arXiv:2602.18662v1 [cs.LG] 20 Feb 2026

Large Causal Models for Temporal Causal Discovery

Nikolaos Kougioulis[1,2] (✉), Nikolaos Gkorgkolis[1,2], MingXue Wang[3], Bora Caglayan[3], Dario Simionato[3], Andrea Tonon[3], and Ioannis Tsamardinos[1,2]

[1] Institute of Applied & Computational Mathematics, FORTH
[2] Computer Science Department, University of Crete
nikolaos.kougioulis@iacm.forth.gr
{gkorgkolis, tsamard}@csd.uoc.gr
[3] Huawei Ireland Research Centre, Dublin, Ireland
{wangmingxue1, bora.caglayan}@huawei.com
dario.simionato1@h-partners.com
andrea.tonon1@huawei-partners.com

Abstract. Causal discovery for both cross-sectional and temporal data has traditionally followed a dataset-specific paradigm, where a new model is fitted for each individual dataset. Such an approach limits the potential of multi-dataset pretraining. The concept of *large causal models (LCMs)* envisions a class of pre-trained neural architectures specifically designed for temporal causal discovery. Prior approaches are constrained to small variable counts, degrade with larger inputs, and rely heavily on synthetic data, limiting generalization. We propose a principled framework for LCMs, combining diverse synthetic generators with realistic time-series datasets, allowing learning at scale. Extensive experiments on synthetic, semi-synthetic and realistic benchmarks show that LCMs scale effectively to higher variable counts and deeper architectures while maintaining strong performance. Trained models achieve competitive or superior accuracy compared to classical and neural baselines, particularly in out-of-distribution settings, while enabling fast, single-pass inference. Results demonstrate LCMs as a promising foundation-model paradigm for temporal causal discovery. Experiments and model weights are available at `github.com/kougioulis/LCM-paper/`.

Keywords: Causal Discovery, Foundation Models, Large Causal Models, Time-Series

1 Introduction

Causal discovery (CD) aims to recover the causal structure of data [40, 47]. In many domains, data are inherently temporal, motivating causal discovery from multivariate time series [43–45].

Classical temporal CD methods provide strong identifiability guarantees under explicit assumptions [43, 47], yet rely on conditional independence testing

or combinatorial search, whose computational and statistical complexity scales poorly with input dimensionality. Methods assuming known functional forms (e.g., linearity) [26, 38] degrade when assumptions are violated.

Neural methods have been introduced that attempt to amortize CD, include variational approaches [1, 19, 34], autoregressive and attention-based models [28], and adversarial-based frameworks [21,22]. While promising, most remain dataset-specific and struggle to generalize across heterogeneous systems, especially in zero-shot settings.

In parallel, foundation models (FMs) [3] have demonstrated that large neural architectures trained on diverse data can generalize effectively across tasks [13, 29,32,33,42,52], thus motivating *Large Causal Models (LCMs)*: pre-trained FMs that can leverage multiple datasets and perform efficient, fast temporal CD. Prior work [48] remains proof-of-concept, limited to low dimensions and narrow synthetic training distributions.

Importantly, a key obstacle to scaling large causal models is the scarcity of large, diverse time-series datasets with ground-truth causal graphs, with prior work relying on synthetic generators [2], which limits generalization.

In this work, we demonstrate that LCMs can be made scalable and robust through data-centric training on a large and heterogeneous corpus, enabling reliable zero-shot generalization and systematic evaluation across diverse temporal settings. Our contributions can be summarized as follows:

1. We formalize *Large Causal Models (LCMs)* as FMs for temporal CD across heterogeneous data distributions, replacing dataset-specific algorithms with foundation model-style pretraining,
2. We demonstrate that scaling failures in prior work stem from insufficient training distribution diversity,
3. We curate a *large-scale collection* of time-series samples and causal graph pairs suitable for both training of LCMs and benchmarking CD methods,
4. We demonstrate superior or competitive performance compared to classical CD methods, with substantially reduced runtime via single-pass CD.

Temporal Structural Causal Models (TSCMs). A Structural Causal Model (SCM) represents each endogenous variable as a function of its causal parents and an exogenous noise term [40,41,47]. For a multivariate time series $\mathbf{V}_t = (V_t^1, \dots, V_t^k)'$, a Temporal SCM (TSCM) is defined as [43] $V_t^j := f^j\Big(\mathrm{Pa}(V_t^j)\Big) + \epsilon_t^j, \quad t \in \mathbb{Z}$, where $\mathrm{Pa}(V_t^j)$ denotes the set of lagged causal parents drawn from the past $\ell_{\max}$ time steps, $\mathrm{Pa}(V_t^j) \subseteq \{\mathbf{V}_{t-1}, \dots, \mathbf{V}_{t-\ell_{\max}}\}$. Here, $\ell_{\max}$ is a fixed *maximum causal lag*. We assume additive noise models, causal stationarity, no latent confounders, and no contemporaneous effects (i.e. effects within a lag $\ell = 0$). Detailed definitions are provided in Appendix .1.

Lagged Causal Graphs. Temporal causal relationships in a TSCM are represented by a lagged causal graph (window graph) [2], a DAG whose nodes correspond to time-indexed variables and whose directed edges are of the form $V_{t-\ell}^i \to V_t^j$ for

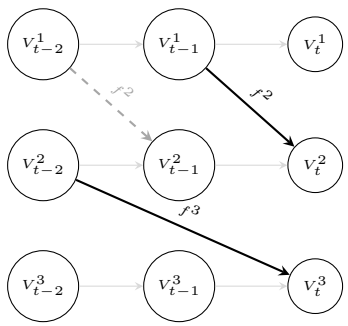

(a) Lagged causal graph

Lag $\ell = 1$

$\mathbb{A}^{(0)}$	V^1	V^2	V^3
V^1	0	0	0
V^2	**1**	0	0
V^3	0	0	0

Lag $\ell = 2$

$\mathbb{A}^{(1)}$	V^1	V^2	V^3
V^1	0	0	0
V^2	0	0	0
V^3	0	**1**	0

(b) Binary lagged adjacency tensor slices Entry $\mathbb{A}^{(\ell-1)}_{j,i} = 1$ represents $V^i_{t-\ell} \rightarrow V^j_t$

Fig. 1. Temporal causal dependencies represented as a (a) lagged causal graph and (b) binary adjacency tensor. Each slice $\mathbb{A}^{(\ell-1)}$ encodes edges at a discrete lag $\ell \leq \ell_{\max}$, where entry $\mathbb{A}^{(\ell-1)}_{j,i} = 1$ denotes $V^i_{t-\ell} \rightarrow V^j_t$.

$\ell \in \{1, \ldots, \ell_{\max}\}$. Edges are oriented forward in time and denote direct causal effects, while indirect effects are encoded via directed paths.

Adjacency Tensor Representation. For V variables and maximum lag $\ell_{\max}$, the lagged causal graph is equivalently represented by a binary adjacency tensor $\mathbb{A} \in \{0,1\}^{V \times V \times \ell_{\max}}$, where $\mathbb{A}^{(\ell)}_{j,i} = 1$ indicates a directed edge $V^i_{t-\ell} \rightarrow V^j_t$ [20, 48](Figure 1).

2 Related Work

Despite advances in deep learning, few works explore pre-trained models for Causality. Most approaches remain task or dataset specific and do not exhibit the broad generalization associated with foundation models.

In *causal inference*, [54] propose CInA, a transformer-based framework for estimating average treatment effects (ATEs) [41] from observational data, which targets treatment effect estimation rather than causal discovery.

Neural approaches to CD include attention-based models for static graph prediction under i.i.d. assumptions [53], supervised sequence formulations in atemporal setups [28] and amortized variational methods under Granger causality assumptions [35]. These methods typically operate on limited synthetic datasets, output summary graphs and do not demonstrate large-scale zero-shot generalization, particularly in temporal or heterogeneous settings.

The closest work is [48], which formulates CD as a multi-class classification problem, and introducing various architectural backbones, notably transformers [50,55]. Although conceptually aligned with our notion for LCMs, it remains

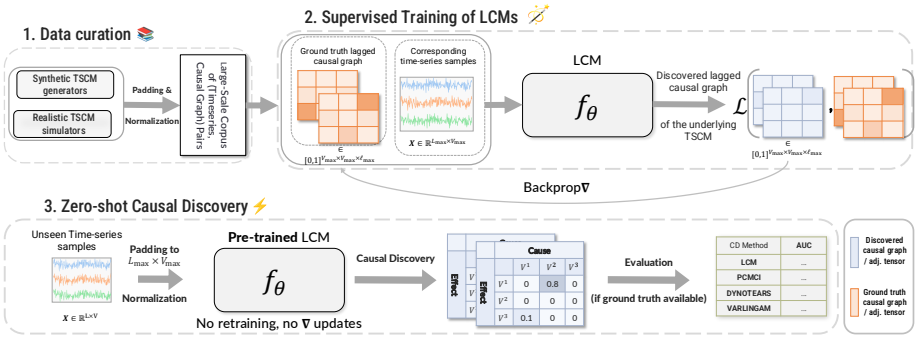

Fig. 2. Overview of the large causal model (LCM) pipeline. (1) Synthetic and realistic TSCM generators produce training pairs of multivariate time series and their lagged causal graphs. (2) The LCM is trained via supervised learning on these pairs to discover a lagged adjacency tensor $\hat{\mathbb{A}}$ for a time series $\mathbf{X} \in \mathbb{R}^{L \times V}$, padded and normalized for stability. (3) At inference (CD phase), the pre-trained LCM predicts causal strengths on unseen datasets in a zero-shot manner.

a proof of concept; limited to small systems, training relies on narrow synthetic distributions, and **performance degrading sharply with increasing dimensionality [48, Figure 3]**.

Positioning of Our Work. We study Large Causal Models as robust, scalable FMs for *temporal* CD. Unlike prior work, we emphasize data diversity as a prerequisite for scaling, train on a large heterogeneous corpus, and evaluate zero-shot and out-of-distribution generalization across diverse benchmarks. This positions LCMs as practical alternatives to existing CD methods.

3 Problem Formulation

We study *temporal CD across heterogeneous datasets and underlying causal mechanisms*. The task is to construct a *Large Causal Model (LCM)*, i.e. an FM that infers lagged causal graphs from multivariate time series, while generalizing across domains, lengths, and causal mechanisms.

