Causal Representation Learning from Unknown Interventions

Kun Zhang

Carnegie Mellon University

Causal Representation Learning from Changes

Kun Zhang

Carnegie Mellon University

Thanks to Biwei Huang, Mingming Gong, Weiran Yao, Feng Xie;
Clark Glymour, Peter Spirtes, Bernhard Schölkopf, Aapo Hyvärinen…

Causality and Modular Changes

- Causal system consists of "irrelevant" modules (Pearl, 2000; Spirtes et al., 1993)

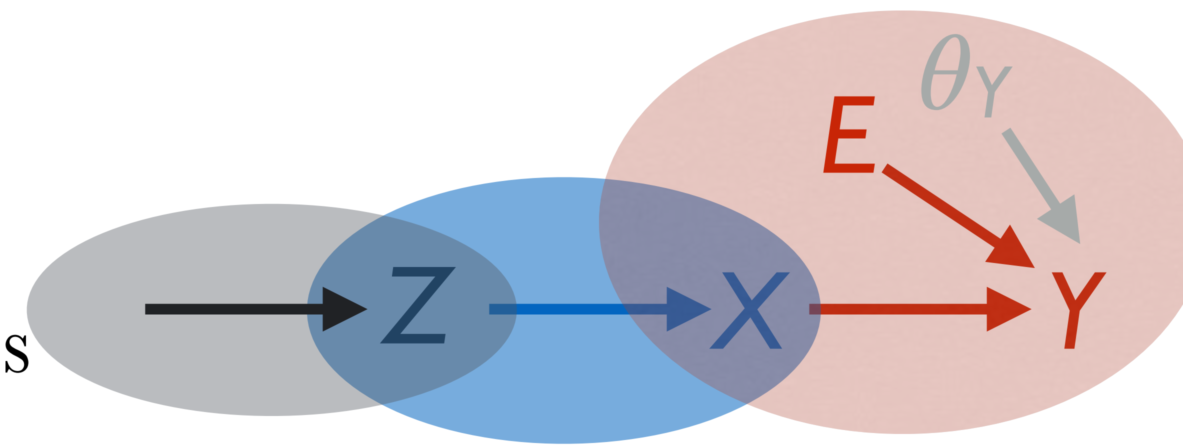

- Independent and minimal changes
  - Changes in values of measured variables or hidden variables
  - A "minimal change" representation explains the CI relations and changeability of the distribution with a minimal number of changing conditional distributions (Ghassami et al., 2018; Huang et al., 2020)
- Causal representation learning: find modular & minimal changes from observational data with identifiability guarantees

- *Huang, Zhang, Zhang, Ramsey, Sanchez-Romero, Glymour, Schölkopf, "Causal Discovery from Heterogeneous/ Nonstationary Data," JMLR, 2020*
- *Ghassami, Huang, Kiyavash, Zhang, "Multi-Domain Causal Structure Learning in Linear Systems," NeurIPS 2018*

Outline

- Causal discovery or causal representation learning: finding causal structure or hidden causal variables of interest from observational data

- Let's consider three possible dimensions of the problem

i.i.d. data?	Parametric constraint?	Latent confounders
Yes	No	No
No	Yes	Yes

Constraint-Based Causal Discovery

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- PC provides asymptotically correct results if there *doesn't* exist latent confounders (Spirtes et al., 1993)
- FCI gives asymptotically correct results even if there *are* latent confounders (Spirtes et al., 1993)
- Outputs equivalence class; might not be informative enough
- Edge minimality: Minimal ways to change the conditional distribution to produce the data dependence

- Spirtes, Glymour, and Scheines. Causation, Prediction, and Search. 1993.

