Autoregressive Variational Auto-Encoder for Causal Modeling

Tom Marty

Setting

Generated by a Structural Causal Model (SCM)

Input: multi-intervention dataset $\mathcal{D} = \{D_{\mathcal{O}}, D_{I_1}, \cdots, D_{I_N}\}$

$$f_i(\mathrm{Pa}(X_i), \epsilon_i) \rightarrow f_i^{\mathrm{I}}(\mathrm{Pa}(X_i), \epsilon_i) \quad \forall i \in \mathcal{F}_{\mathrm{I}}$$

Unknown regime and unknown target

Known DAG, but the problem remains hard

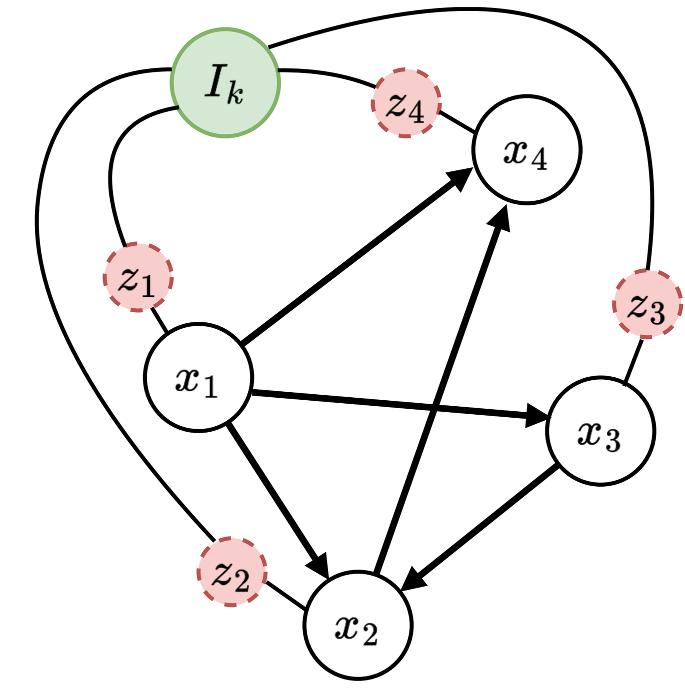

Goal:

- Model the mechanisms f_i under any regime:

$$X_i = f_i(\mathrm{Pa}(X_i), \epsilon_i) \qquad \epsilon_i \sim \mathbb{P}_{\epsilon_i}$$

- Identify the observation regime

Identification requires additionnal assumptions:

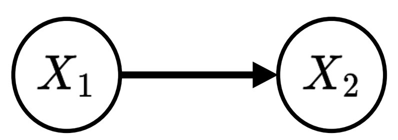

- Interventional-faithfulness / Strong enougth interventions

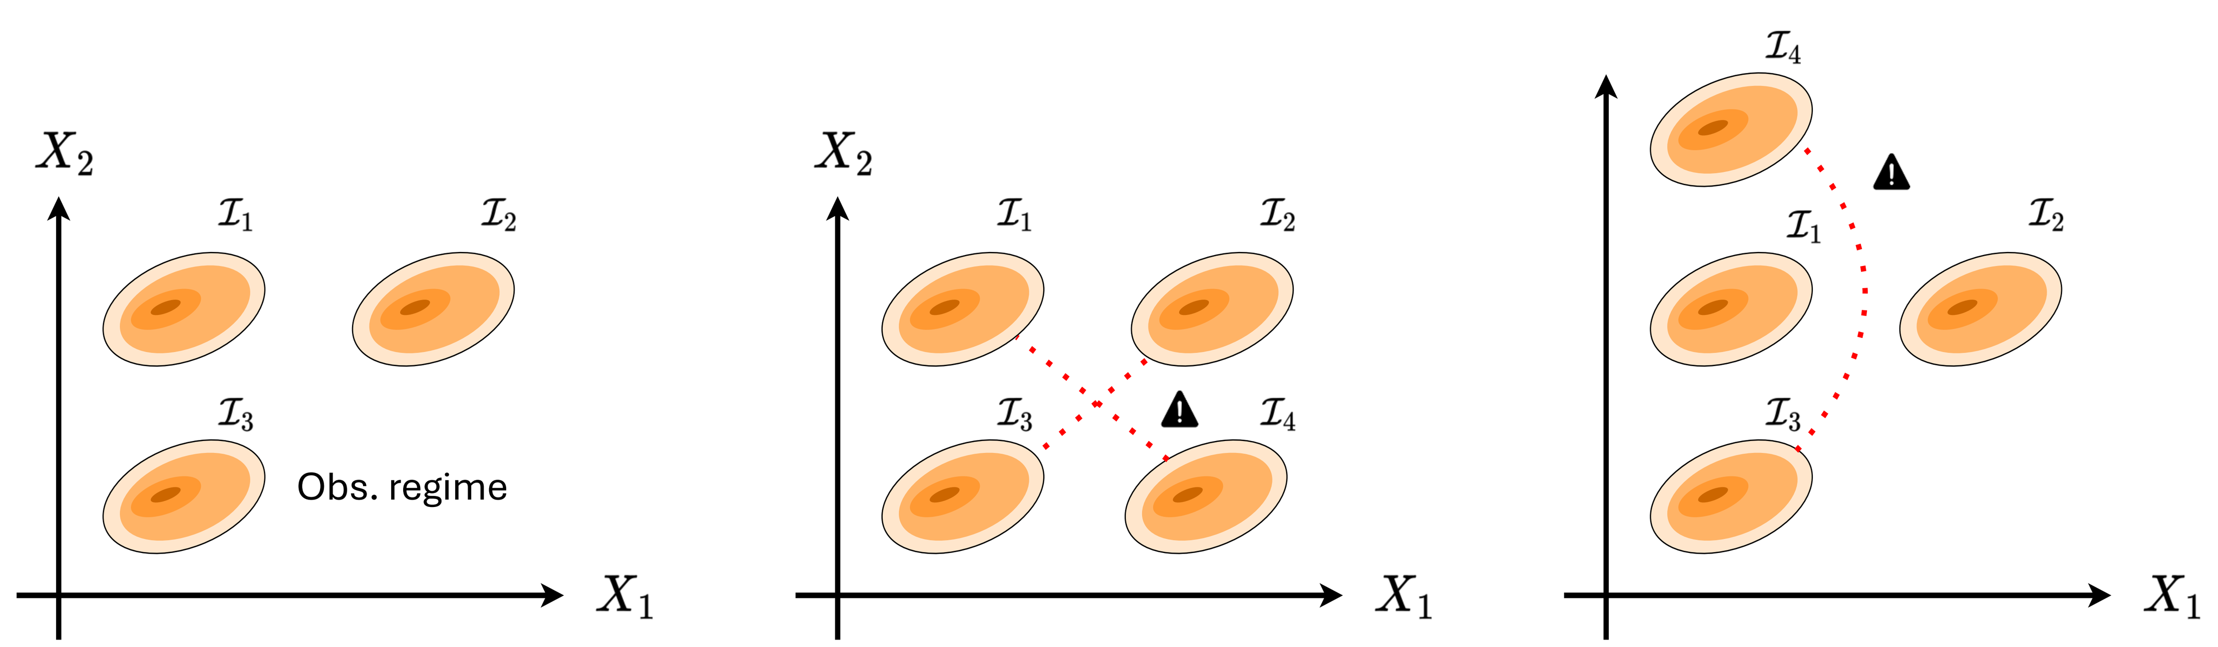

a) single target unique intervention

b) multiple target unique intervention

c) single target non-unique intervention

Experiments

SCM : linear additive noise model

- Autoregressive discrete VAE are hard to train... (*Fu et al.)
- ...but not impossible : latent encoding encodes the right regime on 60% of the dataset (no KL vanishing)
- As soon as the regime is identified, **VAE achieves perfect reconstruction of the conditionnal expectation**.

code report

For the full set of results ➡️ **Please read the report**

https://wandb.ai/dhanya-shridar/VariationalCM/runs/d3dtu3cu?nw=nwuser3rdcore

Method

Idea: formulate this problem as a discrete latent variable probleme where the latent variables are the unknown regimes.

Architecture: Variationnal Auto-Encoder (VAE)

