

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

2-23-2022

Learning Latent Causal Dynamics

Weiran Yao

Carnegie Mellon University

Guangyi Chen

Carnegie Mellon University & Mohamed bin Zayed University of Artificial Intelligence

Kun Zhang

Carnegie Mellon University & Mohamed bin Zayed University of Artificial Intelligence

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Computer Sciences Commons](#), [Data Science Commons](#), and the [Statistics and Probability Commons](#)

Preprint: arXiv

- Archived with thanks to arXiv
- Preprint License: [CC by NC-ND](#)
- Uploaded 24 March 2022

Recommended Citation

W. Yao, G. Chen, and K. Zhang, "Learning latent causal dynamics," 2022, arXiv:arXiv:2202.04828v4

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Learning Latent Causal Dynamics

Weiran Yao¹ Guangyi Chen^{1,2} Kun Zhang^{1,2}

Abstract

One critical challenge of time-series modeling is how to learn and quickly correct the model under unknown distribution shifts. In this work, we propose a principled framework, called *LiLY*, to first recover time-delayed latent causal variables and identify their relations from measured temporal data under different distribution shifts. The correction step is then formulated as learning the low-dimensional change factors with a few samples from the new environment, leveraging the identified causal structure. Specifically, the framework factorizes unknown distribution shifts into transition distribution changes caused by fixed dynamics and time-varying latent causal relations, and by global changes in observation. We establish the identifiability theories of non-parametric latent causal dynamics from their non-linear mixtures under fixed dynamics and under changes. Through experiments, we show that time-delayed latent causal influences are reliably identified from observed variables under different distribution changes. By exploiting this modular representation of changes, we can efficiently learn to correct the model under unknown distribution shifts with only a few samples.

1. Introduction

Unknown distribution shifts or nonstationarity in natural signals have posed a critical challenge for time-series forecasting. For example, if the distribution shifts (e.g., changing dynamics, heterogeneous noise, etc.) are not well captured, models fit on these data samples will either overfit specific kinds of changes, or learn to use only the stationary parts for prediction, both of which significantly degrade the predictive performances. One prevalent approach in the literature is to actively transfer the time-series model under the dis-

¹Carnegie Mellon University, Pittsburgh, PA, USA ²Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), Abu Dhabi, United Arab Emirates. Correspondence to: Kun Zhang <kunz1@cmu.edu>.

Preprint.




Figure 1. **LiLY**: Learning Latent causal dYnamics under modular distribution shift. We exploit distribution changes resulting from fixed causal dynamics, changing causal influences and global observation changes to identify the underlying causal dynamics. The distribution change in a new segment is corrected via learning the low-dimensional change factors in an unsupervised way.

tribution shifts. However, if one is not aware of how the changes influence the underlying data generating process, one may not know how to target specific mechanisms and change the learned model efficiently. In other words, the true underlying processes or a graphical model that encode the change property of the distribution need to be given as a priori. To accomplish this step, causal discovery is often used. Causal discovery aims to identify the underlying structure of the data generation process by exploiting an appropriate class of assumptions (Spirtes et al., 1993; Pearl, 2000). Although distribution changes are not desired by time-series models, these changes, serving as a “soft” way of intervention, have proved to greatly improve the identifiability results for learning the latent causal structure (Yao et al., 2021; Bengio et al., 2019; Ke et al., 2019). For this reason, a natural solution is to first exploit the distribution changes to identify the latent causal dynamics, and then use the uncovered structure to correct the model under changes.

In this work, we propose a principled framework, called *LiLY*, to first recover time-delayed latent causal variables and identify their relations from measured temporal data under different distribution shifts. The transfer step is then formulated as learning the low-dimensional change factors with a few samples from the new environment, leveraging the identified causal structure. For instance, Fig. 1 shows an example of multiple video segments of a physical system (i.e., cartpole) under different mass, gravity and environment rendering settings. With *LiLY*, the differences across segments are characterized by change factors θ_k^{dyn} of segment k that encode changes in transition dynamics, and changes in observation modeled by θ_k^{obs} . We then present a generalized time-series data generative model that takes these change factors as arguments for modeling the distribution changes. Specifically, the proposed latent variable model factorizes distribution changes in fixed dynamics, time-varying transitions and observation by constructing partitioned latent subspaces. If the underlying latent causal dynamics augmented with change factors is uncovered, correcting distribution shifts can be performed efficiently by estimating the change factors in a new segment.

Table 1. Attributes of prior nonlinear ICA theories for time-series. A check denotes that a method has an attribute, whereas a cross denotes the opposite. [†] indicates our approach.

Theory	Time-varying Relation	Causally-related Process	Partitioned Subspace	Nonparametric Transition
PCL	×	×	×	✓
HM-NLICA	×	×	×	×
SlowVAE	×	×	×	×
SNICA	✓	✓	×	×
LEAP	×	✓	×	✓
LiLY[†]	✓	✓	✓	✓

However, estimating latent causal structure from observations is notoriously challenging because the latent variables, even with independent factors of variation (Locatello et al., 2019), are not identifiable in the most general case (Hyvärinen & Pajunen, 1999). There exist several pieces of work aiming to uncover causally related latent variables under linear relations or minimality assumptions. As early as (Spearman, 1928), vanishing Tetrad conditions were exploited to identify latent variables in linear-Gaussian models (Silva et al., 2006). Generalized Independent Noise (GIN) condition estimates linear, non-Gaussian latent variable causal graph (Xie et al., 2020). Recently, in the field of nonlinear Independent Component Analysis (ICA), identifiability results have been established (Hyvärinen & Morioka, 2016; 2017; Hyvärinen et al., 2019; Khemakhem et al., 2020; Sorrenson et al., 2020) by using certain side information, such as class labels, in addition to independence. For time-series data, history information is widely used as side information for nonlinear ICA. However, most existing work that establishes identifiability results considers either stationary independent sources such as PCL (Hyvärinen & Morioka, 2017), SlowVAE (Klindt et al., 2020) or under lin-

ear transition assumptions such as SlowVAE (Klindt et al., 2020) and SNICA (Hälvä et al., 2021), or with certain structure such as Markov properties in HM-NLICA (Hälvä & Hyvärinen, 2020). LEAP (Yao et al., 2021), which is the closest work to ours, has established the identifiability of the nonparametric latent temporal processes in certain non-stationary cases, under the condition that the distribution of the noise terms of the latent processes vary across segments.

Our theory can be differentiated from the prior work in four aspects as summarized in Table 1:

1. We consider a very general nonstationary case which includes changing causal dynamics such as changes in the influencing strength or switching some edges off (Li et al., 2020), nonstationary noise distributions across domains (Yao et al., 2021), and domain shifts in observation (Sahoo et al., 2021), etc., as special cases;
2. We establish the identifiability results for nonparametric, causally-related time-delayed latent processes, instead of independent processes, or under linear assumptions;
3. We allow both stationary and nonstationary latent processes to co-exist by using partitioned latent subspaces in the data generating processes. The identifiability results under this setting are appealing to real-world data where whether and how the distribution changes are unknown. Our identifiability condition is generally much weaker than the ones in (Yao et al., 2021) because we further benefit from distribution changes under fixed causal dynamics, while LEAP assumes all latent processes are changed across contexts.

Through experiments, we show that time-delayed latent causal processes are reliably identified from observed variables under different distribution shifts and dependency structures. By exploiting this modular representation of changes, we can efficiently learn to correct the model under unknown distribution shifts with only a few samples.

2. Related Work

2.1. Nonlinear ICA for Time Series

Temporal structure and nonstationarities were used to achieve identifiability of nonlinear ICA. Hyvärinen & Morioka (2016) proposed time-contrastive learning (TCL) based on the independent sources assumption and leverage variability in variance terms. (Hyvärinen & Morioka, 2017) developed a permutation-based contrastive (PCL) learning framework which discriminates between true independent sources and permuted sources, and the model is identifiable under the uniformly dependent assumption. (Hälvä & Hyvärinen, 2020) combined nonlinear ICA with a Hidden Markov Model (HMM) to automatically model nonstation-

arity without the need for manual data segmentation. (Khemakhem et al., 2020) introduced VAEs to approximate the true joint distribution over observed and auxiliary nonstationary regimes. The conditional distribution in their work is assumed to be within exponential families to achieve identifiability on the latent space. The most recent literature on nonlinear ICA for time-series is the work of (Yao et al., 2021), which proposed both a nonparametric condition leveraging the nonstationary noise terms, and a linear, parametric condition leveraging the functional form with generalized Laplacian properties of the noise terms.

2.2. Causal Discovery from Time Series

Inferring the causal structure from time-series data is critical to many fields including machine learning (Berzuini et al., 2012), econometrics (Ghysels et al., 2016), and neuroscience (Friston, 2009). Most existing work focuses on estimating the temporal causal relations between observed variables. For this task, constraint-based methods (Entner & Hoyer, 2010) apply the conditional independence tests to recover the causal structures, while score-based methods (Murphy et al., 2002; Pamfil et al., 2020) define score functions to guide a search process. Furthermore, Malinsky & Spirtes (2018; 2019) propose to fuse both conditional independence tests and score-based methods. The Granger causality (Granger, 1969) and its non-linear variations (Tank et al., 2018; Löwe et al., 2020) are also widely used.

3. Problem Formulation

We first present a holistic, modular time-series data generative model that factorizes distribution changes in fixed dynamics, time-varying transitions and observation by constructing partitioned latent subspaces. We then define the properties of the generative model which suffice the identifiability of time-delayed causal dynamics. Finally, we formulate the problem of distribution shift correction as learning the change factor representation in a new environment with the identified latent causal structure.

3.1. A Modular Representation of Changes

Following (Huang et al., 2020; 2021), we can represent the changes in a compact way by using a low-dimensional vector θ_k . Specifically, let \mathbf{u} denotes domain or segment index. Suppose there exist K domains or segments, i.e., u_k with $k = 1, 2, \dots, K$. In each segment, we observe time series $\{\mathbf{x}_t\}_{t=1}^T$ generated by an arbitrary invertible nonlinear mixture of the underlying temporal processes \mathbf{z}_t . We partition the latent space into three blocks $\mathbf{z}_t = (\mathbf{z}_t^{\text{fix}}, \mathbf{z}_t^{\text{chg}}, \mathbf{z}_t^{\text{obs}})$ where $z_{s,t}^{\text{fix}}$ denotes the s^{th} component of the fixed dynamics parts, $z_{c,t}^{\text{chg}}$ denotes the c^{th} component of the changing dynamics parts, and $z_{o,t}^{\text{obs}}$ is the o^{th} component of the global

changes (e.g., styles). We assume the data generating process in each segment \mathbf{u}_k can be described by the transition function f for each dimension of \mathbf{z}_t that takes its parent nodes change factors, $(\theta_k^{\text{dyn}}, \theta_k^{\text{obs}})$ and noise terms as inputs, and the mixing function \mathbf{g} as:

$$\begin{cases} z_{s,t}^{\text{fix}} &= f_s(\{z_{i,t-\tau} | z_{i,t-\tau} \in \text{Pa}(z_{s,t}^{\text{fix}})\}, \epsilon_{s,t}), \\ z_{c,t}^{\text{chg}} &= f_c(\{z_{i,t-\tau} | z_{i,t-\tau} \in \text{Pa}(z_{c,t}^{\text{chg}})\}, \theta_k^{\text{dyn}}, \epsilon_{c,t}), \\ z_{o,t}^{\text{obs}} &= f_o(\theta_k^{\text{obs}}, \epsilon_{o,t}), \\ \mathbf{x}_t &= \mathbf{g}(\mathbf{z}_t), \end{cases} \quad (1)$$

where $\text{Pa}(z_{i,t})$ denote the set of (time-delayed) parent nodes of $z_{i,t}$, and $\epsilon_{s,t}, \epsilon_{c,t}, \epsilon_{o,t}$ terms are mutually independent (i.e., spatially and temporally independent) random noises in each segment. The vector $\theta_k = (\theta_k^{\text{dyn}}, \theta_k^{\text{obs}})$ embeds the information of segment u_k . Note that with Eq. 1, we allow a very general class of changes in the model, which includes changing causal dynamics such as changes in the influencing strength or switching some edges off (Li et al., 2020), nonstationary noise distributions across segments (Yao et al., 2021), and domain shifts in observation (Sahoo et al., 2021), etc., all as special cases.

