Discovering Causal Relations by Experimentation: Causal Trees

Gregory D. Weber
Computer Science, Indiana University East
2325 Chester Boulevard, Richmond, Indiana 47374-1289, U.S.A.
gdweber@indiana.edu

Abstract

The Controlled Lesion Method (CLM) is a set of principles for inferring the causal structure of deterministic mechanisms by experimentation. CLM formulates an important part of the common-sense logic of causation and experimentation. Previous work showed that CLM could be used to discover the structure of deterministic chains of binary variables more accurately than statistical methods; however, as implemented, CLM was prone to error when applied to causal trees. A change of knowledge representation, replacing atomic symbols with predicate calculus expressions for representing events, makes it possible to refine the statement of one of the principles of CLM. As a result, CLM can now discover causal tree structures correctly. This suggests that a structured representation of events may be necessary for a causal discovery system.

Introduction

Causation appears to be a fundamental and distinctive category of human thought. Aristotle remarked that humans, alone among the animals, have a sense of wonder, which is excited when we do not understand the cause of something (Aristotle 1984, Book I, Ch. 2). Hume emphasized the role of causation in all of our reasoning about matters of fact beyond the immediate data of perception (Hume 1740). People use causal reasoning to explain why something happened, to make predictions, to be able to control events, and to attribute responsibility to themselves and others (Kim 1973). Not surprisingly, causation also has an important role in artificial intelligence. Causal reasoning is used in expert systems, qualitative physics, planning, solving problems of diagnosis, repair, and control, and understanding narrative discourse (Bobrow 1985; Fikes, Hart, & Nilsson 1972; Hammond 1986; van den Broek 1990).

The causal relation is, of course, something more than a statistical association, or even a constant conjunction, between two kinds of events. Reichenbach distinguished two aspects of the causal concept. First, the causal relation has *direction*: it is one thing to say A causes B and quite another to say B causes A. Second, the causal relation has an *order*, inasmuch as one event B can be causally "between" two other events A and C (Reichenbach 1971). This order can be expressed as $(A - B - C)$, and we can know it independently of the direction.

The combination of order and direction makes directed graphs a very useful representation for causal structures (Pearl 1988; Spirtes, Glymour, & Scheines 1993). An edge $A \rightarrow B$ means that A directly causes B. A ***causal chain*** is a system (pattern) of causal relations in which each cause has one direct effect, such as $X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow \ldots \rightarrow X_n$. A ***causal tree*** is a system of causal relations in which some causes have more than one direct effect, for example, $X_2 \leftarrow X_1 \rightarrow X_3$.

How could our robot make such discoveries? Computational methods for causal discovery may be categorized as experimental, passive non-statistical, and statistical.

In *experimental* methods of causal discovery, the agent actively interferes with the environment and observes what happens. Most experimental approaches to causal knowledge are content to identify the effects of operators without attempting to find the order among these effects (Gil 1991; 1993; 1994; Scott & Markovitch 1989; Shen & Simon 1989). May's CDP does discover causal order, but it relies on causal background knowledge and on the assumption that all input factors can be controlled as independent variables, which is not always realistic (Grasshof & May 1995).

In *passive, non-statistical* approaches, the agent attempts to draw causal conclusions from observations of the environment or from a non-causal model of the environment, without the aid of experimentation or statistical analysis. For example, Simon's method of causal ordering infers causal relations from the order in which subsets of a set of equations can be solved (Simon 1952; Iwasaki & Simon 1986). The system of equations must be "structural"; that is, each equation must express a single "mechanism." It is not clear that a system of structural equations is really a non-causal model; but it is true that the method draws causal conclusions which are not explicitly represented in the model. In any case, several heuristic rules are required to enable the modeler to formulate a system of structural equations. Since the heuristics provided (Iwasaki & Simon 1986) require human interpretation, the method is not really suitable for an autonomous robot.

