Causal Discovery in Heterogenous Environments Under the Sparse Mechanism Shift Hypothesis

Ronan Perry[1]
Julius von Kügelgen*[1,2],
Bernhard Schölkopf*[1],
(* shared last author).

[1]Max Planck Institute for Intelligent Systems, Tübingen, Germany; [2]University of Cambridge

Motivation

- **Causal graphs are typically unidentifiable** from *i.i.d.* data without additional functional constraints.
- Heterogeneous datasets contain **natural distribution shifts** which provide additional information akin to interventions.
- We seek to understand **how we can identify the true causal graph** and **under what distribution shifts** this is possible.

Causal graphs and mechanisms

- Consider a causal system over the random variables $\mathbf{X} = \{X_1, ..., X_d\}$.
- X_i is a stochastic function of its direct causes PA_i which are encoded as the parents of X_i in a directed acyclic graph (DAG) G over $\mathbf{X}$.
- G induces the Markov factorization:

$$\mathbb{P}(\mathbf{X}) = \prod_{j=1}^{d} \underbrace{\mathbb{P}(X_j \mid \mathsf{PA}_j)}_{\text{Causal mechanism}}$$

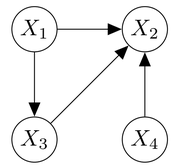

Figure 1. Example DAG

- **Causal discovery goal:** learn the causal DAG (from observational data).
- **Identifiability:** from *i.i.d.* data, methods can only learn a *Markov Equivalence Class (MEC)* of graphs.

Multi-environment data

- We may have multiple datasets $\mathcal{D}^e$ on $\mathbf{X}$ from different environments $e \in \{1, \dots, n_\mathcal{E}\}$ with differing distributions $\mathbb{P}^e_\mathbf{X}$.
- Assume an unobserved base CGM $\mathcal{M} = (G, \mathbb{P}_\mathbf{X})$.

Assumption [Shared mechanisms]: Each environment e independently results from $\mathcal{M}$ by intervening on an (unknown) subset $\mathcal{I}^e \subseteq [d]$ of mechanisms.

$$\mathbb{P}^e_\mathbf{X}(X_1, ..., X_d) = \left(\prod_{j \in \mathcal{I}^e} \underbrace{\mathbb{P}^e_\mathbf{X}(X_j|\mathsf{PA}_j)}_{\text{Changed mechanism}} \right) \prod_{j \in [d] \setminus \mathcal{I}^e} \underbrace{\mathbb{P}_\mathbf{X}(X_j|\mathsf{PA}_j)}_{\text{Base mechanism}} .$$

Definition [Augmented graph]: Let E be a binary random variable indexing a pair of environments. Augment the causal graph G with vertex E and an edge from E to any variable whose causal mechanism changes across those two environments.

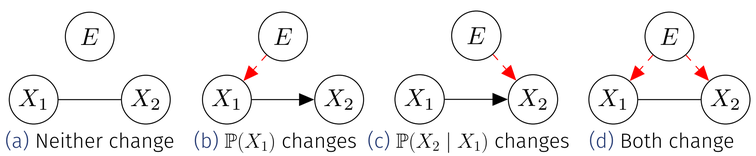

(a) Neither change (b) $\mathbb{P}(X_1)$ changes (c) $\mathbb{P}(X_2 \mid X_1)$ changes (d) Both change

Figure 2. Example augmented graphs

Contributions

- We introduce the **Mechanism Shift Score** causal discovery algorithm.
- We prove convergence rates and **identifiability guarantees** under enough sparsely changing, heterogeneous environments.
- We demonstrate empirical efficacy of the MSS, and flexibility to accommodate existing parametric and **nonparametric** estimators.

Assumptions

- **Causal faithfulness:** distributional independence implies graphical independence.
- **Independent causal mechanisms [1]:** a change in one causal mechanism has no effect on and provides no information on changes to any other causal mechanisms.
- **Pseudo causal sufficiency [2]:** causal sufficiency holds, conditional on an environment.
- **Sparse mechanism shifts (SMS) [3]:** only a sparse number of mechanisms change between any two environments.
- **Additional technical details of note**
  (i) Interventions are *soft*, i.e. an intervention does not change the graph skeleton.
  (ii) The causal graph is the same across all environments.
  (iii) No additional distributional assumptions are made.

