afternoon a meow-like noise was sensed coming from the kitchen, which might (or might not) have been the cat. If it was, then the cat couldn’t possibly have been in the living room at that time, but if the sound came from some other source, then it doesn’t help me determine where the cat was. We estimate at 50% the probability that the noise really was the cat, therefore count the experience as a negative sensory observation with a reliability parameter  $r = 0.5$ .

Mainly the cat’s behavior is inexplicable, but we do have one rule: when the postman comes he scares the cat out of the living room with probability about 0.6. Furthermore, the probability that the postman comes at noon is 0.8.

Finally, background knowledge about the cat leads us to estimate three parameters:

1. the probability that the cat will *enter* the room unobserved is roughly 1/2 in a 90-minute period
2. the probability that the cat will *leave* the room unobserved is roughly 1/2 in a two-hour period
3. the cat is in the room about 60% of the time.

We take the state-change probabilities to be exponential, with a parameter  $\lambda$  such that the above half-lives are realized. That is,

$$\begin{aligned} P(\Delta_{i,j}^{\text{CAT}} | \overline{\text{CAT}}_i) &= 1 - e^{(-.0077)\mathbf{d}(i,j)} \\ P(\Delta_{i,j}^{\text{CAT}} | \text{CAT}_i) &= 1 - e^{(-.00578)\mathbf{d}(i,j)} \end{aligned}$$

Sensory Observations:

$k_1: s_1=1, t_1=9\text{AM}, r_1=1.0$   
 $k_3: s_2=0, t_2=1\text{PM}, r_2=0.5$

Rule:

```
(DEFINE-RULE-SCHEMA 'SCARED-CAT
  '(TRIGGER POSTMAN)
  '(CONSEQUENT CAT)
  '(SIGN 0)
  '(EFFECTIVENESS 0.6))
```

Event:

$P(\text{POSTMAN}_{12\text{N}}) = 0.8$   
 $P(\text{POSTMAN}_t) = 0 \quad (t \neq 12\text{N})$

Background:

$P(\Delta_{i,j}^{\text{CAT}} | \overline{\text{CAT}}_i) = 1 - e^{(-.0077)d(t_i, t_j)}$   
 $P(\Delta_{i,j}^{\text{CAT}} | \text{CAT}_i) = 1 - e^{(-.0077)d(t_i, t_j)}$   
 $P(\text{CAT}_j | \Delta_{i,j}^{\text{CAT}}) = 0.6$

Figure 4: Summary of cat example's parameters




Figure 5: Summary of cat example's scenario

where  $e^{(-.0077)(90)} = 0.5$  and  $e^{(-.00578)(120)} = 0.5$ . (A time unit equals one minute.)

The problem is to estimate the likelihood of CAT at 2PM; we'll demonstrate with a probability threshold  $\tau = 0.4$ . Figure 4 recapitulates the problem's parameters; Figure 5 shows the exact probability of  $P(\text{CAT}_{2\text{PM}}) = 0.376$ .

Analysis proceeds as follows:

Iteration 1 begins with no known observations. First we compute the current estimate, which is just the prior  $E = 0.6$

Since  $E > \tau$ , we search for negative observations. First we determine what the current prior would have to be in order to lower the current estimate (0.6) below the threshold (0.4). Since the current prior *is* the current estimate, the answer is obviously 0.4.

Next we find the time point  $t^*$  such that a perfect negative observation taken at  $t^*$  and persisted to 2PM would yield a posterior value at or below 0.4. Recall that if we have a perfect negative observation at  $t^*$ , then  $P(\text{CAT}_{t^*}) = 0$ . We can then write

$$\begin{aligned} P(\text{CAT}_{2\text{PM}} | \overline{\text{CAT}}_{t^*}) &= P(\text{CAT}_{2\text{PM}} | \Delta_{t^*, 2\text{PM}}^{\text{CAT}}) P(\Delta_{t^*, 2\text{PM}}^{\text{CAT}} | \overline{\text{CAT}}_{t^*}) \\ &= (0.6)(1 - e^{(-.0077)d(t^*, 2\text{PM})}) \end{aligned}$$

and set the left-hand-side to the threshold,

$$\begin{aligned} 0.4 &= 0.6(1 - e^{(-.0077)d(t^*, 2\text{PM})}) \\ 0.667 &= 1 - e^{(-.0077)d(t^*, 2\text{PM})} \\ d(t^*, 2\text{PM}) &= \frac{\ln(0.333)}{-.0077} \\ d(t^*, 2\text{PM}) &= 142 \\ t^* &\approx 11:40\text{AM}. \end{aligned}$$

We next have to compute the bound on  $r$  and  $f$  that will make a negative causal observation interesting. Recall that we do so by postulating a negative observation exactly at the query time 2PM. What can we say about the values of  $r$  and  $f$  that will lower the posterior probability to 0.4 or below? For sensory observations we can set the negative case of update equation (24) equal to  $\tau$  and solve for  $r$ :

$$r^* = \frac{p - \tau}{\tau + p - 2\tau p}$$

and setting the prior parameter  $p$  to the current estimate 0.6 and  $\tau$  to the current threshold 0.4:

$$\begin{aligned} r^* &= \frac{0.6 - 0.4}{0.4 + 0.6 - 2(0.4)(0.6)} \\ r^* &= 0.38 \end{aligned}$$

Similarly for negative causal observations, we set  $\tau = 0.4$  and  $p = 0.6$  in Equation (25)

$$\begin{aligned} f^* &= 1 - \frac{\tau}{p} \\ &= 1 - \frac{0.4}{0.6} \\ f^* &= 0.33 \end{aligned}$$

That is, either a negative sensory observation with reliability greater than 0.38 or a causal rule instance with impact greater than 0.33 would cause us to change our mind and conclude that  $P(\text{CAT}_{2\text{PM}})$  is actually less than 0.4.

