

ManyWaves: A Visual Analytics Approach for Exploring Multivariate Spatiotemporal Data Based on the Graph Wavelet Transform

Giovani Valdrighi¹, Germain Garcia-Zanabria², Mauro Diaz³,
Luis Gustavo Nonato⁴, Nivan Ferreira⁵, and Jorge Poco³

¹Universidade Estadual de Campinas, ²University of Engineering and Technology, ³Getulio Vargas Foundation

⁴University of São Paulo, ⁵Universidade Federal de Pernambuco

giovani.valdrighi@ic.unicamp.br, ggarciaz@utec.edu.pe, dany.espino@fgv.br,
gnonato@icmc.usp.br, nivan@cin.ufpe.br, jorge.poco@fgv.br

Abstract

Georeferenced time series are central to domains such as urban planning and climate studies, yet insights remain difficult to obtain when patterns span multiple variables and scales. Analysts often seek **events**—abrupt changes that may co-occur across variables, space, and time. The Graph Wavelet Transform (GWT) is effective for detecting such changes; however, previous visualization tools have primarily applied it to univariate signals. We introduce *ManyWaves*, a visual analytics system that applies GWT separately to each variable and integrates the resulting event information in a coordinated multivariate environment. Each variable is modeled as a signal on a shared spatiotemporal graph, and its high-frequency wavelet coefficients characterize local spatial and temporal changes. A 1D convolutional autoencoder then learns a low-dimensional representation from these coefficients, which drives linked views: a projection of locations by event similarity, a heatmap summarizing when variables change, and spatial and time-series views for detailed inspection. Through three case studies—US flight disruptions, traffic and crime in São Paulo, and land-use change in the Brazilian Legal Amazon—we show how *ManyWaves* supports the exploration of multivariate events. Expert feedback suggests that the system’s event-oriented views and interactions support sense-making, while also highlighting the need for onboarding to interpret wavelet-based representations.

CCS Concepts

• *Human-centered computing* → *Visualization systems and tools; Visual analytics; Geographic visualization;*

1. Introduction

The growing availability of GPS-enabled sensors has made spatiotemporal data ubiquitous and valuable in applications such as climate studies and urban planning. Such data appear in many forms, including event records, trajectories, and georeferenced time series. In particular, georeferenced time series capture sequential measurements at fixed spatial locations. When measurements are multidimensional, extracting information from such data becomes complex, as relevant patterns may span spatial, temporal, and multivariate dimensions. Visual analytics systems have therefore become important tools for analyzing spatiotemporal data [AMST11, AAG03].

A common analysis approach is to identify abrupt changes, or **events**, such as a sudden increase in measured traffic after an accident. Event-oriented analysis focuses attention on relevant subsets of data, which is essential when dealing with large datasets. Spatiotemporal events can influence neighboring regions and occur at different spatial and temporal resolutions, limiting the effectiveness

of conventional querying [DFD*14]. Moreover, analysts cannot always anticipate an event’s scale in terms of measurement magnitude, spatial extent, or temporal duration. Relationships between variables may also emerge through co-occurring events: an event in one variable (e.g., heavy rain) may precede a spatially adjacent event in another (e.g., a traffic jam).

Existing work has not fully addressed the intersection of multivariate and spatiotemporal event analysis. Valdivia et al. [VDP*15] and Geng et al. [GGZZ22] leveraged the graph wavelet transform to detect multi-scale patterns in univariate spatiotemporal data, but their approaches do not directly support multivariate exploration. Liu et al. [LXR19] introduced TPFlow, a system designed to identify event co-occurrence across locations. Although it can track events with multiple variables, it relies on predefined rules for event classification and does not explicitly capture spatial extent.

We introduce *ManyWaves*, a visual analytics approach for change-oriented exploration of multivariate spatiotemporal data. We model the data as signals on a graph composed of spatial

and temporal connections. As deep learning approaches often require extensive labeled data and training, we instead leverage the Graph Wavelet Transform (GWT), a graph signal processing technique that decomposes a signal into frequencies [HVG11]. By focusing on frequencies associated with abrupt changes, we can identify events in both the spatial and temporal dimensions. Although GWT is designed for univariate signals, ManyWaves applies it independently to each variable and integrates the resulting coefficients through a visual analytics workflow tailored to multivariate exploration. Our approach, inspired by García-Zanabria et al. [GZRP*21], leverages an autoencoder with 1D convolutional layers to learn a low-dimensional representation from the multivariate wavelet coefficients. This encoded representation is visualized in a novel “Wavelet Encoding” view that plots regions based on their temporal event similarity, enabling analysts to identify clusters, outliers, and recurring event patterns. We demonstrate our approach using three real-world usage scenarios (US flights, São Paulo urban data, and Amazon deforestation) and formative expert feedback. In summary, our primary contributions are:

- A GWT-based methodology that encodes high-frequency event patterns to compare multivariate georeferenced time series.
- ManyWaves, a visual analytics system for event-guided exploration of multivariate spatiotemporal data.
- Three real-world use cases and expert feedback demonstrating the usefulness of the proposed approach for analyzing multivariate spatiotemporal events.

2. Related Works

Our approach applies visual analytics to multivariate spatiotemporal data, enhancing both the visualization and analysis of complex patterns. In this section, we discuss related works across three domains and compare them with our proposal. For comprehensive surveys on spatiotemporal data visualization, see [AMST11, KH13, LMW*17, MFS*23].

2.1. Visualization of Multivariate Spatiotemporal Data

Spatiotemporal data are organized along spatial, temporal, and attribute dimensions, requiring analysts to reason about *where*, *when*, and *what* changes occur [AAG03]. Andrienko and Andrienko provide a systematic framework for such analysis, including transformations, aggregation, clustering, and linked visual displays [AAG03, AA20].