Let $\mathbf{X} = \{\mathbf{X}_t\}_{t=1}^{L} \in \mathbb{R}^{L \times V}$ denote a multivariate time series with V variables. Each dataset $\mathcal{D} = \mathbf{X}$ is generated by a TSCM $\mathcal{G}$, $\mathcal{D} \sim \mathcal{P}_{\mathcal{D}}$, where $\mathbb{A}$ is the (ground truth) lagged adjacency tensor of $\mathcal{G}$. The task is to infer $\mathcal{G}$ from $\mathcal{D}$ via a parametric mapping $\mathbf{X} \xrightarrow{f_\theta} \hat{\mathbb{A}}$, where f_θ is a neural model and $\hat{\mathbb{A}} \in \mathbb{R}^{V_{\max} \times V_{\max} \times \ell_{\max}}$ encodes the discovered lagged causal graph (Figure 1). Entry $\hat{\mathbb{A}}_{j,i,\ell_{\max}-\ell}$ represents the confidence in $X_{t-\ell}^{i} \to X_t^{j}$.

Probabilistically, the LCM learns an amortized approximation of the conditional distributions $\mathbb{P}(\mathbb{A}_{j,i,\ell_{\max}-\ell} = 1 | \mathbf{X})$ under the training distribution over

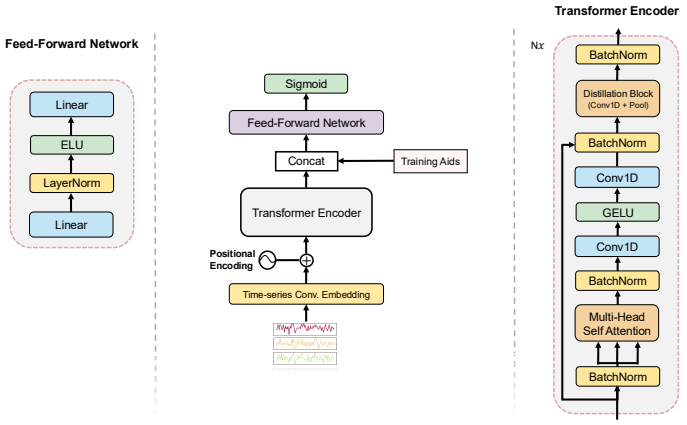

Fig. 3. A multivariate time series is embedded via Conv1D layers and positional encodings, processed through a Transformer encoder stack with optional distillation blocks, and augmented with lagged cross-correlations (training aids). A feedforward head outputs a lagged adjacency tensor representing the discovered temporal causal graph.

the TSCMs. Rather than performing dataset-specific DAG search, the model learns a parametric approximation to the mapping from time-series samples to causal graphs under the training distribution. The discovered lagged adjacency tensor hence represents edge probabilities under the implicit prior induced by the training corpus.

Model Hyperparameters. $L_{\max}$ (max time-series length) and $V_{\max}$ (maximum input variables) are model-specific, akin to token limits in language models [8], while $\ell_{\max}$ (max assumed lag) follows standard assumptions in temporal CD [43].

4 Model Overview

The implemented LCM, illustrated in Figure 3, follows a convolution-enhanced Transformer encoder [48, 55]:

Input Embeddings. Input samples $\mathbf{X} \in \mathbb{R}^{L \times V}$ are normalized and noise-padded (Appendix .2) to $\mathbb{R}^{L_{\max} \times V_{\max}}$, projected via Conv1D to extract local temporal features at dimension d_{model}, with positional encodings added element-wise [50].

Encoder Stack. N layers of multi-head self-attention and feedforward networks with residual connections produce contextualized representations. An optional Conv1D-based distillation layer aims to reduces sequence length for long-horizon processing.

Training Aids. Lagged cross-correlations are concatenated to the final encoder representations to enhance inductive bias (*correlation injection* [48]).

Feedforward Head. A fully-connected network, followed by a sigmoid activation outputs $\hat{\mathbb{A}} \in [0,1]^{V_{\max} \times V_{\max} \times \ell_{\max}}$, where $\mathbb{A}_{j,i,\ell}$ is the probability of $X^i_{t-\ell} \to X^j_t$.

The framework is encoder-agnostic; we adopt Transformers as a scalable, well-understood backbone for multivariate time series [13, 32, 33] rather than as a novel contribution; details are provided in Appendix .4.

4.1 Loss Function

Supervised training of LCMs requires an informative objective that encourages accurate edge predictions while leveraging any observed statistical dependencies. A composite loss is employed, combining supervised edge prediction with a correlation-based regularizing term.

Edge Prediction Loss. Let $\hat{\mathbb{A}} \in \mathbb{R}^{V_{\max} \times V_{\max} \times \ell_{\max}}$ denote the discovered lagged adjacency tensor and $\mathbb{A}$ the ground truth. Each potential edge $X^i_{t-\ell} \to X^j_t$ is treated as a binary classification problem, optimized via binary cross-entropy:

$$\mathcal{L}_{\text{edge}} = \frac{1}{V_{\max}^2 \ell_{\max}} \sum_{i,j,\ell} \text{BCE}\Big(\mathbb{A}_{j,i,\ell}, \hat{\mathbb{A}}_{j,i,\ell}\Big)$$

Regularization term. As observed by [48], causal discovery FMs for time-series can benefit from the inclusion of observed statistics. To improve stability and generalization, an auxiliary term is added, aligning predictions with lagged cross-correlations observed in input time-series. For $\mathbf{X} = \{\mathbf{X}_t\}_{t=1}^n$, correlations up to lag $\ell_{\max}$ are computed and normalized (we denote this tensor $\tilde{\text{CC}}$). The correlation loss corresponds to a weighted mean squared error: $\mathcal{L}_{\text{corr}} = \mathbb{E}\big[(\hat{\mathbb{A}} - \tilde{\text{CC}})^2 \odot \tilde{\text{CC}}^{\gamma}\big]$ where $\gamma > 1$ emphasizes strong correlations. This term encourages confident causal predictions to align with empirical temporal dependencies without imposing hard constraints. Although [48] adopt a different regularizing term, we adopt their terminology *correlation regularization, CR* for clarity.

Overall Objective. The final loss combines the two components: $\mathcal{L} = \lambda_{\text{edge}}\, \mathcal{L}_{\text{edge}} + \lambda_{\text{corr}}\, \mathcal{L}_{\text{corr}}$, where $\lambda_{\text{edge}} = 1$ and $\lambda_{\text{corr}} = 3/4$ are selected by default.

5 Training Setup

5.1 Datasets

LCMs require paired samples $(\mathcal{D}, \mathcal{G})$ for supervised training. To promote zero-shot generalization under distribution shift, we construct a large, diverse corpus combining *synthetic*, *semi-synthetic*, and *realistic* time-series datasets.

Synthetic. We generate random TSCMs by sampling graphs from parametrized random families (e.g., Erdős–Rényi [7,17,23]) and performing ancestral sampling [36, Section 4.2.5], dynamically varying the number of variables, $\ell_{\max}$, graph density, and functional mechanisms (linear and nonlinear additive noise models [27]). Our primary large-scale collection (`Synthetic_2`) contains 230k instances; `CDML` [31] and `Synthetic_1` [48] provide additional synthetic evaluation. Details in Appendix .3.

Semi-synthetic. For OOD evaluation, we use mechanistic simulators with known causal connectivities: *f*MRI-based dynamic nonlinear models (*f*`MRI_5`, *f*`MRI`) [9,35,46] and Kuramoto coupled oscillator systems (`Kuramoto_5`, `Kuramoto_10`) [30, 35]. These are never used for training and serve exclusively as OOD benchmarks.

Realistic. We incorporate realistic TSCMs derived from real-world time series following [20], to the best of our knowledge the only method constructing calibrated causal twins from observed, real data. Input datasets span multiple domains (energy, weather, transportation) [24], yielding 45k instances. Additional realistic datasets are further reserved for OOD evaluation (Appendix .3).

Mixture Collections. To analyze synthetic-to-realistic ratio effects, we construct four mixture datasets (`MIX_1`–`MIX_4`) with proportions from 100/0 to 20/80, each containing 50k instances with identical temporal characteristics.

Large-Scale Training Corpus. Our main corpus (`LS`) combines synthetic and realistic data, totaling 275k instances (**137.5M time points**). All evaluations use holdout test sets for in-distribution scenarios and separate OOD benchmarks. Beyond LCM training, these corpora **constitute a substantial benchmark resource** for systematic evaluation of existing and future temporal CD methods.

6 Experimental Setup

Dataset configurations are described in Section 5.1. Regarding LCMs, we consider multiple variants with increasing model capacity, from $\approx$ 900K to $24M$ parameters. Larger models are used to study scalability in variable count and data size, while smaller models are employed for ablation studies and comparisons (Subsections 7.1, 7.2). The selection of such parameter counts is motivated by established time-series forecasting FMs, such as MOIRAI [51], Timer [33], and TimesFM [13]. Configurations, implementation and training details are provided in Appendix .7.

6.1 Evaluation Metrics

Edge discovery performance is evaluated by the *Area Under the ROC Curve (AUC)* [4], computed directly from the predicted lagged adjacency tensor [11,20].

Edges are ranked by confidence, and the area under the TPR–FPR curve is computed across thresholds. Higher AUC scores indicate better recovery of the ground-truth lagged causal graph. Standard errors are estimated across datasets within each collection.

Statistical significance of AUC differences between LCMs and baseline methods is asserted using the Wilcoxon signed-rank test [12], a non-parametric paired test appropriate for potentially non-Gaussian AUC distributions, under a Bonferroni correction ($\alpha_{\text{corrected}} = \alpha/k$) where k is the number of performed comparisons and α the level of significance.

6.2 Baselines

We compare LCMs against established temporal CD methods: **PCMCI** [44], **DYNOTEARS** [38] and **VARLinGAM** [26], representing constraint-based, score-based, and functional-model approaches respectively. For small datasets ($V = 5$, $\ell_{\max} = 3$), we additionally include the pre-trained Transformer from [48] ($\approx$ 1.4M parameters).

All baselines use official implementations. Outputs are converted to lagged adjacency tensors for fair comparison. Regarding constraint-based methods, edge confidence corresponds to inverse p-values. For methods producing causal effect estimates, edge probabilities are estimated via bootstrap resampling (Appendix .8 & Algorithm 1).

7 Experimental Results

We present representative results in each of the following subsections. Our experimental evaluation spans a broad range of settings, including both in-distribution (holdout, out-of-sample test sets) and out-of-distribution scenarios. Across these settings, we show that LCMs consistently match or outperform classical baselines, which require fitting a new model for each individual dataset. This further highlights the effectiveness of LCMs for efficient CD.

7.1 Observed Statistics Improve LCM Performance

We study the impact of incorporating observed statistics into LCM training through a set of *training aids*, as defined in Section 4. We evaluate a baseline LCM, LCM + CI, and LCM + CI + CR following a grid search procedure [25] over $\lambda_{\text{corr}} \in \{1/4, 1/2, 3/4, 1\}$, keeping all other hyperparameters fixed. Experiments are conducted on the `Synthetic_1` holdout set.

