Marginal Causal Flows for Validation and Inference

Dan Manela*, Laura Battaglia*, and Robin Evans

Department of Statistics, University of Oxford

UNIVERSITY OF OXFORD

DEPARTMENT OF STATISTICS

* equal contribution

Causal Inference

Causal inference practitioners are often interested in **estimating the effect of experimental interventions on an outcome of interest**.
Some example causal questions:

- What is the impact of giving Ozempic on the weight of the Danish population over a 6-month treatment period?
- What is the effect of a new website homepage on the conversion rate of visiting customers?
- How much does a school's average grade change if they change their syllabus?

Confounding

- Randomised Control Trials (RCTs) involve randomly assigning a treatment/intervention protocol to each candidate.
- This guarantees an **unbiased** estimate of the average treatment effect of a treatment over the population.
- However, in some cases it might not be possible or ethical to randomise data.
- If the data is **confounded**, it means that the treatment assignment for an individual is a function of that individual's properties.

Estimating treatment effects on confounded data is hard.

Confounding

- Randomised Control Trials (RCTs) involve randomly assigning a treatment/intervention protocol to each candidate.
- This guarantees an **unbiased** estimate of the average treatment effect of a treatment over the population.
- However, in some cases it might not be possible or ethical to randomise data.
- If the data is **confounded**, it means that the treatment assignment for an individual is a function of that individual's properties.

Estimating treatment effects on confounded data is hard.

Synthetic Benchmarks

Deluge of novel inference methods have hit the “market”. But how good are they really?

- Wittgenstein’s Ruler: *When using a ruler to measure an object, you also measure the accuracy of the ruler itself*.
- This problem exists in ML method development. Trust in our methods is limited by the quality of the synthetic experiments we run.

In Causal Inference, generating good synthetic data is hard.

Frugal Flows

We introduce **Frugal Flows**, a generative model which can target the **marginal causal effect** of an observational joint.

Key capabilities of Frugal Flows include:

- Ability to be trained on complex real-world datasets.
- Flexibility to customise causal properties of the generative model (e.g., the causal effect, confounding).
- Capability to easily be sampled from to generate synthetic benchmarks with *known* causal quantities.

How?

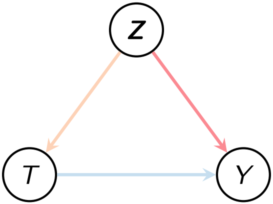

- We build on a frugal parameterization (Evans and Didilez, 2024) of the joint.
- The whole model can be trained using a composition of Normalising Flows.
- We can model the causal quantity independently of the rest.
- The joint dependencies between $Y|\mathrm{do}(X)$ and $\boldsymbol{Z}$ are modelled with a copula.
- The modular structure allows us to freely customise the causal quantity in the benchmark sampling stage.

How?

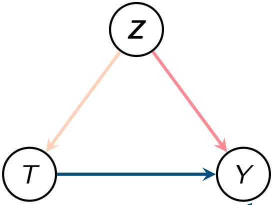

$$\theta_{Y|\mathrm{do}(T)} : F_{Y|\mathrm{do}(T)} = \mathcal{F}^{-1}_{Y|\mathrm{do}(T)}$$

- We build on a frugal parameterization (Evans and Didilez, 2024) of the joint.
- The whole model can be trained using a composition of Normalising Flows.
- We can model the causal quantity independently of the rest.
- The joint dependencies between $Y|\mathrm{do}(X)$ and Z are modelled with a copula.
- The modular structure allows us to freely customise the causal quantity in the benchmark sampling stage.

How?

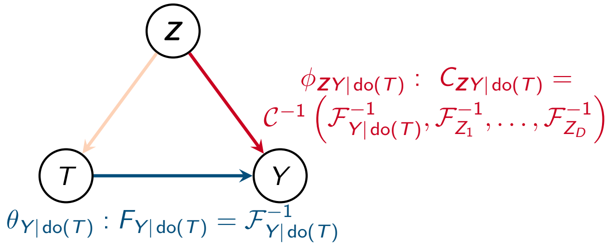

- We build on a frugal parameterization (Evans and Didilez, 2024) of the joint.
- The whole model can be trained using a composition of Normalising Flows.
- We can model the causal quantity independently of the rest.
- The joint dependencies between $Y|\text{do}(X)$ and $\boldsymbol{Z}$ are modelled with a copula.
- The modular structure allows us to freely customise the causal quantity in the benchmark sampling stage.

How?

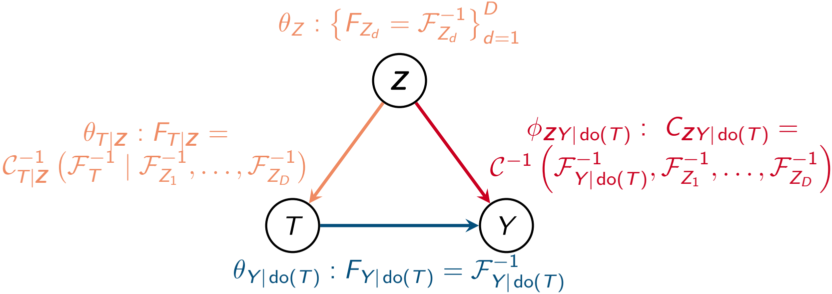

- The frugal parameterization (Evans and Didilez, 2024) of the joint allows us to model the causal quantity independently of the rest.
- The joint dependencies between $Y|\mathrm{do}(X)$ and $\boldsymbol{Z}$ are modelled with a copula.
- The whole model can be trained using a composition of Normalising Flows.
- The modular structure allows us to freely customise the causal quantity in the benchmark sampling stage.

Takeaways

Summary

- We parameterise a frugal single-treatment causal model using Normalizing Flows.
- We can fit a Frugal Flow on real-world data and learn a generative model on the observational joint, whilst targeting the marginal causal effect $p_{Y|\,\mathrm{do}(T)}$.
- One can modify causal properties of the fitted Frugal Flow and customise multiple causal features. Or you can just use the original model. Up to you...

Challenges

- Longitudinal models are hard to parameterise.
- Large datasets are required to accurately infer complex causal margins (hyperparameter tuning is complicated).

Takeaways

Summary

- We parameterise a frugal single-treatment causal model using Normalizing Flows.
- We can fit a Frugal Flow on real-world data and learn a generative model on the observational joint, whilst targeting the marginal causal effect $p_{Y|\,\mathrm{do}(T)}$.
- One can modify causal properties of the fitted Frugal Flow and customise multiple causal features. Or you can just use the original model. Up to you...

Challenges

- Longitudinal models are hard to parameterise.
- Large datasets are required to accurately infer complex causal margins (hyperparameter tuning is complicated).

Thank you!