Remark 1. *The mutual independence of noise terms over space and time imply that: (1) $\epsilon_{i,t}$ are independent from the parent nodes $\text{Pa}(z_{i,t})$, and (2) $z_{i,t}$ are conditional independent given the history information $\mathbf{z}_{Hx} = \{\mathbf{z}_{t-\tau}\}_{\tau=1}^L$ up to maximum time lag L and segment embeddings or index.*

This Remark suggests that the Independent Noise (IN) condition (Pearl et al., 2000) and Conditional Independence (CI) condition (Hyvarinen et al., 2019), both of which are common assumptions for causal discovery, are naturally satisfied in time-series data under random noise.

3.2. Identifiability of Time-Delayed Causal Dynamics

We define the identifiability of time-delayed latent causal dynamics in the representation function space. If the latent variables can be identified at least up to permutation and component-wise invertible nonlinearities, we say that latent causal dynamics are also identifiable because conditional independence relations fully characterize time-delayed causal relations in a time-delayed causally sufficient system.

Definition 1 (Identifiable Latent Causal Dynamics). *Formally let $\{\mathbf{x}_t\}_{t=1}^T$ be a sequence of observed variables generated by the true temporally causal latent processes specified by $(f_i, \theta_k, p(\epsilon_i), \mathbf{g})$ given in Eq. 1. A learned generative model $(\hat{f}_i, \hat{\theta}_k, \hat{p}(\epsilon_i), \hat{\mathbf{g}})$ is observationally equivalent to $(f_i, \theta_k, p(\epsilon_i), \mathbf{g})$ if the model distribution $p_{\hat{f}, \hat{\theta}_k, \hat{p}(\epsilon_i), \hat{\mathbf{g}}}(\{\mathbf{x}_t\}_{t=1}^T)$ matches the data distribution $p_{f, \theta_k, p(\epsilon_i), \mathbf{g}}(\{\mathbf{x}_t\}_{t=1}^T)$ everywhere. We say latent causal processes are identifiable if observational equivalence can lead to identifiability of the latent variables up to permutation π*

and component-wise invertible transformation T :

$$\begin{aligned} p_{f_i, \hat{\theta}_k, \hat{p}_{e_i}, \hat{g}}(\{\mathbf{x}_t\}_{t=1}^T) &= p_{f_i, \theta_k, p_{e_i}, g}(\{\mathbf{x}_t\}_{t=1}^T) \\ \Rightarrow \hat{g}(\mathbf{x}_t) &= g \circ \pi \circ T, \quad \forall \mathbf{x}_t \in \mathcal{X}, \end{aligned} \quad (2)$$

where \mathcal{X} is the observation space.

3.3. Distribution Shift Correction

We first assume that the latent causal dynamics model in Eq. 1 is identified (we will explain under what conditions the identifiability can be achieved in Sec. 4 and how in Sec. 5). Then correcting distribution shifts is equivalent to learning the θ_k embeddings in the new environment. In practice, distributions are often changed in a sparse or local way given the causal disentanglement, i.e., only a few factors may change simultaneously, which is known as sparse mechanism shift assumption (Schölkopf et al., 2021). Following the assumption, we encode θ_k with low-dimensional embeddings, in order to capture only the domain differences.

4. Theory

In this section, we establish the identifiability theory of non-parametric time-delayed latent causal processes under three different types of distribution shifts. W.l.o.g., we consider the latent processes with maximum time lag $L = 1$. In particular, (1) under fixed causal dynamics, we leverage the distribution changes $p(z_{k,t} | \mathbf{z}_{t-1})$ for different values of \mathbf{z}_{t-1} ; (2) under changing causal dynamics, we exploit the changing causal influences on $p(z_{k,t} | \mathbf{z}_{t-1}, u_k)$ under different domain u_k , and (3) under global observation changes, the nonstationarity $p(z_{k,t} | u_k)$ under different values of u_k is exploited. We illustrate through examples that, even without explicit changes in causal influences, the latent causal processes are generally identifiable if the process noises are not perfectly Gaussian. We show that the identifiability results can further benefit from nonstationarity with weaker identifiability conditions. This identifiability theory extends the prior work (Yao et al., 2021; Klindt et al., 2020; Hyvarinen & Morioka, 2017) by allowing causal relations between latent factors, instead of independent sources, and partitioned subspaces with both fixed and/or changing causal influences. The proofs are provided in Appendix S1.

4.1. Identifiability under Fixed Causal Dynamics

Let $\eta_{kt} \triangleq \log p(z_{k,t} | \mathbf{z}_{t-1})$. Assume that η_{kt} is twice differentiable in $z_{k,t}$ and is differentiable in $z_{l,t-1}$, $l = 1, 2, \dots, n$. Note that the parents of $z_{k,t}$ may be only a subset of \mathbf{z}_{t-1} ; if $z_{l,t-1}$ is not a parent of $z_{k,t}$, then $\frac{\partial \eta_{kt}}{\partial z_{l,t-1}} = 0$. Below we provide a *sufficient condition* for the identifiability of \mathbf{z}_t , followed by a discussion of specific unidentifiable and identifiable cases to illustrate how general it is.

Theorem 1 (Identifiability under a Fixed Temporal Causal

Model). Suppose there exists invertible function \hat{g} that maps \mathbf{x}_t to $\hat{\mathbf{z}}_t$, i.e.,

$$\hat{\mathbf{z}}_t = \hat{g}(\mathbf{x}_t) \quad (3)$$

such that the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. Let

$$\begin{aligned} \mathbf{v}_{k,t} &\triangleq \left(\frac{\partial^2 \eta_{kt}}{\partial z_{k,t} \partial z_{1,t-1}}, \frac{\partial^2 \eta_{kt}}{\partial z_{k,t} \partial z_{2,t-1}}, \dots, \frac{\partial^2 \eta_{kt}}{\partial z_{k,t} \partial z_{n,t-1}} \right)^\top \\ \hat{\mathbf{v}}_{k,t} &\triangleq \left(\frac{\partial^3 \eta_{kt}}{\partial z_{k,t}^2 \partial z_{1,t-1}}, \frac{\partial^3 \eta_{kt}}{\partial z_{k,t}^2 \partial z_{2,t-1}}, \dots, \frac{\partial^3 \eta_{kt}}{\partial z_{k,t}^2 \partial z_{n,t-1}} \right)^\top. \end{aligned} \quad (4)$$

If for each value of \mathbf{z}_t , $\mathbf{v}_{1,t}, \hat{\mathbf{v}}_{1,t}, \mathbf{v}_{2,t}, \hat{\mathbf{v}}_{2,t}, \dots, \mathbf{v}_{n,t}, \hat{\mathbf{v}}_{n,t}$, as $2n$ vector functions in $z_{1,t-1}, z_{2,t-1}, \dots, z_{n,t-1}$, are linearly independent, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.

The linear independence condition in Theorem 1 is the core condition to guarantee the identifiability of \mathbf{z}_t from the observed \mathbf{x}_t . To make this condition more intuitive, below we consider specific unidentifiable cases, in which there is no temporal dependence in \mathbf{z}_t or the noise terms in \mathbf{z}_t are additive Gaussian, and two identifiable cases, in which \mathbf{z}_t has additive, heterogeneous noise or follows some linear, non-Gaussian temporal process.

Let us start with two unidentifiable cases. In case N1, \mathbf{t}_t is an independent and identically distributed (i.i.d.) process, i.e., there is no causal influence from any component of \mathbf{z}_{t-1} to any $z_{k,t}$. In this case, $\mathbf{v}_{k,t}$ and $\hat{\mathbf{v}}_{k,t}$ (defined in Eq. 4) are always $\mathbf{0}$ for $k = 1, 2, \dots, n$, since $p(z_{k,t} | \mathbf{z}_{t-1})$ does not involve \mathbf{z}_{t-1} . So the linear independence condition is violated. In fact, this is the regular nonlinear ICA problem with i.i.d. data, and it is well-known that the underlying independent variables are not identifiable (Hyvärinen & Pajunen, 1999).

In case N2, all $z_{k,t}$ follow an additive noise model with Gaussian noise terms, i.e.,

$$\mathbf{z}_t = \mathbf{q}(\mathbf{z}_{t-1}) + \epsilon_t, \quad (5)$$

where \mathbf{q} is a transformation and the components of the Gaussian vector ϵ_t are independent and also independent from \mathbf{z}_{t-1} . Then $\frac{\partial^2 \eta_{kt}}{\partial z_{k,t}^2}$ is constant, and $\frac{\partial^3 \eta_{kt}}{\partial z_{k,t}^2 \partial z_{l,t-1}} \equiv 0$, violating the linear independence condition. In the following proposition we give some alternative solutions and verify the unidentifiability in this case.

Proposition 1 (Unidentifiability under Gaussian Noise). Suppose $\mathbf{x}_t = g(\mathbf{z}_t)$ was generated by Eq. 5, where the components of ϵ_t are mutually independent Gaussian and also independent from \mathbf{z}_{t-1} . Then any $\hat{\mathbf{z}}_t = \mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{z}_t$, where \mathbf{D}_1 is an arbitrary non-singular diagonal matrix, \mathbf{U} is an arbitrary orthogonal matrix, and \mathbf{D}_2 is a diagonal matrix with $\text{Var}^{-1/2}(\epsilon_{k,t})$ as its k^{th} diagonal entry, is a

valid solution to satisfy the condition that the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$.

Roughly speaking, for a randomly chosen conditional density function $p(z_{k,t} | \mathbf{z}_{t-1})$ in which $z_{k,t}$ is not independent from \mathbf{z}_{t-1} (i.e., there is temporal dependence in the latent processes) and which does not follow an additive noise model with Gaussian noise, the chance for its specific second- and third-order partial derivatives to be linearly dependent is slim. Now let us consider two cases in which the latent temporally processes \mathbf{z}_t are naturally identifiable under some technical conditions. First consider case Y1, where $z_{k,t}$ follows a heterogeneous noise process, in which the noise variance depends on its parents:

$$z_{k,t} = q_k(\mathbf{z}_{t-1}) + \frac{1}{b_k(\mathbf{z}_{t-1})} \epsilon_{k,t}. \quad (6)$$

Here we assume $\epsilon_{k,t}$ is standard Gaussian and $\epsilon_{1,t}, \epsilon_{2,t}, \dots, \epsilon_{n,t}$ are mutually independent and independent from \mathbf{z}_{t-1} . $\frac{1}{b_k}$, which depends on \mathbf{z}_{t-1} , is the standard deviation of the noise in $z_{k,t}$. (For conciseness, we drop the argument of b_k and q_k when there is no confusion.) Note that in this model, if q_k is 0 for all $k = 1, 2, \dots, n$, it reduces to a multiplicative noise model. The identifiability result of \mathbf{z}_t is established in the following proposition.

