Curing the Curse of Non-Recursiveness in Structural Causal Models

Stephan Bongers
Informatics Institute
University of Amsterdam
Amsterdam, The Netherlands
s.r.bongers@uva.nl

Jonas Peters
Department of Mathematical Sciences
University of Copenhagen
Copenhagen, Denmark
jonas.peters@math.ku.dk

Bernhard Schölkopf
Max Planck Institute for Intelligent Systems
Tübingen, Germany
bs@tuebingen.mpg.de

Joris M. Mooij
Informatics Institute
University of Amsterdam
Amsterdam, The Netherlands
j.m.mooij@uva.nl

Structural causal models (SCMs) are widely used for causal modeling purposes (Pearl, 2009; Spirtes et al., 2000). In these models, the causal relationships are expressed in the form of deterministic, functional relationships, and the probabilities are introduced through the assumption that certain variables are exogenous. Usually one assumes that the structural equations of the SCM should have a unique solution (Pearl, 2009), which we call the unique solvability property. This property is needed in order to define the SCM in terms of a unique set of random variables. Acyclic SCMs, also known as recursive structural equations models, form a special class for which this property always holds. They have several other interesting properties, one of which is that acyclicity is preserved under perfect interventions, and hence, for this class, unique solvability is preserved under such interventions. However, in the cyclic setting this property of unique solvability is not preserved under interventions, which means that an SCM initially described by a set of random variables may not be defined anymore after performing an intervention. Moreover, in a non-uniquely solvable SCM it may happen that the functions of the variables are well-defined, but they lack measurability, meaning that they do not define a set of random variables. This makes SCMs in the cyclic setting a notoriously more difficult class than the class of acyclic SCMs.

In this work (Bongers et al., 2016) we will give a general treatment of SCMs, where we deal with the measure-theoretic complications that arise in the presence of cyclic relations. Our approach is to strip off the random variables from the definition of the SCM, which gives a well-defined class of SCMs under any perfect intervention. We give conditions under which these SCMs induce a well-defined set of (measurable) random variables. We will take these SCMs as the natural objects from which we derive all the other causal and statistical properties. Moreover, we will show that if one is only interested in a particular subset of the endogenous variables, one can, under some conditions, arrive at a more parsimonious representation of the SCM, by means of marginalizing and reducing the space of exogenous variables.

Intuitively, the idea behind *marginalizing* over a subset of interest is that we consider this subset of endogenous variables as a subsystem, which can be treated as a "black box" that can interact with the rest of the system, see Figure 1. Thereby, we completely remove the representation of the internals of the subsystem, preserving only the essential input-output characteristics of it. This is somewhat analogous to marginalizing a probability distribution on a set of variables down to a distribution on a subset of those variables, hence the name "marginalization", which in the causality literature may also be known as a "latent projection" (Verma, 1993). We show that if the part of the causal mechanism that describes the causal relations in the subsystem satisfies a certain unique solvability condition (intuitively: it gives a unique output for any possible input), we can effectively remove

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

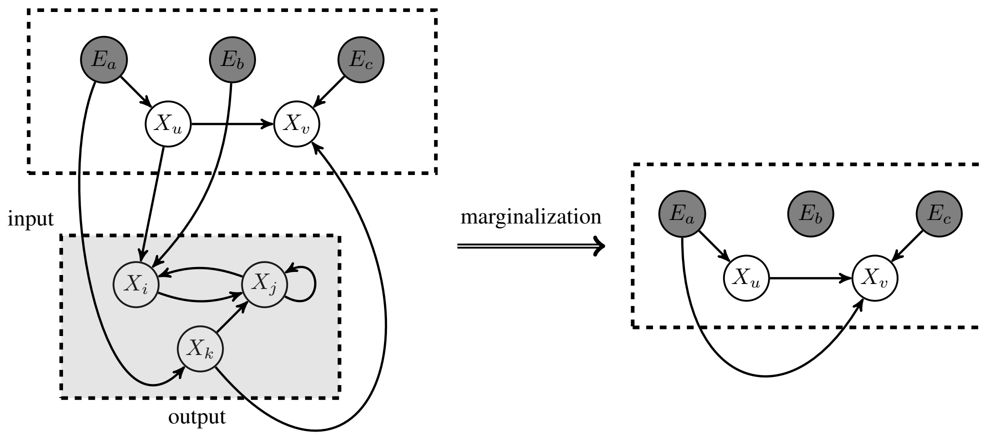

Figure 1: Example of a subsystem of an SCM on the left, depicted as a black box, which for a given input has a unique output. Marginalizing over this subsystem gives the marginal SCM on the right. Exogenous variables are depicted as filled circles, endogenous variables as blank circles.

this subsystem of endogenous variables from the model, treating it as a black box. An important property of this marginalization operation is that it preserves the causal semantics, meaning that the interventional distributions induced by the SCM are identical to those induced by its marginalization.