Early work often focused on single-variable spatiotemporal data, combining spatial representations (*e.g.*, geographical maps) with animations [SHKW20, Cyb22], interactive 3D environments [Kra03, FSN20, BDA*17], coordinated views [FPV*13, LLW*20, TSH*14, PWMB24], or embedded views [MFS*23]. As data complexity grew, techniques evolved to support multiple variables. Maciejewski et al. [MRH*09] combine linked views and statistical models to identify spatiotemporal hotspots, Malik et al. [MME*12] support the exploration of lagged correlations among datasets, VAUD [CHW*17] enables querying heterogeneous urban data to reveal cross-domain relationships, and Bristle Maps [KMM*13] provide compact representations for multi-attribute spatiotemporal event data.

Despite these advances, integrating multiple variables while preserving spatial and temporal context remains challenging. Kim et al. [KCK17] describe common comparative strategies, including *juxtaposition*, *superimposition*, *interchangeable views*, and *explicit encoding*. ManyWaves follows this line of work but focuses specifically on event-oriented exploration: it surfaces abrupt changes across variables and links them to spatial and temporal context for detailed inspection.

2.2. Spatial Graph Visualization

Graphs are a natural abstraction for data with spatial relations, including road networks, transportation systems, and mobility. Visual analytics has therefore used spatial graphs to explore structure, attributes, and temporal dynamics [Rod05, VLKS*11, HSS15, SYPB21, FAM23]. For example, Wang et al. [WLY*13] constructed a jam propagation graph to analyze how congestion evolves over space and time. Von Landesberger et al. [VLBR*15] and Andrienko et al. [AAFW16] used spatial and temporal abstractions to reveal mobility patterns, while Filipov et al. [FSR*21] used dynamic spatial graphs to model historical events. These works demonstrate the value of graph abstractions for reasoning about spatial relations and temporal evolution. ManyWaves differs by using the graph not primarily as the object of visualization, but as the domain on which GWT extracts spatially and temporally localized changes explored through coordinated multivariate views.

2.3. Techniques for Spatiotemporal Event Detection

Identifying spatiotemporal patterns or events in multivariate data is challenging. Events may appear at different spatial and temporal scales, from long-duration changes in small areas to brief events covering large regions. Analytical techniques must therefore support robust pattern discovery for guided data exploration [DFD*14]. Research in spatial statistics has addressed part of this challenge. The Getis-Ord statistic [GO92] identifies spatial outliers by comparing signal values at a location with those in neighboring regions, but incorporating temporal information remains challenging. *Direct differencing* captures temporal change by subtracting values between consecutive timestamps [RAAKR05], but relies on threshold selection and may fail to robustly identify significant events [BP00].

More recent techniques employ computational topology and graph signal processing. Doraiswamy et al. [DFD*14] model spatiotemporal data as a time-varying scalar function on a network and use computational topology to identify persistent extreme points as events. Valdivia et al. [VDP*15] adopt a similar data model but use graph signal processing to detect abrupt local spatiotemporal variations at multiple scales. Li et al. [LCZ*19] guide spatiotemporal analysis by detecting co-occurrence patterns, *i.e.*, events occurring at two locations with temporal regularity. Although this yields a controllable system, it defines events only at specific positions and does not consider spatial neighborhoods. Liu et al. [LXR19] represent the data as a tensor and use tensor decomposition to highlight potentially interesting segments, but treat space and time as discrete indices without incorporating spatial distances or temporal trends into the latent representation.

Building on these advances, *ManyWaves* uses graph signal processing to detect spatially and temporally localized changes in multivariate data. GWT is applied independently to each variable, and the resulting coefficients are aggregated and transformed through an autoencoder to reveal similarities across locations. Thus, *ManyWaves* combines univariate event detection with coordinated multivariate views spanning space, time, and variables.

3. Analysis Tasks

We define a set of analysis tasks that guided the design of *ManyWaves*. These tasks were derived from a synthesis of analytical requirements in the spatiotemporal visual analytics literature.

We consider an **event** as an abrupt change in a time series' behavior [GS99]. Given the spatiotemporal nature of our data, we extend this notion to **spatiotemporal events**: pronounced changes that occur over contiguous spatial regions during a time interval. Despite the difficulty of establishing a unified definition of an event in the literature, this definition aligns with prior work in event-guided visual analytics of spatiotemporal data [DFD*14, VDP*15, GGZZ22].

The task structure follows the analytical framework of Andrienko et al. [AAG03], in which spatiotemporal analysis can be organized into queries based on space, time, or attribute dimensions. For example, fixing a region supports the analysis of *when* changes occur, fixing a time interval supports the analysis of *where* changes occur, and fixing both space and time supports the analysis of *which* attributes are involved. From these considerations, we derive the following tasks:

- T1 Identification of Candidate Events.** The tool should surface candidate events by indicating *when* and *where* abrupt changes occur, supporting the central goal of event-guided spatiotemporal exploration [DFD*14, VDP*15].
- T2 Multivariate Event Context.** For a selected event, the tool should summarize the behavior of the involved variables, enabling users to understand how each variable contributes to or responds to the event. This task corresponds to the attribute-focused (“*what*”) component of spatiotemporal analysis [AAG03].
- T3 Identification of Cross-Variable Synchronization.** The tool should reveal whether events across variables are temporally synchronized or closely aligned, and highlight where such relationships occur. This supports the analysis of multivariate co-occurrence patterns, as considered in systems such as TPFlow [LXR19] and COPE [LCZ*19].
- T4 Local Inspection and Validation.** The tool should provide detailed views for selected locations and time intervals, supporting expert validation of candidate events and hypothesis refinement through detail-on-demand [Shn03].

4. Graph Wavelet Transform for Spatiotemporal Data Analysis

Our methodology leverages the GWT for spatiotemporal event detection. Following prior work on graph-based models of spatiotemporal data [VDP*15, DVP*17], we represent each measurement at a




Figure 1: Example of the Graph Wavelet Transform on a spatiotemporal graph. The graph represents spatial locations across multiple timestamps, with node color encoding the signal value. For the marked nodes, bar plots show wavelet coefficients at different frequencies. Nodes near transitions between signal regions have larger high-frequency coefficients, indicating local spatial or temporal changes.

spatial location and timestamp as a node in a spatiotemporal graph. This defines a signal on a graph that the GWT decomposes into spectral components. Unlike approaches that treat spatial and temporal variation separately, this representation allows the GWT to capture local changes that are simultaneously spatial and temporal.