Functional Causal Model-Based Causal Discovery

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

"Independent changes" renders causal direction identifiable

- Linear non-Gaussian model (Shimizu et al., 2006):
$Y = aX + E$

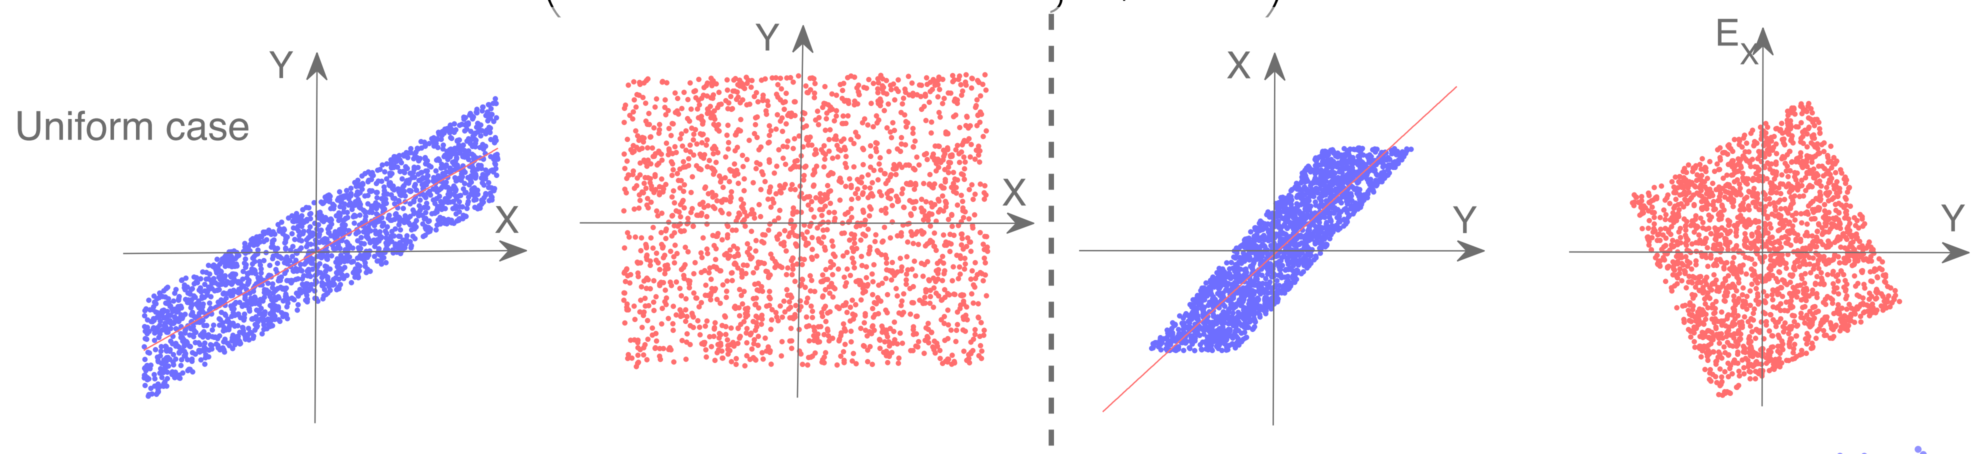

- Post-nonlinear causal model (Zhang & Chan, 2006):
$$Y = f_2\,(\,f_1(X) + E\,)$$

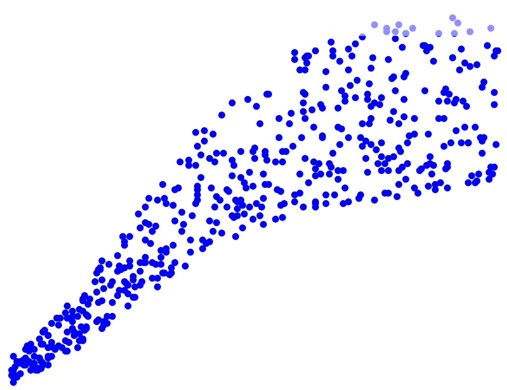

- Additive noise model (Hoyer et al, 2009)
$$Y = f(X) + E$$

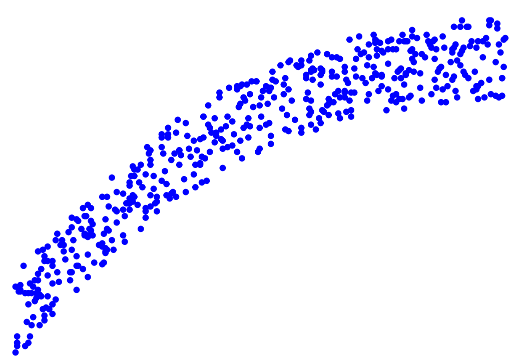

In the Presence of Latent Confounders

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Overcomplete ICA-based approach assumes independent confounders (Hoyer et al., 2008)
- Vanishing "Tetrad" condition-based approach (Silva et al., 2006)
  - Requires >= 3 pure measured variables for each confounder
  - Outputs equivalence class over latent confounders
- Generalized independent noise (GIN) in the linear, non-Gaussian case

- *Hoyer et al., Estimation of causal effects using linear nonGaussian causal models with hidden variables. IJAR, 2008*
- *Salehkaleybar, Ghassami, Kiyavash, Zhang, Learning Linear Non-Gaussian Causal Models in the Presence of Latent Variables, JMLR, 2020*
- *R. Silva et al. (2006). Learning the structure of linear latent variable models, 7:191– 246, 2006*

GIN for Estimating Linear, Non-Gaussian LV Model

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Linear, non-Gaussian *latent variable* causal model
- GIN condition
  - ($\mathbf{Y}$, $\mathbf{Z}$) satisfies GIN iff $\exists \mathbf{w} \neq \mathbf{0}$ such that $\mathbf{w}^\intercal \mathbf{Y}$ is independent from $\mathbf{Z}$
  - Graphical interpretation: exogenous set of parents of $\mathbf{Y}$ d-separate $\mathbf{Y}$ and $\mathbf{Z}$
- Step 1: find causal clusters
- Step 2: find causal order of the latent variables

