

Compass: Towards Better Causal Analysis of Urban Time Series

Zikun Deng, Di Weng, Xiao Xie, Jie Bao, Yu Zheng, Mingliang Xu, Wei Chen, and Yingcai Wu

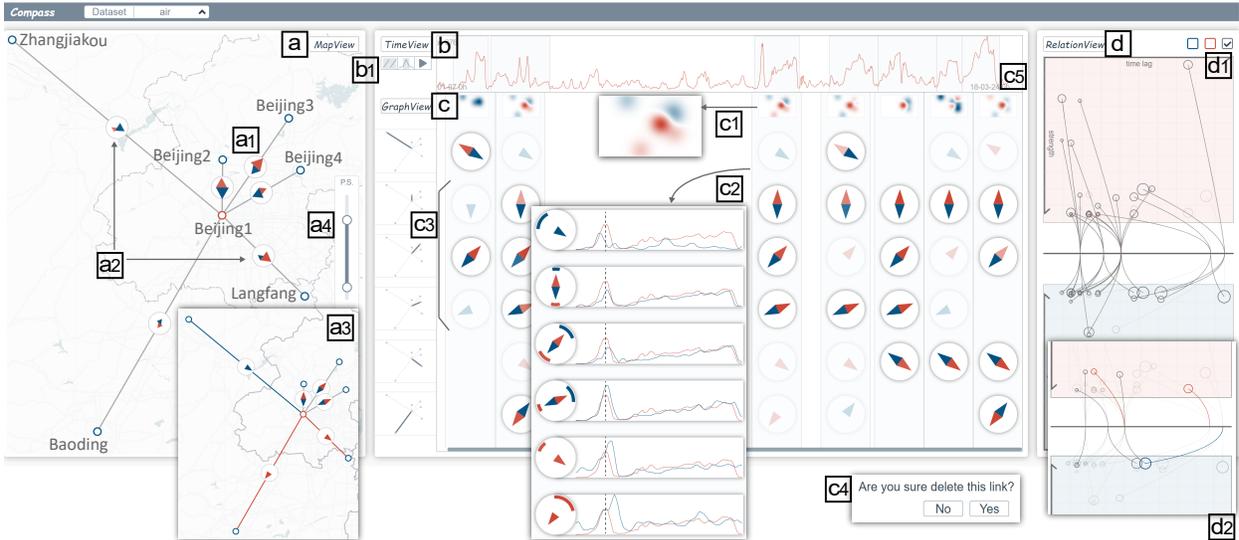


Fig. 1. The interface of Compass. (a) The map view enables users to select a sensor to start analysis and reason causal graphs within a spatial context. (b) The time view presents the time series of the selected sensor and time partitioning results. (c) The graph view visualizes the dynamic causal graphs detected by a causal detection framework along the same timeline of the time view. (d) The relation view presents the causal relations involved in the causal graphs with a multi-dimensional visualization.

Abstract— The spatial time series generated by city sensors allow us to observe urban phenomena like environmental pollution and traffic congestion at an unprecedented scale. However, recovering causal relations from these observations to explain the sources of urban phenomena remains a challenging task because these causal relations tend to be time-varying and demand proper time series partitioning for effective analyses. The prior approaches extract one causal graph given long-time observations, which cannot be directly applied to capturing, interpreting, and validating dynamic urban causality. This paper presents Compass, a novel visual analytics approach for in-depth analyses of the dynamic causality in urban time series. To develop Compass, we identify and address three challenges: detecting urban causality, interpreting dynamic causal relations, and unveiling suspicious causal relations. First, multiple causal graphs over time among urban time series are obtained with a causal detection framework extended from the Granger causality test. Then, a dynamic causal graph visualization is designed to reveal the time-varying causal relations across these causal graphs and facilitate the exploration of the graphs along the time. Finally, a tailored multi-dimensional visualization is developed to support the identification of spurious causal relations, thereby improving the reliability of causal analyses. The effectiveness of Compass is evaluated with two case studies conducted on the real-world urban datasets, including the air pollution and traffic speed datasets, and positive feedback was received from domain experts.

Index Terms—Visual causal analysis, urban time series, causal graph analysis.

1 INTRODUCTION

Detecting and interpreting the causal relations among urban time series have important implications for various urban applications. Such causal

relations can reveal the latent temporal interactions between the sensors where the urban time series are collected, facilitating the effective sense-making of urban phenomena and the informed policy-making for better sustainability [86]. For example, if the pollutant readings at sensors *A* and *B* are identified to *cause* a considerable increase in those at sensor *C* consistently, proper measures should be taken to control the pollution at *A* and *B* rather than *C*.

Many efforts have been devoted to the efficient extraction and extensive analysis of the co-occurrence [15, 36, 75] and correlation relations [45, 63] among urban time series or events. However, the patterns and insights derived in these studies do not imply true causality, and thus the reliability of the results is often limited in applications. Recent studies [38, 89] have investigated applying the Granger causality test to capture the causal relations among urban time series. Based on a set of urban time series, these studies attempt to extract a causal graph, where each cause-effect relation between a pair of sensors is represented with a directed edge. However, one causal graph is insufficient to characterize the dynamic causal relations among multiple urban time series due to the rapidly-changing urban environments. For example, a causal

- Z. Deng, D. Weng, W. Chen, Y. Wu are with the State Key Lab of CAD & CG, Zhejiang University, Hangzhou and Zhejiang Lab, Hangzhou, China. E-mail: {zikun_rain, dweng, ycwu}@zju.edu.cn, chenwei@cad.zju.edu.cn. Y. Wu is the corresponding author.
- X. Xiao is with Department of Sport Science, Zhejiang University, Hangzhou, China. E-mail: xxie@zju.edu.cn.
- J. Bao and Y. Zheng are with JD Intelligent Cities Research, JD Tech, Beijing, China. E-mail: baojie@jd.com, msyuzheng@outlook.com.
- M. Xu is with School of Information Engineering, Zhengzhou University, Zhengzhou, China and Henan Institute of Advanced Technology, Zhengzhou University, Zhengzhou, China. E-mail: iexumingliang@zzu.edu.cn.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

relation between two sensors detected based on their time series can disappear or even be flipped from time to time (Fig. 2a) because of certain external factors, such as the wind fields in the air pollution scenarios. Such a dynamic nature demands fine-grained causality analyses that enable analysts to gain insights into the temporal variations of causal relations in urban contexts. In addition to the limitation of granularity in the prior studies, interpreting and validating these causal relations detected by automated models also require an interactive system to integrate analysts in the urban causality analysis loop [76].

Motivated by the limitations in the prior studies, we propose a visual analytics approach that empowers experts to retrieve and comprehend the dynamic causal relations among urban time series with tailored interactions and visualizations. Developing such an approach poses three major challenges:

Detecting fine-grained urban causality. First, time series should be appropriately partitioned, such as by considering their periodicity and peak occurrences. Second, the causalities detected in every time partition must be stable and not disturbed by noise or dynamics. Finally, the time series’s spatial proximity should be carefully incorporated, thereby improving the reliability of causality detected in urban environments. Addressing the three issues demands an extension of the Granger causality test to make it practical for dynamic urban causality.

Interpreting dynamic causal relations. Gaining insights demands the multifaceted analysis of the dynamic causal relations. The time-varying causal directions between sensors need to be revealed for understanding the dynamics or stability. The causal relations also constitute causal graphs over time. These graphs should be interpreted from spatial and temporal perspectives to uncover how the involved sensors influence each other. Enabling the comprehensive interpretation of such causal relations constitutes the second challenge.

Unveiling suspicious causal relations. Unreasonable causalities may inevitably exist. For example, two remote air quality sensors should not have bi-directional causality, but two close sensors can because air pollutants likely transport back and forth between them. Determining and correcting these suspicious relations requires investigating their multiple dimensions, such as the spatial distance, direction, time lag, and correlation of the time series. A multidimensional visualization tailored for the causal relations should be proposed.

In this study, we first conclude an analytical workflow and six user requirements through careful literature review and close collaboration with experts. Based on the requirements, we further develop Compass, a visual analytics system for the causal analysis of urban time series. Compass addresses the aforementioned three challenges as follows. For the first challenge, we propose a causal detection framework by extending the Granger causality test. This framework includes neighbor index, time partition, and debounce strategies to improve the capability of the causality test for urban time series. For the second challenge, we design a dynamic graph visualization to reveal the temporal variation of the detected causal relations and enable time-oriented drill-down analyses of causal graphs. For the third challenge, we consider the multiple dimensions of causal relations and propose a tailored visualization to expose causality suspiciousness. Finally, two real-world case studies on air pollution and traffic speed datasets, respectively, demonstrate the effectiveness of our approach. In sum, our contributions are as follows:

- ◊ We propose a causal detection framework for urban time series by extending the Granger causality test.
- ◊ We design and implement a novel visual analytics system called Compass for causal analysis of urban time series. Compass incorporates the causal detection framework and a set of effective visualizations for analyzing dynamic causalities.
- ◊ We evaluate our approach through two real-world case studies.

2 RELATED WORK

This section presents relevant studies on causal detection, visual causal analytics, and visual urban analytics.

2.1 Causal Detection

The traditional causal analysis method is controlled variable experiments, e.g., A/B testings in business companies. However, such experiments require stakeholders to test numerous variable pairs tediously and identify the causal relations among variables. By contrast, causal

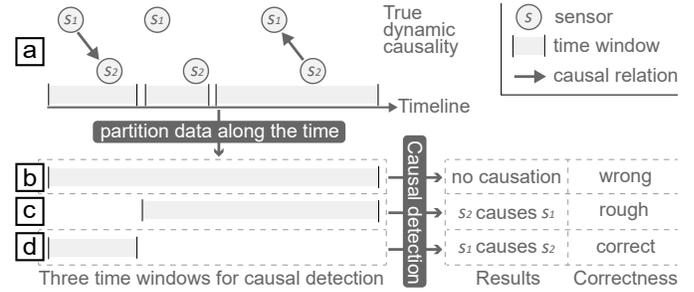


Fig. 2. Motivation illustration. (a) Real-world causal relations between sensors can change over time, shown by the different relations in three time windows. (b, c) Causal detection in large time windows may produce wrong or rough results. (d) Although the result is correct, further interpretation and verification are still required for informed policy making.

detection is a data-driven approach to obtain the causal relations among observed data without setting experimental conditions deliberately.