Mechanism Shift Score (MSS)

- **Intuition:** an incorrect graph in the MEC $\mathcal{G}_{\text{MEC}}$ will have more mechanisms shifts than the true causal graph due to wrong edge directions.
- MSS counts the number of pairwise mechanism changes in graph G:

$$\text{MSS}(G) = \sum_{j=1}^{d} \sum_{e'>e}^{n_\mathcal{E}} \mathbb{I}\left[\mathbb{P}^e(X_j|\mathsf{PA}^G_j) \neq \mathbb{P}^{e'}(X_j|\mathsf{PA}^G_j)\right]$$

- Our **estimand** is the set $\mathcal{G}^{\min}_{\text{MEC}} := \arg\min_{G \in \mathcal{G}_{\text{MEC}}} \text{MSS}(G)$

Main theoretical result

Theorem [Rate of identifiability]: Let G^* be the true DAG

$$\Pr[\mathcal{G}^{\min}_{\text{MEC}} = \{G^*\}] \geq 1 - |\mathcal{G}_{\text{MEC}}|(1-\beta)^{\lfloor \#\text{ environments}/2 \rfloor}.$$

where $\beta \in [0, 1)$ depends on the probabilities of mechanisms shifting.

Corollary [Asymptotic identifiability]: If the probability of each mechanism shifting is bounded away from 0 and 1,

$$\Pr[\mathcal{G}^{\min}_{\text{MEC}} = \{G^*\}] \to 1 \quad \text{as} \quad \#\text{ environments} \to \infty$$

Significance: with enough sparsely changing environments, with high probability we learn exactly the true graph, not just the MEC.

Experiments

Simulation: Random DAGs with non-linear additive noise distributions.

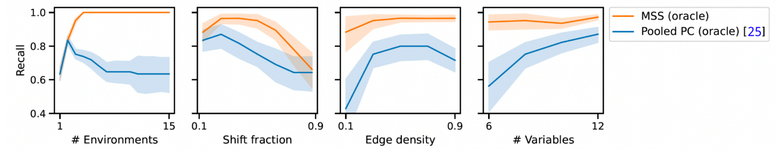

Figure 3. Under an oracle (perfect) test, MSS empirical results match the theory. Pooling all data [2] does not yield identifiability.

Estimator(s): The MSS may use any empirical test of equality of conditional distribution to count changes, parametric [4] or non-parametric [5].

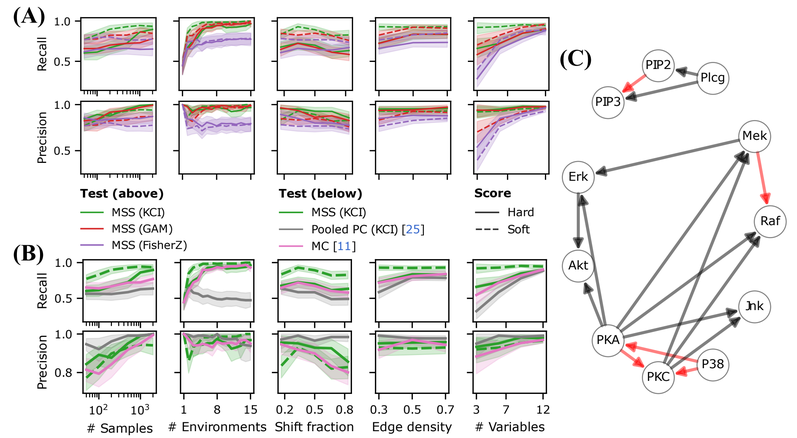

Figure 4. **(A)** The KCI [6] test performs best among considered MSS estimators. **(B)** MSS using the KCI test outpeforms competing methods in recall, while maintaininig good precision. **(C)** Promising MSS results on the Sachs protein network dataset [7].

Paper: [see QR code] `https://arxiv.org/abs/2206.02013`
Code: `https://github.com/rflperry/sparse_shift`

References

[1] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference*. The MIT Press, 2017.
[2] Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *JMLR*, 21(89):1–53, 2020.
[3] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *IEEE - Advances in Machine Learning and Deep Neural Networks*, 109(5):612–634, 2021.
[4] Amir Emad Ghassami, Negar Kiyavash, Biwei Huang, and Kun Zhang. Multi-domain Causal Structure Learning in Linear Systems. In *NeurIPS*, volume 31. Curran Associates, Inc., 2018.
[5] Christina Heinze-Deml, Nicolai Meinshausen, and Jonas Peters. Invariant causal prediction for nonlinear models. *JCI*, 6(2):1–35, 2018.
[6] Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *UAI*, page 804–813. AUAI Press, 2011.
[7] Karen Sachs, Omar Perez, Dana Pe'er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.