




UNIVERSITY OF AMSTERDAM
Faculty of Science




MIT-IBM
Watson
AI Lab






Causal Representation Learning for Instantaneous and Temporal Effects in Interactive Systems

Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M. Asano, Taco Cohen, Efstratios Gavves



Causal Representation Learning



- Given high-dimensional observations of a dynamical system, what are the true causal variables?
- Crucial for reasoning, planning, generalization, and more




Instantaneous Effects in Temporal Sequences

- Common assumption: time resolves causal effects
- But what about observations at low frame rates?

⇒ Instantaneous Effects!



time step t



time step $t + 1$

iCITRIS: Instantaneous Effects in Temporal Sequences

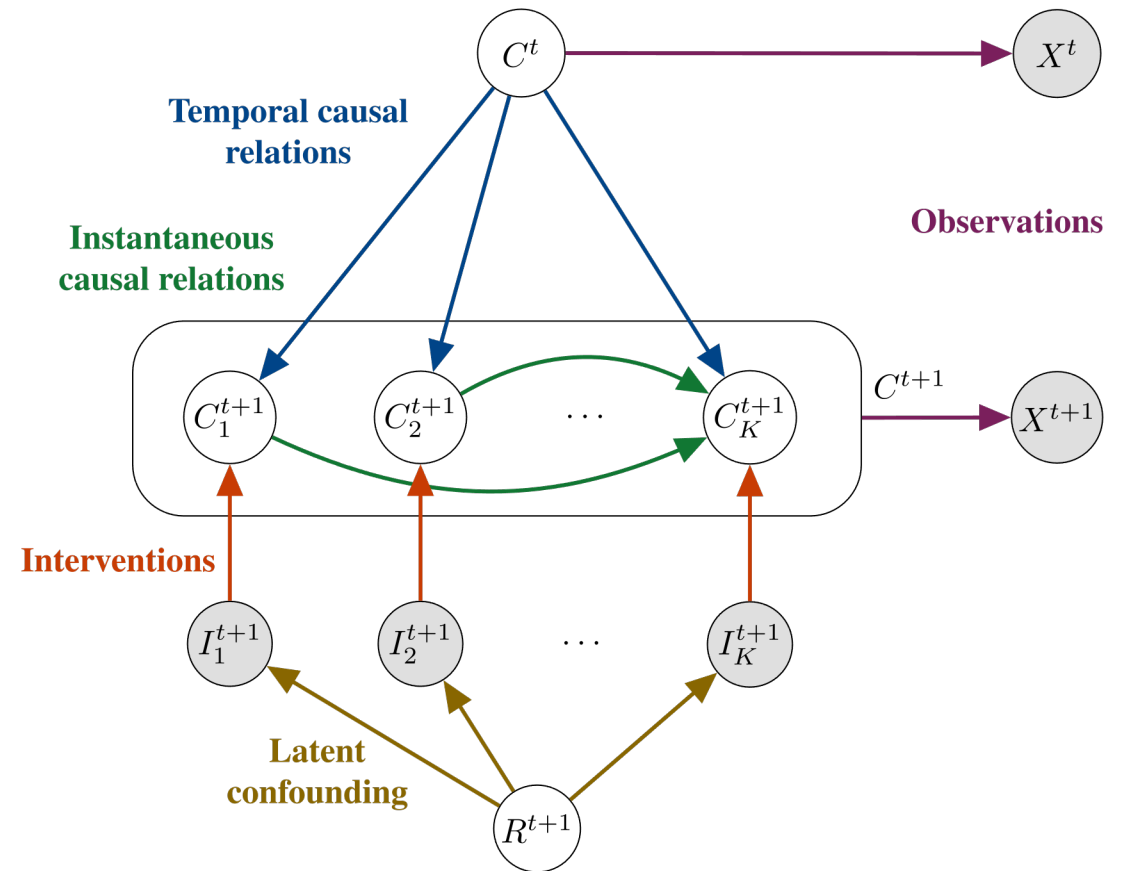
Setup

- Many more pitfalls, e.g.:

