

Causal discovery in urban data with temporal empirical dynamic modeling: The R package *tEDM*[☆]

Wenbo Lyu^{a,b,c}, Yangyang Lei^d, Wen Yi^b, Yongze Song^e, Xiao Li^{f,g,h}, Shaoqing Daiⁱ, Yiming Qin^j, Wufan Zhao^a,^{*}

^a Urban Governance and Design Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

^b Department of Building and Real Estate, The Hong Kong Polytechnic University, Hong Kong, China

^c Shenzhen Research Institute, The Hong Kong Polytechnic University, Shenzhen, China

^d School of Architecture and Design, Harbin Institute of Technology, Harbin, China

^e School of Design and the Built Environment, Curtin University, Perth, Australia

^f Department of Land Surveying and Geo-Informatics, The Hong Kong Polytechnic University, Hong Kong, China

^g Transport Studies Unit, University of Oxford, Oxford, UK

^h Smart Cities Research Institute, The Hong Kong Polytechnic University, Hong Kong, China

ⁱ School of Resource and Environmental Sciences, Wuhan University, Wuhan, China

^j School of the Environment and Sustainable Engineering, Eastern Institute of Technology, Ningbo, China

ARTICLE INFO

Dataset link: https://github.com/SpatLyu/tEDM_CEUS, <https://doi.org/10.5281/zenodo.18247202>, <https://github.com/stscl/tEDM>, <https://doi.org/10.5281/zenodo.18252240>, <https://cran.r-project.org/package=tEDM>

Keywords:

Causal discovery

Empirical dynamic modeling

Time series

Urban data

R package *tEDM*

ABSTRACT

Urban systems are complex, nonlinear, and adaptive, with interactions unfolding across multiple spatial and temporal scales. Traditional correlation-based and regression approaches often fail to capture their dynamic causal structures, limiting understanding of underlying mechanisms. Causal inference provides a more powerful framework, encompassing both causal discovery and causal effect estimation, yet the former remains underexplored in urban studies. Empirical Dynamic Modeling (EDM) is well suited for this task, as it reconstructs system dynamics directly from time series without prespecified models. Despite its potential, EDM has seen limited application in urban contexts. To address this gap, we introduce *tEDM*, an open-source R package that extends EDM to heterogeneous, high-frequency, multivariate, and multi-spatial urban datasets. The package combines a C++ computational backbone with seamless R integration to achieve both efficiency and usability in large-scale urban time series analysis. Its applicability is illustrated through representative case studies on environmental health, carbon-climate dynamics, and epidemic spread, demonstrating how *tEDM* enables nuanced causal discovery and provides methodological as well as practical advances for urban and environmental research. The *tEDM* package itself is publicly available at <https://github.com/stscl/tEDM>.

1. Introduction

Urban systems are inherently complex, comprising interdependent socio-economic, environmental, and infrastructural components that interact across multiple spatial and temporal scales (McPhearson et al., 2016). These interactions are often nonlinear, adaptive, and subject to feedback loops, making them challenging to analyze with traditional correlation-based approaches and linear regression models (Gan et al., 2025), which typically assume fixed functional relationships, linearity, and stationarity. Such assumptions constrain their capacity to capture the dynamic, nonlinear dependencies and feedback mechanisms that characterize real-world urban processes. As a result, much of the current analytical practice in urban science relies on methods that

can identify statistical associations but are not equipped to uncover underlying causal mechanisms.

Causal analysis, in contrast, seeks to move beyond correlation by identifying the drivers and mechanisms that generate observed patterns, thereby providing a stronger basis for effective interventions and informed decision-making (Fekete et al., 2025; Li et al., 2025; Niu & Zhang, 2023). In recent years, it has been successfully applied in urban science to examine the interplay between climate change and human crises (Zhang et al., 2011), assess the impact of air pollution on life expectancy (Chen et al., 2013), investigate the effects of ozone exposure on influenza incidence (Guo et al., 2024), uncover the mechanisms underlying the urban heat island phenomenon (Zhong et al., 2024),

[☆] This article is part of a Special issue entitled: 'Open Urban Data Science' published in Computers, Environment and Urban Systems.

^{*} Corresponding author.

E-mail address: wufanzhao@hkust-gz.edu.cn (W. Zhao).

and explore causality among roads within urban traffic networks (Mao et al., 2024). These examples provide compelling evidence of the transformative potential of causal inference in revealing mechanisms that are not accessible through correlational approaches alone.

Causal inference in quantitative research can be broadly divided into two primary tasks: causal discovery, which seeks to identify the underlying structure of causal relationships among variables (Runge et al., 2019, 2023, 2015), and causal effect estimation, which focuses on quantifying the magnitude of causal impacts once a causal structure is assumed or identified (Pearl, 1995, 2000; Rubin, 1974, 1986). In urban science and related disciplines, causal effect estimation has been widely implemented using quasi-experimental designs such as difference-in-differences (DID), regression discontinuity design (RDD), and instrumental variables, as well as structural equation modeling (SEM) approaches (Bowden et al., 2015; Delgado & Florax, 2015; Kolak & Anselin, 2019; Liu et al., 2021; Rosseel, 2012). These methods have enabled researchers to evaluate the effects of policies, infrastructure investments, and environmental interventions on various urban outcomes (Chen et al., 2013; Jia et al., 2021; Li et al., 2024). In contrast, causal discovery has received comparatively less attention, despite its potential to uncover hidden causal structures in complex urban systems and to inform the design of more robust and targeted effect estimation strategies (Deng et al., 2022; Fekete et al., 2025; Gan et al., 2025).

Among the existing approaches to causal discovery, one prominent category involves predictive models such as Granger causality (Geweke, 1982; Granger, 1969) and its nonlinear extension, transfer entropy (Schreiber, 2000; Vicente et al., 2010), which have been further adapted to multi-spatial observational and nonstationary urban time series (Deng et al., 2022; Gan et al., 2025). Another influential line of methods centers on graphical causal models, such as the Peter and Clark (PC) algorithm, Fast Causal Inference (FCI), and Peter and Clark Momentary Conditional Independence (PCMCI), which rely on conditional independence testing and structural constraints to infer causal structures (Runge et al., 2023; Runge, Nowack, Kretschmer, Flaxman, & Sejdinovic, 2019; Spirtes et al., 1993). A third and increasingly utilized framework is Empirical Dynamic Modeling (EDM), which reconstructs system dynamics from time series data using state-space embedding, enabling causal inference in complex systems without pre-specified model structures (Lorenz, 1969; Mañé, 1981; Sugihara et al., 2012; Takens, 1981). By leveraging the intrinsic properties of dynamic systems, such as nonlinearity and feedback, EDM is particularly well suited for capturing causal interactions in urban systems, which are often temporally evolving, multi-scale, and characterized by complex interdependencies. This makes EDM a compelling tool for advancing causal understanding in urban science.

However, despite its conceptual advantages, the application of EDM for causal discovery in urban contexts remains limited. We argue that harnessing EDM in this domain can uncover causal linkages that are often overlooked by traditional correlation-based or regression approaches, thereby providing deeper insights into the mechanisms governing urban processes. To facilitate such efforts, we introduce *tEDM*, an open-source R package that extends the EDM framework with temporal causality detection capabilities tailored to the heterogeneous, high-frequency, multivariate, and multi-spatial observational nature of urban datasets (Deng et al., 2022; Fekete et al., 2025; Gan et al., 2025). Through three representative case studies, we demonstrate how *tEDM* enables more nuanced and accurate causal discovery in urban time series, contributing both methodological innovation and practical value for urban science.

The remainder of this paper is organized as follows. Section 2 outlines the theoretical foundations of empirical dynamic modeling and its relevance to causal discovery. Section 3 details the design, structure, and core functionalities of the *tEDM* package. In Section 4, we illustrate the practical utility of *tEDM* through three real-world case studies spanning different urban contexts. Finally, Section 5 summarizes the key insights and discusses limitations, as well as directions for future research and development.

