

Temporal Notation and Causal Terminology

Yoav Shoham and Thomas Dean
Yale University
Department of Computer Science

ABSTRACT

We argue that causal reasoning is an essential part of intelligent human behavior, and that discussion of it cannot be divorced from discussion of temporal reasoning. We therefore set out to define causation in three stages. In the first, we present an ontology of time. We then outline a theory of “causal conditionals”, which allows one to reason about multiple possible courses of events. Finally, we define causation in terms of *direct causation* and *causal origins*.

1 Introduction

Philosophers have long disputed the relative merits of one causal theory over another. Some even question whether or not the notion is a useful one at all, suggesting that causality is simply an anachronism that “scientific” man would well be rid of.

- Nothing exists from whose nature some effect does not follow. (Spinoza)
- I assert that nothing ever comes to pass without a cause. (Jonathan Edwards)
- Causality is *to us* the cement of the universe. (David Hume)
- The Law of Causality, I believe, like much that passes among philosophers, is a relic of a bygone age, surviving like the monarchy, only it is erroneously supposed to do no harm ... All philosophers, of every school, imagine that causation is one of the fundamental axioms of science, yet, oddly enough, in advanced science ... the word “cause” never occurs. (Bertrand Russell)

Whether or not causality is respectable in the scientific literature, it is clear that people employ causal terminology in their day to day communication and, we will assume, in their understanding of the world around them. Consider the following excerpts from recent issues of Newsweek magazine (the italics are ours).

- March 4, in an article on the dangers of DDT, the chemical is described as “a pesticide that *silenced* birds, *threatened* West Coast peregrine falcons and bald eagles, and possibly *caused* cancer in humans”.
- March 11, on Page 3, it reads: “Newly disclosed evidence has *reignited* the controversy over Bernard Hugo Goetz ... mounting furor seemed to be *leading towards* a new grand jury investigation ...”
- In the same issue, in a letter to the editor beginning on page 4, one reads: “... tough laws *will not stop* the use of dangerous drugs but will instead *lead to* obscene profits ...”
- and on page 60 an article is titled “A strike cripples Pan Am”.

The italics in the quotations are meant to highlight implicit causal statements. The particular text, *Newsweek* magazine, was not carefully selected for our purposes. The reader is invited to look in any magazine or paper on his local newsstand to verify the ubiquity of causal terminology: bringing about, provoking, instigating, affecting, preventing, enabling, etcetera. Interestingly this tendency is not as evident in the scientific community. Generally the more rigorous or mathematical the subject the more subtle the causality. However, lest the reader be given the impression that causation is a lowly form of reasoning, he is referred to [Lerner 65] for an account of causation's role in scientific research.

Since causal reasoning appears to be an important component of intelligent human behavior, it stands to reason that researchers in AI should strive to understand it and make use of it in their theories. Indeed, in recent years there has been some interest in causation. Medical diagnosis systems stress the "causal reasoning" component of their systems [Pople 82] [Patil et al. 82]. Theories of "qualitative physics" employ causal notions - see the special issue of the *Journal of Artificial Intelligence* on the topic.

In this paper we propose a general theory of causation. A theory intended to capture our intuitive understanding of the term and at the same time be rigorous in its definition. Before we introduce our theory, however, let us explore some of the properties of causation which makes its definition hard.

In the first example, DDT is said to have "caused" various things, including the silence of birds. What the paragraph in *Newsweek* says is that the silence of the birds (without going into its metaphorical meaning) is *explained* by the presence of DDT. But what is the nature of this explanation? That this explanation is distinct from material implication is obvious. To quote Quine, "Whatever the proper analysis of the contrafactual may be, we may be sure in advance that it cannot be truth-functional" [Quine 59] (actually reproduced from [Barwise 85]). Thus X cannot cause X even though X logically implies X, and thunder does not cause the preceding lightning although it implies it. Causality contains a "direction"; there is no causal analog of the logical equivalence, or "if and only if". Causal explanation embodies some notion of a process: of machinery in action.

