

Comment: Graphical Models, Causality and Intervention

Judea Pearl

I am grateful for the opportunity to respond to these two excellent papers. Although graphical models are intuitively compelling for conceptualizing statistical associations, the scientific community generally views such models with hesitancy and suspicion. The two papers before us demonstrate the use of graphs—specifically, directed acyclic graphs (DAGs)—as a mathematical tool of great versatility and thus promise to make graphical languages more common in statistical analysis. In fact, I find my own views in such close agreement with those of the authors that any attempt on my part to comment directly on their work would amount to sheer repetition. Instead, as the editor suggested, I would like to provide a personal perspective on current and future developments in the areas of graphical and causal modeling. A complementary account of the evolution of belief networks is given in Pearl (1993a).

I will focus on the connection between graphical models and the notion of causality in statistical analysis. This connection has been treated very cautiously in the papers before us. In Lauritzen and Spiegelhalter (1988), the graphs were called “causal networks,” for which the authors were criticized; they have agreed to refrain from using the word “causal.” In the current paper, Spiegelhalter et al. deemphasize the causal interpretation of the arcs in favor of the “irrelevance” interpretation. I think this retreat is regrettable for two reasons: first, causal associations are the primary source of judgments about irrelevance, and, second, rejecting the causal interpretation of arcs prevents us from using graphical models for making legitimate predictions about the effect of actions. Such predictions are indispensable in applications such as treatment management and policy analysis. I would like to supplement the discussion with an account of how causal models and graphical models are related.

It is generally accepted that, because they provide information about the dynamics of the system under study, causal models, regardless of how they are discovered or tested, are more useful than associational models. In other words, whereas the joint distribution

tells us how probable events are and how probabilities would change with subsequent observations, the causal model also tells us how these probabilities would change as a result of external interventions in the system. For this reason, causal models (or “structural models” as they are often called) have been the target of relentless scientific pursuit and, at the same time, the center of much controversy and speculation. What I would like to discuss in this commentary is how complex information about external interventions can be organized and represented graphically and, conversely, how the graphical representation can be used to facilitate quantitative predictions of the effects of interventions.

The basic idea goes back to Simon (1977) and is stated succinctly in his foreword to Glymour et al. (1987): “The advantage of representing the system by structural equations that describe the direct causal mechanisms is that if we obtain some knowledge that one or more of these mechanisms has been altered, we can use the remaining equations to predict the consequences—the new equilibrium.” Here, by “mechanism” Simon means any stable relationship between two or more variables that remains invariant to external influences until it falls directly under such influences.

This mechanism-based model was adapted in Pearl and Verma (1991) for defining probabilistic causal theories; each child-parent family in a DAG Γ represents a deterministic function $X_i = f_i(\mathbf{pa}_i, \varepsilon_i)$, where \mathbf{pa}_i are the parents of variable X_i in Γ , and ε_i , $0 < i < n$, are mutually independent, arbitrarily distributed random disturbances. Characterizing each child-parent relationship as a deterministic function, instead of the usual conditional probability $P(x_i | \mathbf{pa}_i)$, imposes equivalent independence constraints on the resulting distributions and leads to the same recursive decomposition

$$(1) \quad P(x_1, \dots, x_n) = \prod_i P(x_i | \mathbf{pa}_i)$$

that appears in Eq. (1) of Spiegelhalter et al.’s article. However, the functional characterization also specifies how the resulting distribution would change in response to external interventions, since, by convention, each function is presumed to remain constant unless specifically altered. This formulation is merely a nonlinear generalization of the usual structural equation models, where function constancy (or stability) is implicitly

Judea Pearl is Professor of Computer Science and Director of the Cognitive Systems Laboratory, University of California Los Angeles, 405 Hilgard Avenue, Los Angeles, California 90024.

assumed. Moreover, the nonlinear character of f_i permits us to treat changes in the function f_i itself as a variable, F_i , by writing

$$(2) \quad X_i = f_i(\mathbf{pa}_i, F_i, \varepsilon_i)$$

where

$$f_i(a, b, c) = f_i(a, c) \text{ whenever } b = f_i.$$

Thus, any external intervention F_i that alters f_i can be represented graphically as an added parent node of X_i , and the effect of such an intervention can be analyzed by Bayesian conditionalization, that is, by simply setting this added parent variable to the appropriate value f_i .