2. Background

2.1. Empirical dynamic modeling and dynamic causality

Traditional causality detection in time series analysis often relies on Granger causality, which assumes that if the historical values of variable X can improve the prediction of variable Y , then X is said to Granger-cause Y (Geweke, 1982; Granger, 1969). In practice, Granger causality is typically implemented using linear vector autoregressive models, which implicitly assume that interacting variables are separable. This separability assumption is frequently violated in nonlinear and dynamically coupled systems, limiting the ability of Granger causality to correctly identify causal relationships (Sugihara et al., 2012). To address this limitation, empirical dynamic modeling (EDM) provides a data-driven, model-free framework capable of capturing nonlinear and inseparable causal relationships directly from time series data. EDM reconstructs the underlying system dynamics through time-delay embedding, enabling robust inference of causal relationships without requiring explicit functional forms or assumptions of separability (Leng et al., 2020; Mañé, 1981; Sugihara et al., 2012; Takens, 1981; Tao et al., 2023).

The theoretical foundations of EDM include the generalized embedding theorem and dynamical systems theory. The generalized embedding theorem states that the dynamics of a deterministic system can be reconstructed from time-delayed observations of a single variable (Lorenz, 1969; Mañé, 1981; Takens, 1981), producing a shadow manifold that preserves the topological properties of the original system's attractor. Dynamical systems theory further indicates that if two variables are causally linked, the geometry of their respective shadow manifolds will exhibit a consistent correspondence, allowing states on one manifold to be used to estimate states on the other through mutual cross mapping (Gao et al., 2023; Sugihara et al., 2012). These two processes are jointly illustrated in Fig. 1, which depicts both manifold reconstruction and the cross mapping between shadow manifolds.

Building on this principle, Sugihara et al. (2012) proposed convergent cross mapping (CCM) to detect causality in complex ecological systems. Subsequent developments extended CCM to address additional methodological challenges. Partial cross mapping (PCM) (Leng et al., 2020) was introduced to distinguish direct from indirect causal influences by conditioning on third-party variables, thereby improving the robustness of causal inference in multivariate settings. Cross mapping cardinality (CMC) (Tao et al., 2023) further advances causality detection by shifting the focus from prediction accuracy to the geometric consistency of shadow manifolds, quantified via the intersectional cardinality of cross-mapped neighborhoods.

Beyond methodological refinements, EDM has evolved from an approach limited to time series analysis to one applicable to spatial cross-sectional data (Gao et al., 2023). To address CCM's requirement for relatively long time series, alternative models such as cross map smoothness (CMS) and multispatial convergent cross mapping (MultispatialCCM) have been developed for shorter time series (Clark et al., 2015; Ma et al., 2014). These advancements have broadened EDM's applicability across diverse empirical scenarios. Building upon these developments, the *tEDM* package is specifically designed to address the structural characteristics of urban time series data rather than merely integrating existing EDM methods. In particular, *tEDM* extends CCM to explicitly leverage spatially replicated observations that frequently arise in urban data, incorporating spatial dependence among replicated units to improve causal detection when individual time series are short or when classical CCM fails due to insufficient temporal information. To accommodate complex interdependencies among high-dimensional urban variables, PCM is further generalized to focus on identifying direct causal influences beyond separating indirect effects, enhancing robustness against confounding and collider structures commonly encountered in multivariate urban systems. For causal strength estimation, CMC is augmented with a robust DeLong placement strategy,




Fig. 1. Schematic illustration of empirical dynamic modeling.

enabling statistically consistent estimation of causal strength together with uncertainty quantification through confidence intervals and significance testing, thereby improving the reliability of causal discovery in noisy urban environments.

2.1.1. Convergent cross mapping (CCM)

Convergent cross mapping tests whether the states of a response variable Y can reliably estimate the states of a potential driver X , based on the assumption that if X causally influences Y , then information about X is encoded in the dynamics of Y (Sugihara et al., 2012). Let $Y_{(t)}$ and $X_{(t)}$ be the two time series variable, the shadow manifolds can be constructed using:

$$\begin{aligned} M_x &= [x(t), x(t-\tau), x(t-2\tau), \dots, x(t-(E-1)\tau)] \\ M_y &= [y(t), y(t-\tau), y(t-2\tau), \dots, y(t-(E-1)\tau)] \end{aligned} \quad (1)$$

where τ denotes the time-delay step, E is the embedding dimension, and M_x and M_y are the reconstructed shadow manifolds of the time series X and Y , respectively.

With the reconstructed shadow manifold M_y , the state of X can be estimated from the state of Y through the cross-mapping procedure:

$$\hat{X}_t | M_y = \sum_{i=1}^k (\omega_{ti} X_{ti} | M_y), \quad (2)$$

where t denotes the time point at which the value of X is to be predicted, $\hat{X}_t$ is the predicted value, k is the number of nearest neighbors

in the manifold used for prediction, t_i is the time index of a neighbor, and X_{ti} is the observed value of X at time t_i . The corresponding weight ω_{ti} is defined as

$$\omega_{ti} | M_y = \frac{\text{weight}(\psi(M_y, t_i), \psi(M_y, t))}{\sum_{i=1}^k \text{weight}(\psi(M_y, t_i), \psi(M_y, t))}, \quad (3)$$

where $\psi(M_y, t_i)$ is the state vector at time t_i in the shadow manifold M_y , and $\text{weight}(\cdot, \cdot)$ is a weight function based on the distance metric between two state-space vectors, defined as

$$\text{weight}(\psi(M_y, t_i), \psi(M_y, t)) = \exp\left(-\frac{\text{dis}(\psi(M_y, t_i), \psi(M_y, t))}{\text{dis}(\psi(M_y, t_1), \psi(M_y, t))}\right). \quad (4)$$

Here, $\exp$ denotes the exponential function, and $\text{dis}(\cdot, \cdot)$ represents the average absolute distance between two state vectors in M_y , defined as

$$\text{dis}(\psi(M_y, t_i), \psi(M_y, t)) = \frac{1}{E} \sum_{m=1}^E |\psi_m(M_y, t_i) - \psi_m(M_y, t)|, \quad (5)$$

where $\psi_m(M_y, t_i)$ denotes the m th coordinate of $\psi(M_y, t_i)$.

The cross-mapping skill ρ is quantified by the Pearson correlation coefficient between the observed and predicted values, with the confidence interval of ρ estimated via the z -transformation under the

normality assumption:

$$\rho = \frac{\text{Cov}(X, \hat{X} | M_y)}{\sqrt{\text{Var}(X) \text{Var}(\hat{X} | M_y)}}, \quad (6)$$

$$t = \rho \sqrt{\frac{n-2}{1-\rho^2}}, \quad (7)$$

where n is the number of predicted observations, and

$$z = \frac{1}{2} \ln \left(\frac{1+\rho}{1-\rho} \right). \quad (8)$$

In CCM, the prediction skill ρ is evaluated across different library sizes L , where L is the number of points used to reconstruct the shadow manifold. Causal influence is inferred when ρ exhibits a positive trend with increasing library size L and shows convergence to a stable, statistically significant value as L approaches its maximum. For the causation from Y to X , the converged cross-mapping skill is expressed as

$$\rho_{X \rightarrow Y} = \lim_{L \rightarrow \infty} \text{cor}(X, \hat{X} | M_y), \quad (9)$$

where $\rho_{X \rightarrow Y}$ measures the strength of causal influence from X to Y after convergence.

Motivated by the multispatial CCM method, *tEDM* extends CCM to explicitly accommodate spatially replicated observations. For time series with repeated measurements across spatial units, each replicated series is first embedded independently using Eq. (1). The resulting embedding matrices are then aligned in temporal order and concatenated to form a combined embedding matrix, together with a combined target vector that preserves the one-to-one correspondence between embedded states and observations. Cross mapping prediction is subsequently performed on the combined embedding using Eqs. (2)–(6). The selection of spatially replicated observations participating in state-space reconstruction and cross mapping can be guided by spatial statistical criteria, such as the bandwidth in geographically weighted regression (Brunsdon et al., 1996), or the range parameter estimated from fitted variograms (Webster & Oliver, 2007), to better coordinate and exploit spatial dependence among replicated units. In this setting, the statistical significance of the cross-mapping skill ρ is assessed via a bootstrap procedure by randomly resampling the target series to be predicted, thereby generating an empirical null distribution of ρ for hypothesis testing. This strategy enables CCM to exploit spatial dependence among replicated units to enhance causal detection when individual time series are short or weakly informative, while preserving the dynamical consistency of the reconstructed state space.