We therefore initiate a fetch in the temporal database, searching for all negative sensory observations from 11:40AM to 2PM for which  $r \geq 0.38$  and for all negative causal observations in the same period for which  $f \geq 0.33$ .

The search nets us  $k_3$ —the negative sensory observation at 1PM, and also  $k_2$ —the negative causal rule instance that occurred at 12N. Since we *did* find evidence as a result of the fetch, we add  $k_2$  and  $k_3$  to our set  $\mathbf{K}$  of relevant events, and initiate another estimate/fetch cycle.

First we recompute the estimate  $E$  on the basis of the newly-enhanced  $\mathbf{K}$ , which produces  $E = 0.354$ . And we find ourselves in the opposite position: the current estimate is *below* the threshold, so we have to look for positive observations that will put it back over 0.4.

Now we recompute the reliability, recency, and firing bounds. First we consider  $E = A(p)$  where  $p$  is the prior at the earliest observation (in this case noon). Recall that  $A(p)$  is non-decreasing in  $p$ . We thus ask the



Evidence <b>K</b>	P(CAT  <b>K</b> )
$\phi$	0.6
$k_1$	0.671
$k_2$	0.498
$k_3$	0.442
$k_1, k_2$	0.530
$k_1, k_3$	0.511
$k_2, k_3$	0.354
$k_1, k_2, k_3$	0.376

Table 1: Probabilities for different collections of evidence

question: how high above 0.6 would  $p$  have to be raised in order to cause  $A(p)$  to move from 0.354 to 0.4. We get the answer through binary search, repeatedly evaluating  $A$  to get  $p^* = 0.88$ . Then we do a state-change analysis as above, asking the question of how recent a perfectly reliable positive observation would have to be in order to raise the probability at noon to at least 0.88.

$$\begin{aligned}
P(\text{CAT}_{12\text{N}} | \text{CAT}_{t^*}) &= P(\text{CAT}_{12\text{N}} | \text{CAT}_t, \Delta_{t^*, 12\text{N}}^{\text{CAT}}) P(\Delta_{t^*, 12\text{N}}^{\text{CAT}} | \text{CAT}_t) + \\
&\quad P(\text{CAT}_{12\text{N}} | \text{CAT}_t, \overline{\Delta}_{t^*, 12\text{N}}^{\text{CAT}}) P(\overline{\Delta}_{t^*, 12\text{N}}^{\text{CAT}} | \text{CAT}_t) \\
0.88 &= (0.6)(1 - e^{(-.00578)d(t^*, 12\text{N})}) + (1)(e^{(-.00578)d(t^*, 12\text{N})}) \\
0.88 &= 0.6 + 0.4e^{(-.00578)d(t^*, 12\text{N})} \\
d(t^*, 12\text{N}) &= \frac{\ln(\frac{0.28}{0.4})}{-.00578} \\
d(t^*, 12\text{N}) &= 61.7 \\
t^* &\approx 11\text{AM}.
\end{aligned}$$

We also derive new reliability and firing bounds:  $r^* = 0.1$ ,  $f^* = 0.07$ . The resulting search for observations nets us nothing, however: we search only back to about 11AM, and thus exclude the 9AM positive observation. Since the search yields no new information we report that our current estimate can't change with respect to the threshold 0.4, and thus  $P(\text{CAT}_{2\text{PM}}) < 0.4$ .

Table 1 shows the probability for all combinations of the three observations, and Table 2 indicates the amount of work that would need to be done in order to process a number of different thresholds. From Table 1 note that we were justified in ignoring observation  $k_1$ : adding it to the set  $\{k_2, k_3\}$  does not change the estimate's position with respect to the threshold 0.4.

From Table 2 note that we tend to do a lot of work if the threshold  $\tau$  is near the prior  $P(\varphi)$ , and also if  $\tau$  is near the exact probability. This reflects two factors: a fixed cost associated with setting up and executing a fetch in the temporal database, and a variable cost associated with how far back in time the fetch covers. The magnitudes of these numbers reflect the actual implementation, of course—we find in [Hanks and McDermott 1992] that most of the time is spent in the temporal database manager module computing distances between time points, so efficiency gains in that module would improve these results considerably.

