We *identify causal variables* and their *causal graph* from temporal sequences with *instantaneous effects*.

iCITRIS: Causal Representation Learning for Instantaneous Temporal Effects

Phillip Lippe[1] (p.lippe@uva.nl), Sara Magliacane[1,2], Sindy Löwe[1], Yuki M. Asano[1], Taco Cohen[3], Efstratios Gavves[1]
[1]University of Amsterdam, Informatics Institute, [2]MIT-IBM Watson AI Lab, [3]Qualcomm AI Research, The Netherlands

PROBLEM SETTING

- Causal effects faster than frame rate cause instantaneous effects
- Joint causal representation learning + causal discovery needed

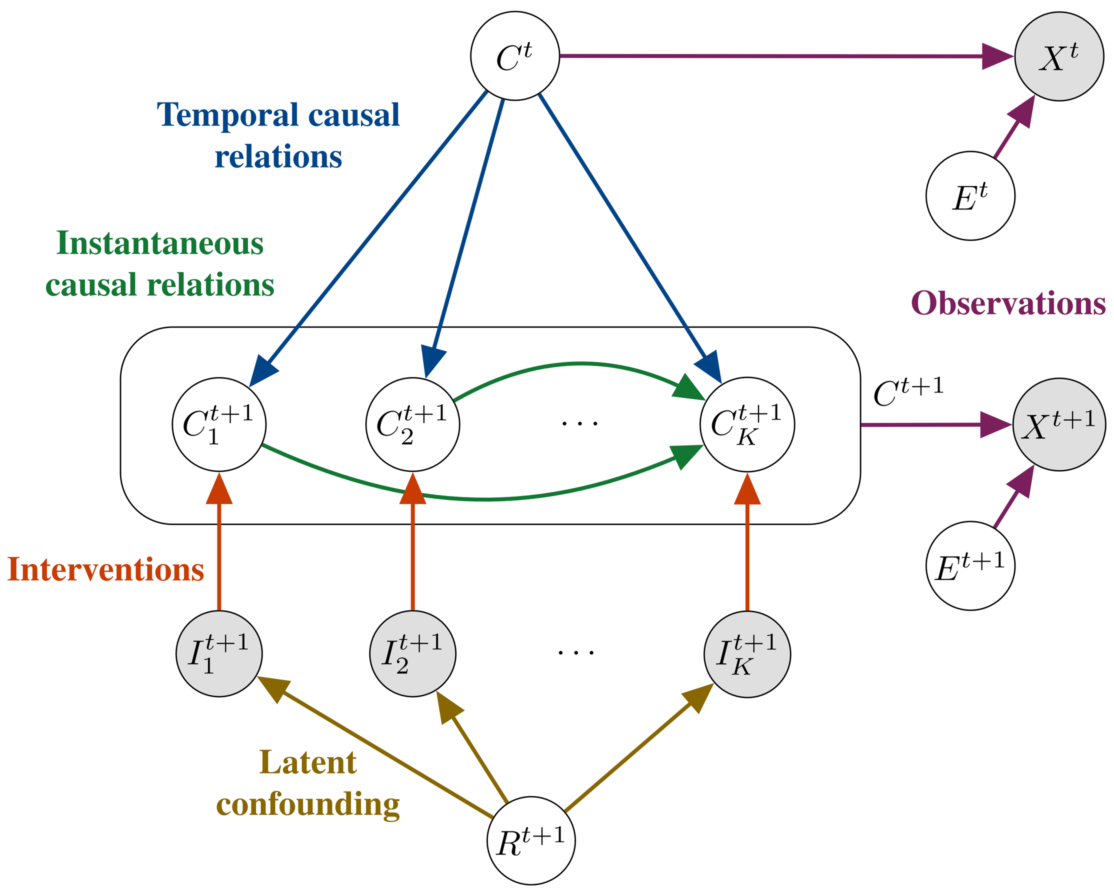

iCITRIS: INSTANTANEOUS EFFECT IN CRL

- Perfect interventions to distinguish instantaneous effects from entanglement in encoding function