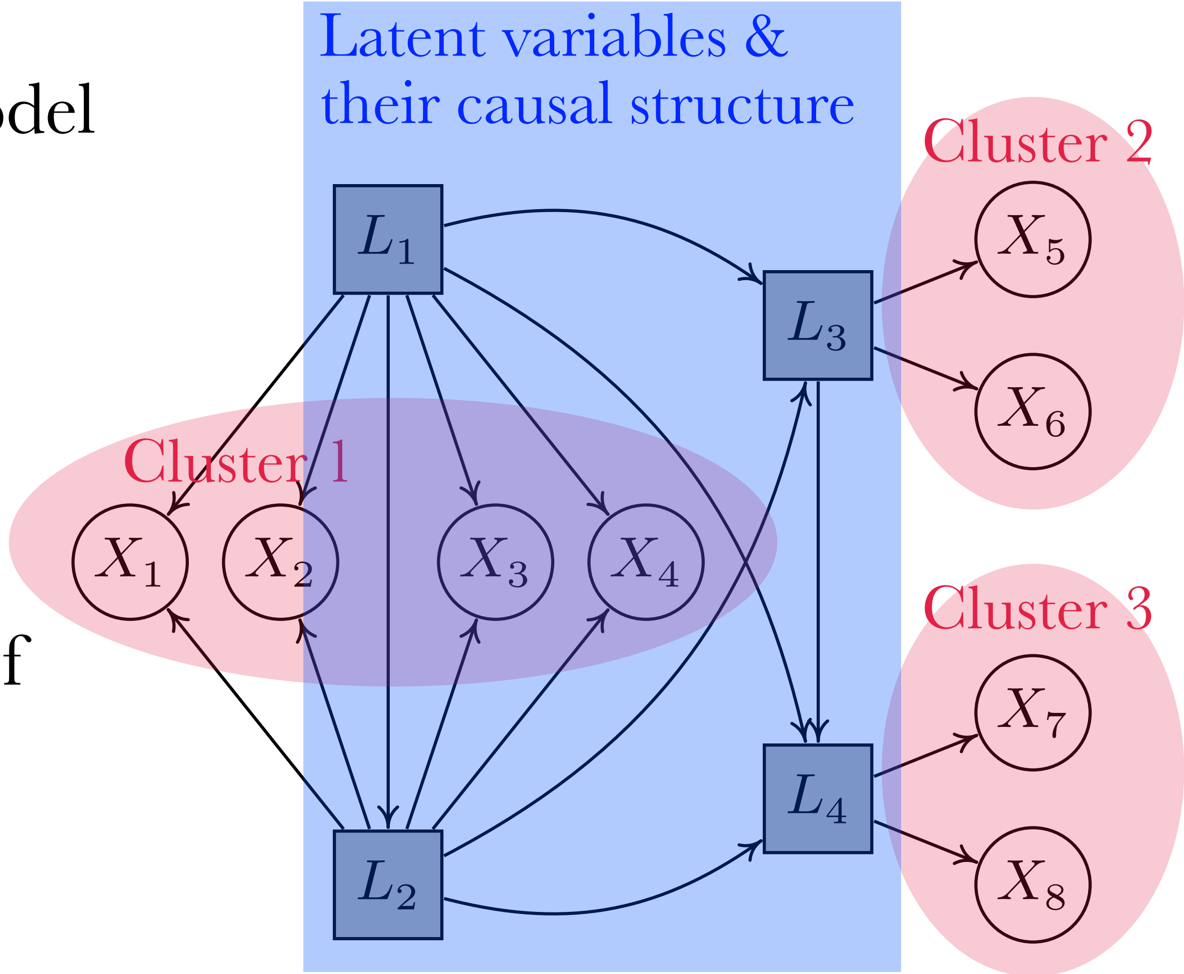

- *Xie, Cai, Huang, Glymour, Hao, Zhang, "Generalized Independent Noise Condition for Estimating Linear Non-Gaussian Latent Variable Causal Graphs," NeurIPS 2020*
- *Cai, Xie, Glymour, Hao, Zhang, "Triad Constraints for Learning Causal Structure of Latent Variables," NeurIPS 2019*

Application to Teacher's Burnout Data

- Contains 28 measured variables
- Discovered clusters and causal order of the latent variables:

Causal Clusters	Observed variables
$\mathcal{S}_1$ (1)	$RC_1, RC_2, WO_1, WO_2, DM_1, DM_2$
$\mathcal{S}_2$ (1)	CC_1, CC_2, CC_3, CC_4
$\mathcal{S}_3$ (1)	PS_1, PS_2
$\mathcal{S}_4$ (1)	$ELC_1, ELC_2, ELC_3, ELC_4, ELC_5$
$\mathcal{S}_5$ (2)	$SE_1, SE_2, SE_3, EE_1, EE_2, EE_3, DP_1, PA_3$
$\mathcal{S}_6$ (3)	DP_2, PA_1, PA_2

$$L(\mathcal{S}_1) \succ L(\mathcal{S}_2) \succ L(\mathcal{S}_3) \succ L(\mathcal{S}_5) \succ L(\mathcal{S}_4) \succ L(\mathcal{S}_6).$$

(from root to leaf)

