

SPACETIME: Causal Discovery from Non-Stationary Time Series

Sarah Mameche¹, Lénaïg Cornanguer¹, Urmi Ninad^{23*}, Jilles Vreeken^{1*}

¹CISPA Helmholtz Center for Information Security

²Technische Universität Berlin, Germany

³German Aerospace Center, Institute of Data Science

sarah.mameche@cispa.de, lenaig.cornanguer@cispa.de, urmi.ninad@tu-berlin.de, jv@cispa.de

Abstract

Understanding causality is challenging and often complicated by changing causal relationships over time and across environments. Climate patterns, for example, shift over time with recurring seasonal trends, while also depending on geographical characteristics such as ecosystem variability. Existing methods for discovering causal graphs from time series either assume stationarity, do not permit both temporal and spatial distribution changes, or are unaware of locations with the same causal relationships. In this work, we therefore unify the three tasks of causal graph discovery in the non-stationary multi-context setting, of reconstructing temporal regimes, and of partitioning datasets and time intervals into those where invariant causal relationships hold. To construct a consistent score that forms the basis of our method, we employ the Minimum Description Length principle. Our resulting algorithm SPACETIME simultaneously accounts for heterogeneity across space and non-stationarity over time. Given multiple time series, it discovers regime changepoints and a temporal causal graph using non-parametric functional modeling and kernelized discrepancy testing. We also show that our method provides insights into real-world phenomena such as river-runoff measured at different catchments and biosphere-atmosphere interactions across ecosystems.

Introduction

Gaining insight into the dynamics of complex real-world processes is closely tied to understanding the causal mechanisms underlying their function. Causal models (Pearl 2009) allow reasoning about the effect of intervention and distributional change, which makes them especially useful for modeling systems under changing environments or over time. At odds with this, however, is that established methods for discovering causal networks from data (Spirtes et al. 2000; Chickering 2002) often make the simplifying assumption that all samples have a fixed data-generating process.

This assumption is especially problematic in time series data. In climate science, for example, measurements often come from geographical regions experiencing different climatic conditions. Even within a single region, weather patterns are not constant but often shift over time due to seasonality, extreme events, or global climate change. Similarly,

*Equal supervision

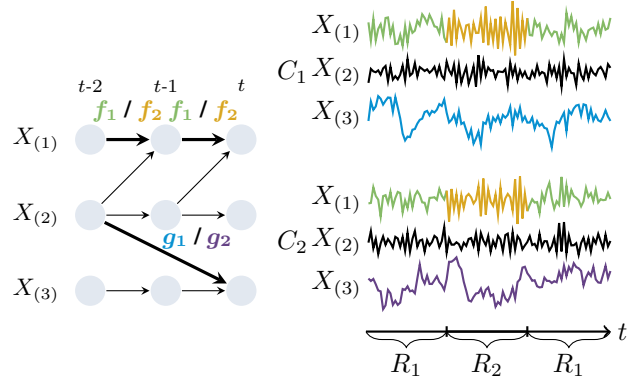


Figure 1: *Left*: A temporal causal graph representing the data-generating causal mechanism of three variables ($X_{(1)}$, $X_{(2)}$, and $X_{(3)}$). The bold edges indicate a local mechanism change across contexts (g_1 or g_2) or over time (f_1 or f_2). *Right*: Variable measurements from different contexts (C_1 and C_2) and under different temporal regimes (R_1 and R_2).

disease trajectories of patients at multiple locations could show variability due to varying healthcare infrastructures or population heterogeneity, and change over time in response to local interventions or changes in public behavior.

To address such scenarios, in this work we consider a collection of non-stationary time series datasets. Fig. 1, right shows this setting for three variables that we observe over time in two contexts. Changes in their data-generating causal mechanisms (colored) occur at unknown points in time ($X_{(1)}$), called *changepoints*, or between contexts ($X_{(3)}$). We aim to jointly discover changes across temporal and spatial scales while explaining the underlying causal interactions, as shown in Fig. 1, left.

Recent methods focus either on changepoint discovery in single time series (Saggioro et al. 2020) or on multiple datasets without detecting changepoints (Günther, Ninad, and Runge 2023). Here, we suggest considering both tasks simultaneously and looking for recurring, invariant causal mechanisms. Examples include cities in geographical regions with similar environmental conditions, or time periods with seasonal trends or economic cycles. Groups of datasets across which the causal mechanisms do not change are said

to belong to the same *context*, and similarly, time intervals when the causal mechanisms of all variables remain stationary are said to belong to the same *regime*. For example, in Fig. 1, $X_{(1)}$ experiences a change at two points in time, but we observe a repeating regime R_1 .

Unlike the mostly constraint-based causal discovery literature for non-stationary time series (Runge 2020; Asaad, Devijver, and Gaussier 2022), we rely on functional modeling to provide direct insights into causal mechanisms and their similarity. We model causal relationships non-parametrically using Gaussian processes (GPs) (Rasmussen and Williams 2005). A constraint-based approach to causal discovery is also limited by the effectiveness of conditional independence testing, which is a provably hard problem (Shah and Peters 2020). This motivates us to adopt instead a score-based approach based on sparsity constraints, where we build on the Minimum Description Length principle (MDL) (Rissanen 1978). We will show how this applies to the time series causal discovery setting and changepoint detection tasks in the remainder of this work.

Contributions We propose a framework for causal modeling of non-stationary time series in heterogeneous contexts. We confirm theoretically that our causal discovery scoring criterion is consistent and develop the SPACETIME algorithm to efficiently discover causal models, regime changepoints, and invariances over time (regimes) and space (contexts). Besides showing that our approach works well in different synthetic data-generation settings, we study two real-world applications exploring drivers of river discharge (Günther et al. 2023) and functional biosphere-atmosphere interactions (Baldocchi 2014). Our approach discovers joint causal networks across the span of multiple years and various locations, and reveals gradients in interaction strength across temporal and spatial scales that match previous findings (Kraemer et al. 2020; Krich et al. 2021).