Correlation is not causation. We focus on those approaches for detecting causality rather than correlation [52, 56]. Many approaches are developed for time series data [19, 20, 54], such as based on information theory [29], conditional independence tests [57], and the Granger causality test [26]. Among these methods, the Granger causality test is the most well-established and has been widely applied to various domains, such as economics [25], geoscience [44], and neuroscience [59].

Granger causality is not limited to traditional time series. To detect the causal relations among events, Jin et al. [31] and Xu et al. [79] transformed event sequences into time series based on Hawkes processes [28] and applied the Granger causality test to these series. Some researchers extended the Granger causality test into urban time series that are associated with geographic positions. Li et al. [38] extracted upstream events from air pollution time series and constructed the causal graphs among these events based on the Granger causality test. Frequent subgraph mining was then applied to extract the propagation pattern of air pollution. Zhu et al. [89] integrated the Granger causality test into Bayesian learning to identify the causal networks of air pollution sensors. However, none of the existing methods can accurately detect the dynamic causality of the urban time series, let alone interpret and validate the causality.

We choose the Granger causality test because of its proven and powerful capability to detect causality based on short partitioned time series. It was extended and integrated into our visual analytics approach.

2.2 Visual Causal Analytics

Causal analysis needs effective causality representations. Visualization studies have investigated causality perception. Xiong et al. [78] and Yen et al. [81] evaluated the performance of some basic charts (e.g., line and bar charts) in conveying causality. These basic charts are not suitable for causal graphs. Bae et al. [6, 7] derived several design guidelines for the graphical representation of causality. These guidelines help us design effective causal visualizations.

Some visualization techniques enable users to reason causality with their knowledge. They all adopted node-link diagrams, a popular representation for relations [85, 87]. Nodes denote variables and links denote causal relations. In the early years, researchers paid more attention to the representation of causality. Elmqvist and Tsigas [21, 22] studied animation techniques and node representations for the causality of event sequences. Also studying event sequences, Kadaba et al. [33] attempted to reveal the causal links in causal graphs from multiple aspects. Recently, many studies have developed visual analytics systems coordinating individual visualizations to make causal analysis practical. Wang and Mueller [64] developed an interactive visual interface for the causal analysis of tabular data. Users can interactively access and verify the causal relations detected by an automatic model. They also further extended their method to allow causal analysis for different data subgroups [65]. Xie et al. [76] proposed a visual analytics approach that addressed the scalability and uncertainty issues of the causal graphs. Jin et al. [31] proposed a visual analytics approach for exploring, verifying, and comparing the causalities in event sequence datasets. In addition,

there were some domain-specific applications. Dang et al. [16] developed a tool for understanding the causality within biological pathways. Pi et al. [53] developed a causal reasoning system based on traffic flow theory to analyze the causality of traffic congestion events.

The existing visualizations and systems cannot accommodate dynamic causal relations. Besides, visual causal analysis of time series has not been well studied. We propose a visual analytics approach for analyzing the dynamic causal relations detected from urban time series.

2.3 Visual Urban Analytics

Visual analytics approaches [4] have been proven effective in leveraging various urban data (e.g., human mobility [2, 3, 58, 60], social media [13, 74], business [41, 71], energy [84] and environmental [17, 55] data) to improve cities [72, 73, 75]. We classified these approaches into two categories, namely, visual urban planning and diagnosis.

Visual urban planning focuses on decision making for improved cities [23, 37, 43, 50], e.g., by optimizing the arrangement of facilities. For example, Liu et al. [40] combined trajectory mining and visualizations to evaluate appropriate locations for billboards. Weng et al. [73] proposed an interactive approach for determining critical locations using computed reachability. Weng et al. [72] developed a visual analytics system for improving bus routes with Monte Carlo search.

Visual urban diagnosis is to obtain and understand valuable knowledge hidden behind urban data for improved cities [18, 30, 35, 42, 47, 51, 53]. For example, Malik et al. [45] proposed a general workflow to analyze the correlation between time series. Cao et al. [12] discovered urban anomalies with a combination of detection algorithms and glyph designs. Li et al. [36] analyzed the potential influences between different regions based on the co-occurrence patterns in spatial time series. Wu et al. [75] studied the visual analysis of co-occurrence patterns of cross-domain urban data and supported comprehensive diagnoses. Deng et al. [17] proposed a visual analytics approach for analyzing the pollution propagation patterns extracted by a specialized model.

Dynamic causality in urban space has not been studied. We propose a visual urban diagnosis method based on the dynamic causality of urban time series to facilitate the in-depth understanding of urban problems.

3 BACKGROUND AND SYSTEM OVERVIEW

This section formulates the problem we study, introduces preliminary concepts, workflow, and the user requirements that guide system design.

3.1 Problem Formulation

This study is the result of cooperation with multiple domains. On the one hand, rapid urbanization has resulted in many notorious urban problems, such as noise [9], traffic congestion [53], and water/air pollution [46]. On the other hand, various applications have proven the utility of causal analysis in providing valuable implications. In the past six months, we collaborated closely with three experts (EA, EB, and EC). EA and EB are urban computing experts who have decades of experience studying data-driven solutions for urban problems. EC is a researcher with expertise in causal analyses. We together attempt to apply causal analysis to solve the urban problems. As a statistical analysis method, causal analysis is not limited to computer scientists but can be accepted by urban practitioners with statistical knowledge.

After preliminary discussions, we realize that the causal relations in urban space can be dynamic. Take Fig. 2 as an example. The causal direction between sensors s_1 and s_2 disappears and is flipped over time (Fig. 2a). Although causal detection models can summarize these causal relations across a whole time span, the detected static results can be rough (partially correct) or even incorrect and provide limited insights into precise and reliable policymaking (Fig. 2b and 2c). No interactive systems are available for interpreting and validating the detected causality in an urban context, which makes the causality unreliable and cannot be fully utilized (Fig. 2d).

In this study, we propose a visual analytics approach to address these limitations. This approach aims to enable analysts to obtain and understand correct urban causalities in dynamic environments.

3.2 Concepts

We introduce the following important concepts defined in this paper:

- A *time series* $\mathbf{v}^s = \langle v_1, v_2, \dots \rangle$ is a chronologically ordered set of values recorded at a geographically fixed *sensor* s . These recorded values reflect the real situation around the sensor. For example, an increase in air pollutant concentration unveils the air quality deterioration around the sensor.
- A *stimulus* in this paper is defined as an event, such as an air pollution event or traffic congestion event, that can propagate over sensors and influence them. It can be reflected by an increase or decrease in time series.
- A *causal link* indicates the one-way causality from a sensor to another (Fig. 6a or 6b). A causal link is detected by our causal detection model based on the Granger causality test (Sec. 4). It comprises 1) a *cause* sensor, 2) an *effect* sensor that receives the effect from the cause, 3) the *causal strength*, and 4) a *lag* implying the time required for the cause to influence the effect. Such a causality, to a degree, can be reflected in the trend relationship of the involved time series under the time lag. For example, the increase of a series causes a similar increase (a positive relationship) or opposite decrease (a negative relationship) of another series. Thus, we further derive a metric that can be visualized as a visual hint of the trend relationship and assist in validating the detected causal links. We use 5) Pearson's r , a widely used correlation metric for time series [34, 45, 63], to capture both positive and negative relationships between time series effectively. We first shift the cause sensor's time series by the lag along the time. The Pearson's r is then computed given this shifted time series and that of the effect sensor. A causal link exists if a stimulus from the cause sensor influences the effect sensor.
- A *causal relation* indicates the causality between two sensors. A causal relation comprises at least one causal link between the two sensors. It is bi-directional if it comprises two opposite causal links (Fig. 6c), i.e., these two sensors are causal to each other.
- A *causal graph* $G_c = (N, E)$ is a directed graph, where each node $n_s \in N$ denotes a sensor s and each edge $(n_i, n_j) \in E$ denotes a causal relation, i.e., there is causality between the sensors i and j .

3.3 Workflow and Requirement Analysis

We carefully conducted literature review [16, 31, 38, 53, 64, 76, 76, 89] and interviewed the experts to compile the user requirements that guide system development. On the one hand, we collected and summarized relevant studies in visualizations and causality detection, derived user requirements, and developed a prototype system. On the other hand, we verified and refined the user requirements together with the prototype system through bi-weekly interviews with the experts. The processes above were carried out iteratively and alternately.

We first conclude a three-step workflow as follows:

WP: Pre-processing data. Previous studies carried out data pre-processing before causal detection according to various strategies, such as clustering [65], key event extraction [31, 38], and pattern mining [89]. This step is also essential in our scenario. Directly detecting and analyzing the causal relations of the entire dataset are intractable because they are intricate in space and dynamic over time.

Our study adopts a sensor-centric strategy in the space dimension. A target sensor is specified as ego first, and then its neighbors that have potential causality will be retrieved. Only the causal relations between them will be detected and analyzed. In this way, the spatial complexity of causality can be reduced, and analysts can easily focus on the target areas of their interests. As for the time dimension, the time is properly partitioned into windows to reveal dynamic causalities. Causal detection works given an ego sensor and partitioned windows.

WA: Analyzing causal graphs. Causal detection models can output causal graphs, each of which is constituted by the causal relations between variables. These graphs are valuable. For example, the roots in graphs imply sources [38], and the pathways reveal how the stimulus propagates over and influences variables [16, 53, 76, 89]. In addition to the overall structure, individual relations are worth studying [31]. In our study, multiple causal graphs centered on the ego sensor can be obtained in the partitioned time windows. A set of effective visualizations are proposed to support the multifaceted analysis of these dynamic causal graphs in a spatiotemporal context.

WI: Improving causal detection results. In practice, data-driven causal detection is not always reliable due to various factors, such as data noise, external variables, and missing consideration of certain rules. Therefore, the detection results usually need to be revised and improved, whether through automated intelligent algorithms [89] or the incorporation of human knowledge [31, 64]. We design interactive visualizations tailored for our problem to assist users in discovering and modifying the incorrect causalities.

We further specify six user requirements that guide visual designs.