Given a signal on a graph, the transform yields, for each node, a vector of coefficients $[w_1, \dots, w_k]$ describing the node's response across frequency bands. Low-frequency coefficients capture smooth behavior among neighboring nodes, whereas high-frequency coefficients highlight abrupt local changes. We next describe the graph construction and the GWT computation.

4.1. Building a Spatiotemporal Graph

The GWT can be applied to georeferenced time series by constructing a graph that encodes spatial and temporal relations [VDP*15, GGZZ22]. We first build a spatial graph G_s , whose nodes represent the locations of the time series and whose edges connect spatially related locations. These spatial relations may be provided by the dataset, derived from spatial adjacency, or specified by domain experts. We then replicate G_s for each timestamp, obtaining $\{G_{s,1}, \dots, G_{s,T}\}$. Nodes in $G_{s,t}$ are connected to their counterparts in $G_{s,t-1}$ and $G_{s,t+1}$, except at temporal boundaries. To also capture spatial relationships across consecutive timestamps, nodes in $G_{s,t}$ are linked to the neighbors of their counterparts in adjacent timestamps. Fig. 1 illustrates these two temporal edge types: red edges connect the same spatial position over time, while blue edges connect a spatial position to its neighbors at adjacent timestamps.

The original time series is then viewed as a signal f over the resulting spatiotemporal graph G : the node representing location i at

time t takes the value measured at that location and timestamp. Applying the GWT to f yields coefficients that jointly reflect spatial and temporal dynamics.

4.2. GWT Fundamentals

Fig. 1 illustrates how GWT coefficients reflect local graph structure. Nodes are colored by signal value, and selected nodes show their wavelet coefficients as bar plots. Nodes inside homogeneous regions exhibit stronger low-frequency responses because their values are similar to those of their neighbors. Nodes near transitions between low- and high-signal regions exhibit stronger high-frequency responses, indicating abrupt local changes. Nodes that are locally homogeneous but distinct at a broader neighborhood scale may exhibit mixed-frequency responses.

We consider an undirected and unweighted graph $G = (V, E)$ with n nodes, adjacency matrix A , and diagonal degree matrix D , where $d_{ii} = \sum_{j=1}^n a_{ij}$. The Laplacian is defined as $L = D - A$, with eigen-decomposition $L = U\Lambda U^\top$, where the columns u_ℓ of U form the graph Fourier basis and Λ contains the non-negative eigenvalues λ_ℓ on the diagonal. Given a signal $f: V \rightarrow \mathbb{R}$, its Graph Fourier Transform (GFT) is $\hat{f}(\ell) = \sum_{i=1}^n u_\ell(i) f(v_i)$, which measures the strength of the frequency λ_ℓ in f , analogous to the classical Fourier transform for time series.

While the GFT provides a global frequency description, the GWT localizes frequency information at each node by applying a spectral kernel g at multiple scales. For a node v_i and scale s , the GWT is defined as

$$W_f(s, v_i) = \sum_{\ell=1}^n g(s\lambda_\ell) \hat{f}(\ell) u_\ell(i).$$

The scale parameter s controls which frequencies are emphasized; smaller s emphasizes higher frequencies. By varying s , we obtain, for each node, a multi-scale description $[W_f(s_1, v_i), \dots, W_f(s_k, v_i)]$. We use the cubic spline kernel from Hammond et al. [HVG11] (Eq. 69), which acts as a band-pass filter, and sample $k = 4$ scales logarithmically to cover the eigenvalue range. Fig. 2 shows the kernel function at four scales. The kernel has the form




Figure 2: Wavelet kernel g .

$$g(\lambda) = \begin{cases} \lambda^2, & \lambda < 1, \\ -5 + 11\lambda - 6\lambda^2 + \lambda^3, & 1 \leq \lambda \leq 2, \\ 4\lambda^{-2}, & \lambda > 2, \end{cases}$$

following prior wavelet-based visual analysis of graph data [VDP*15]. Each kernel will filter a specific frequency range. In the following sections, we refer to $W_f(s_k, v_i)$ as the high-frequency coefficient.

For further details on scale sampling, filter design, and efficient computation, we refer to [HVG11, DVP*17].

Method	ROC AUC $\uparrow$	Time (s) $\downarrow$
G_i^* statistics	0.753 ± 0.02	2.873 ± 0.42
Direct difference	0.767 ± 0.03	$< \mathbf{0.001} \pm \mathbf{0.000}$
GAE	0.503 ± 0.02	42.462 ± 10.24
GWT	$\mathbf{0.910} \pm \mathbf{0.01}$	0.042 ± 0.01

Table 1: ROC AUC of event detection techniques on a synthetic spatiotemporal dataset. Results are the mean and standard deviation over 10 repetitions.

4.3. Comparison with Alternative Methods

To illustrate the advantages of GWT for spatiotemporal event detection, we compare it with three unsupervised approaches: Getis-Ord G_i^* statistics [GO92], direct differencing [RAAKR05], and a graph autoencoder (GAE) [PLC*19]. Getis-Ord G_i^* identifies spatial hot spots at each timestamp but ignores temporal dependencies. Direct differencing captures temporal changes between consecutive timestamps but ignores spatial structure. The GAE encodes the spatiotemporal graph and uses reconstruction error as an anomaly score.

This synthetic experiment is intentionally univariate and is designed to isolate each method's ability to detect spatiotemporal events, independent of the multivariate integration handled later by our visual analytics workflow. We generate a synthetic dataset with 200 time series of length 20 at random positions in $[0, 1]^2$. Each time series follows a half-period sine curve. We then inject 20 events by selecting groups of spatially close nodes, increasing their sine frequency over a fixed duration with random phases, and adding Gaussian noise (illustrated in Appendix B of the supplementary material). The spatial graph G_s is constructed by connecting each location to its 5 nearest neighbors. We assign a label $y_{it} = 1$ if node i at time t participates in an event, and 0 otherwise. Each method returns a score $\hat{y}_{it}$ for node i at time t , with higher values indicating higher event likelihood; for GWT, this score is obtained from the normalized magnitude of the high-frequency wavelet coefficient. We evaluate performance using ROC AUC over 10 repetitions. Details of the GAE implementation and synthetic data generation are provided in Appendix B of the supplementary material.