The presence of DDT, *Newsweek* statement claims, triggered a physiological process whose outcome was "the silence of birds". But what exactly does "triggering" mean? Consider the later example about the Goetz controversy. The new evidence "reignited" the debate - does that mean that if that evidence had not been revealed that the controversy would necessarily have remained dormant? Of course not; there are many other possible factors that could revive it: Goetz could have confessed to be a member of the Ku Klux Klan, a similar case of a "subway vigilante" could have occurred in the NY subway, and so on. So the statement "A caused B" does not mean that A is a necessary condition of B, that if A had not occurred that B would not have occurred. Similarly, it is not a sufficient condition. The newly disclosed information (namely, that Goetz had said to one of his victims "you don't look too bad, here's another" and shot him a second time) would have passed unnoticed if at the same time the USSR had declared war on the U.S., or alternatively if all media went on strike and the news could not spread. Most causal explanations hide many implicit conditions, usually an infinite number of them. In the excerpt from the letter to the editor it says that the use of dangerous drugs would lead to obscene profits, but that is true only if the potential profiteers did not incur equally large losses, which in turn could happen in any number of ways. To give a more macabre example, John's pulling the trigger caused Mary's death, but that is only because she was not wearing a bullet-proof vest, that the

gun was loaded, that the bullets were made of lead rather than marshmallows, and the list could be continued indefinitely.

There is a great deal of philosophical literature on causation spanning several thousand years. For an overview the reader is referred to [Mackey 74], and for the current views of some of the outstanding contemporary philosophers to [Sosa 75]. Philosophy is not, however, constrained by the need for process models and not surprisingly none have surfaced in the literature. Philosophers are concerned with the world *as it really is*, while we will be satisfied with finding a useful notation for expressing knowledge. With some luck our notation should correspond with concepts employed by humans, but at no time do we expect our theories to be “true” in any sense.

We are concerned with the way humans reason about causality because (a) they are relatively good at it and (b) it appears to be a very hard problem. While causality might profitably be excised from physics, we suspect that the basic causal notions play an important computational role in the way humans reason about the world.

In trying to pin down causation there are three main issues to be settled:

1. **Ontology:** What entities participate in the causal interactions? What is the relation between them and how does one derive new entities from old ones?
2. **Organization:** Given a commitment to what things exist, how is the information about such entities stored and how does the storage arrangement impose constraints on the computational processes that will make use of this information.
3. **Computation:** What processes are required to make use of the stored information.

The three issues are in many ways inextricably intertwined. However, the first can be reasonably tackled in isolation and that is precisely what we intend to do in this paper. Our solution will constrain the organization and suggest a process model but our immediate concern is to settle the ontological questions.

Organization of the paper: In section 2 we state our main thesis about causation, and contrast it with two attempts within AI, one by Schank and one by Simon, to deal with causation in a somewhat general way. In Sections 3, 4 and 5 we outline our theory of time and causation, and in Section 6 we step back and attempt some perspective on what we presume to have accomplished.

2 Thesis

In [Simon 65] Herbert Simon offers a specific definition of causality.¹ The situation in which the causality is to be determined is represented by a set of linear equations, and the entities participating in the causal relation are the variables appearing in the equations. Certain restrictions apply to the set of equations, and causation is defined essentially by the effect of perturbing one

¹Since in the following we will criticize Professor Simon's formulation, it is only fair to point out that it is an old one, dating back in its original form to the late fifties. We understand from Professor Simon that he has recently done more work on the topic together with a student of his. We have not yet seen that work and our criticism has no bearing on it.

variable on other variables. The relation defined by Simon and called “causal ordering” is not transitive, and there is no temporal information (for example, an effect may precede the cause).

In [Schank 75] Schank identifies four types of causal relations: result, enable, initiation, and reason. In order to define these relations he distinguishes between “actions” and “states”. To reason about causality Schank constructs “causal chains”, and suggests how such chains can be reconstructed given only partial description of them.