- R1 **Summarize causal graphs across the time (WA).** First, the experts require the system to summarize all causal graphs detected in different time windows and thereby grasp the brief patterns. For example, *which sensors have strong causal relations across the whole period? What are the causal directions between them?*
- R2 **Explore causal graphs along the time (WA).** Multiple graphs in different time windows constitute a dynamic graph where the structure can change over time [8]. The experts need to relate the causal graphs to the time for effective time-oriented exploration [69, 70, 77]. They also want to learn the temporal variations of the causal structures, such as periodicity and stability. Therefore, the system needs to couple a timeline-based organization with a structure-aware representation to visualize these causal graphs.
- R3 **Learn influence propagation via causal graphs (WA).** The experts also need to drill down individual graphs and learn the influence propagation during a specific period. Specifically, they aim to establish *where the stimulus is from* and *how it influences the urban space*. Therefore, in the system, every individual graph should be accessible in the spatial context.
- R4 **Interpret and validate causal relations (WA, WI).** Causal interpretation and validation are required. The experts hope to interpret *why there are causalities between sensors* and learn *what the effects are based on the involved time series*. Furthermore, validation questions may be asked, for instance, *are causal relations including their directions and time lags reasonable?* So, causal relations must be fully encoded regarding the occurrence frequency, involved time series, and comprised multidimensional attributes.
- R5 **Modify incorrect causal relations (WI).** The experts commented that detection results are not completely reliable. Causal interpretation and validation can help identify incorrect causal relations. Afterward, the experts require modifying them interactively.

3.4 System Architecture

Compass is a web-based application constituted by data storage, backend, and frontend (Fig. 3). The data storage indexes urban time series spatially. The backend is written in Python. The backend handles a causal detection framework (Sec. 4.2) including a data pre-processing module (WP) with the Python multiprocessing (a parallel computing package¹). The frontend is written in TypeScript and Vue.js. Users can specify an ego sensor and partition the time in the frontend. The causal relations involving the ego sensor are then detected in the partitioned time windows. The frontend also empowers users to understand urban causality with the causal graph visualizations (WA) and facilitates the discovery and modification of the incorrect causalities (WI).

4 MODEL

This section introduces the causal detection framework we propose to detect the causal relations between urban time series.

4.1 Granger causality

This subsection introduces the core idea of the Granger causality [26]. The Granger causality test mainly builds on the predictability of one time series to another. Let two variables x and y be the cause and effect to be tested, respectively. The Granger causality says x is the cause of y if \mathbf{v}^x , the time series of x , significantly helps predict \mathbf{v}^y under the same controlled condition of a variable collection $\mathbb{C} = \{c_1, c_2, \dots\}$.

Prediction. The Granger causality test applies the vector autoregressive (VAR) model [61] as the prediction model. In VAR, the current state of a system can be predicted by the past K states in different time series across the system. Suppose a time-varying system has N

variables, and each one collects a time series \mathbf{v} . \mathbf{v}_t^n denotes the record at time t in the n -th variable. The VAR model of this system is written as

$$\begin{cases} \mathbf{v}_t^1 = \mu_1 + \sum_{n=1}^N \sum_{k=1}^K \omega_{1,n,k} \mathbf{v}_{t-k}^n + u_{1,t} \\ \mathbf{v}_t^2 = \mu_2 + \sum_{n=1}^N \sum_{k=1}^K \omega_{2,n,k} \mathbf{v}_{t-k}^n + u_{2,t} \\ \dots \\ \mathbf{v}_t^N = \mu_N + \sum_{n=1}^N \sum_{k=1}^K \omega_{N,n,k} \mathbf{v}_{t-k}^n + u_{N,t} \end{cases}$$

The coefficient $\omega_{i,n,k}$ implies the contribution of the value k time-steps ago of the n -th variable to the prediction of the i -th variable, μ denotes constants, and u denotes error terms. The aforementioned regression equation can be solved separately, thereby obtaining the inherent relationships among the variables.

Causality Test. Testing “ $x \rightarrow y$ ” (x causes y) is based on the following two regression equations:

$$\mathbf{v}_t^y = \mu_y + \sum_{n=1}^N \sum_{k=1}^K \omega_{y,n,k} \mathbf{v}_{t-k}^n + u_{y,t} \quad (1)$$

$$\mathbf{v}_t^y = \mu_y + \sum_{n=1, n \neq x}^N \sum_{k=1}^K \omega_{y,n,k} \mathbf{v}_{t-k}^n + u_{y,t} \quad (2)$$

Note that the second equation ignores the contribution of x . All variables except for x and y are viewed as control variables \mathbb{C} . If the prediction performance improved by considering the records of x is significant, “ $x \rightarrow y$ ” holds true. The detection can eliminate the interference of other variables through the control variables \mathbb{C} and focus on the current two variables, which is impossible with correlation methods.

F Test based on the sum of squared residuals (SSR) is commonly used to determine the statistic significance:

$$F = \frac{(SSR_{(1)} - SSR_{(2)})/K}{SSR_{(2)}/(M - KN)}$$

$SSR_{(1)}$ and $SSR_{(2)}$ denote the sum of squared residuals for the regression equations (1) and (2), respectively. M denotes the number of samples for regression. F has an F-distribution with parameters K and $M - KN$, i.e., $F \sim F(K, M - KN)$. The null hypothesis is “ $x \nrightarrow y$.” In practical, the null hypothesis is rejected if the p-value of F is less than the p-value threshold 0.05; otherwise, the null hypothesis is accepted. The causal strength can be measured by $(0.05 - \text{p-value})/0.05$.

Given a maximum time lag K^+ , multiple tests need to be executed under every $K \leq K^+$ to determine the best result based on the p-value.

4.2 Causal Detection of Urban Time Series

This subsection introduces a causal detection framework we propose for detecting dynamic causality from urban time series. It is based on the Granger causality test but addresses the first challenge in Sec. 1.

4.2.1 Data Pre-processing

First, we pre-process the data in space and time dimensions.

Indexing neighbors. A sensor has causal relations only with its neighbor sensors. This neighbor relationship depends on specific application scenarios. Two concrete examples are provided below. A road sensor can be influenced only by those sensors encountered first along forward the traffic direction because traffic congestions propagate backward. An air quality sensor has causal relations only with those closest sensors around it because air pollutants spread continuously in space. For every sensor, we index its neighbors with potential causality and categorize them into downstream and upstream neighbors based on actual applications. The center sensor is denoted as an ego sensor s_e . The downstream and upstream sensors are denoted as $Ego_d(s_e)$ and $Ego_u(s_e)$, respectively. Fig. 4a illustrates $Ego_d(s_x)$ and Fig. 4b illustrates $Ego_u(s_y)$.

Partitioning time. Two strategies for time partitioning can be used based on whether the time series is periodic. For periodic time series, they can be directly partitioned by their period. For example, traffic time series can be naturally partitioned by 24 hours. Otherwise, the peaks of the time series will be extracted automatically, and thereby the time windows can be identified based on these peaks. This procedure is accomplished using Python `find_peaks` imported from a well-

¹<https://docs.python.org/3/library/multiprocessing.html>

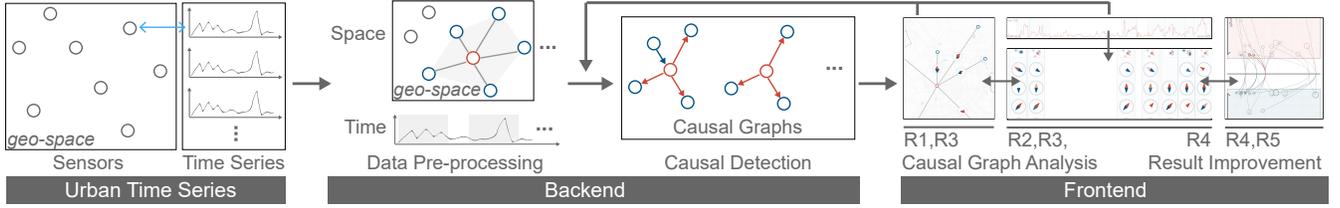


Fig. 3. System architecture. Compass comprises data storage, backend, and frontend. Data storage stores urban time series. Backend runs a causal detection framework. Frontend supports the input specification for causal detection and visual analysis of the detected dynamic causalities.

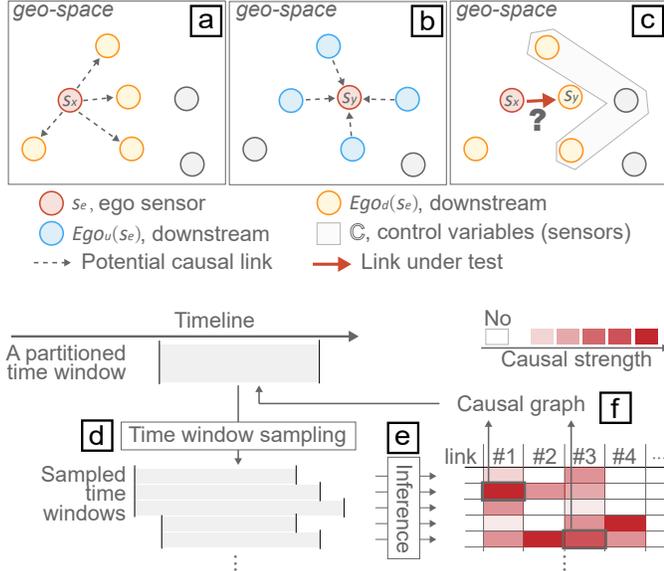


Fig. 4. Causal detection based on neighbor index (top) and debounce strategy for alleviating time sensitivity (bottom). (a) The downstream of s_x can be potentially caused by s_x . (b) The upstream of s_y can potentially cause s_y . (c) Sensors enclosed in the polygon are considered control variables when testing the causal link $s_x \rightarrow s_y$. (d) Sampled time windows are generated by slightly adjusting a partitioned window. (e) Multiple sample causal graphs are detected in every sampled time window. (f) A final causal graph is constructed for the partitioned window.

known package `scipy.signal` [32]. The `find_peaks` has four important parameters: *height*, *prominence*, *wlen*, and *distance*. The *height* and *prominence* limit peaks' minimum heights and prominences. The *wlen* indicates the maximum duration of windows. The *distance* means the minimum horizon distance between peaks. Given these parameters and a time series, `find_peaks` returns the intervals where peaks are located.

4.2.2 Causal Link Testing

Given an ego sensor, we use the Granger causality test (Sec. 4.1) to test the causal links between it and each of its neighbor sensors in every partitioned time window. Before applying the Granger causality test, two input parameters for each test need to be specified, namely, the control variables \mathbb{C} and maximum time lag K^+ .

First, when testing the link " $s_x \rightarrow s_y$ ", $Ego_u(s_y) \setminus \{s_x\}$ are naturally considered control variables \mathbb{C} (Fig. 4c). Second, K^+ can be determined based on actual applications and the spatial distance between sensors. For example, the causality between road sensors results from the backward propagation of traffic waves, while the causality between air quality sensors results from the pollution diffusion. Such physical semantics determines the range of possible lag. The range also includes a maximum time lag. We let the range be specified by system users.