Table 1 summarizes the results. GWT achieves the highest ROC AUC (0.910) with modest computation time. G_i^* and direct differencing reach ROC AUC values around 0.75, reflecting their inability to simultaneously capture spatial and temporal structure. The GAE yielded a ROC AUC of 0.503, performing no better than a random guess. Further analysis revealed that the model reconstructed both event and non-event node representations with small error. This is consistent with the GAE architecture: its graph convolutions aggregate information from spatial and temporal neighbors, and since each event affects a group of nearby nodes over several timestamps, the model learns to reconstruct as well as normal behavior. These observations remained consistent across preliminary experiments varying model capacity and training duration. Despite its simplicity, this experiment illustrates GWT's suitability for detecting local spatiotemporal events, in line with previous visual analytics work [VDP*15, GGZZ22].




Figure 3: Preprocessing pipeline of ManyWaves. From left to right: a spatiotemporal graph is built from the spatial relations between time series and their timestamps; the Graph Wavelet Transform is applied independently to each variable; a convolutional autoencoder then learns a low-dimensional representation from the time series of high-frequency wavelet coefficients.

5. Data Preparation

This section outlines the procedure for preparing data as input to ManyWaves. As shown in Fig. 3, the process has three steps: spatiotemporal graph construction, application of the GWT, and wavelet encoding. We consider n locations, each associated with a time series of T observations $\{x_{i,1}, \dots, x_{i,T}\}$, where $x_{i,t} \in \mathbb{R}^d$ is a d -dimensional feature vector for location i at time t . The process begins with a spatial graph G_s , whose nodes correspond to locations and whose edges encode spatial relations. The construction of G_s is application-dependent; in Sec. 8, we describe the spatial relations used in each case study.

5.1. Spatiotemporal Graph

This step constructs the spatiotemporal graph described in Sec. 4.1. A graph G is built by combining T copies of G_s and adding temporal edges between consecutive timestamps, as illustrated in Fig. 3. For each location i and timestamp t , the feature vector $x_{i,t}$ is attached to the corresponding node in G . This representation allows wavelet coefficients to capture both spatial events, through the structure of G_s , and spatiotemporal events, through temporal and cross-time neighborhood connections.

5.2. Graph Wavelet Transform

Using the graph G , we apply the GWT described in Sec. 4. Since the transform operates on one-dimensional signals, each variable is processed independently. For every node and variable, we compute wavelet coefficients across four scales.

The magnitude of wavelet coefficients depends on the original scale of each variable; therefore, normalization is required before comparing patterns across variables. Following Dal Col et al. [DVP*18], we normalize coefficients separately for each variable. Let $\mathcal{D}_p$ denote, for variable p , the set of wavelet coefficient vectors across all $n \times T$ node-time pairs. We take absolute values and apply robust scaling using the interquartile range, obtaining the set $\mathcal{D}_p$. Each scaled coefficient w at scale s is then log-transformed and globally normalized as $F_s(w) = \log(w + 1) / \max_{u_s \in \mathcal{D}_p} \log(u_s + 1)$.

5.3. Wavelet Encoding

We build a low-dimensional representation from the high-frequency coefficients to compare locations according to their event behavior across variables and time.

For each location i , we construct a sequence $\{\vec{w}_{i,1}, \dots, \vec{w}_{i,T}\}$, where $\vec{w}_{i,t} \in \mathbb{R}^d$ contains the highest-frequency coefficient of each variable at time t and location i . Locations with similar sequences exhibit abrupt changes in similar variables over comparable time intervals.

Inspired by previous approaches for encoding time series [KAA*21] and visualizing similarities [GZRP*21], we use a 1D convolutional autoencoder to encode these d -dimensional sequences. The encoder uses two 1D convolutional layers followed by a fully connected layer, producing a 64-dimensional embedding for each location; implementation details are provided in Appendix A of the supplementary material. We project these embeddings into a 2D space (right-most panel in Fig. 3). Our implementation uses t-SNE, which emphasizes local neighborhood structure, but other dimensionality reduction techniques are also suitable.

6. ManyWaves

ManyWaves supports exploratory analysis of spatiotemporal data by exploiting high-frequency wavelet coefficients (Fig. 4). In this section, we describe its visual components, interaction mechanisms, and implementation details.

6.1. Visual Components

We developed a set of linked visual components so that ManyWaves can support the tasks in Sec. 3. These components provide complementary views of the data: correlations between variables, a 2D projection of wavelet encodings, a heatmap of temporal changes, a spatial view, and a time series view.

Correlation Table. This panel, shown in Fig. 4 (A), provides an initial overview of potential relationships between variables (T2). Given n locations, T timestamps, and d variables, we compute the Pearson correlation for each pair of variables using all $n \times T$ samples, and display the resulting values in a $d \times d$ heatmap. Although this metric overlooks spatial and temporal relations, the panel can incorporate other metrics, such as spatial correlation [Mor50] or cross-correlation [Lew95].

Wavelet Encoding. This panel (Fig. 4 (B)) shows a 2D projection of the high-frequency wavelet encodings described in Sec. 5.3. Each point represents a location, and proximity indicates similarity in the occurrence of abrupt changes across variables and time. To convey spatial context, points are colored according to their geographical positions by mapping normalized coordinates to a 2D colormap with distinct corner colors [BSC*21] (Fig. 4 (B1)). To facilitate analysis of datasets with many spatial positions, the user can filter highlighted points with an adjustable threshold τ (default 0.8), while the remaining points are displayed in gray at low opacity. For each location, we compute the maximum high-frequency wavelet coefficient across time for each variable and highlight locations whose value is above the τ -quantile for at least one variable.