- Consistent with the hypothesized model

Hypothesized model by experts

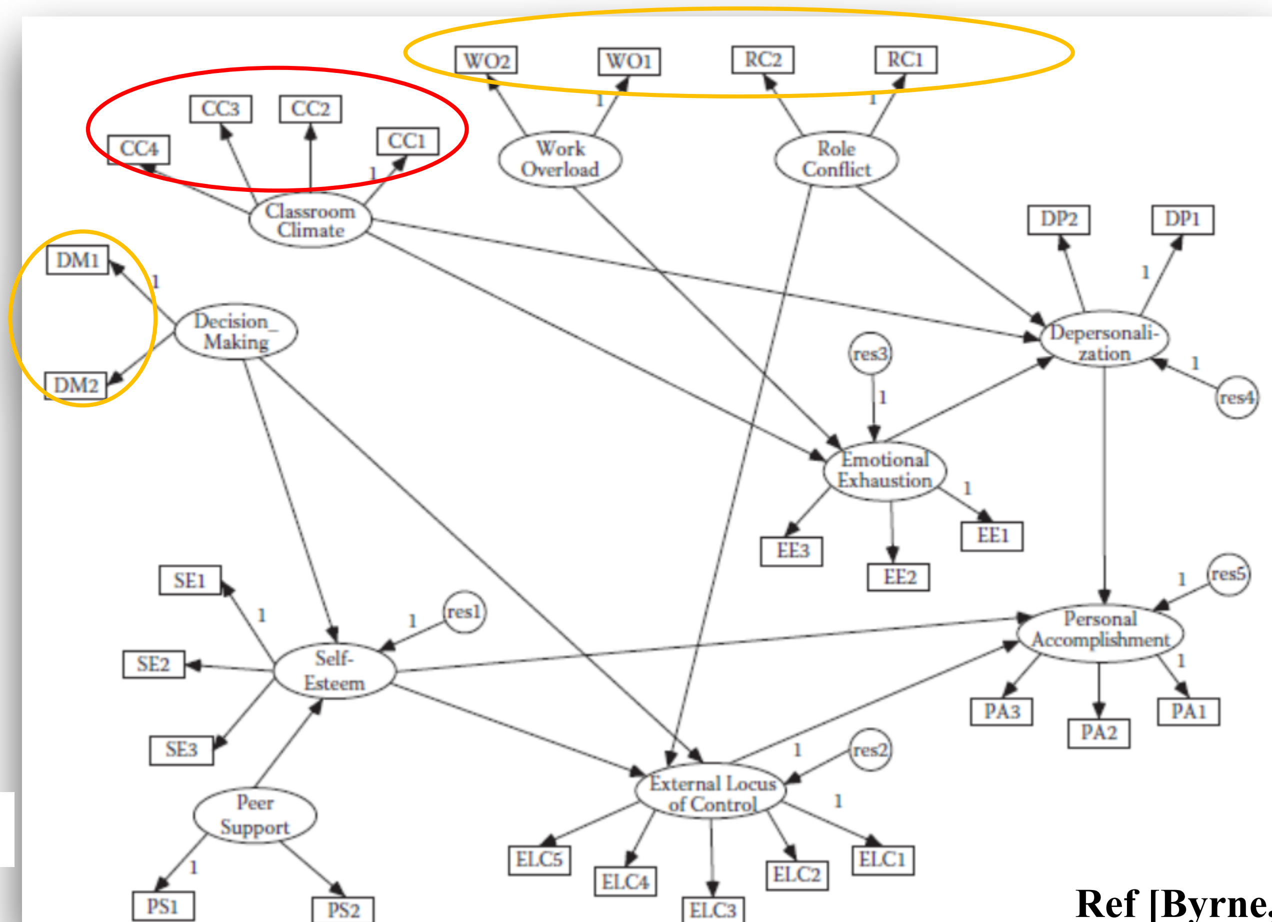

Ref [Byrne, 2010]