2.1.2. Partial cross mapping (PCM)

Partial cross mapping extends the CCM method by accounting for indirect causal influences mediated by third variables. In addition to detecting whether X causally influences Y , PCM seeks to distinguish the *direct* effect of X on Y from any indirect influence transmitted through a mediate variable Z (Leng et al., 2020). Given three time series $X_{(t)}$, $Y_{(t)}$, $Z_{(t)}$, the PCM procedure operates as follows:

First, the shadow manifold M_y of Y is reconstructed, and $\hat{X}$ is obtained from M_y following the standard CCM cross-mapping step in Eq. (2), with weights defined by Eq. (3).

Next, the shadow manifold M_y is used to predict Z in the same manner:

$$\hat{Z}_t | M_y = \sum_{i=1}^k (\omega_{ti} Z_{ti} | M_y). \quad (10)$$

Then, the predicted $\hat{Z}$ is itself used to estimate X via CCM:

$$\hat{X}_t^{(Z)} | M_{\hat{Z}} = \sum_{i=1}^k (\omega_{ti}^{(Z)} X_{ti} | M_{\hat{Z}}), \quad (11)$$

where $M_{\hat{Z}}$ denotes the shadow manifold reconstructed from $\hat{Z}$, and $\omega_{ti}^{(Z)}$ is defined analogously to Eq. (3).

Finally, the causal influence from X to Y is quantified using the *partial correlation* between X and $\hat{X}$ while controlling for $\hat{X}^{(Z)}$:

$$\rho_{X \rightarrow Y|Z} = \frac{\rho_{X, \hat{X}} - \rho_{X, \hat{X}^{(Z)}} \rho_{\hat{X}, \hat{X}^{(Z)}}}{\sqrt{(1 - \rho_{X, \hat{X}^{(Z)}}^2)(1 - \rho_{\hat{X}, \hat{X}^{(Z)}}^2)}}, \quad (12)$$

where $\rho_{A,B}$ denotes the Pearson correlation between variables A and B . A significant $\rho_{X \rightarrow Y|Z}$ suggests a direct causal influence from X to Y after removing the indirect path through Z .

The original PCM formulation can handle more than one potential mediate variable by sequentially considering cumulative indirect effects. For example, when assessing the causal influence of X on Y with two potential mediators Z_1 and Z_2 , one can first predict Z_1 from M_y , then predict Z_2 from $M_{\hat{Z}_1}$, and finally estimate $\hat{X}^{(Z_1, Z_2)}$ from $M_{\hat{Z}_2}$. This procedure captures the causal influence of X on Y transmitted through Z_1 and Z_2 sequentially. Similarly, one can reverse the order — predict Z_2 first, then Z_1 — to account for the alternative indirect pathway.

In *tEDM*, we support this native sequential strategy, while also offering an improved extension referred to as the fully conditioned partial cross mapping. In this formulation, potential third-party variables, whether arising from mediator, collider, or confounding structures or their combinations, are jointly treated as conditioning variables so that the analysis explicitly focuses on direct causal influences. When considering multiple conditioning variables ($Z_1, Z_2, \dots, Z_m$), we first reconstruct M_y and predict each Z_i individually:

$$\hat{Z}_i | M_y = \sum_{j=1}^k (\omega_{ij} Z_{ij} | M_y), \quad i = 1, 2, \dots, m, \quad (13)$$

then use each $\hat{Z}_i$ to predict X :

$$\hat{X}_t^{(Z_i)} | M_{\hat{Z}_i} = \sum_{j=1}^k (\omega_{tj}^{(Z_i)} X_{tj} | M_{\hat{Z}_i}), \quad (14)$$

and finally compute the partial correlation between X and $\hat{X}$ while controlling for all $\hat{X}^{(Z_i)}$:

$$\rho_{X \rightarrow Y|Z_1, \dots, Z_m} = \text{PartialCorr}(X, \hat{X} | \hat{X}^{(Z_1)}, \dots, \hat{X}^{(Z_m)}), \quad (15)$$

where $\text{PartialCorr}(\cdot)$ denotes the standard multivariate partial correlation.

2.1.3. Cross mapping cardinality (CMC)

Cross mapping cardinality evaluates causality based on the structural similarity between the reconstructed shadow manifolds of X and Y . It quantifies how frequently nearest neighbors in one manifold correspond to neighbors in the other, under the assumption that causally linked variables share geometric properties in state space (Tao et al., 2023). Unlike CCM, which focuses on predictive performance, CMC emphasizes topological consistency, offering an alternative metric for assessing the strength of dynamic coupling. Given two time series $X_{(t)}$ and $Y_{(t)}$, we first reconstruct their shadow manifolds M_x and M_y . For each time point i , the intersectional cardinality is defined as the number of shared neighbors between the direct neighborhood of Y and the neighborhood projected via cross mapping from X :

$$I_{C,i,k} = |M_y^{nn}(i, k) \cap M_y^{nn}(M_x^{nn}(i, k), k)|, \quad (16)$$

where $M_x^{nn}(i, k)$ and $M_y^{nn}(i, k)$ denote the sets of the k nearest neighbors of the i th time point in M_x and M_y , respectively. Averaging over all time points yields the global IC value:

$$I_{C,k} = \frac{1}{T} \sum_{i=1}^T I_{C,i,k}, \quad (17)$$

and varying k from 1 to T produces the IC curve:

$$I_{C\text{curve}} = \{I_{C,k} | k = 1, 2, \dots, T\}. \quad (18)$$

To improve the statistical rigor of causal strength (CS) estimation in CMC, we replace the original strategy that directly compares IC curves

Table 1
Open-source toolkits for causal discovery and their applicability to time series.

Package	Language	Implemented methods	Time series ready
<i>dowhy-gcm</i>	Python	Graphical causal modeling	No
<i>causal-learn</i>	Python	Multiple causal discovery algorithms	Partial
<i>cdt</i>	Python	Causal discovery toolbox (hybrid methods)	Partial
<i>lingam</i>	Python	Linear non-Gaussian models	Yes
<i>tigramite</i>	Python	Causal time series analysis via PCMCi and its variants	Yes
<i>pcalg</i>	R	Causal structure learning	Partial
<i>bnlearn</i>	R	Bayesian network structure learning	Partial
<i>lmtest</i>	R	Granger causality tests (GCT)	Yes
<i>NliTS</i>	R	Nonlinear time series causality methods	Yes
<i>rEDM</i> & <i>pyEDM</i>	R & Python	Convergent Cross Mapping (CCM)	Yes
<i>multispatialccm</i>	R	Multispatial Convergent Cross Mapping (MultispatialCCM)	Yes
<i>fastEDM</i>	R & Python	CCM & MultispatialCCM	Yes
<i>crossmapy</i>	Python	Temporal Empirical Dynamic Modeling	Yes
<i>tEDM</i> (Our work)	R	Enhanced Temporal Empirical Dynamic Modeling	Yes

with the DeLong placement method. This non-parametric approach provides statistically consistent estimation of the area under the IC curve (AUC), together with its variance and confidence interval (DeLong et al., 1988; Sun & Xu, 2014). The causal strength from Y to X is therefore quantified by the AUC estimated using the DeLong placement method. Let $\{X_i\}_{i=1}^m$ denote the IC values under the alternative hypothesis H_1 and $\{Y_j\}_{j=1}^n$ denote the IC values under the null hypothesis H_0 . The estimated AUC is given by:

$$\hat{\theta} = \frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n \mathbb{I}(X_i > Y_j) + \frac{1}{2} \mathbb{I}(X_i = Y_j), \quad (19)$$

where $\mathbb{I}(\cdot)$ is the indicator function. The variance of the estimated AUC is

$$\hat{\sigma}^2 = \frac{1}{m} \sum_{i=1}^m (X_i - \hat{\theta})^2 + \frac{1}{n} \sum_{j=1}^n (Y_j - (1 - \hat{\theta}))^2, \quad (20)$$

and the confidence interval is given by

$$\text{CI} = [\hat{\theta} - z_{\alpha/2} \hat{\sigma}, \hat{\theta} + z_{\alpha/2} \hat{\sigma}]. \quad (21)$$

Finally, the statistical significance of the causal strength is assessed via a Z-test:

$$Z = \frac{\hat{\theta} - 0.5}{\hat{\sigma}}, \quad p = 2\Phi(-|Z|), \quad (22)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. A significant p -value indicates that the IC curve under H_1 is significantly different from the null expectation, supporting a causal influence from Y to X .