Related Work

Causal discovery in time series data has long been limited to Granger-causality (Granger 1969), which focuses on predicting whether past values of one variable can help forecast future values of another. We focus here on related work that aims to identify true causal relationships.

Most methods, such as PCMCI+ (Runge 2020), VARLINGAM (Hyvärinen et al. 2010) as well as DYNOTEARS (Pamfil et al. 2020), assume that the data are stationary in time and space. The problem of non-stationarity in time has been studied from different perspectives. R-PCMCI (Saggiaro et al. 2020), like us, assumes the existence of recurrent regimes within each of which the data is stationary. CD-NOD (Huang et al. 2020), on the other hand, can capture temporal non-stationarity if it can be modeled as a smooth function of the time index. PCMCI $_{\Omega}$ (Gao et al. 2024) only considers semi-stationary time series where the causal effects occur periodically. Finally, unlike the aforementioned approaches, J-PCMCI+ (Günther, Ninad, and Runge 2023) can handle multiple heterogeneous time series datasets as well as non-stationarity. However, it does not perform regime detection or context partitioning.

Theoretical Framework

In the following, we introduce the concepts of contexts and regimes and our causal modeling assumptions.

Problem Setting

Throughout our work, we consider multivariate discrete-time stochastic processes over a set of continuous variables $\mathbf{X} = \{X_{(1)}, \dots, X_{(m)}\}$. We assume that we sample each variable at a discrete set of time indices $\mathcal{T} = \{t_1, \dots, t_n\}$. This results in measurements $X_{(i)(t)}$ at each time point t , which we denote as X_t when the variable index can be suppressed. We denote the resulting time series as $(\mathbf{X}_t)_{t \in \mathcal{T}}$, shorthand $\mathbf{X}_{\mathcal{T}}$.

As in our motivating setting, we observe not just one but a collection of time series datasets $((\mathbf{X}_t^d)_{t \in \mathcal{T}})_{d \in \mathcal{D}}$, or $\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}$ for short, for a given set of dataset indices \mathcal{D} . To capture that certain groups of locations share the same data-generating mechanisms while for others they differ, we propose partitioning the index set \mathcal{D} into different contexts.

Definition 1 (Contexts) We define the contexts for a variable X as partitions $\mathcal{C} = \{C_1, \dots, C_K\}$ of \mathcal{D} into disjoint, nonempty index sets $C_k \neq \emptyset$, with $\cup_k C_k = \mathcal{D}$ and $C_k \cap C_{k'} = \emptyset$ for any $k \neq k'$, such that no mechanism change of X occurs between datasets d_1, d_2 in the same C_k .

Analogously, to represent changes over time, we consider a subset $\mathcal{L} \subset \mathcal{T}$ of points that index such changes, so-called *changepoints*. As similar trends might repeat periodically over time, there are not necessarily as many distinct causal mechanisms as changepoints. This motivates partitioning the regions defined by \mathcal{L} into groups, so-called regimes.

Definition 2 (Regimes) We define the regimes for a variable X as set partitions $\mathcal{R} = \{R_1, \dots, R_L\}$ of \mathcal{T} into disjoint, nonempty index sets under a given set of changepoints \mathcal{L} , such that $\mathcal{L} = \{t \in \mathcal{T} \mid \text{regime}(t) \neq \text{regime}(t+1)\}$, where no mechanism change of X occurs in the same R_l .

Above, we associate contexts with datasets and regimes with time indices through the functions *context* : $\mathcal{D} \rightarrow \mathcal{C}$ and *regime* : $\mathcal{T} \rightarrow \mathcal{R}$, which map each dataset (resp. time index) to its corresponding context (resp. regime). Note that we have $l \in \mathcal{L}$ whenever the causal mechanism for *at least* one node X changes. By assuming a fixed set \mathcal{L} over \mathcal{D} , we presume all datasets are identically affected by non-stationarity, for example, due to seasonality.

Together, the contexts and regimes determine in which subsamples of observations no mechanism change occurs. We illustrate these notions in Fig. 1, highlighting each subsample with a different color. We write a subsample in C_k and R_r as $s = (k, r)$ for brevity.

In summary, our input data is as follows.

Definition 3 (Non-stationary Time Series) We consider multivariate time series $\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}$ with known dataset indices $d \in \mathcal{D}$ and unknown changepoints $\mathcal{L} \subset \mathcal{T}$, as well as latent contexts \mathcal{C} over \mathcal{D} and regimes \mathcal{R} over \mathcal{T} for each X .

In the following, we will define more formally what it means that variables undergo causal mechanism changes by introducing our causal model.

Causal Model

As our graphical representation of causal interactions over time, we work with the following temporal causal graphs.

Definition 4 (Temporal Causal Graph (Assaad, Devijver, and Gaussier 2022)) A temporal causal graph (TCG) $\mathcal{G} = (V, E)$ over \mathbf{X} and \mathcal{T} is a directed acyclic graph (DAG) \mathcal{G} , where V includes nodes $X_{(i)(t)}$ for each i and t and E includes two types of edges,

- instantaneous links $X_{(i)(t)} \rightarrow X_{(j)(t)}$ for $i \neq j$ whenever X_i causes X_j at time point $t \in \mathcal{T}$, and
- lag-specific directed links $X_{(i)(t-\tau)} \rightarrow X_{(j)(t)}$ pointing forward in time whenever $X_{(i)(t-\tau)}$ causes $X_{(j)(t)}$ with a time lag of $\tau > 0$.

Note that lag-specific links may include self-transitions so long as $\tau > 0$.