Estimating Latent Hierarchical Structure with GIN

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

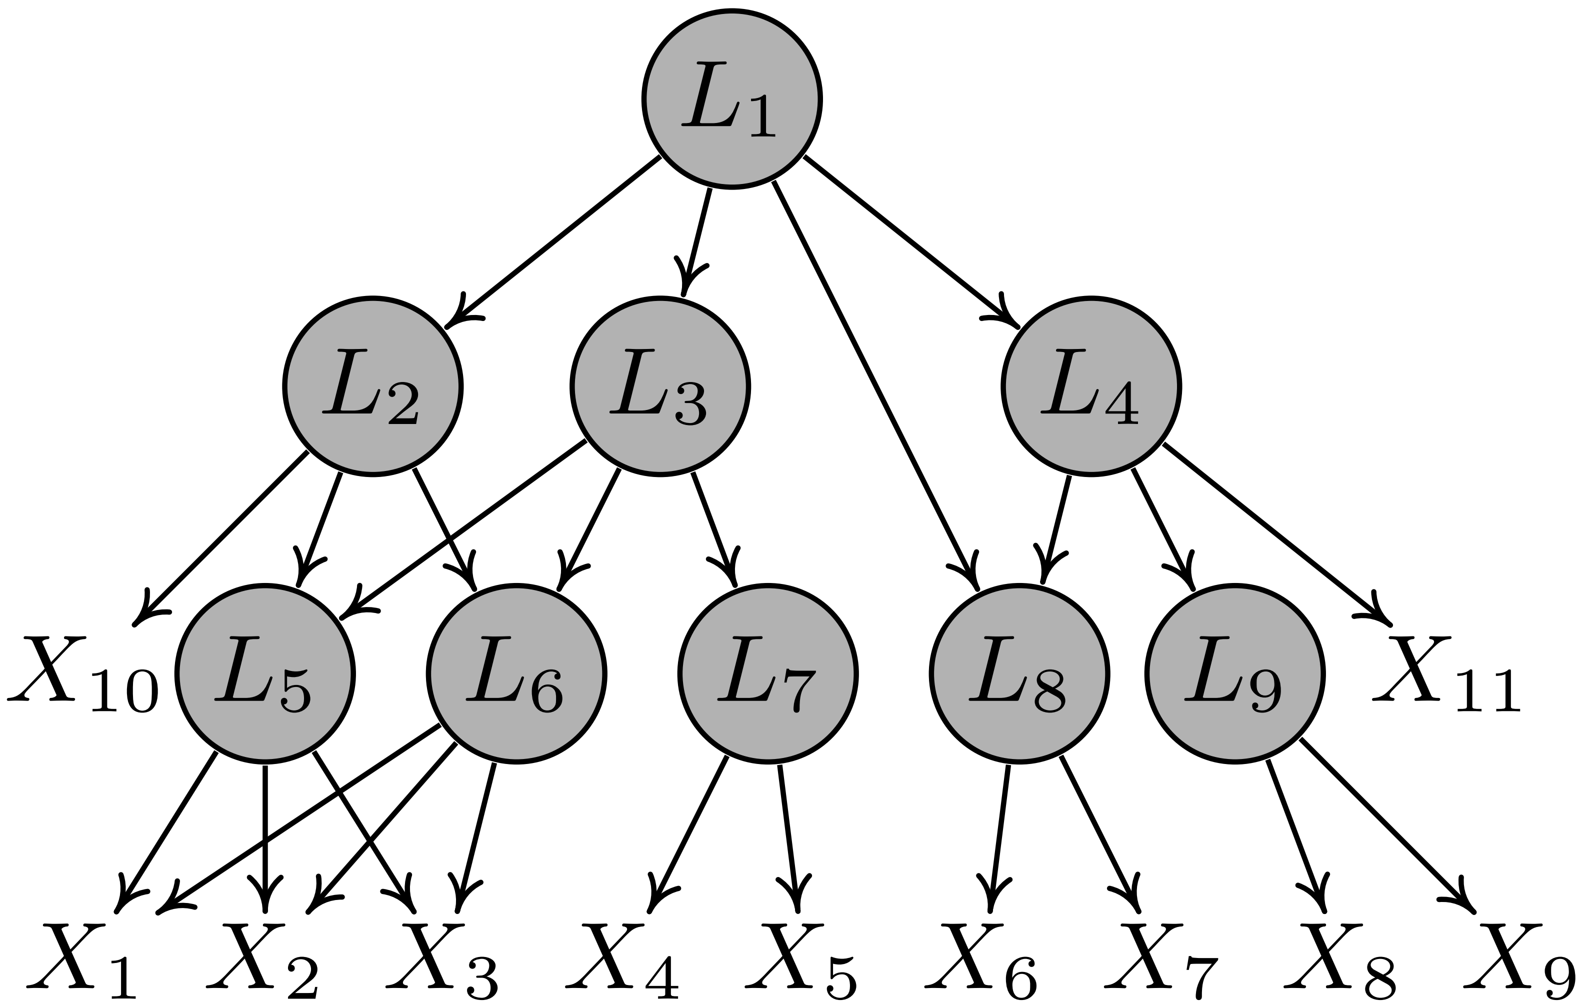

- Transitivity of linear causal influences
- GIN on measured variables

- Easy to estimate
- Linearity! :-(
- Minimality has to be assumed

- *Xie, Huang Chen, He, Geng, Zhang, "Estimation of Linear Non-Gaussian Latent Hierarchical Structure," arxiv 2022*

Necessary and Sufficient Conditions on the Structure

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Allow a large number of latent variables

