Inferring Causal Direction from Observational Data: A Complexity Approach

Nikolaos Nikolaou[1] & Konstantinos Sechidis[2]

[1] Department of Physics & Astronomy, University College London
[2] Advanced Methodology & Data Science, Novartis Pharma AG

ECML PKDD 2020 Workshop on
Machine Learning for Pharma and Healthcare Applications (PharML) 2020

Causal Models from Observational Data

- **Identifying causal relationships** important in biology, medicine & pharmaceutics — and many other fields
  - To do so, ideally, we perform **randomized controlled trials** (**RCT**s)
  - Often impossible for **practical** / **ethical** reasons $\Longrightarrow$ must use **observational data**
- **Can we learn causal directionality from observational data?**
  - **No, if we just test for statistical independence** (most statistical / machine learning methods)
    Multiple causal structures can satisfy same set of statistical independences — e.g. given r.v.'s X & Y, $X \not\perp\!\!\!\perp Y$:

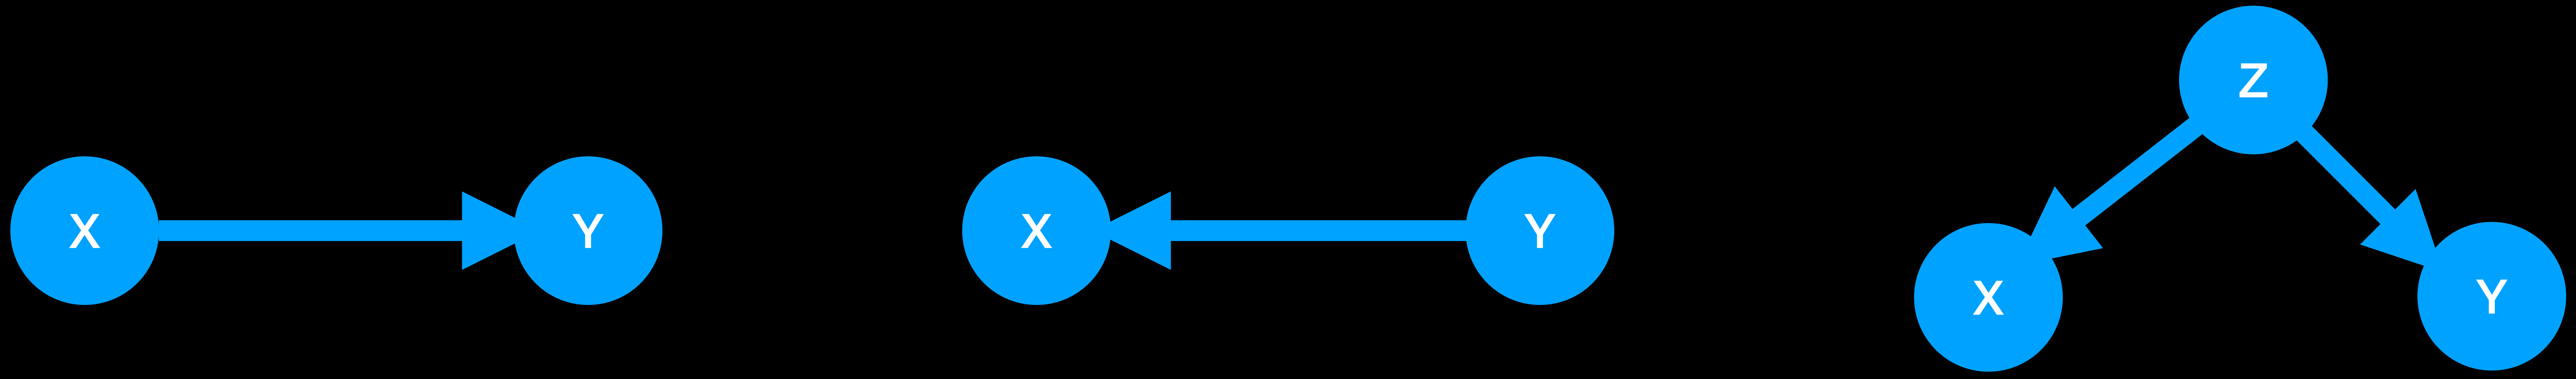

  - **Yes, if extra assumptions are made** (this work)

Complexity & Causal Directionality

- Question: Given r.v.'s X & Y, such that $X \not\perp Y$; Can we tell if $X \rightarrow Y$ or $Y \rightarrow X$?

