

Causal Logic

Glenn Shafer
Faculty of Management
Rutgers University
gshafer@andromeda.rutgers.edu

August 28, 1998

Abstract

This paper proposes a logic for causal based on event trees. Event trees provide a natural and familiar framework for probability and decision theory, but they lack the modularity that would make them convenient models for a logic. Our logic overcomes this deficiency by elaborating the idea of an event tree to that of an event space. This is a space containing situations or instantaneous events, which are related by relations of precedence, refinement, and implication.

1 Introduction

Planning, prediction, diagnosis, and many other tasks in artificial intelligence require a language in which actions and other causes can be represented. Ideally, this language should be a genuine logic, with a workable inference mechanism and a clear semantics. It could then be called a causal logic. In this paper, we join in the current effort to construct causal logics. The most important part of our contribution is a semantics that integrates action, temporal structure, and probability.

Our new semantics—our temporal ontology, as we will call it—is based on the intuition of branching continuous time, but the states that evolve in time, the *situations*, are not fully detailed states of the world. Indeed, the fundamental novelty of the ontology is its modest concept of a situation. Previous temporal ontologies have been formulated in terms of states that fully specify everything about the world, including a metric time. In our ontology, a situation never tells everything about the world. It can be refined so as to tell more, but we consider more refined situations only as the need arises. Probabilities for how a given aspect of a situation will evolve typically depend on certain other aspects of the situation but not on all details that might be added to a refinement of the situation. An action is simply a transition from one situation to a second situation, and hence we can speak also of refining an action; one action is a

refinement of another if it has the same initial situation and a more refined final situation.

Mathematically, our temporal ontology begins with a set of objects called situations. The set is called an *event space*. The situations in it are related by certain binary relations and certain constructions, which satisfy certain axioms. The binary relations include refinement (situation T is more refined than situation S), precedence (situation T is possible in situation S and can only occur after S), and necessity (in situation S, situation T is inevitable). The axioms are derived from a rigorous analysis of the simple and well-understood idea of an event tree, and hence they constitute a clear and rigorous mathematization of the idea of branching time [17].

The event-space ontology has a number of advantages over more familiar mathematical structures for branching time:

1. Assumptions of persistence can be expressed both probabilistically and in terms of preconditions and excluded actions. Thus both probabilistic decision-theory methods and nonmonotonic reasoning methods can be made available within a single language. Within this language, nonmonotonic relations are refinements of probabilistic relations.
2. Since a situation is not necessarily identified with a point on a metric time scale, the problem of the divided instant [21] does not arise.
3. The situations that form the branching time structure can also be thought of as instantaneous events: the situation is the same as the instantaneous event that brings it into existence. Thus a language based on the ontology can also function as a logic of events.
4. Since situations and hence actions can have refinements within the ontology, we can represent both a nondeterministic action (tossing a coin, for example) and a more refined deterministic action (tossing the coin and getting heads, for example).
5. The binary relations in the ontology, which can be reproduced in a language based on it, include relations of necessity and possibility between situations or instantaneous events. A statement that one event is possible or inevitable (or “caused,” as some authors want to say) after another thus has a direct meaning in the ontology.
6. The temporal relations within an event space include several relations that different authors have labeled “cause” [20, 9, 10]. By using all these relations and avoiding labelling any one of them “cause”, we avoid the sense of mystery and metaphysics associated with this word. Our causal relations are specific and down-to-earth, saying something the actual world and how its evolution is best predicted. To use the words of David Lewis [8], these causal relations are supervenient on the bare facts about what happens.

Since our ontology is essentially a mathematical structure for branching time, we can construct a language that uses it in much the same way that established temporal logics use simpler discrete branching-time structures. Such a language can include names for situations, for actions, for individuals who may or may not exist in a certain situation, and for properties of these individuals that may or may not be specified in a certain situation. (In some cases, it may be necessary to refine a given situation in order to get values for a certain property.) Shafer [16] proposes a language similar to CLT* ([2]) except that instead of two kinds of proposition formulae (for propositions that are true or false in situations and propositions that are true or false of paths, respectively), one has proposition formulae (to represent propositions that are true, false, or indefinite in situations) and event formulae (to represent situations themselves, thought of as instantaneous events). Due to constraints of space, we omit event formulae in the language sketched here, but we do consider formulae that represent actions.