Each causality test is independent of the other. We use a parallel computing package called multiprocessing to speed up the process.

4.2.3 Debounce

The causal detection is sensitive to time windows due to the unavoidable noise and dynamics in urban data. We address this issue and obtain

stable causal graphs by using a debounce strategy. When detect a causal graph in a time window, nine windows are sampled by slightly adjusting the time window, such as shifting, shrinking, and enlarging (Fig. 4d). The causal detection will be executed in all these sample windows, generating multiple sample causal graphs (Fig. 4e). If a link appears in all sample graphs (e.g., link #1 and #2 in Fig. 4f), the sample link with the maximum strength is added into the final causal graph (e.g., those links highlighted by a bold border in Fig. 4f).

In the end of the framework, the stable causal graphs involving the ego sensor and its neighbors are detected in every partitioned windows.

5 VISUAL DESIGN

Compass comprises four views, namely, map, time, graph, and relation views. The map view provides a spatial context for selecting the ego sensors of interests and reasoning causal graphs (R1 and R3). The time view shows the time series of selected sensors and the partitioned time windows. It also serves as a timeline for the graph view. The graph view displays the causal graphs detected in each window along the timeline. In this view, users can explore the detected dynamic causal graphs along the time (R2) and interpret each graph with its causal relations (R3 and R4) from a spatiotemporal perspective. The relation view presents every causal relation from multiple dimensions. Thereby, users can interpret causal relations further (R4), discover spurious ones, and improve causal detection results (R5).

5.1 Map View

The map view adopts a geographic map to provide a spatial context. Every sensor is represented as a circle on the map according to its geographic position. Three types of graphs involving these sensors can be displayed on the map: spatial ego-graph, spatial causal graph, and spatiotemporal causal graph. Which type of graph is being displayed depends on the analysis stage and user's interactions (Sec. 5.5).

Visualizing spatial ego-graph (Fig. 5a). A sensor is colored red when users select it as an ego sensor. Its neighbor sensors are colored blue. These two colors are applied throughout the system. These neighbor sensors will connect to the ego sensor with edges, constituting a spatial ego-graph. A question mark is placed on the center of each edge if the causal relation remains to be detected. Three edge color styles are used based on whether the neighbor sensor belongs to the downstream or upstream. 1) If a neighbor belongs only to the downstream, the edge is colored red as the same as the ego sensor to indicate that the causal relation, if exists, must start from the ego (e.g., e1 in Fig. 5a). 2) Similarly, if a neighbor belongs only to the upstream, the edge is colored blue (e.g., e3 in Fig. 5a). 3) If a neighbor belongs to both the downstream and upstream, the edge is colored gray (e.g., e2 in Fig. 5a). Gray is also applied throughout the system to indicate bi-directional.

Visualizing spatial causal graph (Fig. 5c-1). Each causal graph detected in a time window can be displayed on the map as a spatial causal graph (R3). The question mark will be replaced with the arrows that denote the causal relations between sensors. Each arrow represents a causal link and points from the cause sensor to the effect sensor. The arrows are colored according to the cause sensor. The arrow opacity encodes the causal strength. Moreover, we follow a compass metaphor to place the arrows (Fig. 5b). We call such a design *compass glyph*. Each glyph denotes a causal relation. An edge is colored red, blue, or gray according to its direction indicated by the arrows. Sensors are removed to reduce clutter if they have no causality with the ego sensor.

Visualizing spatiotemporal causal graph (Fig. 5f). The causal graphs across the windows can be summarized as a spatiotemporal

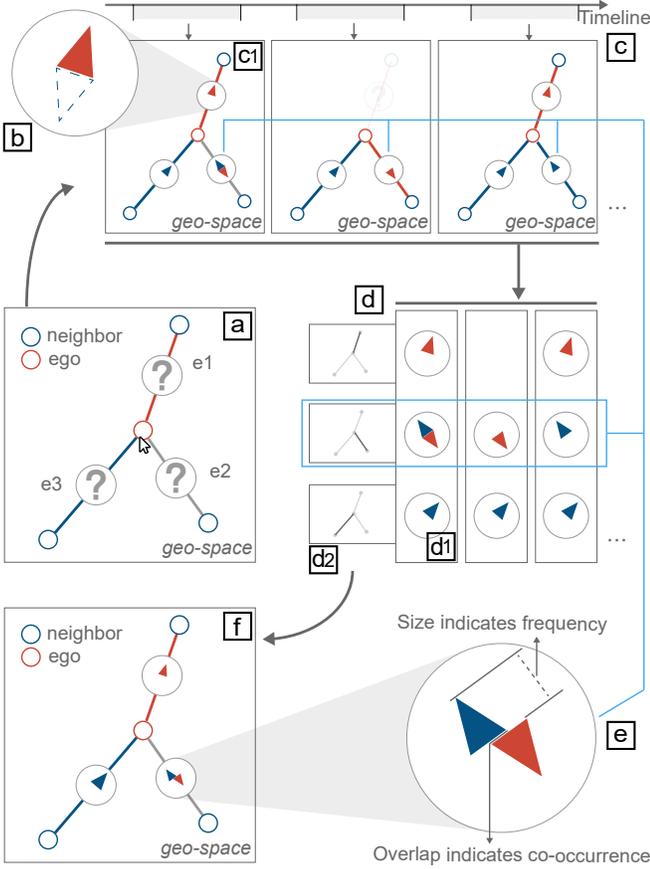


Fig. 5. Visual encodings for causal graphs. (a) An illustrative spatial ego-graph comprises three edges (i.e., e1, e2, and e3) with three different color styles. (b) A compass glyph shows the causal relation where only the causal link from the ego sensor exists. (c) Multiple causal graphs are detected in different time windows and depicted as spatial causal graphs on maps. (d) Graph bands are aligned along the timeline as a compact visualization of dynamic causal graph. (e) A revisited compass glyph summarizes the compass glyphs of the edge e2. (f) A spatiotemporal causal graph summarizes the spatial causal graphs in (c).

causal graph on the map (R1). The causal directions across multiple graphs are first aggregated according to the edge. The aggregation is then encoded with a *revisited compass glyph* (Fig. 5e). For example, Fig. 5e summarizes the causal directions in the bottom-right edge (i.e., e2) across the three causal graphs. In a revisited compass glyph, arrows still convey the causal directions within a spatial context. The glyph is revisited to incorporate the temporal information as follows. The arrow’s size encodes the frequency of the causal links with this direction. An offset between the two opposite arrows can be observed if bi-directional relations exist. The overlapping part of the two arrows encodes the frequency of bi-directional relations.

Justification. The map view uses colors and the compass designs to represent the causal relations. First, color hue is the most effective visual channel for category data second only to spatial region [48]. Second, the compass designs integrate the arrows to one place no matter whether it is bi-directional or not. In this way, users can distinguish the directions at a glance without looking for arrows along the edges. Furthermore, this integration not only supports the encodings of temporal co-occurrences (Fig. 5e) but also enables the structure-aware dynamic causal graph visualization in the graph view (Sec. 5.3).

5.2 Time View

The time view presents the time series of sensors using a line chart, where temporal features, such as peaks and periodicity, are exposed. The time series of the ego sensor is displayed first, while the time series of the neighbor sensors can be displayed on demands. The lines

are colored according to the sensors. Users can specify peak-based and periodicity-based time partitioning strategies. The partitioned time windows are subsequently shown with semitransparent rectangles. Furthermore, this view also serves as a timeline for the graph view (Sec. 5.3) because these two views are juxtaposed up and down.

5.3 Graph View

The graph view adopts a new dynamic graph visualization and supports the effective analysis of multiple causal graphs from graph (R3), relation (R4), and dynamics levels (R2).

Visualizing single causal graph (R3). Simply presenting spatial causal graphs results in low scalability because the maps are space ineffective (Fig. 5c). To this end, we compact each causal graph (e.g., Fig. 5c-1) into a *graph band* (e.g., Fig. 5d-1). The compass glyphs from top to bottom in a graph band correspond to that of the edges in a clockwise order starting from the west direction. These compass glyphs preserve the spatial context for each causal relation.

To further integrate the spatial context of causal graphs, heatmap-based spatial summaries for every causal graph are provided and placed above the bands (Fig. 1c-1). The heatmap is generated from the spatial distribution of the influences of the ego and neighbor sensors. The darker red (or blue) indicates that the area receives more influences from the ego sensor (or the neighbor sensors). To do so, we first generate a red heatmap based on the spatial distribution of the red causal links. A blue heatmap is generated in the same way. We subtract the blue heatmap from the red heatmap and finally derive the summary.

Visualizing the causal relations of a graph. The graphs need to be unfolded to interpret and validate the causal relations further (R4). Fig. 1c-2 shows an unfolded graph band. First, the compass glyphs are shrunk, leaving a ring space to encode the lags with arcs. The arc behind an arrow corresponds to the causal link denoted by the arrow and has the same color as the arrow. The arc length encodes the lag of the causal link. Second, a canvas expands from the right. For every causal relation, the two involved time series are sliced according to the glyph’s right side. The lines are also colored according to the sensor. Users can easily compare the time series of all sensors involved in the graph because these glyphs are aligned vertically in a graph band.

Visualizing multiple causal graphs. The system should expose the temporal variation of causalities and support the time-oriented exploration (R2). Thus, all graph bands are placed along the time view according to the time window in which the graph is detected (Fig. 5d). The bands repel each other to avoid occlusion while keeping their positions along the timeline as much as possible. To further enhance the spatial context, we add minimaps on the left side of this view. Each minimap (Fig. 5d-2) contains the same ego-graph structure as in the map view but removes detailed information such as geographic background, colors, and directions. The bold edge indicates where the glyphs on the right are located. This visualization presents time-varying causal relations in an unobstructed way. Take Fig. 5 as an example. The causal relations of the bottom left edge (i.e., e3) exist in all time windows. These existences may be hidden by the scattered arrows in Fig. 5c but are clearly revealed in the third row in Fig. 5d.