- Central Idea — based on **Occam's Razor**: $Model_{Cause \rightarrow Effect}$ should be '**simpler**' than $Model_{Effect \rightarrow Cause}$

- Application:
    1. Train 2 models: one using X as feature to predict target Y & one using Y as feature to predict target X
    2. Measure the '**complexity**' of the 2 models
    3. Predict causal direction as the one used in the model with lowest complexity

- **Measures of complexity:**
  Several can be used — e.g. if models are decision trees of unbounded depth, can use:

    - **Tree Depth** ('**simpler**' = tree with **smaller depth**)

    - **Residual Entropy** ('**simpler**' = model resulting in **highest decrease in entropy of target variable**)

    - **Interpolation Hardness** ('**simpler**' = model exhibiting **best fit**)

Results

<u>Results of criteria under additive noise & uniformly distributed r.v.'s</u>

Avg. accuracy of using each measure to predict causal direction (with & w/o abstaining)

& avg. value of complexity measure for model respecting causal direction and model that doesn't

	Discrete			Continuous		
Criterion	Depth	Entropy	Fitting	Depth	Entropy	Fitting
Accuracy	0.988	0.974	0.986	0.583	0.976	0.990
Accuracy excl. Abstaining	0.995	0.986	0.998	0.665	0.978	0.997
Avg. Complexity (Causal)	9.252	0.214	0.135	12.871	0.024	427.146
Avg. Complexity (Anti-causal)	37.417	0.819	0.882	14.314	0.113	920.909

Notes:

Results obtained on **artificial data**; for full details on underlying true causal model, please consult the paper

Similar results for **discrete & continuous r.v.'s** (but depth poor measure for continuous)

gaussian, uniform or mixed r.v.'s

additive or multiplicative noise

All code is available at: https://github.com/nnikolaou/CausalDirectionality

Results

<u>Results of criteria under additive noise & uniformly distributed r.v.'s:</u>

Difference in complexity between 2 models, along with **predicted & true causal direction**

Discrete:

Continuous:

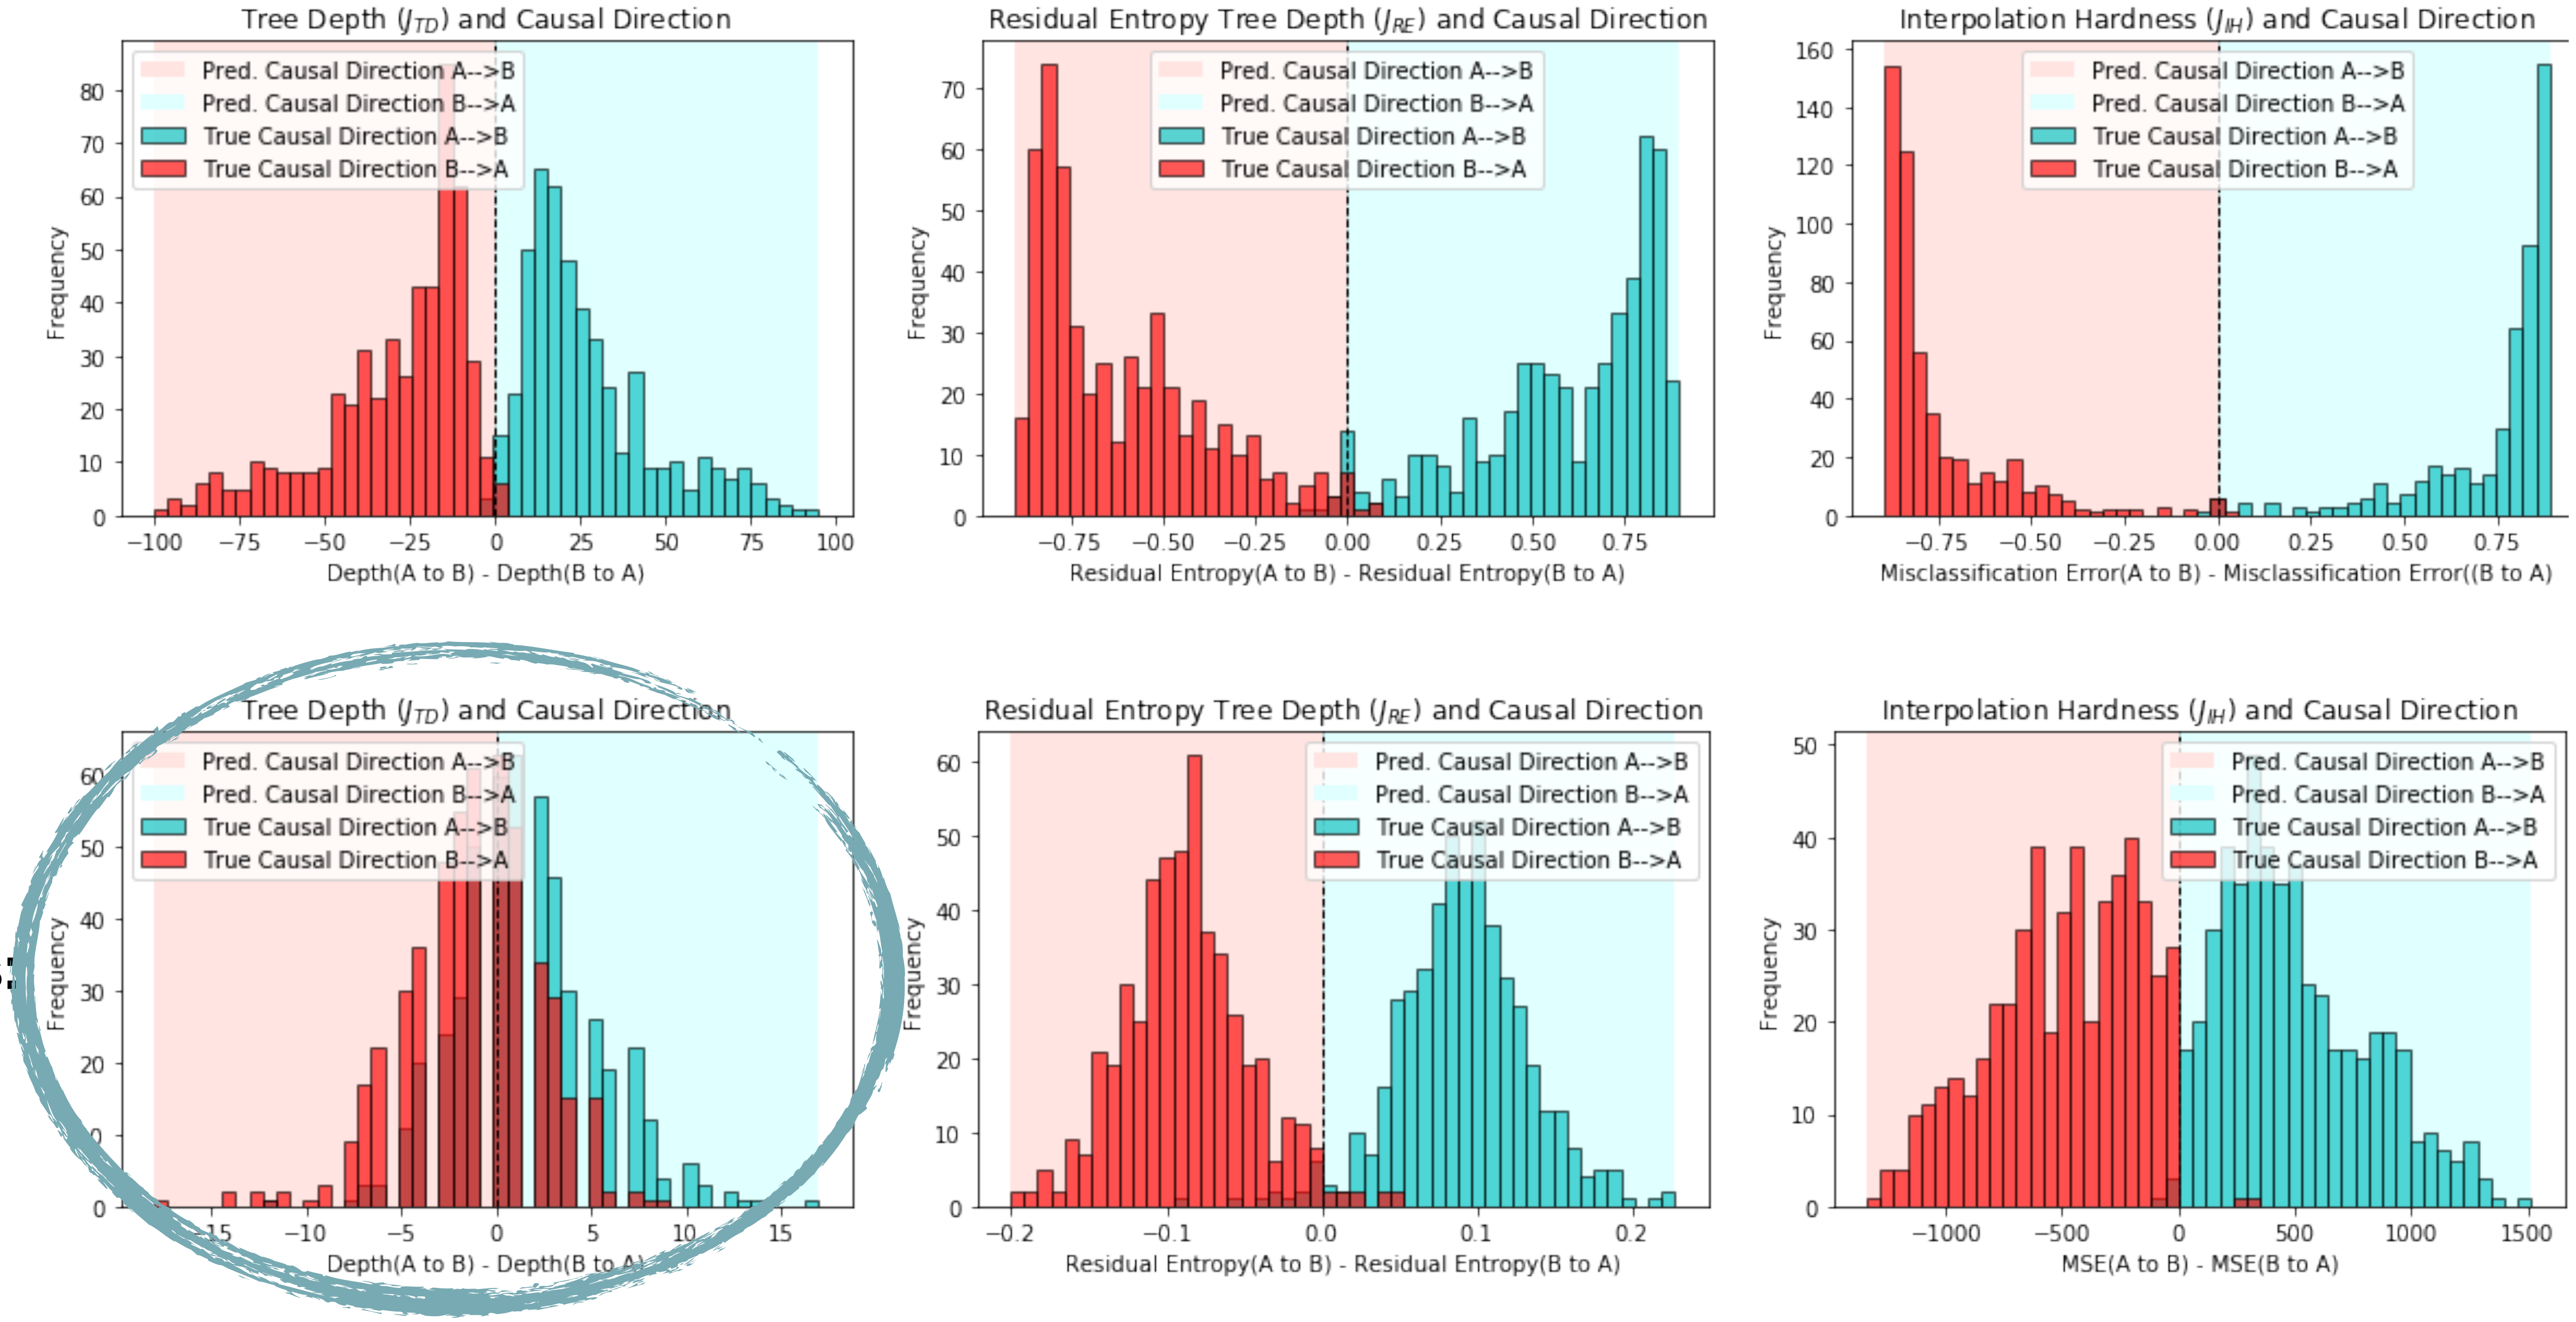

All code is available at: https://github.com/nnikolaou/CausalDirectionality

Take home message:

Can learn causal structure from observational data using fast, easy to code criteria.

Central idea: constructing a model of the data that respects the causal structure should be **simpler** than constructing one that doesn't.

Thank you!

Questions?

All code is available at: https://github.com/nnikolaou/CausalDirectionality