Figure 4: *ManyWaves* interface for the US Airports dataset. *ManyWaves* comprises five linked views: (A) Correlation Table of Pearson correlations between variables, (B) Wavelet Encoding of high-frequency patterns, (C) Heatmap of Changes summarizing temporal change per variable, (D) Map View, and (E) Time Series View. Clicking two heatmap cells (C) highlights the corresponding airports in (D) and (E).

Lower values of τ display more locations, whereas higher values focus on locations with the most pronounced changes. By revealing similarities among locations and emphasizing high-magnitude events, the Wavelet Encoding supports T1 and T3.

Heatmap of Changes. This panel (Fig. 4 (C)) summarizes how strongly each variable changes over time. For each variable and timestamp, we compute the average of the high-frequency wavelet coefficients across locations, after discarding near-zero responses below $\epsilon = 10^{-3}$. Higher averages indicate more pronounced changes in that variable at a given time. The heatmaps, one per variable, are horizontally aligned to facilitate detection of timestamps where multiple variables undergo abrupt changes simultaneously. Users can click a cell to select the corresponding variable and timestamp, highlighting in the Map View the locations that contribute most to that coefficient. By default, we show the highest-frequency wavelet band, but users can switch to other scales; lower-frequency coefficients reveal smoother spatial and temporal variations. The Heatmap of Changes supports T1 and T3.

Map View. Illustrated in Fig. 4 (D), this panel displays the spatial layout of the data, with locations represented as points or polygons depending on the domain. Locations are colored according to the value of a selected variable at the current timestamp, using a consistent color scale across time. Selections in other views highlight corresponding locations on the map, supporting detailed inspection of event locations (T4).

Time Series View. Shown in Fig. 4 (E), this panel presents line plots of the time series for each variable. To provide global context, we plot a gray band with the minimum and maximum values across all locations. When locations are selected in other views, their time series are overlaid on the gray band, enabling comparison of lo-

cal behavior with global trends. A synchronized temporal marker helps users relate patterns across the *Wavelet Encoding*, *Heatmap of Changes*, *Map View*, and *Time Series View*, supporting T4.

6.2. Interaction

All views are linked and respond to user interactions, supporting the analysis tasks in Sec. 3. Initial settings for the *Wavelet Encoding*, *Heatmap of Changes*, and *Map View* can be configured in a control panel.

The *Wavelet Encoding* and *Heatmap of Changes* panels support location selection. Clicking in the projection selects nearby locations in the 2D embedding, while clicking a heatmap cell selects locations whose wavelet coefficients exceed the threshold for the chosen variable and timestamp. The *Map View* highlights the selected locations, and the *Time Series View* shows their time series. Selections are cumulative across views. When multiple heatmap cells are selected, including cells from different variables, *ManyWaves* highlights locations that satisfy the criteria for all selected cells, supporting the inspection of co-occurring events.

6.3. Implementation Details

We implemented *ManyWaves* as a web-based system using a Python Flask server. The backend includes modules for graph construction, GWT computation, and data preprocessing, and uses the PyGSP library [DMPP17] to compute the GWT. All visual components are implemented in JavaScript using D3.js and Mapbox GL.

7. Visual Analysis of Events

To show how the *Heatmap of Changes* and *Wavelet Encoding* support interpreting multivariate spatiotemporal events, we evaluate




Figure 5: Synthetic example with two multivariable spatial events. One event affects features A and B with a two-timestamp lag, while the other affects features A and C with a four-timestamp lag. The Heatmap of Changes reveals event timestamps, and the Wavelet Encoding separates affected regions according to their multivariate event patterns.




Figure 6: Wavelet Encoding with varying latent dimensionality. Projections obtained for the first example of Sec. 7 with latent dimensions ranging from 4 to 1024.

them on synthetic datasets with known patterns. Additionally, we leverage these examples to perform a sensitivity analysis to assess the impact of varying the number of frequencies and latent dimensions. The examples cover three representative cases: 1) spatially localized events involving pairs of variables, 2) temporal events without stable spatial structure, and 3) spatiotemporal events occurring at different frequencies. Because the underlying patterns are known, these examples clarify what each view reveals and how




Figure 7: Synthetic example with temporal periodicity but no consistent spatial structure. Events occur periodically within each feature, but affected locations are randomly selected each time. The Heatmap of Changes reveals distinct temporal periodicities, whereas the Wavelet Encoding yields no clear spatial clusters.

wavelet-based summaries differ from visualizations based only on the original signal values.

All examples use a spatial graph G_S with 400 nodes arranged in a 20×20 grid and $T = 60$ timestamps. Each node is connected to its four immediate spatial neighbors, and each dataset contains three features. We generate a base signal using a linear trend from 0 at $t = 0$ to 1 at $t = 60$, add independent Gaussian noise $\epsilon \sim \mathcal{N}(0, 0.02)$, and then inject the synthetic events described below.

Multivariable Spatial Events. In the first example, we generate two spatially localized events in distinct circular regions. The first affects features A and B at the same location, with feature B delayed by two timestamps; the second affects features A and C at another location, with feature C delayed by four timestamps. Fig. 5 shows the event timestamps, Wavelet Encoding, average-signal heatmap, and Heatmap of Changes. The Heatmap of Changes highlights the timestamps at which each feature undergoes abrupt changes, whereas the average-signal heatmap mainly reflects the global linear trend. The Wavelet Encoding separates the two affected regions, showing that the proposed views can distinguish events that share one feature but differ in spatial location and multivariate composition. Fig. 6 shows the Wavelet Encoding obtained by varying the autoencoder’s latent dimension in $\{4, 16, 32, 64, 128, 1024\}$. In all cases, the locations that participate in the two events are separated from most remaining points. However, with 4 dimensions, additional groups appear. As the latent dimension is small, the representation fails to capture all relevant information. With a large latent dimension (1024), which exceeds the input size ($60 \times 3 = 180$ values), the encoder performs no effective compression, and the two event groups are displayed less clearly separated from the remaining data. This supports our choice of a moderate latent dimension for the case studies.