2.2. Open-source toolkits for time series causal discovery

Causal discovery in time series is a growing area of interest in fields such as econometrics, neuroscience, climate science, and systems biology (Runge et al., 2019, 2023). Numerous open-source toolkits have been developed in both *Python* and *R* to facilitate causal discovery from observational temporal data. These toolkits implement a diverse range of methodologies, including traditional statistical tests such as Granger causality, constraint-based algorithms such as the PC algorithm, score-based and hybrid learning approaches, and more recent techniques derived from empirical dynamic modeling (EDM). Table 1 summarizes a selection of representative toolkits across these categories. Each package is described by its programming language, core methodological focus, and whether it supports direct application to time series data.

Among these, traditional statistical methods are exemplified by the *lmtest* package, which provides implementations of Granger causality tests based on linear autoregressive models (Zeileis & Hothorn, 2002). Constraint-based methods are implemented in *pcalg*, *causal-learn*, and *cdt*, which allow for causal structure learning through conditional independence testing, though they often require additional preprocessing to handle temporal dependencies (Kalainathan et al., 2020; Kalisch et al.,

2012; Zheng et al., 2024). Score-based and hybrid approaches, such as those supported in *bnlearn* and *cdt*, combine structure search with statistical criteria to learn Bayesian networks, but are not inherently designed for time-lagged causality (Kalainathan et al., 2020; Scutari, 2017). Linear non-Gaussian models, as implemented in *lingam*, assume non-Gaussian noise to identify causal directions and are directly applicable to time series under appropriate conditions (Ikeuchi et al., 2023).

Methods designed specifically for causal discovery in time series include *tigramite*, which employs PCMCi and its variants to robustly detect time-lagged causal relations under high-dimensional and auto-correlated conditions (Runge et al., 2023; Runge, Nowack, Kretschmer, Flaxman, & Sejdinovic, 2019). In the context of nonlinear dynamical systems, EDM-based toolkits such as *rEDM*, *pyEDM*, and *crossmapy* implement techniques like convergent cross mapping to infer causality from reconstructed state space (Sugihara et al., 2012; Tao et al., 2023). Packages such as *multispatialccm* and *fastEDM* further extend EDM capabilities by enabling causal inference on short time series with spatial replication, where spatial replicates are leveraged as surrogates for temporal variation, thereby allowing the application of convergent cross mapping even when the available time series data are too short for traditional EDM approaches. (Clark et al., 2015; Li et al., 2021).

Despite these advances, several gaps remain in existing EDM toolkits. Both *rEDM* and *pyEDM* only support convergent cross mapping, limiting their applicability when more nuanced causal inference techniques are required. The *multispatialccm* package incorporates both convergent cross mapping and its multispatial extension, but does not include methods such as partial cross mapping or cross mapping cardinality. The Python library *crossmapy* is the only toolkit that implements convergent cross mapping, partial cross mapping, and cross mapping cardinality simultaneously, yet its interface does not allow users to flexibly specify input parameters. Moreover, its implementation departs from the original definition of convergent cross mapping by returning a single algorithm-determined “optimal” cross mapping ρ value, rather than preserving the core principle that causality is revealed through the convergence of predictive skill with library size. These limitations restrict the flexibility, interpretability, and methodological fidelity of EDM-based causal discovery in complex time series data.

To overcome the limitations of existing empirical dynamic modeling toolkits in terms of flexibility, scalability, and multivariate modeling, we introduce *tEDM*, a novel R package designed for causal discovery in urban time series data. Urban systems are typically characterized by short temporal sequences, multiple spatially replicated observations, and substantial heterogeneity across locations (Gan et al., 2025; Mao et al., 2024). *tEDM* extends and refines core methods, including convergent cross mapping, partial cross mapping, and cross mapping cardinality, to better accommodate these features. In addition, it implements carefully designed numerical computation strategies and efficient algorithms for large-scale datasets. As a result, *tEDM* provides a comprehensive and efficient tool for causal discovery in complex, high-dimensional urban systems, serving as a substantive extension of existing empirical dynamic modeling toolkits.




Fig. 2. Overview of the source code architecture of the *tEDM* package.

3. The *tEDM* package

3.1. Software design philosophy and architecture

The *tEDM* package is developed with an emphasis on computational efficiency, modularity, and tight integration into the R ecosystem for data science and urban analysis (R Core Team, 2026). Fig. 2 presents an overview of the source code organization and architectural layers of the package. At a high level, *tEDM* adopts a layered architecture consisting of a user-facing R interface layer and a high-performance C++ computational backend. The R layer is responsible for user interaction, data preparation, parameter configuration, result formatting, and visualization, whereas all computationally intensive components are implemented in C++ to ensure scalability and runtime efficiency. This separation of concerns improves code maintainability, extensibility, and long-term sustainability of the software.

The computational backend is implemented entirely in C++ and provides the performance required for large-scale time series analysis. Core algorithmic components such as state-space reconstruction, nearest-neighbor search, and cross mapping are implemented using the C++ Standard Template Library (STL). Numerical operations involved in predictive skill estimation for empirical dynamic modeling, including correlation and partial correlation coefficients, are handled by the Armadillo linear algebra library via *RcppArmadillo* (Sanderson & Curtin, 2016). To further improve computational throughput on modern multicore hardware, the backend integrates the *RcppThread* library for multithreaded parallel execution (Nagler, 2021). The C++ backend is exposed to R through the *Rcpp* interface (Eddelbuettel & François, 2011), which enables efficient data exchange while minimizing overhead between the two languages.

At the R interface level, *tEDM* relies on a small set of essential R packages to ensure usability and interoperability. Specifically, *ggplot2* is used for visualization, while *dplyr* supports data manipulation and preprocessing. Together with *Rcpp*, *RcppArmadillo*, and *RcppThread* in the computational layer, these dependencies constitute the mandatory software stack of the package. Core functions accept *data.frame* inputs with explicit column specification and configurable model parameters, and outputs are returned as structured lists with formatted printing in the R console. This design allows seamless integration with existing R-based analytical workflows while maintaining high computational performance and reproducibility.

3.2. Core functionalities

The core functionalities of *tEDM* are summarized in Fig. 3, which illustrates the workflow from data ingestion to causal discovery. The package supports a variety of common data formats, including non-spatial tabular files (*csv*, *tsv*, *xls* and *xlsx*) and spatial datasets that can be pre-processed in R and converted into *data.frame* structures. Once

the data are loaded, the pipeline provides tools to determine the minimal embedding dimension through the *fnn()* function, which applies the false nearest neighbors approach (Kennel et al., 1992). To refine this choice and jointly optimize the embedding dimension E and the number of nearest neighbors k , auxiliary functions such as *simplex()* are employed, with the goal of maximizing predictive skill, or *ic()* to maximize causal strength. These steps establish the foundation for downstream empirical dynamic modeling analyses. A full description of adjustable parameters for each function is available in the package documentation at <https://stsc.github.io/tEDM/reference/index.html>.

After parameter selection, *tEDM* provides several algorithms for causal discovery that quantify and evaluate potential causal relationships. The *ccm()* function applies convergent cross mapping to test for causality by explicitly assessing the convergence of predictive skill across library sizes (Sugihara et al., 2012), thereby retaining the original definition of CCM rather than returning only a single “optimal” strength value as in some existing implementations (the *crossmap* package). When applied to spatially replicated time series, the *multispatialccm()* function incorporates spatial dependencies, as described in Section 2.1.1. The *cmc()* function implements cross mapping cardinality (Tao et al., 2023), and extends prior approaches by internally applying the DeLong placement method to estimate causal strength while also providing corresponding p -values and confidence intervals. This incorporation of formal statistical inference enhances the rigor and interpretability of causal estimates. For cases where direct causation needs to be identified, the *pcm()* function employs partial cross mapping to discriminate direct causal effects from indirect causations (Leng et al., 2020). Beyond reproducing the approach described in the original paper, *tEDM* introduces a new *cumulate* parameter: when set to TRUE, the function follows the original PCM implementation, while the default setting (FALSE) applies our newly defined PCM formulation, thereby offering greater flexibility in handling complex spatiotemporal datasets.