The ramifications of actions can be expressed in two ways in a language based on event spaces. If the ramifications are consequences of constraints among properties that hold in all situations, then they can be expressed directly by such constraints. If they are consequences of conditions of necessity between successive events, then they can be expressed by conditions on event formulae. Nonmonotonic reasoning based on presumptions of persistence of properties not explicitly said to be affected by actions can be carried out by minimizing changes in mentioned properties, as in [20, 9, 10], but this does not require any predicate with the name “cause.” In an event space, there a number of different binary relations that might be named “cause”, but it would be confusing and misleading to bestow the name on any one of them [17].

We believe that our ontology and languages built on it provide an advance that can be used within all the current efforts to found causal reasoning on a branching-time picture, including efforts that already use a mathematical representation of branching time explicitly in their semantics (this is the tradition of temporal logic; see [22]) and those that use a sorted first-order logic and build the branching-time structure syntactically, by specifying permitted sequences of actions (this is the tradition of the situation calculus [11, 4, 12, 7], and also [1]).

Although our logic and ontology have much in common with the various action logics proposed in the A.I. literature [11, 4, 12, 7, 14, 3, 18, 5], our goal is somewhat more general in that the event space is a representation of the various possible ways that the world can turn out and the meaning of our language is defined with respect to an event space. In other words, we handle both actions that are carried out by an agent and so called exogenous actions within a uniform framework.

Interestingly, in [11], some ideas very close to our approach were developed. Borrowing from the work of Prior [13], they use a fluent $F(\pi)$ to make assertions about causality. For example $F(\pi, s)$ asserts that the situation s will be followed (eventually) by a situation that satisfies the ordinary fluent π . The meaning is that π holds in some situation s' in the future of s . Using this notation one can

assert that if a person is out in the rain, he will get wet by:

$$\forall x, p, \text{ Raining}(x) \wedge \text{At}(p, x) \wedge \text{Outside}(p) \rightarrow F(\text{Wet}(p))$$

using a notation in which the situation argument is suppressed. McCarthy and Hayes also define some related special fluents¹.

2 Event Trees and Event Spaces

Event and probability trees have been used since the seventeenth century. An event tree is simply a probability tree without the probabilities. Recently [16, 17, 15] these ideas have been further refined mathematically and developed into a foundation for probability and causality. Event trees show the different ways events may unfold.

As an illustration of a probability tree, consider Figure 1. which is a proba-

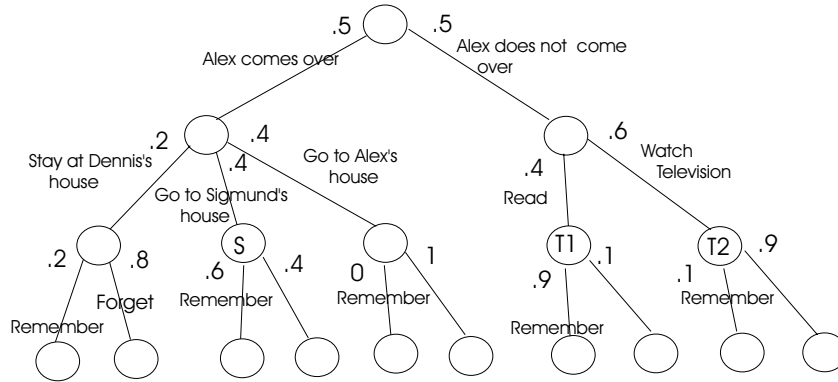


Figure 1: Will Dennis remember to practice his saxophone?

bility tree for whether Dennis, a twelve-year old boy, will remember to practice his saxophone before dinner on a summer afternoon. He is least likely to do so when his friend Alex comes to his house and the two boys then go to Alex's house. He is most likely to do so when he stays at home by himself and reads. The probabilities and contingencies in the tree are in reference to an idealized observer who watches the events as they unfold. At each point the observer's limited ability to predict what will happen next is indicated by the probabilities attached to the branches at that point. Here it is natural to say that watching television contributed to (or caused) Dennis's not practicing. Notice that the nodes (circles) in the tree represent situations. Individuals (such as Dennis) have different properties (such as location) in different situations. The situations can also be thought of as instantaneous events; the node labelled S in Figure 1 can

¹ These include $G(\pi, s)$ which means that π holds in all situations s' in the future of s . Also $Next(\pi, s)$ denotes the situation s' in the future of s in which π holds.

be thought of as the situation where Dennis has just arrived at Sigmund's house or as the event of his arrival.