Temporal Events without Spatial Pattern. In the second example, we create periodic temporal events without consistent spatial structure. Each event is a bump lasting three timestamps: feature A repeats every six timestamps, feature B every twelve, and feature C every twenty. For each occurrence, affected locations are randomly selected. As shown in Fig. 7, the Wavelet Encoding does not produce well-separated spatial clusters, as expected, because locations do not share persistent spatial event patterns. In contrast, the Heatmap of Changes clearly reveals the distinct temporal periodicities of the three features.

Events with Distinct Frequency. In the third example, each fea-




Figure 8: Synthetic example with moving spatial events at distinct frequencies. Each feature contains a circular event moving across the grid at a different velocity. Faster motion produces stronger high-frequency coefficients, while slower motion is more visible at lower frequencies in the *Heatmap of Changes*.

ture contains a circular region moving from the top-left to the bottom-right of the grid, but at different velocities (Fig. 8). Feature A moves fastest, feature B has intermediate velocity, and feature C moves slowest. Because faster motion produces more abrupt changes across neighboring nodes and timestamps, we expect feature A to produce stronger high-frequency coefficients and feature C to be more visible at lower frequencies. The *Heatmap of Changes* confirms this behavior: feature A is strongest at high frequency over a short interval, feature B spans a longer interval with medium-to-high frequencies, and feature C is dominated by lower-frequency coefficients. Therefore, inspecting multiple wavelet scales helps distinguish abrupt events from smoother spatiotemporal changes.

Fig. 8 uses four frequencies, the configuration adopted in our case studies; Fig. 9 extends it to $\{8, 16, 32, 64\}$. Prior work using the GWT for spatiotemporal data analysis has employed either 4 [VDP*15] or 8 [GGZZ22] frequencies. A small number of frequencies keeps the analysis more interpretable while still capturing the main behavior of the phenomena. As shown in Fig. 9, the main interpretation remains stable across configurations: feature A shows a high-frequency event, feature B shows a medium-frequency event, and feature C shows a low-frequency event. The main effect of increasing the number of frequencies is that responses become distributed across more frequency bands. For example, the high-frequency response of feature A is split across multiple adjacent frequencies when more filters are used, which can make the heatmap harder to interpret without revealing substantially different events.

8. Case Studies

To illustrate ManyWaves in practice, we present three case studies on multivariate spatiotemporal data: US flight delays, traffic and crime in São Paulo, and land-use change in the Brazilian Amazon.




Figure 9: Heatmap of Changes with varying number of frequencies. A small number of frequencies is more interpretable while still capturing the relevant patterns from the signal.

8.1. US Flight Delays

We examine the relationship between severe weather and air-traffic disruptions in the United States by combining precipitation data with delays, cancellations, and diversions. This dataset has been used previously for spatiotemporal visualization [PFC19].

Dataset. We analyze US Airports Data [U.S] during the 2017 hurricane season (July 15–October 15), together with NOAA hurricane tracks [LB] and PRISM precipitation data [PRI]. Each of the 272 locations corresponds to an airport, and the time series is daily, with length $T = 93$. In the *Map View*, we display a Voronoi diagram derived from airport coordinates and build G_s by connecting adjacent airports in this diagram. We consider six variables: Hurricane occurrence, average Precipitation (mm), fractions of flights with Cancellations and Diversions, and average Arrival Delay and Departure Delay (minutes). For this dataset, we set $\tau = 0.6$. The *Correlation Table* in Fig. 4 (A) shows a strong correlation between Arrival Delay and Departure Delay, but only modest correlations between Pre-




Figure 10: Wavelet Encoding for the US Airports and Weather dataset. We compare two selections in the Wavelet Encoding: A1 (Southeast) and B1 (West). The associated time series views for Departure Delay and Cancellations show that airports in A1 experience a sharp peak in cancellations related to a hurricane, whereas airports in B1 present consistently higher levels of cancellations and delays over the entire period.

cipitation and Cancellations. This suggests that Pearson correlation alone does not fully capture localized weather-related disruptions.

Hurricanes. The Hurricane row in the *Heatmap of Changes* (Fig. 4 (C)) shows long periods with minimal high-frequency activity and short intervals with strong activity, consistent with the episodic behavior of storms. Three prominent windows, marked C1–C3, correspond to Tropical Storm Emily (C1), Hurricane Irma (C2), and Hurricane Nate (C3). Selecting these regions reveals increases in Delays, Cancellations, and Diversions in the *Time Series View*, while the *Map View* highlights East Coast airports with abrupt cancellations during Hurricane Irma, capturing the localized weather-related disruptions that the global correlation does not reveal.

Regional Event Patterns. The *Wavelet Encoding* (Fig. 10) reveals two notable groups. Selection A1 captures airports in the US Southeast, where Cancellations peak in late August during Hurricane Harvey, leading to a near-shutdown of flights at several airports (A2). Selection B1 highlights airports in the Western US, particularly California, which show persistently higher Cancellations and Departure Delay throughout the period (B2), including disruptions associated with California wildfires around mid-October. By leveraging the high-frequency coefficients, the *Wavelet Encoding* supports T2 and T3 by grouping locations according to the timing and recurrence of abrupt changes, separating a single major disruption (A1) from recurrent disruptions (B1) across multiple variables.

8.2. Analysing Traffic and Crime in São Paulo

This use case focuses on São Paulo, Brazil, a large metropolis with complex traffic patterns and notable crime rates.

Dataset. We use *Waze for Cities Data* [Waz], which contains user-generated alerts with GPS positions and timestamps, and phone robbery/theft records from the São Paulo Public Security Secretariat [Sec]. We associate alerts and crime reports with census

tracts, yielding 5,000 locations, and build G_s by connecting adjacent tracts. For each tract, we build time series with 12-hour resolution from December 1, 2019, to January 31, 2020 ($T = 124$). Each time series has six area-normalized variables: Phone Theft, Phone Robbery, Jam, Accident, Road Closed, and Weather Hazard. Given the large number of locations, we set $\tau = 0.97$ to highlight only the strongest events.