## 7 Assessing joint events

The assessment algorithm estimates the probability of simple propositions. We also need to compute probabilities for composite propositions, however, most critically for conjoined or joint events. Recall that a rule's

Threshold	No. pieces evidence	Time (secs)
0.0	0	0.00
0.1	0	0.75
0.2	0	0.85
0.3	0	0.95
0.36	3	5.50
0.4	2	3.05
0.5	2	2.50
0.61	3	5.90
0.7	0	1.55
0.8	0	1.05
0.9	0	0.80
1.0	0	0.00

Table 2: Process time as a function of threshold

precondition consists of a precondition, which may require that several conditions hold jointly in the world at the time the trigger event occurs. This section points out the difficulties associated with assessing joint events, and provides a method for dealing with one particular class of dependencies.

Suppose I want to assess the probability that my cat and my dog are both in the living room at the same time. Both tend to hang out there some fraction of the time (say 25% each), but the dog typically scares the cat out shortly after he enters the room, so rarely are they in the room together for long. Probabilistically this information is usually conveyed by a conditional probability like  $P(\text{CAT} \mid \text{DOG}) = .01$  or some small number, which in turn implies

$$P(\text{CAT}, \text{DOG}) = P(\text{CAT} \mid \text{DOG})P(\text{DOG}) = (.01)(.25) = .025.$$

Conditional probability constraints comment on interrelationships between propositions at a single point in time, while our focus to this point has been on the *dynamic* behavior of propositions *considered in isolation*.

The relationship between CAT and DOG actually summarizes the “typical” behavior of the theory, taking into account the existence of causal rules like “dog entering the room while cat is in the room causes cat to leave the room,” the relative frequency of “dog enterings” to “cat enterings,” other things that might scare the cat off, and so on. Put more directly, it says that if you were to view the system (look in the living room) at various points of time, under a variety of circumstances, only very rarely would you find the cat in the room if the dog was in there too.

But while we must take into account the typical, static behavior of the system, we must not at the same time ignore what we know about the situation at hand. What should we conclude if we were in the living room not 30 seconds before, and saw both animals in the room? (And what if they were both asleep?) In that case we should consider the possibility that this is one of the rare occasions in which the two animals coexist for a while.

## 7.1 The nature of probabilistic dependencies

Given a causal theory consisting of causal rules and background probabilities, there are two ways that propositions (say  $\varphi$  and  $\psi$ ) might be interdependent:

1. Through dependencies in the causal rules.  $\varphi$  might be a precondition of a rule that affects  $\psi$ , for example, or the same event might tend to cause both propositions. Or perhaps the propositions share a rule precondition. In any case, by examining the structure of the causal rules we can identify those propositions that have an effect on both.

2. Through dependencies in the background probabilities. Recall that the state-change and post-change probabilities depend on the contextual propositions.  $\varphi$  and  $\psi$  can be interdependent because their background probabilities are interdependent.

In the first case we rely on a “scenario” or “hypothetical future” structure generated by the projection algorithm to resolve the dependencies. We will formally introduce the scenario structure in [Hanks and McDermott 1992]<sup>16</sup>; here we will just hint at how such a structure can be used to avoid the over- or under-estimation of joint probabilities in those cases where the propositions are interrelated by sharing causal rule components. We can solve restricted forms of the second problem—interdependent background probabilities—without reference to the scenario structure. We will sketch this solution later in this section.

## 7.2 Dependencies in the causal rules

Returning to the cat and dog example, suppose that we have propositions CAT and DOG, and also the event

E-DOG= “dog enters the living room.”

The relationship between CAT and DOG shows up in the causal rules: the event E-DOG of course causes DOG. Furthermore, we have a rule that says the dog usually chases the cat out of the room: if CAT is true, then E-DOG causes  $\overline{\text{CAT}}$  with some probability. If we compute  $P(\text{CAT}, \text{DOG}) = P(\text{CAT})P(\text{DOG})$  we will overestimate, failing to realize that the situations where DOG is likely are the ones in which E-DOG occurred, but those situations are those in which CAT is unlikely. The scenario structure makes these situations explicit, and reasons about each in isolation.

The way it works is that when  $P(\text{CAT}, \text{DOG})$  is assessed, the system notices that the event E-DOG is relevant to both propositions. Then if, in the assessment process, it is asked to assess the probability that an E-DOG event occurs in a particular interval, it “splits the world” into two hypothetical situations, or chronicles, say  $c_1$  and  $c_2$ . In  $c_1$  an E-DOG event actually occurs, and in  $c_2$  it does not. We can assign a probability to each chronicle. Then, reasoning within each chronicle separately, we notice that in  $c_1$  the probability of DOG is high, but, since the “dog chases cat” causal rule probably fired, the probability of CAT is low. Likewise, in  $c_2$ , the probability of DOG is low (since we have no evidence that he entered). The probability of CAT may be high or low, depending on whether we have other evidence to support CAT.

Since  $c_1$  and  $c_2$  are mutually exclusive and exhaustive we can compute

$$\begin{aligned} P(\text{CAT}, \text{DOG}) &= P(\text{CAT}, \text{DOG} | c_1)P(c_1) + P(\text{CAT}, \text{DOG} | c_2)P(c_2) \\ &= P(\text{CAT} | c_1)P(\text{DOG} | c_1)P(c_1) + P(\text{CAT} | c_2)P(\text{DOG} | c_2)P(c_2) \end{aligned}$$

which will turn out to be low. Note, however, that nowhere do we explicitly represent the probability  $P(\text{CAT} | \text{DOG})$ .

