Universal Causal Inference in a Topos

Sridhar Mahadevan (Adobe Research, UMass Amherst)

NeurIPS 2025

Universal Causality (Mahadevan, Entropy, 2023)

Pollution in New Delhi, India

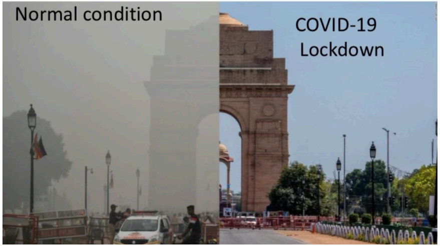

Categorical Framework

- **Objects:** Variables or Causal Models or Sheaves
- **Arrows:** Interventions or Observations
- **Diagrams:** Functors like pullback ($\bullet \rightarrow \bullet \leftarrow \bullet$) that map to concrete causal model.

$$\boxed{\mathcal{C}(\mathrm{T}, \mathrm{P}) \cong \mathbf{Nat}(\mathcal{C}(-, \mathrm{T}), \mathcal{C}(-, \mathrm{P}))}$$

$$\boxed{\textbf{Yoneda embedding: } \mathcal{C} \mapsto \mathbf{Set}^{\mathcal{C}^{op}}}$$

$$\underbrace{P(\text{Pollution} \mid \text{Traffic})}_{\text{observational}} \neq \underbrace{P(\text{Pollution} \mid \mathrm{do}(\text{No-traffic}))}_{\text{interventional}}$$

What is a Topos Causal Model?

- A **topos** is a category that is *Cartesian closed*, has a *terminal object*, and has a *subobject classifer*.
- A structural causal model $\mathcal{M}$ (SCM) defines a unique function $F : U \rightarrow V$ from *exogenous* variables into *endogenous variables*
- SCMs forms a topos, where each object is an SCM model $\langle U, V, F \rangle$, and arrows are given by *commutative diagrams*:

$$
\begin{tikzcd}
U \arrow[r, "h"'] \arrow[d, "f"] & U' \arrow[d, "f'"'] \\
V \arrow[r, "g"] & V'
\end{tikzcd}
$$

Three Examples of Topos Causal Models

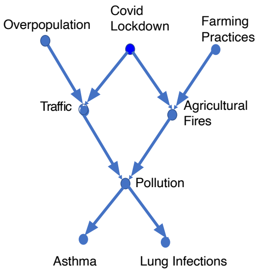

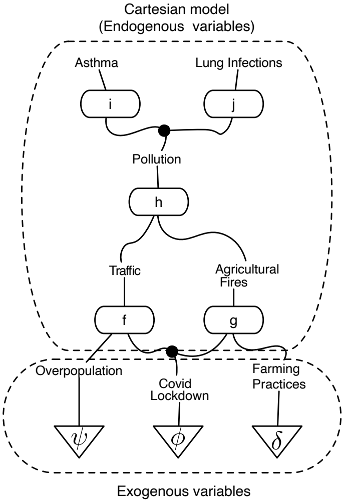

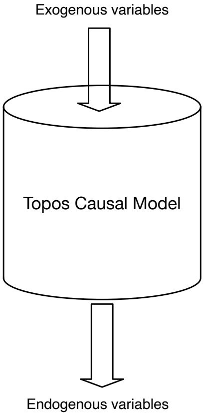

Universal Property of Topos Causal Models

1. A causal model is a functor that maps from a structure domain category to a semantic co-domain category
2. **Structure Category:** Examples include symmetric monoidal categories with a "copy-delete" comonoidal structure on each object (aka Markov category)
3. **Semantics Category:** Examples include **Prob**, where objects are measurable spaces, and arrows are measure-preserving maps.

Theorem

Any causal functor $F : \mathcal{C} \rightarrow \mathcal{E}$ from a structural causal category $\mathcal{C}$ (such as a Markov category) to a semantic cocomplete category $\mathcal{E}$ (such as **Prob***) factors uniquely through a TCM structure defined by the Yoneda embedding.*

Causal Interventions as Subobject Classifiers

$$\begin{tikzcd}
S \arrow[r] \arrow[d, tail, "m"] & \mathbf{1} \arrow[d, tail, "\mathbf{true}"] \\
X \arrow[r, dashed, "\phi"] & \Omega
\end{tikzcd}$$

- A **subobject classifier** is a $\mathcal{C}$-object Ω, and a $\mathcal{C}$-arrow $\mathbf{true} : \mathbf{1} \to \Omega$, such that to every monic arrow $S \hookrightarrow X$ in $\mathcal{C}$, there is a unique arrow ϕ that forms the above pullback square.
- Example: The Covid-19 causal intervention of **No-Traffic** produces a *subobject* of the original causal model of pollution in New Delhi, India.