Traffic Events. The resulting visualizations are shown in Fig. 11. The *Correlation Table* (A) reveals a strong correlation between Jam and Weather Hazard. In the *Wavelet Encoding*, selection B1 highlights tracts around Amaral Gurgel Street and Presidente Maia Avenue; their time series (B2, left) show concurrent peaks in Jam and Weather Hazard around December 12. The *Heatmap of Changes* (C) reveals weekly patterns in Accident and Jam: weekends (C1) and the end-of-year holidays (C2) exhibit lower high-frequency coefficients, indicating more regular traffic conditions. Around mid-December, a band of high-frequency activity in Accident, Jam, and Weather Hazard (D) indicates a localized period of weather-related traffic disruption. The same event is captured by the B1 cluster, illustrating how ManyWaves links temporal summaries and regional patterns.

Traffic and Crime Co-occurrence. On December 31 and January 25, the *Heatmap of Changes* shows elevated coefficients for Phone Theft and Road Closed (Fig. 11 (E1)). Selecting the December 31 cells for both variables makes ManyWaves compute the intersection of the selected locations. The resulting tracts (E2) concentrate along Paulista Avenue, a popular destination for public celebrations that is closed to vehicles during New Year's Eve and the city's anniversary. This relationship is masked in global statistics, where standard Pearson correlation reports a negligible value ($r = 0.01$, Fig. 11 (A)). The *Heatmap of Changes* successfully surfaces this localized co-occurrence, which the *Time Series View* confirms through concurrent peaks in both Road Closed and Phone Theft. Although neither variable reaches its global maximum on December 31, the GWT is well suited to identifying such an event, as its mechanism emphasizes localized changes in signal value across neighboring timestamps and tracts rather than global extremes.

Alternative Configurations. Fig. 12 shows an alternative configuration of the *Heatmap of Changes* that displays all four wavelet scales. Crime variables (Phone Theft and Phone Robbery) present strong coefficients at the highest frequency, confirming they are events localized with short duration that fade when evaluated over larger spatiotemporal scales. Conversely, Jam and Weather Hazard span both high and low frequencies: high frequencies capture localized events (e.g., flooded road segments), whereas lower frequencies reveal systemic congestion propagating across neighboring tracts over multiple consecutive intervals.

8.3. Brazilian Legal Amazon Land Use

We analyze land-use and land-cover dynamics in municipalities of the Brazilian Legal Amazon (BLA). The Amazon rainforest plays a key role in the global climate [ACF*23, Ami20] but has been increasingly affected by the expansion of economic activities [SJPC*21].




Figure 11: Analyzing traffic and crime in São Paulo. Data from 5,000 census tracts are analyzed across six traffic-, weather-, and crime-related variables. The Wavelet Encoding (top) groups tracts by high-frequency patterns; selection B1 highlights a cluster of tracts around Amaral Gurgel Street and Presidente Maia Avenue with concurrent Jam and Weather Hazard activity. The Heatmap of Changes (bottom left) shows the average high-frequency coefficient per variable and timestamp: C1/C2 mark weekends and holidays, D marks a period of heavy rain, and E1 marks New Year's Eve, when Road Closed and Phone Theft jointly spike along Paulista Avenue.




Figure 12: Heatmap of Changes with multiple frequencies for São Paulo. Alternative configuration of the Heatmap of Changes showing all four wavelet scales. Jam and Weather Hazard exhibit strong activity at both low and high frequencies, indicating a mixture of gradual and abrupt events, whereas crime-related variables (Phone Theft and Phone Robbery) show prominent activity mainly at the highest frequency.

Dataset. We use yearly land-use classifications from Map-Biomas [Map22] for 1985–2021 (25 classes at 30 m resolution). For each of the 737 municipalities in the BLA, we compute the yearly change rate of area for each class, yielding a multivariate time series. The change rate is computed as the year-to-year difference in class area within each municipality. We focus on nine representative classes: Forest Formation, Cerrado, Mangrove, Pasture, Soy, Sugar Cane, Mining, Urban In-

frastructure, and Mosaic Agriculture. We use the default threshold $\tau = 0.8$.

Effect of Regulations. The Correlation Table (Fig. 13 (A)) shows a strong negative correlation between Pasture and Forest Formation, consistent with forest loss driven by cattle ranching. We display the Heatmap of Changes with four frequencies. Different land-use categories show changes at different spatiotemporal scales. Soy, Mining, and Cerrado display dominant low-frequency coefficients, while Mangrove and Sugar Cane display high-frequency ones. This relates to the expansion of the agricultural frontier [NSA06] and to whether specific land uses were introduced in localized municipalities (high-frequency signals) or across extensive areas and over longer durations (low-frequency signals). The soy signal shows two low-frequency events. The first, from 1995 to 2004, corresponds to positive signal values (an increase in soy area), and the second, from 2006 to 2010, to negative values; the gap between them coincides with the Soy Moratorium, signed in 2006 [GRM*15]. After 2013, Sugar Cane and Soy show increased high-frequency and low-frequency activity (B), respectively, while Mining shows low-frequency activity starting in 2010 (C). These three changes coincide with weakened forest-protection policies and rising deforestation [SJPC*21].

Fragmentation of the Deforestation Process. The Wavelet Encoding (Fig. 14) groups municipalities by land-use trajectories. Selection A1 corresponds to municipalities with high Urban In-




Figure 13: Correlation and temporal summary for the Brazilian Legal Amazon (BLA) dataset. (A) Correlation Table between land use classes. The Heatmap of Changes displays temporal trends across four frequencies; selections B and C mark periods of increased Soy/Sugar Cane and Mining activity, respectively, coinciding with weakened forest-protection policies in recent years.

frastructure change early in the period, mainly in Maranhão (A2). Selection B1 identifies municipalities in the Central-West and along the Amazon River, where Soy rates were high before 2005 and often accompanied by Sugar Cane; the corresponding map shows spatially contiguous areas of intense soy cultivation (B2). Selection C1 identifies municipalities with strong Soy activity after 2014, which appear more spatially scattered (C2). These patterns align with reported increases in deforestation following relaxed forest protection [SL18, SJPC*21]. Comparing B2 and C2, early soy-related change forms large contiguous patches, whereas later change is more fragmented, matching prior observations of deforestation fragmentation [RSJE12]. Thus, the Wavelet Encoding separates locations that change in the same variable (Soy) but at different times and with distinct spatial structures.