The marginalization operation acts only on the endogenous part of the system and does not touch the exogenous part. This may still lead to very complex models even for systems with only a few endogenous variables. The question of how to represent the influence of the surroundings of the system of interest in a more parsimonious way leads us to introduce the concept of *reduction*, see Figure 2. A reduction of an SCM is an interventionally equivalent SCM (i.e., all interventional distributions induced by the original SCM are identical to those induced by the reduction) defined on the same space of endogenous variables but with a lower-dimensional space of exogenous variables. For certain classes of SCMs, such reductions can be obtained by performing a suitable *exogenous reparametrization*. This has presumably been known for a long time for linear SCMs (Hyttinen et al., 2012) but we are not aware of any general treatment for possibly nonlinear and cyclic SCMs. Interestingly, for SCMs with real-valued variables, it turns out that one can always find an exogenous reparametrization that reduces the exogenous variables down to a single one-dimensional real-valued exogenous variable. However, in general such heavily reduced SCMs should be impossible to estimate from data, as their causal mechanisms are typically very wild.

We observe that when starting with a linear SCM, a reduction exists that is also linear, and hence can be estimated relatively easily from data. Also, when starting with a Markovian SCM, reductions exist that will typically be at least as smooth as the original SCM. These observations suggest that "smooth" reductions, which are easier to estimate from data, might exist in general. However, we provide a counterexample that shows that this is actually not the case. This result may have implications for approaches to estimating SCMs from data.

In summary, we give a rigorous mathematical treatment of SCMs, where we dealt with the measure-theoretic complications that arise in the cyclic setting. We take SCMs as the starting point from which we derive various causal and statistical properties. We show that one can arrive at a more parsimonious representation of the SCM if one is only interested in a certain subset of the endogenous variables, by means of marginalization and reduction (possibly achieved by an exogenous reparameterization).

Acknowledgments

Stephan Bongers and Joris Mooij were supported by NWO, the Netherlands Organization for Scientific Research (VIDI grant 639.072.410 and VENI grant 639.031.036). Joris Mooij was also supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement n$^{\circ}$ 639466).

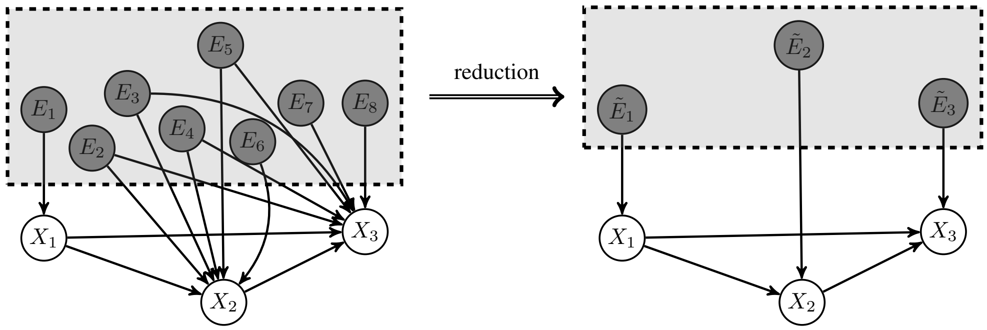

Figure 2: Example of an SCM (left) where the space of exogenous variables, depicted by the box, can be reduced to a reduced SCM (right).

References

Bongers, S., Peters, J., Schölkopf, B., and Mooij, J. M. (2016). Structural causal models: Cycles, marginalizations, exogenous reparametrizations and reductions. *arXiv.org preprint*, arXiv:1611.06221 [stat.ME].

Hyttinen, A., Eberhardt, F., and Hoyer, P. (2012). Learning linear cyclic causal models with latent variables. *Journal of Machine Learning Research*, 13:3387–3439.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition.

Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. The MIT Press, Cambridge, Massachusetts, 2nd edition.

Verma, T. (1993). Graphical aspects of causal models. Technical Report R-191, Computer Science Department, University of California, Los Angeles.