
CausalDynamics: A large-scale benchmark for structural discovery of dynamical causal models

Benjamin Herdeanu*
kausable GmbH
benjamin@kausable.ai

Juan Nathaniel*
Columbia University
jn2808@columbia.edu

Carla Roesch*
Columbia University
cmr2293@columbia.edu

Jatan Buch
Aeolus Labs
jatan@aeolus.earth

Gregor Ramien
kausable GmbH
gregor@kausable.ai

Johannes Haux
kausable GmbH
johannes@kausable.ai

Pierre Gentine
Columbia University
pg2328@columbia.edu

Abstract

Causal discovery for dynamical systems poses a major challenge in fields where active interventions are infeasible. Most methods used to investigate these systems and their associated benchmarks are tailored to deterministic, low-dimensional and weakly nonlinear time-series data. To address these limitations, we present *CausalDynamics*, a large-scale benchmark and extensible data generation framework to advance the structural discovery of dynamical causal models. Our benchmark consists of true causal graphs derived from thousands of both linearly and nonlinearly coupled ordinary and stochastic differential equations as well as two idealized climate models. We perform a comprehensive evaluation of state-of-the-art causal discovery algorithms for graph reconstruction on systems with noisy, confounded, and lagged dynamics. *CausalDynamics* consists of a plug-and-play, build-your-own coupling workflow that enables the construction of a hierarchy of physical systems. We anticipate that our framework will facilitate the development of robust causal discovery algorithms that are broadly applicable across domains while addressing their unique challenges. We provide a user-friendly implementation and documentation on <https://kausable.github.io/CausalDynamics>.

1 Introduction

Understanding causal mechanisms in nonlinear dynamical systems is crucial across many fields. However, direct interventions are often impractical or impossible. An alternate strategy for encoding causal information is through detailed physical modeling, such as Earth System models, which relies on ad-hoc parameterizations and can be computationally expensive. In contrast, data-driven causal discovery frameworks have emerged as a promising avenue to infer cause–effect relationships directly from time series observations, obviating the need for interventions or simulations. Originating from Wright’s path analysis in the 1930s [1] and later formalized by Wiener [2] and Granger [3] in the mid-20th century, causal inference now includes deep learning (DL)-based methods [4, 5], designed to infer complex, nonlinear, and stochastic causal structures.

*Equal contribution




Figure 1: Illustration of the tiered framework in **CausalDynamics** consisting of a plug-and-play, build-your-own coupling workflow that enables the construction of a hierarchy of physical systems with common causal challenges, such as unobserved confounders, time-lags, and noisy time series.

Despite rapid methodological advances, there exists no standardized benchmark for evaluating causal discovery in highly nonlinear, dynamical settings [6, 7, 8], which is critical to understand and predict the behavior of physical systems [9, 10, 11, 12, 13, 14]. Most benchmarks are built on synthetic data generated by static causal graphs (e.g., [15]) or auto-regressive models (e.g., [16, 17]) for domain-specific applications [18, 19]. Limited examples from the real-world are available but lack a fully resolved causal ground truth [20, 21, 22]. As a result, most methods are validated on toy systems or on real-world data failing to capture continuous state-space developments, complex feedback loops, stochasticity, and regime shifts, making it impossible to isolate algorithmic limitations (e.g., varsortability [23, 24]) from dataset characteristics.

We posit that advancing the development of new causal discovery methods capable of capturing high dimensional, nonlinear, time-lagged, sparse, and noisy structures with unobserved confounding, often found in real-world physical systems [25, 26, 27], requires a novel benchmark. In fact, examples from various fields demonstrate how such benchmarks also play a vital role in unlocking methodological innovation for data-driven methods and boosting scientific progress. A prominent example includes CASP13, a protein structure modeling benchmark that enabled the development of AlphaFold, a co-awardee of the 2024 Nobel Prize in Chemistry [28, 29, 30]. Other benchmarks include ImageNet [31] which highlighted the computer vision capabilities of AlexNet [32] or NWPU-RESISC45 which first demonstrated the superior skill of deep learning approaches for remote-sensing classification [33].

In this paper, we present **CausalDynamics**, a tiered benchmark and extensible data generation framework, as shown in Figure 1, to advance the structural discovery of dynamical causal models. The simple tier contains true causal graphs for hundreds of three-dimensional chaotic dynamical systems. In the second tier, we adapt an existing algorithm [34] for graph generation to couple deterministic and stochastic dynamical systems with periodic functions for constructing thousands of complex graph structures. This framework allows us to address the set of challenges often found in real-world systems as highlighted previously. The final tier contains the true causal graphs for two idealized atmosphere-ocean models including multiple coupling experiments of different modes of variability.

Simultaneously generating pseudo-real or idealized data has a two-fold advantage: (i) We can benchmark causal discovery methods with known causal graphs and well-designed challenges, and (ii) once calibrated, methods can be extended to realistic applications. This way, we ensure a comprehensive workflow for developing trustworthy, novel causal discovery methods. **CausalDynamics** provides that framework and builds a robust foundation for testing any novel algorithm’s applicability to real-world scenarios.

Our contributions include:

- **CausalDynamics**, the largest benchmark for causal discovery methods for stochastic chaotic systems containing over 14000 graphs and over 50 million preprocessed samples.
- A novel graph generating algorithm that allows users to easily extend the complexity of the synthetic dataset as more powerful causal discovery methods are developed.
- An all-in-one benchmark for evaluating causal discovery methods such that any new algorithm can be reliably evaluated on synthetic data with known graph structure before being tested on a real-world example within the same framework.
- Evaluation of the performance of several state-of-the-art (SOTA) DL-based and non-DL-based causal discovery algorithms for graph reconstruction on our benchmark.

A quickstart guide to **CausalDynamics** can be found in Appendix A and at <https://kausable.github.io/CausalDynamics/notebooks/quickstart.html>.

2 Theoretical background

In this section, we will give a brief background on the connection between dynamical systems and structural dynamical causal models [35] to introduce the necessary notation.

2.1 Structural dynamical causal model

Dynamical system. For each time $t \in \mathbb{R}_{\geq 0}$, we characterize a dynamical system with an associated state $x(t) \in \mathbb{R}^N$ for $N \in \mathbb{N}$. In general, the description of the system dynamics are given through differential equations of the form:

$$\frac{dx}{dt} = f(t, x) + \delta \frac{dW_t}{dt} \quad (1)$$

where f is a function, the solution $x(t)$ depends on the initial condition $x(t_0) = x_0$ at time t_0 , and δ is the noise amplitude of the Brownian process W_t [36, 37, 38]. When $\delta = 0$, Equation 1 becomes *ordinary differential equations* (ODEs) whereas $\delta > 0$ yields *stochastic differential equations* (SDEs).

Structural causal models. We describe causal mechanisms through structural causal models (SCMs) such that a system of d random variables $\mathbf{x} = \{x_1, \dots, x_d\}$ is expressed as an arbitrary function f^k of its direct parents (causes) $\mathbf{x}_{\text{PA}_k}$ and an exogenous distribution of noise ϵ^k [35]:

$$x_k := f^k(\mathbf{x}_{\text{PA}_k}, \epsilon^k), \text{ for } k = 1, \dots, d \quad (2)$$

For dynamical systems, we can extend Equation 2 with Equation 1 for a collection of d differential equations to define *structural dynamical causal models* (SDCM) [35, 39, 40]:

$$\frac{d}{dt}x_{k,t} := f^k(\mathbf{x}_{\text{PA}_k,t}, \delta), \text{ with } x_{k,0} = x_k(0) \quad (3)$$

where $k \in \{1, \dots, d\}$.

Causal graph. The structural assignment of the SCM induces a *directed acyclic graph* (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ over the variables x_k . $\mathcal{G}$ includes nodes $v_k \in \mathcal{V}$ for every $x_k \in \mathbf{x}$ and directed edges $(k, i) \in \mathcal{E}$ if $x_k \in \mathbf{x}_{\text{PA}_i}$ [23]. Edges are represented in a squared *adjacency matrix* $\mathcal{A} \in \mathbb{R}^{k \times i}$ with each entry $a_{ki} \in \mathcal{A} : a_{ki} = 1$ if x_i is causally impacted by x_k , else $a_{ki} = 0$. A corresponding graphical representation of the DAG is shown in Figure 2.

2.2 Causal challenges

We include a set of causal challenges that are common in real-world physical systems to maximize the applicability of our benchmark. The challenges include: (i) *Noise* can obscure statistical dependencies or identify spurious links [41]; (ii) *Hidden confounders*, i.e., unobserved variables that are a common cause of at least two other unrelated variables, induce correlations that algorithms may misinterpret as direct causal links [42, 43, 25]; (iii) Delays between cause and effect, i.e., *time-lags*, lead to causal




Figure 2: Illustration of the Rössler Oscillator ODEs (left), the associated adjacency matrix $\mathcal{A}$ (center) and the corresponding causal graph (right).

effects at multiple time scales, which obscure conditional independence test and result in spurious links [44]. Lastly, varsortability is an artifact in synthetically generated data from SCMs describing the increase in a variable’s variance along its topological order [45, 46]. Therefore, we also implement the option to (iv) *standardize* [23, 24].

3 Related work

In this section, we review existing benchmarks for causal discovery, highlighting their shortcomings for dynamical systems, and present an overview of existing causal discovery methods.

3.1 Benchmarks for causal discovery

Real-world. Real-world datasets for causal discovery remain rare because true causal graphs are difficult to obtain. A handful of benchmarks provides *time series* data, e.g. human-motion capture (MoCap) [5], S&P 100 stock prices [22], climate variability and teleconnections [47], river discharge (CausalRivers) [20], and ecological observations [48]. However, most observed collections focus on *static*, domain-specific graphs, such as large-scale single-cell RNA perturbations (CausalBench) [18, 19], the Old Faithful geyser eruptions [49] and immune-cell protein networks [50, 51]. Bivariate settings like the Tübingen Cause Effect Pairs [38] further complement these resources.

Synthetic. Due to the domain specifications, limited availability and often low dimensionality of real-world datasets, synthetic data plays an important role in understanding complex systems in general [52, 53, 54, 55] and benchmarking causal discovery methods [6] specifically. In *static* settings, SCMs or structural vector autoregressive (SVAR) models sample causal coefficients and noise to generate graphs [24] often following multivariate linear or nonlinear functions [49, 56, 57], or physical laws [58]. More realistic pipelines derive the causal graph to fit coefficients from observational data (e.g. `causalAssembly` [15], `SynTReN` [59]), which limits their application to the domain and the properties of the underlying observed dataset. For *time series*, idealized dynamical systems, e.g., Lorenz or Rössler Oscillators [60], underpin simpler benchmarks [5, 61, 62, 63], while domain-informed models include fMRI networks (Netsim) [64], gene regulatory networks (DREAM3 & 4) [65, 66] or financial data (FINANCE) [67, 68]. A collection of climate and weather benchmarks can be found on CauseMe (<https://causeme.uv.es/>) [21, 69] with pseudo-realistic climate data generated using the SAVAR model [17]. Going beyond single-domain applications, novel approaches propose pipelines to flexibly generate time series datasets for a range of properties and natural systems [70] from any observational dataset using DL (e.g., `CausalTime` [16]).

Time series data is critical for benchmarking SDCM discovery algorithms. However, as outlined above, relevant benchmarks contain only domain specific datasets consisting of a small number of graphs [64, 17, 16] and weakly nonlinear systems [20]. Further, even though datasets might contain thousands of samples [18], the underlying graphs are but a handful. Although observation-based approaches generate pseudo-real data for a range of domains, they lack a reliable ground truth validation [16, 70].

3.2 Causal discovery methods

In the following section, we provide a summary of the different causal discovery approaches which can be categorized into five classes. A detailed review of existing methods can be found in [4].

(i) *Granger causality* (GC) [3] is one of the oldest concepts in causal inference [71]. GC tests whether a given effect is optimally forecast by its causes under the assumptions of no unobserved confounders. GC often fails for nonlinear dynamics [48], leading to deep-learning variants: Neural GC (NGC) [5] employs regularized multilayer perceptrons (cMLP) or long short-term memory (cLSTM) to learn nonlinear autoregressive links, while CUTS+ uses a coarse-to-fine pipeline with a message-passing graph neural network to recover causal graphs from high-dimensional, irregular time series [72]. Another notable time-series extension is Temporal Causal Discovery Framework (TCDF) [68], which leverages convolutional neural networks with attention mechanisms to identify lagged causal links directly from multivariate sequences.

(ii) *Constraint-based* methods infer causal structure by enforcing the statistical constraints implied by conditional independencies in the data (e.g., PC [41]). PCMCi+ (Peter Clark Momentary Conditional Independence) [44] adapts PC to multivariate time series by preselecting candidate parents and performing Momentary Conditional Independence (MCI) tests over all time-lags to infer contemporaneous and lagged edges. F-PCMCi boosts scalability by prefiltering parents via transfer-entropy before MCI testing [73]. More recently, permutation-based methods such as Greedy Relaxations of the Sparsest Permutation (GRaSP) [74] have been developed, which search over permutations of variables and iteratively prune spurious edges to identify sparse causal graphs.

(iii) *Noise-based* causal discovery methods exploit the fact that in a correctly specified SCM, the noise term is statistically independent of its inputs only in the true causal direction. For example, LiNGAM [75] and its temporal extension VARLiNGAM [76] assume a linear model with non-Gaussian noise and recover a unique DAG and lagged links by analyzing the residuals, while additive noise models (ANMs) (e.g., [49]) generalize this to nonlinear relations by choosing the direction where residuals remain independent of their input. The recent Repetitive Causal Discovery (RCD) algorithm [77] extends the LiNGAM family by incorporating constrained functional causal models with regularization, enabling more robust identification in high-dimensional or noisy settings.

