Some Notes on Generalized Interventional Potentials

Mark Schmidt
Department of Computer Science
University of British Columbia
schmidtm@cs.ubc.ca

1 Generalized Interventional Potentials

It is common to model interventions in DAG models using Pearl's *do*-calculus [Pearl, 2000]. However, for modeling interventional data the assumption of acyclility is often innappropriate; many models of biological networks contain feedback cycles (for example, see Sachs et al. [2005]). In contrast, undirected graphical models allow cycles. However, under most interpretations of the data generating processes associated with undirected graphs there is no difference between conditioning by observation and conditioning by intervention [Lauritzen and Richardson, 2002]; undirected models do not distinguish between observing a variable ('seeing') and setting it by intervention ('doing').

Motivated by the problem of using cyclic models for interventional data, in [Schmidt and Murphy, 2009] we defined the notion of an *interventional* potential. These are undirected potential functions that are augmented with interventional semantics. However, in that previous work we focused on the case of *pairwise* potentials where each potential has a *single* target variable. In this note we consider a generalization of interventional potential that allows *multiple* target variables (or even *no* target variables), and the interpretation of *higher-order* interventional potentials in the special case of hierarchical models.

2 Intervention Semantics by Variable-Partitioning

In causal DAGs, the effect of a (perfect) intervention on node i is to remove the CPD $p(x_i|\mathbf{x}_{\pi(i)})$ from the joint distribution, removing the statistical association between the effect x_i and its causes $\mathbf{x}_{\pi(i)}$. We would like to define similar semantics for the potential functions $\phi_A(\mathbf{x}_A)$ in undirected models, since it might be the case that the statistical relationship represented by ϕ_A might be removed if we intervene on some element of A. However, the nodes in A are treated symmetrically in ϕ_A so we must augment the potential with additional information that defines the effects of possible interventions. Toward this end, we consider using an undirected graphical where we define the probability of a set of p random variables as a globally normalized product of interventional potentials $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$,

$$p(\mathbf{x}) \triangleq \frac{1}{Z}\prod_A \phi_A(\mathbf{x}_B|\mathbf{x}_C). \tag{1}$$

In the interventional potential $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$ we require that B and C from a partition of A, including the case where either B or C is the empty set.[1] Given only observational data, $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$ is simply defined as the usual undirected potential $\phi_A(\mathbf{x}_A)$. Thus, for observational data this representation reduces to the standard undirected graphical model representation. Further, in analogy with the DAG case we say that the effect of a (perfect) intervention on node i in (1) is to remove all potentials $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$ where i is an element of B.

An interesting aspect of these potentials is that we get different types of edges depending on the particular partition of A into B and C. In the case of pairwise edges (where A has two elements), we describe the three possibilities below:

1. **Undirected edges**: If we have $\phi_{ij}(\emptyset|\mathbf{x}_{ij})$, then the statistical relationship between i and j is preserved under intervention on either node. Thus, edges with this partition have the usual semantics of interventions in undirected graphs [Lauritzen and Richardson, 2002], where there is no difference between observation and intervention. Such an edge might reflect a purely associative relationship, or that the statistical relationship between the variables is due to a latent variable that is a common effect.

[1] We obtain the interventional potentials used in [Schmidt and Murphy, 2009] as the special case where we force B to be a singleton.

2. **Directed edges**: If we have $\phi_{ij}(x_i|x_j)$, the statistical relationship between i and j is preserved under intervention on j but lost under intervention on i. Edges with this partition adopt the graphical semantics associated with causal DAGs, and were the type of interventional potential we considered in [Schmidt and Murphy, 2009]. Such an edge might reflect that x_j has a direct causal effect on x_i, or that x_j has a causal influence on x_i through a latent variable.[2]

3. **Unstable edges**: If we have $\phi_{ij}(\mathbf{x}_{ij}|\emptyset)$, the statistical relationship between i and j is lost under intervention on either node. This type of edge does not arise in (purely) directed or undirected models. Such an edge might reflect that the statistical relationship between x_i and x_j exists because of a latent variable that is a common cause.

Note that in causal DAGs we can model indirect causal effects and common effects due to latent variables, without explicitly introducing latent variables. However, it is not possible to model common latent causes using DAGs over the observed variables, while unstable edges allow interventional potentials to represent this case.