Corollary 1 (Identifiability under Heterogeneous Noise). *Suppose $\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$ was generated according to Eq. 6, and Eq. 3 holds true. If $b_k \cdot \frac{\partial b_k}{\partial \mathbf{z}_{t-1}}$ and $b_k \cdot \frac{\partial b_k}{\partial \mathbf{z}_{t-1}}(z_{k,t} - q_k) - b_k^2 \cdot \frac{\partial q_k}{\partial \mathbf{z}_{t-1}}$, with $k = 1, 2, \dots, n$, which are in total $2n$ function vectors in \mathbf{z}_{t-1} , are linearly independent, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.*

Let us then consider another special case, denoted by Y2, with linear, non-Gaussian temporal model for \mathbf{z}_t : the latent processes follow Eq. 5, with \mathbf{q} being a linear transformation and $\epsilon_{k,t}$ following a particular class of non-Gaussian distributions. The following corollary shows that \mathbf{z}_t is identifiable as long as each $z_{k,t}$ receives causal influences from some components of \mathbf{z}_{t-1} .

Corollary 2 (Identifiability under a Specific Linear, Non-Gaussian Model for Latent Processes). *Suppose $\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$ was generated according to Eq. 5, in which \mathbf{q} is a linear transformation and for each $z_{k,t}$, there exists at least one k' such that $c_{k,k'} \triangleq \frac{\partial z_{k,t}}{\partial z_{k',t-1}} \neq 0$. Assume the noise term $\epsilon_{k,t}$ follows a zero-mean generalized normal distribution:*

$$p(\epsilon_{k,t}) \propto e^{-\lambda |\epsilon_{k,t}|^\beta}, \text{ with positive } \lambda \text{ and } \beta > 2 \text{ and } \beta \neq 3. \quad (7)$$

If Eq. 3 holds, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.

4.2. Further Benefits from Changing Causal Influences

Yao et al. (2021) established the identifiability of the la-

tent temporal causal processes \mathbf{z}_t in certain nonstationary cases, under the condition that the noise term in each $z_{k,t}$, relative to its parents in \mathbf{z}_{t-1} , changes across m contexts corresponding to $\mathbf{u} = u_1, u_2, \dots, u_m$. Here we show that the identifiability result shown in the previous section can further benefit from nonstationarity of the causal model, and our identifiability condition is generally much weaker than that in (Yao et al., 2021): we allow changes in the noise term or causal influence on $z_{k,t}$ from its parents in \mathbf{z}_{t-1} , and our ‘‘sufficient variability’’ condition is just a necessary condition for that in (Yao et al., 2021) because of the additional information that one can leverage. Let $\mathbf{v}_{k,t}(u_r)$ be $\mathbf{v}_{k,t}$, which is defined in Eq. 4, in the u_r context. Similarly, Let $\hat{\mathbf{v}}_{k,t}(u_r)$ be $\hat{\mathbf{v}}_{k,t}$ in the u_r context. Let

$$\mathbf{s}_{k,t} \triangleq \left(\mathbf{v}_{k,t}(u_1)^\top, \dots, \mathbf{v}_{k,t}(u_m)^\top, \Delta_2^2, \dots, \Delta_m^2 \right)^\top,$$

$$\hat{\mathbf{s}}_{k,t} \triangleq \left(\hat{\mathbf{v}}_{k,t}(u_1)^\top, \dots, \hat{\mathbf{v}}_{k,t}(u_m)^\top, \Delta_2, \dots, \Delta_m \right)^\top,$$

where $\Delta_i^2 = \frac{\partial^2 \eta_{kt}(u_i)}{\partial z_{k,t}^2} - \frac{\partial^2 \eta_{kt}(u_{i-1})}{\partial z_{k,t}^2}$ and $\Delta_i = \frac{\partial \eta_{kt}(u_i)}{\partial z_{k,t}} - \frac{\partial \eta_{kt}(u_{i-1})}{\partial z_{k,t}}$. As provided below, in our case, the identifiability of \mathbf{z}_t is guaranteed by the linear independence of the whole function vectors $\mathbf{s}_{k,t}$ and $\hat{\mathbf{s}}_{k,t}$, with $k = 1, 2, \dots, n$. However, the identifiability result in (Yao et al., 2021) relies on the linear independence of only the last $m - 1$ components of $\mathbf{s}_{k,t}$ and $\hat{\mathbf{s}}_{k,t}$ with $k = 1, 2, \dots, n$; this linear independence is generally a much stronger condition.

Theorem 2 (Identifiability under Changing Causal Dynamics). *Suppose $\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$ and that the conditional distribution $p(z_{k,t} | \mathbf{z}_{t-1})$ may change across m values of the context variable \mathbf{u} , denoted by u_1, u_2, \dots, u_m . Suppose the components of \mathbf{z}_t are mutually independent conditional on \mathbf{z}_{t-1} in each context. Assume that the components of $\hat{\mathbf{z}}_t$ produced by Eq. 3 are also mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. If the $2n$ function vectors $\mathbf{s}_{k,t}$ and $\hat{\mathbf{s}}_{k,t}$, with $k = 1, 2, \dots, n$, are linearly independent, then $\hat{\mathbf{z}}_t$ is a permuted invertible component-wise transformation of \mathbf{z}_t .*

Theorem 3 (Identifiability under Observation Changes). *Suppose $\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$ and that the conditional distribution $p(z_{k,t} | \mathbf{u})$ may change across m values of the context variable \mathbf{u} , denoted by u_1, u_2, \dots, u_m . Suppose the components of \mathbf{z}_t are mutually independent conditional on \mathbf{u} in each context. Assume that the components of $\hat{\mathbf{z}}_t$ produced by Eq. 3 are also mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. If the $2n$ function vectors $\mathbf{s}_{k,t}$ and $\hat{\mathbf{s}}_{k,t}$, with $k = 1, 2, \dots, n$, are linearly independent, then $\hat{\mathbf{z}}_t$ is a permuted invertible component-wise transformation of \mathbf{z}_t .*

Corollary 3 (Identifiability under Modular Distribution Shifts). *Assume the data generating process in Eq. 1. If the three partitioned latent components $\mathbf{z}_t = (\mathbf{z}_t^{\text{fix}}, \mathbf{z}_t^{\text{chg}}, \mathbf{z}_t^{\text{obs}})$ respectively satisfy the conditions in Theorem 1, Theorem 2, and Theorem 3, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.*




Figure 2. **Framework overview:** Encoder (A) and Decoder (E) using MLP or CNN for specific data types; (B) Factorized inference network that approximates the posteriors of latent variables $\hat{\mathbf{z}}_t$, (C) Change factor representation learning with embedding layer, and (D) Modular prior network that factorizes distribution changes, i.e., fixed causal dynamics, changing causal dynamics, and global changes, into three partitioned latent causal processes constraint (Thm 1,2 and 3).

5. LiLY: Learning Latent Causal Dynamics

Given our identifiability results, we propose Learning Latent causal dYnamics (LiLY) framework to estimate the latent causal dynamics under modular distribution shifts, by extending Sequential Variational Auto-Encoders (Li & Mandt, 2018) with tailored modules to model different distribution shifts, and enforcing the conditions in Sec. 4 as constraints. We give the estimation procedure of the latent causal dynamics model in Eq. 1. The model architecture is showcased in Fig. 2. The framework has the following three major components. The implementation details are in Appendix S3.1.

Modular Prior Network The core component of the estimation is to learn the three types of distribution changes. We leverage the partitioned estimated latent subspaces $\hat{\mathbf{z}}_t = (\hat{\mathbf{z}}_t^{\text{fix}}, \hat{\mathbf{z}}_t^{\text{chg}}, \hat{\mathbf{z}}_t^{\text{obs}})$ and model their distribution changes in conditional priors. In particular, for (1) fixed causal dynamics processes $\hat{\mathbf{z}}_t^{\text{fix}}$, their transition priors are obtained by first learning inverse transition functions f_s^{-1} that take the estimated latent variables and output random noise terms, and applying the change of variables formula to the transformation: $p(\hat{\mathbf{z}}_{s,t}^{\text{fix}} | \hat{\mathbf{z}}_{\text{Hx}}) = p_{\epsilon_s}(f_s^{-1}(\hat{\mathbf{z}}_{s,t}^{\text{fix}}, \hat{\mathbf{z}}_{\text{Hx}})) \left| \frac{\partial f_s^{-1}}{\partial \hat{\mathbf{z}}_{s,t}^{\text{fix}}} \right|$; (2) for changing causal dynamics, we evaluate $p(\hat{\mathbf{z}}_{s,t}^{\text{chg}} | \hat{\mathbf{z}}_{\text{Hx}}, \mathbf{u}_k) = p_{\epsilon_c}(f_c^{-1}(\hat{\mathbf{z}}_{c,t}^{\text{chg}}, \hat{\mathbf{z}}_{\text{Hx}}, \hat{\theta}_k^{\text{dyn}})) \left| \frac{\partial f_c^{-1}}{\partial \hat{\mathbf{z}}_{c,t}^{\text{chg}}} \right|$ by learning a holistic inverse dynamics f_c^{-1} that takes the estimated change factors for dynamics $\hat{\theta}_k^{\text{dyn}}$ as inputs, and similarly for (3) observation changes $\hat{\mathbf{z}}_t^{\text{obs}}$, we learn to project them to invariant noise terms by f_o^{-1} which takes the change factors θ_k^{obs} as arguments, and obtains $p(\hat{\mathbf{z}}_{o,t}^{\text{obs}} | \mathbf{u}_k) = p_{\epsilon_o}(f_o^{-1}(\hat{\mathbf{z}}_{o,t}^{\text{obs}}, \hat{\theta}_k^{\text{obs}})) \left| \frac{\partial f_o^{-1}}{\partial \hat{\mathbf{z}}_{o,t}^{\text{obs}}} \right|$ as the prior. Two immediate benefits are brought by learning the inverse dynamics

function: (1) compared with forward prediction, learning to recover the noise is more appealing for nonparametric processes, in the sense that using forward prediction with a fixed loss function, one cannot model latent processes without parametric forms. For example, heterogeneous noise model in the latent processes (Eq. 6) cannot be estimated by forward prediction with L_2 loss, and (2) (conditional) mutual independence of the estimated latent variables $\hat{\mathbf{z}}_t$ can be easily enforced by summing up all estimated component densities when obtaining the joint $p(\mathbf{z}_t | \mathbf{z}_{\text{Hx}}, \mathbf{u})$. Furthermore, given that the Jacobian is lower-triangular because of mutually-independent noise, we can efficiently compute its determinant as the product of diagonal terms. Learning the inverse dynamics won't cause computational issues as well.