(iv) *Score-based* learning algorithms infer causal relationships between variables by evaluating and ranking potential causal graphs based on a score function such as least-squares error. For instance, DYNOTEARS [22] adopts a score-based SVAR formulation with a penalized least-squares loss and a differentiable acyclicity constraint to jointly recover instantaneous and lagged weights.

(v) *Topology-based* approaches are based on Takens’ theorem [78] to reconstruct the attractor of a dynamical system from delay-embedded time series data. TSCI (Tangent Space Causal Inference) [79] then applies Convergent Cross Mapping [48] within each manifold’s tangent space to test for causal influence by assessing how well the state-space of one variable predicts another.

4 CausalDynamics

In the following, we introduce **CausalDynamics**, a large-scale benchmark and extensible data generation framework as illustrated in Figure 1. We produce three tiers of datasets, each introducing a progressively greater level of complexity as outlined in the rest of the section. Our benchmark consists of true causal graphs derived from thousands of coupled ordinary and stochastic differential equations (tiers 1 & 2) as well as two idealized climate models (tier 3). The code is available at <https://github.com/kausable/CausalDynamics> and the dataset can be downloaded from <https://huggingface.co/datasets/kausable/CausalDynamics>. Full documentation can be found at <https://kausable.github.io/CausalDynamics>. A sample will also be available on CauseMe [44].

4.1 Tier 1 – Simple causal models

In the *simple* complexity tier, we select 59 three-dimensional systems from the *dysts* benchmark [60, 80] and derive their SDCMs as adjacency matrices. For each system, we simulate trajectories with 5 random initial conditions over 1000 time steps.

We perform two experiments in this tier: (i) *unobserved confounding* by excluding the respective time series during evaluation (x in the system in Figure 2), and (ii) we evaluate for *varying Langevin noise amplitude* (δ) exploiting the *dysts* support for both deterministic and stochastic integration schemes.

Algorithm 1 Growing Network with Redirection (GNR) model from [34]

Require: $n \in \mathbb{N}, r \in [0, 1]$

- 1: $\mathcal{A} \leftarrow \text{zeros}(n \times n); K \leftarrow \text{zeros}(n); v_a \leftarrow \text{zeros}(n)$ ▷ adj. matrix; attachm. kernel; ancestors
- 2: **if** $n = 1$ **then** ▷ Case single node
- 3: $\mathcal{A} \leftarrow [[0]]; K \leftarrow [0]; v_a \leftarrow [0]$
- 4: **return** $\mathcal{A}$
- 5: **end if**
- 6: $K[0] \leftarrow 1; v_a[0] \leftarrow 0$ ▷ Initialize attachment kernel and ancestors
- 7: **for** $i \leftarrow 1$ **to** $n - 1$ **do**
- 8: $K^{\text{prob}} \leftarrow K[0 : i] / \sum(K[0 : i])$ ▷ Normalize to get probabilities
- 9: $m \leftarrow \text{Multinomial}(K^{\text{prob}}, 1)$ ▷ Sample node based on probabilities
- 10: **if** $\text{random}(n - 1) < r = \text{True}$ **then**
- 11: $m \leftarrow v_a[m]$ ▷ Redirect to ancestor node
- 12: **end if**
- 13: $\mathcal{A}[i, m] \leftarrow 1$ ▷ Add edge in graph
- 14: $K[i] \leftarrow K[i] + 1; K[m] \leftarrow K[m] + 1$ ▷ Update source & target attachment kernels
- 15: $v_a[i] \leftarrow m$ ▷ Track ancestor
- 16: **end for**
- 17: **return** $\mathcal{A}$

4.2 Tier 2 – Hierarchically coupled causal models

In the second complexity tier, we introduce hierarchically *coupled* causal models, which move beyond isolated dynamical systems and simulate causal interdependence between multiple driving processes. For this, we adapt an algorithm first proposed by [34] and extend it to conduct flexible experiments on pseudo-real causal challenges.

Sampling. Coupled models are sampled using a Growing Network with Redirection (GNR) model following [34], which we refer to as *scale-free* DAGs due to their underlying preferential attachment nature [81] with redirection mechanisms. The redirection probability r controls the balance between preferential attachment and ancestor-based connections. We outline the GNR model for graph generation in Algorithm 1 and visualized in Figure 3.

Causal units. Each node $v_k \in \mathcal{V}$ represents a causal unit consisting of d time series $x_{v_k}(t) \in \mathbb{R}^d$ for $k \in n$, where $d = \{1, 3\}$.

Root nodes, i.e., nodes without incoming edges, are initialized using one of the following drivers (Figure 3a):

- *Dynamical drivers:* Chaotic systems sampled from the `dysts` package (tier 1), such as Lorenz or Rössler ($d = 3$).
- *Periodic drivers:* Sinusoidal functions of the form $x_{v_k}(t) = A \sin(\omega t + \phi)$, capturing seasonal or oscillatory influences ($d = 1$).
- *Linear drivers:* Linear functions $x_{v_k}(t) = mt + b$ modeling continuous temporal changes such as rising global mean temperatures in first order approximation.

Edges. Non-root nodes integrate the transformed signals received from their parent nodes via the information carried by the edges [82, 83]. Each edge $(k, i) \in \mathcal{E}$ is realized as an MLP [84] with optional activations $\phi_{(k,i)}$:

$$\begin{aligned} f_{(k,i)}(x_k(t)) &= \phi_{(k,i)}(W_{(k,i)}x_k(t) + b_{(k,i)}), \\ W_{(k,i)} &\sim \mathcal{N}(0, 1), \quad b_{(k,i)} \sim \mathcal{N}(0, 1), \\ \phi_{(k,i)} &\sim \text{Uniform}(\{\text{identity}, \text{sin}, \text{sigmoid}, \text{tanh}, \text{ReLU}\}) \end{aligned} \tag{4}$$

where $W_{(k,i)} \in \mathbb{R}^{d \times d}$ and $\phi_{(k,i)} = \text{identity}$ in case of no activations. Weights are sparsified with a dropout probability p_{zero} . The value of node v_k at time t is computed by aggregating its incoming signals:

$$x_{v_k}(t) = \sum_{k \in \text{pa}_i} f_{(k,i)}(x_k(t)). \tag{5}$$




Figure 3: (a) Hierarchically coupled graphs are sample as scale-free DAGs, where root nodes can either be driven by period functions (root node 1) or dynamical systems (root node 2). Information is passed to the leaf nodes (leaf node 0) through MLPs. To correct for varsortability node values can optionally be standardized. To create a diverse set of challenges we introduce (b) unobserved confounders (green dotted lines), or (c) time-lag τ (red node). The type of nonlinearity can also be varied by prescribing or randomly assigning different edge-level activation functions. Assigned drivers are plotted by solid black edges. Note that each node is $\in \mathbb{R}^d$ and each edge is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$.

Following this flexible coupling we can generate the following experiments:

- *Confounder*: By sampling two adjacency matrices, i.e. two sets of edges, for a set of nodes, rotating the off-diagonal entries in the second matrix by 90 degrees, and then combining the two, we can introduce confounders as visualized in Figure 3b. This method also ensures that the confounded graph is a scale-free DAG.
- *Time-lag*: To model delayed effects, we introduce time-lagged edges as shown in Figure 3c. For a fixed lag τ , selected edges introduce temporal delay:

$$x_{v_k}(t) = f_{(k,i)}(x_k(t - \tau)), \quad (6)$$

breaking the acyclic constraint over time and introducing temporally cyclic graphs. The delay τ is constant per graph instance and such edges are sampled with probability p_t .

- *Standardized*: Similar to [23, 24] we standardize node values to remove scaling artifacts over the time dimension:

$$\tilde{x}_{v_k}(t) = \frac{x_{v_k}(t) - \mu_{v_k}}{\sigma_{v_k}}, \quad (7)$$

where μ_{v_k} and σ_{v_k} are the mean and standard deviation over time.

More details on data generation for coupled systems, including pseudocode, graphical illustrations, and a detailed sketch of the procedure of standardization are included in Appendix B.

4.3 Tier 3 – Pseudo-real physical systems

In the highest tier, we generate pseudo-real *climate* data of coupled atmosphere–ocean dynamics. The Earth’s climate is a high-dimensional, non-equilibrium, chaotic, complex system and its predictability remains an active area of research [85, 86, 87, 88].

In a low-dimensional setting, we model the El Niño–Southern Oscillation (ENSO) as implemented in the XR0 package [89], which merges the Hasselmann stochastic framework [90] and recharge oscillator dynamics [91]. As XR0 initializes with observed sea surface temperatures and thermocline depth, it reproduces key observational ENSO statistics [92]. The package also links multiple ocean-basin modes, allowing for adjustable coupling strengths to simulate tightly coupled or largely independent behavior. To simulate higher dimensional atmosphere–ocean dynamics, we utilize the qgs package of a simplified quasi-geostrophic two-layer model resolving barotropic and baroclinic interactions derived from the Modular Arbitrary-Order Ocean-Atmosphere Model (MAOOAM) [93].

Between XR0’s observations-based stochastic, oscillatory dynamics and qgs’s high dimensional coupling, our benchmark includes datasets that closely mimic both observed dynamical systems and operational climate model outputs. Appendix C contains more details of the implemented models.

4.4 Dataset summary

We provide a summary of the preprocessed dataset in Table 1 and Figure 13. Unless otherwise stated, each graph constitutes 5 randomly initialized trajectories with 1000 time steps. In total, we generated 585 simple, 14096 coupled, and 12 climate graphs, for a total of 14693 graphs. More details on the selected parameters can be found in Appendix D.

Table 1: Complexity tiers in *CausalDynamics*.

Tier	Model	Challenges	# Graphs
Simple	ODE/SDE	Confounder	585
Coupled	Coupled ODE/SDE ($N=\{3, 5, 10\}$)	Confounder, time-lag, standardized, forcing	14096
Climate	MAOOAM + ENSO models	High dimensionality	12

4.5 Evaluation metrics

Our benchmark is concerned with reconstructing the true causal graph of the underlying dynamical systems. To evaluate the similarity of the graphs inferred by the baselines and the true causal graph we estimate Area Under the Receiver Operating Characteristic (AUROC) [94, 95], Area Under the Precision Recall Curve (AUPRC) [96], and Structural Hamming Distance (SHD) [97, 98] scores by comparing the true and predicted adjacency matrices. Our full evaluation workflow including details on baseline methods is provided in Appendix E.

4.6 Limitations

We only processed data that is based on a fixed set of parameters (e.g., time-lag, noise level) and 3-dimensional systems. This is only a choice reflecting common challenges, as our plug-and-play framework provides options for users to scale the generated data to unrestricted complexity.

5 Experiments

To showcase how SOTA algorithms handle the various challenges of our benchmark, we evaluate the approaches described in Section 3.2, reporting a selection of the results in Tables 2-3 and the full results online at <https://huggingface.co/datasets/kausable/CausalDynamics>. The default setup for the simple case refers to ODE systems ($\delta = 0$) with no unobserved confounder. In the coupled systems tier, the default setup refers to dynamics with $n = 10$ coupled ODEs ($\delta = 0$), no confounder, no time-lag ($\tau = 0$), and no standardization. We performed an initial hyperparameter search for most of the baseline algorithms and provide the results in Section E.1.

We find methods that do not require much fine tuning, such as the non-DL-based algorithms, perform better as illustrated in Tables 2-3 and Figure 4, and observed by [20]. This is striking for the coupled




Figure 4: CDF of AUROC (top) and AUPRC (bottom) for mean values reported in Table 2 for different baselines across coupled system experiments. Models that perform better yield a low area under the curve for both AUROC and AUPRC.

Table 2: Baseline AUROC ($\uparrow$) / AUPRC ($\uparrow$) scores across different experiments in the hierarchy of increasingly complex dynamical systems. The scores are averaged over all generated graphs within a specific experiment. The experiments reported here consider one causal problem at a time, while keeping other factors at the default setting. The full results, including all possible experimental permutations, are found in our data repository.

	PCMCI+	F-PCMCI	VARLiNGAM	DYNOTEARS	NGC	TSCI	CUTS+	RCD	GRaSP	TCDF
Experiments	Simple									
Default	.52 / .71	.51 / .70	.50 / .69	.43 / .67	.50 / .69	.46 / .68	.50 / .69	.50 / .69	.52 / .71	.51 / .70
Noise	.50 / .69	.52 / .70	.53 / .70	.48 / .68	.50 / .68	.49 / .68	.50 / .68	.50 / .68	.49 / .68	.50 / .68
Confounder	.49 / .59	.50 / .59	.48 / .58	.52 / .64	.50 / .58	.53 / .66	.50 / .58	.50 / .58	.55 / .64	.50 / .58
Experiments	Coupled									
Default	.67 / .25	.67 / .27	.60 / .19	.59 / .21	.50 / .15	.60 / .23	.50 / .15	.50 / .15	.49 / .15	.48 / .15
Noise	.64 / .25	.57 / .21	.57 / .18	.57 / .20	.50 / .15	.53 / .17	.50 / .15	.50 / .15	.50 / .50	.50 / .50
Confounder	.58 / .20	.55 / .19	.51 / .17	.49 / .17	.50 / .16	.51 / .18	.49 / .16	.51 / .18	.50 / .17	.50 / .50
Time-lag	.58 / .24	.59 / .24	.54 / .22	.53 / .22	.50 / .20	.53 / .21	.50 / .20	.50 / .20	.50 / .50	.50 / .50
Standardize	.69 / .27	.68 / .28	.60 / .19	.65 / .23	.50 / .15	.65 / .23	.50 / .15	.50 / .16	.50 / .50	.50 / .50
Experiments	Climate									
MAOOAM	.69 / .88	.50 / .81	.50 / .81	.64 / .86	.50 / .81	.58 / .84	.50 / .81	.50 / .81	.48 / .81	.50 / .81
ENSO	.57 / .70	.57 / .70	.56 / .69	.55 / .69	.50 / .67	.50 / .67	.50 / .67	.50 / .67	.50 / .50	.50 / .50

atmosphere-ocean model, where each node represents a 2-dimensional spatial field. Methods that only consider 1-dimensional time-series, perform best, compared to DL-based methods, which claim to have a notion of space. We also find some evidence that topology-based methods, such as TSCI, perform better than purely DL-based approaches, for confounded and higher dimensional systems.