Neither of these formulations, we claim, is acceptable. Our thesis is that in order to define causation one must have a well defined theory of time and a representation and means of reasoning about what we will call causal conditionals. In more detail we posit the following.

- A prerequisite for defining causation is having a precise notation for temporal information, which includes precise definitions of terms like facts and events. We will call it an ontology of time. Simon’s representation contains no temporal information, as he himself points out. Schank does pay attention to entities like “action” and “state” and to the sequencing of such entities in time, but their meaning is left intuitive; he offers no precise definition of them.
- Causation cannot, in principle, be defined solely in terms of constraints on the simultaneous truth value of temporal assertions. For example the constraint $F=MA$ does not contain causal information, nor would that information be obtained by the addition of other constraints. This is in contrast to the approach proposed by Simon.
- Instead it is defined in terms of what we will call a theory of causal conditionals. For example, the rule $F=MA$ is replaced by others, one of which may have the rough form of “if you did this to F then that would happen to A , all other things being equal”. Causally conditional statements do not suffer from the inherent symmetry of constraints.

Based on this we will outline a theory of causation in three stages - develop a theory of time, in terms of that a theory of causal conditionals, and finally use the latter to define causation.

3 An ontology of time

In this section we will briefly outline the theory which is described in more detail in [Shoham 85]. The theory and notation are influenced by the work of James Allen [Allen 84] and Drew McDermott [McDermott 82]. In this paper we will ignore most of the technical details and try to convey the main ideas. These details are, however, important, and the interested reader is encouraged to refer [Shoham 85].

The basic statement about time relates an interval of time, denoted by its two endpoints, and an associated proposition. So, for example, we will write $HOLD(t_1, t_2, LOCATION(ball, loc))$ to represent the fact that the ball *ball* is in location *loc* from t_1 to t_2 . The term $LOCATION(ball, loc)$ is a *fact type*, and the tuple $\langle t_1, t_2, LOCATION(ball, loc) \rangle$ is a *fact token*. It is possible, as several people have done, to similarly define *event types* and *event tokens*. We find that this separation is somewhat arbitrary and tends to obscure important distinctions. For example, *The ball rolled* is an fact and *The ball changed location* is a event, though they both denote the same situation. Also, we will want to speak of facts like *The ball rolled AND there was another ball in*

its path, or *The ball did NOT roll*. How do we conjoin two different entities like facts and events? Furthermore, consider the definition of an *action*. Actions are usually defined in terms of an event or fact, associated with other concepts like an actor and an intention. But that intuitively would call for two kinds of actions: “event actions”, like *John went home*, and “fact actions”, like *John stood still*.

Events and facts have enough in common for us to lump both into the generic category facts. We will separate out “event-like” facts from “fact-like” facts by their particular properties.

We classify facts along several dimensions. A fact type can be *holographic*, meaning that if it holds over an interval I then it holds over every subinterval of I. “John is stood still” is an example of a holographic fact. A fact can be *gestalt*, meaning that if it holds over an interval I it does *not* hold over any subinterval of I. For example, “John walked 3 feet south” is gestalt. Independently, a fact type may or may not be *mergeable*. Intuitively, a fact type is mergeable if whenever it holds over two overlapping intervals it holds over their union (in fact, the technical definition is slightly different). Fact types that are both holographic and mergeable are called *liquid*. For example, “John stood still” is liquid.

Along with the classification of fact types we impose a certain structure on fact types, by recasting James Allen’s [Allen 84] definitions into our notation. This structure allows us to achieve the effect of having logical connectives inside the scope of *HOLD*, so that the meaning of $HOLD(t_1, t_2, AND(f_1, f_2))$, $HOLD(t_1, t_2, NOT(f))$, etcetera are well defined. Quantifiers receive similar treatment, so one can speak of fact types of the form $FORALL(x, f(x))$. The particular definitions have interesting properties such as that for any fact type f , $NOT(f)$ is holographic (these are explored in [Shoham 85]).