Event and probability trees often need to be refined. Figure 2 illustrates this point. There the simpler tree on the right agrees with the probabilities given by the more refined tree on the left; both give the same initial probability for Rick's eventually reading ($.8 \times .8 \times .625 = .4$) and for his eventually watching television ($.2 + .8 \times .2 + .8 \times .8 \times .375 = .6$). Every situation in the simpler tree is also represented in the more refined tree. The situation represented by the node labelled R on the right is also represented by a node labelled R on the left. The situation represented by the node labelled S on the right is represented on the left by a collection of three nodes, S_1, S_2, S_3 . For a more detailed and rigorous account of refinement, see Chapter 13 of [15].

In general, as we have just seen, a situation can be represented in an event tree either by a single node or by a *clade* of nodes—a set of nodes that are divergent, in the sense that no two are on the same path down the tree. The same situation represented by a node in one tree is represented by a clade in another. As we introduce more and more detail, the situation will be represented by a larger and larger clade.

In an *event space*, we consider situations in abstraction from their representation in a particular event tree. This allows us to consider on the same footing situations that are related in a variety of ways. Two situations in an event space might be related by precedence (there is some event tree in which one precedes the other along a path down the tree). And two other situations might be related by refinement (there is some event tree where one is represented by a subset of the clade that represents the other).

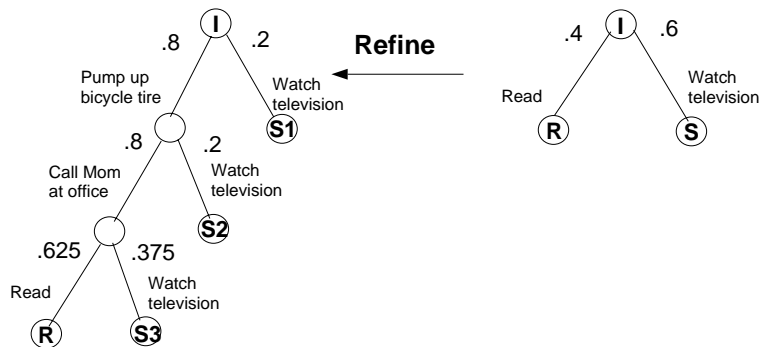


Figure 2: Refinement

Situations or instantaneous events can also be related in much more complicated ways. Consider, for example, Figure 3. This is an event tree for a shopper who sets out to buy wine, cheese, and apples. The green grocer sells all three, the cheese store sells cheese and wine, and the wine store sells only wine. The shopper has time to visit two stores at the most. The clade labelled **T** represents the instantaneous event that the shopper buys apples, while the clade

consisting of the solid nodes (which we may call **S**) represents the instantaneous event that he buys wine and cheese in the same store. How do we explain the relations between these two instantaneous events?

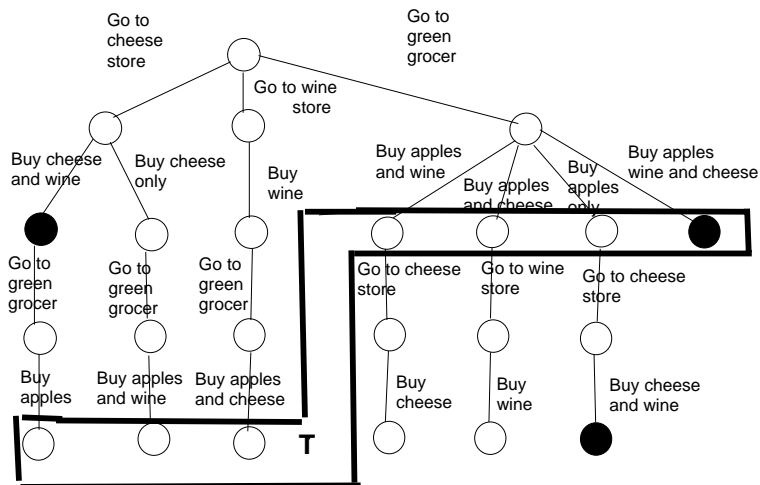


Figure 3: An event tree for a shopper

In order to answer this question, it is helpful first to think about how a single node S^* might be related to a clade **T**. Figure 4 reveals that there are exactly five ways S^* and **T** can be related.

