From Imaging and Stochastic Control to a Calculus of Actions

Judea Pearl
Cognitive Systems Laboratory
Computer Science Department
University of California, Los Angeles, CA 90024
judea@cs.ucla.edu

Abstract

This paper highlights relationships among stochastic control theory, Lewis' notion of "imaging", and the representation of actions in AI systems. We show that the language of causal graphs offers a practical solution to the frame problem and its two satellites: the ramification and concurrency problems. Finally, we present a symbolic machinery that admits both probabilistic and causal information and produces probabilistic statements about the effect of actions and the impact of observations.

1 Representing and Revising Probability Functions

Engineers consider the theory of stochastic control as the basic paradigm in the design and analysis of systems operating in uncertain environments. Knowledge in stochastic control theory is represented by a function $P(s)$, which measures the probability assigned to each state s of the world, at any given time. Given $P(s)$, it is possible to calculate the probability of any conceivable event E, by simply summing up $P(s)$ over all states that entail E. The process of revising $P(s)$ in response to new observations is handled by Bayes conditioning, since by specifying $P(s)$ one specifies not only the probability of any event, but also the conditional probabilities $P(s|e)$, namely, the dynamics of how probabilities would change with any conceivable observation e.

The question naturally arises how one could ever specify, store and revise a probability function $P(s)$, when the number of states is, in effect, astronomical. This problem is the topic of much research in the literature on probabilistic diagnostics, and it has been managed to a large extent by graphical representations; Bayesian networks, influence diagrams, Markov networks etc. [Pearl, 1988].

2 Actions as Transformations of Probability Functions

If an observation e causes an agent to modify its probability from $P(s)$ to $P(s|e)$, one may ask how probabilities should change as a result of actions, rather than observations. This question does not have a clear cut answer since, in principle, actions are not part of standard probability theory; they do not serve as arguments of probability expressions nor as events for conditioning such expressions. Whenever an action is given a formal symbol, that symbol serves merely as an index for distinguishing one probability function from another, but not as a predicate which conveys information about the effect of the action. This means, for example, that the impact of two concurrent actions A and B need not have any connection to the impact of each individual action. Thus, while $P(s)$ tells us everything about responding to new observations, it tells us close to nothing about responding to external actions.

In general, if an action A is to be described as a function that takes $P(s)$ and transforms it to $P_A(s)$, then Bayesian conditioning is clearly inadequate for this transformation. For example, consider the statements: "I have observed the barometer reading to be x" and "I intervened and set the barometer reading to x". If processed by Bayes conditioning on the event "the barometer reading is x" these two reports would have the same impact on our current probabilities, yet we certainly do not consider the two reports equally informative about an incoming storm.

Philosophers [Lewis, 1974] studied another probability transformation called "imaging" (to be distinguished from "conditioning") which was deemed useful in the analysis of subjunctive conditionals. Whereas Bayes conditioning $P(s|e)$ transfers the entire probability mass from states excluded by e to the remaining states (in proportion to their current $P(s)$), imaging works differently; each excluded state s transfers its mass individually to a select set of states $S^*(s)$, which are considered "closest" to s. The reason why imaging is a more adequate representation of transformations associated with actions can be seen more clearly through a representation theorem due to Gardenfors [1988, Theorem 5.2 pp.113] (strangely, the connection to actions never appears in Gardenfors' analysis). Gardenfors' theorem states that a probability update operator $P(s) \rightarrow P_A(s)$ is an imaging operator iff it preserves mixtures, i.e.,

$$[\alpha P(s) + (1-\alpha)P'(s)]_A = \alpha P_A(s) + (1-\alpha)P'_A(s) \quad (1)$$

for all constants $1 > \alpha > 0$, all propositions A , and all probability functions P and P'. In other words, the update of any mixture is the mixture of the updates[1].

This property, called homomorphism, is what permits us to specify actions in terms of *transition probabilities*, as it is usually done in stochastic control. Denoting by $P_A(s|s')$ the probability resulting from acting A on a known state s', homomorphism (1) dictates:

$$P_A(s) = \sum_{s'} P_A(s|s')P(s') \quad (2)$$

saying that, whenever s' is not known with certainty, $P_A(s)$ is given by a weighted sum of $P_A(s|s')$ over s', with the weight being the current probability function $P(s')$. Contrasting Eq. (2) with Bayes conditioning formula,

$$P(s|A) = \sum_{s'} P(s|s', A)P(s'|A), \quad (3)$$

(2) can be interpreted as asserting that s' and A are independent, namely, A acts not as an observed proposition, but as an exogenous force since it does not alter the prior probability $P(s')$ (ordinary propositions cannot be independent of a state).