Other passive, non-statistical approaches use the temporal order as their guide. OCCAM (Pazzani 1990; 1991) accepts conceptual dependency inputs representing natural language descriptions of events, and uses a variety of heuristic rules to infer their causal relations. Unfortunately, the most basic rule depends on a fallacious assumption about the relation between time and causation. Shoham's non-monotonic logic of causation (Shoham 1990) involves a similar error. See (Weber 1997) for a detailed critique of inferring causal relations from temporal relations.

Statistical approaches to causal discovery rely on an analysis of the relations of conditional dependence between variables. Algorithms such as Rebane-Pearl (Pearl 1988), Geiger's (Geiger, Paz, & Pearl 1990), SGS (Glymour, Spirtes, & Scheines 1991), IC (Pearl & Verma 1991), and FCI (Spirtes, Glymour, & Scheines 1993) have been remarkably successful, overturning the conventional wisdom that no causes can be inferred without experiment or prior causal knowledge. Curiously, however, these methods work best when there is plenty of random error or "noise." This fact is acknowledged by their authors, but may be easily overlooked by readers. The IC algorithm, one of the more general, is typical in this respect. It takes longer to find the causal structure when there is less noise (Pearl & Verma 1991, p. 450). Its proof of correctness depends on the assumption that the probability distribution is "faithful" or "stable," a property which is lacking if all the probabilities are 0 and 1 (Pearl & Verma 1991, p. 444). Of course, the algorithm might be more robust than its proof of correctness would lead us to believe; and in fact, IC can sometimes discover causal relations correctly even in deterministic (noise-free) systems. But not always. For example, when IC is run with data from the deterministic causal chain $A \rightarrow B \rightarrow C$, it detects no causal relations at all between A, B, and C (Weber 1996)!

Outside of quantum physics, random error is the effect of unobserved causes: if we observe *all* of the causes, there is no random error. If there is no random error, the statistical algorithms may not find the causal relations. Therefore, if we observe all the variables that are causes, using a statistical causal discovery algorithm, we may fail to learn that they are causes. This conclusion can be stated paradoxically as, "*The more we know about the causes* (i.e., the more variables we observe which are, in fact, causes), *the less we will know about the causes* (i.e., the less we will know about which variables are causes)." Human intelligence has the opposite strength: we are more able to discover causal relations in environments which are deterministic or nearly so—so much more able that, until the twentieth century, nearly all philosophers regarded causation as a necessarily deterministic relation (Anscombe 1993). Evidently this human ability has important survival value, and a robot will need to be able to discover causal relations in noise-free, not just noisy, environments.

Causal discovery methods can also be classified according to the way in which they represent the events or variables which are the terms of the causal relation. Many methods represent causes and effects by means of *atomic symbols*, such as A, B, C; others use *structured representations*, such as $F(x)$, $R(a, b)$. For example, Simon has formulated the causal ordering technique sometimes using propositional variables (Simon 1952) and sometimes using real variables (Simon & Rescher 1966); both are represented as atomic symbols. Most or all of the statistical literature on causal discovery deals with variables represented by atomic symbols (Pearl 1988; Spirtes, Glymour, & Scheines 1993; Heckerman & Shachter 1995). In the philosophical literature, too, causes and effects are most commonly represented by atomic symbols, interpreted as facts or events (Mackie 1974; Scriven 1971; Lewis 1973). The widespread use of atomic symbols apparently derives from Hume's insistence that all objects are distinct and that cause and effect are not intrinsically related, for if they were, they would be related logically, not causally (Hume 1740). Some researchers have used structured representations for cause and effect, but they seem to have done so for reasons unrelated to causality. For example, Pazzani uses conceptual dependency to represent events in OCCAM, but he does so because he approaches causation from a natural language perspective and wants a canonical representation of the meanings of sentences. Others (Gil 1991; 1993; 1994; Scott & Markovitch 1989; Shen & Simon 1989). used structured representations because these are more convenient for robotic planners. Planning is conceptually feasible in a limited domain using only atomic variables, but realistic domains have many regularities which are more easily expressed in the language of, say, predicate calculus or production rules. Shoham uses a structured representation for events, which consists of a pair of time points and a propositional variable, but the propositional variable is atomic. Nowhere do we find an argument that the logic of causation or the nature of causality *requires* structured representations of propositions, as a condition of correctness, although Gil has used structured representations to improve efficiency.