We will call this style of reasoning, in which chronicles are introduced and assigned probabilities, “explicit reasoning by cases.” Although it takes dependencies into account properly, it is computationally expensive: all probability calculations must be undertaken with respect to each separate chronicle, so frequent splitting may result in an exponential increase in the complexity of the probabilistic reasoning. We discuss this point further in Section 8.

## 7.3 Dependencies in the background probabilities

Recall from Section 3.4 that we allow the background probabilities to depend on the contextual (atemporal) propositions. Our algorithm (especially see Section 3.5) handles this case correctly: suppose we are jointly assessing P and Q, and their background probabilities both depend on a contextual proposition C. Further suppose that our current body of evidence is  $\mathbf{E}$ . The algorithm correctly computes

$$P(P, Q | \mathbf{E}) = P(P, Q | \mathbf{E}, C)P(C | \mathbf{E}) + P(P, Q | \mathbf{E}, \overline{C})P(\overline{C} | \mathbf{E})$$

(where  $P(C | \mathbf{E}) = P(C)$  by assumption), thus avoiding the problem of double-counting the effect of C.

---

<sup>16</sup>Also see [Hanks 1990a].

## 8 Assumptions

The assessment algorithm embodies several assumptions both stated and unstated. They divide naturally into three groups: those concerning the temporal model, those concerning the causal rules, and those concerning the probabilistic model and its related calculations. Here we will summarize those assumptions and discuss the advisability and difficulty of relaxing them.

We made two assumptions regarding the temporal model:

- Instantaneous events.  
First we made a distinction between event propositions and fact propositions. Event types were then assumed to occur instantaneously (i.e. at a single point in time). We will relax this assumption in the projection algorithm, thus allowing the agent's actions to consume variable amounts of time. Relaxing the assumption is not problematic, except that the rules' syntax would have to be extended to indicate whether the consequent was realized when the event begins, or when the event ends, or at some indeterminate point during the event's execution. The third option makes the next assumption more problematic.
- Temporal separation known precisely.  
Although the temporal database manager is well equipped to handle imprecision in the time elapsed between time points, we assumed that the time separation among observations, and between the observations and the point of query, were known precisely. This assumption may be violated if we don't know exactly when an event's effects are realized. Relaxing this assumption turns out to be quite problematic. Imprecision in the temporal database is a problem only when it extends to imprecision in the ordering of events. If we are uncertain about the ordering of events we may at worst have to consider all possible total orderings in order to compute probabilities correctly.<sup>17</sup>

Next we made two assumptions about the form of causal rules:

- Rule effects realized immediately.  
We assume that if a rule fires its consequent becomes true at the next instant in time. Relaxing this assumption to allow a delay is no problem except to the extent that the delay's length is imprecisely specified, once again bringing up the problem of imprecision in the ordering of events.
- Events can't cause events.  
A rule's consequent can currently be a fact type but not an event type, which precludes causal chains— an exogenous event triggering a rule, which then triggers other rules subsequently. This assumption allowed us to compute  $\varphi$ 's probability at  $t$  taking into account the world's state prior to  $t$  only.

All of these assumptions are common in the work on causal reasoning within a logical framework (e.g. [Shoham 1988]). Finally we made the following assumptions about the structure of the probability space:

- Sensor reliability not systematically predictable.  
We assumed that an observation's reliability parameter is not correlated with other states of the world, in particular that it is independent of the assessed proposition  $\varphi$  and also of the preconditions of all rules used in assessing  $\varphi$ . This assumption is crucial to our algorithm's behavior, but seems to be an intuitively reasonable one. If a sensor's reliability is correlated with other propositions involved in the assessment we can no longer guarantee a reliability bound, below which an observation has negligible evidential impact. Given the right dependencies, even a very unreliable observation can have an arbitrarily high impact.
- Rule preconditions are short-lived.  
This assumption indicates that a rule's precondition propositions could be expected to change state between potential firings of the rule, and allows us to consider rule firings in isolation. We noted that

---

<sup>17</sup>See, for example, [Dean and Boddy 1988].

rules representing an agent’s intended actions, strung together in a plan, would typically violate this assumption, and this violation forces us to engage in explicit reasoning by cases. We confront this problem in [Hanks and McDermott 1992].

- Events occur dispassionately and predictably.

We assumed that the probability with which an event occurs does not depend on the propositions that figure in the rule for which that event is a trigger. This assumption allowed us to bound the probability that a rule would fire. Obviously there are situations involving adversarial and cooperative behavior that will violate this assumption. We suspect, however, that probabilistic reasoning of the sort we have presented herein will be inappropriate in those situations anyway. Game theory and minimax reasoning provide a better formal framework.

- Evidence flows in a forward direction only.