3 Action Representation: the Frame, Concurrency, and Ramification Problems

Imaging, hence homomorphism, leads to substantial savings in the representation of actions. Instead of specifying a probability $P_A(s)$ for every probability function $P(s)$ – an unboundedly long description, Eq. (2) tells us that for each action A we need to specify just one conditional probability $P_A(s|s')$. This is indeed where stochastic control theory takes off; the states s and s' are normally treated as points in some Euclidean space of real variables or parameters, and the transition probabilities $P_A(s|s')$ are encoded as deterministic equations-of-motion corrupted by random disturbances.

While providing a more adequate and general framework for actions, imaging leaves the precise specification of the transition function almost unconstrained. It does not constrain, for example, the transition associate with a concurrent action relative to those of its constituents. Aside from insisting on $P_A(s|s') = 0$ for every state s satisfying $\neg A$, we must also specify the distribution among the states satisfying A, and the number of such states may be enormous. The task of formalizing and representing these specifications can be viewed as the probabilistic version of the infamous *frame problem* and its two satellites, the ramification and concurrent actions problems.

An assumption commonly found in the literature is that the effect of an elementary action $do(q)$ is merely to change $\neg q$ to q in case the current state satisfies $\neg q$, but, otherwise, to leave things unaltered[2]. We can call this assumption the "delta" rule, variants of which are embedded in STRIPS as well as in probabilistic planning systems. In BURIDAN [Kushmerick et al, 1993], for example, every action is specified as a probabilistic mixture of several elementary actions, each operating under the delta rule.

The problem with the delta rule and its variants, is that they do not take into account the indirect ramifications of an action such as, for example, those triggered by chains of causally related events. To handle such ramifications we must construct a causal theory of the domain, specifying which event chains are likely to be triggered by a given action (the ramification problem) and how these chains interact when triggered by several actions (the concurrent action problem).

A related paper at this symposium [Darwiche & Pearl, 1994] shows how the frame, ramification and concurrency problem can be handled effectively using the language of causal graphs. The key idea is that causal knowledge can efficiently be organized in terms of just a few basic mechanisms, each involving a relatively small number of variables. Each external elementary action overrules just one mechanism leaving the others unaltered. The specification of an action then requires only the identification of the mechanism which are overruled by that action. Once this is identified, the effect of the action (or combinations thereof) can be computed from the constraints imposed by the remaining mechanisms.

The semantics behind probabilistic causal graphs and their relations to actions and belief networks have been discussed in [Goldszmidt & Pearl 1992, Pearl 1993a, Spirtes et al 1993 and Pearl 1993b]. In Spirtes et al [1993] and Pearl [1993b], for example, it is shown how the graphical representation can be used to facilitate quantitative predictions of the effects of interventions, including interventions that were not contemplated during the network construction [3].

[1] Assumption (1) is reflected in the (U8) postulate of [Katsuno and Mendelzon, 1991]: $(K_1 \vee K_2)o\mu = (K_1 o\mu) \vee (K_2 o\mu)$, where o is an update operator

[2] This assumption corresponds to Dalal's [1988] database update, which uses the Hamming distance to define the "closest world" in Lewis' imaging.

[3] Influence diagrams, in contrast, require that actions be considered in advance as part of the network.

The problem addressed in the remainder of this paper is to quantify the effect of interventions when the causal graph is not fully parameterized, that is, we are given the topology of the graph but not the conditional probabilities on all variables. Numerical probabilities will be given to only a subset of variables, in the form of unstructured conditional probability sentences. This is a more realistic setting in AI applications, where the user/designer might not have either the patience or the knowledge necessary for the specification of a complete distribution function. Some combinations of variables may be too esoteric to be assigned probabilities, and some variables may be too hypothetical (e.g., "weather conditions" or "attitude") to even be parameterized numerically.

