University of Amsterdam
Faculty of Science

QUVA Deep Vision Lab

Qualcomm AI research

MIT-IBM Watson AI Lab

UvA - BOSCH DELTA LAB

CITRIS: Causal Identifiability from Temporal Intervened Sequences

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 varying factors?
- Crucial for reasoning, planning, generalization, and more
- Most works consider only scalar causal variables, but can we abstract further to multidimensional causal factors?
  - We may not be able to intervene on all scalar variables separately
  - Represent single causal variable by multiple latent dimensions for better optimization

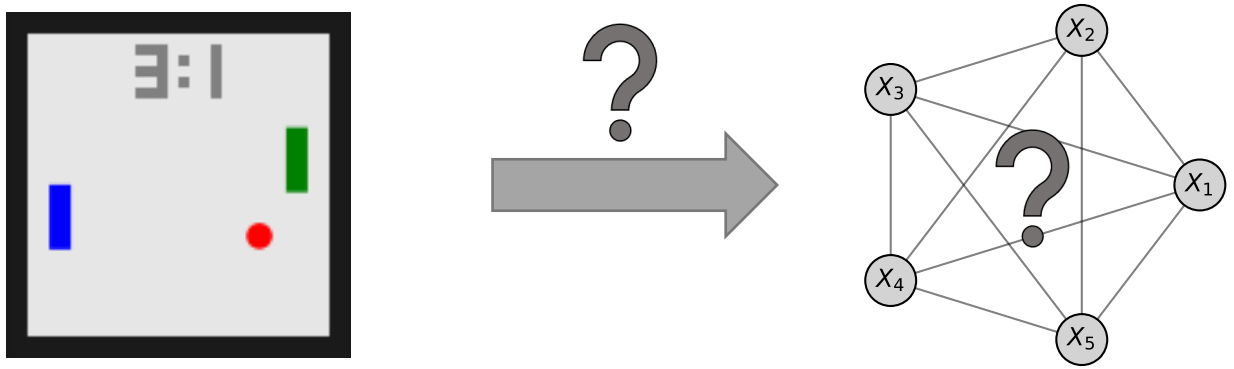

Causal Identifiability from Temporal Intervened Sequences

Setup

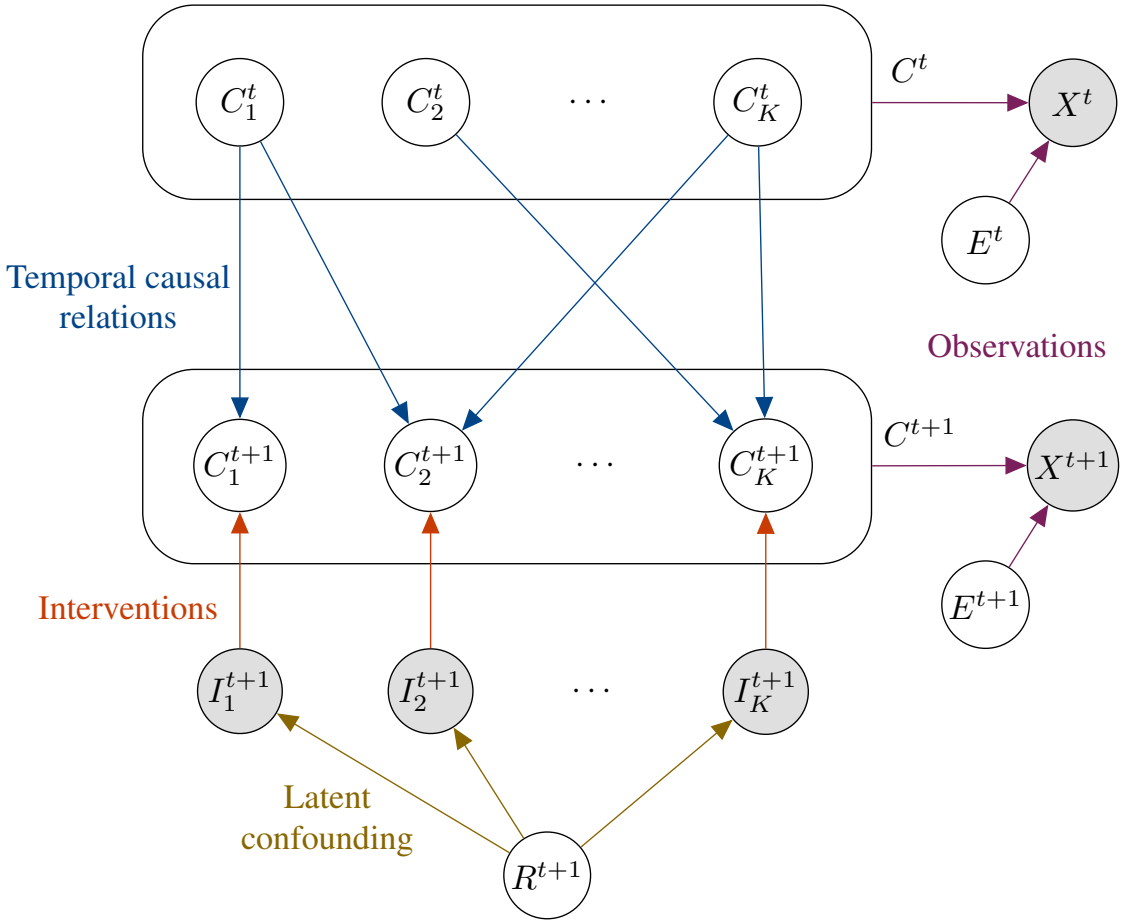

- Multidimensional latent causal variables $C_1^t, \ldots, C_K^t$
- Goal: identify causal variables from observation pairs X^t, X^{t+1} and observed intervention targets I^{t+1}

Causal Identifiability from Temporal Intervened Sequences
Minimal Causal Variables

- Main theoretical result: we can identify the ***minimal causal variables***, i.e. the information/mechanism of a causal variable which strictly depends on the interventions