The simplest type of external intervention is one in which a single variable, say X_i , is forced to take on some fixed value x'_i . Such intervention, which we call *atomic*, amounts to replacing the old functional mechanism $X_i = f_i(\mathbf{pa}_i, \varepsilon_i)$ with a new mechanism $X_i = x'_i$ governed by some external force F_i that sets the value x'_i . If we imagine that each variable X_i potentially could be subject to the influence of such an external force F_i , then we can view the causal network Γ as an efficient code for predicting the effects of atomic interventions and of various combinations of such interventions.

The effect of an atomic intervention $set(X_i = x'_i)$ is encoded by adding to Γ a link $F_i \rightarrow X_i$ (Figure 1), where F_i is a new variable taking values in $\{set(x'_i), idle\}$, x'_i ranges over the domain of X_i , and *idle* represents no intervention. Thus, the new parent set of X_i in the augmented network is $\mathbf{pa}'_i = \mathbf{pa}_i \cup \{F_i\}$, and it is related to X_i by the conditional probability

$$(3) \quad P(x_i | \mathbf{pa}'_i) = \begin{cases} P(x_i | \mathbf{pa}_i), & \text{if } F_i = idle, \\ 0, & \text{if } F_i = set(x'_i) \\ & \text{and } x_i \neq x'_i, \\ 1, & \text{if } F_i = set(x'_i) \\ & \text{and } x_i = x'_i. \end{cases}$$

The effect of the intervention $set(x'_i)$ is to transform the original probability function $P(x_1, \dots, x_n)$ into a new function $P_{x'_i}(x_1, \dots, x_n)$, given by

$$(4) \quad P_{x'_i}(x_1, \dots, x_n) = P'(x_1, \dots, x_n | F_i = set(x'_i)),$$

where P' is the directed Markov field dictated by the augmented network $\Gamma' = \Gamma \cup \{F_i \rightarrow X_i\}$ and (3), with an arbitrary prior distribution on F_i . In general, by adding a hypothetical intervention link $F_i \rightarrow X_i$ to each node in Γ , we can construct an augmented probability function $P'(x_1, \dots, x_n; F_1, \dots, F_n)$ that contains information about richer types of interventions. Multiple interventions would be represented by conditioning P' on a subset of the F_i 's (taking values in their respective $set(x'_i)$), while the preintervention probability function P would be viewed as the posterior distribution induced by conditioning each F_i in P' on the value *idle*.

This representation yields a simple and direct transformation between the preintervention and the postintervention distributions:

$$(5) \quad P_{x'_i}(x_1, \dots, x_n) = \begin{cases} \frac{P(x_1, \dots, x_n)}{P(x_i | \mathbf{pa}_i)}, & \text{if } x_i = x'_i \\ 0, & \text{if } x_i \neq x'_i. \end{cases}$$

This transformation reflects the removal of the term $P(x_i | \mathbf{pa}_i)$ from the product decomposition of (1), since \mathbf{pa}_i no longer influence X_i . Transformations involving conjunctive and disjunctive actions can be obtained by straightforward applications of (4) (Goldszmidt and Pearl, 1992; Pearl, 1993b; Spirtes, Glymour and Scheines 1993). The transformation exhibits the following properties:

1. An intervention $set(x'_i)$ can affect only the descendants of X_i in Γ .
2. For any set S of variables, we have

$$(6) \quad P_{x'_i}(S | \mathbf{pa}_i) = P(S | x'_i, \mathbf{pa}_i).$$

In other words, given $X_i = x'_i$ and \mathbf{pa}_i , it is superfluous to find out whether $X_i = x'_i$ was established by external intervention or not. This can be seen directly from the augmented network Γ' (Figure 1), since $\{X_i\} \cup \mathbf{pa}_i$ *d*-separates F_i from the rest of the network, thus legitimizing the conditional independence $S \perp\!\!\!\perp F_i | (X_i, \mathbf{pa}_i)$.




FIG. 1. Representing external intervention, F_i , by an augmented network $\Gamma' = \Gamma \cup \{F_i \rightarrow X_i\}$.