To manage this problem, we introduce a calculus which operates on whatever probabilistic and causal information is available, and, using symbolic transformations on the input sentences, produces probabilistic assessments of the effect of actions. The calculus admits two types of conditioning operators: ordinary Bayes conditioning, $P(y|X=x)$, and causal conditioning, $P(y|set(X=x))$, that is, the probability of $Y=y$ conditioned on holding X constant (at x) by deliberate external action. Given a causal graph and an input set of conditional probabilities, the calculus derives new conditional probabilities of both the Bayesian and the causal types, and, whenever possible, generates probabilistic formulas for the effect of interventions in terms of the input information.

4 A Calculus of Actions

4.1 Preliminary Notation

Let X, Y, Z, W be four arbitrary disjoint sets of nodes in the dag G. We say that X and Y are independent given Z in G, denoted $(X \perp\!\!\!\perp Y|Z)_G$, if the set Z d-separates all paths from X to Y in G. A *causal theory* is a pair $< P, G >$, where G is a dag and P is a probability distribution compatible with G, that is, P satisfies every conditional independence relation that holds in G. The properties of d-separation are discussed in [Pearl 1988] and are summarized in [Pearl 1993b]. We denote by $G_{\overline{X}}$ ($G_{\underline{X}}$, respectively) the graph obtained by deleting from G all arrows pointing to (emerging from, respectively) nodes in X.

Finally, we replace the expression $P(y|do(x), z)$ by a simpler expression $P(y|\hat{x}, z)$, using the $\hat{}$ symbol to identify the variables that are kept constant externally. In words, the expression $P(y|\hat{x}, z)$ stands for the probability of $Y=y$ given that $Z=z$ is observed and X is held constant at x.

4.2 Inference Rules

Armed with this notation we are now able to formulate the three basic inference rules of the proposed calculus.

Theorem 1 *Given a causal theory $< P, G >$, for any sets of variables X, Y, Z, W we have:*

Rule 1 *Insertion/deletion of Observations (Bayes conditioning)*

$$P(y|\hat{x}, z, w) = P(y|\hat{x}, w) \;\; \textit{if} \;\; (Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}}} \tag{4}$$

Rule 2 *Action/observation Exchange*

$$P(y|\hat{x}, \hat{z}, w) = P(y|\hat{x}, z, w) \;\; \textit{if} \;\; (Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}\underline{Z}}} \tag{5}$$

Rule 3 *Insertion/deletion of actions*

$$P(y|\hat{x}, \hat{z}, w) = P(y|\hat{x}, w) \;\; \textit{if} \;\; (Y \perp\!\!\!\perp Z|X, W)_{G_{\overline{X}\overline{Z}}} \tag{6}$$

Each of the inference rules above can be proven from the basic interpretation of the "do(x)" operation as a replacement of the causal mechanism which connects X to its parents prior to the action by a new mechanism $X = x$ introduced by the intervention. Graphically, the replacement of this mechanism is equivalent to removing the links between X and its parents in G, while keeping the rest of the graph intact. This results in the graph $G_{\overline{X}}$.

Rule 1 reaffirms d-separation as a legitimate test for Bayesian conditional independence in the distribution determined by the intervention $do(X = x)$, hence the graph $G_{\overline{X}}$.

Rule 2 provides conditions for an external intervention $do(Z = z)$ to have the same effect on Y as the passive observation $Z = z$. It is equivalent to the "back-door" criterion of [Pearl, 1993b].

Rule 3 provides conditions for introducing (or deleting) an external intervention $do(Z = z)$ without affecting the probability of $Y = y$. The validity of this rule stems, again, from simulating the intervention $do(Z = z)$ by severing all relations between Z and its parent (hence the graph $G_{\overline{XZ}}$).

4.3 Example

We will now demonstrate how these inference rules can be used to quantify the effect of actions, given partially specified causal theories. Consider the causal theory

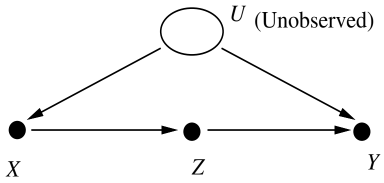

Figure 1

$< P(x, y, z), G >$ where G is the graph given in Figure 1 below, and $P(x, y, z)$ is the distribution over the observed variables X, Y, Z. Since U is unobserved, the theory is only partially specified; it will be impossible to infer all required parameters such as $P(u)$, or $P(y|z, u)$. We will see however that this structure still permits us to quantify, using our calculus, the effect of every action on every observed variable.

The applicability of each inference rule requires that certain d-separation conditions hold in some graph, the structure of which would vary with the expressions to be manipulated. Figure 2 displays the graphs that will be needed for the derivations that follow.