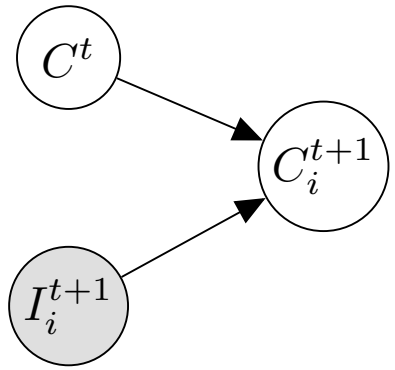

(a) Original causal graph of C_i

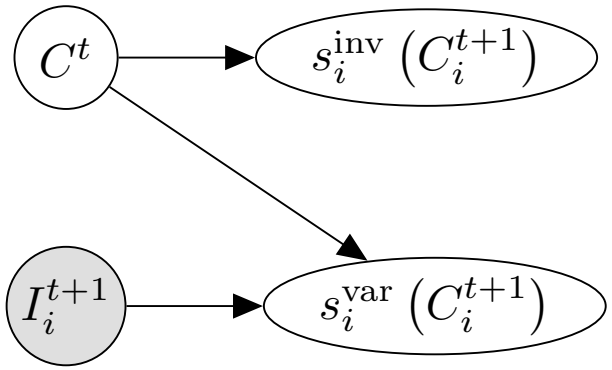

(b) Minimal causal split graph of C_i

CITRIS Architecture

CITRIS-VAE

VAE with learned latent-to-causal variable assignment

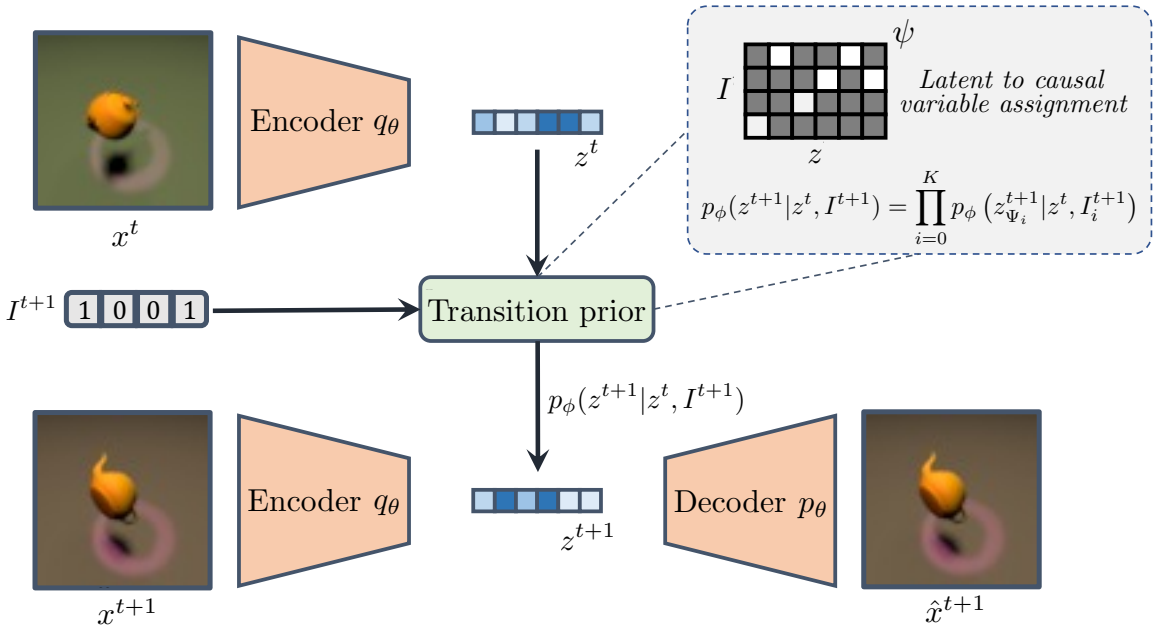

CITRIS-NF

Normalizing Flow on pretrained autoencoder