Using these training aids substantially improves performance. CI alone yields a significant improvement over the baseline, and adding CR further increases AUC across all tested weights. The highest mean AUC is achieved for $\lambda_{\text{corr}} \in [0.25, 0.75]$, with no statistically significant difference among them. We adopt $\lambda_{\text{corr}} = 0.75$ for all subsequent experiments. In the remainder of this paper, we refer to the combination of CI and CR collectively as *training aids*.

Table 1. Ablation of training aids on `Synthetic_1`. Asterisks (*) indicate statistical significance over the preceding model (left column).

	LCM	+CI	+CR (λ)			
			.25	.5	.75	1.0
AUC	.868	.914*	.926*	.918*	**.926**	.925

Table 2. Out-of-distribution performance of LCMs trained on varying synthetic/realistic mixtures, evaluated on semi-synthetic fMRI benchmarks and the OOD `Power` & `Climate` benchmarks. Statistical significance versus the 80/20% reference (**bold**) is indicated by an asterisk (*****).

Synth/Sim%	**fMRI_5**	**fMRI**	**Power**	**Climate**
100/0%	.924* ± .050	.874* ± .142	.966* ± .011	.938* ± .035
80/20%	**.966** ± .028	**.960** ± .028	**.981** ± .006	**.982** ± .011
50/50%	.951 ± .050	.944* ± .047	.979 ± .008	.988 ± .010
20/80%	.968 ± .004	.966* ± .036	.975* ± .010	.986 ± .015

7.2 LCMs benefit from training on mixtures of synthetic and realistic data

Training LCMs solely on synthetic data limits generalization beyond narrowly defined causal structures. To mitigate this, we adopt a *dataset mixing* approach, combining synthetic with realistic datasets during training (Subsection 5.1).

Table 2 shows OOD CD performance on semi-synthetic fMRI benchmarks and the real-data–derived `Power` and `Climate` collections. Models trained solely on synthetic cases are consistently outperformed by mixed-dataset models.

Adding a modest fraction of realistic data (20%) yields systematic, statistically significant performance gains. Higher proportions provide marginal or inconsistent improvements. Thus, an 80/20 synthetic-to-realistic mixture balances structural diversity and real-world fidelity, consistent with recent time-series FM findings [13], and is adopted for all subsequent large-scale experiments.

7.3 Rich Training Data Enable Scalable, Large-Scale LCMs

Prior work suggests that *scaling LCMs to higher-dimensional settings leads to near-random performance, even with increased model capacity [48, Figure 3], implying the need for prohibitively deep models.* Our results indicate that these scaling failures are inherent to limitations of the training data distribution rather than architectural constraints. By training on rich datasets, LCMs can scale beyond previously studied small systems (e.g., 3–5 variables), demonstrating stable convergence and generalization up to $V_{\max} = 12$ in our experiments.

Performance on Synthetic Data Table 3 shows that LCMs achieve state-of-the-art performance on synthetic benchmarks, outperforming all baselines. We attribute their lower performance to their restrictive assumptions, as our models benefit from exposure to a diverse large-scale training corpus.

Table 3. Causal discovery performance (AUC) of large-scale LCMs and baseline methods across *synthetic* benchmarks. Results are reported as mean ± standard deviation across datasets. An asterisk indicates a statistically significant difference with respect to the best-performing model (LCM-`12.2M`), which achieves the highest mean AUC across all semi-synthetic collections.

Model	Synthetic_1	Synthetic_2	CDML
$\text{LCM}_{\texttt{2.5M}}$	.995* ± .013	.903* ± .107	.772 ± .177
$\text{LCM}_{\texttt{9.4M}}$	.995* ± .013	.906* ± .106	.768 ± .171
$\text{LCM}_{\texttt{12.2M}}$	**.996** ± .013	**.909** ± .103	**.773** ± .172
$\text{LCM}_{\texttt{24M}}$	.990* ± .012	.865* ± .133	.727* ± .183
CP (Stein et al.)	.943* ± .151	–	–
PCMCI	.671* ± .444	.800 ** ± .181	.521* ± .108
DYNO	.551* ± .216	.605* ± .152	490* ± .087
VARLiNGAM	.922* ± .145	.861* ± .136	.517* ± .099

Out-of-distribution / Zero-shot Performance We evaluate zero-shot transfer on semi-synthetic and real-world datasets (Subsection 5.1), unseen during training, to assess generalization under distribution shifts, in Tables 4 & 5 respectively. An asterisk indicates a statistically significant difference with respect to the best-performing model ($\text{LCM}_{\texttt{24M}}$ for Table 4, $\text{LCM}_{\texttt{2.5M}}$ for Table 5), achieving the highest mean AUC across most collections.

Results demonstrate that all model sizes achieve superior or competitive zero-shot performance compared to baseline methods. This highlights that *richer, diverse training data enable strong generalization*, confirming that scaling failures stem from insufficient training distributions rather than architectural limitations.

7.4 Running Times

We compare computational runtimes of LCMs and conventional CD algorithms (Section 6.2) to highlight scalability advantages. For an LCM, this corresponds to a single forward pass, effectively replacing iterative search of classic CD methods with parametric approximation. Each method is run 10 times per dataset, and results are averaged (Figure 4).

A key observation is that LCM runtimes remain effectively *independent of input dimensionality*, as computations operate on learned representations rather

Table 4. Out-of-distribution (zero-shot) causal discovery performance (AUC) of large-scale LCMs and baseline methods across *semi-synthetic* benchmarks. Results are reported as mean ± standard deviation across datasets.

Model	*f*MRI_5	*f*MRI	Kuramoto_5	Kuramoto_10
$\text{LCM}_{\texttt{2.5M}}$	.981 ± .017	.978 ± .017	.956* ± .017	.919* ± .018
$\text{LCM}_{\texttt{9.4M}}$	.983 ± .007	.979 ± .012	.972 ± .009	.924* ± .008
$\text{LCM}_{\texttt{12.2M}}$	.986 ± .007	.982 ± .012	.953* ± .016	.921* ± .014
$\text{LCM}_{\texttt{24M}}$	**.987** ± .006	**.984** ± .010	**.973** ± .007	**.931** ± .010
CP (Stein et al.)	.883* ± .078	–	.541* ± .108	–
PCMCI	.746* ± .111	.768* ± .110	.476* ± .103	.635* ± .046
DYNO	.528* ± .152	.483* ± .159	.498* ± .077	.519* ± .027
VARLiNGAM	.793* ± .095	.781* ± .101	.479* ± .105	.647* ± .053

Table 5. Out-of-distribution (zero-shot) causal discovery performance (AUC) of large-scale LCMs and baseline methods across *realistic* benchmarks. Results are mean ± std. across datasets. AQ = AirQualityMS, Clim = Climate, Gar = Garments, ETT = ETTm2.

Model	AQ	Clim	Gar	ETT	Gear	Pow
$\text{LCM}_{\texttt{2.5M}}$	$\mathbf{.985}_{\pm.014}$	$\mathbf{.989}_{\pm.007}$	$.995_{\pm.004}$	$\mathbf{.977}_{\pm.007}$	$.998_{\pm.007}$	$.980^{*}_{\pm.007}$
$\text{LCM}_{\texttt{9.4M}}$	$.982_{\pm.013}$	$.983_{\pm.012}$	$.996_{\pm.004}$	$.965_{\pm.024}$	$.998_{\pm.005}$	$\mathbf{.983}_{\pm.005}$
$\text{LCM}_{\texttt{12.2M}}$	$.977^{*}_{\pm.018}$	$.987_{\pm.010}$	$\mathbf{.996}_{\pm.003}$	$.975_{\pm.010}$	$\mathbf{.999}_{\pm.003}$	$.983^{*}_{\pm.005}$
$\text{LCM}_{\texttt{24M}}$	$.970^{*}_{\pm.020}$	$.982_{\pm.014}$	$.995_{\pm.006}$	$.972_{\pm.053}$	$.996_{\pm.010}$	$.979^{*}_{\pm.008}$
PCMCI	$.565^{*}_{\pm.209}$	$.927_{\pm.096}$	$.972^{*}_{\pm.053}$	$.898^{*}_{\pm.035}$	$.980_{\pm.040}$	$.966_{\pm.028}$
DYNO	$.706^{*}_{\pm.218}$	$.708^{*}_{\pm.118}$	$.536^{*}_{\pm.051}$	$.566^{*}_{\pm.037}$	$.693^{*}_{\pm.083}$	$.551^{*}_{\pm.053}$
VARLiNGAM	$.539^{*}_{\pm.146}$	$.938^{*}_{\pm.075}$	$.984^{*}_{\pm.020}$	$.906^{*}_{\pm.024}$	$.995^{*}_{\pm.009}$	$.941^{*}_{\pm.030}$

than fitting a new model per input dataset. In contrast, traditional methods exhibit superlinear scaling with variable count and temporal lag, possibly limiting applicability in real-time, high-dimensional settings.

8 Model Complexity

We derive a closed-form expression for the number of trainable parameters of our LCM models. Let B denote the number of encoder blocks, d_{model} the model dimension, n_{heads} the number of attention heads, and d_{ff} the feedforward dimension. We introduce the binary indicators $I_{\text{train}}, I_{\text{distil}} \in \{0, 1\}$ for the use of our correlation-based training aids and attention distillation respectively. The kernel size of the convolutional embedding is denoted by k.