Ideally, a causal discovery system for an autonomous robot would satisfy the following criteria:

1. It would be completely automated, not requiring any human intervention.
2. It would not depend on too much background knowledge. Generally, the less background knowledge needed, the better; the robot should be able to start

out with the "mind of an infant" and learn everything it needs.

3. It would not depend on fallacious assumptions about the relation between temporal and causal order.
4. It would work well in deterministic environments.
5. It would work well in indeterministic environments.

Statistical methods satisfy all of the criteria except for the second last. Unfortunately, it doesn't seem possible to modify the statistical algorithms so that they would work reliably in the deterministic case, because their correctness depends critically on the assumption of a "faithful" probability distribution, which is inconsistent with determinism. However, if we had a system which worked well in the deterministic case, it might be possible to generalize it so that it also worked in noisy environments. This could be done, perhaps, by broadening the system's notion of when two variables are associated, from being "always the same value" to "usually the same value."

This paper describes a technique for causal discovery, called the Controlled Lesion Method (CLM), which satisfies all but the last of the above criteria. An earlier version of CLM, based on an atomic (propositional) representation of events, successfully discovered the structures of deterministic causal chains, but it failed when applied to causal trees (Weber 1996). The new version of CLM, based on a predicate calculus representation of events, is also able to discover the structures of causal trees. This paper shows that the change of knowledge representation enabled CLM to be correctly formulated for causal trees. This result suggests that atomic representations of cause and effect may be inadequate for a general causal discovery system.

Principles

CLM depends on two key principles. Both principles are quite simple; indeed they are nothing but formulations of common sense knowledge of causation and experimentation. The *Random Operator Principle* (**R**) is used to infer the direction of causation; it states that operators, by which the experimenter controls variables, are randomly determined, i.e., causation flows from operators to state variables. Thus, one of the few items of background knowledge needed is that the agent must know which of the variables are operators. The operators are *decision variables*; the other variables are *state variables.* The *Lesion Principle* (**L**) is used to infer the order of causation, i.e., undirected causal betweenness relations ($A - B - C$). It is based on the idea that "damaging" parts of the system will disrupt patterns of associations in a characteristic way which reveals causal structure. CLM combines the order information from **L** and the directional information from **R** to infer a directed graph representing the causal structure.

The Random Operator Principle expresses the fact that the experimenter *randomly* decides the values of the independent state variables and chooses appropriate operators to bring these values about. This implies that the operators are causes of the state variables, not vice-versa. Moreover, an operator and a state variable cannot have a common cause, even an unobserved common cause. More precisely, the operator cannot be caused by anything which also causes the state variable by a causal path not leading through the operator (Scriven 1971). For example, if the experimenter decides by tossing a coin, the coin toss cannot affect the mechanism which is being studied except through the operator which is chosen. The outcome of the coin toss is considered *not* to be a state variable, because it has no influence on the mechanism except through the selection of an operator. Formally, the **Random Operator Principle** is:

R: Let X be any state variable. Let O be any operator. Then, X does not cause O.

R eliminates causal structures with arrows into O (such as $X \rightarrow O \rightarrow Y$). By Occam's razor, we may normally eliminate complex structures (multiply-connected and cyclic graphs). Thus, the second piece of background knowledge is the assumption that the causal structure is relatively simple.

For example, consider a lighting system with two bulbs and one switch. It can be described using three propositional variables: A = the switch is up, B = light 1 is shining, C = light 2 is shining. Ignoring C for the moment, we observe A if and only if B. As experimenters, our options include two operators, or patterns of movements: O_1 is, roughly, "move hand to switch, grasp, and pull up"; O_2 is "move hand to switch, grasp, and push down." This pair of operators controls both A and B, but we don't yet know in what way. There are 64 possible graphs with the variables A, B, and the pair (O_1, O_2) as vertices. Applying **R** and Occam's razor rules out all but three of these: (1) $(O_1, O_2) \rightarrow A \rightarrow B$; (2) $(O_1, O_2) \rightarrow B \rightarrow A$; (3) $A \leftarrow (O_1, O_2) \rightarrow B$.