4 A theory of causal conditionals

To summarize the previous section, all temporal information is expressed by fact tokens, and inference rules which refer to the classification of fact types and allow the deduction of new fact tokens from old ones. This information contains no modality. Each tuple presumably represents an actual fact in the world. We say “the world”, since our notation so far only allows us to describe one state of affairs. The theory of causal conditionals, we said, would represent and reason about information of the rough form “if X occurs then Y will occur, all other things being equal”. We must define what the X and the Y are, and above all spell out the meaning of “all other things being equal”. Here too we will only sketch the main ideas of what we call “potential facts theory”, keeping formalism down to a minimum.

Imagine that our world - past, present and future - is inhabited by a tremendous number of *potential facts* which try to manifest themselves, as it were. For example, consider a billiard ball rolling across the table. There is the potential fact of its rolling, which in this case also manifests itself - call it ROLLING1. Now suppose that we roll another ball that collides with the first rolling one. The laws of physics tell us that at that point the balls will stop rolling in the “expected” direction and begin rolling in a new direction. However we do not forget where each ball *would have rolled* were it not for the collision. In particular we remember ROLLING1, only it is not a fact any longer - it manifested only an initial segment of itself. We call it a *potential fact* - it starts out as a fact but due to interactions with other potential facts it becomes a counterfact. In the first scenario, where the second ball was not rolled, ROLLING1 actually managed to manifest

all of itself. In the second scenario it managed to manifest only the part of it which preceded the collision. If, for example, we rolled yet a third ball which collided with the second ball before the second ball has a chance to collide with the first one, ROLLING1 would again manage to manifest all of itself.

The image to keep in mind then is of a world full of potential facts fighting for existence and very few “real” facts which are selected by some central arbiter - reality is but the tip of an iceberg.

Without going into too much detail, the way we represent potential fact tokens is by replacing the previous 3-tuples $\langle t_1, t_2, \text{fact-type} \rangle$ by 5-tuples $\langle \text{world}, \text{token-name}, t_1, t_2, \text{fact-type} \rangle$. *Token-name* is a unique name associated with each such token, and *world* defines the world to which that fact token belongs. We will return to world definers soon, but first let us look at what these fact tokens mean. They can mean one of four things, depending on whether the token is *potential*, *materialized*, *ghost* or *real*. By convention, the token-name determines which type the token is. For each potential fact token $\langle w, \text{pid}, t_1, t_2, f \rangle$ there exist unique $\langle w, \text{mid}, t_1, t_3, f \rangle$ and $\langle w, \text{gid}, t_3, t_2, f \rangle$ such $\text{mid} = \text{MATERIAL}(\text{pid})$ and $\text{gid} = \text{GHOST}(\text{pid})$. We will say more about real fact tokens later.

For example, consider two billiard balls colliding. We have the two potential fact tokens (for convenience we replace logical functions by English text in the fact type):

$\langle w, \text{rolling1}, 1, 9, \text{“ball1 rolls from loc1 in direction dir1”} \rangle$
 $\langle w, \text{rolling1}, 2, 8, \text{“ball2 rolls from loc2 in direction dir2”} \rangle$

Now suppose we know that the balls will collide at time $t=4$ and location *loc3* and roll in new directions. We have not yet said *how* we know that the balls will collide - we will soon - but here we are just demonstrating the representation. We then also have the following:

$\langle w, \text{MATERIAL}(\text{rolling1}), 1, 4, \text{“ball1 rolls from loc1 in direction dir1”} \rangle$
 $\langle w, \text{MATERIAL}(\text{rolling2}), 2, 4, \text{“ball2 rolls from loc2 in direction dir2”} \rangle$
 $\langle w, \text{GHOST}(\text{rolling1}), 4, 9, \text{“ball1 rolls from loc1 in direction dir1”} \rangle$
 $\langle w, \text{GHOST}(\text{rolling2}), 4, 8, \text{“ball2 rolls from loc2 in direction dir2”} \rangle$
 $\langle w, \text{rolling3}, 4, t_1, \text{“ball1 rolls from loc3 in direction dir3”} \rangle$
 $\langle w, \text{rolling4}, 4, t_2, \text{“ball2 rolls from loc3 in direction dir4”} \rangle$

where the *dir3*, *dir4*, t_1 and t_2 are again determined in an as yet unknown way.