We make the standard assumptions of Causal Markov Condition, Faithfulness, and Sufficiency (Pearl 2009).

Assumption 1 (Causal Markov Condition, Faithfulness, and Sufficiency) We assume the Causal Markov Condition and Faithfulness stating that d -separations in \mathcal{G} correspond to conditional independencies in the data distribution P . We further assume that no unobserved confounding or selection variables exist.

We also adopt a common assumption that the causal edge directions in the TCG persist over time.

Assumption 2 (Repeating Edge Property (Gerhardus 2024)) We associate the time series $\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}$ with a TCG \mathcal{G} , with all edges remaining constant in direction across datasets \mathcal{D} and through time \mathcal{T} , in the sense that if $X_{(i)(t)} \rightarrow X_{(j)(t')}$ then $X_{(i)(t+\tau)} \rightarrow X_{(j)(t'+\tau)}$ is in \mathcal{G} for any $\tau \geq 0$.

Under this property of edge consistency over time, we can condense the information in the TCG into a single time window of the length of the maximum time lag τ .

Definition 5 (Window Causal Graph (Assaad, Devijver, and Gaussier 2022)) For the maximum time lag τ in a TCG \mathcal{G} , the window causal graph (WCG) $\mathcal{G}_{\tau} = (V, E)$ over vertices $X_{(i)(t)}$ at each time $t - \tau, \dots, t$ contains lag-specific directed links $X_{(i)(t-\tau)} \rightarrow X_{(j)(t)}$ pointing forward in time whenever $X_{(i)(t-\tau)}$ causes $X_{(j)(t)}$ in \mathcal{G} with time lag τ , where $0 \leq \tau < \tau$ if $j \neq i$ and $0 < \tau < \tau$ if $j = i$.

Under edge consistency, the causal mechanism changes preserve the causal parent sets. Their differences lie, therefore, in the structural mechanisms themselves as follows.

Assumption 3 (SCM with Contexts and Regimes) Given a time series $\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}$, we assume a Structural Causal Model (SCM) $\mathcal{M} = (\mathcal{G}, \mathcal{L}, \mathcal{C}, \mathcal{R})$ with TCG \mathcal{G} , changepoints \mathcal{L} , and contexts and regimes \mathcal{C}, \mathcal{R} for each X . Each variable X is generated through a set of structural equations of the form

$$X_t^d = f^{(k,r)}(pa(X_t^d), N_t^d) \quad (1)$$

where X_t^d a sample at time t in set d with context $(d) = k$, regime $(t) = r$, $f^{(k,r)}$ is its causal mechanism in C_k and R_r , $pa(X_t^d)$ the set of its contemporaneous and lagged parents in \mathcal{G} and $N_r^k \perp\!\!\!\perp X_r^k$ an independent noise variable.

That is, all datasets and time points in a subsample (k, r) have identical generating process, where each variable has its mechanism $f^{(k,r)}$ and noise distribution N_t^k . The above keeps variable indices implicit but can be similarly stated for all mechanisms $f_{(i)}$ of $X_{(i)}$ in the system. Note that our model can therefore show *which* variables induce temporal or spatial changes at *which* locations.

Discovering this model from data will be our next focus.

Causal Discovery with Contexts and Regimes

Here, we describe our kernelized approach combined with minimum description lengths to discover the causal model underlying non-stationary time series.

MDL for Causal Discovery

We base our approach on the Algorithmic Model of Causation (AMC) (Janzing and Schölkopf 2010), which postulates that the factorization of the joint distribution with the lowest Kolmogorov complexity (Li and Vitányi 2009) is the true causal model. In this paradigm, we reason about causal mechanisms as programs that compute effects from their parents. Kolmogorov complexity is not computable, but, can be approximated from above in a statistically sound manner via the Minimum Description Length (MDL) (Rissanen 1978; Grünwald 2007) principle. Marx and Vreeken (2021) formally connect two-part MDL and the AMC.

MDL systematically balances model complexity and goodness of fit. The description length L of the data \mathbf{X} together with the model, here its causal graph \mathcal{G} , is given by

$$L(\mathbf{X}; \mathcal{G}) = \sum_{i \in \mathcal{G}} L(X_{(i)} \mid pa(X_{(i)}); \mathcal{H})$$

under a fixed model class \mathcal{H} . Each summand measures the length L , in bits, of both encoding the functional model class \mathcal{H} itself as well as the data for $X_{(i)}$ given $pa(X_{(i)})$ under its optimal model. Our model class of choice are hereby Gaussian processes (GPs), also used by Mameche, Kaltenpoth, and Vreeken (2023), as they are both non-parametric and have refined MDL scores (Grünwald 2007), that is, there is a principled way of defining the model cost.

A GP models a distribution over functions $f \sim \text{GP}(m(x), \kappa(x, x'))$ where $m(x) = \mathbb{E}[f(x)]$ is a mean function and $\kappa(x, x') = \mathbb{E}[(f(x) - m(x))(f(x') - m(x')))]$ a covariance kernel (Rasmussen and Williams 2005).

The refined MDL description length (Grünwald 2007) of the data for a variable $X_{(i)}$ given its causal parents $pa(X_{(i)})$ under its optimal GP model is given by

$$\begin{aligned} L(X_{(i)} \mid pa(X_{(i)}); \mathcal{H}_{\kappa}) \\ = \min_{f \in \mathcal{H}_{\kappa}} \left(-\log P(X_{(i)} \mid pa(X_{(i)})) + \|f\|_{\kappa}^2 \right) \\ + \frac{1}{2} \log \det(\sigma^{-2} K_S + I), \end{aligned}$$

where K is the Gramian for κ over $pa(X_{(i)})$.