1. The single node S^* is inside **T**. In this case, we say that the situation S^* *refines* the situation **T**, or that whenever the instantaneous event S^* happens, the instantaneous event **T** happens as well.
2. The single node S^* is below **T**. In this case, we say that S^* *requires* **T**, or that S^* can only happen if **T** has already happened.
3. The single node S^* is above **T**. In this case, we say that S^* *foretells* **T**, or that if S^* happens, then **T** will certainly happen.
4. The single node S^* is partially above **T**. Some, but not all, branches emanating from S^* lead to a situation in **T**. We say that S^* *forebears* **T**, or that if S^* happens, then **T** may happen.
5. The single node S^* is outside of **T**. No branches emanating from S^* lead to a situation in **T**. We say that S^* *diverges* **T**, or that if S^* happens, then **T** will definitely not happen.

Returning now to Figure 3, we see that in general the relation between two situations **T** and **S** can be understood in terms of how many and which of the relations in Figure 4 hold between **T** and parts of **S**. In the case of the shopping story, there are parts of **S** in relations 1, 2, and 3. There are no parts

S 1 T	S overlaps T (S overlaps T)	S may happen in such a way that T happens at the same time.
S 2 T	S may.Require T (S may strictly require T)	S may happen in such a way that T has already happened.
S 3 T	S may.Foretell T (S may strictly foretell T)	S may happen in such a way that T must happen later.
S 4 T	S may.forbear T (S may strictly forbear T)	S may happen in such a way that T remains only possible.
S 5 T	S may.diverge T (S may diverge from T)	S may happen in such a way that T cannot happen.

Table 1: The Five Primitive Relations

in relations 4 or 5. We can express this in terms of the five binary relations in Table 1: **S overlaps T** holds, **S may.Require T** holds, **S may.Foretell T** holds, **S may.forbear T** does not hold, and **S may.diverge T** does not hold.




Figure 4: Five ways a single node S^* can relate to a clade T

Formally, an event space is a set \mathcal{S} and five primitive binary relations with the names and intuitive meanings listed in Table 1. Certain axioms must be satisfied and certain constructions must be allowed. For lack of space, we refer the reader to [16, 17] for these axioms and constructions.

The five primitive binary relations can be combined in many possible ways to yield further relations. For example, the disjunction

$$\mathbf{E1} \quad S 1 T \text{ or } S 2 T \text{ or } S 3 T \text{ or } S 4 T$$

specifies that S and T are *compatible*, meaning that events can turn out in such

S \neg 2345 T	S refines T (S refines T)	Whenever S happens, T happens at the same time.
S \neg 1345 T	S Requires T (S strictly requires T)	Whenever S happens, T has already happened.
S \neg 1245 T	S Foretells T (S strictly foretells T)	Whenever S happens, T must happen later.
S \neg 1235 T	S forbears T (S forbears T)	Whenever S happens, T remains only possible.
S \neg 1234 T	S diverges T (S diverges from T)	Whenever S happens, T cannot happen.
S \neg 345 T	S requires T (S requires T)	Whenever S happens, T happens at the same time or has already happened.
S \neg 245 T	S foretells T (S foretells T)	Whenever S happens, T happens at the same time or must happen later.
S \neg 25 T	S allows T (S allows T)	Whenever S happens, T happens at the same time, must happen later, or remains only possible.

Table 2: The Negative Relations

a way that they both happen. The negative version

E2 neither S 1 T nor S 2 T nor S 3 T nor S 4 T

states that S and T are *divergent*, meaning that if one happens the other does not.

We will use a concatenation of numerals to indicate a disjunction and a negation sign preceding the concatenated numerals to indicate the negation of the disjunction. Thus **E1** is abbreviated as S 1234 T, and **E2** is abbreviated as S \neg 1234 T.

Table 2 lists a number of relations defined in terms of the primitive relations. The first five strengthen the corresponding relations in Table 1. For example, S 1 T says that S may happen in such a way that T happens at the same time, while S \neg 2345 T says that this is the only way S may happen. In terms of these relations, we define S **precedes** T to mean S **allows** T and T **requires** S. When S **precedes** T, we also say that T follows S.

We saw in Figure 1 how probabilities can be added to an event tree, and we saw in Figure 2 how these probabilities can be elaborated as the event tree is refined. How do we think about probability when we abstract to an event space? This question is answered in part by Shafer [15]. In general, we cannot

always speak, in an arbitrary situation in an event space, of the probability of a possible future event. This is because that probability may depend on details that are only resolved in refinements of the situation. In Figure 1, for example, we cannot speak of the probability that Dennis will practice in the situation represented by the clade \mathbf{T}_1 , \mathbf{T}_2 because this probability depends on whether we are in \mathbf{T}_1 or \mathbf{T}_2 .