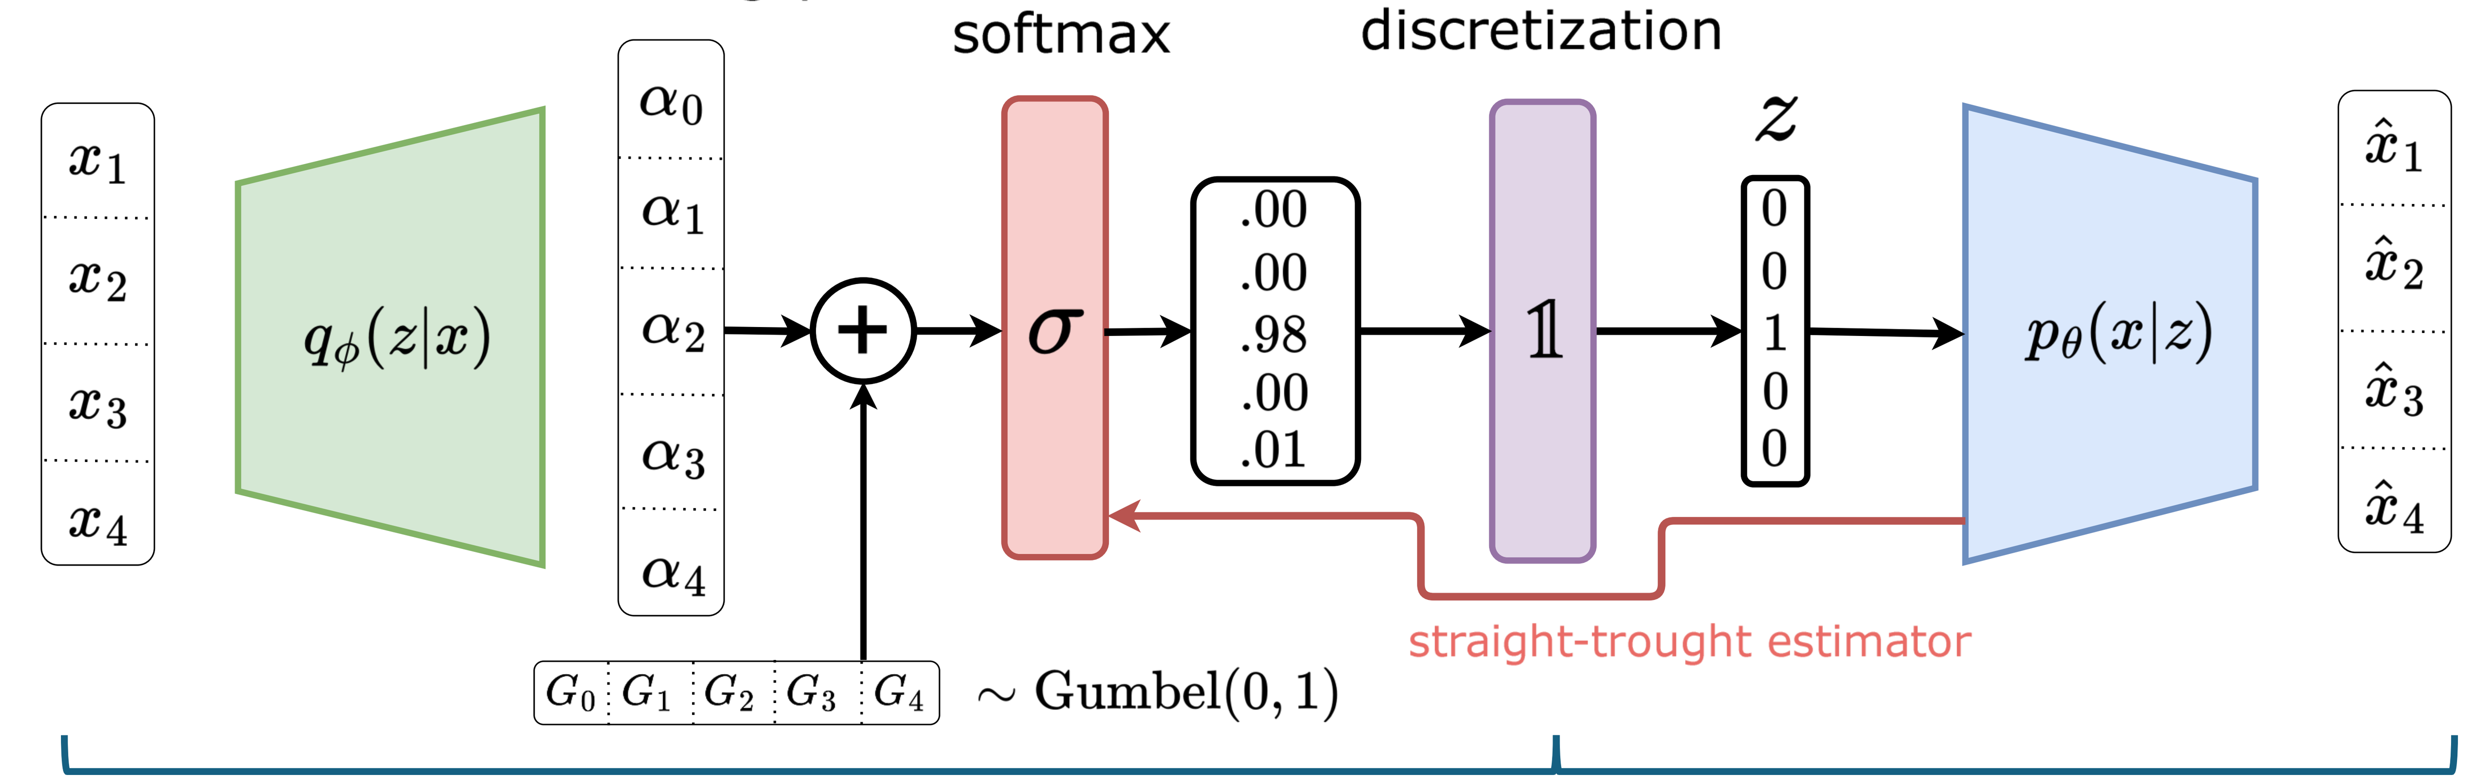

Regime learning (Categorical encoder)