The score assesses the fit of the GP via the negative log-likelihood component, as well as its complexity via the squared norm $\|f\|_{\kappa}^2 = \alpha^{\top} K \alpha$ of the space \mathcal{H} which can be seen as a measure of the smoothness of a given function, where $\alpha = (K + \sigma^2 I)^{-1} y$, and σ^2 is a scaling coefficient.

MDL for Non-stationary Time Series

To encode our causal model, we consider the causal mechanisms per variable for each subsample under given contexts and regimes. We approximate each using a GP with additive Gaussian noise to obtain their description lengths.

Given time series datasets, the MDL principle suggests the true causal model is the one minimizing this score,

$$\begin{aligned} & \arg \min_{\mathcal{M}} L(\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}; \mathcal{M}) \\ = & \arg \min_{\mathcal{G}_{\tau}, \mathcal{L}, (\mathcal{C}, \mathcal{R})} \sum_{X \in \mathcal{G}_{\tau}} \sum_{(r \in \mathcal{R}, k \in \mathcal{C})} L(X_r^k | pa(X_r^k); \mathcal{H}_{\kappa}). \quad (2) \end{aligned}$$

This raises the question of whether this score identifies the correct model components, including the WCG \mathcal{G}_{τ} , the regime changepoint indices \mathcal{L} , and partitions \mathcal{C} and \mathcal{R} . To this end, we impose a notion of persistence of regimes similar to the literature (Saggioro et al. 2020).

Assumption 4 (Sufficient Capacity and Persistence) For each variable X and subsample $s = (k, r) \in \mathcal{C} \times \mathcal{R}$, we have a large enough sample size

$$|s| = |\{X_t^d \mid \text{context}(d) = k, \text{regime}(t) = r\}|$$

so that there exists a GP $f^{(k,r)} \in \mathcal{H}_{\kappa}$ such that Eq. (1) holds.

For each interval $w = [t_{\min}, t_{\max}] \subset \mathcal{T}$ subject to

$$\begin{cases} \text{regime}(t) = r & \forall t \in w \\ \text{regime}(t_{\min} - 1) \neq r \text{ and } \text{regime}(t_{\max} + 1) \neq r \end{cases}$$

we have $|w| \geq d_{\min}$, ensuring persistent regimes.

It is important to note that time series are subject to auto-correlation and therefore not i.i.d., therefore we assume above that we access a large enough sample for each causal mechanism such that a GP can capture it. We also assume that each process persists for a minimal duration d_{\min} through time. Furthermore, we assume the following.

Assumption 5 (Shift Faithfulness) For distinct pairs $i \neq j, k \neq k'$ and $r \neq r'$ we have $f_{(i)}^{(k,r)} \neq f_{(j)}^{(k',r')}$.

Last, to allow us to distinguish causal directions we assume the principle of independent changes of causal mechanisms (Schölkopf et al. 2021; Perry, Kügelgen, and Schölkopf 2022; Mameche, Vreeken, and Kaltenpoth 2024). We state it informally as follows and also express it through a latent-variable formulation of our causal model in the Appendix.

Assumption 6 (Independent Changes (Schölkopf et al. 2021)) The causal mechanisms $f_{(i)}, f_{(j)}$ change independently of each other for each pair of variables $X_{(i)}, X_{(j)}$.

Under these assumptions, we can show that our score is consistent, with the proof provided in the Appendix.

Theorem 1 (Consistency) Let Assumptions 1-6 hold. Then Eq. (2) is minimised for the true causal model \mathcal{M}^* with WCG \mathcal{G}_{τ}^* , changepoints \mathcal{L}^* , and partitions \mathcal{C}^* and \mathcal{R}^* in the limit of \mathcal{D} and \mathcal{T} .

Given the large search space over models in Eq. (2) and the inherent complexity in identifying the causal relationships and mechanisms while the regime changepoints and context and regimes partition are unknown, we develop a practical algorithm in the next section.

Algorithm 1: SPACETIME

Require: Multivariate time series datasets $\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}$, minimum distance between changepoints d_{\min} , maximum time lag τ

Ensure: A window causal graph \mathcal{G}_{τ} , a set of regime changepoints \mathcal{L} , regime and context partitions \mathcal{R}, \mathcal{C}

- 1: $\mathcal{G}_{\tau} = \text{empty graph}$
 - 2: **while** not converged **do**
 - 3: $\mathcal{L} = \text{CHANGEPOINTDETECTION}(\mathcal{G}_{\tau}, \mathbf{X}_{\mathcal{T}}^{\mathcal{D}}, d_{\min})$
 - 4: $\mathcal{R}, \mathcal{C} = \text{PARTITIONING}(\mathcal{G}_{\tau}, \mathcal{L}, \mathbf{X}_{\mathcal{T}}^{\mathcal{D}})$
 - 5: $\mathcal{G}_{\tau} = \text{EDGEGREEDYSEARCH}(\mathcal{L}, \mathcal{R}, \mathcal{C}, \mathbf{X}_{\mathcal{T}}^{\mathcal{D}}, \tau)$
 - 6: **end while**
-

Algorithm

Here, we introduce our algorithm SPACETIME¹ for causal discovery over space (contexts) and time (regimes). Following our theory, we develop our algorithm around discovering the causal model that minimizes the summed description lengths in Eq. (2). Our algorithm’s modular elements are therefore discovering the graph, changepoints, and partitions. We first describe how we discover a WCG \mathcal{G}_{τ} for fixed \mathcal{L}, \mathcal{C} , and \mathcal{R} ; then how we leverage changes of the functional models in \mathcal{G}_{τ} to reveal contexts \mathcal{C} and regimes \mathcal{R} ; and last, how to identify the correct changepoints \mathcal{L} given \mathcal{G}_{τ} . We iterate these steps until convergence.