Overall, existing algorithms show shortcomings when confronted with coupled dynamical and physical systems as shown in Figure 5 for a random single graph realization. Across methods we find that autocorrelation is inferred where none exists. The baselines also appear conservative, predicting dense adjacency in the presence of nonstationary dynamics. In the climate case, methods perform well but do not recover the full coupling between ocean basins. Nevertheless, scores are the highest for the climate examples, possibly due to the noise in the generated data which might benefit some methods in recovering the underlying graph structure. Surprisingly in Figure 4, we observe minimal improvement in the standardized case, especially for methods like VARLiNGAM that exploit topological ordering. We performed additional ablation studies that consider the effect of time subsampling, partial observability, and varying edge-level activation function in Appendix F.

Finally, we note that for sparse graphs, AUROC can be deceptively high because the false-positive rate remains low due to a higher number of true-negatives. Thus, for coupled systems AUPRC scores might be more representative of the true performance due to the penalization of false positives. Other metrics like SHD [97, 98] or Structural Intervention Distance [99], suffer from different limitations

Table 3: Baseline SHD ($\downarrow$) score across different experiments in the hierarchy of increasingly complex dynamical systems. The scores are averaged over all generated graphs within a specific experiment. The experiments reported here consider one causal problem at a time, while keeping other factors at the default setting. The full results, including all possible experimental permutations, are found in our data repository.

	PCMCI+	F-PCMCI	VARLiNGAM	DYNOTEARS	NGC	TSCI	CUTS+	RCD	GRaSP	TCDF
Experiments	Simple									
Default	41.04	35.30	35.96	52.37	28.91	52.50	48.11	61.85	59.04	59.46
Noise	183.90	149.60	248.90	180.95	842.55	173.50	150.50	155.65	842.55	842.55
Confounder	23.02	21.07	22.04	21.74	19.96	22.02	24.04	26.74	27.09	26.61
Experiments	Coupled									
Default	224.80	192.90	311.45	181.50	840.95	174.70	152.00	157.05	215.94	252.00
Noise	183.90	149.60	248.90	180.95	842.55	173.50	150.50	155.65	842.55	842.55
Confounder	324.63	195.74	159.63	248.32	670.53	265.26	272.68	136.53	136.00	670.53
Time-lag	327.72	350.61	449.33	261.44	793.67	244.56	247.22	201.11	793.67	793.67
Standardize	228.32	201.79	349.63	243.84	840.26	244.42	310.63	159.84	840.26	840.26
Experiments	Climate									
MAOOAM	80.00	130.00	130.00	94.00	31.00	108.00	130.00	130.00	126.00	130.00
ENSO	529.36	530.27	453.00	589.36	337.09	666.27	608.73	665.36	666.27	666.27




Figure 5: Example of baseline performance for coupled systems ($n = 5$) for causal challenges and ENSO in the decoupled Atlantic setting (climate). Inference is performed using the best performing algorithm. Grey nodes and edges represent unobserved confounder, and dashed lines denote time-lagged relationships.

as they generally only return absolute error counts independent from the edge density or graph size, thus making a comparison across graph hierarchies difficult.

6 Conclusion

We present **CausalDynamics**, an extensible data generation framework that we use to construct the largest benchmark dataset with over 14000 preprocessed graphs of increasing complexity. We structure our benchmark as a tiered system ranging from simple three-dimensional dynamical systems to pseudo-real physical systems. We provide a plug-and-play workflow to facilitate the development of novel causal discovery methods across various domains. Evaluating a set of state-of-the-art causal discovery algorithms on **CausalDynamics** shows that many advanced DL-based algorithms are outperformed by simpler methods, notably on high-dimensional datasets, highlighting a need for future method development. We believe our work provides the necessary foundation for the advancement of causal discovery algorithms that are applicable in high-dimensional, nonlinear and dynamical settings.

Acknowledgment

We would like to thank Lars Lorch and Johannes Veith for discussions and their valuable insights. The authors acknowledge funding, computing, and storage resources from the NSF Science and Technology Center (STC) Learning the Earth with Artificial Intelligence and Physics (LEAP) (Award #2019625), and Department of Energy (DOE) Advanced Scientific Computing Research (ASCR) program (DE-SC0022255). JN also acknowledges funding from Columbia-Dream Sports AI Innovation Center. The project on which this report is based was funded by the Federal Ministry for Economic Affairs and Climate Action under the funding code 03EGTBW076. The responsibility for the content of this publication lies with the authors. Finally, the project was funded by the BW Pre-Seed program of the State of Baden-Württemberg, funded by the Ministry of Economics, Labour and Tourism and the L-Bank.

References

- [1] Sewall Wright. The method of path coefficients. *The annals of mathematical statistics*, 5(3):161–215, 1934.
- [2] Norbert Wiener. Modern mathematics for engineers. *New York: McGraw-Hill*, 1956.
- [3] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- [4] Charles K Assaad, Emilie Devijver, and Eric Gaussier. Survey and evaluation of causal discovery methods for time series. *Journal of Artificial Intelligence Research*, 73:767–819, 2022.
- [5] Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily B Fox. Neural granger causality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4267–4279, 2021.
- [6] Philippe Brouillard, Chandler Squires, Jonas Wahl, Konrad P Kording, Karen Sachs, Alexandre Drouin, and Dhanya Sridhar. The landscape of causal discovery data: Grounding causal discovery in real-world applications. *arXiv preprint arXiv:2412.01953*, 2024.
- [7] Adam Rupe, Derek DeSantis, Craig Bakker, Parvathi Kooloth, and Jian Lu. Causal discovery in nonlinear dynamical systems using koopman operators. *arXiv preprint arXiv:2410.10103*, 2024.
- [8] Juan Nathaniel, Carla Roesch, Jatan Buch, Derek DeSantis, Adam Rupe, Kara Lamb, and Pierre Gentine. Deep koopman operator framework for causal discovery in nonlinear dynamical systems. *arXiv preprint arXiv:2505.14828*, 2025.
- [9] Carla M Roesch, Emilie Fons, Andrew P Ballinger, Jakob Runge, and Gabriele C Hegerl. Decreasing aerosols increase the european summer diurnal temperature range. *npj Climate and Atmospheric Science*, 8(1):47, 2025.
- [10] Emilie Fons, Jakob Runge, David Neubauer, and Ulrike Lohmann. Stratocumulus adjustments to aerosol perturbations disentangled with a causal approach. *npj Climate and Atmospheric Science*, 6(1):130, 2023.
- [11] Juan Nathaniel, Sara Schwetschenau, and Upmanu Lall. Inferring failure risk of on-site wastewater systems from physical and social factors. *npj Clean Water*, 7(1):126, 2024.
- [12] Andrew Jesson, Peter Manshausen, Alyson Douglas, Duncan Watson-Parris, Yarin Gal, and Philip Stier. Using Non-Linear Causal Models to Study Aerosol-Cloud Interactions in the Southeast Pacific. *arXiv preprint arXiv:2110.15084*, 2021.
- [13] Mathew Williams, David T Milodowski, T Luke Smallman, Kyle G Dexter, Gabi C Hegerl, Iain M McNicol, Michael O’Sullivan, Carla M Roesch, Casey M Ryan, Stephen Sitch, et al. Precipitation–fire functional interactions control biomass stocks and carbon exchanges across the world’s largest savanna. *Biogeosciences*, 22(6):1597–1614, 2025.
- [14] Sebastian Hickman, Ilija Trajkovic, Julia Kaltenborn, Francis Pelletier, Alex Archibald, Yaniv Gurwicz, Peer Nowack, David Rolnick, and Julien Boussard. Causal climate emulation with bayesian filtering. *arXiv preprint arXiv:2506.09891*, 2025.
- [15] Konstantin Göbler, Tobias Windisch, Mathias Drton, Tim Pychynski, Martin Roth, and Steffen Sonntag. causalAssembly : Generating Realistic Production Data for Benchmarking Causal Discovery. In *Conference on Causal Learning and Reasoning*, pages 609–642. PMLR, 2024.
- [16] Yuxiao Cheng, Ziqian Wang, Tingxiong Xiao, Qin Zhong, Jinli Suo, and Kunlun He. Causaltime: Realistically generated time-series for benchmarking of causal discovery. In *The Twelfth International Conference on Learning Representations*, 2024.
- [17] Xavier-Andoni Tibau, Christian Reimers, Andreas Gerhardus, Joachim Denzler, Veronika Eyring, and Jakob Runge. A spatiotemporal stochastic climate model for benchmarking causal discovery methods for teleconnections. *Environmental Data Science*, 1:e12, 2022.
- [18] Mathieu Chevalley, Jacob Sackett-Sanders, Yusuf Roohani, Pascal Notin, Artemy Bakulin, Dariusz Brzezinski, Kaiwen Deng, Yuanfang Guan, Justin Hong, Michael Ibrahim, et al. The CausalBench challenge: A machine learning contest for gene network inference from single-cell perturbation data. *arXiv preprint arXiv:2308.15395*, 2023.

- [19] Mathieu Chevalley, Yusuf H Roohani, Arash Mehrjou, Jure Leskovec, and Patrick Schwab. A large-scale benchmark for network inference from single-cell perturbation data. *Communications Biology*, 8(1):412, 2025.
- [20] Gideon Stein, Maha Shadaydeh, Jan Blunk, Niklas Penzel, and Joachim Denzler. CausalRivers—Scaling up benchmarking of causal discovery for real-world time-series. *arXiv preprint arXiv:2503.17452*, 2025.
- [21] Jakob Runge, Sebastian Bathiany, Erik Bollt, Gustau Camps-Valls, Dim Coumou, Ethan Deyle, Clark Glymour, Marlene Kretschmer, Miguel D Mahecha, Jordi Muñoz-Marí, et al. Inferring causation from time series in Earth system sciences. *Nature Communications*, 10(1):2553, 2019.
- [22] Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. DYNOTEARS: Structure Learning from Time-Series Data. In *International Conference on Artificial Intelligence and Statistics*, pages 1595–1605, 2020.
- [23] Weronika Ormaniec, Scott Sussex, Lars Lorch, Bernhard Schölkopf, and Andreas Krause. Standardizing structural causal models. *arXiv preprint arXiv:2406.11601*, 2024.
- [24] Rebecca J. Herman, Jonas Wahl, Urmi Ninad, and Jakob Runge. Unitless Unrestricted Markov-Consistent SCM Generation: Better Benchmark Datasets for Causal Discovery. *arXiv preprint arXiv:2503.17037*, 2025.
- [25] Valeria Aloisi, Andrea Gatto, Gabriele Accarino, Francesco Donato, and Giovanni Aloisio. The effect of known and unknown confounders on the relationship between air pollution and Covid-19 mortality in Italy: A sensitivity analysis of an ecological study based on the E-value. *Environmental research*, 207:112131, 2022.
- [26] Juan Nathaniel, Jiangong Liu, and Pierre Gentine. Metaflux: Meta-learning global carbon fluxes from sparse spatiotemporal observations. *Scientific Data*, 10(1):440, 2023.
- [27] Siyeon Kim, Juan Nathaniel, Zhewen Hou, Tian Zheng, and Pierre Gentine. Spatiotemporal upscaling of sparse air-sea pco2 data via physics-informed transfer learning. *Scientific data*, 11(1):1098, 2024.
- [28] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.
- [29] Mohammed AlQuraishi. AlphaFold at CASP13. *Bioinformatics*, 35(22):4862–4865, 2019.
- [30] Jenny Straiton. The path to solving the protein folding problem. *Taylor & Francis*, 2023.
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [32] Md Zahangir Alom, Tarek M Taha, Christopher Yakopcic, Stefan Westberg, Paheding Sidike, Mst Shamima Nasrin, Brian C Van Esesn, Abdul A S Awwal, and Vijayan K Asari. The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches. *arXiv preprint arXiv:1803.01164*, 2018.
- [33] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, October 2017.
- [34] Paul L Krapivsky and Sidney Redner. Organization of growing random networks. *Physical Review E*, 63(6):066123, 2001.
- [35] Jonas Peters, Stefan Bauer, and Niklas Pfister. Causal models for dynamical systems. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 671–690, 2022.
- [36] Stephan Bongers, Tineke Blom, and Joris M Mooij. Causal modeling of dynamical systems. *arXiv preprint arXiv:1803.08784*, 2018.
- [37] Stephan Bongers and Joris M. Mooij. From random differential equations to structural causal models: the stochastic case. *arXiv preprint arXiv:1304.7920*, 2018.