Factorized Inference We infer the posteriors of each time step $q(\hat{\mathbf{z}}_t | \mathbf{x}_t, \hat{\theta}_k^{\text{obs}})$, using only the observation at that time step and with the change factors θ_k^{obs} allocated for global changes. Recurrent inference (Yao et al., 2021) is not used because in Eq. 1, \mathbf{x}_t preserves all the information of the current system states so the joint probability $q(\hat{\mathbf{z}}_{1:T} | \mathbf{x}_{1:T})$ can be factorized into product of these terms.

Change Representation To allow transfer in a new environment, instead of learning separate transformation or dynamics in each domain (Yao et al., 2021), we learn to embed domain index \mathbf{u}_k into low-dimensional change factors $(\hat{\theta}_k^{\text{dyn}}, \hat{\theta}_k^{\text{obs}})$ and insert them as inputs to the inverse dynamics function, or the encoder/decoder function, respectively. We make the best use of the multiple domains created by distribution changes, and efficiently correct distribution shifts with only a few samples from the target domain, since only the low-dimensional change factors need to be re-estimated.

Algorithm 1 CORRECTSHIFT($\hat{f}_i, \hat{\theta}_k, \hat{p}(\epsilon_i), \hat{g}$)

Input: Data $\{\mathbf{x}_t\}_{t=1}^T$ for a new environment k^* .
Initialize change factors $\hat{\theta}_{k^*}$ for environment k^* and set as model parameters. Freeze other parameters.
while $\hat{\theta}_{k^*}$ does not converge **do**
 $\hat{\theta}_{k^*} \leftarrow \hat{\theta}_{k^*} + \alpha \nabla_{\hat{\theta}_{k^*}} \mathcal{L}_{\text{ELBO}}(\{\mathbf{x}_t\}_{t=1}^T, \hat{\theta}_{k^*})$
end while
return the corrected model $(\hat{f}_i, \hat{\theta}_{k^*}, \hat{p}(\epsilon_i), \hat{g})$

We train the VAE using the ELBO objective $\mathcal{L}_{\text{ELBO}} = \frac{1}{N} \sum_{i \in N} \mathcal{L}_{\text{Recon}} - \beta \mathcal{L}_{\text{KLD}}$, in which we use mean-squared error (MSE) for the reconstruction likelihood $\mathcal{L}_{\text{Recon}}$, and use sampling approach to estimate the \mathcal{L}_{KLD} , since with a learned modular prior, KLD does not have an explicit form.

6. Experimental Results

We comparatively evaluate LiLY on a number of time-series datasets. We aim to answer the following questions:

1. Does LiLY reliably learn temporally-causal latent processes from scratch under different distribution changes?
2. Is history/nonstationary information necessary for the identifiability of latent causal variables?
3. How do assumptions in previous nonlinear ICA (i.e., independent sources, nonstationary noises) distort the identifiability results if there are time-delayed causal relations between the latent factors, or if the nonstationary cases involve time-varying causal relations?

Evaluation Metrics We report Mean Correlation Coefficient (MCC) on the validation dataset. MCC is a standard metric in the ICA literature for continuous variables which measure the identifiability of the learned latent causal processes. MCC is close to 1 when latent variables are identifiable up to permutation and componentwise invertible transformation in the noiseless case. To evaluate the transfer performances, we visualize the learned change factors θ_k with a few samples (16 in our experiments) from the target segment, and compare them with the true change factors.

Baselines We used nonlinear ICA baselines: (1) BetaVAE (Higgins et al., 2016) which ignores both history and nonstationarity information; (2) iVAE (Khemakhem et al., 2020) and TCL (Hyvarinen & Morioka, 2016) which leverage nonstationarity to establish identifiability but assumes independent factors, and (3) SlowVAE (Klindt et al., 2020) and PCL (Hyvarinen & Morioka, 2017) which exploits temporal constraints but assumes independent sources and stationary processes, and (4) LEAP (Yao et al., 2021) which assumes nonstationary, causal processes but models all types of distribution changes in nonstationary noise.

6.1. Synthetic Experiments

We generate synthetic datasets that satisfy our identifiability conditions in the theorems following the procedures described in Appendix S2.1.1. In particular, we consider three representative simulation settings respectively to validate the identifiability results under fixed causal dynamics, changing causal dynamics, and modular distribution shift which contains fixed dynamics, changing dynamics and global changes together in the latent processes. For synthetic datasets with fixed and changing causal dynamics, we set latent size $n = 8$. For the modular shift dataset, we add one dimension for global observation changes. The lag number of the process is set to $L = 2$. The mixing function g is a random three-layer MLP with LeakyReLU units.

6.1.1. COMPARISONS WITH BASELINES

As presented in Table 2, our framework can recover the latent processes under fixed dynamics (heterogeneous noise model), under changing causal dynamics and under modular distribution shifts. The baselines that do not exploit history or nonstationarity cannot recover the latent processes. SlowVAE and PCL distort the results due to independent source assumptions. LEAP, although exploits both the history and nonstationarity information, considers only limited nonstationary cases in which the noise distributions are changing, thus cannot recover the causal latent processes under changing causal dynamics and modular distribution shifts.

Table 2. MCC scores and their standard deviations for the three simulation settings over 3 random seeds. Note: The symbol “-” represents that this method is not applicable to this dataset.

	Simulation Setting		
	Fixed Causal Dynamics	Changing Causal Dynamics	Modular Shift
LiLY	0.954 \pm 0.009	0.958 \pm 0.017	0.993 \pm 0.001
LEAP	-	0.726 \pm 0.187	0.657 \pm 0.108
SlowVAE	0.411 \pm 0.022	0.511 \pm 0.062	0.406 \pm 0.045
PCL	0.516 \pm 0.043	0.599 \pm 0.041	0.564 \pm 0.049
i-VAE	-	0.581 \pm 0.083	0.557 \pm 0.005
TCL	-	0.399 \pm 0.021	0.297 \pm 0.078
β -VAE	0.353 \pm 0.001	0.523 \pm 0.009	0.433 \pm 0.045

6.1.2. FIXED CAUSAL DYNAMICS

We use a heterogeneous noise model in Eq. 6 with standard Gaussian noise terms as the latent processes. The transition function q_k is a random two-layer MLP with LeakyReLU units, and the b_k a mean function of its history \mathbf{z}_{t-1} . The latent dynamics is identifiable since the history modulates both the mean and variance of the transition distribution. Fig. 3 gives the results on heterogeneous noise datasets. The latent processes are successfully recovered, as indicated by high MCC for the casually-related factors. Panel (b) suggests that the latent causal variables are estimated up to




Figure 3. Fixed Causal Dynamics.




Figure 4. Changing Causal Dynamics.

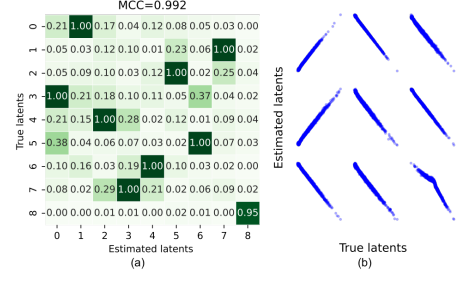


Figure 5. Modular Distribution Shifts.

permutation and componentwise invertible transformation.

6.1.3. CHANGING CAUSAL DYNAMICS

We use a Gaussian additive noise model with changes in the influencing strength as the latent processes. It has been shown in **Proposition 1** that this latent dynamics, if without changes, is unidentifiable. We use a random two-layer MLP with LeakyReLU units as the transition function. To add changes, we vary the values of the first layer of the MLP across segments. In Fig. 4, the latent processes are successfully recovered, as indicated by high MCC and the recovered latent factors, which shows the further benefits brought by distribution changes.

6.1.4. MODULAR DISTRIBUTION SHIFT

To partition the latent space, we allocate six latent processes (variables) to fixed heterogeneous noise dynamics as in Sec. 6.1.2, two latent processes following changing causal dynamics in Sec. 6.1.3, and the last latent variable is sampled from Gaussian distribution whose mean and variance vary by segment index. The results for the modular shift dataset are in Fig. 5. Similarly, the latent processes are recovered, as indicated by (a) high MCC and (b) invertible mappings between the estimated and true latent variables.

6.2. Causal Discovery from Videos

We evaluate LiLY on the modified cartpole (Huang et al., 2021) video dataset and compare the performances with LEAP. Modified Cartpole is a nonlinear dynamical system with cart positions x_t and pole angles θ_t as the true state variables. The dataset has 6 source domains with different gravity values $g = \{5, 10, 15, 20, 25, 30\}$. Together with the 2 discrete actions (i.e., left and right), we have 12 segments of data with changing causal dynamics. We fit LiLY with two-dimensional change factors θ_k^{dyn} . We set the latent size $n = 8$ and the lag number $L = 2$. In Fig. 6, the latent causal processes are recovered, as seen from (a) high MCC for the latent causal processes; (b) the latent factors are estimated up to componentwise transformation; (c) the latent traversals confirm that the two latent causal variables

correspond to the position and pole angle, and (d) our model outperforms LEAP, which only considers nonstationarity in noise distributions.

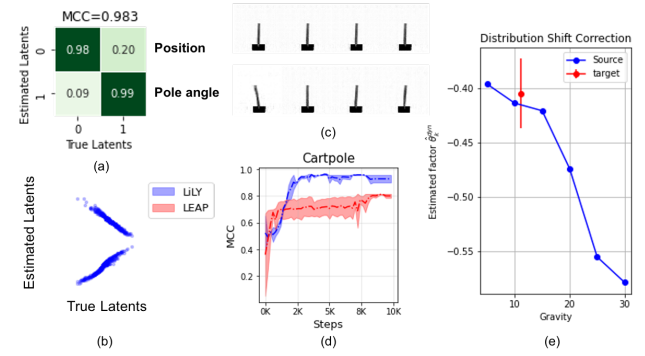


Figure 6. Modified Cartpole dataset results: (a) MCC for causally-related factors; (b) scatterplots between estimated and true factors; (c) latent traversal on a fixed video frame; (d) baseline comparisons with LEAP (Yao et al., 2021), and (e) domain shift correction.

We visualize the learned change factors in panel (e). We also plot the estimated change factor in a new environment ($g = 11.15$) with 16 samples over 3 random seeds. The positive association between the true and estimated change factors indicates that meaningful change representation in the source environment is learned. The mean of the estimated change factor in the target environment is aligned with the true value, showing that we can efficiently adapt with a few samples from the target environment.

7. Conclusion

We established theories for the identifiability of latent causal dynamics from their observed nonlinear mixtures under fixed dynamics and unknown distribution shifts. The main limitation of this work is that instantaneous causal influences between latent causal processes are not allowed. We acknowledge that there exist situations where the resolution is low and there appears to be instantaneous dependence. However, there are several pieces of work dealing with causal discovery from measured time series in such situa-

tions; see. e.g., (Granger, 1987; Gong* et al., 2015; Danks & Plis, 2013; Gong et al., 2017). Extending our theories and framework to address the issue of instantaneous dependency or instantaneous causal relations in the latent structure will be one line of our future work.