What determines which potential fact tokens exist, and how each is decomposed into its materialized and ghost part? The answer to the first question is what we call *the generating function*, and the answer to the second is what we call *the clipping function*. Together they represent the “physics” of the world, the set of rules that are believed to govern the world.

The generating function is a collection of statements of the form “if X, Y,... materialize then A,B,... potentially hold”. For example, if you (actually) fire a gun loaded with n bullets then from then on it (potentially) has $n - 1$ bullets. Or if you (actually) put a block on the table then from then on it (potentially) is on the table. Notice that the generating function only adds potential facts to existing ones - so where do the “first” potential facts come from? This is where what we called the “world definer” comes into play. A “world definer” is just a set of potential facts: a set of initial conditions so to speak. Together, the world definer and the generating function




Figure 1: The potential fact tokens describing John's double shooting

determine what potential facts exist². For example, in a world defined by the set {“the gun had n bullets as of $t=1$ ”, “John fired the gun from $t=2$ to $t=3$ ”, “John fired the gun from $t=5$ to $t=6$ ”}, and assuming that the only generating rule is the one about gun firing described above, we have exactly potential facts shown in Figure 1.

Notice that the generating function says nothing about there not being n bullets in the gun after it was fired. This is the responsibility of the clipping function. The clipping function is a collection of statements of the form “if Y, Z, \dots ever materialize then the ghost part of X will begin no later than t ”. For this to make sense, the type of the fact token X must be holographic. For example, the clipping function may include the rule “if you fire a loaded gun it no longer has the number of bullets it used to”. Adding this clipping rule to the scenario just described, we get situation described in Figure 2.

This is the basic outline of the theory of causal conditionals. The precise definitions of the generating and clipping functions appear in [Shoham 85] but are too lengthy to include here.

5 Causation

Let's summarize the theory of causal conditionals presented thus far. A world is defined by a set of initial potential facts called the world definer. Given a world, the generating function determines what potential fact tokens exist. For each such potential fact token, the clipping function determines its decomposition into materialized and ghost parts.

5.1 Direct causation

Each potential fact token X that is not in the world definer was derived by the generating function applied to a set of other potential fact tokens. Each token in that set is said to be a *direct cause* of X . Notice that in general a token has several direct causes. Referring back to Figure 2, the potential fact token “there are $n - 2$ bullets in the gun from $t=6$ onwards” has two direct causes:

²This is not strictly true as we shall see; the clipping function will also influence which facts materialize.