- [38] Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. From ordinary differential equations to structural causal models: the deterministic case. *arXiv preprint arXiv:1304.7920*, 2013.
- [39] Philip Boeken and Joris M Mooij. Dynamic structural causal models. *arXiv preprint arXiv:2406.01161*, 2024.
- [40] Paul K. Rubenstein, Stephan Bongers, Bernhard Schoelkopf, and Joris M. Mooij. From Deterministic ODEs to Dynamic Structural Causal Models. *arXiv preprint arXiv:1608.08028*, 2018.
- [41] P Spirtes, C Glymour, and R Scheines. *Causation, prediction, and search*. The MIT Press, 2000.
- [42] Christian Reiser. Causal discovery for time series with latent confounders. *arXiv preprint arXiv:2209.03427*, 2022.
- [43] Andreas Gerhardus and Jakob Runge. High-recall causal discovery for autocorrelated time series with latent confounders. *Advances in neural information processing systems*, 33:12615–12625, 2020.
- [44] Jakob Runge. Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In *Conference on uncertainty in artificial intelligence*, pages 1388–1397. Pmlr, 2020.
- [45] Alexander Reisach, Christof Seiler, and Sebastian Weichwald. Beware of the simulated DAG! causal discovery benchmarks may be easy to game. *Advances in Neural Information Processing Systems*, 34:27772–27784, 2021.
- [46] Alexander Reisach, Myriam Tami, Christof Seiler, Antoine Chambaz, and Sebastian Weichwald. A scale-invariant sorting criterion to find a causal order in additive noise models. *Advances in Neural Information Processing Systems*, 36:785–807, 2023.
- [47] Jakob Runge, Peer Nowack, Marlene Kretschmer, Seth Flaxman, and Dino Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science advances*, 5(11):eaau4996, 2019.
- [48] George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting Causality in Complex Ecosystems. *Science*, 338(6106):496–500, October 2012.
- [49] Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21, 2008.
- [50] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.
- [51] Karen Sachs, Omar Perez, Dana Pe’er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.
- [52] Lukas Riedel, Benjamin Herdeanu, Harald Mack, Yunus Sevinchan, and Julian Weninger. Utopia: a comprehensive and collaborative modeling framework for complex and evolving systems. *Journal of Open Source Software*, 5(53):2165, 2020.
- [53] Benjamin Herdeanu. *Emergence of Cooperation in Evolutionary Social Interaction Networks*. Doctoral dissertation, Heidelberg University, 2021.
- [54] Yunus Sevinchan, Benjamin Herdeanu, and Jeremias Traub. dantro: a python package for handling, transforming, and visualizing hierarchically structured data. *Journal of Open Source Software*, 5(52):2316, 2020.
- [55] Yunus Sevinchan, Benjamin Herdeanu, Harald Mack, Lukas Riedel, and Kurt Roth. Boosting group-level synergies by using a shared modeling framework. In Valeria V. Krzhizhanovskaya, Gábor Závodszy, Michael H. Lees, Jack J. Dongarra, Peter M. A. Sloot, Sérgio Brissos, and João Teixeira, editors, *Computational Science – ICCS 2020*, pages 442–456, Cham, 2020. Springer International Publishing.

- [56] Joris M Mooij, Dominik Janzing, Tom Heskes, and Bernhard Schölkopf. On causal discovery with cyclic additive noise models. *Advances in neural information processing systems*, 24, 2011.
- [57] Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social science computer review*, 9(1):62–72, 1991.
- [58] Edward De Brouwer, Adam Arany, Jaak Simm, and Yves Moreau. Latent convergent cross mapping. In *International Conference on Learning Representations*, 2020.
- [59] Tim Van den Bulcke, Koenraad Van Leemput, Bart Naudts, Piet van Remortel, Hongwu Ma, Alain Verschoren, Bart De Moor, and Kathleen Marchal. Syntren: a generator of synthetic gene expression data for design and analysis of structure learning algorithms. *BMC bioinformatics*, 7:1–12, 2006.
- [60] William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling. *arXiv preprint arXiv:2110.05266*, 2021.
- [61] Yuxiao Cheng, Runzhao Yang, Tingxiong Xiao, Zongren Li, Jinli Suo, Kunlun He, and Qionghai Dai. Cuts: Neural causal discovery from irregular time-series data. *arXiv preprint arXiv:2302.07458*, 2023.
- [62] Saurabh Khanna and Vincent YF Tan. Economy statistical recurrent units for inferring nonlinear granger causality. *arXiv preprint arXiv:1911.09879*, 2019.
- [63] Alexis Bellot, Kim Branson, and Mihaela van der Schaar. Neural graphical modelling in continuous-time: consistency guarantees and algorithms. *arXiv preprint arXiv:2105.02522*, 2021.
- [64] Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, 2011.
- [65] Robert J Prill, Daniel Marbach, Julio Saez-Rodriguez, Peter K Sorger, Leonidas G Alexopoulos, Xiaowei Xue, Neil D Clarke, Gregoire Altan-Bonnet, and Gustavo Stolovitzky. Towards a rigorous assessment of systems biology models: the DREAM3 challenges. *Public Library of Science*, 5(2):e9202, 2010.
- [66] Robert J. Prill, Julio Saez-Rodriguez, Leonidas G. Alexopoulos, Peter K. Sorger, and Gustavo Stolovitzky. Crowdsourcing network inference: The dream predictive signaling network challenge. *Science Signaling*, 4(189):mr7–mr7, 2011.
- [67] Eugene F Fama and Kenneth R French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992.
- [68] Meike Nauta, Doina Bucur, and Christin Seifert. Causal discovery with attention-based convolutional neural networks. *Machine Learning and Knowledge Extraction*, 1(1):19, 2019.
- [69] Jakob Runge, Xavier-Andoni Tibau, Matthias Bruhns, Jordi Muñoz-Marí, and Gustau Camps-Valls. The causality for climate competition. In *NeurIPS 2019 Competition and Demonstration Track*, pages 110–120. Pmlr, 2020.
- [70] Andrew R Lawrence, Marcus Kaiser, Rui Sampaio, and Maksim Sipos. Data generating process to evaluate causal discovery techniques for time series data. *arXiv preprint arXiv:2104.08043*, 2021.
- [71] David Hume. *A treatise of human nature*. Oxford University Press, 1738.
- [72] Yuxiao Cheng, Lianglong Li, Tingxiong Xiao, Zongren Li, Jinli Suo, Kunlun He, and Qionghai Dai. Cuts+: High-dimensional causal discovery from irregular time-series. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(10):11525–11533, 2024.
- [73] Luca Castri, Sariah Mghames, Marc Hanheide, and Nicola Bellotto. Enhancing causal discovery from robot sensor data in dynamic scenarios. In *Conference on Causal Learning and Reasoning (CLEaR)*, 2023.
- [74] Wai-Yin Lam, Bryan Andrews, and Joseph Ramsey. Greedy relaxations of the sparsest permutation algorithm. In *Uncertainty in Artificial Intelligence*, pages 1052–1062. PMLR, 2022.

- [75] Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- [76] Aapo Hyvärinen, Kun Zhang, Shohei Shimizu, and Patrik O Hoyer. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research*, 11(5), 2010.
- [77] Takashi Nicholas Maeda and Shohei Shimizu. Rcd: Repetitive causal discovery of linear non-gaussian acyclic models with latent confounders. In *International Conference on artificial intelligence and statistics*, pages 735–745. PMLR, 2020.
- [78] Floris Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980: proceedings of a symposium held at the University of Warwick 1979/80*, pages 366–381. Springer, 1981.
- [79] Kurt Butler, Daniel Waxman, and Petar Djuric. Tangent space causal inference: Leveraging vector fields for causal discovery in dynamical systems. *Advances in Neural Information Processing Systems*, 37:120078–120102, 2024.
- [80] William Gilpin. Model scale versus domain knowledge in statistical forecasting of chaotic systems. *arXiv preprint arXiv:2303.08011*, 2023.
- [81] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [82] Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeyer, and Frank Hutter. Accurate predictions on small data with a tabular foundation model. *Nature*, 637(8045):319–326, 2025.
- [83] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A transformer that solves small tabular classification problems in a second. *arXiv preprint arXiv:2207.01848*, 2022.
- [84] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [85] Michael Ghil and Valerio Lucarini. The physics of climate variability and climate change. *Rev. Mod. Phys.*, 92:035002, Jul 2020.
- [86] Juan Nathaniel, Yongquan Qu, Tung Nguyen, Sungduk Yu, Julius Busecke, Aditya Grover, and Pierre Gentile. Chaosbench: A multi-channel, physics-based benchmark for subseasonal-to-seasonal climate prediction. *arXiv preprint arXiv:2402.00712*, 2024.
- [87] Dmitrii Kochkov, Janni Yuval, Ian Langmore, Peter Norgaard, Jamie Smith, Griffin Mooers, Milan Klöwer, James Lottes, Stephan Rasp, Peter Düben, Sam Hatfield, Peter Battaglia, Alvaro Sanchez-Gonzalez, Matthew Willson, Michael P. Brenner, and Stephan Hoyer. Neural general circulation models for weather and climate. *Nature*, 632(8027):1060–1066, August 2024.
- [88] Juan Nathaniel and Pierre Gentile. Generative emulation of chaotic dynamics with coherent prior. *Computer Methods in Applied Mechanics and Engineering*, 448:118410, 2026.
- [89] Sen Zhao. Extended nonlinear Recharge Oscillator (XRO) model for "Explainable El Niño predictability from climate mode interactions". *Zenodo*, 2024.
- [90] Klaus Hasselmann. Stochastic climate models part i. theory. *tellus*, 28(6):473–485, 1976.
- [91] Fei-Fei Jin. An equatorial ocean recharge paradigm for ENSO. Part I: Conceptual model. *Journal of the atmospheric sciences*, 54(7):811–829, 1997.
- [92] Sen Zhao, Fei-Fei Jin, Malte F Stuecker, Philip R Thompson, Jong-Seong Kug, Michael J McPhaden, Mark A Cane, Andrew T Wittenberg, and Wenju Cai. Explainable El Niño predictability from climate mode interactions. *Nature*, 630(8018):891–898, 2024.
- [93] Jonathan Demayer, Lesley De Cruz, and Stéphane Vannitsem. qgs: A flexible python framework of reduced-order multiscale climate models. *Journal of Open Source Software*, 5(56):2597, 2020.
- [94] WWTG Peterson, T Birdsall, and We Fox. The theory of signal detectability. *Transactions of the IRE professional group on information theory*, 4(4):171–212, 1954.
- [95] Andrew P Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.

- [96] Jesse Davis and Mark Goadrich. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240, 2006.
- [97] David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- [98] Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65:31–78, 2006.
- [99] Jonas Peters and Peter Bühlmann. Structural intervention distance (SID) for evaluating causal graphs. *arXiv preprint arXiv:1306.1043*, 2013.
- [100] R. K. Rew, G. P. Davis, S. Emmerson, and H. Davies. *NetCDF User’s Guide for C, An Interface for Data Access, Version 3*, April 1997.
- [101] S. Hoyer and J. Hamman. xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, 5(1), 2017.
- [102] Hao Zuo, Magdalena Alonso Balmaseda, Steffen Tietsche, Kristian Mogensen, and Michael Mayer. The ECMWF operational ensemble reanalysis–analysis system for ocean and sea ice: a description of the system and assessment. *Ocean science*, 15(3):779–808, 2019.
- [103] Soon-Il An, Hyo-Jin Park, Soong-Ki Kim, Wenju Cai, Agus Santoso, Daehyun Kim, and Jong-Seong Kug. Main drivers of Indian Ocean Dipole asymmetry revealed by a simple IOD model. *npj Climate and Atmospheric Science*, 6(1):93, 2023.
- [104] Soon-Il An and Fei-Fei Jin. Nonlinearity and asymmetry of ENSO. *Journal of Climate*, 17(12):2399–2412, 2004.
- [105] Soon-Il An, Eli Tziperman, Yuko M Okumura, and Tim Li. ENSO irregularity and asymmetry. *El Niño Southern Oscillation in a changing climate*, pages 153–172, 2020.
- [106] In-Sik Kang and Jong-Seong Kug. El Niño and La Niña sea surface temperature anomalies: Asymmetry characteristics associated with their wind stress anomalies. *Journal of Geophysical Research: Atmospheres*, 107(D19):ACL–1, 2002.
- [107] Tao Geng, Wenju Cai, and Lixin Wu. Two types of ENSO varying in tandem facilitated by nonlinear atmospheric convection. *Geophysical Research Letters*, 47(15):e2020GL088784, 2020.
- [108] Kit-Yan Choi, Gabriel A Vecchi, and Andrew T Wittenberg. ENSO transition, duration, and amplitude asymmetries: Role of the nonlinear wind stress coupling in a conceptual model. *Journal of Climate*, 26(23):9462–9476, 2013.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We clearly outline our contribution to the paper and scope of the work in the abstract and introduction. We further separate our work from related work in Section 3.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We provide a description of what is included in our benchmark and what baseline methods we evaluate. We further discuss limitations of the benchmark in Section 4.6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include any theoretical results. Where necessary published theoretical background has been included in the paper and properly referenced (see Section 2 and 4.2).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All code required to reproduce the paper and its results is publicly available at <https://kausable.github.io/CausalDynamics/README.html>. We give a detailed description of the methods and data generation process including pseudo-code of the algorithms and instructive code examples in Section 4 and in Appendix B, C, D, and E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our code is available at <https://github.com/kausable/CausalDynamics/tree/main> and the dataset can be downloaded from <https://huggingface.co/datasets/kausable/CausalDynamics> as mentioned in the abstract and Section 4. We further provide a description on how to use the code and access the dataset in Appendix A and on <https://kausable.github.io/CausalDynamics/notebooks/quickstart.html>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Yes hyperparameters for method reproducibility are provided in Appendix D, E, and the code is made available at <https://github.com/kausable/CausalDynamics/tree/main/src/causaldynamics/baselines>.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We include a critical discussion of the experimental results in Section 5 and visualize the results further in Figure 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide computing resources in Section E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research conforms with the NeurIPS Code of Ethics, as we do not present any research involving human subjects or participants, follow norms of fair use and copyright for any algorithms or datasets we use, and present a benchmark with no risk of potential harmful consequences or societal impact, such as privacy, surveillance, discrimination, human rights, etc.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.

- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: As we provide a benchmark for foundational research into causal discovery methods, we do not see any opportunity for potential negative societal impacts such as disinformation, deepfakes, surveillance, unfair impacts of specific groups, privacy considerations, or security considerations.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not provide any data or models that are at high risk of misuse such as LLMs, image generators or from the internet scraped datasets. Our data is generated synthetically for idealized and purely physical systems. Methods we include have been previously published.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Where necessary we cite the original authors of adapted open-source algorithms and refer to the respective open-sourced GitHub repositories for benchmarks. Specifically, we acknowledge [80, 60] in Section 4.1, [34] in Section 4.2 and [89, 92, 93] in Section 4.3. We further refer to the original GitHub repositories and references for the evaluated baselines in Section 3.2 and on <https://kausable.github.io/CausalDynamics/baseline.html>.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: A detailed description of the benchmark is provided in Section 4. More details on the methods, data generation process and a general summary are provided in Appendix B, C, and D, respectively. The associated open-srouce code is documented and available under <https://github.com/kausable/CausalDynamics>. We provide additional documentation on a project website available at <https://kausable.github.io/CausalDynamics/README.html>.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not include any crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not include any crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We do not use any LLM for the core methodology, scientific rigorousness, or originality of the research

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

CausalDynamics: A large-scale benchmark for structural discovery of dynamical causal models

Supplementary Material

Benjamin Herdeanu^{1,*}, Juan Nathaniel^{2,*}, Carla Roesch^{2,*},
Jatan Buch², Gregor Ramien¹, Johannes Haux¹, Pierre Gentine²

¹kausable GmbH, ²Columbia University

Table of Contents

A	Getting started	1
A.1	Installation	2
A.2	Download data	2
A.3	Generate data	2
A.4	Plotting functions	5
A.5	Baseline evaluation	6
B	Generate coupled causal models	8
B.1	System initialization	8
B.2	MLP propagation	11
B.3	Standardization	12
C	First-order climate models	13
C.1	XRO: eXtended nonlinear Recharge Oscillator model	13
C.2	qgs: A flexible Python framework of reduced-order multiscale climate models . .	15
D	Benchmark details	17
D.1	Summary	17
D.2	Hierarchy of Differential Equations	17
E	Experiment details	18
E.1	Baselines	18
E.2	Metrics	19
F	Additional results	21
F.1	Time discretisation ablation	21
F.2	Partial observability ablation	21
F.3	Edge-level activation ablation	21

A Getting started

We provide an overview, including code examples, on how to use the **CausalDynamics** Python package. The open-sourced code is available at <https://github.com/kausable/CausalDynamics/> and the latest documentation is published at <https://kausable.github.io/CausalDynamics/README.html>.

All results, code examples and descriptions rely on version 1.0.0 of the **CausalDynamics** Python package.

*Corresponding authors: benjamin@kausable.ai, jn2808@columbia.edu, cmr2293@columbia.edu

A.1 Installation

The easiest way to install the Python package is via PyPi, see <https://pypi.org/project/causaldynamics/>, which currently requires Python=3.10.

```
$ pip install causaldynamics
```

It is also possible to install the package locally. Further installation instructions are available in the project repository at <https://kausable.github.io/CausalDynamics/README.html>.

A.2 Download data

We provide a pre-generated dataset at <https://huggingface.co/datasets/kausable/CausalDynamics> that can be directly downloaded using the following commands:

```
$ wget https://huggingface.co/datasets/kausable/CausalDynamics/resolve/main/process_causaldynamics.py
$ python process_causaldynamics.py
```

The dataset was generated using the scripts published in the repository under: <https://github.com/kausable/CausalDynamics/tree/main/scripts>.

A.3 Generate data

Simple causal models. Here, we showcase how to generate data for the simple complexity tier that consists of individual dynamical systems. As an example, we choose the Lorenz system and show how to (i) get the adjacency matrix of the system, (ii) solve the system for 1000 time steps resulting in time series trajectory data, and (iii) store the data.

```
import xarray as xr
from causaldynamics.systems import get_adjacency_matrix_from_jac
from causaldynamics.systems import solve_system

# Define the system name, ...
system_name = "Lorenz"

# ... get the adjacency matrix of the system,...
A = get_adjacency_matrix_from_jac(sys_name)

# ... integrate the Lorenz ODE,...
data = solve_system(
    num_timesteps=1000,
    num_systems=1,
    system_name=system_name
)
data = xr.DataArray(data, dims=['time', 'node', 'dim'])

# ... and store the generated data
ds = create_dynsys_dataset(adjacency_matrix=A, time_series=data)
save_xr_dataset(ds, "out_path.nc")
```

We introduce function `solve_system`, a lightweight wrapper around the `dysts` package [60, 80], including ODE/SDE integration schemes that return the trajectories of the integrated system. The function `get_adjacency_matrix_from_jac` function then takes the Jacobian computed by `dysts` and extracts its adjacency matrix. For a visual walkthrough, see our explanatory notebook: https://github.com/kausable/CausalDynamics/blob/main/notebooks/adj_mat_

from_dysts.ipynb. Within our framework, a single dynamical system can correspond to a single root node, hence one of the dataset's dimension is referred to as node. Multiple systems can be simulated in parallel via the num_systems argument in solve_system.

Finally, we combine the generated time series (data) and its corresponding adjacency matrix A in a single dataset to save as a NetCDF file [100] using the Xarray package [101]. We followed the same procedure for the preprocessed benchmark data described in Appendix A.2.

More information and a detailed description of the simple causal model data generation can be found at https://kausable.github.io/CausalDynamics/notebooks/simple_causal_models.html.

Coupled causal models. We showcase example code of how to (i) create an SCM, (ii) simulate the system to generate the time series data, and (iii) plot the data. To create the SCM, we use the create_scm function which returns the corresponding adjacency matrix A , the weights W and biases b of all MLPs and the root_nodes that act as temporal system drivers. We then simulate the system consisting of num_nodes nodes for num_timesteps driven by the dynamical systems system_name located on the root nodes. We provide a simple example for basic functionality and an advanced example that shows the modular plug-and-play feature configurability.

```
# Simple example: Generate data of coupled causal models
from causaldynamics.creator import create_scm, simulate_system,
    create_plots

# Define system parameters
num_nodes = 2
node_dim = 3
num_timesteps = 1000
system_name='Lorenz'
confounders = False

# Create a coupled causal model,...
A, W, b, root_nodes, _ = create_scm(num_nodes, node_dim, confounders=
    confounders)

# ... simulate the system,...
data = simulate_system(A, W, b,
    num_timesteps=num_timesteps,
    num_nodes=num_nodes,
    system_name=system_name)

# ... and plot the results.
create_plots(
    data,
    A,
    root_nodes=root_nodes,
    out_dir='',
    show_plot=True,
    save_plot=False,
    create_animation=False,
)
```

The create_scm and simulate_system function interface can be used to generate more complex graphs and causal challenges providing the option to introduce confounders, time-lag, noise, periodic functions as root nodes, and to standardize (i.e., correct for varsortability), which are visualized at <https://kausable.github.io/CausalDynamics/notebooks/causaldynamics.html>.

A more detailed introduction to coupled causal models with visualizations can also be found at https://kausable.github.io/CausalDynamics/notebooks/coupled_causal_models.html.

```

# Advanced example: Generate more complex systems
from matplotlib import pyplot as plt
from causaldynamics.creator import create_scm, simulate_system
from causaldynamics.scm import create_scm_graph
from causaldynamics.plot import (
    plot_scm,
    plot_trajectories,
    plot_3d_trajectories
)
from causaldynamics.data_io import create_output_dataset,
    save_xr_dataset

# Define system parameters
num_nodes = 5
node_dim = 3
num_timesteps = 1000

confounders = False      # Set to True to add confounders
standardize = False      # Set to True to standardize the data
init_ratios = [1, 1, 1]  # Set ratios of dynamical systems,
                        # periodic, and linear drivers at root nodes.
                        # Here: equal ratio.
system_name='random'     # Sample random dyn. sys. for root nodes
activations_names = ['identity', 'sin', 'sigmoid', 'tanh', 'relu']
                        # Activation names sampled uniformly per edge
noise = 0.5              # Set noise for the dynamical systems

time_lag = 10            # Set time lag for time-lagged edges
time_lag_edge_probability = 0.1 # Set probability of time-lagged edges

# Create a coupled causal model,...
A, W, b, root_nodes, _ = create_scm(num_nodes,
                                     node_dim,
                                     confounders=confounders,
                                     time_lag=time_lag,
                                     time_lag_edge_probability=
                                         time_lag_edge_probability)

# ... simulate the system, ...
data = simulate_system(A, W, b,
                      num_timesteps=num_timesteps,
                      num_nodes=num_nodes,
                      system_name=system_name,
                      init_ratios=init_ratios,
                      time_lag=time_lag,
                      standardize=standardize,
                      activations_names=activations_names,
                      make_trajectory_kwargs={'noise': noise})

# ... visualize the results,...
plot_scm(G=create_scm_graph(A), root_nodes=root_nodes)
plot_3d_trajectories(data, root_nodes)
plot_trajectories(data, root_nodes=root_nodes, sharey=False)
plt.show()

# ... and save the data.
dataset = create_output_dataset(
    adjacency_matrix=A,
    weights=W,
    biases=b,
    time_lag=time_lag,
    time_series=data,
    root_nodes=root_nodes,
    verbose=False,
)
save_xr_dataset(dataset, "out_path.nc")

```

A.4 Plotting functions

Plotting functions have been implemented to visualize the generated graph structures and system trajectories given the adjacency matrix $\mathcal{A}$, the time series data and optionally the root_nodes (see Figure 6).

```
# Plotting functions
from causaldynamics.scm import create_scm_graph
from causaldynamics.plot import (plot_scm, plot_trajectories,
    plot_3d_trajectories, animate_3d_trajectories)


# Plot SCM, root nodes are plotted in grey, dashed edges are time-
# lagged, dash-dotted edges as both, regular and time-lagged
plot_scm(G=create_scm_graph(A), root_nodes=root_nodes)

# Plot time-series at nodes, root nodes are plotted in grey
plot_trajectories(data, root_nodes=root_nodes, sharey=False)


# 3D plot of the systems at each node
plot_3d_trajectories(data, root_nodes, line_alpha=1.)

# 3D animation of the time-dependent trajectory buildup with rotating
# axes.
anim = animate_3d_trajectories(data,
    root_nodes=root_nodes,
    plot_type="subplots")



display(anim)
```



(a) Plot SCM for a graph consisting of 2 nodes (autocorrelated root node in grey).



(b) Plotted trajectories for the SCM from (a).



(c) 3D trajectories for the SCM from (a) and the time series from (b). These trajectories can also be animated to visualize their temporal development with rotating axes to intuitively grasp the 3D structure. The last frame of this animation is the visualized figure.

Figure 6: Plotting function in CausalDynamics at the example of a SCM consisting of 2 nodes, with the root node (grey) driven by a 3-dimensional dynamical system.

A.5 Baseline evaluation

The following example demonstrates how to evaluate a single dynamical system. First, we load the dataset along with its generating causal graph (ground truth), represented as adjacency matrix.

```
import xarray as xr
import copy
import numpy as np
from tqdm import tqdm
from causaldynamics.baselines import CUTSPlus

# Load dataset
ds = xr.open_dataset(DATA_DIR / "<SYSTEM_NAME>.nc")

# Extract timeseries and adjacency matrix as target
timeseries = ds['time_series'].to_numpy().transpose(1, 0, 2)
adj_matrix = ds['adjacency_matrix'].to_numpy()
```

For coupled systems, slight changes to data loading are required since there are several new adjacency matrices, such as the `adjacency_matrix_time_edges` for the lagged connection.

```
# Select the first dimension of each multidimensional node
timeseries = ds['time_series'].to_numpy()[..., 0].transpose(1, 0, 2)

# 'adjacency_matrix_summary' is the true adjacency matrix here
adj_matrix = ds['adjacency_matrix_summary'].to_numpy()
```

Finally, using CUTS+ as an example, we provide sample evaluation script and visualize the predicted SCM in Figure 7.

```
# Define parameters
tau_max = 1
corr_thres = 0.8

# Estimate adjacency matrix
cuts_adj_matrix = []
for x in tqdm(timeseries):

    # Initialize model
    cuts_model = CUTSPlus(tau_max=tau_max, corr_thres=corr_thres)

    # Fit the model
    cuts_model.run(X=x)

    # Extract inferred matrix
    cuts_adj_matrix.append(
        copy.deepcopy(cuts_model.adj_matrix)
    )

# Compute scores
score(preds= cuts_adj_matrix, labs= adj_matrix, name='CUTS+')

# Plot recovered SCM
G = create_scm_graph(cuts_model.adj_matrix)
plot_scm(G);
```




Figure 7: Predicted graph for a single dynamical system using CUTS+.

The complete baseline pipeline is available at https://kausable.github.io/CausalDynamics/notebooks/eval_pipeline.html.

We also provide the complete metrics for each graph in our HuggingFace data repository <https://huggingface.co/datasets/kausable/CausalDynamics>, but analyzing each of the 14k+ graphs can be too granular. As a result, we have added a built-in, easy-to-use diagnosis script that allows users to quickly analyze results on the experiment-level. For instance, if users are interested in understanding the effects of unobserved confounders (`confounders=True`) on the performance of different causal discovery algorithms, especially in light of high internal noise/stochasticity (`noise=2.0`), they can run the following:

```
# Diagnostic tool to evaluate each experiment
python diagnose.py --exp_dir data/simple/noise=2.00_confounder=True
```

B Generate coupled causal models

Here, we elaborate on the propagation of information via MLPs in the DAG and provide the key algorithms used to generate coupled causal models.

B.1 System initialization

To generate data for coupled causal models we need to initialize the number of time steps, the number of nodes (e.g., 5), the ratio of dynamical systems to periodic functions as root node drivers, the system name of dynamical system drivers, the dimensionality of the nodes (e.g., 3), and optional time-lags, as outlined in Algorithm 2 and sketched out in Figure 8.

If the dynamical system drivers are selected at random, the algorithm chooses randomly from the set of available 3D systems (see Algorithm 3). Users can also choose whether to include periodic or linear functions as drivers by initializing r appropriately. In this case, sinusoidal or linear trajectories are generated as outlined in Algorithm 4 and Algorithm 5. The generated data of the system's drivers is randomly assigned to the respective root nodes, and all non-root nodes initial states are set to zero (see Algorithm 6).