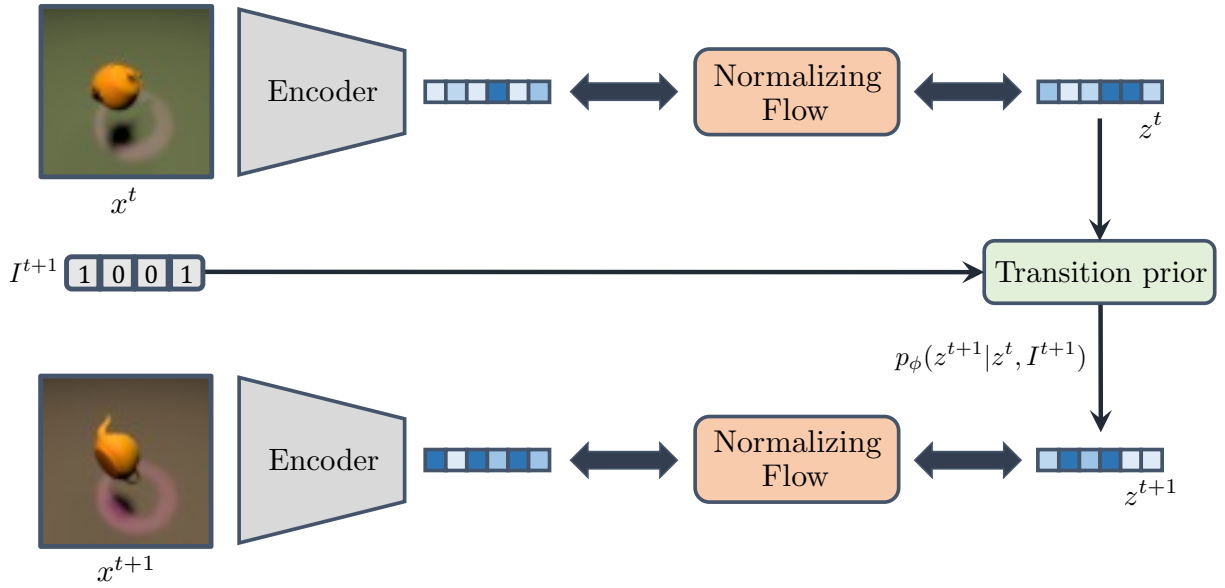

Experiments
Temporal Causal3DIdent dataset

Novel combinations of causal factors

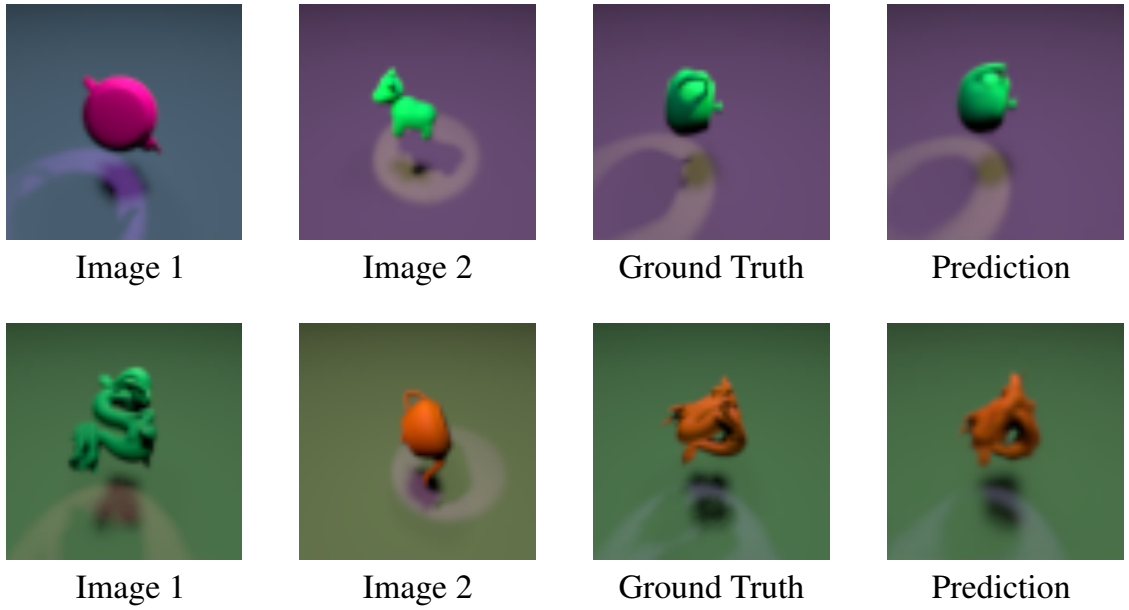

Causal Graph

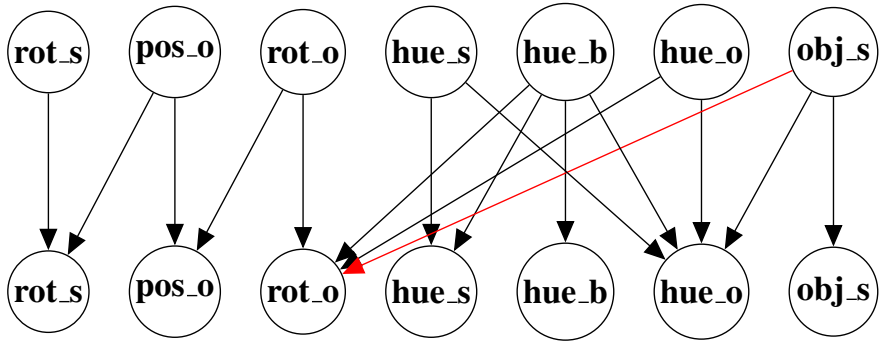

Conclusion

- CITRIS: finding multidimensional causal variables from temporal sequences with interventions
- Main characteristics of approach:
  - Identifiable minimal causal variables depend on provided interventions
  - Supports modeling causal variables in arbitrary number of latent dimensions
  - Disentangling latent representation of pretrained autoencoder with normalizing flow
  - Scales to visually complex 3d scenes

Hope to see you at our poster later!

Paper and code