References

- Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., and Pal, C. A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*, 2019.
- Berzuini, C., Dawid, P., and Bernardinell, L. *Causality: Statistical perspectives and applications*. John Wiley & Sons, 2012.
- Danks, D. and Plis, S. Learning causal structure from under-sampled time series. In *JMLR: Workshop and Conference Proceedings*, pp. 1–10, 2013.
- Entner, D. and Hoyer, P. O. On causal discovery from time series data using fci. *Probabilistic graphical models*, pp. 121–128, 2010.
- Friston, K. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS biology*, 7(2):e1000033, 2009.
- Ghysels, E., Hill, J. B., and Motegi, K. Testing for granger causality with mixed frequency data. *Journal of Econometrics*, 192(1):207–230, 2016.
- Gong*, M., Zhang*, K., Tao, D., Geiger, P., and Schölkopf, B. Discovering temporal causal relations from subsampled data. In *Proc. 32th International Conference on Machine Learning (ICML 2015)*, 2015.
- Gong, M., Zhang, K., Schölkopf, B., Glymour, C., and Tao, D. Causal discovery from temporally aggregated time series. In *Proc. Conference on Uncertainty in Artificial Intelligence (UAI’17)*, 2017.
- Granger, C. W. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pp. 424–438, 1969.
- Granger, C. W. Implications of aggregation with common factors. *Econometric Theory*, 3(02):208–222, 1987.
- Hälvä, H. and Hyvarinen, A. Hidden markov nonlinear ica: Unsupervised learning from nonstationary time series. In *Conference on Uncertainty in Artificial Intelligence*, pp. 939–948. PMLR, 2020.
- Hälvä, H., Corff, S. L., Lehéricy, L., So, J., Zhu, Y., Gassiat, E., and Hyvarinen, A. Disentangling identifiable features from noisy data with structured nonlinear ica. *arXiv preprint arXiv:2106.09620*, 2021.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- Huang, B., Zhang, K., Zhang, J., Ramsey, J. D., Sanchez-Romero, R., Glymour, C., and Schölkopf, B. Causal discovery from heterogeneous/nonstationary data. *J. Mach. Learn. Res.*, 21(89):1–53, 2020.
- Huang, B., Feng, F., Lu, C., Magliacane, S., and Zhang, K. Adarl: What, where, and how to adapt in transfer reinforcement learning. *arXiv preprint arXiv:2107.02729*, 2021.
- Hyvarinen, A. and Morioka, H. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. *Advances in Neural Information Processing Systems*, 29: 3765–3773, 2016.
- Hyvarinen, A. and Morioka, H. Nonlinear ica of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pp. 460–469. PMLR, 2017.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks*, 12(3):429–439, 1999.
- Hyvarinen, A., Sasaki, H., and Turner, R. Nonlinear ica using auxiliary variables and generalized contrastive learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 859–868. PMLR, 2019.
- Ke, N. R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Schölkopf, B., Mozer, M. C., Pal, C., and Bengio, Y. Learning neural causal models from unknown interventions. *arXiv preprint arXiv:1910.01075*, 2019.
- Khemakhem, I., Kingma, D., Monti, R., and Hyvarinen, A. Variational autoencoders and nonlinear ica: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pp. 2207–2217. PMLR, 2020.
- Klindt, D., Schott, L., Sharma, Y., Ustyuzhaninov, I., Brendel, W., Bethge, M., and Paiton, D. Towards nonlinear disentanglement in natural data with temporal sparse coding. *arXiv preprint arXiv:2007.10930*, 2020.
- Li, Y. and Mandt, S. Disentangled sequential autoencoder. *arXiv preprint arXiv:1803.02991*, 2018.
- Li, Y., Torralba, A., Anandkumar, A., Fox, D., and Garg, A. Causal discovery in physical systems from videos. *arXiv preprint arXiv:2007.00631*, 2020.
- Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled

- representations. In *international conference on machine learning*, pp. 4114–4124. PMLR, 2019.
- Löwe, S., Madras, D., Zemel, R., and Welling, M. Amortized causal discovery: Learning to infer causal graphs from time-series data. *arXiv preprint arXiv:2006.10833*, 2020.
- Malinsky, D. and Spirtes, P. Causal structure learning from multivariate time series in settings with unmeasured confounding. In *Proceedings of 2018 ACM SIGKDD Workshop on Causal Discovery*, pp. 23–47. PMLR, 2018.
- Malinsky, D. and Spirtes, P. Learning the structure of a nonstationary vector autoregression. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2986–2994. PMLR, 2019.
- Murphy, K. P. et al. Dynamic bayesian networks. *Probabilistic Graphical Models, M. Jordan*, 7:431, 2002.
- Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Georgatzis, K., Beaumont, P., and Aragam, B. Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*, pp. 1595–1605. PMLR, 2020.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
- Pearl, J. et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 19, 2000.
- Sahoo, A., Shah, R., Panda, R., Saenko, K., and Das, A. Contrast and mix: Temporal contrastive video domain adaptation with background mixing. In *Advances in Neural Information Processing Systems*, 2021.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634, 2021.
- Silva, R., Scheines, R., Glymour, C., and Spirtes, P. Learning the structure of linear latent variable models. *Journal of Machine Learning Research*, 7:191–246, 2006.
- Sorrenson, P., Rother, C., and Köthe, U. Disentanglement by nonlinear ica with general incompressible-flow networks (gin). *arXiv preprint arXiv:2001.04872*, 2020.
- Spantini, A., Bigoni, D., and Marzouk, Y. M. Inference via low-dimensional couplings. *J. Mach. Learn. Res.*, 19: 66:1–66:71, 2018.
- Spearman, C. Pearson’s contribution to the theory of two factors. *British Journal of Psychology*, 19:95–101, 1928.
- Spirtes, P., Glymour, C., and Scheines, R. *Causation, Prediction, and Search*. Springer-Verlag Lectures in Statistics, 1993.
- Tank, A., Covert, I., Foti, N., Shojaie, A., and Fox, E. Neural granger causality. *arXiv preprint arXiv:1802.05842*, 2018.
- Xie, F., Cai, R., Huang, B., Glymour, C., Hao, Z., and Zhang, K. Generalized independent noise condition for estimating latent variable causal graphs. *arXiv preprint arXiv:2010.04917*, 2020.
- Yao, W., Sun, Y., Ho, A., Sun, C., and Zhang, K. Learning temporally causal latent processes from general temporal data. *arXiv preprint arXiv:2110.05428*, 2021.

Supplement to: “Learning Latent Causal Dynamics”

S1. Identifiability Theory

The observed variables were generated according to :

$$\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t), \quad (8)$$

in which \mathbf{g} is invertible, and z_{it} , as the i th component of \mathbf{z}_t , is generated by (some) components of \mathbf{z}_{t-1} and noise E_{it} . $E_{1t}, E_{2t}, \dots, E_{nt}$ are mutually independent. In other words, the components of \mathbf{z}_t are mutually independent conditional on \mathbf{z}_{t-1} . Let $\eta_{kt} \triangleq \log p(z_{kt} | \mathbf{z}_{t-1})$. Assume that $\eta_k(t)$ is twice differentiable in z_{kt} and is differentiable in $z_{l,t-1}$, $l = 1, 2, \dots, n$. Note that the parents of z_{kt} may be only a subset of \mathbf{z}_{t-1} ; if $z_{l,t-1}$ is not a parent of z_{kt} , then $\frac{\partial \eta_k}{\partial z_{l,t-1}} = 0$.

S1.1. Proof for Theorem 1

Theorem S1 (Identifiability under a Fixed Temporal Causal Model). *Suppose there exists invertible function \mathbf{f} , which is the estimated mixing function (i.e., we use \mathbf{f} and $\hat{\mathbf{g}}$ interchangeably in Appendix) that maps \mathbf{x}_t to $\hat{\mathbf{z}}_t$, i.e.,*

$$\hat{\mathbf{z}}_t = \mathbf{f}(\mathbf{x}_t) \quad (9)$$

such that the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. Let

$$\mathbf{v}_{kt} \triangleq \left(\frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{1,t-1}}, \frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{2,t-1}}, \dots, \frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{n,t-1}} \right)^\top, \quad \hat{\mathbf{v}}_{kt} \triangleq \left(\frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{1,t-1}}, \frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{2,t-1}}, \dots, \frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{n,t-1}} \right)^\top. \quad (10)$$

If for each value of \mathbf{z}_t , $\mathbf{v}_{1t}, \hat{\mathbf{v}}_{1t}, \mathbf{v}_{2t}, \hat{\mathbf{v}}_{2t}, \dots, \mathbf{v}_{nt}, \hat{\mathbf{v}}_{nt}$, as $2n$ vector functions in $z_{1,t-1}, z_{2,t-1}, \dots, z_{n,t-1}$, are linearly independent, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.

Proof. Combining Eq. 8 and Eq. 9 gives $\mathbf{z}_t = \mathbf{g}^{-1}(\mathbf{f}^{-1}(\hat{\mathbf{z}}_t)) = \mathbf{h}(\hat{\mathbf{z}}_t)$, where $\mathbf{h} \triangleq \mathbf{g}^{-1} \circ \mathbf{f}^{-1}$. Since both \mathbf{f} and \mathbf{g} are invertible, \mathbf{h} is invertible. Let \mathbf{H}_t be the Jacobian matrix of the transformation $\mathbf{h}(\hat{\mathbf{z}}_t)$, and denote by \mathbf{H}_{kit} its (k, i) th entry.

First, it is straightforward to see that if the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$, then for any $i \neq j$, \hat{z}_{it} and \hat{z}_{jt} are conditionally independent given $\hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\})$. Mutual independence of the components of $\hat{\mathbf{z}}_t$ conditional on $\hat{\mathbf{z}}_{t-1}$ implies that \hat{z}_{it} is independent from $\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\}$ conditional on $\hat{\mathbf{z}}_{t-1}$, i.e.,

$$p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1}) = p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\})).$$

At the same time, it also implies \hat{z}_{it} is independent from $\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}\}$ conditional on $\hat{\mathbf{z}}_{t-1}$, i.e.,

$$p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1}) = p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}\})).$$

Combining the above two equations gives $p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}\})) = p(\hat{z}_{it} | \hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\}))$, i.e., for $i \neq j$, \hat{z}_{it} and \hat{z}_{jt} are conditionally independent given $\hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\})$.

