October 2019

**Causality for Program Synthesis**

A 2 years post-doc position is available at TAU, INRIA Saclay,
in collaboration with LACODAM, INRIA Rennes
within the HyAIAI (Hybrid Approaches for Interpretable AI) Inria Project Lab

**Contact**: Send CV, publications, and references to Marc dot Schoenauer at inria dot fr

Most ML studies are concerned with learning predictive models, that allow to predict the value of some variables, based on the values of others. However, this amounts to learn 'only' correlations: another approach to understanding and explaining a complex system is to discover the causal relationships between variables. While randomized controlled experiments will remain the gold standard to determine causal relationships, such experiments are hampered due to practical, economical, or ethical reasons. In such cases, determining causal relationships from observational data is of primary importance [PM18, PJS17].

A first approach to pairwise causal relation learning has been proposed as an ML task: given enough examples of pairs for which the relation is known (independance, causality, or presence of confonders), it amounts to image classification, considering the available data for both variables as a 2D images [GSB19]. However, too few examples (pairs of variables with known causality relation) are available except in some very speficic domains, like biological regulatory networks, where controlled experiments are possible. Unfortunately, what can be learned from these examples tends to behave poorly on completely different domains. On the other hand, domain adaptation has now reached a mature state, transferring knowledge from one domaine to another (sufficiently related) one, e.g. using adversarial learning approaches [GUA+16]. One idea forward is to apply domain adaptation to causation learning, in order to adapt the causation model learned from biological networks to other domains.

Another possible direction is related to recent advances in causal learning in the TAU team, namely SAM [KGG+18]. SAM also uses the idea of adversarial learning. Considering $D$ variables, SAM builds $D$ neural nets, each aimed to learn the $i$-th variable from all other variables (and a noise variable). The set of these NNs is used to reconstruct full 'fake' samples, that are jointly learned with a discriminator. Some slack variables enforce the sparsity of the NNs, under the assumption that the true underlying causal model achieves a best trade-off between data-fitting and structural simplicity. SAM discovers the full causality graph at once, without the need to start from pairwise causation strengths and conditional independence. However, in the case where latent confounders control several variables, these would need to be identified several times. Hence another direction is to consider a brand new neural architecture, inspired from auto-encoders, that would ensure that all confounders are found and prevent the discovery of trivial relations, while nevertheless imposing sparsity.

In both research directions above, a crucial issue is that of existing domain knowledge: in many domains, some partial causality information is available, that it would be a waste to try to rediscover. Incorporating such (symbolic) information is key to incremental learning of causality, and hence to scalability. In particular, constraints due to time-dependencies are often present, that can greatly limit the search for causal links. Incorporating such knowledge is yet another research direction.

In particular, time constraints are crucial in the use case proposed by the LACODAM team that is proposed in this work, aiming to investigate the practical interest of causality discovery when providing user assistance in the everyday use of a computer, that involves many repetitive tasks. Reasearch in the field of Program Synthesis have already proposed ways to automatically detect simple repetitive tasks and assist the user in their completion (ass e.g., in the FlashFill functionality of Microsoft Excel [Gul11] The goal here will be to extend this approach to a more general use of multiple applications at once. Repetitive actions can be detected thanks to pattern mining techniques [YHA03, TV12] applied to traces of user actions. However, to propose a relevant assistance to the user, it is especially important to determine *precisely* the events/conditions that "trigger" an often repeated behavior. Using causality detection techniques will allow to filter the most likely causes of repetitive behaviors, and propose more accurate suggestions to end users. This application is rich with domain knowledge, both from the application and from the user preferences, that can be used to help scale up causality learning. The LACODAM team is already building a trace collection system and will provide data and pattern extraction techniques.

Depending on the experience of the candidate, one or several above research directions will be tackled during the two proposed two years.

## References

[GSB19] Isabelle Guyon, Alexander Statnikov, and Berna Bakir Batu, editors. *Cause Effect Pairs in Machine Learning.* The Springer Series on Challenges in Machine Learning. Springer Verlag, 2019.

[GUA+16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.

[Gul11] Sumit Gulwani. Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*, pages 317–330, 2011.

[KGG+18] Diviyan Kalainathan, Olivier Goudet, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Structural agnostic modeling: Adversarial learning of causal graphs, 2018.

[PJS17] J. Peters, D. Janzing, and B. Schoölkopf. *Elements of Causal Inference - Foundations and Learning Algorithms.* MIT Press, 2017.

[PM18] Judea Pearl and Dana Mackenzie. *The Book of Why. The New Science of Cause and Effect.* Basic Books, 2018.

[TV12] Nikolaj Tatti and Jilles Vreeken. The long and the short of it: summarising event sequences with serial episodes. In *18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 462–470, 2012.

[YHA03] Xifeng Yan, Jiawei Han, and Ramin Afshar. Clospan: Mining closed sequential patterns in large datasets. In *Proceedings of the Third SIAM International Conference on Data Mining, San Francisco, CA, USA, May 1-3, 2003*, pages 166–177, 2003.