When probability is understood causally (as when we say that a certain action will produce a certain event with a certain probability), it requires a certain stability (in frequencies in similar situations). Since it is implausible to assume such stability always holds, we do not suppose that every event has a well-defined probability, no matter how refined our situation. But even when not all events have probabilities, many events or variables may have well-defined prices, or at least lower prices. This can be made rigorous (see Chapter 12 of [15]), and hence a language based on event spaces can legitimately use lower probability statements, such as “P is necessary to degree .75.”

3 Modeling

A number of issues that have come up in the A.I. literature are readily modeled in the event space framework. This section considers two of these, in particular non-deterministic actions and also the qualification problem.

Event spaces easily represent non-deterministic actions. This is illustrated in Figure 5. On the far left is a representation of the action of flipping a coin. This action takes us from situation \mathbf{I} to situation \mathbf{S} at a level of detail in which the outcome of heads or tails is undefined. The situation \mathbf{S} can be refined into the situations $\mathbf{S1}$ and $\mathbf{S2}$ in which the two possible outcomes are represented. This diagram can be further refined to the one on the far right where the action *flip* takes us from situation \mathbf{I} to \mathbf{T} , a situation where the coin is in the air and can still land heads or tails. Then there are two possible events of the coin ending up as heads or as tails.




Figure 5: Non-Deterministic Actions: Flipping a Coin

The handling of an instance of the qualification problem is illustrated in Figure 6. Following [19], the qualification problem is the problem of “trading off the amount of knowledge that is required in order to make inferences on the

one and, and the accuracy of those inferences on the other hand.” One wants to be able to predict that after shooting a loaded gun there will be a loud noise, but without having to verify that the gun has a firing pin, the bullets are not made out of marshmallows etc. At a more abstract level the concepts of there

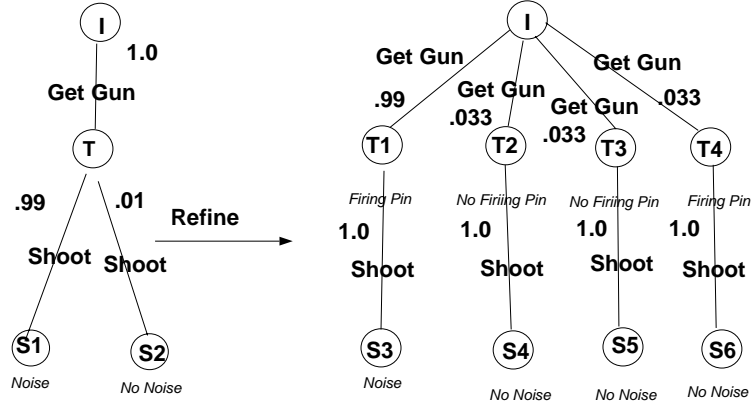


Figure 6: Qualification Problem: Shooting a Gun

being marshmallow bullets, no firing pin etc. are not defined. At this level we can only predict that there is a .99 probability that the loud noise will be heard after the gun is fired. This level of refinement supports reasoning that does not consider guns without firing pins or bullets made of marshmallows. But at a more refined level these predicates are defined and need to be considered.

4 The Language

In this section, we sketch a language \mathcal{CAL} for talking about actions and their effects in the context of event spaces. This language is based on the language \mathcal{A} [3, 6]. Its well formed formulae represent either actions or propositions. Some interesting formulae and their intuitive meanings are given below. The symbols P, Q etc. can represent either an atomic symbol or an arbitrary formula.

1. *Poss a if P_1, \dots, P_n .* This is equivalent to the action precondition axioms of [14].
2. *initially P .*
3. *P after a_1, \dots, a_n*
4. *a makes true P If P_1, \dots, P_n* (where P, P_1, \dots, P_n are predicates and a is an action symbol) This is equivalent to the (positive) effect axioms of [14].