We mentioned in Section 3.6 the fact that our algorithm did not perform the usual Bayesian propagation of evidence forward and backward through its graph. This difference amounts to saying that we are doing temporal projection only, and manifests itself as the assumption that the evidence brought to bear in assessing  $P(\varphi_t)$  involves only events that occur temporally prior to  $t$ . Our algorithm thus “walks forward in time” to compute probabilities. The assumption means, for example, that one could simultaneously believe that  $\varphi$  was very likely true at  $t$ , and that a very reliable sensor reported  $\varphi$  false at  $t^>$ . We are currently exploring both the conceptual and the computational consequences of relaxing this assumption. Note, however, that the assumption is quite reasonable when applied to the problem of plan projection.

The last group of assumptions mainly allow us to avoid explicit reasoning by cases, as we defined it in Section 7.2—introducing hypothetical courses of events and reasoning about each in parallel. These assumptions allow us to summarize the computation’s state, at each stage of the probability calculation, using a single number  $p$  which is passed from one stage to the next (as in Figure 5, for example).

## 9 Related Work

Our survey of related work starts with the logical approaches to reasoning about change, then touches on several probabilistic approaches: Markov processes, expert-systems frameworks, a rule-based approach, and a similar approach to approximate reasoning with probabilities.

### 9.1 Logical approaches

We will not attempt to summarize here the literature on logical reasoning about change—[Shoham and Goyal 1988] provides a good starting point. Our formulation most closely resembles Shoham’s [1988] *causal theories*, to which we have added probabilistic information in various places. Our algorithm for computing probability given a *fixed* set of information corresponds to Shoham’s definition of a *chronologically minimal model*.<sup>18</sup>

### 9.2 Markov processes

Before discussing alternative proposals in the AI literature we should mention briefly the literature on Markov processes, e.g. [Howard 1971]. The idea is to represent the domain by a state space, in which every state represents a complete static description of the world (or that portion of it under study). Change is described by a *transition matrix*: for a domain with  $n$  possible states a transition matrix is an  $n \times n$  matrix whose  $(i, j)^{th}$  entry contains the probability that the system will be in state  $s_j$  at time  $t + \delta$  given that it was in

---

<sup>18</sup>In fact it corresponds exactly to chronological minimality in the case that (1) all priors are either 0 or 1, (2) all rules have an effectiveness parameter of 1, (3) all observations have reliability 1, and (4) the probability of an unobserved state change over an interval (the  $\Delta^\varphi$  parameter) is 0.

state  $s_i$  at time  $t$ . We can then compute a probability distribution over possible states at time  $n$  by (1) providing a probability distribution over world states at the initial time  $t_0$ , and (2) multiplying this vector  $n$  times by the transition matrix.

While the representational power of these systems is the same as for our symbolic graphical models (see [Nunez 1989], for example), the problem with the standard Markov-process representation is that it is inconvenient. We are generally more comfortable building symbolic state descriptions than defining the system's state space as a whole, and can more easily define our causal theory using symbolic rules of change than using numeric transition matrices. Our rule-based representation is simply more appropriate for the larger task of maintaining a world model that will be used for plan generation, refinement, and repair.

### 9.3 Expert-system approaches

[Cooper *et al.* 1988] confronts the problem of probabilistic diagnostic reasoning, in particular determining the probability that a patient has a certain disease given that he has displayed a certain set of observable symptoms (evidence). The temporal dimension enters in that the system processes probabilities like the likelihood that a disease persists to the present given that it began at some time  $t_x$  and the likelihood that a particular profile of symptoms would be observed between  $t_x$  and the present given that the disease began at that time.

Apart from the explicit mention of time and probabilities, there is little similarity between this work and our own. Our work is oriented primarily toward two goals: (1) to use an internal causal model (rules and event probabilities) to predict future states of the system, and (2) to limit inference in a manner appropriate to the particular problem by instantiating only relevant parts of the model.

Cooper's system addresses neither of these issues. In the first place, the probabilities relate diseases and symptoms (evidence) directly, and diseases are assumed to be mutually exclusive. Therefore there is no way to reason about a scenario in which disease  $d_1$  begins at time  $t_1$ , eventually causes  $d_2$  at some later time, and so on. In fact Assumption 7 explicitly rules out this sort of reasoning by saying essentially that the probability that disease  $d_i$  begins at  $t_x$  does *not* depend on the patient's state prior to time  $t_x$ .

Decisions as to what evidence to consider are made by the system designer ahead of time rather than by the system itself dynamically. The system builder fixes a time granularity and a time horizon. The system then considers all evidence occurring within the horizon. There is no explicit model of noisy evidence (in fact Assumption 5 states that the values of all evidential variables are known at all times), and there is no internal causal model corresponding to our rules, therefore no way for the system to ignore evidence that is too unreliable or too far in the past, nor is there any way to abandon causal chains as their inferences become too tenuous.

Berzuini's work—[Berzuini and Quaglini 1989], [Berzuini 1990]—likewise provides a discussion of how an expert might build a graphical model that explicitly describes the time course of the system. The (human) expert describes the domain's causal structure by providing the model's state space and transition probabilities directly.

An interesting aspect of Berzuini's work is his adoption of a state-based temporal model. He suggests a continuous-time semi-Markov process to represent situations in which temporal information (when events occur, how long do propositions remain true) is imprecise, but does not provide a method for building such a model automatically. A promising area for future research is to establish the relationship between these semi-Markov models and symbolic rule-based theories in which a rule's consequent has an indeterminate delay and an event's time of occurrence is uncertain.