The total number of trainable parameters decomposes into terms for (i) input token embeddings, (ii) B encoder blocks, (iii) $(B - 1)$ distillation layers when enabled, and (iv) the final causal prediction head. Summing the above,

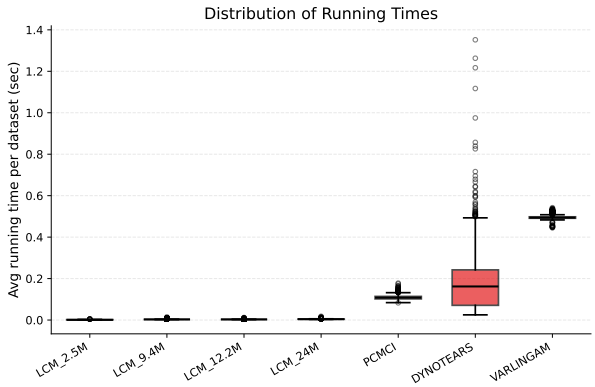

Fig. 4. Running times (in seconds) for LCMs and baseline algorithms on the `Synthetic_2` holdout set, averaged over 10 runs. Traditional methods (e.g., PCMCI & DYNOTEARS) scale superlinearly with lag and variable count, while Transformer-based LCMs remain effectively independent of input dimensionality due to their constant-time forward pass.

$$
\begin{aligned}
P_{\text{total}} = {} & (d_{\text{model}} V_{\max} \cdot k + d_{\text{model}}) \\
& + B\big(4 \cdot d_{\text{model}}^2 + 2d_{\text{ff}} \cdot d_{\text{model}} + d_{\text{ff}} + 9 \cdot d_{\text{model}}\big) \\
& + I_{\text{distil}}(B-1)\big(d_{\text{model}}^2 \cdot k + 3 \cdot d_{\text{model}}\big) \\
& + \big((d_{\text{model}} + I_{\text{train}} \cdot A)d_{\text{ff}} + d_{\text{ff}}\big) + (d_{\text{ff}} \cdot A + A).
\end{aligned}
$$

where $A = V_{\max}^2 \ell_{\max}$ and d_{model} is assumed to be divisible by n_{heads}. Dominant term grows as $\Theta(V_{\max}^2 \ell_{\max})$, i.e., pairwise causal edges in the feedforward head. For $V_{\max} = 25, \ell_{\max} = 3$ with `24M` hyperparameters, the head alone contributes $\approx$ 2M parameters, total $\approx$ 75M. Our considered input dimensionality ($V_{\max} = 12, \ell_{\max} = 3$) represents a careful balance between scalability, generalization, and applicability, with higher-dimensional extensions possible via time-series selection [49].

9 Training Data–Model Size Convergence

We study the interaction between model capacity and training data size by training 500K, 1M, and 2M parameter LCMs on subsampled datasets of sizes $\{10\text{K}, 25\text{K}, 50\text{K}, 100\text{K}\}$ from our large-scale corpus (five random seeds; fixed validation/test sets of 2×10^3 samples).

As shown in Figure 5, smaller models saturate early with diminishing returns from additional data, while larger models continue to improve and increasingly outperform them at scale. Results confirm that LCM performance is jointly constrained by model capacity and data availability: scaling either alone yields limited gains, but scaling both together leads to consistent improvements. We do

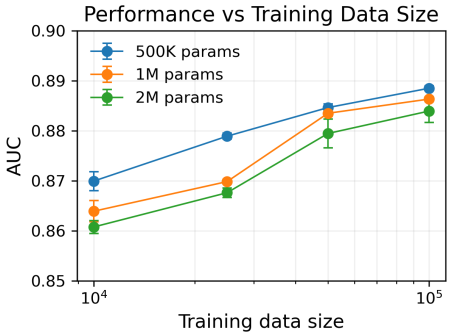

Fig. 5. Empirical convergence of LCMs with increasing training data. Test AUC for 500K, 1M, and 2M parameter models trained on subsampled datasets. Validation/test sets are fixed to isolate the effect of data scale.

not claim a universal scaling law; rather, these findings inform practical design choices for large-scale temporal CD.

10 Concluding Remarks & Future Work

Large Causal Models (LCMs) are foundation models for temporal causal discovery that amortize CD under a supervised scheme, enabling fast, single-pass graph prediction without any optimization. Trained on mixed data collections, LCMs exhibit strong in- and out-of-distribution generalization, consistently matching or outperforming established baselines. Future work includes relaxing causal sufficiency to handle latent confounders, contemporaneous effects, and scaling to higher-dimensional systems and longer horizons.

11 Limitations

LCMs operate under the causal assumptions defined in Section 1 and do not detect or correct violations. When assumptions are violated, predictions may reflect training distribution biases rather than true causal structure, and should be interpreted accordingly.

12 Impact Statement

LCMs enable scalable, zero-shot temporal causal discovery without dataset-specific retraining, potentially accelerating scientific discovery in fields such as biology and marketing by reducing reliance on costly randomized control trials & A/B testing.

References

1. Ashman, M., Ma, C., Hilmkil, A., Jennings, J., Zhang, C.: Causal reasoning in the presence of latent confounders via neural ADMG learning. arXiv preprint arXiv:2303.12703 (2023)
2. Assaad, C.K., Devijver, E., Gaussier, E.: Survey and evaluation of causal discovery methods for time series. Journal of Artificial Intelligence Research **73**, 767–819 (2022)
3. Bommasani, R.: On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021)
4. Bradley, A.P.: The use of the area under the ROC curve in the evaluation of machine learning algorithms. Pattern Recognition **30**(7), 1145–1159 (1997)
5. Breiman, L.: Random forests. Machine Learning **45**(1), 5–32 (2001)
6. Brockwell, P.J., Davis, R.A.: Time series: Theory and Methods. Springer Science & Business Media (1991)
7. Brouillard, P., Lachapelle, S., Lacoste, A., Lacoste-Julien, S., Drouin, A.: Differentiable causal discovery from interventional data. Advances in Neural Information Processing Systems (2020)
8. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. Advances in Neural Information Processing Systems **33**, 1877–1901 (2020)
9. Buxton, R.B., Wong, E.C., Frank, L.R.: Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magnetic Resonance in Medicine **39**(6) (1998)
10. Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining. pp. 785–794 (2016)
11. Cheng, Y., Wang, Z., Xiao, T., Zhong, Q., Suo, J., He, K.: Causaltime: Realistically generated time-series for benchmarking of causal discovery. International Conference on Learning Representations (2024)
12. Conover, W.J.: Practical Nonparametric Statistics. Wiley & Sons (1999)
13. Das, A., Kong, W., Sen, R., Zhou, Y.: A decoder-only foundation model for time-series forecasting. In: International Conference on Machine Learning. vol. 41 (2024)
14. Dickey, D.A., Fuller, W.A.: Distribution of the estimators for autoregressive time-series with a unit root. Journal of the American statistical association **74**(366a), 427–431 (1979)
15. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. In: International Conference on Learning Representations (2017)
16. Durkan, C., Bekasov, A., Murray, I., Papamakarios, G.: Neural spline flows. Advances in neural information processing systems **32** (2019)
17. Fienberg, S.E.: A brief history of statistical models for network analysis and open challenges. Journal of Computational and Graphical Statistics **21**(4), 825–839 (2012)
18. Fix, E.: Discriminatory analysis: nonparametric discrimination, consistency properties, vol. 1. USAF school of Aviation Medicine (1985)
19. Geffner, T., Antoran, J., Foster, A., Gong, W., Ma, C., Kiciman, E., Sharma, A., Lamb, A., Kukla, M., Pawlowski, N., et al.: Deep end-to-end causal inference. Transactions of Machine Learning Research (2024)
20. Gkorgkolis, N., Kougioulis, N., Wang, M., Caglayan, B., Tonon, A., Simionato, D., Tsamardinos, I.: Temporal causal-based simulation for realistic time-series generation. arXiv preprint arXiv:2506.02084 (2025)

21. Goudet, O., Kalainathan, D., Caillou, P., Guyon, I., Lopez-Paz, D., Sebag, M.: Causal generative neural networks. arXiv preprint arXiv:1711.08936 (2017)
22. Goudet, O., Kalainathan, D., Caillou, P., Guyon, I., Lopez-Paz, D., Sebag, M.: Learning functional causal models with generative neural networks. Explainable and Interpretable Models in Computer Vision and Machine Learning pp. 39–80 (2018)
23. Hagberg, A., Swart, P.J., Schult, D.A.: Exploring network structure, dynamics, and function using networkx. Tech. rep., Los Alamos National Laboratory (LANL) (2007)
24. Hahn, Y., Langer, T., Meyes, R., Meisen, T.: Time series dataset survey for forecasting with deep learning. Forecasting **5**(1), 315–335 (2023)
25. Hutter, F., Kotthoff, L., Vanschoren, J.: Automated Machine Learning: Methods, Systems, Challenges. Springer Nature (2019)
26. Hyvärinen, A., Zhang, K., Shimizu, S., Hoyer, P.O.: Estimation of a structural vector autoregression model using non-gaussianity. Journal of Machine Learning Research **11**(5) (2010)
27. Kalainathan, D., Goudet, O., Guyon, I., Lopez-Paz, D., Sebag, M.: Structural agnostic modeling: Adversarial learning of causal graphs. Journal of Machine Learning Research **23**(219), 1–62 (2022)
28. Ke, N.R., Chiappa, S., Wang, J.X., Bornschein, J., Goyal, A., Rey, M., Weber, T., Botvinick, M., Mozer, M.C., Rezende, D.J.: Learning to induce causal structure. In: International Conference on Learning Representations (2023)
29. Kolesnikov, A., Dosovitskiy, A., Weissenborn, D., Heigold, G., Uszkoreit, J., Beyer, L., Minderer, M., Dehghani, M., Neil, H., Sylvain, G., Unterthiner, T., Zhai, X.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021)
30. Kuramoto, Y.: Self-entrainment of a population of coupled non-linear oscillators. In: International Symposium on Mathematical Problems in Theoretical Physics: Kyoto University, Kyoto/Japan. Springer (1975)
31. Lawrence, A.R., Kaiser, M., Sampaio, R., Sipos, M.: Data generating process to evaluate causal discovery techniques for time series data. NeurIPS Causal Discovery & Causality-Inspired Machine Learning Workshop (2020)
32. Liu, Y., Qin, G., Huang, X., Wang, J., Long, M.: Timer-xl: Long-context transformers for unified time series forecasting. arXiv preprint arXiv:2410.04803 (2024)
33. Liu, Y., Zhang, H., Li, C., Huang, X., Wang, J., Long, M.: Timer: Generative pretrained transformers are large time series models. arXiv preprint arXiv:2402.02368 (2024)
34. Lorch, L., Sussex, S., Rothfuss, J., Krause, A., Schölkopf, B.: Amortized inference for causal structure learning. Advances in Neural Information Processing Systems **35**, 13104–13118 (2022)
35. Löwe, S., Madras, D., Zemel, R., Welling, M.: Amortized causal discovery: Learning to infer causal graphs from time-series data. In: Conference on Causal Learning and Reasoning. pp. 509–525. PMLR (2022)
36. Murphy, K.P.: Probabilistic Machine Learning: Advanced Topics. MIT Press (2023)
37. Nauta, M., Bucur, D., Seifert, C.: Causal discovery with attention-based convolutional neural networks. Machine Learning and Knowledge Extraction **1**(1), 19 (2019)
38. Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Georgatzis, K., Beaumont, P., Aragam, B.: DYNOTEARS: Structure learning from time-series data. In: International Conference on Artificial Intelligence and Statistics. pp. 1595–1605. Proceedings of Machine Learning Research (2020)

39. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems **32** (2019)
40. Pearl, J.: Causality. Cambridge University Press, 2nd edn. (2009)
41. Peters, J., Janzing, D., Schölkopf, B.: Elements of causal inference: foundations and learning algorithms. The MIT Press (2017)
42. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training. OpenAI (2018)
43. Runge, J.: Causal network reconstruction from time series: From theoretical assumptions to practical estimation. Chaos: An Interdisciplinary Journal of Nonlinear Science **28**(7) (2018)
44. Runge, J., Bathiany, S., Bollt, E., Camps-Valls, G., Coumou, D., Deyle, E., Glymour, C., Kretschmer, M., Mahecha, M.D., Muñoz-Marí, J., et al.: Inferring causation from time series in earth system sciences. Nature communications **10**(1), 2553 (2019)
45. Runge, J., Gerhardus, A., Varando, G., Eyring, V., Camps-Valls, G.: Causal inference for time series. Nature Reviews Earth & Environment **4**(7), 487–505 (2023)
46. Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W.: Network modelling methods for fMRI. Neuroimage **54**(2), 875–891 (2011)
47. Spirtes, P., Glymour, C., Scheines, R.: Causation, Prediction, and Search. MIT Press (2001)
48. Stein, G., Shadaydeh, M., Denzler, J.: Embracing the black box: Heading towards foundation models for causal discovery from time series data. AAAI Workshop (AI4TS) (2024)
49. Vareille, E., Linardi, M., Tsamardinos, I., Christophides, V.: Chronoepilogi: Scalable time series selection with multiple solutions. Advances in Neural Information Processing Systems **37**, 134287–134316 (2024)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in Neural Information Processing Systems **30** (2017)
51. Woo, G., Liu, C., Kumar, A., Xiong, C., Savarese, S., Sahoo, D.: Unified training of universal time series forecasting transformers. PMLR (2024)
52. Wu, D., He, Y., Cao, Y., Fan, J., Liu, H.: Transformers and their roles as time series foundation models. arXiv preprint arXiv:2502.03383 (2025)
53. Wu, M., Bao, Y., Barzilay, R., Jaakkola, T.: Sample, estimate, aggregate: A recipe for causal discovery foundation models. In: ICLR Workshop on Machine Learning for Genomics Explorations (2024)
54. Zhang, J., Jennings, J., Hilmkil, A., Pawlowski, N., Zhang, C., Ma, C.: Towards causal foundation model: On duality between causal inference and attention. arXiv preprint arXiv:2310.00809 (2023)
55. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: Beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of AAAI. pp. 11106–11115 (2021)

.1 Causal Assumptions & Brief Definitions

All Causal Discovery algorithms are governed by a set of assumptions, regarding the statistical properties of the data and the underlined structure of the causal model. The following assumptions are defined in a way to be applied to either a (standard) causal graph or a temporal causal graph.

Definition 1 (Causal Markov Condition, [40,47]). *Let $\mathcal{G}$ be a causal graph with vertex set $\mathbf{V}$ and $\mathbb{P}$ a probability distribution over $\mathbf{V}$, generated by the causal structure induced by $\mathcal{G}$. Then $\mathcal{G}$ and $\mathbb{P}$ satisfy the Markov Condition if and only if $\forall\ W \in \mathbf{V}$, W is independent of its non-descendants on the causal DAG (non causal effects) given its parents (direct causes) on $\mathcal{G}$.*

Definition 2 (Faithfulness, [40,47]). *Let $\mathcal{G}$ be a causal graph and $\mathbb{P}$ a probability distribution over $\mathbf{V}$. We say that $\mathcal{G}$ and $\mathbb{P}$ are faithful to each other if and only if the all and only the independence relations of $\mathbb{P}$ are entailed by the Causal Markov condition of $\mathcal{G}$. Specifically, $\mathcal{G}$ and $\mathbb{P}$ satisfy the Faithfulness Condition if-f every conditional independence relation true in $\mathbb{P}$ is entailed by the Causal Markov Condition applied to $\mathcal{G}$.*

Definition 3 (Causal Sufficiency, [40,47]). *A set $\mathbf{V}$ of variables is causally sufficient for a population if and only if in the population every common cause of any two or more variables in $\mathbf{V}$ is in V, or has the same value for all units in the population. The common cause Z of two or more variables in a DAG $X \leftarrow Z \rightarrow Y$ is called a confounder of X and Y. Hence Causal Sufficiency implies no unobserved confounders. The notion of causal sufficiency is being used without explicitly mentioning the population.*

Definition 4 (Causal Stationarity, [43]). *Consider the SCM description from Section 1. If the causal relationships between variables $\left(V^i_{t-\tau}, V^j_t\right)$ for time lag $\tau > 0$ also hold for all time-shifted versions $\left(V^i_{t'-\tau}, V^j_{t'}\right)$, the described process is causally stationary. Informally, the graph structure and noise distribution of the SCM are time-invariant.*

Definition 5 (Time-series Stationarity, [6]). *Let $\{X_t\}_{t\in\mathbb{Z}}$ be a time-series with index set $\mathbb{Z} = \{0, \pm1, \pm2, \ldots\}$. The series is said to be stationary if:*

1. $\mathbb{E}|X_t|^2 < \infty, \quad \forall t \in \mathbb{Z}$
2. $\mathbb{E}X_t = m, \quad \forall t \in \mathbb{Z}$
3. $\gamma_X(r,s) = \gamma_X(r+t, s+t), \quad \forall r,s,t \in \mathbb{Z}$

where $\gamma_X(r,s) = Cov(X_r, X_s) = \mathbb{E}\left[(X_r - \mathbb{E}X_r)(X_s - \mathbb{E}X_s)\right],\ r,s \in \mathbb{Z}$ is the autocovariance function of the stochastic process $\{X_t\}_{t\in\mathbb{Z}}$.

We now give definitions on used terminology for our causal graphs.

Definition 6 (Lagged Causal Graph - Window Causal Graph). *A Lagged Causal Graph is a directed acyclic graph (DAG) that represents the causal relationships between time-shifted variables in a multivariate time-series. Formally, let* $\mathbf{V}_t = \{V_t^1, V_t^2, \ldots, V_t^k\}$ *denote the set of observed variables at time step* t*. A lagged causal graph* $\mathcal{G}$ *contains directed edges of the form* $V_{t-\tau}^i \to V_t^j$*, where* $\tau \in \{1, 2, \ldots, \ell_{\max}\}$ *is a positive time lag. An edge* $V_{t-\tau}^i \to V_t^j$ *indicates that* V^i *is a direct cause of* V^j *with a time lag of* τ*.*

Definition 7 (Summary Causal Graph). *Given a lagged causal graph* $\mathcal{G}$ *with edges of the form* $V_{t-\tau}^i \to V_t^j$ *for various lags* τ*, the corresponding summary graph* $\mathcal{S}$ *contains an edge* $V^i \to V^j$ *if there exists at least one lag* τ *such that* $V_{t-\tau}^i$ *is a direct cause of* V_t^j *in* $\mathcal{G}$*.*

Essentially, a summary graph is a simplified representation of the lagged causal structure in a time-series, where the temporal information about lags is abstracted away. The summary graph captures whether a causal relationship exists between two variables, but it does not specify the lag or delay at which the causal influence occurs.

.2 Input Handling and Padding Strategies

	W/o Training Aids	**W/ Training Aids**
Input dimensions	d_{model}	$d_{\text{model}} + 2 \cdot V_{\max}^2 \times \ell_{\max}$
Hidden state	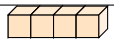	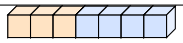

Fig. 6. Hidden state representations with and without auxiliary training aids.

In order to handle variable-length time-series, variable numbers of time-series and differing maximum lags across datasets, we adopt a set of standardized padding strategies, both for the input time-series and the ground truth lagged adjacency tensor. This enables consistent batch processing and generalization of trained models across heterogeneous dataset configurations, as the model is able to handle an input of fixed dimensions, as in conventional foundation models.

Time-Series Padding. When the number of observed time-steps (samples) L is less than the configured maximum number of samples $L_{\max}$, the time-series is *padded with Gaussian noise* $\mathcal{N}(0, 0.01)$ *along the time-step dimension.* This prevents introducing zero artifacts, preserving the marginal statistics, and makes the model robust to handling varying sequence lengths. When inputs are longer than the configured maximum (either in time-steps L or the number of time-series $d_{\max}$), the excess is *truncated* to fit the maximum dimensions. This is conceptually similar to how Foundation models like large language models - LLMs truncate input lengths that exceed the maximum allowed tokens [8].

Causal Graph Padding. A similar procedure is carried out for padding of the lagged causal graph. When a dataset has fewer variables than $V_{\max}$, or the lagged causal graph assumes a lower maximum lag $l < \ell_{\max}$, the lagged adjacency tensor is *zero-padded* alongside the unused variable and lag dimensions. Such padding explicitly encodes the absence of nodes or time dependencies outside the observed configuration. Such a step, together with time-series padding described previously, is crucial for both training of the LCMs, as well as during evaluation against the ground truth causal graph, as it allows handling of time-series and causal graph pairs of variable time-step, node and lag dimensions. An illustration of the above is shown in Figure .2.

Input Concatenation. Implemented LCMs, unless specified otherwise, concatenate lagged crosscorrelations (training aids), which are flattened along the batch dimension to the last hidden state (as illustrated in Figure 3 of the main text). We depict the above in Figure 6.

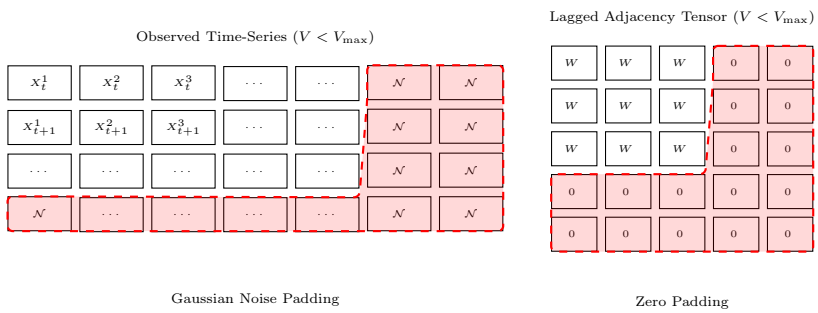

Fig. 7. Illustration of the padding strategies in LCMs. *Left:* Gaussian noise padding for time-series where $V < V_{\max}$ and $L < L_{\max}$. *Right:* Zero padding of lagged adjacency tensors when $V < V_{\max}$ or $\ell < \ell_{\max}$. Red dashed boxes indicate padded regions. $\mathcal{N}$ denotes $\mathcal{N}(0, 0.01)$ for brevity.