The Lesion Principle allows us to discriminate between these three. The name alludes to Lashley's method of studying the relation between brain and behavior in rats (Lashley 1963); there is a similarity of spirit, though not of detail. If we break or damage light 1, it is clear that whatever was affected by the light's being on will no longer be so affected. Likewise, if we break or damage the switch, whatever was caused by moving the switch will no longer happen. However, it is not always necessary to be violent. If we want to know whether A causes B, we can just do something to the switch that restricts A to one value. For example, we can tape up the switch. If we want to know whether B causes A, we can do something to the light bulb that restricts B to one value; for example, unscrew the bulb.

Using only propositional variables, we can express most of the "non-violent" concept of lesion which has just been developed in

Definition 1: *Lesion* of a variable Y means an alteration of the mechanism which restricts Y to one value (Weber 1996).

The **Lesion Principle** can now be stated:

L: Lesion of Y in a mechanism M disrupts an association $(X \cong Z)$ between variables X and Z if and only if Y lies on the causal path between X and Z in M.

Examples: in the mechanism $X \rightarrow Y \rightarrow Z$, lesion of Y disrupts $(X \cong Z)$; but in $X \rightarrow Z \rightarrow Y$, lesion of Y fails to disrupt $(X \cong Z)$.

In our lighting problem, experimenting with lesions of both A and B enables us to eliminate all but one hypothesis, namely, $A \rightarrow B$. Similarly, experimenting with lesions of B and C would show that neither B nor C causes the other. Table 1 summarizes the inferences we can draw from lesion experiments.

L_A disrupts $(O \cong B)$?	L_B disrupts $(O \cong A)$?	Conclusion
true	*true*	Error
true	*false*	$O \rightarrow A \rightarrow B$
false	*true*	$O \rightarrow B \rightarrow A$
false	*false*	$A \leftarrow O \rightarrow B$

Table 1: Logic of Lesion Studies. A and B are state variables; L_A and L_B represent lesions of A and B, respectively. O represents a pair of operators which control A and B.

The basic step in CLM is to apply the logic of this table (which follows trivially from **L**) to establish the causal structure connecting any similar trio (one operator-pair and two state variables). If there is an extended causal chain or tree, the results for each pair of state variables can be integrated to establish the total causal structure of the mechanism.

Following the prevalent custom in research on causation (Glymour, Spirtes, & Scheines 1991; Pearl & Verma 1991; Simon 1952), the first implementation of CLM, called SCLM-1, used only atomic symbols (such as X and Y) to represent events; it accurately discovered the structure of causal chains but fell into error when applied to causal trees. It appeared that the errors occurred because SCLM-1 would sometimes choose an inappropriate operator for lesion. For example, in the lighting system, it would choose flipping the switch to lesion the variable B (light 1 is shining). Under these circumstances the association between B and C would be disrupted, so SCLM-1 inferred that B caused C. Flipping the switch is not an appropriate way to "lesion" B; in fact, if we define "lesion" correctly, it is not a lesion of B at all. But why not?

Definition 1 failed to capture one important aspect of the concept of lesion. When we tape the switch to lesion A ("the switch is up"), we are doing something to the *subject* of A (the object that A is "about"). When we unscrew the light bulb to lesion B, we are doing something to the subject of B. We are not just doing something that restricts a variable to one value; we are doing so by operating *on the subject* of the variable.

To make this clearer, we need to use the notation of predicate calculus. Let $F(a)$ and $G(b)$ be events. If we want to know if $F(a)$ causes $G(b)$, we can *lesion* a so as to restrict $F(a)$, and we can lesion b so as to restrict $G(b)$. We now redefine *lesion*, replacing Definition 1 with

Definition 2: Let Y be a binary variable representing an event $F(a)$. *Lesion of* Y, written L_Y, means an operation on a, or a condition of a, which restricts Y to one value.