- Minimality has to be assumed
- Estimation is generally difficult

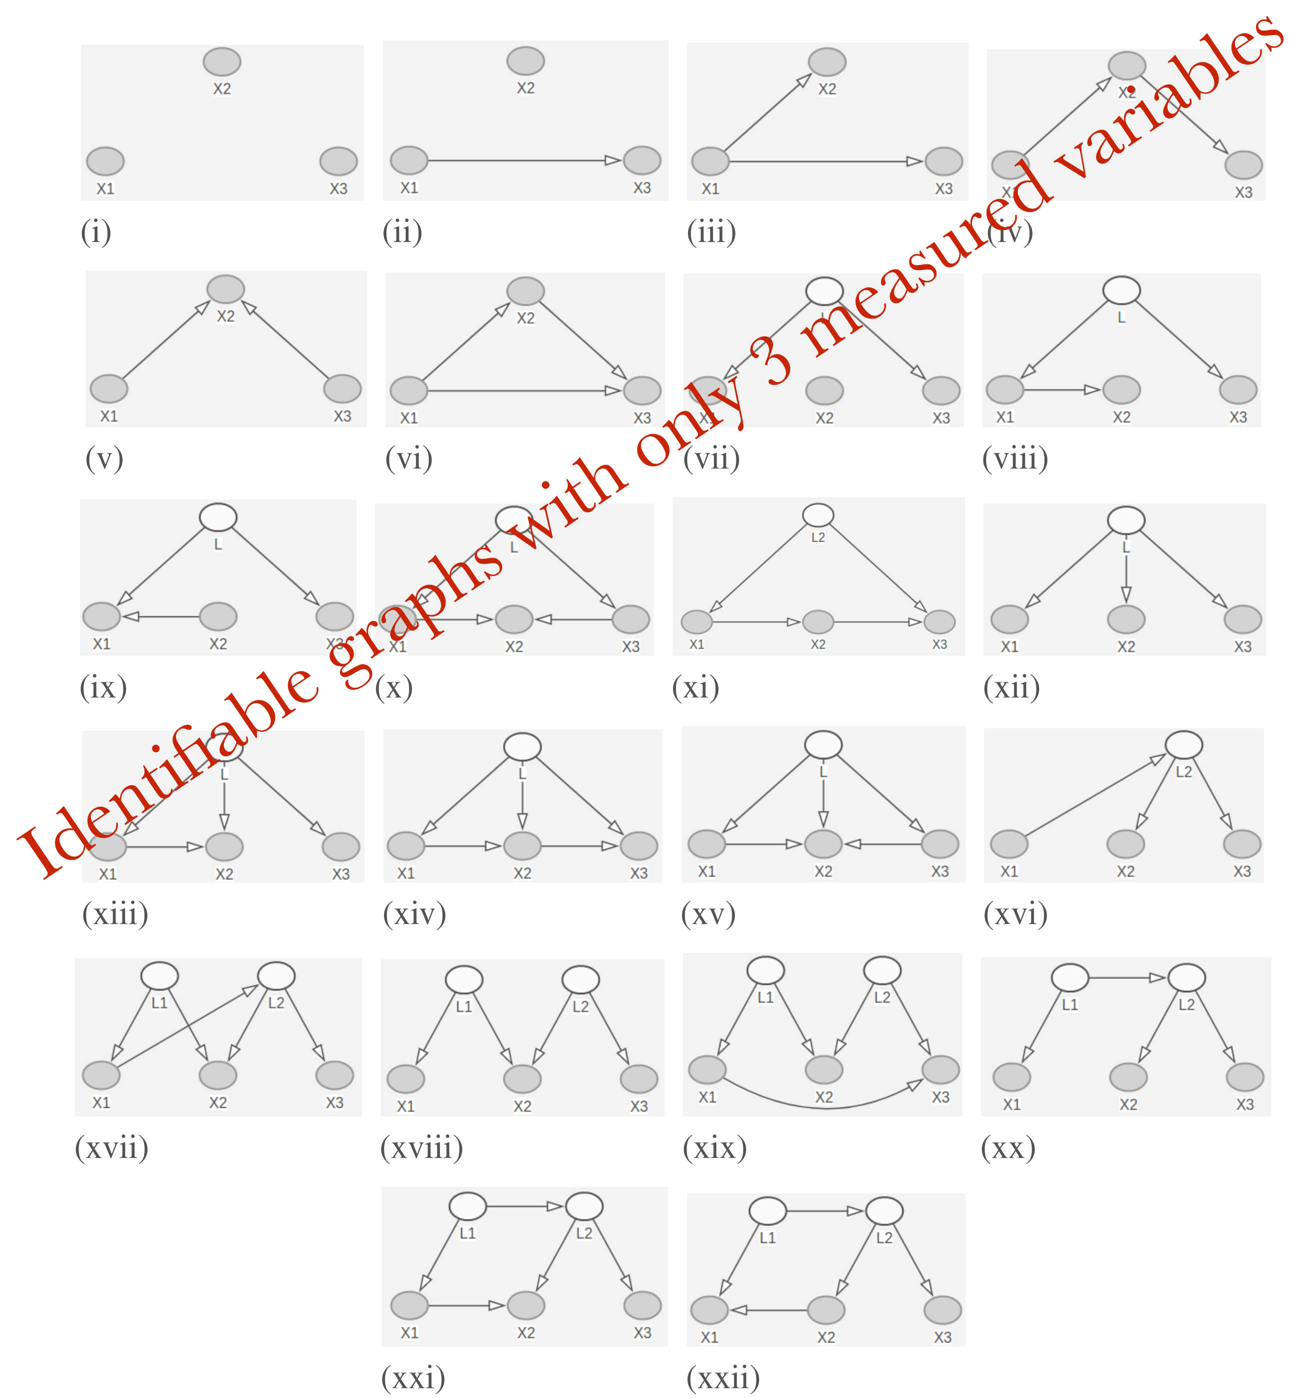

- Adams, Hansen, Zhang, "Identification of Partially Observed Linear Causal Models: Graphical Conditions for the Non-Gaussian and Heterogeneous Cases," NeurIPS 2021

Estimating Fixed *Time-Delayed* Causal Model

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Granger causality: Conditional independence-based approach + temporal constraints
- Further with instantaneous causal relations
  - Conditional independence-based approach for instantaneous relations (Swanson & Granger, 1997)
  - With linear, non-Gaussian model (Hyvärinen et al, 2010)

- *Swanson, Granger. Impulse response functions based on a causal approach to residual orthogonalization in vector autoregression. J. of the Americal Statistical Association, 1997*
- *Hyvärinen, Zhang, Shimizu, Hoyer, "Estimation of a structural vector autoregression model using non-Gaussianity," JMLR, 2010*

Learning Latent Causal Dynamics

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

"Time-delayed" influence generally renders latent processes & their relations identifiable

Time-series Inputs $\{x_t\}_{t=0}^{T}$

Unsupervised Representation Learning

$$\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$$

Latent processes

Temporal VAE with causal prior

Latent processes z_{it} follow temporal causal model with
- completely **nonparametric** model; or furthermore,
- **non-stationary** noise; or
- **non-stationary** causal influence, or
- **Parametric** constraints

Causal Skeleton Recovery

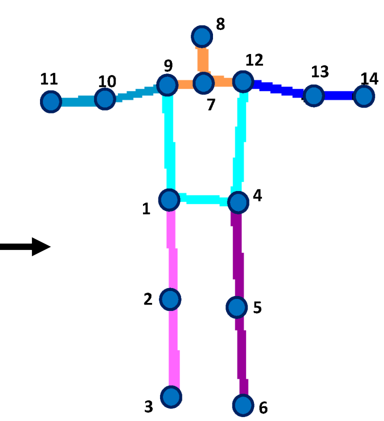

Recovered latent processes

- Why? $\cdots$ $Z_{1t} \rightarrow Z_{1,t+1}$, $Z_{1t} \rightarrow Z_{2,t+1}$, $Z_{2t} \rightarrow Z_{1,t+1}$, $Z_{2t} \rightarrow Z_{2,t+1}$ for $\mathbf{Z}'_t = h(\mathbf{Z}_t)$: $\cdots$ $Z'_{1t} \rightarrow Z'_{1,t+1}$, $Z'_{1t} \rightarrow Z'_{2,t+1}$, $Z'_{2t} \rightarrow Z'_{1,t+1}$, $Z'_{2t} \rightarrow Z'_{2,t+1}$

- *Yao, Chen, Zhang, "Learning Latent Causal Dynamics," arXiv 2022*
- *Yao, Sun, Ho, Sun, Zhang, "Learning Temporally causal latent processes from general temporal data," arxiv 2021*

Independent but Not Identically Distributed Data

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Estimation of *instantaneous* causal relations from heterogeneous/nonstationary data (Huang et al., 2020)
  - Directly benefit from minimal & independent changes
  - Statistically more efficient approaches under the linearity assumption (Ghassami et al., 2018; Huang et al., 2019)

- *Huang, Zhang, Zhang, Ramsey, Sanchez-Romero, Glymour, Schölkopf, "Causal Discovery from Heterogeneous/Nonstationary Data," JMLR, 2020*
- *Zhang, Huang, Glymour, Schölkopf, Discovery and visualization of nonstationary causal models, arxiv 2015*
- *Ghassami, Huang, Kiyavash, Zhang, "Multi-Domain Causal Structure Learning in Linear Systems," NeurIPS 2018*
- *Huang, Zhang, Gong, Glymour, "Causal Discovery and Forecasting in Nonstationary Environments with State-Space Models," ICML 2019*

Causal Discovery from Nonstationary/ Heterogeneous Data

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

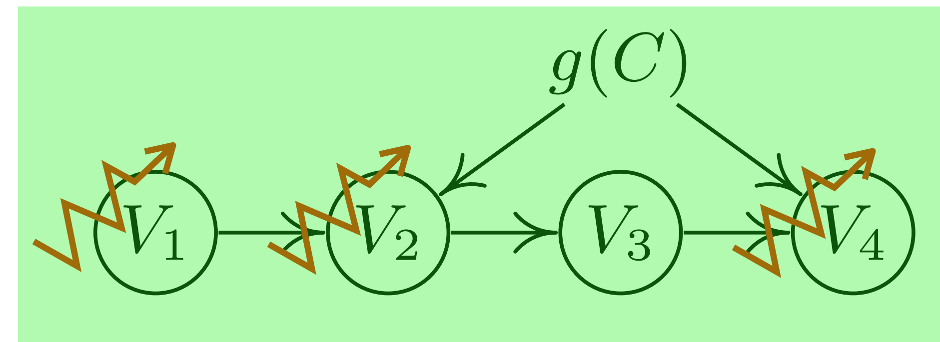

- Task:
  - Determine changing causal modules & estimate skeleton
  - Causal orientation determination benefits from independent changes in P(cause) and P(effect | cause), including invariant mechanism/ cause as special cases
  - Visualization of changing modules over time/ across data sets?