### 9.4 Rule-based approaches

[Dean and Kanazawa 1988] presents a probabilistic model essentially identical to the one that we built in Section 3.6: they too take the idea that causal rules are triggered by events and effect changes in facts. Their model has no explicit model of noisy observations, but it could easily be added.

Where the work differs from ours is in the method for building the full network. Dean and Kanazawa take the “time-based” approach: they choose a time increment  $\delta$  and instantiate the model at each time point,  $t_0, t_0 + \delta, t_0 + 2\delta, \dots$ . They suggest using approximate simulation methods to solve the network, but provide no insight into the time required to guarantee convergence.<sup>19</sup> [Tatman and Shachter 1990] propose an alternative approach that involves applying dynamic-programming techniques.

There are thus two fundamental differences between the work of Dean and Kanazawa and our own: they rely on a time-based approach (see Section 3.6), requiring them to choose an appropriate time mesh and instantiate the model at all time points, whereas we try to instantiate the model only at those time points at which the system changes significantly. Second, they propose a forward simulation solution technique that does not take into account the specific inferences required by the problem solver. The simulation solves the entire network (approximately), computing the probability of every proposition at every point in time. We try to reason only about those parts of the network that are of interest to the problem solver and only to the degree of accuracy it requires.

## 9.5 Belief updating by assessing impact

Weber’s work ([Weber 1989a], [Weber 1989b]) is a computational theory of how to form beliefs based on statistical information. The idea is that estimating the likelihood of some proposition (say whether my car will start) suggests the problem of what information is relevant and significant to the prediction. One may have statistical information about how the car starts under a variety of circumstances: when it is cold, wet, cold and wet, parked on the street as opposed to in the garage, on weekdays as opposed to weekends, in the summer as opposed to in the winter, and so on. Some of these factors will be irrelevant, some will be redundant; as a computational matter it may be infeasible to consider all possible information.

Suppose we are estimating the likelihood of a hypothesis  $h$ , and have currently considered a set of evidence  $r$ . Now we consider whether to incorporate a new proposition  $f$  into the conditioning set—i.e. our new estimate of  $h$  will be based on  $r \cup f$  instead of on  $r$ . Weber defines a quantity called the *impact* of  $f$  on  $h$  given  $r$ , which is related to the ratio of the relative frequency of  $f$  given  $h \cap r$  to the relative frequency of  $f$  given  $\bar{h} \cap r$ . High positive values mean that  $f$  offers strong positive evidence for  $h$  in the context of  $r$ ; high negative values indicate strong evidence that  $h$  is false, and impact of 0 means that incorporating  $f$  will not change the estimate.

Weber’s algorithm takes as input a proposition and a number  $a$  called *accuracy*, and performs the following loop:

1. Initial estimate is  $h$ ’s prior; initial  $r$  is  $\emptyset$ .
2. For all evidence  $f$  not in  $r$ , compute the impact of  $f$  on  $h$  given  $r$ .
3.  $r \leftarrow r \cup f^*$  where  $f^*$  is that single piece of evidence with the highest impact (in absolute value).
4. Recompute estimate on the basis of new evidence set  $r$
5. If  $|\text{impact}(f^*)| < a$  then terminate and return current estimate.

This algorithm is similar in spirit to our own (although it is not applied explicitly to temporally qualified propositions) in that it tries to consider only that set of evidence relevant to a particular probabilistic assessment task. He also adopts the heuristic of considering only a single piece of evidence in deciding when the algorithm should terminate.

One difference is that Weber uses an accuracy value instead of a threshold. Accuracy  $a$  defines an interval  $[e - a, e + a]$  where  $e$  is the estimate returned by the algorithm,  $a$  is the input accuracy value, and the interval is guaranteed (or very likely) to contain the exact probability<sup>20</sup>. Deciding degree of belief with

<sup>19</sup>Kanazawa [1991] in subsequent work explores the use of state-based approaches for temporal reasoning, but does not address the problem of building the graphical network from a set of rules.

<sup>20</sup>Things are a little more complicated than that because the updates are defined in terms of logarithms.

respect to a threshold  $\tau$  amounts to asking which of the intervals  $[-\infty, \tau]$  and  $[\tau, \infty]$  the exact probability lies in. [Haddawy and Hanks 1990] shows a decision-theoretic justification for queries of this form. One advantage of the threshold approach is that it allows us to assess the impact of certain pieces of information extremely quickly: if the current estimate is over the threshold we can ignore all positive evidence, even *strongly* positive evidence. Weber’s algorithm would continue incorporating positive evidence, even though it would only continue to confirm the current hypothesis.

[Weber 1989b] also provides classes of “statistical causal rules” that correspond to the causal information we code in causal rules, observations, and background probabilities. He does not go on to explore whether his highest-impact-first algorithm would provide an effective inference mechanism for doing predictive inference on these theories (although he expresses pessimism in Chapter 8).