Justification. Many timeline-based and scalable visualizations for dynamic graphs have been proposed [24]. In Beck et al.’s hierarchical taxonomy [8], juxtaposed node-link diagrams [11] are the most suitable for our problem. We have created an alternative design using juxtaposed node-link diagrams (Fig. 7a-1). However, all edges are connected to the red node (i.e., the ego sensor), causing visual clutter. The variations of the causal directions cannot be easily revealed. We have also created an alternative (Fig. 7a-2). Each row represents a neighbor sensor or its causal relations with the ego sensor. Each rectangle is divided into two parts to indicate the opposite directions. Although the causal relations are revealed without clutter, the spatial structures are still lost. We enhance this design and derive our final design (Fig. 1c) by using the compass glyphs and adding the heatmaps and the minimaps.

5.4 Relation View

The relation view (Fig. 6d) shows the multidimensional details of causal relations (R4). Causality suspiciousness is visually encoded to help identify and correct suspicious causal links and relations (R5).

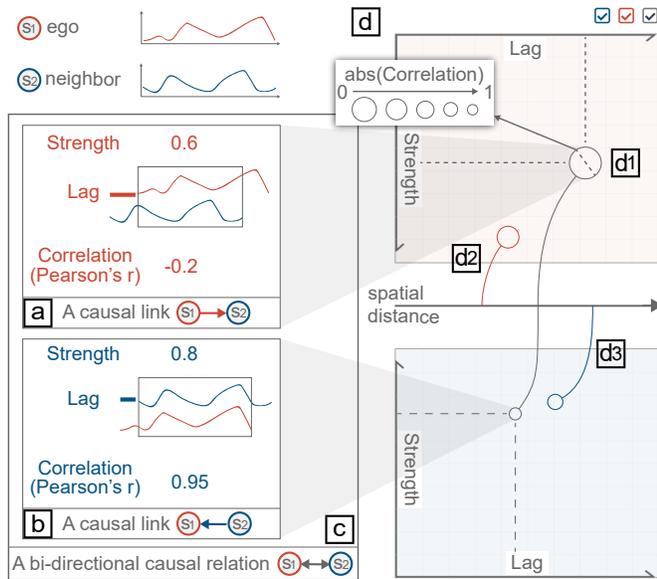


Fig. 6. A causal relation (left) and the relation view for visualizing causal relations (right). (a) A causal link may be incorrect because the Pearson’s r is near 0. (b) A causal link opposite to (a) has a large Pearson’s r . (c) A bi-directional causal relation comprises two causal links with opposite directions. (d) The relation view adopts a multidimensional visualization and presents (d1, d2, and d3) three illustrative causal relations.

Showing multidimensional details. Causal links that start from the ego sensors are plotted as the top scatterplot according to their lags and strengths. The circle size encodes the absolute value of the Pearson’s r . Please refer to Sec. 3.2 for how to compute the Pearson’s r . These circles also connect to the middle axis with curves according to the two involved sensors’ spatial distance. Causal links that start from the neighbor sensors are visualized at the bottom in the same way. If two links are comprised by the same causal relation, their curves are regarded as one and correspond to the relation (Fig. 6d-1). We call these curves *relation curves*. The curves are colored by red, blue, or gray according to the causal directions.

Unveiling suspicious causalities. While showing the details, the well-designed visual encoding helps reveal two types of suspiciousness. The first one is the bi-directional relations with a large spatial distance. Bi-directional means that either the stimulus returns after arriving at one sensor or the two sensors receive the same stimulus from another sensor. It is suspicious if these two situations happen between two remote sensors. The opposite causal links are unfolded in two different charts so that users can compare them and determine the correct one. The second one is the causal links with low Pearson’s r , i.e., large circles. The time series of two sensors intuitively portray the causality. If a low r is observed, then the causal link needs to be further validated. Because these visual encodings have provided sufficient visual hints, we do not use extra algorithms for suspiciousness recommendations.

Justification. Each causal relation is an essentially multidimensional datum with seven dimensions (Fig. 7b-1). Various visualizations have been proposed for this kind of data, such as dimension reduction [68, 88], scatterplot matrix [49], parallel coordinate plot (PCP) [10], and glyphs [82]. This view should satisfy two requirements, showing details and revealing suspiciousness. We have created two alternatives after exploring the design space. The first one is the spreadsheet that is commonly seen as a detail view (Fig. 7b-1). However, it cannot handle many relations. The second one is the PCP (Fig. 7b-2) better than the spreadsheet. Numerous causal relations can be displayed. The values of each dimension can be seen in the parallel axes. The first type of suspicious relations can be indicated by the gray lines passing the top of the distance axis (Fig. 7b-2). Built on the PCP, we adopt scatterplots to visualize the dimensions related to the same causal direction (Fig. 6d). Scatterplots allow another suspiciousness metric, the Pearson’s r , to be encoded using the size that can be accurately perceived. Besides, scatterplots are more familiar to the experts than the PCP.

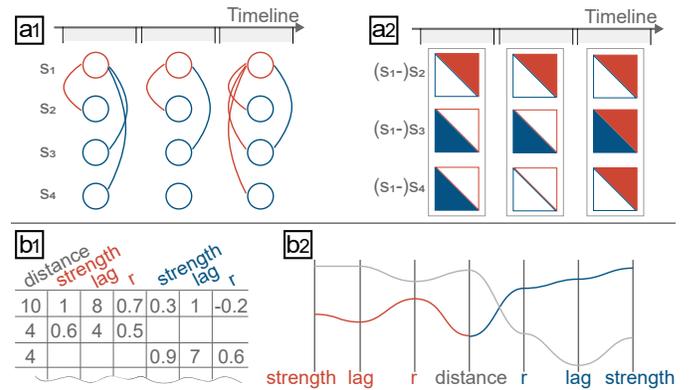


Fig. 7. The alternative designs are designed for (a1, a2) visualizing dynamic causal graphs (b1, b2) presenting causal relations.

5.5 User Interactions

The following interactions are incorporated into the system.

Parameters setting. Users can specify the input parameters of the causal detection model. First, the ego sensor of users’ interests can be selected after users click it in the map view. Second, the range of the influence speed can be adjusted via the slider of Fig. 1a-4, which determines the range of lag. Third, users can determine the time partitioning strategies via the buttons in Fig. 1b1. The first is periodicity-based, and the second is peak-based. The parameters have defaulted. Clicking the button can generate windows. These parameters can be changed via the panel triggered by right-clicking the buttons. The system also allows users to resize and move each time window directly.

Details-on-demand inspection. The time view can show the time series of the neighbor sensors if users right-click them. The graph band can be unfolded when users click the band. Then, more information about the causal relations is shown.

Causal relations hovering and highlighting. Causal relations can be highlighted throughout the system to respond to graph, edge, and individual level hovering. First, the relations comprised by a causal graph are highlighted when users hover over or unfold the graph band. The map view will present the corresponding spatial causal graph. Second, multiple relations between two sensors are highlighted when users hover over the minimap (in the graph view) or the revised compass glyph. Finally, a single relation is highlighted when users hover over the compass glyph or the relation curve. Besides, the global highlighting based on the causal directions is also supported (Fig. 1d-1).

6 EVALUATION

This section presents the evaluation of the system. We first invited the experts and performed a training session to introduce the system. We ensured that the experts understood the system, including the visual encodings and user interactions. Afterward, the experts freely used our system to perform causal analyses on two real-world time series datasets, namely, air pollution and traffic datasets. We also interviewed the experts one-on-one to collect their feedback after the case studies.

6.1 Causal Analysis of Air Pollution Time Series

Data description. The air pollution dataset we used comprises the hourly readings of the PM_{2.5}² concentration from 448 major air quality sensors in China between January 8 and March 23, 2018. Each sensor comprises a time series with 1800 recorded values. This dataset has a total of 806,400 records and a size of 2.99 MB. We constructed a neighbor index based on the sensors’ spatial proximity (Sec. 4.2.1).

Beijing is the largest city in China and the capital of China. The experts aimed to disclose how did air pollution influence Beijing and its surrounding areas. They first selected the sensor located in the urban area of Beijing as an ego sensor. Six neighbor sensors were connected to the ego sensor with edges, which constitutes a spatial

²PM_{2.5} refers to the atmospheric particulate matter (PM) that has a diameter of fewer than 2.5 micrometers

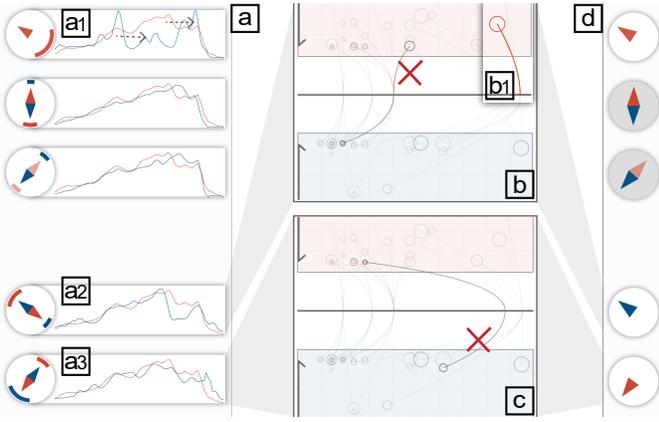


Fig. 8. (a) A graph band is unfolded to show much information of five causal relations. (b, c) Two suspicious causal relations are highlighted in the relation view. (d) A graph band shows the correct causal relations.

ego-graph. All edges were gray. These neighbor sensors belong to both the downstream and upstream because the spread of air pollutants has no explicit directions. The time series of the ego sensor was shown in the time view (Fig. 1b). The experts applied the automatic peak identification because of no periodicity. We defaulted a set of suitable parameters after multiple trials to alleviate the user’s interaction burden: $height=180$, $distance=100$, $prominence=100$, and $wlen=210$. Seven time windows were generated in Fig. 1b. These windows each cover the prominent peak with a proper duration. Thus, the experts were satisfied with this result and did not adjust them. They set the range of pollution influence speed as [6 km/h, 38 km/h] based on the typical wind speeds in Beijing via the slider of Fig. 1a-4. Seven causal graphs were then accordingly detected in these windows by the model.

The experts first obtained the spatiotemporal summary (Fig. 1a). Fig. 1a-1 showed that the sensors of Beijing strongly influenced each other because the large arrows and obvious overlaps were observed in the revised compass glyphs. EB explained that the pollutants transported back and forth in Beijing because of its terrain surrounded by mountains on three sides. EA commented “*all districts in Beijing should be treated as a whole for better air pollution control.*” The experts also learned that air pollutants of Beijing could be from Zhangjiakou and could propagate to Langfang (Fig. 1a-2). They decided to analyze the fine-grained causality revealed by causal graphs after briefly understanding the detected causalities through this summary.