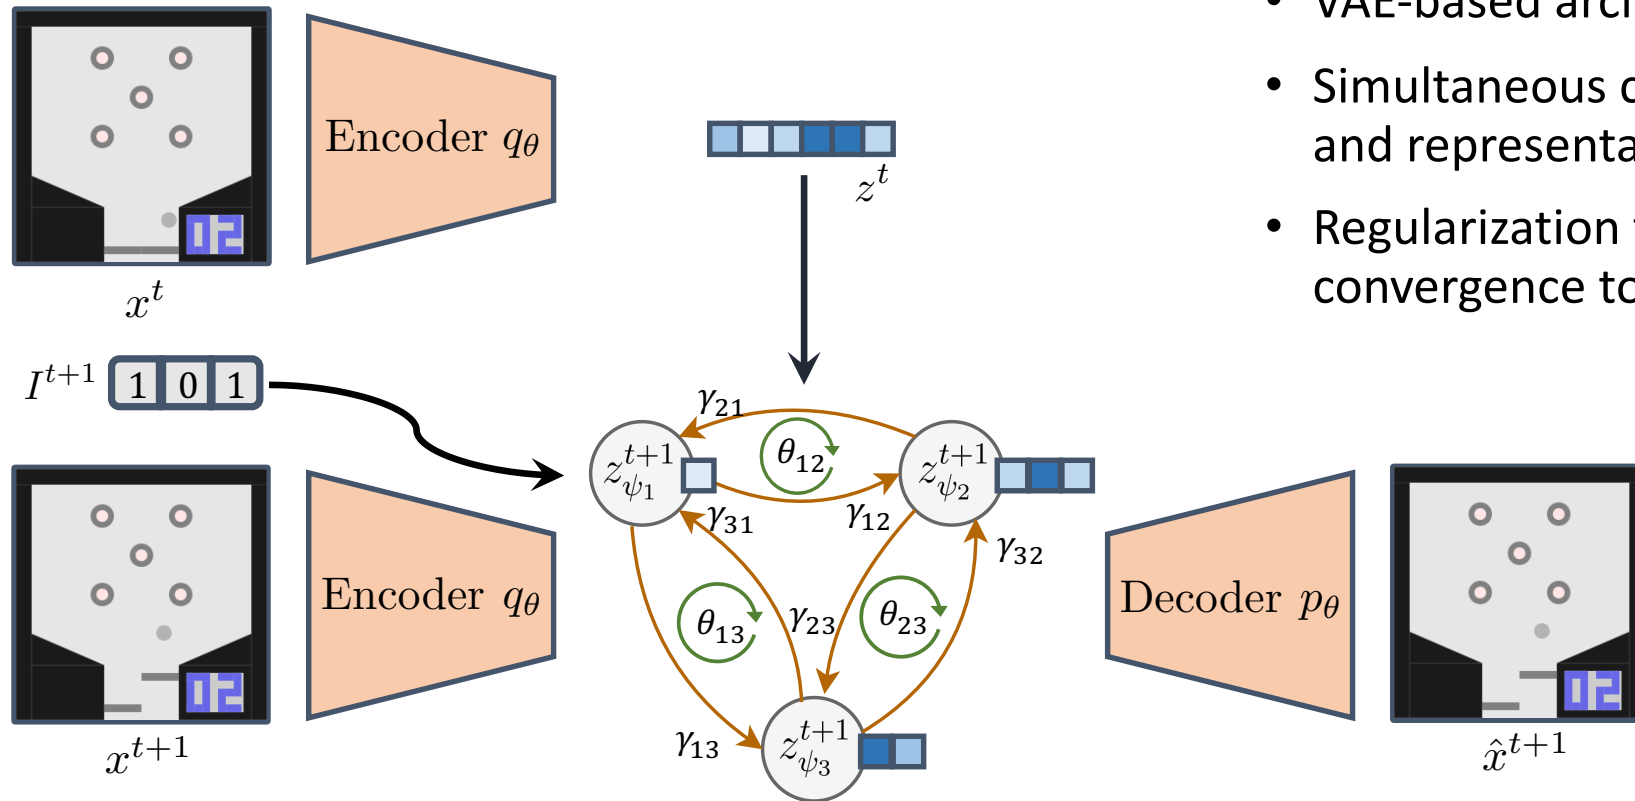
$$p_1(C_1)p_2(C_2) \text{ vs } p_1(C_1)\hat{p}_2(C_2 + C_1|C_1)$$

- Solution: *partially-perfect* interventions that remove instantaneous parents
 \Rightarrow Minimal causal variables [Lippe et al., 2022] become identifiable

- Chicken-and-egg situation:
 - Without graph, no causal variables
 - Without causal variables, no graph



iCITRIS Architecture






- VAE-based architecture
- Simultaneous causal discovery and representation learning
- Regularization to prevent early convergence to local minima

Experiments

Causal Pinball

Learned Instantaneous Causal Graphs



Model	R^2 (diag \uparrow / sep \downarrow)	SHD (instant \downarrow / temp \downarrow)
iCITRIS-ENCO	0.99 / 0.12	0.67 / 3.00
iCITRIS-NOTEARS	0.98 / 0.18	3.33 / 4.67
CITRIS	0.90 / 0.39	3.00 / 7.67
iVAE	0.44 / 0.05	4.33 / 4.67
iVAE-AR	0.47 / 0.15	8.00 / 3.67

Conclusion

- Causal Representation Learning tries to find the latent causal variables and their relations
- **Instantaneous Effects:** Effects that occur faster than the frame rate
- We prove the identifiability of causal variables under partially-perfect interventions
- **iCITRIS:** End-to-end learning of causal representations from temporal sequences
 - Joint causal discovery and causal representation learning
 - Regularization to prevent early convergence to local minima

Paper and code