mechanisms learning (Autoregressive decoder)

Objective : maximise data log-likelihood

$$\log p_\theta(x) \geq -\mathrm{KL}\big(q_\phi(z|x) \parallel p_\theta(z)\big) + \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)]$$

softmax(α) uniform prior $\mathcal{N}(f_\theta(Pa(x_i), z), 1)$

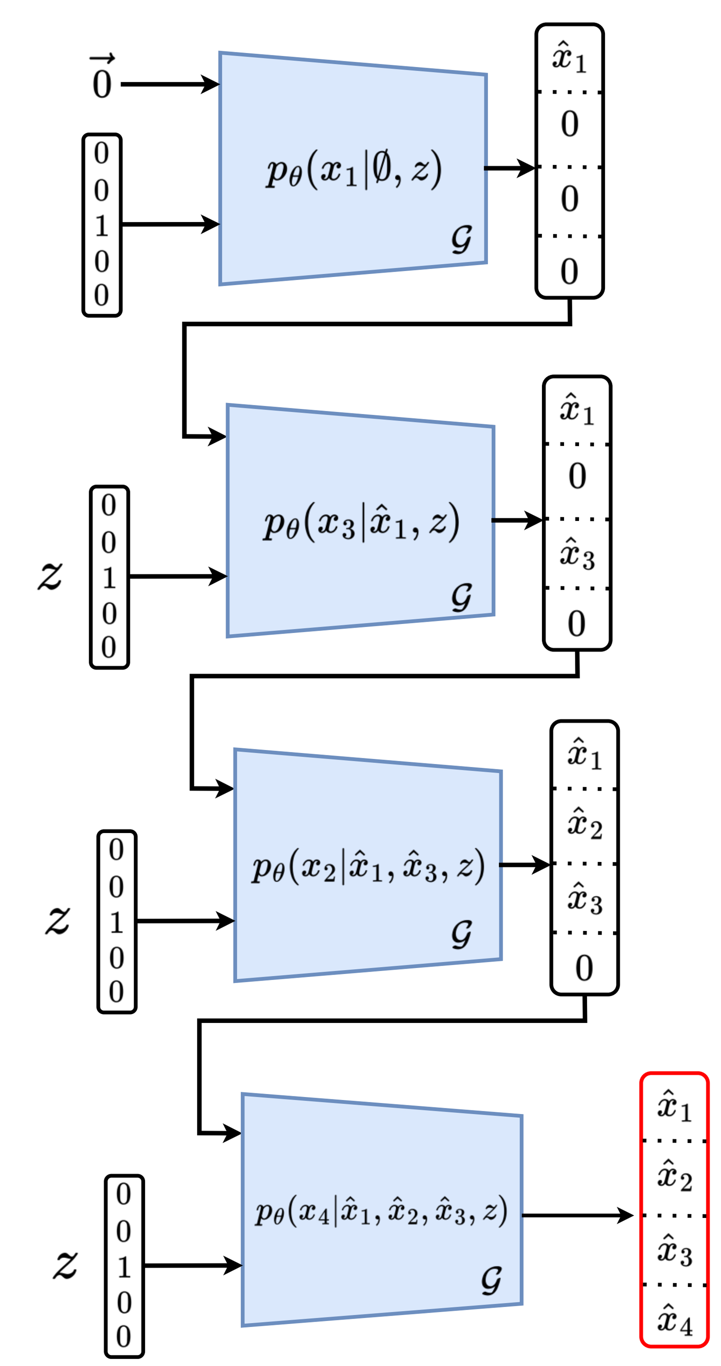

Regularization term :

- Maximise entropy of the posterior

Reconstruction term :

- Predict the full sample $\hat{x}$
- Can't learn to perfectly achieve reconstruction because of exogenous noise

Why not Expectation-Maximization ?
Closed form update does not always exist for realistic SCM

PyTorch Lightning

*Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, & Lawrence Carin. (2019). Cyclical Annealing Schedule: A Simple Approach to Mitigating KL Vanishing.