The graph view presented all causal graphs and causal relations along the time (Fig. 1c). The experts first wanted to confirm the conclusion drawn by Fig. 1a-1 further. They highlighted all bi-directional causal relations by checking the gray check box (Fig. 1d-1). Not surprisingly, the causal relations among Beijing sensors enclosed in Fig. 1c-3 were frequently bi-directional.

The experts then explored the causal graphs and drilled down to specific graphs of their interests based on the ego’s time series features.

A sharp peak in the third time window was first noticed. The heatmap below showed that the stimulus of an air pollution event spread from north to south (Fig. 1c-1). After hovering over this band, the map view displayed the spatial causal graph (Fig. 1a-1). However, some involved causal links were suspicious because they had low correlation coefficients indicated by the circle sizes in the relation view (Fig. 1d-2). Given that, the experts unfolded this graph band to verify the causalities (Fig. 1c-2). The orders in which the peaks appear in the two time series supported the causal directions detected by the model. The time differences of the peaks were also consistent with the lags. EB said, “*although the time series were not highly correlated, the model captured the key features, i.e., the peaks, and obtained reasonable causalities.*”

The highest peak in the last time window attracted the experts’ attention (Fig. 1c-5). They unfolded the graph band below. Three suspicious causal links were identified as follows.

First, the causal link with Zhangjiakou comprised two dissimilar time series (Fig. 8a-1). The experts hovered over it and highlighted its relation curve in the relation view (Fig. 8b-1). The large circle

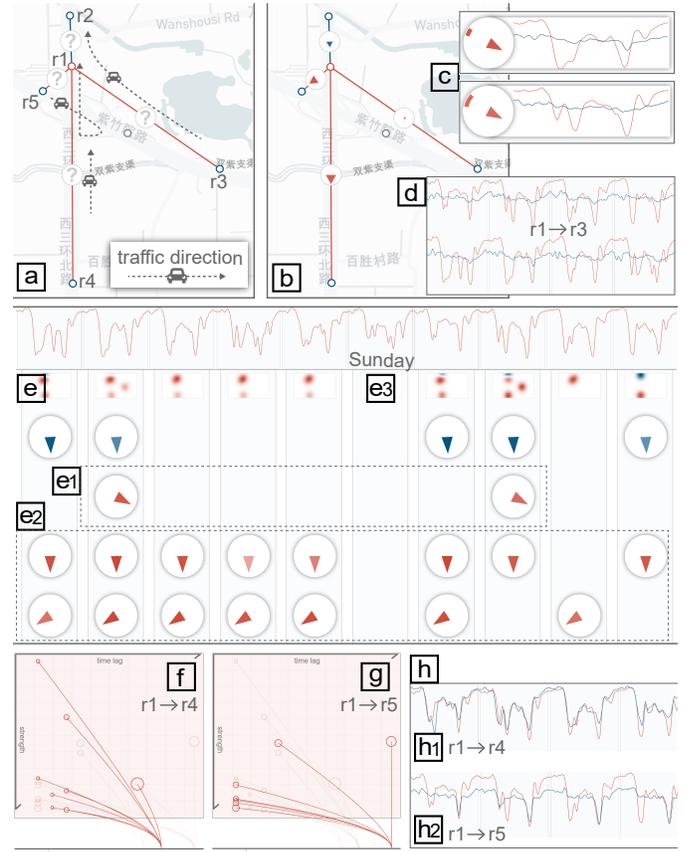


Fig. 9. Causal analysis of traffic time series. (a) The spatial ego-network of r_1 is shown on the map. (b) A spatiotemporal causal graph summarizes the dynamic causal graphs. (c) Two causal links of “ $r_1 \rightarrow r_3$ ” are unfolded in the graph view. (d) The time series of r_3 is shown together with that of r_1 . (e) The graph view presents ten causal graphs. (f, g) The causal relations of “ $r_1 \rightarrow r_4$ ” and “ $r_1 \rightarrow r_5$ ” are presented in the relation view. (h) The time series of r_4 and r_5 are shown together with that of r_1 .

also implied a low correlation. Nonetheless, the expert scrutinized the time series carefully (Fig. 8a-1) and learned that the time differences between the peaks in these two time series generally corresponded to the large lag. EA explained that the low correlation might result from the long distance and geographical terrain between them.

Second, the causal relation with Langfang (Fig. 8a-2) was associated with the two circles with different sizes in the relation view (Fig. 8b). The smaller circle indicated the higher correlation of the blue causal link that started from Langfang. The experts also learned that the ups and downs of the blue time series were ahead of that of the red one after examining the two time series. This finding was consistent with that in the relation view, and the red link was considered incorrect.

Third, the bi-directional relation with Baoding was also suspicious because of the remote distance. The experts considered the blue link incorrect based on the correlations indicated by the relation curve (Fig. 8c) and the ups and downs in the time series (Fig. 8a-3).

Finally, the experts deleted the incorrect causal links and obtained the causal relations as Fig. 8d. They further knew that the air pollutants generally propagated from the east to the west. EA and EB commented, “*identifying the sources of air pollution is a highly complex task, and these causal relations between sensors can be of great help.*”

We showcased that the two most serious air pollution in Beijing were caused by the northwest and southeast areas, respectively. The time-oriented analysis disclosed these fine-grained and dynamic causalities that shed light on precise source identification and better policymaking that a static causal graph cannot support.

6.2 Causal Analysis of Traffic Time Series

Data description. The traffic dataset we used was from the released Q-Traffic dataset [39]. This dataset includes subparts of the road network

in Beijing and the traffic speed for every 15 minutes for every road segment. We first sampled 511 high-level roads, such as highways. Each road segment is regarded as a road sensor. We also sampled the readings from May 9, 2017 to May 18, 2017. Finally, this dataset totally comprises 490,560 records and has a size of 1.89 MB.

The experts stated that traffic congestion events always propagate backward and cause other congestion events. Thus, the causal directions are opposite to the traffic directions. Based on this prior knowledge, we recovered the downstream and upstream relations of the roads based on the road network and constructed a neighbor index (Sec. 4.2.1). A busy traffic area with elevated roads was selected for this case study.

After the experts selected the road segment $r1$ (abbreviated as $r1$ below) as the ego sensor, an ego-graph was shown on the map view and implied possible causal relations (Fig. 9a). We denoted the four neighbor sensors as $r2$ to $r5$ for convenience. The experts set the range of traffic influence speed as [0.1 km/15 min, 1.5 km/15min] via the slider of Fig. 1a-4. Multiple causal graphs were detected in these time windows and summarized in the map view (Fig. 9b). The summary showed that $r2$ caused $r1$, and $r1$ caused $r4$ and $r5$ but rarely $r3$. The experts then analyzed the causal graphs in the graph view (Fig. 9e).

Pattern #1. “ $r1 \rightarrow r3$ ” ($r1$ caused $r3$) occurred only twice (Fig. 9e-1), while others occurred many times. To figure out the reason, the experts unfolded the related graph bands where “ $r1 \rightarrow r3$ ” exits and examined the two causal relations (Fig. 9c). The raw time series showed that the ups and downs in the red time series caused only slight fluctuations in the blue one. Given these observations and the frequency, the experts thought that the causality between $r1$ and $r3$ can be ignored. EA required to inspect the whole time series of $r3$ further. The traffic conditions of $r3$ were more stable compared with that of $r1$ (Fig. 9d), which supported the experts’ decision.

Pattern #2. “ $r1 \rightarrow r4$ ” and “ $r1 \rightarrow r5$ ” frequently occurred (Fig. 9e-2). The experts obtained the details of the relevant causal links in the relation view. Fig. 9f shows the causal links of “ $r1 \rightarrow r4$ ”. The correlation of these links were generally high, indicated by the small circle sizes. By contrast, the causal links of “ $r1 \rightarrow r5$ ” had larger circles and lower correlations (Fig. 9g). The experts figured out such a difference by inspecting the whole time series of $r4$ and $r5$ (Fig. 9h). The time series of $r1$ and $r5$ differed obviously in each morning time period. EB explained that $r1$ influenced $r4$ throughout the days (Fig. 9h-1), while $r1$ influenced $r5$ only in the afternoon (Fig. 9h-2).

Pattern #3. No causality was detected on Sunday (Fig. 9e-2). EA explained that this was because the traffic volume on Sundays was small. On Sundays, residents tend to rest at home instead of working overtime or traveling. EC also stated that even if residents wanted to go out, their departure times could be more varied than working days.

These patterns can promote informed traffic management. For instance, managers should focus more on the vehicles from $r4$ and $r5$ to take timely measures, such as early warning and routing. Furthermore, traffic control may be omitted on Sundays, thereby saving resources.

6.3 Expert Interviews

We conducted informal interviews with the experts and collected their feedback after the case studies. The feedback is summarized as follows.

Visual design. All three experts agreed that the visual designs of Compass were easy to learn and understand. They praised the compass glyph, “*it is interesting and intuitive to indicate causal directions because the purpose of a compass is to show directions.*” EA and EB can quickly understand the graph view although they have not seen dynamic graph visualizations before because “*the causal relations in the graph view can be clearly associated with the time and space.*”

Usability and improvements. All experts confirmed the system usability. EA commented, “*Compass allows us to obtain fine-grained and dynamic urban causalities, which cannot be supported before.*” EC said, “*It would be very interesting if the proposed methods are applied to other fields.*” While the system has been confirmed to be satisfied with the user requirements, the experts provided suggestions on improving the usability further. EB required to modify or add neighbor sensors based on users’ knowledge. We have implemented this function. EA and EC said that accessing actual video helps confirm the identified knowledge [66, 67, 77], e.g., the traffic causality in Sec. 6.2. Domain-specific views can be added when deploying to the actual scenarios.

7 DISCUSSION

Implications. This study is the first to extend the visual causal analysis to urban scenarios. First, we are aware of the dynamics of causality in urban environments. A causal detection framework and temporal visualizations are combined to unveil these dynamic causalities. The case studies demonstrate the dynamic causalities in air pollution phenomena and traffic situations. The static causal graphs detected by the previous methods will ignore these time-varying cause-effect relationships and thus cannot support accurate and informed decisions. Second, we observe the causalities detected by automatic models are not always reliable in urban domains. This observation supports other researchers’ opinion [31, 76] that human knowledge needs to be integrated into the analysis loop for reliable analyses.