Thus, when SCLM-1 chose flipping the switch as a lesion of B, it was incorrect because flipping the switch was not something done to the *subject* of B, namely, light 1. SCLM-1 was unable to rule out such operators because it did not know what they were about, nor did it know what B was about.

A new implementation, SCLM-2, uses predicate calculus to represent events; for example, instead of B, we have `Shining(light-1)`. The next sections describe experiments which tested whether this change of knowledge representation enabled SCLM-2 to get correct results for causal trees.

Method

The domain of gears turned by a motor provides a convenient platform for experimenting with a variety of causal structures. The motor can directly turn one or more gears, and each gear can turn additional gears. The motor can be switched on and off. Both the motor and the gears can be mounted and dismounted; dismounting a motor or gear disables it from turning, and mounting it enables it to turn again.

In this study, an SCLM-2 agent was given nine problems in the gears-and-motor domain. In each problem it attempted to discover the causal structure of a different simulated environment (or "mechanism"). The operators, observable events, and "laws of nature" for the domain are shown in Figures 1, 2, and 3. The problems differed only in the configuration of gears and motor, as shown in Figure 4. The central causal structure is a chain in problems 1 and 3; in the remaining problems, it is a tree.

For the motor m: `Start!(m), Stop!(m), Mount!(m), Dismount!(m)`
For each gear g: `Mount!(g), Dismount!(g)`

Figure 1: Operators Available to the SCLM-2 Agent. Each operator is expressed as a predicate calculus formula which can be made true or false at the agent's option.

For the motor m: `Mounted(m), Turning(m)`
For each gear g: `Mounted(g), Turning(g)`

Figure 2: Events Observable by the SCLM-2 Agent. Each event is expressed as a predicate calculus formula which can be observed to be true or false. The agent observes all the events at all times.

For the motor m:

1. `if Mount!(m) then Mounted(m) := true`
2. `if Dismount!(m) then Mounted(m) := false`
3. `if Start!(m) and Mounted(m) then Turning(m) := true`
4. `if Stop!(m) then Turning(m) := false`

For each gear g:

1. `if Mount!(g) then Mounted(g) := true`
2. `if Dismount!(g) then Mounted(g) := false`
3. `Turning(g) := Mounted(g)` and `Turning(driver(g))`, where `driver(g)` is the wheel that "drives" g, i.e., the motor, if g is directly connected to it, or the gear that is connected to g and is between it and the motor, otherwise.

Figure 3: Laws of the Gear-and-Motor World.

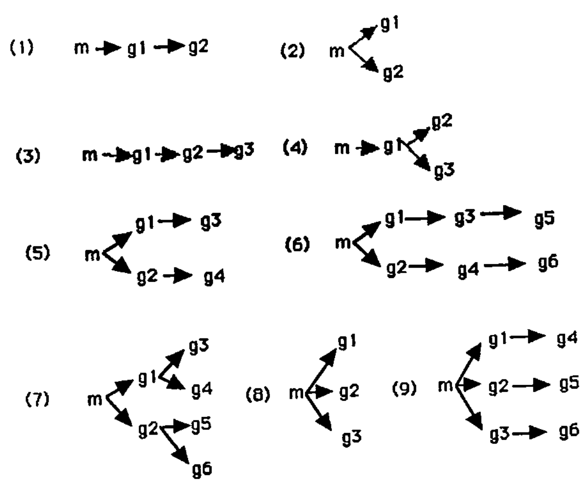

Figure 4: The nine gear-and-motor problems.

The agent was run with 40 trials for each problem. A switch in the program, called *same-subject-required?*, controls whether the agent is able to make use of its predicate calculus representation of knowledge to determine that an operator has the same subject as an observable event (variable), and so that it is a suitable operator for the lesion of that variable. If *same-subject-required?* is *false*, the agent is unable to make use of this knowledge, just as SCLM-1 was unable to know this because events were represented as atomic variables. The value of *same-subject-required?* was set to *true* for 20 trials and *false* for the other 20 trials.