In the pairwise case, we can visualize the dependencies encoded by the interventional potentials as a graph with directed, undirected, and unstable edges. The conditional independence properties in the observational distribution are simply given by ignoring the edge types and using the conditional independence properties of the resulting undirected graph. When we intervene on node i, we first remove unstable edges involving i and directed edges into i, and then we remove the edge types and use the conditional independence properties of the resulting undirected graph. Note that directed cycles are allowed in both the observational and interventional distributions.

3 Higher-Order Potentials

The independence properties encoded in a model with pairwise interventional potentials can be represented by a directed cyclic graph. However, when we have higher-order factors it will prove more useful to represent the independencies using a factor graph [Koller and Friedman, 2009, §4.4.1] augmented with additional information about the effects of intervention. In particular, we first ignore the partition of each potential ϕ_A into B and C and draw a standard factor graph. We then add an arrow from each factor to the nodes in the first component of the partition B, and in this case we will refer to nodes in B as children and nodes in C as parents. In the observational distribution, we simply ignore these arrows and apply graph separation in the factor graph. If we intervene on node i, we can visualize all dependencies present in the interventional distribution by removing all factors with an arrow into node i.

We discussed the three types of edges that arise from different partitions of a pairwise interventional potential, and how they can be interpreted. Here, we use latent variables to give an interpretation to arbitrary higher-order interventional potentials. First, we will make an assumption that is similar to the hierarchical inclusion assumption present in hierarchical log-linear models: if we have a factor $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$, then we also have the factors $\phi_A(\mathbf{x}_D|\mathbf{x}_C)$ for all $D \subset B$. Then, to give an interpretation to the factor $\phi_A(\mathbf{x}_B|\mathbf{x}_C)$, we draw a causal DAG over the variables B, C, and an additional latent variable y. We draw directed edges from all variables in C to y, and from y to all variables in B. We give the four possible cases for a three-way factor below:

1. **Three causes (undirected edge)**: If we have $\phi_{ijk}(\emptyset|\mathbf{x}_{ijk})$, then all three nodes point to the common effect y, and the potential remains after intervention on any node. This is the usual semantics of intervention in undirected graphical models [Lauritzen and Richardson, 2002].

2. **Two causes, one effect (directed edge)**: If we have $\phi_{ijk}(x_i|\mathbf{x}_{jk})$, then j and k point to y and y points to i. Here, intervention on the effect i removes the potential but intervening on the causes j or k leaves the potential unaffected. This is the interventional semantics associated with a causal DAG model where i is the head of a v-structure between the three nodes.

3. **One cause, two effects**: If we have $\phi_{ijk}(\mathbf{x}_{ij}|x_k)$, then k points to y and y points to i and j. Here, intervention on k leaves the potential unaffected but intervention on i or j removes the potential. However, note that if we only intervene on one of the nodes, say i, then the lower-order potential $\phi_{jk}(x_j|x_k)$ will remain that preserves the dependency between j and k.

4. **Three effects (unstable edge)**; If we have $\phi_{ijk}(\mathbf{x}_{ijk}|\emptyset)$, then y points to all three nodes. Here, intervention on any node removes the potential. This is similar to unstable pairwise edges. However, if we only intervene on one of the nodes, say i, then the lower order potential $\phi_{jk}(\mathbf{x}_{jk}|\emptyset)$ preserves the dependency between the other two. However, if we intervene on two of the nodes then the dependency on the third node is removed.

[2] We could also have both $\phi_{ij}(x_i|x_j)$ and $\phi_{ij}(x_j|x_i)$, and this may represent a feedback cycle.

4 Other Possible Effects of Intervetions

In this note we considered a generalization of interventional potentials where for each potential we choose a subset of the variables in the potential, and intervention on *any* variable in the subset removes the potential from the interventional distribution. Another alternative would be to have a potential where we must intervene on *all* variables in the subset to remove the potential, or some other combinatorial structure.

References

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press, 2009.

S. Lauritzen and T. Richardson. Chain graph models and their causal interpretations. *Journal of the Royal Statistical Society: Series B*, 64(3):321–361, 2002.

J. Pearl. *Causality: Models, reasoning, and inference.* Cambridge Univ Press, 2000.

K. Sachs, O. Perez, D. Pe'er, D. Lauffenburger, and G. Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308 (5721):523–529, 2005.

M. Schmidt and K. Murphy. Modeling Discrete Interventional Data using Directed Cyclic Graphical Models. *Conference on Uncertainty in Artificial Intelligence*, 2009.