Generalizability. The generalizability of Compass is two-fold. In terms of *application*, the two real-world case studies have demonstrated that Compass can be easily applied to the urban time series data from different domains. Other potential applications include analyzing noise pollution [9], water pollution [1], and complex physical systems (e.g., IoT systems of tree ecosystem services [46]). For example, a river can be polluted by its upstream tributaries. Compass can help identify when the river is polluted and reveal which tributaries pollute it for each period when the pollution events are observed. In terms of *visual analytics*, Compass can be regarded as a model-free visual analytics approach and is not limited to the Granger causality test. Granger causality test is currently the most suitable method for our problem because of its wide use, interpretability, and ability to detect causality based on short partitioned time series. However, some recognized issues of the Granger causality test still exist. For example, it is a linear model [27]. Any better model proposed in the future can be easily integrated into our visual analytics system.

Limitations and future work. Three limitations are observed in this study. First, the system cannot support the analysis of complex causal graphs. Although ego graph-like causal graphs support an effective analysis for the target area, causal analyses can be enhanced with more complex causal graphs, such as by involving the relations between neighbor sensors. However, analyzing the complex graphs in a spatiotemporal context is challenging. For example, visual clutter is hard to avoid because the geo-fixed nodes make the causal graph layout methods inapplicable [65, 76]. It also becomes difficult to clearly reveal the temporal variations of the detected causalities as the number of edges increases. Utilizing the third dimension can be a potential solution [5, 14, 62, 80] and deserves further study. Secondly, due to the lack of multi-source datasets in the same temporal and spatial scope, the effectiveness of our method in cross-domain causal analysis has not been proven. In fact, causality may exist between different urban datasets. For example, as the traffic volume increases, the air quality will deteriorate because cars emit exhaust gas and pollute the air [83]. In the future, we will collect richer datasets and conduct studies, thereby evaluating and improving our approaches. Finally, we found that when the number of links to be tested in all time windows is more than 150, the calculation time exceeds ten seconds. In the future, we will optimize the detection model, implement it with a more high-performance programming language, and deploy it to a distributed computing server.

8 CONCLUSION

This study presents a novel visual analytics approach that assists analysts in detecting and analyzing dynamic causalities in urban domains. We first extend the Granger causality test and propose a causal detection framework for dynamic urban causality. We then develop Compass by coordinating a set of well-designed visualizations. Compass facilitates analysts to interpret dynamic causal relations and improve causal detection results. We have demonstrated its effectiveness with case studies on real-world datasets and expert interviews. In the future, we will enhance the capability of Compass and deploy it to specific fields.

ACKNOWLEDGMENTS

The work was supported by NSFC (62072400), Zhejiang Provincial Natural Science Foundation (LR18F020001), and the Collaborative Innovation Center of Artificial Intelligence by MOE and Zhejiang Provincial Government (ZJU).

REFERENCES

- [1] P. Accorsi, N. Lalande, M. Fabrègue, A. Braud, P. Poncelet, A. Sallaberry, S. Bringay, M. Teisseire, F. Cernesson, and F. L. Ber. HydroQual: Visual analysis of river water quality. In *Proc. of IEEE VAST*, pp. 123–132, 2014.
- [2] G. L. Andrienko, N. V. Andrienko, W. Chen, R. Maciejewski, and Y. Zhao. Visual analytics of mobility and transportation: State of the art and further research directions. *IEEE Transactions on Intelligent Transportation Systems*, 18(8):2232–2249, 2017.
- [3] G. L. Andrienko, N. V. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel. From movement tracks through events to places: Extracting and characterizing significant places from mobility data. In *Proc. of IEEE VAST*, pp. 161–170, 2011.
- [4] N. V. Andrienko, G. L. Andrienko, S. Miksch, H. Schumann, and S. Wrobel. A theoretical model for pattern discovery in visual analytics. *Visual Informatics*, 5(1):23–42, 2021.
- [5] B. Bach, P. Dragicic, D. Archambault, C. Hurter, and S. Carpendale. A review of temporal data visualizations based on space-time cube operations. In *Proc. of EuroVis 2014*, 2014.
- [6] J. Bae, T. Helldin, and M. Riveiro. Understanding indirect causal relationships in node-link graphs. *Computer Graphics Forum*, 36(3):411–421, 2017.
- [7] J. Bae, E. Ventocilla, M. Riveiro, T. Helldin, and G. Falkman. Evaluating multi-attributes on cause and effect relationship visualization. In *Proc. of VISIGRAPP*, pp. 64–74, 2017.
- [8] F. Beck, M. Burch, S. Diehl, and D. Weiskopf. A taxonomy and survey of dynamic graph visualization. *Computer Graphics Forum*, 36(1):133–159, 2017.
- [9] J. P. Bello, C. T. Silva, O. Nov, R. L. DuBois, A. Arora, J. Salamon, C. Mydlarz, and H. Doraiswamy. SONYC: a system for monitoring, analyzing, and mitigating urban noise pollution. *Communications of the ACM*, 62(2):68–77, 2019.
- [10] J. Bok, B. Kim, and J. Seo. Augmenting parallel coordinates plots with color-coded stacked histograms. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [11] M. Burch, C. Vehlou, F. Beck, S. Diehl, and D. Weiskopf. Parallel edge splatting for scalable dynamic graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2344–2353, 2011.
- [12] N. Cao, C. Lin, Q. Zhu, Y. Lin, X. Teng, and X. Wen. Voila: Visual anomaly detection and monitoring with streaming spatiotemporal data. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):23–33, 2018.
- [13] S. Chen, X. Yuan, Z. Wang, C. Guo, J. Liang, Z. Wang, X. L. Zhang, and J. Zhang. Interactive visual discovering of movement patterns from sparsely sampled geo-tagged social media data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):270–279, 2016.
- [14] Z. Chen, Y. Wang, T. Sun, X. Gao, W. Chen, Z. Pan, H. Qu, and Y. Wu. Exploring the design space of immersive urban analytics. *Visual Informatics*, 1(2):132–142, 2017.
- [15] F. Chirigati, H. Doraiswamy, T. Damoulas, and J. Freire. Data Polygamy: The many-many relationships among urban spatio-temporal data sets. In F. Özcan, G. Koutrika, and S. Madden, eds., *Proc. of ACM SIGMOD*, pp. 1011–1025, 2016.
- [16] T. N. Dang, P. Murray, J. Aurisano, and A. G. Forbes. ReactionFlow: An interactive visualization tool for causality analysis in biological pathways. In *BMC proceedings*, vol. 9, pp. 1–18. Springer, 2015.
- [17] Z. Deng, D. Weng, J. Chen, R. Liu, Z. Wang, J. Bao, Y. Zheng, and Y. Wu. AirVis: Visual analytics of air pollution propagation. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):800–810, 2020.
- [18] Z. Deng, D. Weng, Y. Liang, J. Bao, Y. Zheng, T. Schreck, M. Xu, and Y. Wu. Visual cascade analytics of large-scale spatiotemporal data. *IEEE Transactions on Visualization and Computer Graphics*, 2021.
- [19] J.-M. Dufour and E. Renault. Short run and long run causality in time series: theory. *Econometrica*, pp. 1099–1125, 1998.
- [20] M. Eichler. Granger causality and path diagrams for multivariate time series. *Journal of Econometrics*, 137(2):334–353, 2007.
- [21] N. Elmquist and P. Tsigas. Causality visualization using animated growing polygons. In *Proc. of IEEE Symposium on Information Visualization*, pp. 189–196, 2003.
- [22] N. Elmquist and P. Tsigas. Growing squares: Animated visualization of causal relations. In *Proc. of ACM Symposium on Software Visualization*, pp. 17–26, 2003.
- [23] N. Ferreira, M. Lage, H. Doraiswamy, H. T. Vo, L. Wilson, H. Werner, M. Park, and C. T. Silva. Urbane: A 3D framework to support data driven decision making in urban development. In *Proc. of IEEE VAST*, pp. 97–104, 2015.
- [24] V. A. Filipov, V. Schetinger, K. Raminger, N. Soursos, S. Zapke, and S. Miksch. Gone full circle: A radial approach to visualize event-based networks in digital humanities. *Visual Informatics*, 5(1):45–60, 2021.
- [25] J. Geweke. Inference and causality in economic time series models. *Handbook of econometrics*, 2:1101–1144, 1984.
- [26] C. W. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pp. 424–438, 1969.
- [27] C. W. Granger. Time series analysis, cointegration, and applications. *American Economic Review*, 94(3):421–425, 2004.
- [28] A. G. Hawkes. Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 33(3):438–443, 1971.
- [29] K. Hlaváčková-Schindler, M. Paluš, M. Vejmelka, and J. Bhattacharya. Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports*, 441(1):1–46, 2007.
- [30] X. Huang, Y. Zhao, C. Ma, J. Yang, X. Ye, and C. Zhang. TrajGraph: A graph-based visual analytics approach to studying urban network centralities using taxi trajectory data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):160–169, 2016.
- [31] Z. Jin, S. Guo, N. Chen, D. Weiskopf, D. Gotz, and N. Cao. Visual causality analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1343–1352, 2021.
- [32] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001.
- [33] N. R. Kadaba, P. Irani, and J. Leboe. Visualizing causal semantics using animations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1254–1261, 2007.
- [34] M. Kay and J. Heer. Beyond weber’s law: A second look at ranking visualizations of correlation. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):469–478, 2016.
- [35] C. Lee, Y. Kim, S. Jin, D. Kim, R. Maciejewski, D. S. Ebert, and S. Ko. A visual analytics system for exploring, monitoring, and forecasting road traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3133–3146, 2020.
- [36] J. Li, S. Chen, K. Zhang, G. L. Andrienko, and N. V. Andrienko. COPE: Interactive exploration of co-occurrence patterns in spatial time series. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2554–2567, 2019.
- [37] Q. Li, Q. Q. Liu, C. F. Tang, Z. W. Li, S. C. Wei, X. R. Peng, M. H. Zheng, T. J. Chen, and Q. Yang. Warehouse Vis: A visual analytics approach to facilitating warehouse location selection for business districts. *Computer Graphics Forum*, 39(3):483–495, 2020.
- [38] X. Li, Y. Cheng, G. Cong, and L. Chen. Discovering pollution sources and propagation patterns in urban area. In *Proc. of ACM SIGKDD*, pp. 1863–1872, 2017.
- [39] B. Liao, J. Zhang, C. Wu, D. McIlwraith, T. Chen, S. Yang, Y. Guo, and F. Wu. Deep sequence learning with auxiliary information for traffic prediction. In *Proc. of ACM SIGKDD*, pp. 537–546, 2018.
- [40] D. Liu, D. Weng, Y. Li, J. Bao, Y. Zheng, H. Qu, and Y. Wu. SmartAdP: Visual analytics of large-scale taxi trajectories for selecting billboard locations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):1–10, 2017.
- [41] D. Liu, P. Xu, and L. Ren. TPFlow: Progressive partition and multidimensional pattern extraction for large-scale spatio-temporal data analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1–11, 2019.
- [42] L. Liu, H. Zhang, J. Liu, S. Liu, W. Chen, and J. Man. Visual exploration of urban functional zones based on augmented nonnegative tensor factorization. *Journal of Visualization*, 24(2):331–347, 2021.
- [43] Q. Q. Liu, Q. Li, C. F. Tang, H. Lin, X. Ma, and T. Chen. A visual analytics approach to scheduling customized shuttle buses via perceiving passengers’ travel demands. In *Proc. of IEEE VIS*, pp. 76–80, 2020.
- [44] A. C. Lozano, H. Li, A. Niculescu-Mizil, Y. Liu, C. Perlich, J. Hosking, and N. Abe. Spatial-temporal causal modeling for climate change attribution. In *Proc. of ACM SIGKDD*, pp. 587–596, 2009.
- [45] A. Malik, R. Maciejewski, N. Elmquist, Y. Jang, D. S. Ebert, and W. Huang. A correlative analysis process in a visual analytics environment. In *Proc. of IEEE VAST*, pp. 33–42, 2012.
- [46] V. Matasov, L. Belelli Marchesini, A. Yaroslavtsev, G. Sala, O. Fareeva, I. Seregin, S. Castaldi, V. Vasenev, and R. Valentini. Lot monitoring of urban tree ecosystem services: Possibilities and challenges. *Forests*, 11(7), 2020.
- [47] F. Miranda, H. Doraiswamy, M. Lage, K. Zhao, B. Gonçalves, L. Wilson,