5. *a makes_false P If P_1, \dots, P_n* (where P, P_1, \dots, P_n are predicates and a is an action symbol) This is equivalent to the (negative) effect axioms of [14].
6. *Necessarily_Eventually Q If P_1, \dots, P_n* This is McCarthy's [11] fluent F . A situation that satisfies the fluents P_1, \dots, P_n will eventually be followed by one that satisfies Q .
7. *Possibly_Eventually Q if P_1, \dots, P_n* . For every situation \mathbf{S} satisfying $P_1 \dots P_n$, there exists a \mathbf{T} such that \mathbf{T} may follow \mathbf{S} and \mathbf{T} satisfies Q .
8. *Necessarily_Always Q if $P_1 \dots P_n$* . For every situation \mathbf{S} that satisfies $P_1 \dots P_n$, for every \mathbf{T} such that \mathbf{T} follows \mathbf{S} , \mathbf{T} satisfies Q .
9. *Possibly_Always Q if P_1, \dots, P_n* . For every situation \mathbf{S} satisfying P_1, \dots, P_n and every situation \mathbf{T} that is inevitable after \mathbf{S} , there is a refinement \mathbf{T}' of \mathbf{T} that satisfies Q .
10. *Q only if P first*. In every situation that satisfies P , Q can only be true if P was true first.
11. *a_1 allows a_2 if P_1, \dots, P_n* . For every situation \mathbf{S} satisfying $P_1 \dots P_n$, a_1 is possible and a_2 is possible in the situation resulting from a_1 .
12. *IF a_1 HAPPENS, THEN a_2 HAPPENS*. If a_1 happens then a_2 must happen on every later path.
13. *a_1 HAS_AS_PREREQUISITE a_2* . If a_1 happens then a_2 happened earlier.

Further expressions are possible as well.

5 Semantics

A model \mathbf{M} for our causal action language \mathcal{CAL} is a pair $\langle \mathcal{S}, I \rangle$, where \mathcal{S} and I are as follows:

- \mathcal{S} is an event space with appropriate relations and constructions.
- I is the interpretation function, defined on the action terms and propositional symbols in \mathcal{CAL} . The function I takes two arguments, a situation S from \mathcal{S} and an element e that is either an action term or propositional symbol of \mathcal{CAL} , as in $I_S(e)$.
- There is a distinguished root situation Ω .
- If e is a propositional symbol, then $I_S(e)$ is either TRUE, FALSE, or unspecified. Thus the logic is multivalued. Note that in a particular situation S , we allow $P \vee Q$ to be true while both P and Q are unspecified as long as either P or Q is true in every situation of the clade resulting at some level of refinement of S . This is important in properly handling

nondeterministic actions. (If a situation S refines into S_1 and S_2 , we say that $S = S_1 \oplus S_2$.)

If a propositional symbol has a truth value at s_2 and s_1 **refines** s_2 then that symbol must have the same truth value at s_1 .

If A is an action formula, then $I_S(A)$ is a pair $\langle S, T \rangle$ where S **precedes** T . The function $I_S(A)$ is partially defined. When $I_S(A)$ is defined we say that A is possible at S .

Given a model \mathbf{M} for the language \mathcal{CAL} , a situation S , the semantic value of basic expressions, indicated by $\llbracket \varphi \rrbracket^{\mathbf{M}, S}$, are given below:

1. $\llbracket \neg \varphi \rrbracket^{\mathbf{M}, S} = TRUE$ iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = FALSE$,
 $= FALSE$ iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = TRUE$,
 $=$ unspecified otherwise.
2. $\llbracket \varphi \wedge \psi \rrbracket^{\mathbf{M}, S} = TRUE$
iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = TRUE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = TRUE$
 $= FALSE$ iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = FALSE$ or
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = FALSE$ or
 $S = S_1 \oplus S_2$ and
 $\llbracket \varphi \rrbracket^{\mathbf{M}, S_1} = FALSE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S_2} = FALSE$
 $=$ unspecified otherwise.
3. $\llbracket \varphi \vee \psi \rrbracket^{\mathbf{M}, S} = FALSE$
iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = TRUE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = TRUE$
 $= TRUE$ iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = TRUE$ or
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = TRUE$ or
 $S = S_1 \oplus S_2$ and
 $\llbracket \varphi \rrbracket^{\mathbf{M}, S_1} = TRUE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S_2} = TRUE$
 $=$ unspecified otherwise.
4. $\llbracket \varphi \rightarrow \psi \rrbracket^{\mathbf{M}, S} = FALSE$
iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = TRUE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = FALSE$
 $= TRUE$ iff $\llbracket \varphi \rrbracket^{\mathbf{M}, S} = FALSE$ or
 $\llbracket \psi \rrbracket^{\mathbf{M}, S} = TRUE$ or
 $S = S_1 \oplus S_2$ and
 $\llbracket \varphi \rrbracket^{\mathbf{M}, S_1} = FALSE$ and
 $\llbracket \psi \rrbracket^{\mathbf{M}, S_2} = TRUE$
 $=$ unspecified otherwise.