## 10 Conclusion

An agent’s model of the world has to take into account the world’s propensity to change and the agent’s ignorance about its state. Doing so involves integrating a theory of change with a theory of uncertainty. In this paper we developed an algorithm for computing the *projected probability* of a proposition  $\varphi$  at a point in time  $t$  based on a body of evidence known as a probabilistic causal theory. A causal theory consists of:

1. sensory observations—reports from some external sensor about  $\varphi$ ’s state at some point in time,
2. rules governing change—statements that certain conditions in the world, namely the occurrence of a *trigger event* in the presence of a certain *precondition*, will change  $\varphi$ ’s state,
3. probabilities quantifying confidence in the model—expressing the likelihood that the changes predicted by the rules above are sufficient to predict *all* changes in  $\varphi$ ’s state.

This system represents an advance over deterministic theories of change ([McDermott 1982], [McCarthy 1987], [Shoham 1988]) both due to its greater expressive power and due to its ability to limit its inference according to the demands of the task at hand.

The gain in expressive power has to do with enriching the notion of “persistence” or “inertia” captured by the deterministic systems, to include the following notions:

- that if we reason about the state of  $\varphi$  over a long interval of time, we may be reasonably sure that  $\varphi$ ’s state changes during the interval, even though we have no *explicit* evidence that it has changed,
- that our causal model of  $\varphi$  may be incomplete, failing to predict the occurrence of causally relevant events, failing to express correctly the rules that govern  $\varphi$ ’s changes over time, or failing to capture the exact circumstances under which a particular known rule will actually effect a change.

Our system for probabilistic projection captures all these concepts, using the following information:

1. the known ways in which  $\varphi$ ’s state can change (the causal rules)
2. the extent to which each rule correctly captures a particular causal relationship (the rule’s effectiveness parameter)
3. the extent to which we believe that our causal model is complete (the  $\Delta^\varphi$  parameter).

In addition to this increased expressive power—the ability to quantify the perceived accuracy of our causal model—the probabilistic approach affords us the advantage of limiting our inferential effort in those cases in which approximate answers will suffice. Instead of computing  $\varphi$ ’s exact probability at  $t$ , we instead decide only whether that probability is less than, or greater than, an application-supplied threshold. By doing so we can consider only that evidence significant enough to change our estimate with respect to the threshold, ignoring sufficiently unreliable or old sensory evidence, and abandoning tenuous lines of causal inference.



Another way to view the work is that it takes a causal theory that defines a graphical structure over the system's state space, implicitly quantified over all time points. Instead of trying to evaluate the structure at all times the algorithm tries to identify the times at which the system changes state in relevant ways, and instantiates the graph only at those points.

Future work in this area takes three paths:

1. Projecting plans

As we mentioned, various assumptions we made in this theory tend to be violated regularly when the events encountered comprise a deliberate plan. Relaxing these assumptions involves more reasoning by cases, which tends to be computationally expensive. Additionally, planning involves reasoning about a richer collection of phenomena than can be conveniently modeled using our restriction of rule consequents to binary random variables. We report on this work in [Hanks and McDermott 1992].

2. Relaxing structural assumptions

Relaxing the “projection” assumption—that all evidence bearing on  $\varphi$ 's state at  $t$  occurs prior to  $t$ —complicates the probability computation, but makes the algorithm amenable to a wider variety of tasks, such as failure diagnosis and replanning. Adding the idea of “induced causality” (rules causing other events to occur, thus generating causal chains) and of “domain constraints” (allowing states of propositions to interact by constraints other than the causal rules) makes the system amenable to reasoning about complex physical systems.

3. Incorporating temporal imprecision

We have assumed that the temporal distance between any two points is known precisely, but temporal information can be incomplete and inaccurate in exactly the same sense as can propositional information. We need to integrate probabilistic information about time, e.g. the amount of time an action will take, into our model. Conceptually there is no difficulty; the trick will be to do so in a manner that preserves fast computation.

4. Planning, replanning, and learning

We have yet to integrate the model manager into a planning system. Interesting questions arise as to how a planner might use decision-theoretic methods to guide the process of plan selection [Haddawy and Hanks 1992], and how that process affects the interface between planner and world model. Then when the plans are actually executed we confront the problem of how to deal with failure—failure of the world to conform to the agent's model, and thus failure of the plan to achieve its intended effects. Replanning is the problem of how subsequent actions should change to cope with the failure; learning involves how the agent's world model—probabilities, rules, and plan library—should be modified to prevent subsequent failures.

## References

- [Berzuini and Quaglini 1989] Riccardo Bellazzi Berzuini, Carlo and Silvana Quaglini. Temporal Reasoning with Probabilities. In *Proceedings of the Uncertainty in Artificial Intelligence Workshop*, 1989.
- [Berzuini 1990] Carlo Berzuini. Representing Time in Causal Probabilistic Networks. In Max Henrion, editor, *Uncertainty in Artificial Intelligence 5*, pages 15–28. Elsevier North Holland, 1990.
- [Cooper *et al.* 1988] Gregory F. Cooper, Eric J. Horvitz, and David E. Heckerman. A Method for Temporal Probabilistic Reasoning. Technical Report KSL 88-30, Knowledge Systems Laboratory, Stanford University, July 1988. Working paper.
- [de Kleer and Williams 1987] Johan de Kleer and Brian Williams. Diagnosing Multiple Faults. *Artificial Intelligence*, 32(1):97–130, 1987.

