

Causal Domain Adaptation: extending causal effect from a domain to another with generalization bounds

September 29, 2026

1 Context and description

Out-of-distribution (OOD) generalization is a central challenge in machine learning and application fields: a model trained on data from one distribution may be deployed in another setting, but previously predictive associations can become spurious [Ben-David et al., 2010, Zhou et al., 2021]. To address this challenge, we consider the notion of environments, which correspond to distinct data-generating processes. These environments can represent different experimental conditions, populations, time periods, or data-generating regimes.

Causal perspectives provide a principled framework for tackling OOD generalization in both settings. Under a structural causal model (SCM, [Pearl, 2009]), where each variable is determined by a functional relationship with its parents and an exogenous noise term, an environment corresponds to a specific intervention or distribution shift on the exogenous variables or mechanisms.

2 Internship Goals

During the internship, **the objectives** will be to investigate the problem of the transportability of causal mechanisms, that is, extending a causal effect from one domain to another, and to provide generalization bounds. The aim is to take into account

- the *covariate shift*, i.e., the shift in the marginal distribution of the explanatory variables from one domain to another;
- the so-called *graphical shift*, accounting for the fact that the causal graphs may differ from one domain to another; d
- the *mechanism shift*, accounting for the fact that the structural equations (or conditional distributions $P(Y \mid \text{Pa}(Y))$) associated with the nodes of the causal graph may themselves change from one domain to another, even when the graph structure is preserved.

Relevant examples illustrating these different cases can be found for example when analyzing multi-hospital / multi-site medical data. Indeed, several studies, such as [Wang et al., 2025] compare treatment effects (e.g., ICU interventions, sepsis treatment protocols) across hospitals with different patient demographics, equipment, and clinical practices. Mechanism shift here could be different hospital protocols (e.g., how quickly a treatment is escalated) even if the same variables are recorded.

A concrete example of graphical shift can be found in studies of targeted temperature management (TTM) after cardiac arrest, where treatment assignment is randomized in a clinical trial (HYPERION) but chosen by the clinician based on the patient's condition in observational databases (eICU-CRD, MIMIC-IV) [Wang et al., 2025]. This is exactly a change in the causal graph: in the randomized trial, no edge connects patient severity to treatment, since randomization guarantees that treatment assignment is independent of patient characteristics; in the observational registries, however, an edge from severity to treatment appears, since the physician's decision is informed by how sick the patient is. The causal graph structure therefore differs between the source and target domains, even though

the same variables (severity, treatment, outcome) are recorded in both — a textbook case of what the transportability literature formalizes as a selection diagram [Bareinboim and Pearl, 2013].

Ranging from covariate shift to graphical shift possibly passing by mechanism shift, the intern shall understand step by step what is their impact on generalization bounds.

3 Candidate profile

The ideal candidate is passionate about artificial intelligence and motivated by applied research issues. He has a solid background in machine learning, statistics, applied mathematics, computer science, modeling, and optimization, and good scientific writing and oral communication skills. Proven experience and taste in computer programming (Python) and data analysis are expected. Past experience in Causal Inference would be a major advantage.

Applicants must be in their final year of Master (M2 or Msc) or last year of engineering school (3A) in one of the following fields: Applied Mathematics, Machine Learning, AI, Computer Science, Statistics or related fields. A good command of the English language will be appreciated.

4 Supervision team, Funding and location

The internship will be in an academical and industrial research environment as it is the result of a collaboration between the [Eki.Lab](#) (the AI research lab of [Ekimetrics](#)), the CRAN (Centre de Recherche en Automatique de Nancy, Université de Lorraine and CNRS) and the the LIG (Laboratoire d’Informatique de Grenoble, Université Grenoble Alpes).

The internship duration will be **5 or 6 months**, with an expected start in February or March 2027. The remuneration will be higher than a classical university internship gratification. More details will be given during the interview. The intern will officially be hired by the CRAN. However, the internship does not necessarily need to take place in Nancy, Paris could also be an option.

Finally, this internship offers the **opportunity to pursue a Ph.D. in causal inference** within these research labs.

5 Application

Application must be sent to emilie.devijver@univ-grenoble-alpes.fr, marianne.clausel@univ-lorraine.fr and audrey.poinsot@ekimetrics.com as soon as possible and must include:

- Cover letter (can be the email itself);
- Resume / CV ;
- GitHub link containing projects completed during their studies (optional but recommended)
- Transcript for the most recent academic year

Incomplete applications will not be considered.

References

- [Bareinboim and Pearl, 2013] Bareinboim, E. and Pearl, J. (2013). A general algorithm for deciding transportability of experimental results. *Journal of causal Inference*.
- [Ben-David et al., 2010] Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. (2010). A theory of learning from different domains. *Machine Learning*.

- [Pearl, 2009] Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition.
- [Wang et al., 2025] Wang, A., Wentzlof, K., Rajala, J., Green, M., Zhang, Y., and Yang, S. (2025). Transfer learning for individualized treatment rules: Application to sepsis patients data from eicu-crd and mimic-iii databases.
- [Zhou et al., 2021] Zhou, K., Yang, Y., Qiao, Y., and Xiang, T. (2021). Domain generalization with mixstyle.