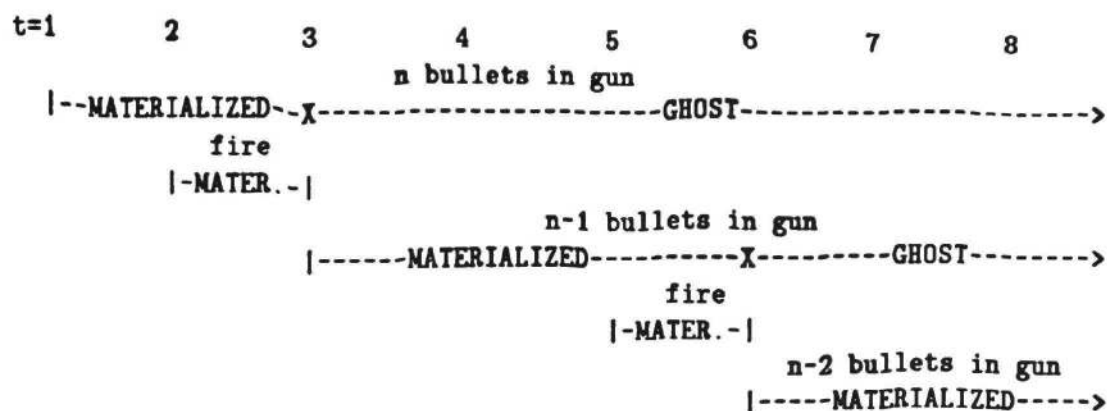


Figure 2: The potential, materialized and ghost fact tokens describing John's double shooting

"there are $n - 1$ bullets in the gun from $t=3$ onwards" and "John fired the gun from $t=5$ to $t=6$ ". In turn, "there are $n - 1$ bullets in the gun from $t=3$ onwards" has two direct causes of its own: "there are n bullets in the gun from $t=1$ onwards" and "John fired the gun from $t=1$ to $t=2$ ". However, "there are n bullets in the gun from $t=1$ onwards" has no direct causes.

The direct causes of materialized fact tokens are defined to be simply the direct causes of the potential tokens to which they belong. So in the previous example, the materialized fact token "there are $n - 1$ bullets in the gun from $t=3$ to $t=6$ " has two direct causes: "there are n bullets in the gun from $t=1$ onwards" and "John fired the gun from $t=1$ to $t=2$ ".

Definition of the direct cause of ghost fact tokens is slightly more complex and we will not give it here. We will only make a note of the fact that "causing a ghost fact token X " is not the same as "preventing X ". If " X prevented Y " then if X had not occurred then Y would have, but we have already seen that " X caused Y not to occur" does not mean that if X had not occurred then Y would have.

Finally, we define the direct causes of real fact tokens. Real fact tokens were mentioned earlier on but not explained. We will explain them through an example. Consider a world in which John and Bill both shot Mary simultaneously. Assume that there is one generating rule, stating that when the trigger of a loaded gun is pulled then the gun has one less bullet and the shot person is dead. There is also one clipping rule stating that under the same conditions the gun no longer has as many bullets as it did and that the person no longer lives.

In this case there are two materialized tokens asserting Mary's being dead over the same interval. However, to us, observers of reality, there is only one death going on. We therefore speak of real fact tokens, which are the result of projecting materialized fact tokens onto reality, so to speak. In particular there is exactly one real fact token asserting Mary's death over the interval. In general, $\langle w, rid, t_1, t_2, f \rangle$ is a real fact token exactly when there exists (at least one) materialized fact token $\langle w, mid, t_1, t_2, f \rangle$. In ordinary speech, when we speak about the causes of X , X is a real fact token. When we speak about the causes of Mary's death, there is exactly one death we are referring to. The fact that in our notation two potential deaths materialized is transparent.

For each real fact token there is the set of materialized fact tokens that gave rise to it, namely

all the materialized fact tokens with which it shares the temporal extent and fact type (and of course the world definer). In the last example, the real fact token asserting Mary's death has two materialized tokens that gave rise to it, all the rest have exactly one. We now define the set of direct causes of a real fact token X to be the set of direct causes of all the materialized fact tokens that gave rise to X . Notice that this is a set of sets. In the same example the direct cause of Mary's death is the set $\{\{\text{"Bill's gun was loaded"}, \text{"Bill pulled the trigger of his gun"}\}, \{\text{"John's gun was loaded"}, \text{"John pulled the trigger of his gun"}\}$.

As we mentioned, the set of direct causes of a real fact token is a set of sets. There are a few special cases. If the set of direct causes of a real fact token Y is $\{\{X\}, S_1, S_2, \dots, S_n\}$ then we will say that X is a *direct disjunctive cause* of Y . If the set of direct causes of a real fact token Y is $\{\{X, Z, \dots, T\}\}$ then we will say that X is a *direct conjunctive cause* of Y . If the set of direct causes of Y is $\{\{X\}\}$ then we will say that X is *the unique direct cause* of Y .

