Weak Interventions and Instrumental Variables

Frederick Eberhardt
University of California, Berkeley
fde@berkeley.edu

1 Extended Abstract

Traditional experimental design has focused on experimental interventions that take full control of the distribution of treatment variables by means of randomization or clamping. The underlying motiviation, going back to R.A. Fisher, is that such interventions make the treatment variable independent of its causes, including potential latent confounders of the treatment and outcome, and therefore enable unbiased estimates of the causal effect of the treatment on the outcome. In many cases, especially in the social sciences, such "surgical" interventions (Pearl) are not possible or not feasible. However, this does not imply that discovery of the the causal structure among a set of variables is limited to passive observational data. Interventions can be weaker, influencing the conditional distribution of the intervened variable given its causes (graphical parents), without making the intervened variable independent of its normal causes. The implications for learning the causal structure under these circumstances do not revert to results available for "surgical" interventions or passive observation.

Given that a weak intervention introduces a known external influence into the system under investigation, but does not destroy causal connections, under particular assumptions more can be learned faster about causal structure than with "surgical" interventions or passive observation – even if there are latent variables. That is, in fewer experiments, more details of the causal structure can be determined than when using "surgical" interventions or not intervening at all. The results rely on the assumption that the intervention is not confounded by a latent variable and not caused by any variable in the system under investigation, and that the distribution over the variables under investigation is faithful to a directed acyclic graph. The results do not depend on the model being linear or continuous.

Formally, these conditions match the assumptions made about instrumental variables in economics. Instrumental variables are used for the estimation of the effect of X on Y when the error of Y is correlated with X. To obtain an unbiased estimate a further variable, Z, called an instrument, is used. Z is independent of the error and only related to Y through X. By performing weak interventions, one can, so to speak, create an instrumental variable to obtain unbiased estimators of the causal influence of the treatment on the outcome variables.

In graphical terms, the weak intervention is represented by an additional intervention variable with a direct arrow from the intervention variable into the intervened variable. If the weak intervention only influences one variable and is not caused by any variable in the system under investigation, then the intervention variable creates an unshielded collider (V-structure) with the intervened variable and any of its causes. This unshielded collider implies a particular signature in the independence constraints that can be identified. For some approaches of identifying V-structures, discovery does not even depend on a marginal distribution over the intervention variable, but can be computed purely in terms of conditional distributions. Furthermore, weak interventions can be combined indepedently of one another. While a simultaneous randomization of two variables implies that nothing can be learned about the relation between the two, this is not the case for weak interventions. Since the intervenened variables are not made independent of their normal causes by the weak intervention, some aspects of the causal relation between the two variables remain and can be discovered. This leads to a faster search for graphical structure – faster, in the sense of fewer experiments, but not necessarily in terms of convergence. Since higher order conditional independence tests are required under these circumstances, more samples are needed. But in at least some social science cases, sample size is not the issue, and so this trade-off might be acceptable.

We will be presenting results, both positive and negative, under a variety of different assumptions, of what can and cannot be learned about the causal structure among a set of variables using weak interventions. The results are aimed to build a bridge connecting research on the search for caual structure with interventions to the literature on instrumental variables in economics.