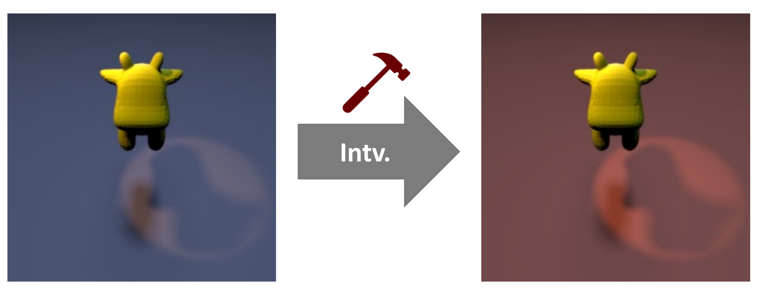

- Theoretical result: the minimal causal variables [1] are identifiable under a directed acyclic causal graph
- Extend CITRIS [1] with common causal discovery methods (ENCO [2] or NOTEARS [3])

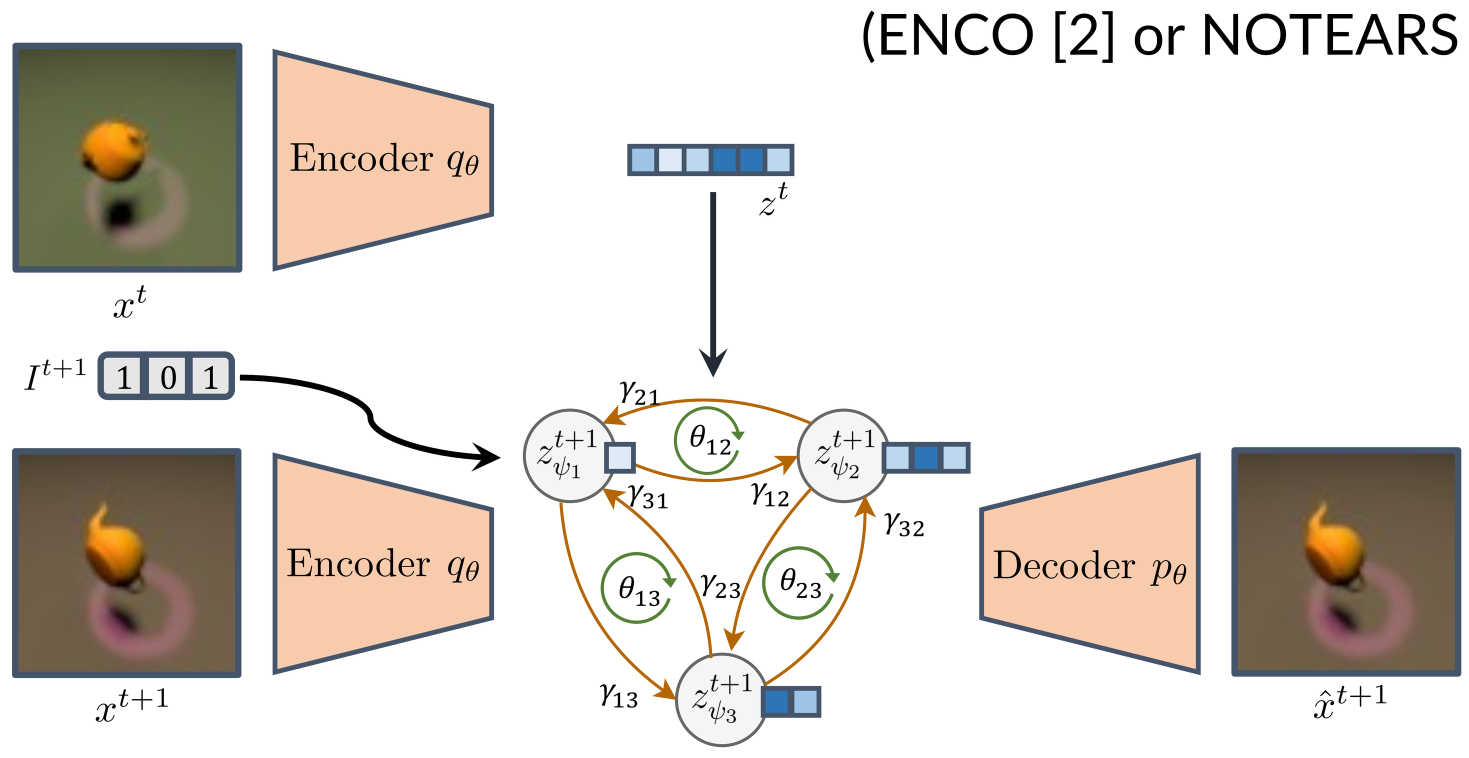

OPTIMIZATION STABILIZATION

- Chicken-and-egg situation: without graph, no disentanglement; without variables, no graph
- Our solution:
  - *Graph Learning Scheduling*: freeze graph parameters for first several iterations
  - *Mutual Information Estimator:* no MI between intervened variables and previous time step

EXPERIMENTS

- *Instantaneous Temporal Causal3DIdent*: 7 causal variables with temp. and instantaneous effects

Model	R^2 (diag ↑ / sep ↓)	SHD (instant ↓ / temp ↓)
iCITRIS-ENCO	**0.96 / 0.05**	**1.33 / 5.00**
iCITRIS-NOTEARS	0.95 / 0.09	4.00 / **5.00**
CITRIS	0.92 / 0.19	4.67 / 10.00
iVAE	0.82 / 0.20	6.67 / 15.33
iVAE-AR	0.79 / 0.29	11.00 / 12.67

