Stability-based causal discovery: an adversarial approach

Topic: Causal Inference, Neural Network Identifiability of Causal Graph
Team: AO/TAU, CNRS − INRIA − LISN, Université Paris-Saclay
Advisors: Michèle Sebag (sebag@lri.fr)
Alessandro Leite (alessandro.leite@inria.fr)
Duration: 6 months
Location: LISN, Paris-Sud University – Building 660 – Shannon
Level: Master M2

1 Context & Motivation

Critical applications require understanding the underlying causal mechanism to enable decision-makers to integrate them into their high-stake decisions. Likewise, knowing the cause or the causal structure often helps in interpreting the outputs of the system. As randomized control trials (RCT), the golden standard to establish the causal relationship from data, can be forbidden due to ethical or practical difficulties, increasing attention is paid to observational causal discovery [1].

The challenges are manifold w.r.t. Machine Learning. Control and treated samples (e.g. patients receiving a drug) might follow different distributions (treatment bias); estimating the average treatment effect then faces a problem of prediction with missing values (one only knows one outcome, depending on whether the patient receives the drug or not) not at random. Another issue is due to the presence of hidden confounders, possibly biasing the causal relationships among the observed variables. A third issue is that the sought models are learned in a small n large p context (few samples, many variables). Lastly, the causal theory focuses on the identifiability property (existence and uniqueness) of the sought model.

The state of the art involves diverse methods, depending on the assumptions done (e.g., no unobserved confounders; linear models) and the considered methods, e.g. based on (conditional) independence tests and/or generative models optimized w.r.t. a causal loss. The internship will build upon an adversarial causal model (Structural Agnostic Modeling, SAM [2]) that adapts the famed Generative Adversarial Network framework [3] to jointly learn the distribution of each variable, conditional to its causes (Markov kernel).

2 Goal

Two goals will be considered. The first one consists in adapting SAM to the unobserved confounder case. Formally, each Markov kernel is expressed as:

$$X_i \sim f_i(Pa(X_i), \epsilon_i)$$

with f_i a causal mechanism (neural net), $Pa(X_i)$ the causes (parents) of X_i and ϵ_i a random noise independent of any variable (the known unknowns).

The presence of hidden confounders can be modeled as the fact that two variables $X_i \neq X_j$ share the same noise variable. The causal loss will be extended to allow several variables to share the same noise input; and regularise the structure to enforce a sparse confounder graph.

A more ambitious goal is to revisit the identifiability requirement. In the case where the causal graph $\mathcal{G}$ can be used as a generative model, an interesting question is: if I learn $\mathcal{G}$ from some dataset, and if I use $\mathcal{G}$ to generate new data $\mathcal{D}'$, will I recover the same causal graph $\mathcal{G}$ from these new data ?
If this is actually the case, the stability of the algorithm used to learn $\mathcal{G}$ is (empirically) established, paving the way toward

3 Profile

The internship requires some curiosity about causality, besides excellent machine learning skills and programming expertise, preferred in Python.

References

[1] M. J. Vowels, N. C. Camgoz, and R. Bowden, "D'ya like dags? a survey on structure learning and causal discovery," *ACM Computing Surveys*, vol. 55, no. 4, pp. 1–36, 2022.

[2] D. Kalainathan, O. Goudet, I. Guyon, D. Lopez-Paz, and M. Sebag, "Structural agnostic modeling: Adversarial learning of causal graphs," *Journal of Machine Learning Research*, vol. 23, no. 219, pp. 1–62, 2022.

[3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.