The semantic value of the expressions of \mathcal{CAL} listed in the preceding section, are defined ² as follows:

1. $\llbracket \text{Poss a if } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$
 $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, then there is some $T \in \mathcal{S}$ such that $I_S(a) = \langle S, T \rangle$.
2. $\llbracket \text{initially } P \rrbracket^{\mathbf{M}} = \text{TRUE}$ if $\llbracket P_1 \rrbracket^{\mathbf{M}, \Omega} = \text{TRUE}$.
3. $\llbracket P \text{ after } a_1, \dots, a_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if there is some $S_1 \in \mathcal{S}$ such $I_\Omega(a_1) = \langle \Omega, S_1 \rangle$, and ... $I_{S_{n-1}}(a_n) = \langle S_{n-1}, S_n \rangle$, and $\llbracket P \rrbracket^{\mathbf{M}, S_n} = \text{TRUE}$.
4. $\llbracket a \text{ makes_true } P \text{ If } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that $I_S(a) = \langle S, T \rangle$ and $\llbracket P \rrbracket^{\mathbf{M}, T} = \text{TRUE}$.
5. $\llbracket a \text{ makes_false } P \text{ If } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that $I_S(a) = \langle S, T \rangle$ and $\llbracket P \rrbracket^{\mathbf{M}, T} = \text{FALSE}$.
6. $\llbracket \text{Necessarily_Eventually } Q \text{ If } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that S **foretells** T and $\llbracket Q \rrbracket^{\mathbf{M}, T} = \text{TRUE}$.
7. $\llbracket \text{Possibly_Eventually } Q \text{ if } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that S **allows** T and $\llbracket Q \rrbracket^{\mathbf{M}, T} = \text{TRUE}$.
8. $\llbracket \text{Necessarily_Always } Q \text{ if } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, then for every $T \in \mathcal{S}$ such that S **precedes** T , $\llbracket Q \rrbracket^{\mathbf{M}, T} = \text{TRUE}$.
9. $\llbracket \text{Possibly_Always } Q \text{ if } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and ... and ... $\llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, then for all $T \in \mathcal{S}$ such that S **precedes** T and S **foretells** T , there is some $T' \in \mathcal{S}$ such that T **refines** T' and $\llbracket Q \rrbracket^{\mathbf{M}, T'} = \text{TRUE}$.
10. $\llbracket Q \text{ onlyif } P \text{ first} \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, and $\llbracket Q \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that S **requires** T and $\llbracket P \rrbracket^{\mathbf{M}, T} = \text{TRUE}$

²Due to lack of space only the TRUE truth value is given.

11. $\llbracket a_1 \text{ allows } a_2 \text{ if } P_1, \dots, P_n \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $\llbracket P_1 \rrbracket^{\mathbf{M}, S} = \text{TRUE}$ and \dots and $\dots \llbracket P_n \rrbracket^{\mathbf{M}, S} = \text{TRUE}$, there is some $T \in \mathcal{S}$ such that $I_S(a_1) = \langle S, T \rangle$, and there is some T' such that $I_T(a_2) = \langle T, T' \rangle$.
12. $\llbracket \text{IF } a_1 \text{ HAPPENS, THEN } a_2 \text{ HAPPENS} \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $I_S(a_1) = \langle S, T \rangle$, there is some $T', T'' \in \mathcal{S}$ such that $I_{T'}(a_2) = \langle T', T'' \rangle$ and T **Foretells** T' .
13. $\llbracket a_1 \text{ HAS_AS_PREREQUISITE } a_2 \rrbracket^{\mathbf{M}} = \text{TRUE}$ if for all $S \in \mathcal{S}$ such that $I_S(a_1) = \langle S, T \rangle$, there is some $S', S'' \in \mathcal{S}$ such that $I_{S''}(a_2) = \langle S'', S' \rangle$ and S **Requires** S'' .

Other expressions of \mathcal{CAL} can be defined in a similar fashion.

6 Conclusions and Future Work

In this paper, we have proposed an ontology that integrates action, temporal structure, and probability. The ontology is a mathematical structure called an event space. The space is inhabited by situations or instantaneous events with binary relations over them. The temporal relations within an event space include several relations that different authors have labeled “cause.”