- *Causal Pinball*: game dynamics with 5 causal vars

Model	R^2 (diag ↑ / sep ↓)	SHD (instant ↓ / temp ↓)
iCITRIS-ENCO	**0.98 / 0.04**	**0.67 / 3.67**
iCITRIS-NOTEARS	**0.98** / 0.06	2.33 / **3.67**
CITRIS	**0.98 / 0.04**	2.67 / 4.00
iVAE	0.55 / 0.04	2.33 / 4.33
iVAE-AR	0.53 / 0.15	4.33 / 6.33

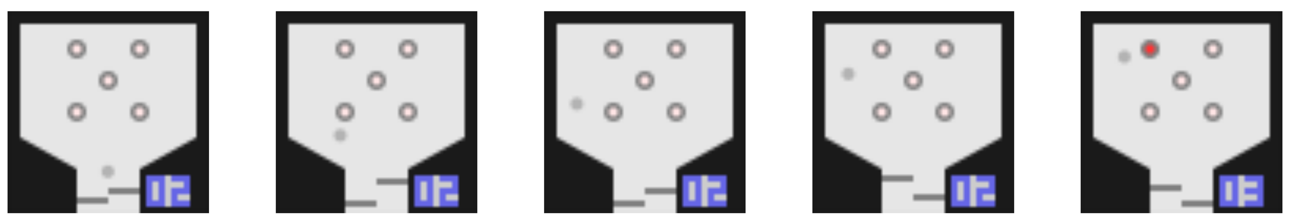

iCITRIS identifies the causal variables and their temp.+instant. graph well in both datasets

References
[1] Lippe, Phillip, et al. "CITRIS: Causal Identifiability from Temporal Intervened Sequences." *International Conference on Machine Learning*. PMLR, 2022.
[2] Lippe, Phillip, et al. "Efficient Neural Causal Discovery." *ICLR* 2022
[3] Zheng, Xun et al. "DAGs with NO TEARS: Continuous Optimization for Structure Learning." *Advances in Neural Information Processing Systems*, 2018.

QUVA Deep Vision Lab
Qualcomm AI research
MIT-IBM Watson AI Lab
DELTA LAB

Check out the full paper!