Edge-Greedy Search MDL naturally allows ordering potential edges in the causal graph by their causal strength. Therefore, we adopt an edge-greedy search as proposed by Mian, Marx, and Vreeken (2021). Starting with an empty WCG \mathcal{G}_{τ} for a prespecified time lag τ , we first consider all pairwise edges among its nodes. We define the edge strength of a directed edge $E : X_{(i)(t)} \rightarrow X_{(j)(t')}$ under the current model \mathcal{M} as the relative gain in compression

$$\delta E := L(\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}; \mathcal{M}) - L(\mathbf{X}_{\mathcal{T}}^{\mathcal{D}}; \mathcal{M}')$$

where \mathcal{M}' differs from \mathcal{M} only by the additional edge E .

We proceed with a forward phase where we populate the graph \mathcal{G}_{τ} , adding edges in decreasing order of causal strength. Because this process is greedy, we may add edges that later become obsolete, so we follow it by a pruning phase. In this backward phase, we refine each variable’s parent set and time lags based on our score and finally return the resulting WCG \mathcal{G}_{τ} .

Regime- and Context Partitioning To discover regimes and contexts, we need to test for differences in causal mechanisms. Given two subsamples $s = (d, w), s' = (d', w')$ for any time intervals w, w' and $d, d' \in \mathcal{D}$, we test for (in)equality in conditional distribution under our model,

$$H_0 : P(X \mid pa(X), \theta^{(d,w)}) \neq P(X \mid pa(X), \theta^{(d',w')})$$

where $pa(X)$ are the causal parents of X in \mathcal{G}_{τ} and θ the GP parameters for the respective samples s, s' . For this purpose, we apply kernelized hypothesis testing (Huang et al. 2020;

¹Code available at <https://eda.group/spacetime>

Perry, Kügelgen, and Schölkopf 2022). Specifically, introducing a discrete variable S that labels samples in s with $S = 0$ resp. s' with $S = 1$ we test the conditional independence $X \perp\!\!\!\perp S \mid pa(X)$ in the pooled data over s and s' using the kernel test HSiC (Zhang et al. 2011).

We recall that we define the contexts in a spatial dimension and not constant over time. Therefore, we can perform this test pairwise over the datasets in \mathcal{D} , considering measurements aggregated over all of \mathcal{T} . To arrive from pairwise tests to a partition \mathcal{C} , we group the datasets such that no group contains two datasets for which the test indicated a mechanism change.

Regime shifts, on the other hand, may occur over time irrespective of the context. Thus, we obtain the partition \mathcal{R} by aggregating the pairwise tests performed over the time windows induced by \mathcal{L} with all datasets combined.

Changepoint Detection Unlike the dataset indices for the context partition, the regime change points must be identified beforehand. We propose an efficient method that takes advantage of the causal knowledge to identify change points without having to enumerate and test for mechanism change over all possible combinations over \mathcal{T} .

We recall that two regimes are distinguished by a change in the functional relationship between at least one variable and one of its parents. A functional model fitted on data from one regime should poorly predict data from another. Hence, we identify the change points \mathcal{L} by looking for changes in the GPs prediction error over time, instead of directly in the raw data where it might not be obvious.

We begin by fitting GPs to each variable and context within a time window where the regime is assumed constant, starting from the last known change point and of length equal to the minimum distance d_{min} between change points. We use the resultant GPs to predict the variable values across the entire dataset. Then, we apply a statistical change point detection technique to the prediction errors across all variables and contexts simultaneously, identifying the next change point. Many different techniques can be used here; a good choice is PELT (Killick, Fearnhead, and Eckley 2012), an efficient algorithm that does not require the number of change points.

Iterative Model Learning We apply an iterative strategy to combine the above steps, as summarized in Alg. 1. Initially, we segment the temporal space \mathcal{T} into change points under an empty graph (Line 1), given that this captures marginal distribution shifts. An iteration consists of detecting change points \mathcal{L} (Line 3) using changes in residual distribution of the functions $f^{(k,r)}$ under the current graph; learning the regime and context partitions (Line 4) using the current graph and updated change points; and finally, discovering \mathcal{G}_τ (Line 5) with a greedy strategy, where we need the current change points and partitions to obtain the scores in Eq. (2). We repeat these steps until convergence. Pseudocodes of the components are deferred to the Appendix.

Experiments

In this section, we evaluate our approach on synthetic data and real-world datasets in hydrology and meteorology.

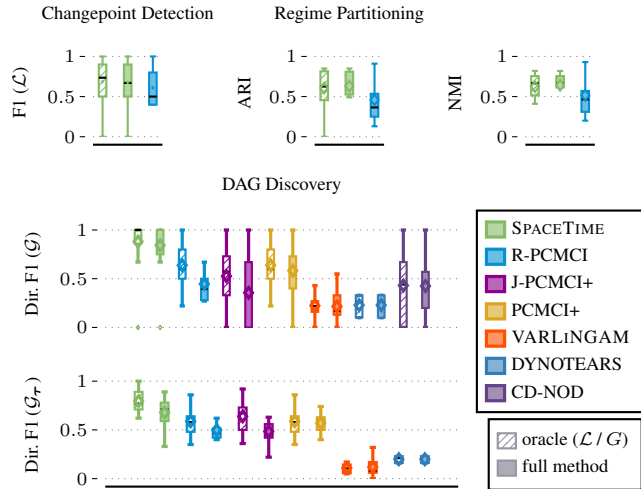


Figure 2: *Causal Discovery, Changepoint Detection and Regime Partitioning.* We evaluate the methods on multiple time series with causal mechanism shifts across time and datasets, with non-linear functional form, Gaussian noise, where $|\mathcal{G}_\tau| = 5$, $|\mathcal{T}| = 200$, $|\mathcal{C}| = 2$, $|\mathcal{R}| = 3$, $|\mathcal{L}| = 2$, and a fraction $s = \frac{1}{2}$ of intervened edges.