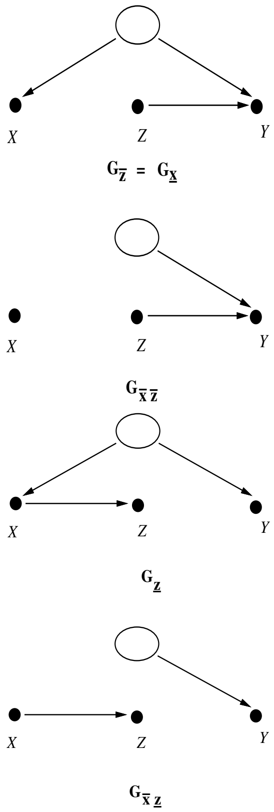

Figure 2

Task-1, compute $P(z|\hat{x})$
This task can be accomplished in one step, since G satisfies the applicability condition for Rule 2, namely $X \perp\!\!\!\perp Z$ in $G_{\underline{X}}$ (because the path $X \leftarrow U \rightarrow Y \leftarrow Z$ is blocked by the collider at Y) and we can write

$$P(z|\hat{x}) = P(z|x) \tag{7}$$

Task-2, compute $P(y|\hat{z})$
Here we cannot apply Rule 2 to exchange $\hat{z}$ by z, because $G_{\underline{Z}}$ contains a path from Z to Y (so called a "back-door" path [Pearl, 1993b]). Naturally, we would like to "block" this path by conditioning on variables (such as X) that reside on that path. Symbolically, this operation involves conditioning and summing over all values of X,

$$P(y|\hat{z}) = \sum_x P(y|x, \hat{z}) P(x|\hat{z}) \tag{8}$$

We now have to deal with two expressions involving $\hat{z}$, $P(y|x, \hat{z})$ and $P(x|\hat{z})$. The latter can be readily computed by applying Rule 3 for action deletion.

$$P(x|\hat{z}) = P(x) \text{ if } (Z \perp\!\!\!\perp X)_{G_{\overline{Z}}} \tag{9}$$

noting that, indeed, X and Z are d-separated in $G_{\overline{Z}}$. (This can be seen immediately from Figure 1; manipulating Z will have no effect on X.) To reduce the former quantity, $P(y|x, \hat{z})$, we consult Rule 2

$$P(y|x, \hat{z}) = P(y|x, z) \text{ if } (Z \perp\!\!\!\perp Y|X)_{G_{\underline{Z}}} \tag{10}$$

and note that X d-separates Z from Y in $G_{\underline{Z}}$. This allows us to write Eq. (8) as

$$P(y|\hat{z}) = \sum_x P(y|x, z) P(x) = E_x P(y|x, z) \tag{11}$$

which is a special case of the "back-door" formula [Pearl, 1993b, Eq. (14)] with $S = X$. This formula appears in a number of treatments on causal effects (see for example [Rosenbaum & Rubin, 1983; Rosenbaum, 1989; Pratt & Schlaifer, 1988]) where the legitimizing condition, $(Z \perp\!\!\!\perp Y|X)_{G_{\underline{Z}}}$ was given a variety of names, all based on conditional-independence judgments of one sort or another. Action calculus replaces such judgments by formal tests (d-separation) on a single graph (G) which represents the domain knowledge.

We are now ready to tackle a harder task, the evaluation of $P(y|\hat{x})$, which cannot be reduced to an observational expression by direct application of any of the inference rules.

Task-3, compute $P(y|\hat{x})$

Writing

$$P(y|\hat{x}) = \sum_z P(y|z, \hat{x}) P(z|\hat{x}) \tag{12}$$

we see that the term $P(z|\hat{x})$ was reduced in Eq. (7) while no rule can be applied to eliminate the manipulation symbol $\hat{}$ from the term $P(y|z,\hat{x})$. However, we can add a $\hat{}$ symbol to this term via Rule 2

$$P(y|z,\hat{x}) = P(y|\hat{z},\hat{x}) \tag{13}$$

since Figure 2 shows:

$$(Y \perp\!\!\!\perp Z|X)_{G_{\overline{X}\underline{Z}}}$$

We can now delete the action $\hat{x}$ from $P(y|\hat{z},\hat{x})$ using Rule 3, since $Y \perp\!\!\!\perp X|Z$ holds in $G_{\overline{XZ}}$. Thus, we have

$$P(y|z,\hat{x}) = P(y|\hat{z}) \tag{14}$$

which was calculated in Eq. (11). Substituting, (11), (14, and (7) back in (12), finally yields

$$P(y|\hat{x}) = \sum_z P(z|x) \sum_{x'} P(y|x',z)P(x') \tag{15}$$

Eq. (15) was named the Mediating Variable formula in [Pearl, 1993c], where it was derived by algebraic manipulation of the joint distribution and taking the expectation over U.