We then make use of the fact that if \hat{z}_{it} and \hat{z}_{jt} are conditionally independent given $\hat{\mathbf{z}}_{t-1} \cup (\hat{\mathbf{z}}_t \setminus \{\hat{z}_{it}, \hat{z}_{jt}\})$, then

$$\frac{\partial^2 \log p(\hat{\mathbf{z}}_t, \hat{\mathbf{z}}_{t-1})}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} = 0,$$

assuming the cross second-order derivative exists (Spantini et al., 2018). Since $p(\hat{\mathbf{z}}_t, \hat{\mathbf{z}}_{t-1}) = p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})p(\hat{\mathbf{z}}_{t-1})$ while $p(\hat{\mathbf{z}}_{t-1})$ does not involve \hat{z}_{it} or \hat{z}_{jt} , the above equality is equivalent to

$$\frac{\partial^2 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} = 0. \quad (11)$$

The Jacobian matrix of the mapping from $(\mathbf{x}_{t-1}, \hat{\mathbf{z}}_t)$ to $(\mathbf{x}_{t-1}, \mathbf{z}_t)$ is $\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ * & \mathbf{H}_t \end{bmatrix}$, where $*$ stands for a matrix, and the (absolute value of the) determinant of this Jacobian matrix is $|\mathbf{H}_t|$. Therefore $p(\hat{\mathbf{z}}_t, \mathbf{x}_{t-1}) = p(\mathbf{z}_t, \mathbf{x}_{t-1}) \cdot |\mathbf{H}_t|$. Dividing both sides of this equation by $p(\mathbf{x}_{t-1})$ gives

$$p(\hat{\mathbf{z}}_t | \mathbf{x}_{t-1}) = p(\mathbf{z}_t | \mathbf{x}_{t-1}) \cdot |\mathbf{H}_t|. \quad (12)$$

Since $p(\mathbf{z}_t | \mathbf{z}_{t-1}) = p(\mathbf{z}_t | \mathbf{g}(\mathbf{z}_{t-1})) = p(\mathbf{z}_t | \mathbf{x}_{t-1})$ and similarly $p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}) = p(\hat{\mathbf{z}}_t | \mathbf{x}_{t-1})$, Eq. 12 tells us

$$\log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}) = \log p(\mathbf{z}_t | \mathbf{z}_{t-1}) + \log |\mathbf{H}_t| = \sum_{k=1}^n \eta_{kt} + \log |\mathbf{H}_t|. \quad (13)$$

Its partial derivative w.r.t. \hat{z}_{it} is

$$\begin{aligned} \frac{\partial \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})}{\partial \hat{z}_{it}} &= \sum_{k=1}^n \frac{\partial \eta_{kt}}{\partial z_{kt}} \cdot \frac{\partial z_{kt}}{\partial \hat{z}_{it}} - \frac{\partial \log |\mathbf{H}_t|}{\partial \hat{z}_{it}} \\ &= \sum_{k=1}^n \frac{\partial \eta_{kt}}{\partial z_{kt}} \cdot \mathbf{H}_{kit} - \frac{\partial \log |\mathbf{H}_t|}{\partial \hat{z}_{it}}. \end{aligned}$$

Its second-order cross derivative is

$$\frac{\partial^2 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} = \sum_{k=1}^n \left(\frac{\partial^2 \eta_{kt}}{\partial z_{kt}^2} \cdot \mathbf{H}_{kit} \mathbf{H}_{kjt} + \frac{\partial \eta_{kt}}{\partial z_{kt}} \cdot \frac{\partial \mathbf{H}_{kit}}{\partial \hat{z}_{jt}} \right) - \frac{\partial^2 \log |\mathbf{H}_t|}{\partial \hat{z}_{it} \partial \hat{z}_{jt}}. \quad (14)$$

The above quantity is always 0 according to Eq. 11. Therefore, for each $l = 1, 2, \dots, n$ and each value $z_{l,t-1}$, its partial derivative w.r.t. $z_{l,t-1}$ is always 0. That is,

$$\frac{\partial^3 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})}{\partial \hat{z}_{it} \partial \hat{z}_{jt} \partial z_{l,t-1}} = \sum_{k=1}^n \left(\frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{l,t-1}} \cdot \mathbf{H}_{kit} \mathbf{H}_{kjt} + \frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{l,t-1}} \cdot \frac{\partial \mathbf{H}_{kit}}{\partial \hat{z}_{jt}} \right) \equiv 0, \quad (15)$$

where we have made use of the fact that entries of \mathbf{H}_t do not depend on $z_{l,t-1}$.

If for any value of \mathbf{z}_t , $\mathbf{v}_{1t}, \hat{\mathbf{v}}_{1t}, \mathbf{v}_{2t}, \hat{\mathbf{v}}_{2t}, \dots, \mathbf{v}_{nt}, \hat{\mathbf{v}}_{nt}$ are linearly independent, to make the above equation hold true, one has to set $\mathbf{H}_{kit} \mathbf{H}_{kjt} = 0$ or $i \neq j$. That is, in each row of \mathbf{H}_t there is only one non-zero entry. Since h is invertible, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$. \square

The linear independence condition in Theorem S1 is the core condition to guarantee the identifiability of \mathbf{z}_t from the observed \mathbf{x}_t . Roughly speaking, for a randomly chosen conditional density function $p(z_{kt} | \mathbf{z}_{t-1})$, the chance for this constraint to hold on its second- and third-order partial derivatives is slim. For illustrative purposes, below we make this claim more precise, by considering a specific unidentifiable case, in which the noise terms in \mathbf{z}_t are additive Gaussian, and two identifiable cases, in which \mathbf{z}_t has additive, heterogeneous noise or follows some linear, non-Gaussian temporal process.

Let us start with an unidentifiable case. If all z_{kt} follow the additive noise model with Gaussian noise terms, i.e.,

$$\mathbf{z}_t = \mathbf{q}(\mathbf{z}_{t-1}) + \mathbf{E}_t, \quad (16)$$

where \mathbf{q} is a transformation and the components of the Gaussian vector \mathbf{E}_t are independent and also independent from \mathbf{z}_{t-1} . Then $\frac{\partial^2 \eta_{kt}}{\partial z_{kt}^2}$ is constant, and $\frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{l,t-1}} \equiv 0$, violating the linear independence condition. In the following proposition we give some alternative solutions and verify the unidentifiability in this case.

Proposition S1 (Unidentifiability under Gaussian noise). *Suppose \mathbf{x}_t was generated according to Eq. 8 and Eq. 16, where the components of \mathbf{E}_t are mutually independent Gaussian and also independent from \mathbf{z}_{t-1} . Then any $\hat{\mathbf{z}}_t = \mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{z}_t$, where \mathbf{D}_1 is an arbitrary non-singular diagonal matrix, \mathbf{U} is an arbitrary orthogonal matrix, and \mathbf{D}_2 is a diagonal matrix with $\text{Var}^{-1/2}(E_{kt})$ as its k th diagonal entry, is a valid solution to satisfy the condition that the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$.*

Proof. In this case we have

$$\hat{\mathbf{z}}_t = \mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{q}(\mathbf{z}_{t-1}) + \mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{E}_t.$$

It is easy to verify that the components of $\mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{E}_t$ are mutually independent and are independent from $\mathbf{D}_1 \mathbf{U} \mathbf{D}_2 \cdot \mathbf{q}(\mathbf{z}_{t-1})$. As a consequence, $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. \square

Now let us consider some cases in which the latent temporally processes \mathbf{z}_t are naturally identifiable under some technical conditions. Let us first consider the case where z_{kt} follows a heterogeneous noise process, in which the noise variance depends on its parents:

$$a_{kt} = q_k(\mathbf{z}_{t-1}) + \frac{1}{b_k(\mathbf{z}_{t-1})} E_{kt}. \quad (17)$$

Here we assume E_{kt} is standard Gaussian and $E_{1t}, E_{2t}, \dots, E_{nt}$ are mutually independent and independent from \mathbf{z}_{t-1} . $\frac{1}{b_k}$, which depends on \mathbf{z}_{t-1} , is the standard deviation of the noise in z_{kt} . (For conciseness, we drop the argument of b_k and q_k when there is no confusion.) Note that in this model, if q_k is 0 for all $k = 1, 2, \dots, n$, it reduces to a multiplicative noise model. The identifiability result of \mathbf{z}_t is established in the following proposition.

Corollary S1 (Identifiability under Heterogeneous Noise). *Suppose \mathbf{x}_t was generated according to Eq. 8 and Eq. 17. Suppose Eq. 9 holds true. If $b_k \cdot \frac{\partial b_k}{\partial \mathbf{z}_{t-1}}$ and $b_k \cdot \frac{\partial b_k}{\partial \mathbf{z}_{t-1}}(z_{kt} - q_k) - b_k^2 \cdot \frac{\partial q_k}{\partial \mathbf{z}_{t-1}}$, with $k = 1, 2, \dots, n$, which are in total $2n$ function vectors in \mathbf{z}_{t-1} , are linearly independent, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.*

Proof. Under the assumptions, one can see that

$$\eta_{kt} = \log p(z_{kt} | \mathbf{z}_{t-1}) = -\frac{1}{2} \log(2\pi) + \log b_k - \frac{b_k^2}{2} (z_{kt} - q_k)^2.$$

Consequently, one can find

$$\begin{aligned} \frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{l,t-1}} &= -b_k \cdot \frac{\partial b_k}{\partial z_{l,t-1}}, \\ \frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{l,t-1}} &= -b_k \cdot \frac{\partial b_k}{\partial z_{l,t-1}} (z_{kt} - q_k) + b_k^2 \cdot \frac{\partial q_k}{\partial z_{l,t-1}}. \end{aligned}$$

Then the linear independence of \mathbf{v}_{kt} and $\hat{\mathbf{v}}_{kt}$ (defined in Eq. 10), with $k = 1, 2, \dots, n$, reduces to the linear independence condition in this proposition. Theorem S1 then implies that \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$. \square

Let us then consider another special case, with linear, non-Gaussian temporal model for \mathbf{z}_t : the latent processes follow Eq. 16, with \mathbf{q} being a linear transformation and E_{kt} following a particular class of non-Gaussian distributions. The following corollary shows that \mathbf{z}_t is identifiable as long as each z_{kt} receives causal influences from some components of \mathbf{z}_{t-1} .

Corollary S2 (Identifiability under a Specific Linear, Non-Gaussian Model for Latent Processes). *Suppose \mathbf{x}_t was generated according to Eq. 8 and Eq. 16, in which \mathbf{q} is a linear transformation and for each z_{kt} , there exists at least one k' such that $c_{kk'} \triangleq \frac{\partial z_{kt}}{\partial z_{k',t-1}} \neq 0$. Assume the noise term E_{kt} follows a zero-mean generalized normal distribution:*

$$p(E_{kt}) \propto e^{-\lambda |e_{kt}|^\beta}, \text{ with positive } \lambda \text{ and } \beta > 2 \text{ and } \beta \neq 3. \quad (18)$$

Suppose Eq. 9 holds true. Then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.

Proof. In this case, we have

$$\begin{aligned} \frac{\partial^3 \eta_{kt}}{\partial z_{kt}^2 \partial z_{k',t-1}} &= -\lambda \cdot \text{sgn}(e_{kt}) \cdot \alpha(\beta - 1)(\beta - 2) |e_{kt}|^{\beta-3} c_{kk'}, \\ \frac{\partial^2 \eta_{kt}}{\partial z_{kt} \partial z_{k',t-1}} &= -\lambda \beta (\beta - 1) |e_{kt}|^{\beta-2} c_{kk'}. \end{aligned} \quad (19)$$

We know that $|e_{lt}|^{\beta-2}$ and $|e_{lt}|^{\beta-3}$ are linearly independent (because their ratio, $|e_{lt}|$, is not constant). Furthermore, $|e_{lt}|^{\beta-2}$ and $|e_{lt}|^{\beta-3}$, with $l = 1, 2, \dots, n$, are $2n$ linearly independent functions (because of the different arguments involved).