.3 Datasets

Synthetic In line with the synthetic temporal data generators provided in `Tigramite` [45], we implement a temporal SCM-based time-series generator to produce synthetic datasets for training and validating our LCMs. The generator produces both observational samples by simulating temporal structural equation models with additive noise, as defined in Section 1. Synthetic data allow us to evaluate and compare algorithms under controlled conditions while ensuring reproducibility and flexibility.

The data generation process consists of several steps. Initially, a lagged temporal causal graph is constructed by sampling a random directed acyclic graph (DAG) following the Erdős-Rényi scheme [7]. Similar to [31], our implementation

offers extensive configurability, including the number of variables and samples, the minimum and maximum lags $\ell_{\min}$ and $\ell_{\max}$, as well as the graph density. To model a variety of realistic causal mechanisms, we incorporate a diverse set of (equiprobable) functional dependencies in the SCM, including both linear and nonlinear relationships, and additive Gaussian noise; examples are summarized in Table 7. Observational samples are generated via ancestral sampling, similar to the procedure used in latent variable models [36, Section 4.2.5], by traversing the topological order of the lagged causal graph. The overall sampling procedure is summarized in Algorithm 2. For ensuring stability of the generated time-series samples, we employ a warm-up period as in [45], discarding an initial set of samples before collecting data.

Table 6. Synthetic time-series dataset collections. L = linear, NL = non-linear functional relationships.

Collection	Pairs	Variables	Timesteps	$\ell_{\max}$	Func. Rels.
`Synthetic_1`	100k	3–5	500	3	L, NL
`Synthetic_2`	270k	3–12	500	3	L, NL
`CDML`	240	4–11	500	3	L, NL

An important consideration in synthetic data generation is the stationarity of the resulting time-series. The generator in `Tigramite` addresses this by performing eigenvalue tests on the stability matrix and aborting the generation if instability is detected, without adjusting the parameters. In contrast, [48] handle instability by repeatedly resampling until stationarity is achieved. Both approaches, however, become impractical when generating large, high-dimensional datasets needed for training our LCMs. We mitigate this issue by wrapping the outcomes of unbounded functional dependencies with bounded functions, such as the hyperbolic tangent tanh or the sigmoid σ. This strategy caps the growth of unbounded functions, ensuring stationarity without altering the underlying causal relationships, and thereby improves the stability and scalability of our synthetic data generation while preserving causal structure.

Semi-synthetic Semi-synthetic datasets combine analytically defined structural equations with generative processes and experimental knowledge (e.g., Markov processes of physical systems). These datasets simulate complex systems where the underlying causal graph is known, allowing for controlled benchmarking under quasi-realistic dynamics. An overview is provided in Table 8.

The fMRI collection comprises 27 datasets with variable counts of 5, 10, and 15, and observation lengths ranging from 50 to 5000, simulating BOLD (Blood-Oxygen-Level Dependent) responses from different regions of interest in the brain. The simulation process follows a non-linear balloon model [9], translating neural activity into hemodynamic responses, and is designed to mimic

Table 7. Functional dependencies used in synthetic SCM generation.

Category	Example Function	Description
Linear	$f(x) = ax + b$	Additive linear effect
	$f(x_1, x_2) = a_1x_1 + a_2x_2 + b$	Linear combination of parents
Non-linear	$f(x) = \alpha_n x^n + \dots \alpha_0 x_0$	Polynomial
	$f(x) = e^x$	Exponential Transformation
	$f(x) = \sin(x)$	Sinusoidal
	$f(x_1, x_2) = x_1x_2$	Multiplicative interaction
	$f(x) = \log(1 + \lvert x \rvert)$	Logarithmic transformation
Bounded	$f(x) = \tanh(x)$	Hyperbolic tangent (bounded in $[-1, 1]$)
	$f(x) = \sigma(x)$	Sigmoid $\left(\frac{1}{1+e^{-x}}\right)$

Table 8. Overview of semi-synthetic time-series dataset collections.

Collection	Datasets	Vars	Timesteps	Max Lag	Rels.
fMRI5	21	5	1200–5000	1	Non-linear
fMRI	26	5–10	1200–5000	1	Non-linear
Kuramoto_5	1000	5	500	1	Non-linear
Kuramoto_10	1000	10	500	1	Non-linear

realistic neural connectivity patterns based on known structural priors. Due to our model input limitations, the 15-variable case is discarded.

Realistic To generate realistic (simulated) time-series datasets paired with ground-truth temporal causal graphs, we employ the methodology of [20]. TCS takes as input real multivariate time-series data, learns a Temporal Structural Causal Model (TSCM) consisting of a directed causal graph, functional mechanisms, and noise models, and synthesizes arbitrary numbers of time-series samples from the learned causal model. This procedure produces fully supervised dataset–graph pairs suitable for training LCMs.

For any of these input time-series, we verify for stationarity using the *Augmented Dickey-Fuller test* [14] (ADF). In case non-stationarity is observed, it is normalized using finite differences up to second order. This step assures that data fed into training of LCMs do not possess extreme trends, possibly resulting in unstable training. We follow the python implementation of `statsmodels` with maximum lag order $12 \cdot \left(\frac{T}{100}\right)^{1/4}$ where T is the number of observed timesteps.

Table 9. Overview of real time-series datasets used for generating realistic data pairs, using the methodology by [20].

Collection	Variables	Timesteps	Domain
WTH	12	35064	Weather
ETTh1	7	17420	Power
ETTm1	7	69680	Power
AirQualityUCI	12	9357	Weather
Bike-usage	5	52584	Transportation
Outdoors	3	1440	Environmental
ETTm2	7	69680	Power
Climate	4	1463	Environmental
Garments	7	692	Manufacturing
power-consumption	8	52416	Power
Gearbox-fault	3	2000	Automotive
AirQualityMS	36	4937	Weather

Used configurations & algorithms We implement the same methods for CD, functional dependency and noise estimation for the generation of simulated data as provided by the official implementation of [20] [4], apart from Random Forest Regressors. We instead use kNN Regressors [18] and Gradient Boosting [10] for improved computational efficiency of generating large amounts of data pairs. An overview is shown in Table 10.

Apart from the generation of different time-series samples and causal graph data pairs generated from different configurations of algorithms in Phases 1-3 of TCS, we apply two augmentation strategies before processing them through the TCS pipeline to increase the number of simulated datasets: i) time window sub-sampling and ii) node subsets sub-sampling.

Time Window Sub-sampling Given a multivariate time-series of minimum length $L = 400$, we extract multiple non-overlapping windows of fixed length $w = 200$. Each window is treated as an independent dataset and passed through the full TCS pipeline (causal discovery, function approximation, and noise modeling), resulting in a distinct simulated time-series and corresponding lagged causal graph.

Node sub-sampling For each time window, we additionally generate subsets of variables of size $k \in \{3, 5, 7\}$. Each subset is independently simulated via TCS, with their total count dependent on the dimension of the original dataset. The

[4] `https://github.com/gkorgkolis/TCS`

Table 10. Implemented Methods for Temporal Functional Dependency Estimation, Noise Density Estimation in Temporal Causal Discovery.

Causal Discovery	
Method	**Description**
PCMCI [43]	Constraint-Based Algorithm
DYNOTEARS [38]	Constraint Optimization Algorithm
CP [48]	Pre-trained Transformer
Functional Dependencies	
Random Forest Regressor [5]	Ensemble Learning Regressor
kNN Regressors [18]	Non-parametric Supervised Method
XGBoost [10]	Gradient Boosting Tree Algorithm
AD-DSTCN [37]	Time-series Forecasting Model (TCDF module)
TimesFM [13]	Foundational Time-series Forecaster
Noise Density	
Normal Distribution	Fitted Parametric Distribution
Uniform Distribution	Fitted Parametric Distribution
Neural Splines [16]	Normalizing Flows
RealNVP [15]	Normalizing Flows

above strategies increase variability of training pairs by exposing the model to smaller-scale graphs.

The final collections for training and evaluation are provided in Table 11.

.4 Architecture of LCMs

This section provides a detailed description of the implemented architecture [48, 55] used throughout experiments, consistent with the formulation presented in the main text.

Input Handling & Embeddings

Input Embeddings. The model receives an observational multivariate time series $X \in \mathbb{R}^{B \times L \times V}$, where B denotes the batch dimension, L the number of timesteps, and V the number of variables. During inference, $B = 1$. All inputs are min–max normalized per variable. To allow training across datasets with heterogeneous dimensionality and sequence length, inputs are padded (or truncated) to fixed

Table 11. Overview of realistically generated time-series collections for training and OOD evaluation of LCMs. Collections in **bold** notate inclusion in the `SIM` training corpus and the large scale, `LS` corpus.

Collection	Variables	$\ell_{\max}$	Num. Datasets
WTH	3-12	3	20200
ETTh1	3-7	3	9000
ETTm1	3-7	3	9000
AirQualityUCI	3-12	3	3900
Bike-usage	3-5	3	1800
Outdoors	3	3	800
ETTm2	3-7	3	15
Climate	3,4	3	17
Garments	3-7	3	24
Power	3-8	3	18
Gearbox-fault	3	3	20
AirQualityMS	3-12	3	54

upper bounds $L_{\max} = 500$ and $V_{\max} = 12$, as described in Subsection .4. Temporal padding is performed using i.i.d. Gaussian noise $z \sim \mathcal{N}(0, 0.01)$.

.5 Min–Max Normalization

To ensure consistent scaling across heterogeneous datasets and to stabilize optimization, each variable is min–max normalized prior to being processed. This transformation is applied during both training and inference. For a variable x_t^v, the normalized value $\tilde{x}_t^v$ is defined as

$$\tilde{x}_t^v = \frac{x_t^v - \min(x^v)}{\max(x^v) - \min(x^v) + \epsilon}, \tag{1}$$

where $\epsilon > 0$ ensures numerical stability.

After min–max normalization, sequences are padded along both temporal and variable dimensions to fixed sizes $L_{\max}$ and $V_{\max}$ with sequences longer than $L_{\max}$ truncated, as in Appendix .2.

Temporal Convolutional Embedding. The normalized and padded input is then projected into a latent representation using a one-dimensional convolution operating along the temporal dimension. After permuting the tensor to channel-first format $[B, V_{\max}, L_{\max}]$, a convolution is applied with $c_{\text{in}} = V_{\max}, \quad c_{\text{out}} = d_{\text{model}}$, where d_{model} denotes the model dimension. The resulting embeddings are $X_{\text{tok}} = \text{Conv1D}(X) \in \mathbb{R}^{B \times L_{\max} \times d_{\text{model}}}$.

This convolutional projection serves as a learnable feature extractor capturing local temporal dependencies prior to global attention. It replaces the purely linear token embedding of the vanilla Transformer, following the Informer design.