Experimental Setup

To simulate time series datasets, we generate random WCGs \mathcal{G}_τ with maximum lag $\tau = 2$, with corresponding summary DAG \mathcal{G} . We sample regime change points \mathcal{L} uniformly at random using a pre-set minimal duration, set to $d_{min} = 30$ unless otherwise specified. In each context and regime, we re-sample the edge weights of a fraction of all edges, and finally sample data similarly to Günther, Ninad, and Runge (2023) using TIGRAMITE Runge (2020), with linear or non-linear functional form and Gaussian or uniform noise.

Our first real-world case study investigates the effect of meteorological drivers on river discharge over different gauged catchments across Europe, given data derived from the Global Runoff Data Centre (GRDC) datasets (Cornes et al. 2018). Secondly, we study biosphere–atmosphere fluxes in the FLUXNET dataset (Baldocchi 2014). Both cases include multiple datasets \mathcal{D} from diverse geographical locations where the respective time series span multiple years, with daily time resolution. After pre-processing (Günther et al. 2023; Krich et al. 2021), we obtain $|\mathcal{D}| = 307$ locations for the river runoff, resp. $|\mathcal{D}| = 63$ locations for the fluxes data, and set \mathcal{T} to a period of one year (2006 resp. 2010) during DAG discovery given that the available years are non-overlapping. We provide the datasets \mathcal{D} from all locations to our method, using fixed hyperparameters $d_{min} = 30$ and $\tau = 2$ for consistency.

Throughout our experiments, we are interested in how the modular components of SPACETIME work together, which leads to the following questions: whether it discovers causal directions correctly over time and contexts; whether it accurately detects change points over time; and whether it finds meaningful partitions of similar causal mechanisms.

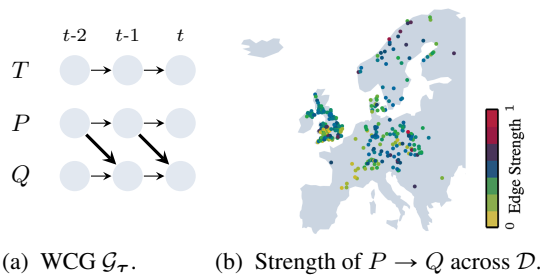


Figure 3: *Drivers of River Discharge*. Given temperature (T), precipitation (P) and river discharge (Q) in 307 catchments across Europe, we discover a common WCG \mathcal{G}_τ for $\tau = 2$ (a) with a lagged edge $P \rightarrow Q$, and illustrate its variability in causal strength between catchments (b).

Synthetic Data

We show our results on synthetic data in Fig. 2. We report F1 scores over \mathcal{L} (top) for changepoint detection, F1 scores over directed edges in \mathcal{G}_τ (middle) and the DAG \mathcal{G} (bottom) for causal discovery, and the Adjusted Rand Index (ARI) and Normalized Mutual Information (NMI) for regime partitioning. We also report changepoint detection with known \mathcal{G}_τ resp. causal discovery with known \mathcal{L} (hatched). How we apply the methods depends on whether they can handle multiple datasets (ours, JPCMCI) resp. multiple regimes (ours, R-PCMCI). We apply methods that assume a single dataset to each d in turn. Across time, we apply methods that assume a single regime either to each subsample s with known \mathcal{L} (hatched), or all regimes together (solid). We aggregate all resulting graphs using a majority vote for each edge in the graph. We find that SPACETIME consistently performs well compared to its competitors, exemplified in Fig. 2 for the non-linear Gaussian case, and with varying functional form and noise distributions which we postpone to the Appendix.

Causal Discovery Regarding causal discovery, the methods DYNOTEARS and VARLINGAM make linearity assumptions and the latter in addition assumes that the noises in the model are non-Gaussian. Therefore, it is not surprising that they show sub-par performance here. Even though they and PCMCI+ assume stationarity, i.e. the absence of regime shifts, the effect of known (hatched) compared to unknown changepoints (solid) is not overly pronounced. This could be due to a slight advantage due to increased sample size when combining all regimes and the fact that changes in edge strength are not as drastic as changes in the graph. While CD-NOD can address temporal shifts, it only discovers a summary DAG \mathcal{G} . As all the aforementioned methods assume a single time series, note that we apply them to each dataset and combine the resulting graphs using majority voting. Only J-PCMCI+ can consider multiple contexts, but does not achieve a significant performance gap here. This may be due to mechanism shifts being too subtle, given that J-PCMCI+ relies on changing graphs across contexts, or because it may need a larger number of different contexts to effectively outperform single-context methods, in which case conditional independence tests however




Figure 4: *Changepoints of the Interaction between Precipitation and Runoff*. We show the regime change points \mathcal{L} and partition \mathcal{R} that we discover for $P \rightarrow Q$ with SPACETIME for selected catchments d_1 (CH), d_2 (PL), d_3 (GB). The colors denote different regimes.

become increasingly expensive due to a high dimensional one-hot encoding.

Changepoint and Regime Identification In the changepoint detection and partitioning experiments (Fig. 2 top), we can see that SPACETIME is robust to a small number of mistakes in the causal graph as the full results (solid) are close to the oracle version (hatched). This also matches our observation that the interleaving steps in Alg. 1 converge quickly, typically within three iterations. Reasons that we outperform R-PCMCI (Fig. 2) may include the fact that we impose regime persistence explicitly, whereas R-PCMCI can only impose a given number of regime-switches here provided as background knowledge.

Case Study: Identifying Drivers of River Discharge

As a real-world application, we explore the effects of meteorological variables on river discharge using data from so-called catchments, that is, an area where water drains into an outlet. We focus on temperature (T), precipitation (P) as well as river runoff (Q) in our analysis.