Finally, the initial time series data can be standardized in order to prevent varsortability as illustrated in Appendix B.3.




Figure 8: Initialize graph.

Algorithm 2 Initialize Coupled Causal Models

Require: T (num_timesteps), N (num_nodes), r (init_ratios), name (system_name), D (node_dim), l (time_lag)

Ensure: $\text{init} \in \mathbb{R}^{T \times N \times D}$ (initial values tensor)

```

1: if  $l > 0$  then
2:    $T_{\text{ext}} \leftarrow T + l$  ▷ Extended time for lag calculation
3: else
4:    $T_{\text{ext}} \leftarrow T$ 
5: end if
6:  $n_{\text{sys}}, n_{\text{sin}}, n_{\text{lin}} \leftarrow \text{allocate\_elements\_based\_on\_ratios}(N, r)$  ▷ Divide nodes by type
7: if system_name = "random" then
8:    $d_{\text{sys}} \leftarrow \text{solve\_random\_systems}(T_{\text{ext}}, n_{\text{sys}})$  ▷ Generate random system trajectories
9: else
10:   $d_{\text{sys}} \leftarrow \text{solve\_system}(T_{\text{ext}}, n_{\text{sys}}, \text{system\_name})$  ▷ Generate named system trajectories
11: end if
12:  $d_{\text{lin}} \leftarrow \text{drive\_lin}(T_{\text{ext}}, n_{\text{lin}}, D)$  ▷ Generate linear trajectories
13:  $d_{\text{sin}} \leftarrow \text{drive\_sin}(T_{\text{ext}}, n_{\text{sin}}, D)$  ▷ Generate sinusoidal trajectories
14:  $\text{init} \leftarrow \text{concat}(d_{\text{sys}}, d_{\text{sin}}, d_{\text{lin}}, \text{dim} = 1)$  ▷ Combine trajectories
15:  $\text{init} \leftarrow \text{init}[:, \text{randperm}(N), :]$  ▷ Randomly permute nodes
16: if  $l > 0$  then
17:    $\text{init\_now} \leftarrow \text{init}[0 : T, :, :]$  ▷ Current time steps
18:    $\text{init\_future} \leftarrow \text{init}[l : T_{\text{ext}}, :, :]$  ▷ Future time steps for lagged edges
19:    $\text{init} \leftarrow \text{concat}(\text{init\_now}, \text{init\_future}, \text{dim} = 1)$  ▷ Combine current and future
20: end if
21: return init

```

Algorithm 3 Generate random system trajectories

Require: T (num_timesteps), N (num_nodes)
Ensure: $\text{data} \in \mathbb{R}^{T \times N \times D}$ (trajectory tensor)

- 1: $\text{names} \leftarrow \text{get_3d_chaotic_system_names}()$ ▷ Get list of available 3D systems
- 2: $s \leftarrow \text{random_sample}(\text{names}, \min(N, |\text{names}|))$ ▷ Select systems
- 3: $\text{data} \leftarrow []$ ▷ Initialize empty trajectory list
- 4: **for** $i \leftarrow 0$ to $N - 1$ **do**
- 5: $\text{sol} \leftarrow \text{None}$
- 6: $\text{counter} \leftarrow 0$
- 7: **while** $\text{sol} = \text{None}$ **and** $\text{counter} < \text{max_retry}$ **do** ▷ Handle failing integration
- 8: $\text{system} \leftarrow \text{random_choice}(\text{names} \setminus \{s[i \bmod |s|]\})$
- 9: $\text{sol} \leftarrow \text{solve_single_system}(\text{system}, T)$ ▷ Solve system ODEs
- 10: $\text{counter} \leftarrow \text{counter} + 1$
- 11: **end while**
- 12: **if** $\text{counter} \geq \text{max_retry}$ **then**
- 13: **raise** Exception ▷ Failed to integrate system
- 14: **end if**
- 15: $\text{data.append}(\text{sol})$ ▷ Add trajectory to list
- 16: **end for**
- 17: $\text{data} \leftarrow \text{convert_to_tensor}(\text{data})$ ▷ Convert to tensor
- 18: $\text{data} \leftarrow \text{data.permute}(1, 0, 2)$ ▷ Reshape to $[T, N, D]$
- 19: **return** data

Algorithm 4 Generate Sinusoidal Driver Trajectories

Require: T (num_timesteps), N (num_nodes), D (node_dim), P_{max} (max_num_periods)
Ensure: $\text{data} \in \mathbb{R}^{T \times N \times D}$ (sinusoidal trajectory tensor)

- 1: $\text{amplitude} \leftarrow 2 \times \text{rand}(N, D) - 1$ ▷ Random amplitudes in $[-1, 1]$
- 2: $\text{phase_shift} \leftarrow 2\pi \times \text{rand}(N, D)$ ▷ Random phase shifts in $[0, 2\pi]$
- 3: $\text{data} \leftarrow \text{zeros}(T, N, D)$ ▷ Initialize output tensor
- 4: **if** $N > 0$ **then**
- 5: $\text{max_time} \leftarrow P_{\text{max}} \times 2\pi \times \text{rand}(N, D)$ ▷ Random max times
- 6: $\text{time} \leftarrow \text{linspace}(0, \text{max_time}, T)$ ▷ Generate time points
- 7: **for** $i \leftarrow 0$ to $T - 1$ **do**
- 8: $\text{data}[i, :, :] \leftarrow \text{amplitude} \times \sin(\text{time}[i] + \text{phase_shift})$ ▷ Compute sine values
- 9: **end for**
- 10: **end if**
- 11: **return** data

Algorithm 5 Generate Linear Driver Trajectories

Require: T (num_timesteps), N (num_nodes), D (node_dim), $m_{\text{min}}, m_{\text{max}}$ (slope range), $b_{\text{min}}, b_{\text{max}}$ (intercept range)
Ensure: $\text{data} \in \mathbb{R}^{T \times N \times D}$ (linear trajectory tensor)

- 1: $m \leftarrow (m_{\text{max}} - m_{\text{min}}) \times \text{rand}(N, D) + m_{\text{min}}$ ▷ Random slopes in $[m_{\text{min}}, m_{\text{max}}]$
- 2: $b \leftarrow (b_{\text{max}} - b_{\text{min}}) \times \text{rand}(N, D) + b_{\text{min}}$ ▷ Random intercepts in $[b_{\text{min}}, b_{\text{max}}]$
- 3: $\text{data} \leftarrow \text{zeros}(T, N, D)$ ▷ Initialize output tensor
- 4: **if** $N > 0$ **then**
- 5: $\text{max_time} \leftarrow \text{rand}(N, D)$ ▷ Random max times in $[0, 1]$
- 6: $\text{time} \leftarrow \text{linspace}(0, \text{max_time}, T)$ ▷ Generate time points
- 7: **for** $i \leftarrow 0$ to $T - 1$ **do**
- 8: $\text{data}[i, :, :] \leftarrow m \times \text{time}[i] + b$ ▷ Compute linear values
- 9: **end for**
- 10: **end if**
- 11: **return** data

Algorithm 6 Initialize state x for Coupled Causal Models

Require: $\text{init} \in \mathbb{R}^{T \times N \times D}$ (initial values tensor), $A \in \mathbb{R}^{N \times N}$ or $\mathbb{R}^{2N \times N}$ (adjacency matrix),
 standardize (boolean)

Ensure: $x \in \mathbb{R}^{T \times N \times D}$ (initialized tensor with values for root nodes and zero else)

```

1:  $x \leftarrow \text{zeros\_like}(\text{init})$  ▷ Initialize output tensor with zeros
2: if  $A.\text{shape}[0] = A.\text{shape}[1]$  then ▷ No time lag edges
3:    $\text{root\_nodes} \leftarrow \text{get\_root\_nodes\_mask}(A)$  ▷ Identify nodes with no incoming edges
4:    $x[:, \text{root\_nodes}, :] \leftarrow \text{init}[:, \text{root\_nodes}, :]$  ▷ Set values only for root nodes
5: else ▷ With time lag edges
6:    $A_{\text{now}} \leftarrow A[:, A.\text{shape}[0]/2 :]$  ▷ Extract current time step connections
7:    $A_{\text{past}} \leftarrow A[A.\text{shape}[0]/2 :]$  ▷ Extract past time step connections
8:    $\text{root\_now} \leftarrow \text{get\_root\_nodes\_mask}(A_{\text{now}})$  ▷ Current time root nodes
9:    $\text{root\_past} \leftarrow \text{get\_root\_nodes\_mask}(A_{\text{past}})$  ▷ Past time root nodes
10:   $\text{root\_nodes} \leftarrow \text{concat}(\text{root\_now}, \text{root\_past}, \text{dim} = 0)$  ▷ Combine masks
11:   $x[:, \text{root\_nodes}, :] \leftarrow \text{init}[:, \text{root\_nodes}, :]$  ▷ Set values only for root nodes
12: end if
13: if  $\text{standardize}$  then
14:    $\mu \leftarrow \text{mean}(x, \text{dim} = 0)$  ▷ Compute mean over time
15:    $\sigma^2 \leftarrow \text{var}(x, \text{dim} = 0)$  ▷ Compute variance over time
16:    $x \leftarrow \frac{x - \mu}{\sqrt{\sigma^2}}$  ▷ Standardize results
17: end if
18: return  $x$ 

```

B.2 MLP propagation

In the following, we outline how the information is passed through the DAG starting from the root nodes (see Figure 9). A detailed description of the MLP implementation to forward information through the DAG is giving in Algorithm 7.

Note, that this process follows the reverse order, i.e., from the highest node number to zero. This ensures that information is processed in the right order through the network because by construction, a node cannot have incoming edges from a node with higher degree (see Algorithm 1).

To correct for varsortability, we include the option to standardize the propagated time-series to maintain the variance of the original drivers along the causal structure (see Appendix B.3).




Figure 9: MLP propagation.

Algorithm 7 MLP Propagation through the DAG

Require: $A \in \{0, 1\}^{n \times n}$ (adjacency matrix), $W \in \mathbb{R}^{n \times d \times d}$ (weight tensor), $b \in \mathbb{R}^{n \times d}$ (bias tensor), $\text{init} \in \mathbb{R}^{T \times n \times d}$ (initial values), $\phi \in \mathcal{A}^n$, $\text{standardize} \in \{\text{True}, \text{False}\}$

Ensure: $x \in \mathbb{R}^{T \times n \times d}$ (propagated values)

```

1:  $T, n, d \leftarrow \text{shape}(\text{init})$  ▷ Get dimensions
2:  $x \leftarrow \text{initialize\_x}(\text{init}, A)$  ▷ Initialize state tensor, see Algorithm 6
3:  $A \leftarrow A^T$  ▷ Transpose adjacency matrix for easier indexing
4: for  $i \leftarrow n - 1$  to  $0$  do ▷ Process nodes in reverse topological order
5:   if  $A.\text{shape}[0] \neq A.\text{shape}[1]$  then ▷ Handle time lag edges
6:      $i \leftarrow i \bmod A.\text{shape}[0]$ 
7:   end if
8:    $m_i \leftarrow A[i].\text{bool}()$  ▷ Get incoming edge mask
9:   if  $m_i.\text{any}()$  then ▷ If node has any incoming edges
10:     $W_{\text{sel}} \leftarrow W[m_i]$  ▷ Select weights for incoming edges
11:     $b_{\text{sel}} \leftarrow b[m_i]$  ▷ Select biases for incoming edges
12:     $x_{\text{sel}} \leftarrow x[:, m_i]$  ▷ Select input values
13:     $y \leftarrow \text{matmul}(W_{\text{sel}}, x_{\text{sel}}) + b_{\text{sel}}$  ▷ MLP transformation
14:     $y \leftarrow \phi[i](y)$  ▷ Apply activation (optional)
15:     $y \leftarrow \text{sum}(y, \text{dim} = 1)$  ▷ Aggregate incoming signals
16:    if  $\text{standardize}$  then
17:       $\mu \leftarrow \text{mean}(y, \text{dim} = 0)$  ▷ Compute mean over time
18:       $\sigma^2 \leftarrow \text{var}(y, \text{dim} = 0)$  ▷ Compute variance over time
19:       $y \leftarrow \frac{y - \mu}{\sqrt{\sigma^2}}$  ▷ Standardize results
20:    end if
21:     $x[:, i] \leftarrow x[:, i] + y$  ▷ Update node values
22:  end if
23: end for
24: return  $x$ 

```

B.3 Standardization

In this section we present our approach to standardize the generated data by estimating the mean and standard deviation over time in order to account for varsortability, following the idea of [24, 23]. Our approach is similar to standardizing data in for causal discovery on time series to account for confounders such as seasonality (e.g., [21, 9]). More details and examples can be found at <https://github.com/kausable/CausalDynamics/blob/main/notebooks/standardization.ipynb>.

As shown at the example of the SCM in Figure 10a, the variance increases with each node v_k for $k = [0, 1, 2, 3, 4]$ (compare orange line in (b), and the corresponding time series in (c)). By standardizing the data following:

$$\tilde{x}_{v_k}(t) = \frac{x_{v_k}(t) - \mu_{v_k}}{\sigma_{v_k}}, \quad (8)$$

where μ_{v_k} and σ_{v_k} are the mean and standard deviation over time, we can control for the variance increase (see blue line in (b), and the corresponding time series in (d)).




Figure 10: In (a) structural causal model consisting of 5 nodes. In (b) variance comparison for the non-standardized and standardized causal model along decreasing order of nodes (i.e., causal direction). In (c) trajectory of the non-standardized nodes and in (d) for the standardized nodes.

C First-order climate models

In the following, we outline two pseudo-realistic examples of real world dynamical physical systems. For more details, please refer to the documentation at https://kausable.github.io/CausalDynamics/notebooks/climate_causal_models.html.