- [Dean and Boddy 1988] Thomas Dean and Mark Boddy. Reasoning about Partially Ordered Events. *Artificial Intelligence*, 36(3):375–400, 1988.
- [Dean and Kanazawa 1988] Thomas Dean and Keiji Kanazawa. Probabilistic Temporal Reasoning. In *Proceedings AAAI*, 1988.
- [Dean and Kanazawa 1989] Thomas Dean and Keiji Kanazawa. A Model for Reasoning about Persistence and Causation. *Computational Intelligence*, 5:142–150, 1989.
- [Dean 1985] Thomas Dean. Temporal Imagery: An Approach to Reasoning about Time for Planning and Problem Solving. Technical Report 433, Yale University, Department of Computer Science, October 1985. Ph.D. thesis.
- [Haddawy and Hanks 1990] Peter Haddawy and Steve Hanks. Issues in Decision-Theoretic Planning: Symbolic Goals and Numeric Utilities. In *Workshop on Innovative Approaches to Planning, Scheduling, and Control*, November 1990.
- [Haddawy and Hanks 1992] Peter Haddawy and Steve Hanks. Representations for Decision-Theoretic Planning: Utility Functions for Deadline Goals. In *Third International Conference on Knowledge Representation and Reasoning*, October 1992.
- [Hanks and Firby 1990] Steve Hanks and R. James Firby. Issues and Architectures for Planning and Execution. In *Workshop on Innovative Approaches to Planning, Scheduling, and Control*, November 1990.
- [Hanks and McDermott 1987] Steven Hanks and Drew McDermott. Nonmonotonic Logic and Temporal Projection. *Artificial Intelligence*, 33(3):379–412, 1987.
- [Hanks and McDermott 1992] Steven Hanks and Drew McDermott. Modeling a Dynamic and Uncertain World II: Projecting Courses of Action. *Artificial Intelligence*, 1992. Submitted.
- [Hanks 1987] Steven Hanks. Temporal Reasoning About Uncertain Worlds. In *Proceedings of the Uncertainty in Artificial Intelligence Workshop*, pages 114–122, 1987.
- [Hanks 1990a] Steven Hanks. Practical Temporal Projection. In *Proceedings AAAI*, 1990.
- [Hanks 1990b] Steven Hanks. Projecting Plans for Uncertain Worlds. Technical Report 756, Yale University, Department of Computer Science, January 1990. Ph.D. thesis.
- [Howard 1971] Ronald A. Howard. *Dynamic Probabilistic Systems—Volume I: Markov Models*. John Wiley & Sons, 1971.
- [Hunter 1989] D. Hunter. Causality and Maximum Entropy Updating. *International Journal of Approximate Reasoning*, 1989.
- [Jaynes 1979] E.T. Jaynes. Where Do We Stand on Maximum Entropy? In R.D Levine and M. Tribus, editors, *The Maximum Entropy Formalism*. MIT Press, 1979.
- [Kanazawa 1991] Keiji Kanazawa. A Logic and Time Nets for Probabilistic Inference. In *Proceedings AAAI*, 1991.
- [Lifschitz 1987] Vladimir Lifschitz. Formal Theories of Action. In Frank Brown, editor, *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*. Morgan-Kaufmann, 1987.
- [McCarthy 1987] John McCarthy. Applications of Circumscription to Formalizing Common Sense Knowledge. In Matthew L. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*, pages 153–166. Morgan-Kaufmann, 1987.

- [McDermott 1982] Drew McDermott. A Temporal Logic for Reasoning About Processes and Plans. *Cognitive Science*, 6:101–155, 1982.
- [Nunez 1989] Linda I. Mensinger Nunez. Relationship Between Temporal Bayes Networks and Markov Random Process Transition Tables. Technical Report CS-89-M2, Brown University Department of Computer Science, 1989.
- [Pearl 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan-Kaufmann, 1988.
- [Shoham and Goyal 1988] Yoav Shoham and Nita Goyal. Temporal Reasoning in Artificial Intelligence. In Howard Shrobe, editor, *Exploring Artificial Intelligence*. Morgan Kaufmann, 1988.
- [Shoham 1988] Yoav Shoham. *Reasoning about Change and Time from the Standpoint of Artificial Intelligence*. MIT Press, 1988.
- [Tatman and Shachter 1990] Joseph A. Tatman and Ross D. Shachter. Dynamic Programming and Influence Diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):365–379, March 1990.
- [Weber 1989a] Jay Weber. A Parallel Algorithm for Statistical Belief Refinement and its use in Causal Reasoning. In *Proceedings IJCAI*, pages 900–905, 1989.
- [Weber 1989b] Jay Weber. Principles and Algorithms for Causal Reasoning under Uncertainty. Technical Report 287, University of Rochester, Department of Computer Science, May 1989.
- [Wellman 1988] Michael P. Wellman. Formulation of Tradeoffs in Planning under Uncertainty. Technical Report MIT/LCS/TR-427, MIT Laboratory for Computer Science, August 1988. Ph.D. thesis.