Predictive skill and intersectional cardinality are incorporated as complementary measures to enhance the robustness of inference. Results can be visualized directly using *ggplot2*, which ensures that causal discovery outcomes in urban time series can be examined in a transparent and interpretable manner. These functionalities provide a coherent workflow that integrates data preparation, parameter selection, causal quantification, and visualization within a unified framework. By extending CCM, CMC, and PCM with flexible parameterization, statistical inference, and new formulations, *tEDM* addresses the key gaps of prior EDM toolkits and establishes a comprehensive and rigorous platform for causal discovery in urban systems.

4. Case studies

To demonstrate the practical utility of the *tEDM* package, we present three representative case studies drawn from prior work in




Fig. 3. Schematic representation of the standard workflow for causal discovery using the *tEDM* package.

urban and environmental research, each exemplifying different characteristics of urban time series data and corresponding causal discovery tasks. The first case study investigates the relationship between air pollution and cardiovascular health in Hong Kong (Leng et al., 2020), illustrating how *tEDM* can be employed to analyze multivariate time series and infer the underlying causal interaction network between environmental exposures and health outcomes. The second case study examines county-level carbon emissions and temperature dynamics in the contiguous United States (Gan et al., 2025), showcasing the capability of *tEDM* to uncover causal dependencies within spatially replicated multivariate time series, a data scenario frequently encountered in large scale urban and climate research. The third case study explores the spread of COVID-19 across Japanese prefectures (Tao et al., 2023), demonstrating how *tEDM* can reconstruct causal structures from spatially replicated univariate time series, where repeated temporal observations across locations (spatial units) enable causal inference even in single-variable settings. For transparency and reproducibility, the datasets corresponding to all three case studies are archived within the *tEDM* package, together with complete analysis code available in the package vignette at <https://stsc1.github.io/tEDM/articles/tEDM.html>. These case studies not only highlight the methodological versatility and empirical relevance of *tEDM*, but also provide a foundation for its broader application in advancing causal discovery across diverse domains of urban and environmental research.

4.1. Air pollution and cardiovascular health in Hong Kong

The first case study examines the causal relationships between air pollutants and cardiovascular disease (CVD) incidence in Hong Kong, where the fully conditioned partial cross mapping extension implemented in the *tEDM* package was applied. Fig. 4a presents the

time series of the variables, including daily counts of CVD hospital admissions and the concentrations of SO_2 , NO_2 , O_3 , and respirable suspended particulates (RSP). The dynamics exhibit pronounced seasonal variability together with substantial short-term fluctuations, reflecting the complex interactions between atmospheric processes and population health responses. Fig. 4b reports the cross-mapping skill (ρ) evaluated at the maximum library size for each candidate causal direction. The variation of cross-mapping skill with library size for all tested links is further reported in Supplementary Figure S1, which confirms the convergence behavior and robustness of these estimates. Fig. 4c summarizes the reconstructed causal network linking air pollutants to CVD incidence. The results indicate strong dynamical coupling among the four pollutants, consistent with known atmospheric chemical and physical interactions. In contrast, among the pollutant-to-health links, only NO_2 exhibits a statistically significant and sufficiently large cross-mapping skill toward CVD, suggesting a direct causal influence, whereas the effects of SO_2 , O_3 , and RSP on CVD are likely indirect. Compared with the original PCM study, our implementation in *tEDM* further detects a plausible bidirectional interaction between O_3 and NO_2 , which is consistent with photochemical smog formation and nitric oxide cycling (Ravina et al., 2022). For comparison, the PC algorithm incorrectly identifies a causal influence from CVD to NO_2 and yields an oversimplified internal structure among SO_2 , NO_2 , O_3 , and RSP (see Supplementary Figure S2). This contrast exemplifies the difficulty of recovering reliable causal directions in strongly coupled and nonlinear systems using conditional independence tests alone.

4.2. US county carbon emissions and temperature dynamics

Fig. 5 reports the causal strengths between county-level carbon emissions (CO_2 concentrations) and annual mean temperature (tem)




Fig. 4. Causal relationships between air pollutants and cardiovascular disease incidence in Hong Kong. **a** Time series of the relevant variables. **b** Cross mapping skill (ρ value) among variables estimated via partial cross mapping (values not statistically significant at $p > 0.05$ are marked with [#]). **c** Causal network inferred from panel **b**, where ρ values that are statistically significant and exceed 0.2 are identified as causal influences.

across the contiguous United States, revealing a pronounced asymmetry in the directional causal influence. The causal influence from carbon emissions to temperature (carbon $\rightarrow$ tem) is substantially stronger than the reciprocal pathway (tem $\rightarrow$ carbon). Specifically, the distribution of causal strengths for carbon $\rightarrow$ temperature has quartiles of 0.2099, 0.2515, and 0.2917 (with a full range of 0.0972–0.4151), whereas the corresponding quartiles for temperature $\rightarrow$ carbon are 0.0864, 0.0926, and 0.0988 (range: 0.0648–0.1620). These results indicate that anthropogenic carbon emissions exert a dominant influence on temperature variations, while the reverse effect is comparatively weak.

The broader spread of causal strengths in the carbon $\rightarrow$ temperature direction highlights spatial heterogeneity across U.S. counties, potentially reflecting variations in local climatic sensitivities or emission profiles. By contrast, the consistently small causal strengths in the temperature $\rightarrow$ carbon pathway suggest that this relationship is better interpreted as a secondary feedback effect rather than a primary driver. Notably, these findings are consistent with prior results obtained using the feature-adjusted transfer entropy (FsATE) method (Gan et al., 2025), while the cross mapping cardinality approach applied here to spatially replicated county-level time series offers a more compelling delineation between dominant causal effects and weaker feedback mechanisms. This provides quantitative evidence reinforcing the well-established role of carbon emissions as a principal driver of temperature dynamics in the context of anthropogenic climate change (Madden & Ramanathan, 1980).

Because this case study involves a bivariate dataset with spatially replicated observations, we further benchmarked alternative causal discovery approaches. Granger causality provides little evidence for a causal relationship between carbon emissions and temperature, as most tests fail to reject the null hypothesis with p -values exceeding 0.05 (Supplementary Figure S3). Standard CCM exhibits stronger cross-mapping skill for the carbon $\rightarrow$ temperature direction than for the reverse pathway (Supplementary Figure S4a), but the limited temporal length of individual county-level series prevents statistically stable inference. In contrast, when the *tEDM* multispatial CCM extension is applied by selecting spatially replicated observations using a k -nearest-neighbor strategy ($k = 8$), the causal signal becomes substantially clearer and the carbon $\rightarrow$ temperature influence is successfully recovered (Supplementary Figure S4b). Together, these comparisons indicate that while conventional methods may struggle under short and spatially fragmented time series, the multispatial extensions in *tEDM* effectively exploit spatial replication and spatial dependence to enhance causal detectability, supporting its suitability for urban and environmental applications with dense spatial sampling but limited temporal depth.

4.3. COVID-19 spread across Japanese prefectures

The third case study investigates the dynamics of COVID-19 transmission among Japanese prefectures during the pandemic, utilizing convergent cross mapping to reconstruct the influence of the Tokyo




Fig. 5. Boxplots of causal strengths between annual mean temperature and carbon emissions across contiguous U.S. counties. Causal strengths were estimated using cross mapping cardinality by the *tEDM* package.