C.1 XRO: eXtended nonlinear Recharge Oscillator model

The XRO model is a Python implementation [89] of a recharge oscillator (RO) model for El Niño-Southern Oscillation (ENSO) [91] coupled to stochastic-deterministic models for other climate modes (M) (as shown in Figure 11) which allows for a two-way interaction. A detailed description of the model is given by [92]. In the following we will provide an overview.

The system can be described through the following set of equations:

$$\frac{d}{dt} \begin{pmatrix} \mathbf{X}_{\text{ENSO}} \\ \mathbf{X}_{\text{M}} \end{pmatrix} = \mathbf{L} \cdot \begin{pmatrix} \mathbf{X}_{\text{ENSO}} \\ \mathbf{X}_{\text{M}} \end{pmatrix} + \begin{pmatrix} \mathbf{N}_{\text{ENSO}} \\ \mathbf{N}_{\text{M}} \end{pmatrix} + \sigma_{\xi} \xi \quad (9)$$

with

$$\frac{d}{dt} \xi = -r_{\xi} \xi + w(t) \quad (10)$$

where $\mathbf{X}_{\text{ENSO}} = [T_{\text{ENSO}}, h]$ and $\mathbf{X}_{\text{M}} = [T_{\text{NPMM}}, T_{\text{SPMM}}, T_{\text{IOB}}, T_{\text{IOD}}, T_{\text{SIOD}}, T_{\text{TNA}}, T_{\text{ATL3}}, T_{\text{SASD}}]$ are the state vectors of averaged sea surface temperatures (T) over the respective regions (see Table 4). To describe the oscillatory behavior of ENSO the thermocline depth (h) averaged over the ENSO region is additionally included in $\mathbf{X}_{\text{ENSO}}$.




Figure 11: Anomalies of observed sea surface temperatures from the detrended ORAS5 reanalysis [102] for the period 1979-2019. Colored circles represent index regions for ENSO and other modes. Arrows visualize the associated causal graph with constantly coupled modes shown in black and potentially decoupled modes in grey. Figure inspired by [92].

Governing dynamics of $\mathbf{X} = [\mathbf{X}_{\text{ENSO}}, \mathbf{X}_{\text{M}}]$ can be decomposed into linear ($\mathbf{L}$), nonlinear ($\mathbf{N}$) and stochastic (ξ) terms. Linear dynamics contain four submatrices:

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}_{\text{ENSO}} & \mathbf{C}_1 \\ \mathbf{C}_2 & \mathbf{L}_{\text{M}} \end{pmatrix} \quad (11)$$

where $\mathbf{L}_{\text{ENSO}}$ describes ENSO internal dynamics, $\mathbf{L}_{\text{M}}$ internal processes and interactions of other climate modes, and $\mathbf{C}_1$ and $\mathbf{C}_2$ represent the coupling matrices summarizing the feedback of ENSO on other modes and reversed, respectively. To conduct decoupling experiments we therefore intervene on $\mathbf{C}_1$ and $\mathbf{C}_2$. Note, that IOD cannot be decoupled as it is essential for the prediction of ENSO [103] (black arrows in Figure 11) and its asymmetry is therefore included in the nonlinear dynamics

N along with quadratic terms describing ocean advection and sea surface temperature-wind stress feedbacks in the ENSO region [104, 105, 106, 107, 108]. Stochastic forcing ξ is composed of weather and high-frequency noise for example the Madden-Julian Oscillation or westerly wind bursts. Due to a strong seasonality across climate modes, periodic parameters are added to the linear and nonlinear terms:

$$\mathbf{L} = \mathbf{L}_0 + \sum_{j=0}^2 (\mathbf{L}_j^c \cos(j\omega t) + \mathbf{L}_j^s \sin(j\omega t)), \quad (12)$$

$$\mathbf{N} = \mathbf{N}_0 + \sum_{j=0}^2 (\mathbf{N}_j^c \cos(j\omega t) + \mathbf{N}_j^s \sin(j\omega t)) \quad (13)$$

where $\omega = 2\pi/(12 \text{ months})$, and the subscripts $j = [0, 1, 2]$ refer to the mean, annual cycle and the semi-annual components, respectively.

Table 4: Definition of SST indices for climate modes used in the study.

Climate Mode	Abbr.	Geographic region
El Niño–Southern Oscillation	ENSO	Niño3.4 region (170°–120°W, 5°S–5°N)
North Pacific Meridional Mode	NPMM	160°–120°W, 10°–25°N
South Pacific Meridional Mode	SPMM	110°–90°W, 25°–15°S
Indian Ocean Basin mode	IOB	40°–100°E, 20°S–20°N
Indian Ocean Dipole mode	IOD	50°–70°E, 10°S–10°N minus 90°–110°E, 10°S–0°N
Southern Indian Ocean Dipole mode	SIOD	65°–85°E, 25°–10°S minus 90°–120°E, 30°–10°S
Tropical North Atlantic variability	TNA	55°–15°W, 5°–25°N
Atlantic Niño	ATL3	20°W–0°E, 3°S–3°N
South Atlantic Subtropical Dipole	SASD	60°–0°W, 45°–35°S minus 40°W–20°E, 30°–20°S

C.2 qgs: A flexible Python framework of reduced-order multiscale climate models

The `qgs`[†] library provides a Python implementation of a simplified, first-order coupled atmosphere–ocean model, based in part on the Modular Arbitrary-Order Ocean-Atmosphere Model (MAOOAM) [93]. It represents a two-layer quasi-geostrophic atmosphere (ψ^1, ψ^3), represented by barotropic ψ_a and baroclinic θ_a streamfunctions, mechanically and thermally coupled to a shallow-water ocean component (ψ_o), with interactions driven by wind forcing alongside radiative and heat exchanges (see Figure 12a).




Figure 12: In (a) schematic of the atmosphere (grey) and ocean (blue) components of the simplified model, where dotted lines represent periodic and solid lines closed boundary conditions; figure inspired by https://qgs.readthedocs.io/en/latest/files/model/maooam_model.html. In (b) associated causal graph showing the coupling of the individual model components represented by the set of ODEs in Equations 14-17.

The core PDEs are the barotropic and baroclinic streamfunctions for the atmospheric layer and the ocean, plus anomalies of ocean δT_o and atmospheric temperature δT_a . Each field is defined by a finite basis with a zonally periodic channel with no-flux boundary conditions in the meridional direction for the atmosphere and a closed basin with no-flux boundary conditions for the ocean. The fields are then projected on Fourier modes respecting these boundary conditions and the truncated model can be summarized by the following system of ODEs (see Figure 12b):

[†]<https://github.com/Climdyn/qgs/>

$$\begin{aligned}
\dot{\psi}_{a,i} = & -a_{i,i}^{-1} \sum_{j,m=1}^{n_a} b_{i,j,m} (\psi_{a,j} \psi_{a,m} + \theta_{a,j} \theta_{a,m}) - \beta a_{i,i}^{-1} \sum_{j=1}^{n_a} c_{i,j} \psi_{a,j} \\
& - \frac{k_d}{2} (\psi_{a,i} - \theta_{a,i}) + \frac{k_d}{2} a_{i,i}^{-1} \sum_{j=1}^{n_o} d_{i,j} \psi_{o,j},
\end{aligned} \tag{14}$$

$$\begin{aligned}
\dot{\theta}_{a,i} = & \frac{\frac{\sigma}{2}}{a_{i,i} \frac{\sigma}{2} - 1} \left\{ - \sum_{j,m=1}^{n_a} b_{i,j,m} (\psi_{a,j} \theta_{a,m} + \theta_{a,j} \psi_{a,m}) - \beta \sum_{j=1}^{n_a} c_{i,j} \theta_{a,j} \right. \\
& \left. + \frac{k_d}{2} a_{i,i} (\psi_{a,i} - \theta_{a,i}) - \frac{k_d}{2} \sum_{j=1}^{n_o} d_{i,j} \psi_{o,j} - 2 k'_d a_{i,i} \theta_{a,i} \right\} \\
& + \frac{1}{a_{i,i} \frac{\sigma}{2} - 1} \left\{ \sum_{j,m=1}^{n_a} g_{i,j,m} \psi_{a,j} \theta_{a,m} + (\lambda'_a + S_{B,a}) \theta_{a,i} \right. \\
& \left. - \left(\frac{\lambda'_a}{2} + S_{B,o} \right) \sum_{j=1}^{n_o} s_{i,j} \delta T_{o,j} - C'_{a,i} \right\}.
\end{aligned} \tag{15}$$

$$\begin{aligned}
\dot{\psi}_{o,i} = & \frac{1}{M_{i,i} + G} \left\{ - \sum_{j,m=1}^{n_o} C_{i,j,k} \psi_{o,j} \psi_{o,k} - \beta \sum_{j=1}^{n_o} N_{i,j} \psi_{o,j} - (d + r) \sum_{j=1}^{n_o} M_{i,j} \psi_{o,j} \right. \\
& \left. + d \sum_{j=1}^{n_a} K_{i,j} (\psi_{a,j} - \theta_{a,j}) \right\},
\end{aligned} \tag{16}$$

$$\begin{aligned}
\delta \dot{T}_{o,i} = & - \sum_{j,m=1}^{n_o} O_{i,j,m} \psi_{o,j} \delta T_{o,m} - (\lambda'_o + s_{B,o}) \delta T_{o,i} \\
& + (2\lambda'_o + s_{B,a}) \sum_{j=1}^{n_a} W_{i,j} \theta_{a,j} + \sum_{j=1}^{n_a} W_{i,j} C'_{o,j}.
\end{aligned} \tag{17}$$

Here, W , K , d and s denote the coupling coefficients governing ocean–atmosphere interactions; a , g , b and c are the inner-product coefficients for the atmospheric Fourier modes and M , O , C , N those for the ocean. A detailed model description is given in [93] and https://qgs.readthedocs.io/en/latest/files/model/maoam_model.html.

D Benchmark details

D.1 Summary

The final distribution of our preprocessed dataset is illustrated in Figure 13. For the SDE cases, we combine all dynamics with noise level $\delta > 0$. The nonlinear and periodic attributes refer to graphs where at least one of the root nodes represents only a simple dynamics or also mixed with periodic, nonlinear forcing.




Figure 13: Distribution of graphs across hierarchy of dynamical systems, and their associated causal tasks meant to be solved.

D.2 Hierarchy of Differential Equations

The data-generating code for the three tiers of dynamical system complexity described below can be found at <https://github.com/kausable/CausalDynamics/tree/main/scripts>.

Simple. From over 130 ordinary and stochastic models, we filter the `dysts` package[60] for three state variables and available Jacobian, which enables us to extract the adjacency matrix, yielding our 59 target systems. We generate different combination of graphs from these 59 uncoupled (simple) differential equations and perform the following combination of experiments: (i) varying the Langevin noise amplitude ($\delta \in \{0.0, 0.5, 1.0, 1.5, 2.0\}$, where $\delta > 0.0$ generates SDEs integrated with Euler-Maruyama integration scheme), and (ii) mimicking unobserved confounding scenarios. In total, 585 unique simple graphs are generated. Finally, the success rate for evaluating simple cases is 90% (or 3612 out of 4095 possible evaluations). One of the most common failure reasons involves a singular value of the true adjacency matrix where score metrics such as AUROC fail.

Coupled. We then couple the simple systems and perform the following combination of experiments: (i) varying the Langevin noise amplitude ($\delta \in \{0.0, 0.5, 1.0, 1.5, 2.0\}$), (ii) mimicking unobserved confounding scenarios, (iii) changing the number of coupled nodes ($N_{node} \in \{3, 5, 10\}$), (iv) introducing periodic drivers (Coupled: Periodic) at equal ratio as dynamical system drivers or only having the latter, i.e., nonlinear chaotic drivers (Coupled: Nonlinear), (v) time-signal standardization, (vi) time-lags ($\tau \in \{0, 1\}$), (vii) nonlinear edge-level activations. Though 18000 coupled graphs are supposed to be generated, we keep 14096 of them, as the rest exhibit undesired outcomes, e.g., trajectories diverge, even after multiple retries.

Climate. We generated 11 graphs for the different coupling strategies in the coupled ENSO model, and 1 graph for the coupled atmosphere-ocean model. Specific hyparameters follow the original papers and discussed in Section C. All evaluation succeeded for the climate case.

E Experiment details

In this section, we describe additional experimental details, including the baseline models, hyperparameter choices, and evaluation metrics used. Unless otherwise stated, all experiments are conducted on 1xA100 NVIDIA GPU on a 100GB memory node with 12 CPU cores.

E.1 Baselines

As part of our benchmark, we deploy 10 methods. We also summarize the final set of hyperparameters after a comprehensive search on a held-out dataset from the simple case, and attempt to set them to their default configurations. Nevertheless, we find that methods with little to no fine-tuning often outperform those that do so intensively. In all algorithms, we set the maximum lag $\tau_{\max} = 1$, unless otherwise specified. This choice is motivated by the fact that dynamical systems governed by differential equations typically evolve based on a single discretization step, which suffices to capture the system’s causal structure in the resulting time series.

PCMCI+ (Peter Clark Momentary Conditional Independence) [44] extends the PC-algorithm to multivariate time series by first preselecting candidate parents and then applying momentary conditional independence (MCI) tests across all lags to estimate contemporaneous and lagged edges under controlled false discovery. We set the maximum lagged time ($\tau_{\max} = 1$) and the critical p-value ($\alpha_{crit} = 0.01$).

F-PCMCI (Filtered PCMCI) [73] extends PCMCI+ by incorporating transfer-entropy-based feature selection to prefilter candidate parents before momentary conditional independence testing, improving scalability and robustness in high-dimensional time series. We set the maximum lagged time ($\tau_{\max} = 1$) and the critical p-value ($\alpha_{crit} = 0.01$).