Two kinds of error are possible: the agent can infer a causal relation that does not exist (an *extra edge* in the graph), and it can fail to infer a causal relation that does exist (a *missing edge*). The number of extra edges and the number of missing edges were recorded for each trial. Some edges are more important than others. The central causal structure of each mechanism consists of the causal relations between the `Start(m)`, `Stop(m)`, and `Turning(x)` events; the program will be considered a success if it is able to infer all of these relations correctly, without either any extra or missing edges of this kind. The other causal relations are in some respects peripheral. Motors and gears are normally left mounted, so an observer would not find an association between `Mounted(x)` events and other events unless something strange were going on. If there is no observed association, there is no need of a causal relation to explain it.[1]

Efficiency is also a concern, although secondary to correctness. A CLM agent must design, execute, and analyze lesion experiments, following the logic of Table 1. To avoid repeating the same experiment, the lesion experiments are stored in a table and looked up as needed. The measures of efficiency recorded for each trial were the number of queries of the table, the number of experiments executed, and the CPU time.

Results

The use of predicate calculus, enabling SCLM-2 to require that a proposed lesion operator have the same subject as the variable it is supposed to lesion, resulted in significant improvements both in correctness, when applied to causal trees, and in efficiency.

With *same-subject-required?* = *true*, SCLM-2 inferred no extra edges. In all cases the central causal structure involving `Start`, `Stop`, and `Turning` variables was completely correct. With *same-subject-required?* = *false*, there were extra edges (except for causal chains), and there were errors in the central causal structure. The value of *same-subject-required?*

[1] Moreover, the condition not `Mounted(x)` is the lesion condition for the variable `Turning(x)`. To apply CLM to find the causal relations involving `Mounted(x)` would require finding a lesion for a lesion, and this invites an infinite regress. A rational agent does not have the time or the interest to find the causes of the causes of ... *everything!*

appeared to make no significant difference in the number of missing edges. The numbers of missing and extra edges are shown in Table 2. The actual and inferred structure for a typical causal tree are shown in Figure 5.

With *same-subject-required? = true*, SCLM-2 made fewer queries, executed fewer experiments, and used less CPU time, indicating that this rule is not only required for correctness but provides good search control by eliminating many possible lesion operators from consideration. With *same-subject-required? = false*, the agent made, in some cases, more than 10 times the number of queries, executed more than twice as many experiments, and required more than twice the CPU time. In a real (not simulated) environment, performance would be dominated by the number of experiments performed. The performance measures are shown in Table 3.

Conclusion

The experiments demonstrate that the Controlled Lesion Method, when the knowledge of events is represented in predicate calculus formulas and lesion operators are required to have the same subject as the variables being lesioned, is an effective technique for discovering the structure of causal trees as well as causal chains. It is completely correct in finding the central causal structure, and it makes only errors of omission (missing edges) in the peripheral structure. The "same-subject" test and the knowledge representation to support it are *required* for correct discovery of causal trees, but not required for causal chains. It also enhances the efficiency of search, for both causal trees and chains.

Further research is needed to extend the method to indeterministic environments, non-binary variables, and more general causal structures (multiply-connected graphs).

Acknowledgments

I am grateful to two anonymous referees for their comments. Part of this research was supported by an Indiana University East Summer Faculty Fellowship.

References

Anscombe, G. E. M. 1993. Causality and determination. In Sosa, E., and Tooley, M., eds., *Causation*. Oxford University Press.

Aristotle. 1984. Metaphysics. In Barnes, J., ed., *The Complete Works of Aristotle: The Revised Oxford Translation*. Princeton University Press.

Bobrow, D. G., ed. 1985. *Qualitative Reasoning about Physical Systems*. MIT Press. Reprints the special issue of *Artificial Intelligence* devoted to this subject in 1984 (vol. 24).

Fikes, R. E.; Hart, P. E.; and Nilsson, N. J. 1972. Learning and executing generalized robot plans. *Artificial Intelligence* 3:251–288.

Geiger, D.; Paz, A.; and Pearl, J. 1990. Learning causal trees from dependence information. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 770–776.