Positional Encoding. Because self-attention is permutation-invariant with respect to temporal order, explicit positional information is injected using sinusoidal encodings [50]:

$$\begin{aligned} \mathrm{PE}_{(t,2i)} &= \sin\left(\frac{t}{10000^{2i/d_{\text{model}}}}\right), \\ \mathrm{PE}_{(t,2i+1)} &= \cos\left(\frac{t}{10000^{2i/d_{\text{model}}}}\right). \end{aligned} \tag{2}$$

The final embedding representation is thus $X_{\text{emb}} = X_{\text{tok}} + \mathrm{PE}$, followed by dropout for regularization.

.6 Transformer Encoder

The embedding representation X_{emb} is processed by a stack of Informer-style encoder layers. For a layer l:

$$\begin{aligned} E'^{(l)} &= \mathrm{LayerNorm}\left(E^{(l)} + \mathrm{SelfAttn}(E^{(l)})\right), \\ E^{(l+1)} &= \mathrm{LayerNorm}\left(E'^{(l)} + \mathrm{FFN}(E'^{(l)})\right). \end{aligned} \tag{3}$$

Where $\mathrm{SelfAttn}(\cdot)$ denotes scaled multi-head self-attention and

$$\mathrm{FFN}(x) = \mathrm{Conv1D}_2\left(\mathrm{GELU}\big(\mathrm{Conv1D}_1(x)\big)\right) \tag{4}$$

is a two-layer convolutional feed-forward network applied position-wise. Residual connections and layer normalization are also applied after both sublayers to ensure stable propagation of gradients.

Temporal Distillation. Between encoder blocks, a self-attention distillation layer is inserted (included in all trained models). It performs convolutional downsampling followed by normalization, activation, and max pooling:

$$X_{\text{distil}} = \mathrm{MaxPool}\left(\mathrm{ELU}\big(\mathrm{BatchNorm}(\mathrm{Conv1D}(E^{(l)}))\big)\right). \tag{5}$$

This reduces the temporal dimension and lowers computational cost while preserving dominant temporal patterns.

Let the final encoder output be $E_{\text{enc}} \in \mathbb{R}^{B \times L' \times d_{\text{model}}}$ ($L' \leq L_{\max}$) where only the final timestep embedding is retained: $h = E_{\text{enc}}[:, -1, :] \in \mathbb{R}^{B \times d_{\text{model}}}$.

This design relies on the global receptive field of self-attention: the final token representation can attend to all preceding timesteps and therefore serves as a global summary of the sequence.

Auxiliary Correlation Injection To assist the model in inferring directed dependencies, empirical Pearson cross-correlations are computed for each ordered variable pair (X^i, X^j) and lag $\tau \in \{1, \ldots, \ell_{\max}\}$, which are flattened and concatenated to the learned neural representations of the final encoder block.

$$\rho_{X^i,X^j}(\tau) = \frac{\mathrm{Cov}(X^i_{1:T-\tau}, X^j_{\tau+1:T})}{\sqrt{\mathrm{Var}(X^i_{1:T-\tau})\,\mathrm{Var}(X^j_{\tau+1:T})}}.$$

These values form a tensor $\tilde{CC} \in \mathbb{R}^{B \times V_{\max} \times V_{\max} \times \ell_{\max}}$, which is flattened and concatenated with the encoder representation: $Z = \mathrm{Concat}\big(\mathrm{Flatten}(\bar{X}), \mathrm{Flatten}(\tilde{CC})\big)$. This correlation injection acts as a statistical training aid, aligning learned representations with empirically observable lagged dependencies.

Feedforward Head and Output Interpretation The concatenated representation Z is passed through a two-layer feedforward head with layer normalization and SwiGLU activation: $H = \mathrm{SwiGLU}\big(\mathrm{LN}(W_1 Z + b_1)\big)$, $\hat{\mathbb{A}}_{\mathrm{flat}} = W_2 H + b_2$.

After reshaping and a sigmoid activation for the transformation of logits, the lagged adjacency tensor corresponding to the discovered lagged causal graph is

$$\hat{\mathbb{A}} = \sigma\Big(\mathrm{Reshape}(\hat{\mathbb{A}}_{\mathrm{flat}})\Big) \in [0,1]^{B \times V_{\max} \times V_{\max} \times \ell_{\max}}$$

Each entry $\hat{\mathbb{A}}_{j,i,\tau}$ represents the predicted probability that a directed lagged edge $X^i_{t-\tau} \to X^j_t$ exists. Since multiple edges and lags may simultaneously be active, sigmoid activation is used instead of softmax, thus framing temporal causal discovery as a multi-label classification problem.

Limitations. The flattened output head does not enforce permutation equivariance by design. Future directions can explore equivariant neural models that meet this desiderata by design, while retaining expressivity.

.7 Model Sizes & Training Setup

Table 12 summarizes the implemented models and their configurations used in Section 7.

All models are implemented in Python 3.10 using PyTorch 2.2 [39]. Training uses an effective batch size of 64, with gradient accumulation when necessary to accommodate larger models. Learning rate scheduling and early stopping based on validation AUC are employed to ensure stable convergence. Training was conducted on a dedicated workstation equipped with a NVIDIA RTX 4090 GPU (32 GB VRAM) and 128 GB physical memory. Training of large-scale models took around two weeks to complete. Causal Discovery (LCM Inference step) is performed on an AMD Ryzen 5 CPU.

Table 12. Model configurations and parameter sizes for the LCM variants used in the ablations and additional results in Section 7.

Model	#Params	Enc. Layers	d_{model}	Attn. Heads	d_{ff}	Usage
-	905K-914K	1	256	2	128	Subsection 7.1
-	1M	1	256	2	256	Subsection 7.2
LCM-2.5M	2.5M	4	256	4	256	Subsection 7.3
LCM-9.4M	9.4M	4	512	4	512	Subsection 7.3
LCM-12.2M	12.2M	4	512	4	512	Subsection 7.3
LCM-24M	24M	8	1024	8	1024	Subsection 7.3

Algorithm 1 Bootstrap-based Estimation of Edge Probabilities

Require: Time-series dataset $\mathcal{D}$, maximum lag $\ell_{\max}$, number of bootstraps n, causal discovery algorithm configuration $\mathbf{B}_{\text{CD}}$
Ensure: Soft adjacency matrix $\mathbb{A} \in [0,1]^{V \times V \times \ell_{\max}}$
Initialize accumulator $\mathbf{S} \leftarrow \mathbf{0}^{V \times V \times \ell_{\max}}$
for $b = 1$ to n **do**
  {Bootstrap Sampling Phase}
  Draw bootstrap sample $\mathcal{D}^{(b)}$ from $\mathcal{D}$ by resampling with replacement
  {Causal Discovery Phase}
  $\mathbf{A}^{(b)} \leftarrow \text{causalAlg}(\mathcal{D}^{(b)}, \mathbf{B}_{\text{CD}})$
  {Edge Confidence Update}
  $\mathbf{S} \leftarrow \mathbf{S} + f(\mathbf{A}^{(b)})$
  {f extracts edge indicators or confidence scores}
end for
{Normalization Phase}
$\mathbf{A} \leftarrow \mathbf{S}/n$
return $\mathbf{A}$

.8 Model Evaluation Methodology

To ensure a fair and consistent evaluation across all causal discovery methods, we compute performance metrics based on each model's native output format while adapting our evaluation pipeline accordingly. For models that produce real-valued edge confidences (our LCMs and CausalPretraining models), the values are used verbatim as continuous scores for computing the Area Under the ROC Curve (AUC) and the (binarized) ground-truth adjacency labels. For constraint-based methods like PCMCI, we use the inverse of the edge p-values as confidence scores. For models that output an adjacency matrix of causal effect coefficients (like VARLiNGAM and DYNOTEARS), results are not directly comparable to our LCMs that output soft adjacency probabilities and computing AUCs is not directly applicable. For a fair comparison, we follow a bootstrap-based procedure to estimate the probability of each causal edge, as shown in the official

documentation[5]. Specifically, we run VARLiNGAM $n = 10$ times with resampled datasets and compute the proportion of times each edge is discovered, resulting in a soft adjacency matrix of edge probabilities. We follow the same procedure for DYNOTEARS and the same number of boostrapped datasets, as in Algorithm 1.

.9 Pseudocodes

Algorithm 2 Temporal SCM Ancestral Sampling (ANCESTRAL)

Input: Temporal Causal Graph $\mathcal{G}$ of time-series $\{\mathbf{X_t}\}_{t=1}^{\ell_{\max}}$, causal discovery algorithm configuration $\mathbf{B}_{\text{CD}}$, predictive model configuration $\mathbf{B}_{\text{pred}}$, noise estimator configuration $\mathbf{B}_{\text{noise}}$, max lag $\ell_{\max} \in \mathbf{B}_{\text{CD}}$, number of timesteps T, number of warmup steps W

Output: Generated time-series sample $\{\mathbf{X}_t\}_{t=1}^{T}$

Initialize $X_{t=0}^i$ for all $i \in \{1, \ldots, N\}$ with random noise
`// Forward Sampling`
for $t = W + \ell_{max} + 1$ **to** T **do**
  for *each variable* X_t^i *in topological order based on* $\mathcal{G}$ **do**
    Determine lagged parents $\text{Pa}_{X_t^i} \leftarrow \{X_{t-k}^j \mid (X^j, X^i) \in \mathcal{G}, 1 \leq k \leq \ell_{\max}\}$;
    `// Lagged parents` X^j `of` X^i `at time` t
    Compute $X_t^i := f_i(\text{Pa}_t^i) + \epsilon_t^i$
  end
end
Return generated time-series dataset $\{\mathbf{X}_t^i\}_{W+\ell_{\max}}^{T}$

Algorithm 2 describes the ancestral sampling procedure for obtaining multivariate time-series samples from a known TSCM. Initial warming steps are discarded as a way to ensure stability [43].