VARLiNGAM [76] fits a vector autoregressive model (VAR) and then applies the LiNGAM (Linear Non-Gaussian Acyclic Model) [75] algorithm to the non-Gaussian residuals to recover a directed acyclic graph of contemporaneous effects alongside the estimated lagged coefficients. We set the maximum lagged time to $\tau_{\max} = 1$.

DYNOTEARS [22] adopts a score-based structural VAR (SVAR) formulation to jointly estimate contemporaneous and time-lagged weight matrices by minimizing a penalized least-squares objective expended with a differentiable penalty that forbids any directed cycles in the instantaneous effects. We set the maximum lagged time to $\tau_{\max} = 1$.

TSCI (Tangent Space Causal Inference) [79] builds on the idea that nonlinear dynamics locally resemble linear systems by estimating the data manifold’s tangent spaces. TSCI first learns a continuous vector-field model for each variable’s dynamics, e.g., via neural ODEs or Gaussian processes, then applies Convergent Cross Mapping (CCM) within each tangent to test for causal influence by assessing how well the state-space of one variable predicts another. This yields a single, interpretable causal graph for deterministic dynamical systems. We set the maximum lagged time ($\tau_{\max} = 1$), the delay embedding dimension ($h_{embed} = 2$), and the correlation threshold ($\rho_{crit} = 0.8$).

Neural GC (Neural Granger Causality) [5] employs sparse-input multilayer perceptrons (cMLP) and long short-term memory (cLSTM) to model nonlinear autoregressive dependencies in multivariate time series, using input-weight regularization to identify directed Granger-causal links. In our evaluation, we focus solely on the cLSTM variant. We set the maximum lagged time ($\tau_{\max} = 1$), number of hidden LSTM dimension ($h_{lstm} = 16$), with a learning rate of ($lr = 10^{-3}$) fitted over 100 epochs.

CUTS+ [72] builds on the Granger-causality framework of CUTS [61] by combining a two-stage coarse-to-fine strategy with a message-passing graph neural network to discover causal structure in high-dimensional, irregularly-sampled time series. First, lightweight Granger-causality tests prefilter candidate parents for each variable, dramatically reducing the search space. Then, a GNN simultaneously imputes missing or unevenly spaced data and learns a sparse adjacency matrix via a penalized reconstruction loss that incorporates temporal encoding. This alternating imputation–learning loop makes CUTS+ both scalable and robust to missing data. We set the maximum lagged time ($\tau_{\max} = 1$), the number of hidden MLP dimension ($h_{mlp} = 16$), the number of gated recurrent unit (GRU) ($n_{gru} = 1$), with a learning rate of $lr = 10^{-3}$ fitted over 10 epochs.

RCD (Repetitive Causal Discovery) [77] extends LiNGAM to settings with *latent confounders* under the linear, non-Gaussian, acyclic model (LiNGAM) assumption. Classic LiNGAM [75] identifies a unique DAG by exploiting non-Gaussianity and independence of error terms, but its basic form assumes causal sufficiency. RCD relaxes this by repeatedly partitioning variables into (approximate) causal orders and pruning edges using statistical tests, while allowing for latent common causes. Concretely, RCD alternates between: (i) *screening and ordering* steps that use correlation/independence diagnostics to propose parent sets under linear models with non-Gaussian residuals, and (ii) *refinement* steps that remove spurious links via independence checks on residuals (testing that, in the correct direction, residuals are independent of putative causes) and normality checks to ensure non-Gaussianity (a key LiNGAM identifiability condition). The repetitive loop iteratively updates ordering and adjacency until convergence, yielding a sparse adjacency estimate that is robust to latent confounding under the RCD model assumptions. We set the maximum lagged time ($\tau_{max} = 1$) and the critical p-value ($\alpha_{crit} = 0.01$).

GRaSP (Greedy Sparsest Permutations) [74] is a permutation-based causal discovery method that learns DAGs by searching over variable orderings. The method is grounded in the permutation-DAG correspondence: any variable ordering defines a unique minimal DAG consistent with the observed conditional independencies. Identifying the true DAG is then equivalent to finding the sparsest permutation, i.e., the one that yields the fewest edges. GRaSP employs a greedy local search strategy to approximate this combinatorial optimization. Starting from a random initial ordering, it iteratively applies adjacent swaps of variables. At each step, the algorithm updates the candidate DAG using conditional independence tests and accepts swaps that reduce the number of edges, gradually moving toward a sparser representation. This approach avoids the exhaustive search over all $n!$ permutations, making it scalable to moderately high-dimensional problems. By focusing on sparsity, GRaSP is particularly effective when the true underlying causal graph is relatively sparse and conditional independencies can be reliably tested. We set the maximum lagged time ($\tau_{max} = 1$).

TCDF (Temporal Causal Discovery Framework) [68] uses convolutional neural networks (CNN) with an attention mechanism to detect lagged causal relationships in multivariate time series, combining temporal filters and statistical pruning to handle nonlinear, high-dimensional, and noisy data. We set the maximum lagged time ($\tau_{max} = 1$), the CNN kernel size ($h_{cnn} = 4$), the dilation coefficient of 4 with a learning rate of $lr = 10^{-2}$ fitted over 100 epochs.

Note that most of these models, in the standard implementation, only infer a single and stationary causal graph. For this benchmark, we focus on summary graph estimation. Though some algorithms, such as PCMCi+, are able to infer a lagged adjacency matrix, this remains less common and will be the focus of the next version of the benchmark when the need arises and the choice of baseline algorithms capable of this proliferates.

E.2 Metrics

We describe the metrics used in this work, including AUROC [94, 95], AUPRC [96], and SHD [97, 98].

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the true DAG on n nodes $v_i \in \mathcal{V}$ with edge $(i, j) \in \mathcal{E}$ and let $\hat{s} : V \times V \rightarrow \mathbb{R}$ be a scoring function so that an edge (i, j) is predicted whenever $\hat{s}(i, j) > \tau$. For each threshold τ , we define a true-positive rate (TPR) and false-positive rate (FPR):

$$\text{TPR}(\tau) = \frac{|\{(i, j) \in \mathcal{E} : \hat{s}(i, j) > \tau\}|}{|\mathcal{E}|}, \quad \text{FPR}(\tau) = \frac{|\{(i, j) \notin \mathcal{E} : \hat{s}(i, j) > \tau\}|}{n(n-1) - |\mathcal{E}|}, \quad (18)$$

The Area Under the receiver operating characteristic curve (AUROC) is then:

$$\text{AUROC} = \int_0^1 \text{TPR}(\text{FPR}^{-1}(u)) du, \quad (19)$$

which defines the probability that a true edge ranks above a false one, i.e., 0.5 means random, < 0.5 means worse than random, 1.0 means a perfect prediction.

Similarly, we define precision and recall:

$$\text{Precision}(\tau) = \frac{|\{(i, j) \in \mathcal{E} : \hat{s}(i, j) > \tau\}|}{|\{(i, j) : \hat{s}(i, j) > \tau\}|}, \quad \text{Recall}(\tau) = \text{TPR}(\tau). \quad (20)$$

The area under the precision–recall curve (AUPRC) is:

$$\text{AUPRC} = \int_0^1 \text{Precision}(\text{Recall}^{-1}(r)) \, dr. \quad (21)$$

which unlike AUROC depends on the edge-density $\frac{\mathcal{E}}{n(n-1)}$ and thus more accurately reflects performance in the sparse-graph regime. For AUROC the optimal score is 1, i.e., perfect prediction.

Finally, SHD is defined as the minimum number of edge additions, deletions, or reversals required to transform $\hat{\mathcal{G}}$ into $\mathcal{G}$:

$$\text{SHD}(\hat{\mathcal{G}}, \mathcal{G}) = |\{(i, j) \in \mathcal{E} \setminus \hat{\mathcal{E}}\}| + |\{(i, j) \in \hat{\mathcal{E}} \setminus \mathcal{E}\}| + \#\{(i, j) : (i, j) \in \hat{\mathcal{E}}, (j, i) \in \mathcal{E}\}. \quad (22)$$

which a smaller SHD indicates a closer match to the ground truth, with 0 corresponding to exact recovery. Unlike AUROC and AUPRC, which are ranking-based metrics, SHD directly evaluates the correctness of the learned graph structure.

F Additional results

F.1 Time discretisation ablation

We investigate the sensitivity of causal discovery algorithms to different sampling frequencies by sub-sampling every 1st, 5th, or 10th time step. As an illustration, we consider one of the most difficult experiments: `coupling=nonlinear_noise=2.00_systems=10_confounder=True_standardize=True_timelag=1`. Results are shown in Table 5. Overall, we find that reducing sampling frequency tends to improve detection accuracy. This is potentially due to the reduction of data redundancy that can noise the algorithm.

Table 5: Ablating time discretisation: AUROC, AUPRC, and SHD scores at different sampling frequencies (every 1st, 5th, or 10th step).

Model	AUROC (1 / 5 / 10)	AUPRC (1 / 5 / 10)	SHD (1 / 5 / 10)
PCMCi+	0.486 / 0.550 / 0.553	0.382 / 0.415 / 0.425	51.9 / 44.4 / 40.5
F-PCMCi	0.547 / 0.557 / 0.566	0.406 / 0.418 / 0.429	45.7 / 44.2 / 38.9
VARLiNGAM	0.519 / 0.518 / 0.548	0.389 / 0.389 / 0.407	51.9 / 51.3 / 47.2
DYNOTEARS	0.546 / 0.539 / 0.574	0.426 / 0.416 / 0.444	38.9 / 41.6 / 37.9
NGC-LSTM	0.497 / 0.500 / 0.500	0.377 / 0.378 / 0.378	56.0 / 34.0 / 34.0
TSCI	0.533 / 0.576 / 0.567	0.416 / 0.458 / 0.443	41.4 / 38.6 / 36.4
CUTS+	0.495 / 0.500 / 0.500	0.376 / 0.378 / 0.378	39.0 / 34.0 / 34.0
RCD	0.500 / 0.501 / 0.505	0.378 / 0.382 / 0.385	34.0 / 34.1 / 34.1
GRaSP	0.498 / 0.484 / 0.490	0.385 / 0.379 / 0.384	35.8 / 36.2 / 36.2
TCDF	0.500 / 0.500 / 0.500	0.378 / 0.378 / 0.378	34.0 / 34.0 / 34.0

F.2 Partial observability ablation

We also test algorithm performance under partial versus full observability, by sampling only one variable per multidimensional node versus averaging across all variables at each node. As an illustration, we consider one of the most difficult experiments: `coupling=nonlinear_noise=2.00_systems=10_confounder=True_standardize=True_timelag=1`. Results are shown in Table 6. We find that fully observing the system yields a modest improvement in detection performance. One possible explanation is that, even when some variables are hidden, their influence still propagates through the variables we do observe, so the loss of information is less severe than it might appear.

Table 6: Ablating partial observability: AUROC, AUPRC, and SHD scores for partially vs. fully observed systems.

Model	AUROC (partial / full)	AUPRC (partial / full)	SHD (partial / full)
PCMCi+	0.486 / 0.500	0.382 / 0.386	51.9 / 50.2
F-PCMCi	0.545 / 0.546	0.405 / 0.406	45.8 / 47.4
VARLiNGAM	0.519 / 0.503	0.389 / 0.381	51.9 / 51.7
DYNOTEARS	0.546 / 0.559	0.426 / 0.426	38.9 / 37.2
NGC-LSTM	0.499 / 0.500	0.377 / 0.378	55.9 / 34.0
TSCI	0.533 / 0.552	0.416 / 0.421	41.4 / 38.5
CUTS+	0.496 / 0.500	0.376 / 0.378	39.2 / 34.0
RCD	0.500 / 0.499	0.378 / 0.377	34.0 / 34.0
GRaSP	0.498 / 0.499	0.385 / 0.388	35.8 / 36.4
TCDF	0.500 / 0.500	0.378 / 0.378	34.0 / 34.0

F.3 Edge-level activation ablation

We also perform additional experiment where the edge-level activation functions vary. For each identical `experiment_id`, we pass a new argument `activation=mixed`, which by default ran-

domly samples activations_names=[‘identity’, ‘sin’, ‘sigmoid’, ‘tanh’, ‘relu’] as the graph’s edge-level activation. We compare results for the following experiments with coupling=nonlinear_noise=0.00_systems=3_confounder=False_standardize=False_timelag=0:

1. Linear activation (identity)
2. Mixed activation

Table 7: Ablating different (linear vs. mixed nonlinear) activations: AUROC, AUPRC, and SHD scores.

Model	AUROC (linear / mixed)	AUPRC (linear / mixed)	SHD (linear / mixed)
CUTS+	0.500 / 0.478	0.378 / 0.417	37.200 / 45.667
DYNOTEARS	0.532 / 0.646	0.407 / 0.541	39.500 / 30.167
F-PCMCI	0.546 / 0.604	0.405 / 0.493	46.650 / 36.333
GRaSP	0.521 / 0.488	0.396 / 0.437	35.500 / 42.667
NGC-LSTM	0.498 / 0.503	0.377 / 0.428	55.900 / 51.278
PCMCI+	0.526 / 0.560	0.398 / 0.462	47.900 / 41.167
RCD	0.500 / 0.495	0.378 / 0.424	34.000 / 38.333
TCDF	0.500 / 0.491	0.378 / 0.430	34.000 / 41.167
TSCI	0.537 / 0.689	0.414 / 0.601	39.850 / 24.889
VARLiNGAM	0.503 / 0.554	0.380 / 0.458	51.700 / 43.667

We find that mixed activations tend to improve causal discovery performance, potentially because directionality is more detectable in nonlinear cases, i.e., the signal is less invertible (see LiNGAM paper [75]). This effect is especially pronounced for deep learning-based methods such as TSCI, which shows an improvement of the SHD score by 40% (Table 7). This suggests that DL-based methods are better at capturing nonlinearities in the data.