Gil, Y. 1991. A domain-independent framework for effective experimentation in planning. In *Proceedings of the Eighth International Workshop on Machine Learning*, 13–17.

Gil, Y. 1993. Efficient domain-independent experimentation. In *Proceedings of the Tenth International Conference on Machine Learning*.

Gil, Y. 1994. Learning by experimentation: Incremental refinement of incomplete planning domains. In *Proceedings of the Eleventh International Conference on Machine Learning*, 87–95.

Glymour, C.; Spirtes, P.; and Scheines, R. 1991. Causal inference. *Erkenntnis* 35:151–189.

Grasshof, G., and May, M. 1995. From historical case studies to systematic methods of discovery. In *Systematic Methods of Scientific Discovery: Papers from the 1995 AAAI Symposium*, 46–57. AAAI Press. Technical Report SS-95-03.

Hammond, K. 1986. CHEF: A model of case-based reasoning. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 267–271.

Heckerman, D., and Shachter, R. 1995. Decision-theoretic foundations for causal reasoning. *Journal of Artificial Intelligence Research* 3:405–430.

Hume, D. 1740. *A Treatise of Human Nature*. Oxford University Press. Reprinted 1888.

Iwasaki, Y., and Simon, H. A. 1986. Causality in device behavior. *Artificial Intelligence* 29:3–32.

Kim, J. 1973. Causes and counterfactuals. *Journal of Philosophy* 70:570–572.

Lashley, K. S. 1963. *Brain mechanisms and intelligence: a quantitative study of injuries to the brain*. Dover Publications.

Lewis, D. 1973. Causation. *Journal of Philosophy* 70:556–567.

Mackie, J. L. 1974. *The Cement of the Universe: A Study of Causation*. Oxford University Press.

Pazzani, M. J. 1990. *Creating a Memory of Causal Relationships*. Lawrence Erlbaum Associates.

Pazzani, M. 1991. A computational theory of learning causal relationships. *Cognitive Science* 15:401–424.

Pearl, J., and Verma, T. S. 1991. A theory of inferred causation. In *Proceedings of KR-91*, 441–452. Morgan Kaufmann.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

	Same-subject-required?	Problem 1	2	3	4	5	6	7	8	9	Total for all problems
Mean extras	true	0	0	0	0	0	0	0	0	0	0
	false	0	0.7	0	0.9	0.7	0.8	2.6	1.3	1.5	8.5
Max. extras	true	0	0	0	0	0	0	0	0	0	0
	false	0	1	0	1	2	3	4	2	4	17
Mean missing	true	4.1	3.9	6.3	6.8	7.5	12.8	11.5	6.0	10.8	69.7
	false	4.6	4.8	4.6	6.4	6.2	10.7	12.9	7.6	12.6	70.4
Min. missing	true	3	3	4	4	5	7	7	4	7	44
	false	3	3	4	4	5	7	8	4	7	45
Max. missing	true	7	7	10	10	13	19	19	10	19	114
	false	7	8	8	11	8	21	22	12	22	119

Table 2: Correctness Comparison. The mean and maximum number of *extra edges*, and the mean, minimum, and maximum number of *missing edges*, are shown for the two conditions *same-subject-required?* = *true* and *same-subject-required?* = *false*, for each of nine problems.