Suppose there exist α_{l1} and α_{l2} , with $l = 1, 2, \dots, n$, such that

$$\sum_{l=1}^n (\alpha_{l1} \mathbf{v}_{lt} + \alpha_{l2} \hat{\mathbf{v}}_{lt}) = 0. \quad (20)$$

It is assumed that for each $k = 1, 2, \dots, n$, there exists at least one k' such that $c_{kk'} \neq 0$. Eq. 20 then implies that for any k we have

$$\alpha_{k1} c_{kk'} |e_{kt}|^{\beta-2} + \alpha_{k2} c_{kk'} |e_{kt}|^{\beta-3} + \sum_{l \neq k} (\alpha_{l1} c_{lk'} |e_{lt}|^{\beta-2} + \alpha_{l2} c_{lk'} |e_{lt}|^{\beta-3}) = 0. \quad (21)$$

Since $|e_{lt}|^{\beta-2}$ and $|e_{lt}|^{\beta-3}$, with $l = 1, 2, \dots, n$, are linearly independent and $c_{kk'} \neq 0$, to make the above equation hold, one has to set $\alpha_{k1} = \alpha_{k2} = 0$. As this applies to any k , we know that for Eq. 20 to be satisfied, α_{l1} and α_{l2} must be 0, for all $l = 1, 2, \dots, n$. That is, $\mathbf{v}_{1t}, \hat{\mathbf{v}}_{1t}, \mathbf{v}_{2t}, \hat{\mathbf{v}}_{2t}, \dots, \mathbf{v}_{nt}, \hat{\mathbf{v}}_{nt}$ are linearly independent. The linear independence condition in Theorem S1 is satisfied. Therefore \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$. \square

S1.2. Proof for Theorem 2 and 3

Let $\mathbf{v}_{kt}(u_r)$ be \mathbf{v}_{kt} , which is defined in Eq. 10, in the u_r context. Similarly, Let $\hat{\mathbf{v}}_{kt}(u_r)$ be $\hat{\mathbf{v}}_{kt}$ in the u_r context. Let

$$\begin{aligned} \mathbf{s}_{kt} &\triangleq \left(\mathbf{v}_{kt}(u_1)^\top, \dots, \mathbf{v}_{kt}(u_m)^\top, \frac{\partial^2 \eta_{kt}(u_2)}{\partial z_{kt}^2} - \frac{\partial^2 \eta_{kt}(u_1)}{\partial z_{kt}^2}, \dots, \frac{\partial^2 \eta_{kt}(u_m)}{\partial z_{kt}^2} - \frac{\partial^2 \eta_{kt}(u_{m-1})}{\partial z_{kt}^2} \right)^\top, \\ \hat{\mathbf{s}}_{kt} &\triangleq \left(\hat{\mathbf{v}}_{kt}(u_1)^\top, \dots, \hat{\mathbf{v}}_{kt}(u_m)^\top, \frac{\partial \eta_{kt}(u_2)}{\partial z_{kt}} - \frac{\partial \eta_{kt}(u_1)}{\partial z_{kt}}, \dots, \frac{\partial \eta_{kt}(u_m)}{\partial z_{kt}} - \frac{\partial \eta_{kt}(u_{m-1})}{\partial z_{kt}} \right)^\top. \end{aligned}$$

As provided below, in our case, the identifiability of \mathbf{z}_t is guaranteed by the linear independence of the whole function vectors \mathbf{s}_{kt} and $\hat{\mathbf{s}}_{kt}$, with $k = 1, 2, \dots, n$. However, the identifiability result in Yao et al. (2021) relies on the linear independence of only the last $m - 1$ components of \mathbf{s}_{kt} and $\hat{\mathbf{s}}_{kt}$ with $k = 1, 2, \dots, n$; this linear independence is generally a much stronger condition.

Theorem S2 (Identifiability under Changing Causal Dynamics). *Suppose the observed processes \mathbf{x}_t was generated by Eq. 8 and that the conditional distribution $p(z_{kt} | \mathbf{z}_{t-1})$ may change across m values of the context variable \mathbf{u} , denoted by u_1, u_2, \dots, u_m . Suppose the components of \mathbf{z}_t are mutually independent conditional on \mathbf{z}_{t-1} in each context. Assume that the components of $\hat{\mathbf{z}}_t$ produced by Eq. 9 are also mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. If the $2n$ function vectors \mathbf{s}_{kt} and $\hat{\mathbf{s}}_{kt}$, with $k = 1, 2, \dots, n$, are linearly independent, then $\hat{\mathbf{z}}_t$ is a permuted invertible component-wise transformation of \mathbf{z}_t .*

Proof. As in the proof of Theorem S1, because the components of $\hat{\mathbf{z}}_t$ are mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$, we know that for $i \neq j$,

$$\frac{\partial^2 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}; \mathbf{u})}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} = \sum_{k=1}^n \left(\frac{\partial^2 \eta_{kt}(\mathbf{u})}{\partial z_{kt}^2} \cdot \mathbf{H}_{kit} \mathbf{H}_{kjt} + \frac{\partial \eta_{kt}(\mathbf{u})}{\partial z_{kt}} \cdot \frac{\partial \mathbf{H}_{kit}}{\partial \hat{z}_{jt}} \right) - \frac{\partial^2 \log |\mathbf{H}_t|}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} \equiv 0. \quad (22)$$

Compared to Eq. 14, here we allow $p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1})$ to depend on \mathbf{u} . Since the above equations is always 0, taking its partial derivative w.r.t. $z_{l,t-1}$ gives

$$\frac{\partial^3 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}; \mathbf{u})}{\partial \hat{z}_{it} \partial \hat{z}_{jt} \partial z_{l,t-1}} = \sum_{k=1}^n \left(\frac{\partial^3 \eta_{kt}(\mathbf{u})}{\partial z_{kt}^2 \partial z_{l,t-1}} \cdot \mathbf{H}_{kit} \mathbf{H}_{kjt} + \frac{\partial^2 \eta_{kt}(\mathbf{u})}{\partial z_{kt} \partial z_{l,t-1}} \cdot \frac{\partial \mathbf{H}_{kit}}{\partial \hat{z}_{jt}} \right) \equiv 0. \quad (23)$$

Similarly, Using different values for \mathbf{u} in Eq. 22 take the difference of this equation across them gives

$$\begin{aligned} &\frac{\partial^2 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}; u_{r+1})}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} - \frac{\partial^2 \log p(\hat{\mathbf{z}}_t | \hat{\mathbf{z}}_{t-1}; u_r)}{\partial \hat{z}_{it} \partial \hat{z}_{jt}} \\ &= \sum_{k=1}^n \left[\left(\frac{\partial^2 \eta_{kt}(u_{r+1})}{\partial z_{kt}^2} - \frac{\partial^2 \eta_{kt}(u_r)}{\partial z_{kt}^2} \right) \cdot \mathbf{H}_{kit} \mathbf{H}_{kjt} + \left(\frac{\partial \eta_{kt}(u_{r+1})}{\partial z_{kt}} - \frac{\partial \eta_{kt}(u_r)}{\partial z_{kt}} \right) \cdot \frac{\partial \mathbf{H}_{kit}}{\partial \hat{z}_{jt}} \right] \equiv 0. \end{aligned} \quad (24)$$

Therefore, if \mathbf{s}_{kt} and $\mathring{\mathbf{s}}_{kt}$, for $k = 1, 2, \dots, n$, are linearly independent, $\mathbf{H}_{kit}\mathbf{H}_{kjt}$ has to be zero for all k and $i \neq j$. Then as shown in the proof of Theorem S1, $\hat{\mathbf{z}}_t$ must be a permuted component-wise invertible transformation of \mathbf{z}_t . \square

Theorem S3 (Identifiability under Observation Changes). *Suppose $\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$ and that the conditional distribution $p(z_{k,t} | \mathbf{u})$ may change across m values of the context variable \mathbf{u} , denoted by u_1, u_2, \dots, u_m . Suppose the components of \mathbf{z}_t are mutually independent conditional on \mathbf{u} in each context. Assume that the components of $\hat{\mathbf{z}}_t$ produced by Eq. 3 are also mutually independent conditional on $\hat{\mathbf{z}}_{t-1}$. If the $2n$ function vectors $\mathbf{s}_{k,t}$ and $\mathring{\mathbf{s}}_{k,t}$, with $k = 1, 2, \dots, n$, are linearly independent, then $\hat{\mathbf{z}}_t$ is a permuted invertible component-wise transformation of \mathbf{z}_t .*

Proof. As in the proof of Theorem S2, because \mathbf{z}_t is not dependent on the history \mathbf{z}_{t-1} so are the components of $\hat{\mathbf{z}}_t$, the conditioning on $\hat{\mathbf{z}}_t$ in Eq. 22 and the following equations can be removed because of the independence. This directly leads to the same conclusion as in Theorem S2. \square

Corollary S3 (Identifiability under Modular Distribution Shifts). *Assume the data generating process in Eq. 1. If the three partitioned latent components $\mathbf{z}_t = (\mathbf{z}_t^{\text{fix}}, \mathbf{z}_t^{\text{chg}}, \mathbf{z}_t^{\text{obs}})$ respectively satisfy the conditions in Theorem 1, Theorem 2, and Theorem 3, then \mathbf{z}_t must be an invertible, component-wise transformation of a permuted version of $\hat{\mathbf{z}}_t$.*

Proof. Because the three partitioned subspaces $(\mathbf{z}_t^{\text{fix}}, \mathbf{z}_t^{\text{chg}}, \mathbf{z}_t^{\text{obs}})$ are conditional independent given the history and domain index, it is straightforward to factorize the joint conditional log density into three components. By using the proof in Theorem 1, 2, and 3, we can directly derive the same quantity as in Eq. 15 or Eq. 24. Therefore, if \mathbf{s}_{kt} and $\mathring{\mathbf{s}}_{kt}$, for $k = 1, 2, \dots, n$, are linearly independent, $\mathbf{H}_{kit}\mathbf{H}_{kjt}$ has to be zero for all k and $i \neq j$. Then as shown in the proof of Theorem S1, $\hat{\mathbf{z}}_t$ must be a permuted component-wise invertible transformation of \mathbf{z}_t . \square

S1.3. Comparisons with Existing Identifiability Theories

The closest work to ours includes (1) LEAP (Yao et al., 2021), which leverage the changes in noise distribution to disentangle the nonparametric causal processes, (2) PCL (Hyvarinen & Morioka, 2017), which exploited temporal constraints to separate independent sources, and (3) SlowVAE (Klindt et al., 2020), which leveraged sparse transition of adjacent video frames to separate independent sources..

LEAP The sources z_{it} in LEAP can have nonparametric, time-delayed causal relations in between. However, LEAP only considers a special case of nonstationarity caused by changes in noise distributions across domains, while our work can allow the causal influencing strength to change across segments. Furthermore, because LEAP assumes all sources' conditional distributions are changed by nonstationary noise, it doesn't exploit the fixed dynamics parts for identifiability. On the other hand, since our work exploits both the distribution changes from the fixed causal dynamics and changing dynamics (also from observation changes), our derived identifiability conditions are generally weaker than (Yao et al., 2021).

$$\underbrace{\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)}_{\text{Nonlinear mixing}}, \quad \underbrace{z_{it} = f_i(\{z_{j,t-\tau} | z_{j,t-\tau} \in \mathbf{Pa}(z_{it})\}, \epsilon_{it})}_{\text{Nonparametric transition}} \quad \text{with} \quad \underbrace{\epsilon_{it} \sim p_{\epsilon_i | \mathbf{u}}}_{\text{Nonstationary noise}}. \quad (25)$$

PCL The sources z_{it} in PCL were assumed to be mutually independent (see Assumption 1 of Theorem 1 in PCL). In contrast, our identifiability conditions under fixed causal dynamics allow the sources to have time-delayed causal relations in between, which is much more realistic in real-world applications. The underlying processes of PCL are described by Eq. 26:

$$\log p(z_{i,t} | z_{i,t-1}) = G(z_{i,t} - \rho z_{i,t-1}) \quad \text{or} \quad \log p(z_{i,t} | z_{i,t-1}) = -\lambda (z_{i,t} - r(z_{i,t-1}))^2 + \text{const}. \quad (26)$$

where G is some non-quadratic function corresponding to the log-pdf of innovations, $\rho < 1$ is regression coefficient, r is some nonlinear, strictly monotonic regression, and λ is a positive precision parameter.