Task-4, compute $P(y,z|\hat{x})$

$$P(y,z|\hat{x}) = P(y|z,\hat{x})P(z|\hat{x}) \tag{16}$$

The two terms on the r.h.s. were derived before in Eqs. (7) and (14), from which we obtain

$$\begin{aligned} P(y,z|\hat{x}) &= P(y|\hat{z})P(z|x) \\ &= P(z|x)\textstyle\sum_{x'} P(y|x',z)P(x') \end{aligned}$$

5 Discussion

In this example we were able to compute answers to all possible queries of the form $P(y|z,\hat{x})$ where Y, Z, and X are subsets of observed variables. In general, this will not be the case. For example, there is no general way of computing $P(y|\hat{x})$ from the observed distribution whenever the causal model contains the subgraph shown in Figure 3, where X and Y are adjacent, and the dashed line represents a path traversing unobserved variable[4]. Similarly, our ability to compute $P(y|\hat{x})$ for every pair of singleton variables does not ensure our ability to compute joint distributions, e.g. $P(y_1,y_2|\hat{x})$. Figure 4, for example, shows a causal graph where both $P(y_1|\hat{x})$ and $P(y_2|\hat{x})$ are computable, but $P(y_1,y_2|\hat{x})$ is not. Consequently, we cannot compute $P(z|\hat{x})$. Interestingly, the graph of Figure 4 is the smallest graph which does not contain the pattern of Figure 3 and still presents an uncomputable causal effect.

[4] One can calculate upper and lower bounds on $P(y|\hat{x})$ and these bounds may coincide for special distributions, $P(x,y,z)$ [Balke & Pearl, 1993] but there is no way of computing $P(y|\hat{x})$ for *every* distribution $P(x,y,z)$.

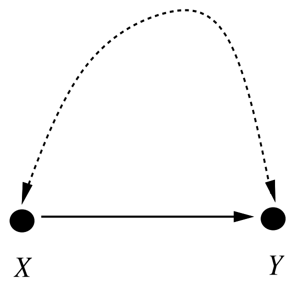

Figure 3

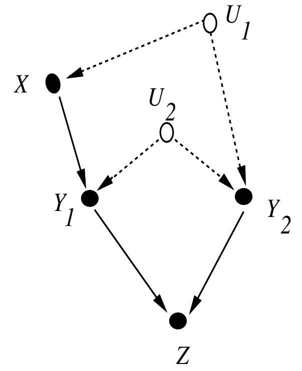

Figure 4

Another interesting feature demonstrated by the network in Figure 4 is that it is often easier to compute the effect of a joint action than the effects of its constituent singleton actions[5]. In this example, it is possible to compute $P(z|\hat{x},\hat{y}_1)$, yet there is no way of computing $P(z|\hat{x})$. For example, the former can be evaluated by invoking Rule 2, writing

$$P(z|\hat{x},\hat{y}_2) = \sum_{y_1} P(z|y_1,\hat{x},\hat{y}_2)P(y_1|\hat{x},\hat{y}_2)$$

$$= \sum_{y_1} P(z|y_1,x_1,y_2)P(y_1|x)$$

On the other hand, Rule 2 cannot be applied to the computation of $P(y_1|\hat{x},y_2)$ because, conditioned on Y_2, X and Y_1 are d-connected in $G_{\underline{X}}$ (through the dashed lines). We conjecture, however, that whenever $P(y|\hat{x}_i)$ is computable for every singleton x_i, then $P(y|\hat{x}_1,\hat{x}_2,...\hat{x}_l)$ is computable as well, for any subset of variables $\{X_1,...X_l\}$.

