Statistical Decidability in Confounded, Linear Non-Gaussian Models

Throughout the 1990s, causal discovery algorithms were largely *constraint-based*; such methods inferred causal structure exclusively on the basis of conditional independence relationships among the observed variables [Spirtes et al., 2000]. Beginning in the early 2000s, it was realized that algorithms could harness other observable features of the underlying joint distribution (e.g., that noise is additive) to infer causal structure [Hoyer et al., 2009, Peters et al., 2012, Loh and Bühlmann, 2014]. For our purposes, especially impressive advances have been made in the study of linear, non-Gaussian causal models (LiNGAMs) [Shimizu et al., 2011, Hoyer et al., 2008]. When the true model is LiNGAM, then the causal structure among all measured variables can be learned in the limit [Shimizu et al., 2006]. With further assumptions, variants of maximum-likelihood estimation are *uniformly* consistent [Bühlmann et al., 2014].

Our main result is that the orientation of every edge in a LiNGAM is *statistically decidable*, even in the presence of confounding variables. This result improves upon previous findings by (1) [Hoyer et al., 2008] and [Salehkaleybar et al., 2020] who show that edge orientation in LiNGAMs is *identifiable* in the presence of confounding and (2) [Genin and Mayo-Wilson, 2020], who show that orientation is statistically decidable in the absence of confounding.

Statistical decidability is a reliability concept that is stronger than consistency (simpliciter) but weaker than uniform consistency [Genin, 2018]. A set of models is statistically decidable if, for any $\alpha > 0$, there is a consistent procedure that, *at every sample size*, hypothesizes a false model with chance less than α. Roughly, a statistical decision procedure exists when one can produce shrinking $(1-\alpha)$-confidence sets around the true parameter. Uniform consistency is the stronger requirement that one be able to determine the sample size *a priori* at which one's chances of identifying the true model (to a certain degree of approximation) are at least $1-\alpha$; statistical decidability requires no such pre-experimental guarantees.

Although our result may appear to be a mere technical improvement upon the work of [Hoyer et al., 2008, Salehkaleybar et al., 2020], it may have profound implications for methodology in the biomedical sciences. In particular, our main result raises serious questions about standard arguments for randomized controlled-trials (RCTs), which are considered by many as the "gold standard" of evidence for causal claims.

To defend RCTs, researchers often make two claims: (1) when an intervention on a variable X is performed, the causal effect of X on any other measured variable is *identifiable* because the influence of confounders is eliminate (at least in expectation) [Hernán and Robins, 2020] and (2) the causal effect of one variable on another is, in general, not identifiable from passive observation alone, even if confounding is assumed not to exist. In this context, identifiability means that if two models M_1 and M_2 disagree about the effect that X has on some variable Y, then M_1 and M_2 assign different probabilities to the possible data sets that might be observed when an intervention on X is performed. In general, a partition of statistical models is called identifiable if, whenever two models M_1 and M_2 belong to different partition cells, they assign different probability distributions to the observable data.

Identifiability is a *necessary* condition for discovery: if identifiability fails, two distinct models may be indistinguishable no matter how much data is collected, and those models may provide radically different answers to some question of interest. But identifiability is far from sufficient, either in a theoretical sense or practical one. When a set of statistical models is identifiable, there may be no statistically *consistent* procedure [Gabrielsen, 1978], let alone a statistical decision procedure in the sense we have described above. Thus, our main result suggests that, in the absence of stronger reliability guarantees, RCTs cannot be defended as superior to passive observation if researchers have good reason to believe the underlying model is LiNGAM.

References

Peter Bühlmann, Jonas Peters, and Jan Ernest. CAM: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42(6):2526–2556, 2014. Publisher: Institute of Mathematical Statistics.

Arne Gabrielsen. Consistency and identifiability. *Journal of Econometrics*, 8 (2):261–263, October 1978. ISSN 0304-4076. doi: 10.1016/0304-4076(78) 90035-0. URL https://www.sciencedirect.com/science/article/pii/0304407678900350.

Konstantin Genin. *The Topology of Statistical Inquiry.* Doctoral Dissertation, Carnegie Mellon University, Pittsburgh, PA, 2018.

Konstantin Genin and Conor Mayo-Wilson. Statistical Decidability in Linear, Non-Gaussian Models. In *NeurIPS 2020:Workshop on Causal Discovery and Causality-Inspired Machine Learning*, 2020.

Miguel A. Hernán and James M. Robins. *Causal Inference: What If.* Chapman & Hall/CRC, Boca Raton, December 2020. URL https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/.

Patrik O. Hoyer, Shohei Shimizu, Antti J. Kerminen, and Markus Palviainen. Estimation of causal effects using linear non-Gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49(2): 362–378, 2008. Publisher: Elsevier.

Patrik O. Hoyer, Dominik Janzing, Joris M. Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, pages 689–696, 2009.

Po-Ling Loh and Peter Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *The Journal of Machine Learning Research*, 15(1):3065–3105, 2014. Publisher: JMLR. org.

Jonas Peters, Joris Mooij, Dominik Janzing, and Bernhard Schölkopf. Identifiability of causal graphs using functional models. *arXiv preprint arXiv:1202.3757*, 2012.

Saber Salehkaleybar, AmirEmad Ghassami, Negar Kiyavash, and Kun Zhang. Learning Linear Non-Gaussian Causal Models in the Presence of Latent Variables. *Journal of Machine Learning Research*, 21(39):1–24, 2020.

Shohei Shimizu, Patrick O. Hoyer, Aapo Hyvarinen, and Antti J. Kerminen. A linear non-Gaussian acyclic model for causal discovery. *The Journal of Machine Learning Research*, 7:2003–2030, 2006.

Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O. Hoyer, and Kenneth Bollen. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *The Journal of Machine Learning Research*, 12:1225–1248, 2011. Publisher: JMLR. org.

Peter Spirtes, Clark N. Glymour, and Richard Scheines. *Causation, Prediction, and Search*. The MIT Press, 2000.