SlowVAE Inspired by slow feature analysis, SlowVAE assumes the underlying sources to have identity transitions with generalized Laplacian innovations described in Eq. 27:

$$p(\mathbf{z}_t | \mathbf{z}_{t-1}) = \prod_{i=1}^d \frac{\alpha \lambda}{2\Gamma(1/\alpha)} \exp(-(\lambda |z_{i,t} - z_{i,t-1}|^\alpha) \quad \text{with } \alpha < 2. \quad (27)$$

Our Corollary 2 completes to the Laplacian innovation model above by allowing time-delayed vector autoregressive transitions in the latent process with multiple time lags, and with generalized Gaussian noises. Consequently, temporally causally-related latent processes with linear transition dynamics can thus be modeled and recovered from their nonlinear mixtures with our condition.

S2. Experiment Settings

S2.1. Datasets

S2.1.1. SYNTHETIC DATASET GENERATION

To evaluate the identifiability of our method under different conditions, we generate the synthetic data with 1) fixed causal dynamics; 2) changing causal dynamics and 3) the modular distribution shift.

Fixed Causal Dynamics For the fixed causal dynamics. We generate 100,000 data points according to Eq. (6), where the latent size is $n = 8$, lag number of the process is $L = 2$. We apply a 2-layer MLP with LeakyReLU as the state transition function. The process noise are sampled from i.i.d. Gaussian distribution ($\sigma = 0.1$). The process noise terms are coupled with the history information through multiplication with the average value of all the time-lagged latent variables.

Changing Causal Dynamics We use a Gaussian additive noise model with changes in the influencing strength as the latent processes. To add changes, we vary the values of the first layer of the MLP across the 20 segments and generate 7,500 samples for each segment. The entries of the kernel matrix of the first layer are uniformly distributed between $[-1, 1]$ in each domain.

Modular Distribution Shifts The latent space of this dataset is partitioned into 6 fixed dynamics components under the heterogeneous noise model, 2 changing components with changing causal dynamics and 1 component modulated by domain index only. The fixed and changing dynamics components follow the same generating procedures above. The global change component is sampled from i.i.d Gaussian distribution whose mean and variance are modulated by domain index. In particular, distribution mean terms are uniformly sampled between $[-1, 1]$ and variance terms are uniformly sampled between $[0.01, 1]$.

S2.1.2. REAL-WORLD DATASET

Modified Cartpole The Cartpole problem (Huang et al., 2021) “consists of a cart and a vertical pendulum attached to the cart using a passive pivot joint. The cart can move left or right. The task is to prevent the vertical pendulum from falling by putting a force on the cart to move it left or right. The action space consists of two actions: moving left or right.”

The original dataset (Huang et al., 2021) introduces “two change factors respectively for the state transition dynamics θ_k^{dyn} : varying gravity and varying mass of the cart, and a change factor in the observation function θ_k^{obs} that is the image noise level. Fig. S1 gives a visual example of Cartpole game, and the image with Gaussian noise. The images of the varying gravity and mass look exactly like the original image. Specifically, in the gravity case, we consider source domains with gravity $g = \{5, 10, 20, 30, 40\}$. We take into account both interpolation (where the gravity in the target domain is in the support of that in source domains) with $g = \{15\}$, and extrapolation (where it is out of the support w.r.t. the source domains) with $g = \{55\}$. Similarly, we consider source domains where the mass of the cart is $m = \{0.5, 1.5, 2.5, 3.5, 4.5\}$, while in target domains it is $m = \{1.0, 5.5\}$. In terms of changes on the observation function, we add Gaussian noise on the images with variance $\sigma = \{0.25, 0.75, 1.25, 1.75, 2.25\}$ in source domains, and $\sigma = \{0.5, 2.75\}$ in target domains. The detailed settings in both source and target domains are in Table S1.”




Figure S1. Visual examples of Cartpole game and change factors. (a) Cartpole game; (b) Modified Cartpole game with Gaussian noise on the image. The light blue arrows are added to show the direction in which the agent can move. Figure source: (Huang et al., 2021).

	Gravity	Mass	Noise
Source domains	{5, 10, 20, 30, 40}	{0.5, 1.5, 2.5, 3.5, 4.5}	{0.25, 0.75, 1.25, 1.75, 2.25}
Interpolation set	{15}	{1.0}	{0.5}
Extrapolation set	{55}	{5.5}	{2.75}

Table S1. The settings of source and target domains for modified Cartpole experiments (Huang et al., 2021).

S2.2. Mean Correlation Coefficient

MCC is a standard metric for evaluating the recovery of latent factors in ICA literature. MCC first calculates the absolute values of the correlation coefficient between every ground-truth factor against every estimated latent variable. Depending on whether componentwise invertible nonlinearities exist in the recovered factors, Pearson correlation coefficients or Spearman’s rank correlation coefficients can be used. The possible permutation is adjusted by solving a linear sum assignment problem in polynomial time on the computed correlation matrix.

S3. Implementation Details

S3.1. Network Architecture

We summarize our network architecture below and describe it in detail in Table S2 and Table S3.

S3.2. Hyperparameter and Training

Hyperparameter Selection The hyperparameters of LiLY include $[\beta, \gamma]$, which are the weights of KLD terms for the time-lagged variables and the latent variables at the current time step, as well as the latent size n and maximum time lag L . We use the ELBO loss on the validation dataset to select the best pair of $[\beta, \gamma, \sigma]$ because low ELBO loss always leads to high MCC. We always set a larger latent size than the true latent size. This is critical in video datasets because the image pixels contain more information than the annotated latent causal variables, and restricting the latent size will hurt the reconstruction performances. For the maximum time lag L , we set it by the rule of thumb. For instance, we use $L = 2$ for temporal datasets with a latent physics process.

Training Details The models were implemented in PyTorch 1.8.1. The VAE network is trained using AdamW optimizer for a maximum of 50 epochs and early stops if the validation ELBO loss does not decrease for five epochs. A learning rate of 0.002 and a mini-batch size of 64 are used. We have used three random seeds in each experiment and reported the mean performance with standard deviation averaged across random seeds. We have used several standard tricks to improve training stability: (1) we use a slightly larger latent size than the true latent size for real-world datasets in order to make sure the meaningful latent variables are among the recovered latent variables, and (2) we use AdamW optimizer as a regularizer

Table S2. Architecture details. BS: batch size, T: length of time series, i.dim: input dimension, z.dim: latent dimension, LeakyReLU: Leaky Rectified Linear Unit.

Configuration	Description	Output
1. MLP-Encoder Encoder for Synthetic Data		
Input: $\mathbf{x}_{1:T}$	Observed time series	$\text{BS} \times \text{T} \times \text{i.dim}$
Dense	128 neurons, LeakyReLU	$\text{BS} \times \text{T} \times 128$
Dense	128 neurons, LeakyReLU	$\text{BS} \times \text{T} \times 128$
Dense	128 neurons, LeakyReLU	$\text{BS} \times \text{T} \times 128$
Dense	Temporal embeddings	$\text{BS} \times \text{T} \times \text{z.dim}$
2. MLP-Decoder Decoder for Synthetic Data		
Input: $\hat{\mathbf{z}}_{1:T}$	Sampled latent variables	$\text{BS} \times \text{T} \times \text{z.dim}$
Dense	128 neurons, LeakyReLU	$\text{BS} \times \text{T} \times 128$
Dense	128 neurons, LeakyReLU	$\text{BS} \times \text{T} \times 128$
Dense	i.dim neurons, reconstructed $\hat{\mathbf{x}}_{1:T}$	$\text{BS} \times \text{T} \times \text{i.dim}$
5. Factorized Inference Network Bidirectional Inference Network		
Input	Sequential embeddings	$\text{BS} \times \text{T} \times \text{z.dim}$
Bottleneck	Compute mean and variance of posterior	$\mu_{1:T}, \sigma_{1:T}$
Reparameterization	Sequential sampling	$\hat{\mathbf{z}}_{1:T}$
6. Modular Prior Nonlinear Transition Prior Network		
Input	Sampled latent variable sequence $\hat{\mathbf{z}}_{1:T}$	$\text{BS} \times \text{T} \times \text{z.dim}$
InverseTransition	Compute estimated residuals $\hat{\epsilon}_{it}$	$\text{BS} \times \text{T} \times \text{z.dim}$
JacobianCompute	Compute $\log(\det(\mathbf{J}))$	BS

to prevent training from being interrupted by overflow or underflow of variance terms of VAE.

Computing Hardware We used a machine with the following CPU specifications: Intel(R) Core(TM) i7-7700K CPU @ 4.20GHz; 8 CPUs, four physical cores per CPU, a total of 32 logical CPU units. The machine has two GeForce GTX 1080 Ti GPUs with 11GB GPU memory.

Reproducibility We’ve included the code for our framework and all experiments in the supplementary materials. We plan to release our code under the MIT License after the ICML paper review period.

Table S3. Architecture details on CNN encoder and decoder. BS: batch size, T: length of time series, h_dim: hidden dimension, z_dim: latent dimension, F: number of filters, (Leaky)ReLU: (Leaky) Rectified Linear Unit.

Configuration	Description	Output
3.1.1 CNN-Encoder		
Input: $\mathbf{x}_{1:T}$	RGB video frames	$BS \times T \times 3 \times 64 \times 64$
Conv2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 64 \times 64$
Conv2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 32 \times 32$
Conv2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 16 \times 16$
Conv2D	F: 64, BatchNorm2D, LeakyReLU	$BS \times T \times 64 \times 8 \times 8$
Conv2D	F: 64, BatchNorm2D, LeakyReLU	$BS \times T \times 64 \times 4 \times 4$
Conv2D	F: 128, BatchNorm2D, LeakyReLU	$BS \times T \times 128 \times 1 \times 1$
Dense	F: $2 * z_dim$ = dimension of hidden embedding	$BS \times T \times 2 * z_dim$
4.1 CNN-Decoder		
Input: $\mathbf{z}_{1:T}$	Sampled latent variable sequence	$BS \times T \times z_dim$
Dense	F: 128, LeakyReLU	$BS \times T \times 128 \times 1 \times 1$
ConvTranspose2D	F: 64, BatchNorm2D, LeakyReLU	$BS \times T \times 64 \times 4 \times 4$
ConvTranspose2D	F: 64, BatchNorm2D, LeakyReLU	$BS \times T \times 64 \times 8 \times 8$
ConvTranspose2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 16 \times 16$
ConvTranspose2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 32 \times 32$
ConvTranspose2D	F: 32, BatchNorm2D, LeakyReLU	$BS \times T \times 32 \times 64 \times 64$
ConvTranspose2D	F: 3, estimated scene $\hat{\mathbf{x}}_{1:T}$	$BS \times T \times 3 \times 64 \times 64$