In the next two subsections we introduce broader notions of causation. Unfortunately, due to length limitations, we can do little more than introduce them.

5.2 Indirect causation

We will say that X is an *indirect cause*, or simply a cause, of Y if one of the two holds:

1. X is a direct cause of Y , or
2. Z is a direct cause of Y and X is an (indirect) cause of Z .

This is a generalization of the naive notion of a "causal chain". The latter is a special case of a list of fact tokens each being the unique direct cause of the following one.

5.3 Causal origin

Given our (informal) definition of indirect causation, we now (informally) define the causal origins of fact tokens, "the original sin". The causal origin of a potential fact token are all its indirect causes which are in the world definer. For example, again referring back to Figure 2, the causal origin of "there are $n - 2$ bullets in the gun from $t=3$ onwards" is the set consisting of the three tokens "there are n bullets in the gun from $t=1$ onwards", "John fired the gun from $t=2$ to $t=3$ " and "John fired the gun from $t=5$ to $t=6$ ".

The significance of causal origins derives from the significance of world definers. The latter completely characterize the state of affairs, assuming a given physics (that is, fixed generating and clipping functions). This means that in order to control the world an agent need only control the facts and events described in the world definer. If he wishes to eliminate an undesirable phenomenon he need only eliminate some of its causal origins, or alternatively add some tokens to the world definer that will prevent the phenomenon. Similarly if he notices desirable manifestations he can set about eliminating unnecessary work which does not contribute to bringing about those manifestations.

6 Summary

We have defined what it means for one thing to cause another. In doing so we developed an ontology of time in which potential facts about the world are decomposed into materialized and ghost parts. To capture the actual physics of the world we presented a theory of causal conditionals. This introduced the idea of generating and clipping functions which suggest an organization for our causal knowledge. This organization, we claim, fits well with our intuitions about how people represent and reason about causality.

Bibliography

- [Allen 84] Allen, J.F., *Towards a General Theory of Action and Time*, Artificial Intelligence, 23/2 July (1984), pp. 123-154.
- [Barwise 85] Barwise, Jon, *The Situation in Logic-II: Conditionals and Conditional Information*, Technical Report CSLI-85-21, Stanford, 1985.
- [Lerner 65] Lerner, D. (ed.), *Cause and effect*, Collier-Macmillan, Toronto, 1965.
- [Mackey 74] Mackey J.L., *The Cement of the Universe: a Study of Causation*, Oxford University Press, 1974.
- [McDermott 82] McDermott, D.V., *A Temporal Logic for Reasoning about Processes and Plans*, Cognitive Science, 6 (1982), pp. 101-155.
- [Patil et al. 82] Patil, R.S., Szolovitz, P. and Schwartz, W.B, Causal Understanding of Patient Illness in Patient Diagnosis, *Proc. AAAI-82*, AAAI, 1982.
- [Pople 82] Pople, H. E., Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnosis, Szolovits, P. ed., *Artificial Intelligence in Medicine*, Westview Press, 1982, pages 119-190.
- [Quine 59] Quine, Willard Van Orman, *Methods of Logic*, Holt, Rinehard & Winston, 1959.
- [Schank 75] Schank, R.C., The Structure of Episodes in Memory, Bobrow, D.G. and Collins, A. eds., *Representation and Understanding*, Academic Press, 1975, pages 237-272.
- [Shoham 85] Shoham, Y., Time and Causality from the Standpoint of AI1985. In preparation.
- [Simon 65] Simon, H. A., Causal Ordering and Identifiability, Lerner, D. ed., *Cause and Effect*, Collier-Macmillan, Toronto, 1965.
- [Sosa 75] Sosa, E., *Causation and Conditionals*, Oxford University Press, 1975.