.10 Illustrative Example & Visualizations

To provide an intuitive demonstration of LCMs, we construct a synthetic temporal causal process consisting of three variables V^1, V^2 and V_3. The example is governed by a simple, linear temporal structural causal model:

$$V_t^1 := \epsilon_1(t), \quad V_t^2 := 3V_{t-1}^1 + \epsilon_t^2, \quad V_t^3 := V_{t-2}^2 + 5V_{t-3}^1 + \epsilon_t^3,$$

where $\epsilon_t^i \sim \mathcal{N}(0,1)\ \forall i \in \{1,2,3\}, t \in \mathbb{Z}$. The TSCM thus induces the following lagged causal relationships:

$$V^1 \xrightarrow[\text{lag}=1]{} V^2, \quad V^1 \xrightarrow[\text{lag}=3]{} V^3, \quad V^2 \xrightarrow[\text{lag}=2]{} V^3.$$

[5] `https://lingam.readthedocs.io/en/latest/tutorial/var.html`

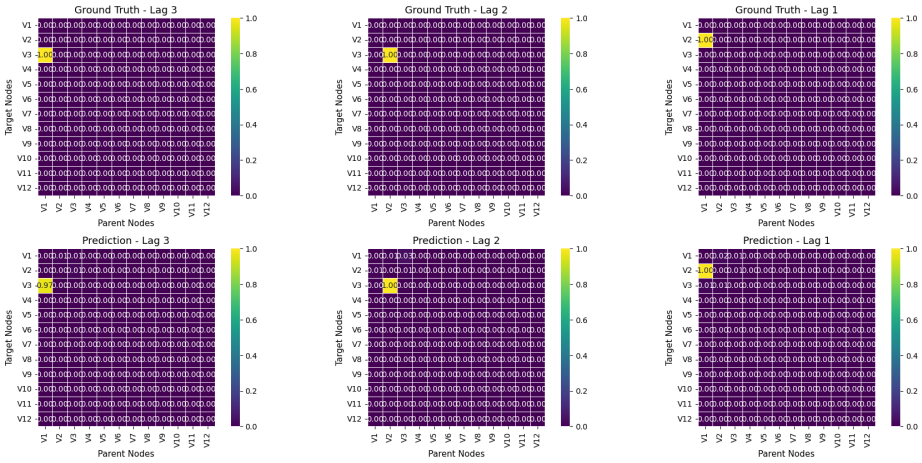

Fig. 8. Heatmap visualization of the discovered lagged adjacency matrix (bottom row) compared to the ground truth (top row) in an example with $n_{\max} = 12$ and $\ell_{\max} = 3$ using the pretrained `LCM-2.4M` model. Brighter colors in the predicted adjacency indicate stronger confidence for edge existence.

In Figure 8 we visualize the discovered causal graph of the final feedforward head against the ground truth; the pretrained LCM successfully captures the underlying causal structure of this simple temporal dataset, resulting in a perfect AUC score of 1.00.

.11 Statistical Significance of Results

We evaluate statistical significance of AUC differences between LCMs and baseline methods using the Wilcoxon signed-rank test [12], a non-parametric paired test appropriate for potentially non-Gaussian AUC distributions. To control the family-wise error rate across multiple pairwise comparisons, we apply Bonferroni correction ($\alpha_{\text{corrected}} = \alpha/k$) where k is the number of performed comparisons.

For each dataset collection, we report the median paired difference in AUC (Δ_{AUC}) relative to a common reference model (the highest average AUC per collection).

Figures 9–11 visualize the paired effects from Subsection 7.3. Asterisks denote statistically significant differences after correction.

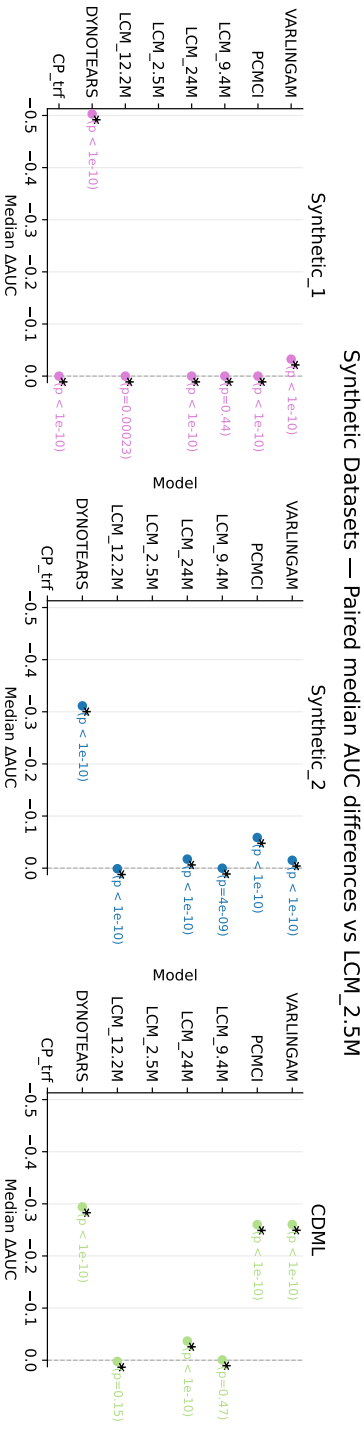

Fig. 9. Paired median AUC differences relative to the reference model across *Synthetic* collections. Positive values indicate improved performance. Statistical significance is determined via Wilcoxon signed-rank test with Bonferroni correction.

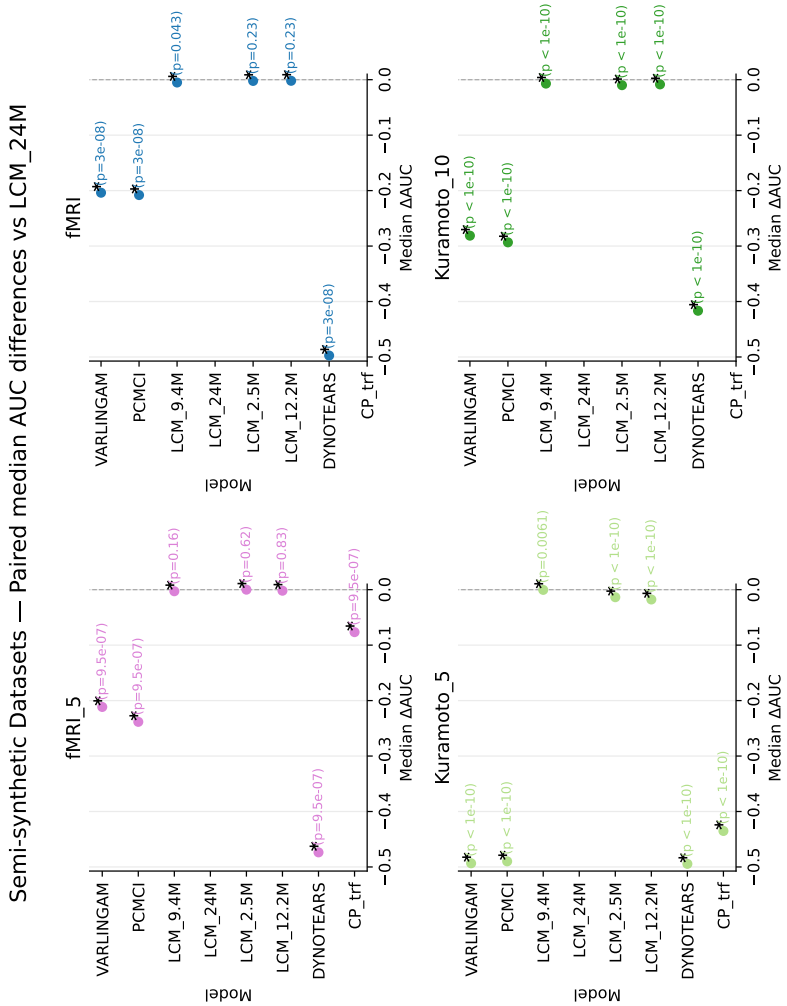

Fig. 10. Paired median AUC differences across *Semi-Synthetic* collections. Positive values indicate improved performance. Statistical significance is determined via Wilcoxon signed-rank test with Bonferroni correction.

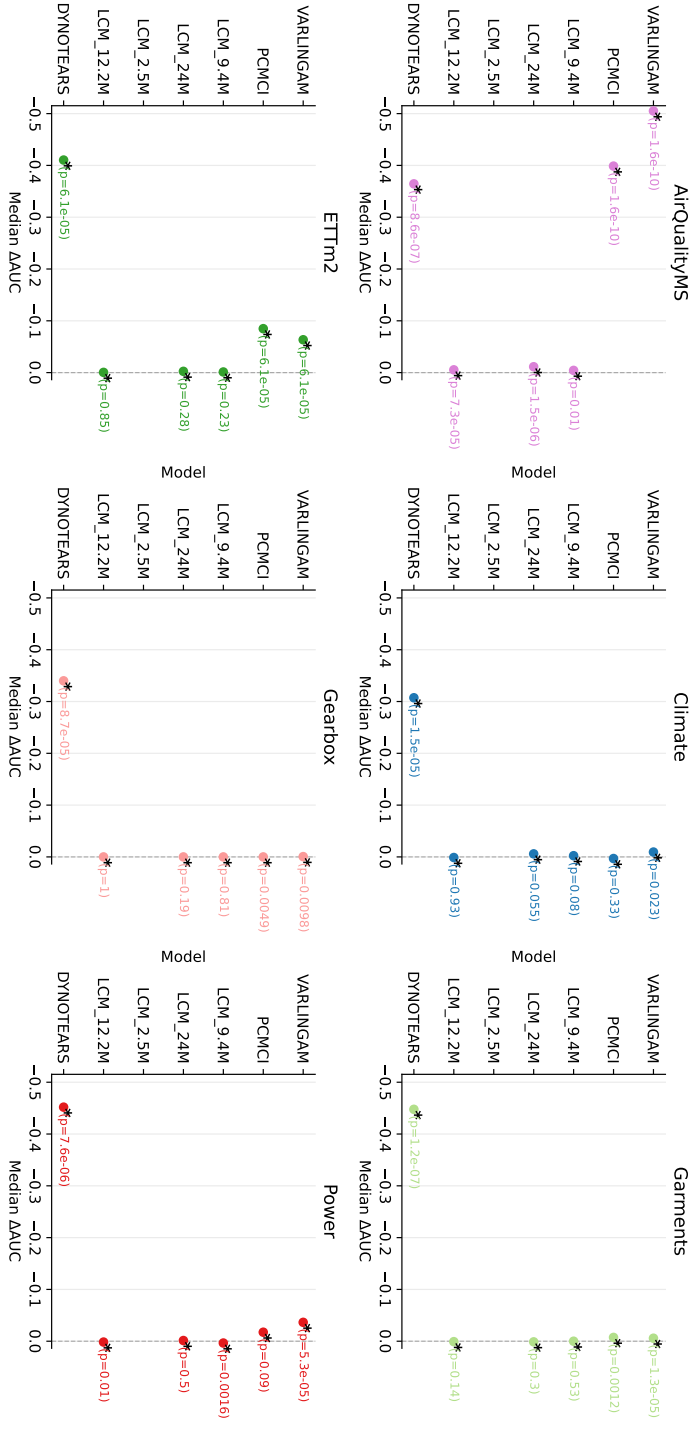

Fig. 11. Paired median AUC differences across *Realistic* collections. Positive values indicate improved performance. Statistical significance is determined via Wilcoxon signed-rank test with Bonferroni correction.