Computing the effect of actions from partial theories in which probabilities are specified on a select subset of (observed) variables is an extremely important task

[5] The fact that the two tasks are not equivalent was brought to my attention by James Robins who has worked out many of these computations in the context of sequential treatment management [Robins 1989].

in statistics and socio-economic modeling, since it determines when a parameter of a causal theory are (so called) "identifiable" from non-experimental data, hence, when randomized experiments are not needed. The calculus proposed above, indeed uncovers possibilities that have remained unnoticed by economists and statisticians. For example, the structure of Figure 4 uncovers a class of observational studies in which the causal effect of an action (X) can be determined by measuring a variable (Z) that mediates the interaction between the action and its effect (Y). The relevance of such structures in practical situations can be seen, for instance, if we identify X with smoking, Y with lung cancer, Z with the amount of tar deposits in one's lung and U with an unobserved carcinogenic genotype which, according to the tobacco industry also induces an inborn crave for nicotine. Eq. (20) would provide us in this case with the means for quantifying, from non-experimental data, the causal effect of smoking on cancer. (Assuming, of course, that the data $P(x, y, z)$ is made available, and that we believe that smoking does not have a direct effect on lung cancer except that mediated by tar deposits).

However, our calculus is not limited to the derivation of causal probabilities from non-causal probabilities; we can reverse the role, and derive conditional and causal probabilities from causal expressions as well. For example, given the graph of figure 3 together with the quantities $P(z|\hat{x})$ and $P(y|\hat{z})$, we can derive an expression for $P(y|\hat{x})$,

$$P(y|\hat{x}) = \sum_z P(y|\hat{z}) P(z|\hat{x}) \qquad (17)$$

using the steps that led to Eq. (19). Note that the derivation is still valid when we add a common cause to X and Z, which is the most general condition under which the transitivity of causal relationships holds.

BIBLIOGRAPHY

Balke, A. and J. Pearl (1993) Nonparametric bounds on treatment effects in partial compliance studies. em Technical Report R-199, UCLA Computer Science Department. (Submitted.)

Dalal, M. (1988) Investigations into a Theory of Knowledge Base Revision: Preliminary Report, in *Proceedings of the Seventh National Conference on Artificial Intelligence*, 475-479.

Darwiche, A. and J. Pearl (1994) Symbolic Causal Networks for Reasoning about Actions and Plans. In this volume.

Gardenfors, P. (1988) *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, 1988.

Goldszmidt, M. and J. Pearl (1992) Default ranking: A practical framework for evidential reasoning, belief revision and update. in *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, Morgan Kaufmann, San Mateo, CA, 661-672.

Katsuno, H. and A. Mendelzon (1991) On the Difference Between Updating a Knowledge Base and Revising it, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, 387-394.

Kushmerick, N., S. Hanks, and D. Weld (1993) An Algorithm for Probabilistic Planning. *Technical Report 93-06-03.* Department of Computer Science and Engineering, University of Washington.

Lewis, D.K. (1974) Probability of Conditionals and Conditional Probabilities. *The Philosophical Review*, 85, 297-315.

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Mateo, CA: Morgan Kaufmann.

Pearl, J. (1993a) From Conditional Oughts to Qualitative Decision Theory. in *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence* (D. Heckerman and A. Mamdani (eds.)), Morgan Kaufmann, San Mateo, CA,

Pearl, J. (1993b) Graphical Models, Causality, and Intervention. *Statistical Science*, **8**, (3), 266-273.

Pearl, J. (1993c) Mediating Instrumental Variables. *Technical Report R-210*, Cognitive Systems Laboratory, UCLA Computer Science Department.

Pratt, J., and R. Schlaifer (1988) On the Interpretation and Observation of Laws. *Journal of Economics*, **39**, 23-52.

Robins, J.M. (1989) The Analysis of Randomized and Non-Randomized AIDS Treatment Trials using a New Approach to Causal Inference in Longitudinal Studies, in *Health Service Research Methodology: A Focus on AIDS*, (eds. L. Sechrest, H. Freeman, and A. Mulley), NCHSR, U.S. Public Health Service, 113-159, 1989.

Rosenbaum, P.R. (1989) The Role of Known Effects in Observational Studies. *Biometrics*, **45**, 557-569.

Rosenbaum, P. and D. Rubin (1983) The Central Role of Propensity Score in Observational Studies for Causal Effects. *Biometrica*, **70**, 41-55.

Spirtes, P., C. Glymour and R. Schienes (1993) *Causation, Prediction, and Search*, Springer-Verlag, New York.