- M. Hsieh, and C. T. Silva. Urban Pulse: Capturing the rhythm of cities. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):791–800, 2017.
- [48] T. Munzner. *Visualization Analysis and Design*. A.K. Peters visualization series. A K Peters, 2014.
- [49] Q. V. Nguyen, N. Miller, D. Arness, W. Huang, M. L. Huang, and S. Simoff. Evaluation on interactive visualization data with scatterplots. *Visual Informatics*, 4(4):1–10, 2020.
- [50] T. Ortner, J. Sorger, H. Steinlechner, G. Hesina, H. Piringer, and M. E. Gröller. Vis-A-Ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. *IEEE Transactions on Visualization and Computer Graphics*, 23(2):1139–1151, 2017.
- [51] C. Palomo, Z. Guo, C. T. Silva, and J. Freire. Visually exploring transportation schedules. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):170–179, 2016.
- [52] J. Pearl et al. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009.
- [53] M. Pi, H. Yeon, H. Son, and Y. Jang. Visual cause analytics for traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2186–2201, 2021.
- [54] D. A. Pierce and L. D. Haugh. Causality in temporal systems: Characterization and a survey. *Journal of econometrics*, 5(3):265–293, 1977.
- [55] D. Qu, X. Lin, K. Ren, Q. Liu, and H. Zhang. AirExplorer: Visual exploration of air quality data based on time-series querying. *Journal of Visualization*, 23(6):1129–1145, 2020.
- [56] J. D. Ramsey, M. Glymour, R. Sanchez-Romero, and C. Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics*, 3(2):121–129, 2017.
- [57] J. Runge, P. Nowack, M. Kretschmer, S. Flaxman, and D. Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11):eaau4996, 2019.
- [58] R. Scheepens, C. Hurter, H. van de Wetering, and J. J. van Wijk. Visualization, selection, and analysis of traffic flows. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):379–388, 2016.
- [59] A. K. Seth, A. B. Barrett, and L. Barnett. Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297, 2015.
- [60] L. Shi, C. Huang, M. Liu, J. Yan, T. Jiang, Z. Tan, Y. Hu, W. Chen, and X. Zhang. UrbanMotion: Visual analysis of metropolitan-scale sparse trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [61] C. A. Sims. Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, pp. 1–48, 1980.
- [62] C. Tominski, H. Schumann, G. L. Andrienko, and N. V. Andrienko. Stacking-based visualization of trajectory attribute data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2565–2574, 2012.
- [63] H. Wang, Y. Lu, S. T. Shutter, M. Steptoe, F. Wang, S. Landis, and R. Maciejewski. A visual analytics framework for spatiotemporal trade network analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):331–341, 2019.
- [64] J. Wang and K. Mueller. The Visual Causality Analyst: An interactive interface for causal reasoning. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):230–239, 2016.
- [65] J. Wang and K. Mueller. Visual causality analysis made practical. In *Proc. of IEEE VAST*, pp. 151–161, 2017.
- [66] J. Wang, J. Wu, A. Cao, Z. Zhou, H. Zhang, and Y. Wu. Tac-Miner: Visual tactic mining for multiple table tennis matches. *IEEE Transactions on Visualization and Computer Graphics*, 27(6):2770–2782, 2021. doi: 10.1109/TVCG.2021.3074576
- [67] J. Wang, K. Zhao, D. Deng, A. Cao, X. Xie, Z. Zhou, H. Zhang, and Y. Wu. Tac-Simur: Tactic-based simulative visual analytics of table tennis. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):407–417, 2020.
- [68] M. Wang, J. E. Heskovich, L. House, N. F. Polys, and C. North. Bridging cognitive gaps between user and model in interactive dimension reduction. *Visual Informatics*, 5(2):13–25, 2021.
- [69] Y. Wang, H. Liang, X. Shu, J. Wang, K. Xu, Z. Deng, C. D. Campbell, B. Chen, Y. Wu, and H. Qu. Interactive visual exploration of longitudinal historical career mobility data. *IEEE Transactions on Visualization and Computer Graphics*, 2021.
- [70] Y. Wang, T.-Q. Peng, H. Lu, H. Wang, X. Xie, H. Qu, and Y. Wu. Seek for success: A visualization approach for understanding the dynamics of academic careers. *IEEE Transactions on Visualization and Computer Graphics*, 2022.
- [71] D. Weng, R. Chen, Z. Deng, F. Wu, J. Chen, and Y. Wu. SRVis: Towards better spatial integration in ranking visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):459–469, 2018.
- [72] D. Weng, C. Zheng, Z. Deng, M. Ma, J. Bao, Y. Zheng, M. Xu, and Y. Wu. Towards Better Bus Networks: A visual analytics approach. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):817–827, 2021.
- [73] D. Weng, H. Zhu, J. Bao, Y. Zheng, and Y. Wu. HomeFinder Revisited: Finding ideal homes with reachability-centric multi-criteria decision making. In *Proc. of ACM CHI*, p. 247, 2018.
- [74] Y. Wu, Z. Chen, G. Sun, X. Xie, N. Cao, S. Liu, and W. Cui. Stream-Explorer: A multi-stage system for visually exploring events in social streams. *IEEE Transactions on Visualization and Computer Graphics*, 24(10):2758–2772, 2018.
- [75] Y. Wu, D. Weng, Z. Deng, J. Bao, M. Xu, Z. Wang, Y. Zheng, Z. Ding, and W. Chen. Towards better detection and analysis of massive spatiotemporal co-occurrence patterns. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [76] X. Xie, F. Du, and Y. Wu. A visual analytics approach for exploratory causal analysis: Exploration, validation, and applications. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1448–1458, 2021.
- [77] X. Xie, J. Wang, H. Liang, D. Deng, S. Cheng, H. Zhang, W. Chen, and Y. Wu. PassVizor: Toward better understanding of the dynamics of soccer passes. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1322–1331, 2021.
- [78] C. Xiong, J. Shapiro, J. Hullman, and S. Franconeri. Illusion of causality in visualized data. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):853–862, 2020.
- [79] H. Xu, M. Farajtabar, and H. Zha. Learning granger causality for hawkes processes. In *Proc. of ICML*, vol. 48, pp. 1717–1726, 2016.
- [80] S. Ye, Z. Chen, X. Chu, Y. Wang, S. Fu, L. Shen, K. Zhou, and Y. Wu. ShuttleSpace: Exploring and analyzing movement trajectory in immersive visualization. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):860–869, 2021.
- [81] C. Yen, A. G. Parameswaran, and W. Fu. An exploratory user study of visual causality analysis. *Computer Graphics Forum*, 38(3):173–184, 2019.
- [82] L. Ying, T. Tang, L. Luo, L. Shen, X. Xie, L. Yu, and Y. Wu. GlyphCreator: Towards example-based automatic generation of circular glyphs. *IEEE Transactions on Visualization and Computer Graphics*, 2022.
- [83] K. Zhang and S. Batterman. Air pollution and health risks due to vehicle traffic. *Science of the Total Environment*, 450:307–316, 2013.
- [84] T. Zhang, Z. Chen, Z. Zhao, X. Luo, W. Zheng, and W. Chen. Faulttracer: interactive visual exploration of fault propagation patterns in power grid simulation data. *Journal of Visualization*, pp. 1–14, 2021.
- [85] Y. Zhao, H. Jiang, Q. Chen, Y. Qin, H. Xie, Y. Wu, S. Liu, Z. Zhou, J. Xia, and F. Zhou. Preserving minority structures in graph sampling. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1698–1708, 2021.
- [86] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Urban Computing: Concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology*, 5(3):38:1–38:55, 2014.
- [87] Z. Zhou, C. Shi, X. Shen, L. Cai, H. Wang, Y. Liu, Y. Zhao, and W. Chen. Context-aware sampling of large networks via graph representation learning. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1709–1719, 2021.
- [88] H. Zhu, M. Zhu, Y. Feng, D. Cai, Y. Hu, S. Wu, X. Wu, and W. Chen. Visualizing large-scale high-dimensional data via hierarchical embedding of knn graphs. *Visual Informatics*, 5(2):51–59, 2021.
- [89] J. Y. Zhu, C. Zhang, H. Zhang, S. Zhi, V. O. K. Li, J. Han, and Y. Zheng. pg-Causality: Identifying spatiotemporal causal pathways for air pollutants with urban big data. *IEEE Transactions on Big Data*, 4(4):571–585, 2018.