3. A necessary and sufficient condition for an external intervention $set(X_i = x'_i)$ to have the same effect on X_j as the passive observation $X_i = x'_i$ is that X_i d -separates \mathbf{pa}_i from X_j , that is,

$$(7) \quad P_{x'_i}(x_j) = P(x_j | x'_i) \text{ iff } X_j \perp\!\!\!\perp \mathbf{pa}_i | X_i.$$

Equation (4) explains why randomized experiments are sufficient for estimating the effect of interventions even when the causal network is not given: because the intervening variable F_i enters the networks as a root node (i.e., independent of all other ancestors of X_i) it is equivalent to a treatment-selection policy governed by a random device.

The immediate implication of (5) is that, given the structure of the causal network Γ , one can infer postintervention distributions from preintervention distributions; hence, we can reliably estimate the effects of interventions from passive (i.e., nonexperimental) observations. Of course, (5) does not imply that we can always substitute observational studies for experimental studies, as this would require an estimation of $P(x_i | \mathbf{pa}_i)$. The mere identification of \mathbf{pa}_i (i.e., the direct causal factors of X_i) requires substantive causal knowledge of the domain which is often unavailable. Moreover, even when we have sufficient substantive knowledge to structure Γ , some members of \mathbf{pa}_i may be unobservable, or *latent*. Fortunately, there are conditions for which an unbiased estimate of $P_{x'_i}(x_j)$ can be obtained even when the \mathbf{pa}_i variables are latent and, moreover, a simple graphical criterion can tell us when these conditions are satisfied.

Assume we are given a causal network Γ together with nonexperimental data on a subset X_o of observed variables in Γ and we wish to estimate what effect the intervention $set(X_i = x'_i)$ would have on some response variable X_j . In other words, we seek to estimate $P_{x'_i}(x_j)$ from a sample estimate of $P(X_o)$. Applying (4), we can write

$$(8) \quad \begin{aligned} P_{x'_i}(x_j) &= P'(x_j | F_i = set(x'_i)) \\ &= \sum_{\mathbf{S}} P'(x_j | \mathbf{S}, X_i = x'_i, F_i = set(x'_i)) \\ &\quad \times P'(\mathbf{S} | F_i = set(x'_i)), \end{aligned}$$

where \mathbf{S} is any set of variables. Clearly, if \mathbf{S} satisfies

$$(9) \quad \mathbf{S} \perp\!\!\!\perp F_i \text{ and } X_j \perp\!\!\!\perp F_i | (X_i, \mathbf{S}),$$

then (8) can be reduced to

$$(10) \quad \begin{aligned} P_{x'_i}(x_j) &= \sum_{\mathbf{S}} P(x_j | \mathbf{S}, x'_i) P(\mathbf{S}) \\ &= E_{\mathbf{S}}[P(x_j | \mathbf{S}, x'_i)]. \end{aligned}$$

Thus, if we find a set $\mathbf{S} \subseteq X_o$ of observables satisfying (9), we can estimate $P_{x'_i}(x_j)$ by taking the expectation (over \mathbf{S}) of $P(x_j | \mathbf{S}, x'_i)$, and the latter can easily be

estimated from nonexperimental data. It is also easy to verify that (9) is satisfied by any set \mathbf{S} that meets the following *back-door criterion*:

1. No node in \mathbf{S} is a descendant of X_i , and
2. \mathbf{S} d -separates X_i from X_j along every path containing an arrow into X_i .

The name “back-door” echos condition 2, which requires that only indirect paths from X_i to X_j be d -separated; these paths can be viewed as entering X_i through the back door.

In Figure 2, for example, the sets $\mathbf{S}_1 = \{X_3, X_4\}$ and $\mathbf{S}_2 = \{X_4, X_5\}$ would qualify under the back-door criterion, but $\mathbf{S}_3 = \{X_4\}$ would not because X_4 does not d -separate X_i from X_j along the path $(X_i, X_3, X_1, X_4, X_2, X_5, X_j)$. Thus, we have obtained a simple graphical criterion for finding a set of observables for estimating (by conditioning) the effect of interventions from purely nonexperimental data.

It is interesting that the conditions formulated in (9) are equivalent to those known as *strongly ignorable treatment assignment* (SITA) conditions in Rubin’s model for causal effect (Pearl, 1993c; Rosenbaum and Rubin, 1983). [The graphical translation of Rubin’s model invokes the mechanism $X_i \rightarrow X_j \leftarrow r$, where X_i represents the treatment-assignment, X_j the observed response, and r represents the causal-effect variable. Indeed, following the counterfactual interpretation of r , X_j is a deterministic function of X_i and r , and r plays the role of f_i in (2) (Pearl, 1993c)]. Reducing the SITA conditions to the graphical back-door criterion facilitates the search for an optimal conditioning set \mathbf{S} and significantly simplifies the judgments required for ratifying the validity of such conditions in practical situations.