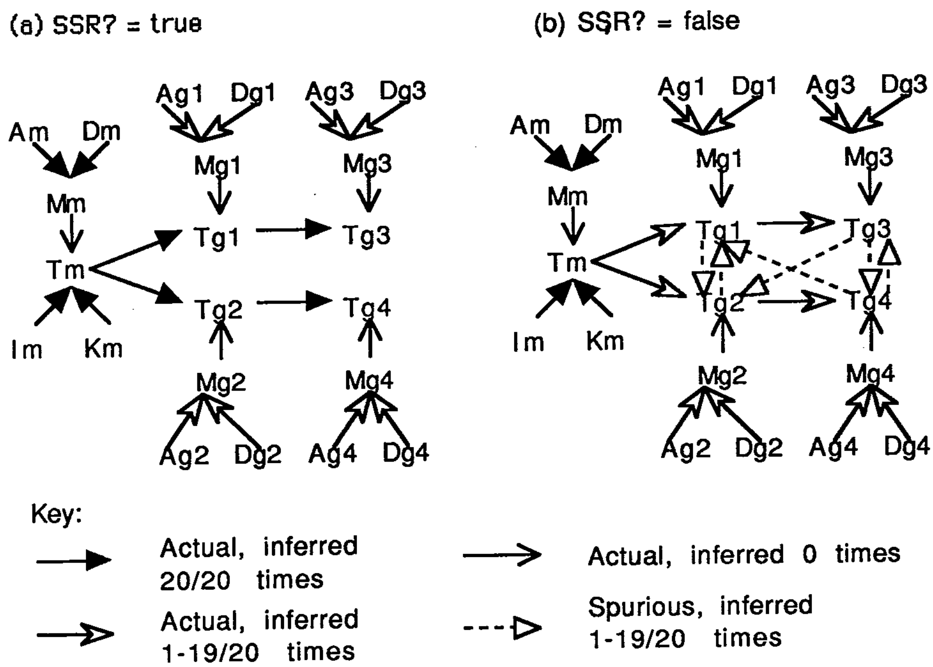

Figure 5: Experimental Results for Problem 5. (a) *same-subject-required?* = *true*; (b) *same-subject-required?* = *false*. Notation: `Im = Start!(m)`, `Km = Stop!(m)`, `Tx = Turning(x)`, `Mx = Mounted(x)`, `Ax = Mount!(x)`, `Dx = Dismount!(x)`

	Same-subject-required?	Problem 1	2	3	4	5	6	7	8	9	Total for all problems
Mean queries	true	33	32	66	66	100	213	209	61	197	977
	false	107	83	221	259	468	2298	2309	273	2377	8395
Mean executions	true	9.1	8.8	12.5	12.9	14.5	23.0	22.2	11.7	19.7	134.4
	false	14.3	13.1	15.0	18.5	18.3	42.3	45.4	22.1	50.7	239.7
Mean CPU sec	true	0.8	0.8	1.9	1.9	3.3	12.1	12.6	1.7	10.6	45.7
	false	1.3	1.1	2.3	2.6	4.5	26.4	27.5	3.3	30.5	99.5

Table 3: Performance Comparison. The mean number of experiments queried, mean number of experiments executed, and mean CPU time are shown for the two values of *same-subject-required?*, for each of nine problems. CPU times are for Chez Scheme 5.0b on a DEC AXP 7720.

Reichenbach, H. 1971. *The Direction of Time.* University of California Press.

Scott, P. D., and Markovitch, S. 1989. Learning novel domains through curiosity and conjecture. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 669–674.

Scriven, M. 1971. The logic of cause. *Theory and Decision* 2:49–66.

Shen, W.-M., and Simon, H. A. 1989. Rule creation and rule learning through environmental exploration. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 675–680.

Shoham, Y. 1990. Nonmonotonic reasoning and causation. *Cognitive Science* 14:213–252.

Simon, H. A., and Rescher, N. 1966. Cause and counterfactual. *Philosophy of Science* 33:323–340.

Simon, H. A. 1952. On the definition of the causal relation. *Journal of Philosophy* 49:517–528.

Spirtes, P.; Glymour, C.; and Scheines, R. 1993. *Causation, Prediction, and Search.* Springer-Verlag.

van den Broek, P. 1990. Causal inferences and the comprehension of narrative texts. In Graesser, A. C., and Bower, G. H., eds., *Psychology of Learning and Motivation*, volume 25. Hillsdale, NJ: Lawrence Erlbaum. 175–196.

Weber, G. D. 1996. Acquiring causal knowledge: An experimental approach. In *Proceedings of the Ninth Florida Artificial Intelligence Research Symposium*, 251–255. Florida AI Research Society.

Weber, G. D. 1997. Causation, similarity-based learning, and time. Forthcoming.