Fig. 6. Reconstruction of COVID-19 transmission dynamics from Tokyo to other prefectures in Japan using convergent cross mapping. Results are based on our reanalysis of the original dataset, with visualization style adapted from Tao et al. (2023).

outbreak. As illustrated in Fig. 6, seven prefectures exhibit a causal strength greater than 0.9 from Tokyo, indicating strong transmission linkages. These prefectures, which include both immediate neighbors in the Tokyo metropolitan area and other major urban centers, represent the regions most substantially affected by the Tokyo outbreak. This pattern reflects the hub-like role of Tokyo as the primary source of infection spread, consistent with its status as the earliest and largest

epicenter of confirmed cases in Japan. The result also resonates with prior studies employing cross mapping cardinality approaches, which have shown that causal influence from Tokyo tends to decrease with increasing geodesic distance, while metropolitan areas with higher population density and mobility networks remain particularly susceptible (Tao et al., 2023). For comparison, when the PC algorithm is applied to the same dataset, it only identifies causal transmission links

from Tokyo to Kanagawa and Chiba, while failing to detect influences on other highly connected but geographically more distant metropolitan regions (Supplementary Figure S5), suggesting limited capacity to nonlinear and mobility-driven transmission dynamics. Overall, the reconstructed transmission network recovered by *tEDM* highlights the need for coordinated public health strategies across prefectures to effectively manage and mitigate epidemic spread beyond immediate geographic proximity. These results contribute to a deeper understanding of regional transmission mechanisms and provide actionable insights for future preparedness in large-scale public health crises.

5. Discussion and conclusion

In this study, we have introduced *tEDM*, an open-source R package designed to implement and extend empirical dynamic modeling for temporal causal discovery in complex urban systems. Through three representative case studies, which investigate the relationship between air pollution and cardiovascular health in Hong Kong, the dynamics between county-level carbon emissions and temperature in the United States, and the spread of COVID-19 across Japanese prefectures, we demonstrated the versatility and practical utility of *tEDM* in uncovering dynamic causal linkages from multivariate and spatially replicated time series. Our results highlight the advantages of EDM-based causal inference in capturing nonlinear interactions, feedback mechanisms, and temporally evolving dependencies that are often overlooked by traditional correlation-based or quasi-experimental approaches. The *tEDM* package provides both methodological innovation and actionable insights, offering a reliable toolkit for advancing causal understanding in urban science and related disciplines.

A central premise of *tEDM* is that causality in complex systems should be understood in a dynamical sense, rather than as a static association between variables. Dynamic causality emphasizes whether the evolution of one variable leaves a persistent and directionally identifiable imprint on the state evolution of another variable through the underlying system dynamics (Harnack et al., 2017; Sugihara et al., 2012). Unlike conventional correlation analysis, which quantifies instantaneous or marginal dependence, EDM-based approaches exploit state-space reconstruction to test whether information about a potential driver is embedded in the temporal evolution of a response variable. Cross mapping therefore evaluates directional information recovery from reconstructed shadow manifolds, and the convergence of cross-mapping skill with increasing library size provides an operational criterion to distinguish genuine causal signal from spurious correlation driven by shared trends, synchronization, or external forcing. In this sense, the causal relationships inferred by *tEDM* represent temporal and dynamical causality encoded in system trajectories, rather than merely statistical co-variation.

Within this dynamical causality perspective, the quantities reported as “cross-mapping skill” (the ρ value) in CCM and PCM, and “causal strength” in CMC quantify the influence of data-driven deterministic processes that govern the evolution of the underlying system (Leng et al., 2020; Sugihara et al., 2012; Tao et al., 2023). Specifically, they measure the consistency and reliability with which the state of one variable can be recovered from another in the reconstructed state space, reflecting the extent to which dynamical information is shared directionally between variables. These metrics should be interpreted as relative measures of directional dynamical dependence within a given system, rather than as counterfactual causal effects comparable to quantities such as the average treatment effect in intervention-based causal inference. Consequently, their magnitudes are system-dependent and primarily support causal discovery tasks, including ranking, comparison, and screening of potential causal influences within a dataset, rather than direct transferability across unrelated systems or policy scenarios.

From a practical perspective, different EDM-based methods implemented in *tEDM* are suited to different data characteristics and analytical objectives. Standard CCM is most appropriate when sufficiently

long and high-quality time series are available and when dominant causal pathways are expected to be relatively low-dimensional. PCM is preferable in multivariate systems where indirect effects, confounding structures, or complex interdependencies may obscure direct causal relationships. CMC is advantageous when geometric consistency of reconstructed manifolds provides more stable evidence than prediction accuracy, particularly under noisy or partially observed conditions. In applied analyses, combining evidence from these complementary perspectives often yields more reliable causal interpretation than reliance on a single criterion.

Because the true causal structure is typically unknown in real-world systems, causal discovery inevitably involves uncertainty and the risk of false positives. *tEDM* mitigates these risks through multiple layers of validation, including statistical significance testing, convergence diagnostics with respect to library size, and consistency checks across spatial replications when available. Joint use of multiple EDM methods further improves robustness by requiring agreement across independent causal criteria. In practice, effect magnitude is also considered in addition to statistical significance. For example, in the presented case studies, only statistically significant cross-mapping skill values exceeding a threshold (e.g., $\rho > 0.2$) are interpreted as meaningful causal influence. This criterion reflects the fact that strong coupling, synchronization, or mirror dynamics between variables may produce high reconstruction skill without true causal influence (Gao et al., 2023; Mur, 2013), and thus additional effect-size filtering helps reduce spurious detections. Compared with constraint-based causal discovery algorithms such as PC, EDM-based approaches have intrinsic advantages in handling nonlinear, non-separable, and strongly coupled dynamical systems without requiring explicit model specification. Nevertheless, algorithmic inference alone cannot fully resolve causal ambiguity. Domain knowledge, physical constraints, and contextual interpretation remain essential for validating inferred causal relationships and for translating dynamical causality into actionable scientific or policy insight.

Despite its strengths, the use of *tEDM* for causal discovery is subject to several important limitations that relate to data quality, statistical interpretation, and underlying dynamical assumptions. The reliability of state-space reconstruction depends critically on the presence of sufficient deterministic structure in the observed system. In empirical settings, short or irregular time series (Clark et al., 2015; Ma et al., 2014), measurement noise (Monster et al., 2017), and partial observability can degrade the fidelity of the reconstructed attractor and, consequently, the stability of cross-mapping estimates. More fundamentally, the causal relationships identified by *tEDM* are defined with respect to the specified set of observed variables. Inferred directional influences are therefore conditional on the chosen covariate set and may change as additional relevant variables are incorporated. In this sense, the method operates at the level of observational causal discovery and does not, by itself, resolve ambiguities arising from incomplete system specification.

Second, the statistical interpretation of EDM-based causal measures requires careful qualification. Significance testing in CCM, PCM, and related approaches is typically formulated against a null hypothesis of no causal association (Tao et al., 2023). However, in complex systems, the assumption of complete causal independence is rarely tenable (Meehl, 1990), as system components are often entangled through complex interactions and shared latent drivers. Consequently, rejection of the null should be understood as evidence against specific null-generating processes rather than as confirmation of a well-defined causal mechanism. In this context, cross-mapping skill and related metrics are more appropriately interpreted as measures of directional dynamical dependence within the observed system, rather than as estimates of causal effects in the interventionist sense. For this reason, robust inference in *tEDM* relies on the joint consideration of convergence behavior, effect size, consistency across methods, and compatibility with domain knowledge, rather than on statistical significance alone. In addition, evaluating the performance of EDM-based causal discovery under controlled or

simulated systems remains important for understanding its statistical power and failure modes, particularly in settings characterized by noise, limited observations, or model misspecification.

Third, the applicability of *tEDM* is constrained by assumptions from delay embedding theory and, more broadly, by the extent to which the system admits a coherent dynamical representation. The framework builds upon results such as the Takens–Mañé theorem (Mañé, 1981; Takens, 1981), which apply to deterministic and (approximately) autonomous systems evolving on a stable attractor (Leng et al., 2020). In practice, this implies that the method is most reliable in settings where the underlying dynamics are governed by relatively stable mechanisms and where the effective dimensionality of the system is limited. Such conditions are more likely to be satisfied in domains where prior knowledge, including physical or mechanistic constraints, provides guidance on plausible system structure and variable selection. In contrast, for systems with rapidly changing dynamics, strong exogenous forcing, or poorly understood coupling mechanisms, causal interpretations based on reconstructed state spaces become less stable (Butler et al., 2023; Park et al., 2023).