We have also proposed a logic to talk about event spaces. It is a logic of action and causality that has much in common with the various action logics proposed in the A.I. literature [11, 4, 12, 7, 14, 3, 18, 5]. But note that its primary goal is different in that the event space is a representation of the various possible ways that the world can turn out. Furthermore, our language (unlike for example [20, 9, 10]) does not single out any one of them under the name “causality” and therefore avoids any appearance of reliance on some mysterious relation of causality that goes beyond the facts about the world and about how its evolution is best predicted.

A number of further topics are subjects of our ongoing research.

- Currently our logic does not make use of the probabilities found in the event space. To cover the full range of “causal” concepts and relations, the language needs to be extended to make use of this information.
- The logic is a propositional logic. Variables and quantifiers need to be added to the language and individuals and relations over them are to be added to the semantics.
- Inference rules and also automatic methods for reasoning with the logic need to be developed.
- The addition of clock time (discrete or continuous) to the language will raise a number of issues. Additionally, it may be useful to further extend the language to introduce symbols for situations.

With these additions the logic and the ontology will have much more in common with the various logics of action that have been developed within the A.I. literature.

References

- [1] Baral, C.; Gelfond, M.; and Proveti, A. 1997. Representing actions: Laws, observations and hypotheses. *Journal of Logic Programming*. to appear.
- [2] Emerson, E. 1990. Temporal and modal logic. In van Leeuwen, J., ed., *Handbook of Theoretical Computer Science*, volume B. Amsterdam: Elsevier. 995–1072.
- [3] Gelfond, M., and Lifschitz, V. 1992. Describing action and change by logic programs. In *Proceedings, Joint International Conference and Symposium on Logic Programming*.
- [4] Gelfond, M.; Lifschitz, V.; and Rabinov, A. 1996. What are the limitations of the situation calculus? In *Automated Reasoning: Essays in Honor of Woody Bledsoe*. Dordrecht: Kluwer Academic Publishers. 167–179.
- [5] Haddawy, P. 1996. A logic of time, chance, and action for representing plans. *Artificial Intelligence* 80(2):243–308.
- [6] Kartha, G. N. 1993. Soundness and completeness theorems for three formalizations of action. In *Proceedings, Thirteenth International Joint Conference on Artificial Intelligence*, 724–729.
- [7] Levesque, H.; Reiter, R.; Lespérance, Y.; Lin, F.; and Scherl, R. B. 1997. Golog: A logic programming language for dynamic domains. *Journal of Logic Programming*.
- [8] Lewis, D. 1986. *Philosophical Papers*, volume II. New York: Oxford University Press.
- [9] Lin, F. 1995. Embracing causality in specifying the indirect effects of actions. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, 1985–1991.
- [10] McCain, N., and Turner, H. 1997. Causal theories of action and change. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*.
- [11] McCarthy, J., and Hayes, P. 1969. Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B., and Michie, D., eds., *Machine Intelligence 4*. Edinburgh, UK: Edinburgh University Press. 463–502.
- [12] Pinto, J. 1994. *Temporal Reasoning in the Situation Calculus*. Ph.D. Dissertation, Department of Computer Science, University of Toronto, Toronto, Ontario. Available as technical report KRR-TR-94-1.

- [13] Prior, A. 1968. *Past, present and future*. Oxford: Clarendon Press.
- [14] Reiter, R. 1991. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Lifschitz, V., ed., *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. San Diego, CA: Academic Press. 359–380.
- [15] Shafer, G. 1996. *The Art of Causal Conjecture*. Cambridge: MIT Press.
- [16] Shafer, G. 1997a. *Causality*. In preparation.
- [17] Shafer, G. 1997b. Mathematical foundations for probability and causality. In *Mathematical Aspects of Artificial Intelligence*. Providence, Rhode Island: American Mathematical Society.
- [18] Shanahan, M. 1995. A circumscriptive calculus of events. *Artificial Intelligence* 77(2):249–284.
- [19] Shoham, Y. 1988. *Reasoning About Change*. Cambridge, MA: MIT Press.
- [20] Stein, L., and Morgenstern, L. 1994. Motivated action theory: a formal theory of causal reasoning. *Artificial Intelligence* 71:1–42.
- [21] van Benthem, J. 1983. *The Logic of Time*. Dordrecht: Kluwer Academic.
- [22] Vila, L. 1994. A survey on temporal reasoning in artificial intelligence. *AICOM (Artificial Intelligence Communications)* 7(1):4–28.