Topos Causal Models define an internal intuitionistic logic

- The internal language of a TCM is intuitionistic (constructive): law of the excluded middle does not hold.
- The semantics of the logic is defined by a topology on the arrows:
  - *Grothendieck topology*: open sets map to *sieves*
  - *Lawvere-Tierney topology*: specified by a modal operator $j : \Omega \to \Omega$ on the subobject classifier that defines "local" truth.
- **Syntax:** defined by the Mitchell-Bénabou Language
- **Semantics:** defined by Kripke-Joyal possible worlds

Judo Calculus: Intuitionistic j-do-calculus in TCM

Characteristic	**Classical do-calculus**	**Judo Calculus**
Logic	Causal claims are globally true or false	Intuitionistic logic: truth is *local*
Context	Uses "average" treatment effect	Local truth is "glued" together
Interventions	"Surgery" of a causal graph	Subobject classifier
Identification	Axioms define three rules	More general axiomatic framework

Table: Some of the salient differences between classical do-calculus and judo calculus.

Sridhar Mahadevan, *Intuitionistic j-do-calculus for Topos Causal Models*, Arxiv

Rules of Judo Calculus

Each premise means: there exists a *j*-cover $\mathcal{S} = \{S_i \to U\}_i$ such that the stated CI holds on *every* chart S_i after the indicated graph surgery.

[*j*-Rule 1: insert/delete observations]

$$\Big(Y \perp Z \mid X, W \text{ in } \mathcal{G}_{\overline{X}} \text{ on a } j\text{-cover of } U\Big) \implies \mathrm{P}(y \mid \mathrm{do}(x), z, w) = \mathrm{P}(y \mid \mathrm{do}(x), w).$$

[*j*-Rule 2: action/observation exchange]

$$\Big(Y \perp Z \mid X, W \text{ in } \mathcal{G}_{\overline{X},\, \underline{Z}} \text{ on a } j\text{-cover of } U\Big) \implies \mathrm{P}(y \mid \mathrm{do}(x), \mathrm{do}(z), w) = \mathrm{P}(y \mid \mathrm{do}(x), z, w).$$

[*j*-Rule 3: insert/delete actions]

$$\Big(Y \perp Z \mid X, W \text{ in } G_{\overline{X},\, \overline{Z(W)}} \text{ on a } j\text{-cover of } U\Big) \implies \mathrm{P}(y \mid \mathrm{do}(x), \mathrm{do}(z), w) = \mathrm{P}(y \mid \mathrm{do}(x), w).$$

Decentralized Causal Discovery in Judo Calculus

- A significant advantage of judo calculus is that it is *sheaf*-based and highly decentralized.
- A J-cover $\mathcal{S} = \{\, V_i \hookrightarrow U \,\}_{i=1}^{E}$ turns a global causal discovery problem into E *independent* subproblems plus a light-weight aggregation. This matches a map–reduce pattern:

$$\underbrace{\textsc{Discover}(U)}_{\text{pooled}} \quad \rightsquigarrow \quad \{\, \underbrace{\textsc{Discover}(V_i)}_{\text{per-env/chart}} \,\}_{i=1}^{E} \ \text{ then } \ \underbrace{\textsc{Glue}(\{A_i\})}_{j\text{-aggregation}}.$$

- Preliminary Experiments with j-stable GES, ψ-FCI and DCDI show significant benefits of TCM framework.

Sridhar Mahadevan, *Decentralized Causal Discovery in Judo Calculus*, Arxiv

Experimental Results with TCM-enabled Methods

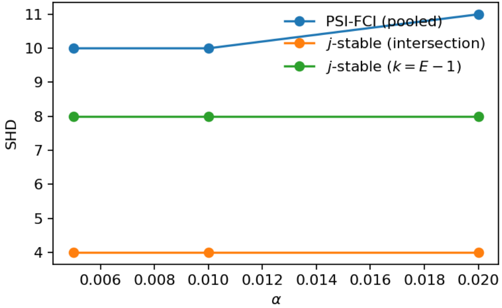

Figure: j-stable ψ-FCI outperforms pooled version by a wide margin.

Experimental Results with TCM-enabled Methods

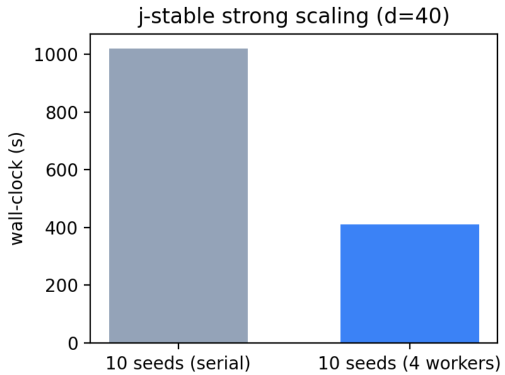

Figure: *j*-stable DCDI scales far better than standard DCDI.

Summary

- Topos Causal Models (TCM) introduce a new framework for causal inference.
- Causal interventions are modeled as *subobjects* and induce an intutionistic j-do calculus (aka "judo" calculus)
- Judo calculus has an axiomatic set of rules for drawing inferences.
- Preliminary experimental results show significant benefits of judo calculus over classical do-calculus
- Sridhar Mahadevan, "Intutionistic j-do-calculus for Topos Causal Models (Arxiv)
- Sridhar Mahadevan, "Decentralized Causal Discovery with Judo Calculus" (Arxiv)