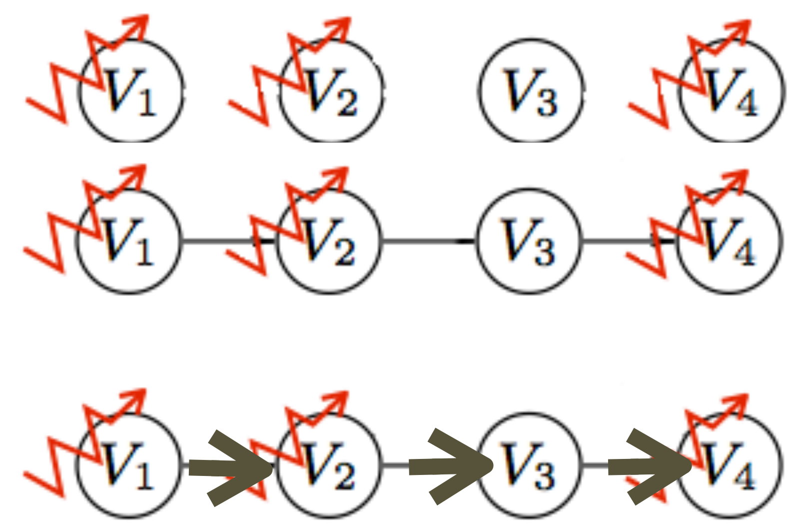

Kernel nonstationary driving force estimation

- Huang, Zhang, Zhang, Ramsey, Sanchez-Romero, Glymour, Schölkopf, "Causal Discovery from Heterogeneous/ Nonstationary Data," JMLR, 2020

Nonlinear ICA with Multiple Domains

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Nonlinear ICA: observed variables follow $\mathbf{X} = \mathbf{g}(\mathbf{Z})$, in which the components of $\mathbf{Z}$, Z_i, are mutually independent
  - Solutions to nonlinear ICA high non-unique
  - If the distributions of Z_i change across multiple domains, generally their are identifiable (up to component-wise transformations)
  - Why?

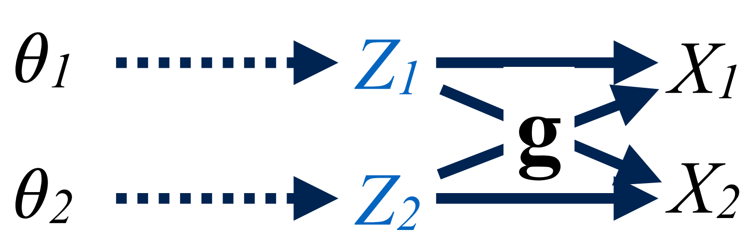

for $\mathbf{Z}'=h(\mathbf{Z})$:

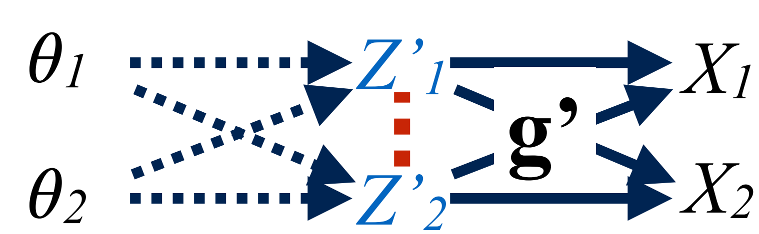

- *Hyvarinen, Sasaki, Turner, "Nonlinear ICA using auxiliary variables and generalized contrastive learning," In The 22nd International Conference on Artificial Intelligence and Statistics, 2019.*

With Changing Causal Relations among Latent Variables

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Measured variables follow $\mathbf{X} = \mathbf{g}(\mathbf{Z})$, in which the components of $\mathbf{Z}$, Z_i, are causally related and some causal relations change
- Fixed causal relations and the the involved variables are not identifiable