Causal Discovery over Different Catchments While a previous study of this data using PCMCI+ was limited in considering each location and month separately (Günther et al. 2023), we are interested in a unified analysis of all catchments. As we show in Fig. 3a, we discover a direct influence of precipitation on discharge through the lagged edge $E : P_{(t-1)} \rightarrow Q_{(t)}$ which we consider reasonable. The joint graph aside, we also visualize the causal strength δE of this edge for each of the 307 locations in the year 2010 in Fig. 3b. We observe regional similarities but also geographical variation in causal strength. The latter is likely due to differences in topology and climate in the different parts of Europe, plus distinct characteristics of each catchment such as area, altitude, and vegetation (Günther et al. 2023).

Changepoints and Regimes over Time Fig. 4 shows for selected examples the changepoints and regimes with similar causal relationships that SPACETIME discovers.

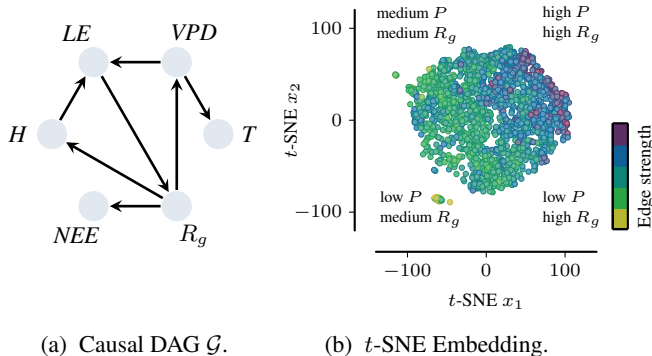


Figure 5: *Biosphere-Atmosphere Interactions*. For the FLUXNET data from 64 locations \mathcal{D} in multiple years \mathcal{T} , we show the DAG \mathcal{G} (a) that SPACETIME discovers. We visualize the causal strengths of the seven edges in \mathcal{G} in two dimensions using t -SNE (b), where each sample (x_1, x_2) corresponds to a fixed location d_i , month m_j , and year y_k . We recover regions that correspond to distinct underlying meteorological conditions (precipitation P , global radiation R_g).


Case Study: Reconstructing Gradients of Biosphere-Atmosphere Interactions

In the FLUXNET datasets, we study interactions of meteorological and atmospheric variables, with a primary focus on air temperature (T), global radiation (R_g), net ecosystem exchange (NEE), sensible heat (H), latent heat flux (LE) and vapor pressure deficit (VPD).

Causal Strengths across Datasets We show the common graph jointly over all locations and months in Fig. 5a. As there is no known ground truth regarding the network, we replicate the qualitative analysis in Krich et al. (2021) to see whether the causal strengths across locations and time correspond to plausible interactions between atmospheric and meteorological variables. To this end, we consider the MDL edge strengths δE for the seven edges in \mathcal{G}_τ for each location $d \in \mathcal{D}$ and each month $m \in \mathcal{T}_d$, and project them to two dimensions using the t -distributed stochastic neighbor embedding (t -SNE). The result is shown in Fig. 5b, where each sample (x_1, x_2) represents a location d and month m .

Perhaps surprisingly, the grouping by causal strength is not strongly tied to the actual geographic location and corresponding ecosystem, but rather correlates with water availability and global radiation. For example, we find strong edge connections in \mathcal{G} (upper right) to be associated with high radiation and precipitation, even though these variables are not directly used during causal discovery. As noted in prior work (Kraemer et al. 2020; Krich et al. 2021), these conditions create an optimal growing environment in the respective ecosystems and months. This opposes a low-energy output state often reached in the winter time, which we reconstruct with weak edge connectivity in \mathcal{G} (lower left).

Grouping the biosphere-atmosphere interactions by causal strength therefore reveals different meteorological



(a) Regimes \mathcal{R} in d_1 : De-Hai. (b) Regimes \mathcal{R} in d_2 : De-Tha.

Figure 6: *Regime partitioning reveals distinct trends*. The regime partitions that SPACETIME discovers reveal abnormal conditions, such as the effects of a European heatwave (2003) on the FLUXNET ecosystems Hainich (De-Hai) and Tharandt (De-Tha) (2000-2009).

states that the ecosystems traverse over the year. In particular, matching previous insights (Krich et al. 2021) these findings support the hypothesis that distinct ecosystems can traverse very similar meteorological states over time.

Distinct Trends over Time A limitation of the t -SNE visualization is that the resulting space is not metric and thus does not provide a direct way to perform comparisons (Krich et al. 2021). With SPACETIME, we can test for statistical differences in causal mechanism, and, for example, investigate how the causal interactions evolve over the years at a given location. To illustrate, we show in Fig. 6 for the locations Hainich (DeHai) and Tharandth (DeThai) the regime partitions over multiple years (here 2000-2009). We discover a common causal mechanism in most years (greyed out) and distinct regimes in 2003, resp. 2006 (colored). Such changes likely result from abnormal meteorological conditions; in this case, a drought event in Europe in 2003, which is reflected in the monthly precipitation P (Fig. 6).

Conclusion

While causal discovery often treats cause-effect relationships as fixed, most real-world processes change and evolve, often in multiple environments and over time. Thinking of such settings, we study causal discovery and changepoint detection for a collection of multivariate time series. We model causal mechanism shifts across datasets and time jointly using the notions of contexts and regimes. We propose discovering temporal causal graphs using Minimum Description Length (MDL) encodings for Gaussian processes (GPs), where we detect regime changepoints from changes in residual distributions. Alongside this, we discover contexts and regimes with similar causal mechanisms using kernelized hypothesis testing. We confirmed that our strategy for discovering causal models, changepoints, and partitions, summarized in the SPACETIME algorithm, works well in our simulations. On real-world data, our method can be used to discover summary networks, varying edge strengths, and meaningful groups of time spans and locations with similar mechanisms. This supports a research direction where we base cluster analysis not only on the marginal distributions, but on causal interaction structures between variables of interest (Günther et al. 2023), which future continuations of this work could explore further.