These limitations are not specific to individual implementations, but reflect the fundamental constraints of EDM-based causal inference. Consequently, the methodological extensions introduced in *tEDM*, including multispatial embedding, generalized PCM conditioning, and enhanced CMC-based inference, remain subject to the same assumptions and potential sources of bias. In addition, these extensions introduce their own practical considerations. For example, multispatial embeddings require a sufficient degree of comparability and dynamical consistency across replicated units, PCM-based conditioning depends on the inclusion of an adequately informative set of variables, and the robustness of CMC-based inference remains sensitive to reconstruction quality and sampling variability. While these extensions improve applicability under realistic data conditions, they do not eliminate challenges arising from noise, incomplete system specification, or limited observability.

From the perspective of the causal hierarchy (Pearl & Mackenzie, 2018), the methods implemented in *tEDM* primarily operate at the level of observational causal discovery, identifying directional dynamical dependencies without explicit intervention or counterfactual structure. Recent work has begun to address some of these limitations by developing diagnostic tools to assess whether data satisfy the assumptions required for EDM-based inference (Butler et al., 2023), as well as by exploring extensions that improve causal identification in the presence of incomplete observations or latent confounding (Yan et al., 2026). These developments suggest that the scope of EDM-based causal analysis is actively evolving, particularly with respect to assessing applicability conditions and strengthening robustness under realistic data constraints. At the same time, extending these approaches toward interventional and counterfactual settings would require additional structural assumptions or integration with complementary causal frameworks. In this context, *tEDM* is best understood as a principled tool for observational causal discovery that can support hypothesis generation and guide subsequent analysis, rather than as a standalone solution for fully identified causal effects.

Building on these methodological considerations, the long-term value of *tEDM* lies not only in its current capabilities but also in its continued development as an extensible and research-oriented platform. The sustainability of the *tEDM* package is ensured through its open-source development model and adherence to best practices in research software engineering (Boettiger et al., 2015). The source code is maintained under version control on GitHub, where continuous integration workflows enable automated testing, code review, and platform-independent reproducibility. Stable releases are regularly archived on CRAN, which guarantees long-term accessibility, standardized distribution, and interoperability within the broader R ecosystem (R Core Team, 2026). In addition, sustainability is supported

by a community-driven governance model, where issues, feature requests, and contributions are openly tracked and managed, and by the active involvement of the core development team in ongoing maintenance. These practices ensure that *tEDM* remains a reliable, extensible, and transparent software platform.

The development roadmap for *tEDM* emphasizes methodological, computational, and usability advances. On the methodological side, future releases will incorporate extensions of empirical dynamic modeling, including adaptive cross mapping strategies and integration with network-based representations of urban dynamics. Computational advances will focus on GPU acceleration, parallelization, and distributed implementations in order to accommodate the rapidly increasing volume of urban sensing data. From a usability perspective, the package will be expanded with richer vignettes, reproducible workflow templates, and advanced visualization capabilities to further lower barriers to adoption while maintaining analytical rigor. These developments will further position *tEDM* as a flexible platform for causal discovery in complex urban systems.

CRediT authorship contribution statement

Wenbo Lyu: Writing – original draft, Validation, Software, Methodology, Conceptualization. **Yangyang Lei:** Writing – original draft, Visualization, Software, Conceptualization. **Wen Yi:** Writing – review & editing, Methodology, Conceptualization. **Yongze Song:** Writing – review & editing, Validation, Software, Methodology. **Xiao Li:** Writing – review & editing, Methodology, Conceptualization. **Shaoqing Dai:** Software, Conceptualization. **Yiming Qin:** Writing – review & editing, Visualization, Data curation. **Wufan Zhao:** Writing – review & editing, Supervision, Software, Resources, Project administration.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (42401567, 42405095, 72201229, 72361137006), AI Research and Learning Base of Urban Culture, Guangdong Provincial Department of Education (2023WZJD008), Guangdong Provincial Project (2024QN11G095), Tertiary Education Scientific research project of Guangzhou Municipal Education Bureau (2024312159), and Guangzhou Municipal Science and Technology Bureau Program (2025A03J3640).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.compenvurbysys.2026.102435>.

Data availability

The datasets and analysis scripts used in the case studies are openly available in the GitHub repository at https://github.com/SpatLyu/tEDM_CEUS. An archived and versioned release of this repository is available via Zenodo with DOI: <https://doi.org/10.5281/zenodo.18247202>. The *tEDM* package itself is released under the GPL-3 license, with source code hosted on GitHub at <https://github.com/stscl/tEDM>. A corresponding archived release of the package source code is also available via Zenodo with DOI: <https://doi.org/10.5281/zenodo.18252240>, ensuring long-term accessibility and reproducibility of the software version referenced in this study. Stable versions of the package are additionally distributed through CRAN at <https://cran.r-project.org/package=tEDM>.