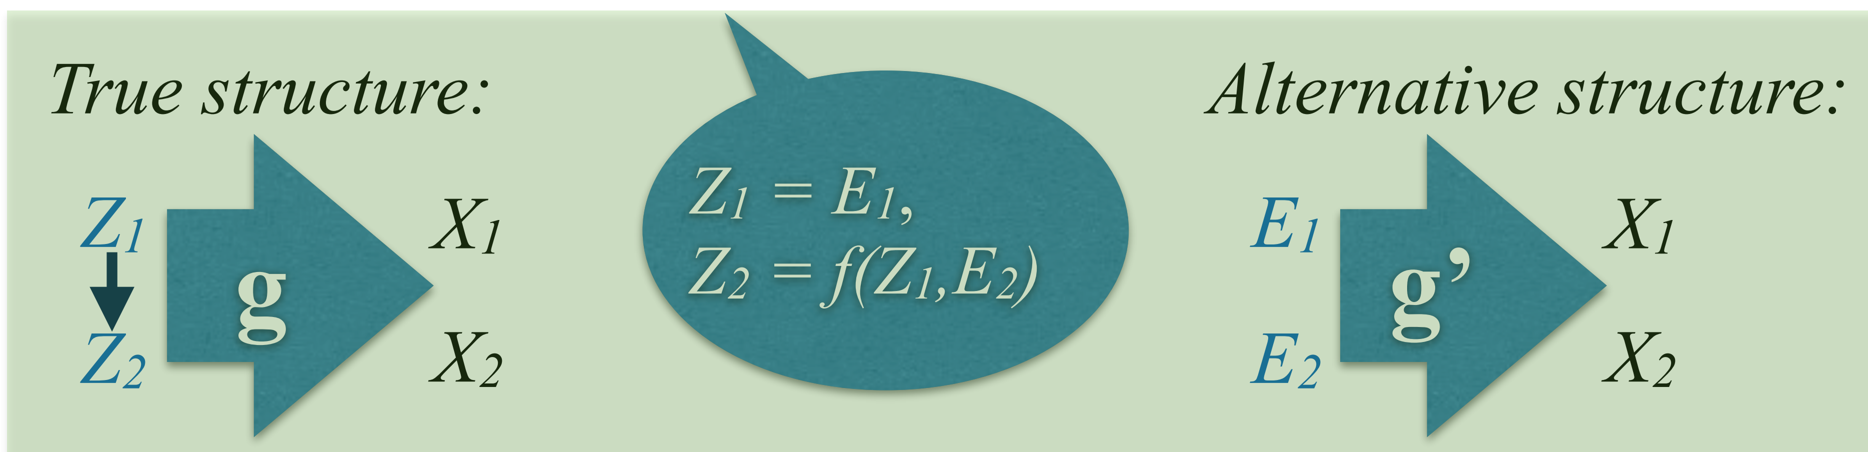

- What if some causal relations (over latent variables) change?

- *Zhang, Yao "Causal Disentanglement with Minimal Changes from Multiple Distributions," available upon request*

With Changing Causal Relations among Latent Variables: Partial Identifiability

i.i.d. data?	Parametric constraint?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Canonical representation:

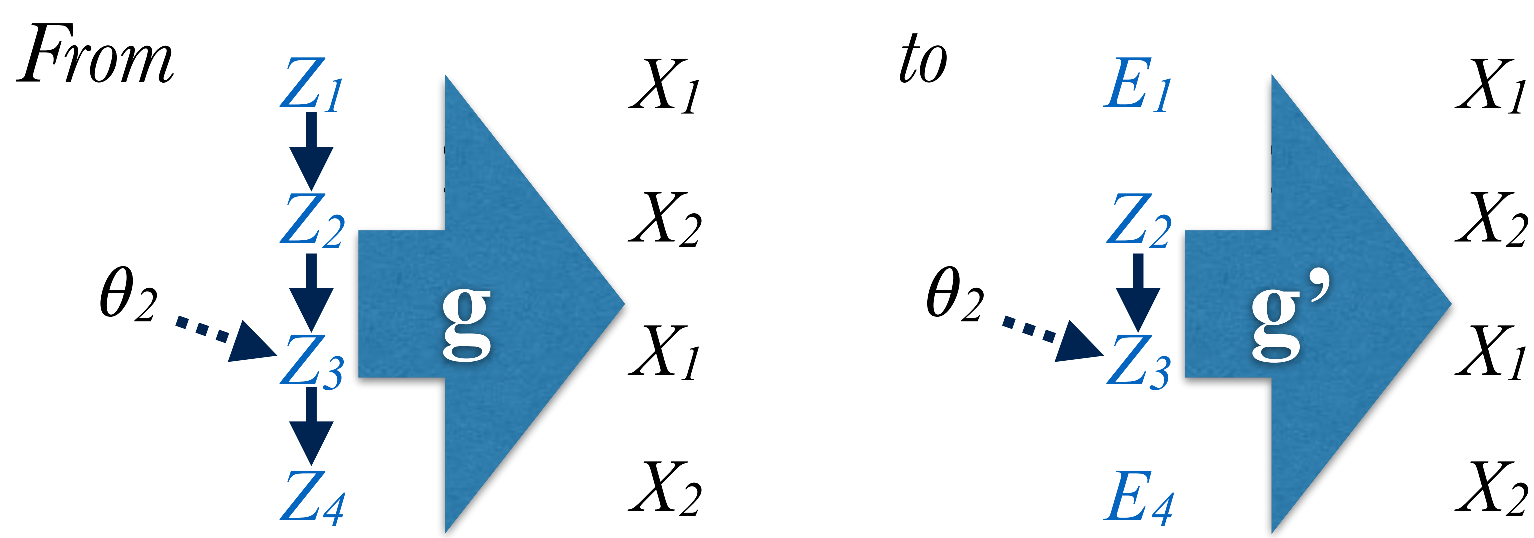

- Invariant part E_i are identifiable up to its subspace (estimated E_i do not receive contribution from Z_2 or Z_3)
- Variables involved in changing causal influence, Z_2 and Z_3, are identifiable up to their transformations
    - Z_2 is further identifiable in the linear-Gaussian case

Summary

- Causal representation learning: identifiable structure/ variables under modular/minimal changes in the data
- Different levels of changes
  - Changes in values of variables
  - Changes in hidden variables/ parameters
- Latent variables and their relations involved in changing influences are generally identifiable

Summary

i.i.d. data?	Parametric constraint?	Latent confounders?	What can we get?
Yes	No	No	(Different types of) equivalence class
		Yes	
	Yes	No	Unique identifiability (under structural conditions)
		Yes	
Non-I, but I.D.	No/Yes	No	(Extended) regression
		Yes	Latent temporal causal processes identifiable!
I., but non-I.D.	No	No	More informative than MEC (CD-NOD)
	Yes		May have unique identifiability
	No	Yes	Changing subspace identifiable
	Yes		Variables in changing relations identifiable