Acknowledgments

U.N. has received funding from the European Research Council (ERC) Starting Grant CausalEarth under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 948112).

We thank the anonymous reviewers for their helpful comments.

References

- Assaad, C.; Devijver, E.; and Gaussier, E. 2022. Survey and Evaluation of Causal Discovery Methods for Time Series. *Journal of Artificial Intelligence Research*, 73: 767–819.
- Baldocchi, D. 2014. Measuring fluxes of trace gases and energy between ecosystems and the atmosphere – the state and future of the eddy covariance method. *Global Change Biology*, 20(12): 3600–3609.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3: 507–554.
- Cornes, R.; Schrier, G.; Van den Besselaar, E.; and Jones, P. 2018. An Ensemble Version of the E-OBS Temperature and Precipitation Data Sets. *Journal of Geophysical Research Atmospheres*, 123.
- Gao, S.; Addanki, R.; Yu, T.; Rossi, R. A.; and Kocaoglu, M. 2024. Causal discovery in semi-stationary time series. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23. Red Hook, NY, USA: Curran Associates Inc.
- Gerhardus, A. 2024. Characterization of causal ancestral graphs for time series with latent confounders. *The Annals of Statistics*, 52(1): 103 – 130.
- Granger, C. W. J. 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37(3): 424–438.
- Grünwald, P. 2007. *The Minimum Description Length Principle*. MIT Press.
- Günther, W.; Ninad, U.; and Runge, J. 2023. Causal Discovery for time series from multiple datasets with latent contexts. In Evans, R. J.; and Shpitser, I., eds., *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, 766–776. PMLR.
- Günther, W.; Miersch, P.; Ninad, U.; and Runge, J. 2023. Clustering of causal graphs to explore drivers of river discharge. *Environmental Data Science*, 2.
- Huang, B.; Zhang, K.; Zhang, J.; Ramsey, J.; Sanchez-Romero, R.; Glymour, C.; and Schölkopf, B. 2020. Causal Discovery from Heterogeneous/Nonstationary Data. *Journal of Machine Learning Research*, 21(1).
- Hyvärinen, A.; Zhang, K.; Shimizu, S.; and Hoyer, P. O. 2010. Estimation of a Structural Vector Autoregression Model Using Non-Gaussianity. *Journal of Machine Learning Research*, 11(56): 1709–1731.
- Janzing, D.; and Schölkopf, B. 2010. Causal inference using the Algorithmic Markov Condition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32: 5168–5194.
- Killick, R.; Fearnhead, P.; and Eckley, I. 2012. Optimal Detection of Changepoints With a Linear Computational Cost. *Journal of the American Statistical Association*, 107: 1590–1598.
- Kraemer, G.; Camps-Valls, G.; Reichstein, M.; and Mahecha, M. D. 2020. Summarizing the state of the terrestrial biosphere in few dimensions. *Biogeosciences*, 17(9): 2397–2424.
- Krich, C.; Migliavacca, M.; Miralles, D. G.; Kraemer, G.; El-Madany, T. S.; Reichstein, M.; Runge, J.; and Mahecha, M. D. 2021. Functional convergence of biosphere-atmosphere interactions in response to meteorological conditions. *Biogeosciences*, 18(7): 2379–2404.
- Li, M.; and Vitányi, P. 2009. *An introduction to Kolmogorov complexity and its applications*. Springer Science & Business Media.
- Mameche, S.; Kaltenpoth, D.; and Vreeken, J. 2023. Learning Causal Models under Independent Changes. In *Advances in Neural Information Processing Systems*, volume 36, 75595–75622. Curran Associates, Inc.
- Mameche, S.; Vreeken, J.; and Kaltenpoth, D. 2024. Identifying Confounding from Causal Mechanism Shifts. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, 4897–4905. PMLR.
- Marx, A.; and Vreeken, J. 2021. Formally Justifying MDL-based Inference of Cause and Effect. arXiv:2105.01902.
- Mian, O. A.; Marx, A.; and Vreeken, J. 2021. Discovering Fully Oriented Causal Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(10): 8975–8982.
- Pamfil, R.; Sriwattanaworachai, N.; Desai, S.; Pilgerstorfer, P.; Georgatzis, K.; Beaumont, P.; and Aragam, B. 2020. DYNOTEARS: Structure Learning from Time-Series Data. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, 1595–1605. PMLR.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition.
- Perry, R.; Kügelgen, J. V.; and Schölkopf, B. 2022. Causal Discovery in Heterogeneous Environments Under the Sparse Mechanism Shift Hypothesis.
- Rasmussen, C. E.; and Williams, C. K. I. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press. ISBN 026218253X.
- Rissanen, J. 1978. Modeling by shortest data description. *Automatica*, 14(1): 465–471.
- Runge, J. 2020. Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In Peters, J.; and Sontag, D., eds., *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124 of *Proceedings of Machine Learning Research*, 1388–1397. PMLR.

Saggioro, E.; de Wiljes, J.; Kretschmer, M.; and Runge, J. 2020. Reconstructing regime-dependent causal relationships from observational time series. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(11): 113115.

Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward Causal Representation Learning. *Proceedings of the IEEE*, 109(5): 612–634.

Shah, R. D.; and Peters, J. 2020. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3): 1514 – 1538.

Spirtes, P.; Glymour, C. N.; Scheines, R.; and Heckerman, D. 2000. *Causation, prediction, and search*. MIT Press.

Zhang, K.; Peters, J.; Janzing, D.; and Schölkopf, B. 2011. Kernel-Based Conditional Independence Test and Application in Causal Discovery. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI'11*, 804–813. Arlington, Virginia, USA: AUAI Press. ISBN 9780974903972.