References

- Boettiger, C., Chamberlain, S., Hart, E., & Ram, K. (2015). Building software, building community: Lessons from the rOpenSci project. *Journal of Open Research Software*, 3(1), 8.
- Bowden, J., Davey Smith, G., & Burgess, S. (2015). Mendelian randomization with invalid instruments: effect estimation and bias detection through Egger regression. *International Journal of Epidemiology*, 44(2), 512–525.
- Brunsdon, C., Fotheringham, A. S., & Charlton, M. E. (1996). Geographically weighted regression: A method for exploring spatial nonstationarity. *Geographical Analysis*, 28(4), 281–298.
- Butler, K., Feng, G., & Djurić, P. M. (2023). On causal discovery with convergent cross mapping. *IEEE Transactions on Signal Processing*, 71, 2595–2607.
- Chen, Y., Ebenstein, A., Greenstone, M., & Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy. *Proceedings of the National Academy of Sciences*, 110(32), 12936–12941.
- Clark, A. T., Ye, H., Isbell, F., Deyle, E. R., Cowles, J., Tilman, G. D., & Sugihara, G. (2015). Spatial convergent cross mapping to detect causal relationships from short time series. *Ecology*, 96(5), 1174–1181.
- Delgado, M. S., & Florax, R. J. (2015). Difference-in-differences techniques for spatial data: Local autocorrelation and spatial interaction. *SSRN Electronic Journal*.
- DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, 44(3), 837–845.
- Deng, Z., Weng, D., Xie, X., Bao, J., Zheng, Y., Xu, M., Chen, W., & Wu, Y. (2022). Compass: Towards better causal analysis of urban time series. *IEEE Transactions on Visualization and Computer Graphics*, 28(1), 1051–1061.
- Eddelbuettel, D., & François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40(8), 1–18.
- Fekete, T., Mengistu, G., & Wicaksono, H. (2025). Leveraging causal AI to uncover the dynamics in sustainable urban transport: A bike sharing time-series study. *Sustainable Cities and Society*, 122, Article 106240.
- Gan, T., Succar, R., Macri, S., Marín, M. R., & Porfiri, M. (2025). Causal discovery from city data, where urban scaling meets information theory. *Cities*, 162, Article 105980.
- Gao, B., Yang, J., Chen, Z., Sugihara, G., Li, M., Stein, A., Kwan, M. P., & Wang, J. (2023). Causal inference from cross-sectional earth system data with geographical convergent cross mapping. *Nature Communications*, 14(1).
- Geweke, J. (1982). Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association*, 77(378), 304.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424.
- Guo, F., Zhang, P., Do, V., Runge, J., Zhang, K., Han, Z., Deng, S., Lin, H., Ali, S. T., Chen, R., Guo, Y., & Tian, L. (2024). Ozone as an environmental driver of influenza. *Nature Communications*, 15(1).
- Harnack, D., Laminski, E., Schünemann, M., & Pawelzik, K. R. (2017). Topological causality in dynamical systems. *Physical Review Letters*, 119(9).
- Ikeuchi, T., Ide, M., Zeng, Y., Maeda, T. N., & Shimizu, S. (2023). Python package for causal discovery based on LiNGAM. *Journal of Machine Learning Research*, 24(14), 1–8.
- Jia, R., Shao, S., & Yang, L. (2021). High-speed rail and CO2 emissions in urban China: A spatial difference-in-differences approach. *Energy Economics*, 99, Article 105271.
- Kalinathan, D., Goudet, O., & Dutta, R. (2020). Causal discovery toolbox: Uncovering causal relationships in Python. *Journal of Machine Learning Research*, 21(37), 1–5.
- Kalisch, M., Mächler, M., Colombo, D., Maathuis, M. H., & Bühlmann, P. (2012). Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11), 1–26.
- Kennel, M. B., Brown, R., & Abarbanel, H. D. I. (1992). Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Physical Review A*, 45(6), 3403–3411.
- Kolac, M., & Anselin, L. (2019). A spatial perspective on the econometrics of program evaluation. *International Regional Science Review*, 43(1–2), 128–153.
- Leng, S., Ma, H., Kurths, J., Lai, Y. C., Lin, W., Aihara, K., & Chen, L. (2020). Partial cross mapping eliminates indirect causal influences. *Nature Communications*, 11(1).
- Li, G., Fang, C., Watson, J. E. M., Sun, S., Qi, W., Wang, Z., & Liu, J. (2024). Mixed effectiveness of global protected areas in resisting habitat loss. *Nature Communications*, 15(1).
- Li, X., Zhan, W., Deng, F., Liang, X., & Luo, P. (2025). Causal discovery and analysis of global city carbon emissions based on data-driven and hybrid intelligence. *Computers, Environment and Urban Systems*, 115, Article 102206.
- Li, J., Zyphur, M. J., Sugihara, G., & Laub, P. J. (2021). Beyond linearity, stability, and equilibrium: The edm package for empirical dynamic modeling and convergent cross-mapping in stata. *The Stata Journal: Promoting Communications on Statistics and Stata*, 21(1), 220–258.
- Liu, T., Ungar, L., & Kording, K. (2021). Quantifying causality in data science with quasi-experiments. *Nature Computational Science*, 1(1), 24–32.
- Lorenz, E. N. (1969). Atmospheric predictability as revealed by naturally occurring analogues. *Journal of the Atmospheric Sciences*, 26(4), 636–646.
- Ma, H., Aihara, K., & Chen, L. (2014). Detecting causality from nonlinear dynamics with short-term time series. *Scientific Reports*, 4(1).
- Madden, R. A., & Ramanathan, V. (1980). Detecting climate change due to increasing carbon dioxide. *Science*, 209(4458), 763–768.
- Mañé, R. (1981). On the dimension of the compact invariant sets of certain nonlinear maps. In D. A. Rand, & L. S. Young (Eds.), *Lecture notes in mathematics: Vol. 898, Dynamical systems and turbulence* (pp. 230–242).
- Mao, J., Huang, H., Gu, Y., Lu, W., Tang, T., & Ding, F. (2024). A convergent cross-mapping approach for unveiling congestion spatial causality in urban traffic networks. *Computer-Aided Civil and Infrastructure Engineering*, 40(3), 301–322.
- McPhearson, T., Haase, D., Kabisch, N., & Gren, Å. (2016). Advancing understanding of the complex nature of urban systems. *Ecological Indicators*, 70, 566–573.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66(1), 195–244.
- Mønster, D., Fusaroli, R., Tylén, K., Roepstorff, A., & Sherson, J. F. (2017). Causal inference from noisy time-series data — Testing the convergent cross-mapping algorithm in the presence of noise and external influence. *Future Generation Computer Systems*, 73, 52–62.
- Mur, J. (2013). Causality, uncertainty and identification: Three issues on the spatial econometrics agenda. *Scienze Regionali: Italian Journal of Regional Science*, 12, 1, 2013, 5–27.
- Nagler, T. (2021). R-friendly multi-threading in C++. *Journal of Statistical Software, Code Snippets*, 97(1), 1–18.
- Niu, C., & Zhang, W. (2023). Causal effects of mobility intervention policies on intracity flows during the COVID-19 pandemic: The moderating role of zonal locations in the transportation networks. *Computers, Environment and Urban Systems*, 102, Article 101957.
- Park, S. H., Ha, S., & Kim, J. K. (2023). A general model-based causal inference method overcomes the curse of synchrony and indirect effect. *Nature Communications*, 14(1).
- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4), 669–688.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge University Press.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. New York, NY: Basic Books.
- R Core Team (2026). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing, URL <https://www.R-project.org/>.
- Ravina, M., Caramitti, G., Panepinto, D., & Zanetti, M. (2022). Air quality and photochemical reactions: analysis of NOx and NO2 concentrations in the urban area of Turin, Italy. *Air Quality, Atmosphere & Health*, 15(3), 541–558.
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2), 1–36.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688–701.
- Rubin, D. B. (1986). Statistics and causal inference: Comment: Which ifs have causal answers. *Journal of the American Statistical Association*, 81(396), 961.
- Runge, J., Bathiany, S., Bollt, E., Camps-Valls, G., Coumou, D., Deyle, E., Glymour, C., Kretschmer, M., Mahecha, M. D., Muñoz-Marí, J., van Nes, E. H., Peters, J., Quax, R., Reichstein, M., Scheffer, M., Schölkopf, B., Spirtes, P., Sugihara, G., Sun, J., ... Zscheischler, J. (2019). Inferring causation from time series in Earth system sciences. *Nature Communications*, 10(1).
- Runge, J., Gerhardus, A., Varando, G., Eyring, V., & Camps-Valls, G. (2023). Causal inference for time series. *Nature Reviews Earth & Environment*, 4(7), 487–505.
- Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., & Sejdinovic, D. (2019). Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11).
- Runge, J., Petoukhov, V., Donges, J. F., Hlinka, J., Jajcay, N., Vejmelka, M., Hartman, D., Marwan, N., Paluš, M., & Kurths, J. (2015). Identifying causal gateways and mediators in complex spatio-temporal systems. *Nature Communications*, 6(1).
- Sanderson, C., & Curtin, R. (2016). Armadillo: a template-based C++ library for linear algebra. *The Journal of Open Source Software*, 1(2), 26.
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–464.
- Scutari, M. (2017). Bayesian network constraint-based structure learning algorithms: Parallel and optimized implementations in the bnlearn R package. *Journal of Statistical Software*, 77(2), 1–20.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. Springer New York.
- Sugihara, G., May, R., Ye, H., Hsieh, C. h., Deyle, E., Fogarty, M., & Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338(6106), 496–500.
- Sun, X., & Xu, W. (2014). Fast implementation of DeLong's algorithm for comparing the areas under correlated receiver operating characteristic curves. *IEEE Signal Processing Letters*, 21(11), 1389–1393.
- Takens, F. (1981). Detecting strange attractors in turbulence. In *Lecture notes in mathematics* (pp. 366–381). Springer Berlin Heidelberg.
- Tao, P., Wang, Q., Shi, J., Hao, X., Liu, X., Min, B., Zhang, Y., Li, C., Cui, H., & Chen, L. (2023). Detecting dynamical causality by intersection cardinal concavity. *Fundamental Research*.

- Vicente, R., Wibral, M., Lindner, M., & Pipa, G. (2010). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30(1), 45–67.
- Webster, R., & Oliver, M. A. (2007). *Geostatistics for environmental scientists*. Wiley.
- Yan, J., Zhang, S. W., Zhang, C., Huang, W., Shi, J., & Chen, L. (2026). Dynamical causality under latent confounders for biological network reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–18.
- Zeileis, A., & Hothorn, T. (2002). Diagnostic checking in regression relationships. *R News*, 2(3), 7–10.
- Zhang, D. D., Lee, H. F., Wang, C., Li, B., Pei, Q., Zhang, J., & An, Y. (2011). The causality analysis of climate change and large-scale human crisis. *Proceedings of the National Academy of Sciences*, 108(42), 17296–17301.
- Zheng, Y., Huang, B., Chen, W., Ramsey, J., Gong, M., Cai, R., Shimizu, S., Spirtes, P., & Zhang, K. (2024). Causal-learn: Causal discovery in python. *Journal of Machine Learning Research*, 25(60), 1–8.
- Zhong, Y., Li, S., Liang, X., & Guan, Q. (2024). Causal inference of urban heat island effect and its spatial heterogeneity: A case study of wuhan, China. *Sustainable Cities and Society*, 115, Article 105850.