Equation (4) was derived under the assumption that the preintervention probability P is given by the product of (1), which represents a joint distribution prior to making any observations. To predict the effect of action F_i after observing C , we must also invoke assumptions about persistence, so as to distinguish prop-




FIG. 2. A DAG representing the back-door criterion, adjusting for variables $\{X_3, X_4\}$ (or $\{X_4, X_5\}$) yields an unbiased estimate of $P(x_j | set(x'_i))$.

erties that will terminate as a result of F_i from those that will persist despite of acting F_i . Such a model of persistence was invoked in (Pearl, 1993b); there, it was assumed that only those properties should persist that are not under any causal influence to terminate. This assumption yields formulas for the effect of *conditional interventions* (conditioned on the observation C) which, again, given Γ , can be estimated from nonexperimental data.

A more ambitious task has been explored by Spirtes, Glymour and Scheines, (1993) – estimation of the effect of intervention when the structure of Γ is not available and must also be inferred from the data. Recent developments in graphical models (Pearl and Verma, 1991; Spirtes, Glymour and Scheines, 1993) have produced methods that, under certain conditions, permit us to infer plausible causal structures from nonexperimental data, albeit with a weaker set of guarantees than those obtained through controlled randomized experiments. These guarantees fall into two categories: minimality and stability (Pearl and Verma, 1991). Minimality guarantees that any other structure compatible with the data is necessarily more redundant, and hence less trustworthy, than the one(s) inferred. Stability ensures

that any alternative structure compatible with the data must be less stable than the one(s) inferred; namely, slight fluctuations in the distributions of the disturbances ε_i (2) will render that structure no longer compatible with the data.

When the structure of Γ is to be inferred under these guarantees, the formulas governing the effects of interventions and the conditions required for estimating these effects become rather complex (Spirtes, Glymour and Scheines, 1993). Alternatively, one can produce bounds on the effect of interventions by taking representative samples of inferred structures and estimating $P_{x_i}(x_j)$ according to (10) for each such sample.

In summary, I hope my comments convince the reader that DAGs can be used not only for specifying assumptions of conditional independence but also as a formal language for organizing claims about external interventions and their interactions. I hope to have demonstrated as well that DAGs can serve as an analytical tool for predicting, from nonexperimental data, the effect of actions (given substantive causal knowledge), for specifying and testing conditions under which randomized experiments are not necessary and for aiding experimental design and model selection.

Comment

Michael E. Sobel

It is a pleasure to discuss these excellent papers. Spiegelhalter, Dawid, Lauritzen and Cowell nicely put together a number of themes, demonstrating, in a Bayesian context, the utility of graphical modelling in the construction of probabilistic expert systems. The authors show how graphs can be used heuristically to solicit expert opinion, and in Section 6, how the theory of conditional independence graphs can be used to make tractable (while maintaining reasonable substantive assumptions) the calculation of probabilistic features of the system (monitors). For example, the authors want to apply to the directed independence graph of their Figure 2 the decomposability theorem for undirected conditional independence graphs, which permits a full factorization of the probability distribution. To do so, they associate the graph of Figure 2 with its moral graph (an undirected conditional independence

graph) and use the fact that the separation properties of the moral graph apply to the directed independence graph. They then embed the moral graph into a triangulated graph, enabling use of the desired theorem; further simplifications come from organizing the cliques of the triangulated graph into junction trees.

My vantage point is that of a social statistician: as such, there is more for me to say about the paper by Cox and Wermuth. In particular, I want to expand on and further tie several themes in this paper to research in the social and behavioral sciences. Thus, discussion focuses primarily on this paper; I shall often freely borrow notation from there.

TYPES OF INDEPENDENCE GRAPHS

Cox and Wermuth nicely characterize various types of dependencies among random variables. Prior work has focused attention on two types of independence graphs. If no ordering is imposed on the variables, undirected graphs are used; here, the absence of an edge between two vertices denotes conditional independence of the variables associated with the vertices,

Michael E. Sobel is Professor of Sociology and of Applied Mathematics, Department of Sociology, University of Arizona, Tucson, Arizona 85721.