9. Expert Feedback

To collect formative feedback on the usefulness and limitations of ManyWaves, we invited three external researchers with 3–12 years of experience in spatiotemporal data analysis. The participants were not involved in the development of the system.

Each expert participated in a 40-minute individual semi-structured session. The session followed a common protocol: we first introduced the motivation and main concepts of the method-

ology, including the use of high-frequency graph wavelet coefficients to identify abrupt spatiotemporal changes. We then demonstrated ManyWaves using the São Paulo use case (Sec. 8.2). The walkthrough was organized around the tasks in Sec. 3: identifying events with the Heatmap of Changes and Wavelet Encoding (T1), inspecting the multivariate context of selected events in the Time Series View (T2), comparing simultaneous events across variables such as Phone Theft and Road Closed (T3), and validating selected locations through the Map and Time Series views (T4). After the walkthrough, we conducted an open discussion covering interaction usability, response time, the relevance of multivariate spatiotemporal analysis, the usefulness of São Paulo insights, the perceived added value relative to conventional exploration, limitations, and suggestions. Sessions were recorded and consolidated by grouping recurring comments into common themes.

Overall, experts reacted positively to the system. They particularly appreciated the linked interaction between the Wavelet Encoding and Heatmap of Changes, and how selections propagate to the Map and Time Series views. They found the example linking Phone Theft and Road Closed events (Fig. 11 (E1-E2)) to be a compelling demonstration of how ManyWaves integrates spatial, temporal, and multivariate information. One expert noted that effective use of ManyWaves requires some familiarity with wavelet-based analysis and graph-based modeling. We agree that interpreting wavelet coefficients has a learning curve, but the distinction between low- and high-frequency behavior can be explained qualitatively and supported with short onboarding material or tutorials. Experts also raised concerns about scaling to datasets larger than those used in our case studies, especially with many more timestamps or locations. Most of the computational cost lies in preprocessing (graph construction and GWT computation), which in our experiments completed within a few minutes on a modern workstation. The current system is a research prototype; optimizing data structures, caching strategies, and client-side rendering is an important avenue for improving scalability.

10. Limitations

Scalability of Visual Representations. ManyWaves can highlight spatiotemporal events across several variables, but some visual components do not scale well as dimensionality increases. In the Brazilian Legal Amazon usage scenario, the system produced useful insights with nine variables; however, the Correlation Table grows quadratically with the number of variables, and the Heatmap of Changes requires one row per variable and scale (e.g., $9 \times 4 = 36$ rows in Fig. 13). As a result, these views may become cluttered for high-dimensional datasets. In such cases, users can analyze subsets of variables, for example, by applying subspace clustering [TMF*12] or correlation-based feature selection [TSL01] to identify groups of variables for joint exploration.

Graph Construction. ManyWaves assumes a spatial graph G_s as input, and the resulting wavelet coefficients depend on how this graph is constructed. In our case studies, we used spatial adjacency because the target phenomena are geographically localized, such as storms, traffic disruptions, crime events, and land-use changes. However, other analytical goals may require different graph definitions. For example, studying delay propagation in air traffic




Figure 14: Wavelet Encoding for the BLA dataset with three selections. Selection A1 groups municipalities in Maranhão with early Urban Infrastructure activity (map A2). Selection B1 highlights municipalities along the Amazon River with high Soy values before 2005 (map B2), and selection C1 identifies more spatially fragmented Soy activity after 2014 (map C2).

would require a flight-connectivity graph rather than geographic adjacency. Thus, graph construction should be guided by the phenomenon under analysis, and sensitivity to alternative graph definitions remains an important direction for future work.

Identification of Relationships. ManyWaves provides visual components for exploring potential relationships among variables, but the discovery of such relationships currently relies on manual interaction. Users select regions and time intervals in the *Wavelet Encoding* and *Heatmap of Changes* and interpret the resulting patterns. Future work could automate the detection of multivariate events by deriving auxiliary signals that capture local correlations or similarities between pairs or groups of variables [NNN11], and then applying our methodology to those signals.

Projection and Visual Selection. The *Wavelet Encoding* uses a 2D projection of high-dimensional embeddings to support exploration. Although this view helps users identify candidate groups of locations, the projection should not be interpreted as an automatic clustering result. Dimensionality reduction may distort global distances, and visual selection can bias attention toward visually salient or spatially coherent groups. ManyWaves mitigates this by linking the projection to the *Map View* and *Time Series View*, allowing users to validate selected regions in the original spatial and temporal context. Algorithmic clustering in the high-dimensional embedding space is a promising extension, but its effectiveness depends on hyperparameter choices such as the number of clusters or distance thresholds.

Data Format and Preprocessing. Our approach was designed for georeferenced time series and assumes that: 1) each spatial position has measurements for the full analysis period, and 2) measurements are available at regular time intervals. These assumptions are necessary for constructing the spatiotemporal graph and are also used in prior work [VDP*15, GGZZ22]. In our case studies, datasets were preprocessed to fit this format, for example by aggregating events into fixed intervals and setting values to 0 when no observation was recorded. Such preprocessing choices should be made carefully because they define the spatial and temporal scale at which events can be detected. For instance, events occurring within minutes may be obscured by daily aggregation, whereas very fine temporal resolution may introduce noise or sparsity.

11. Conclusion and Future Work

ManyWaves applies the Graph Wavelet Transform to individual variables and uses a convolutional autoencoder to encode high-frequency patterns, enabling event-oriented exploration of multivariate spatiotemporal data through coordinated visual views. By linking the *Wavelet Encoding*, *Heatmap of Changes*, *Map View*, and *Time Series View*, the system helps analysts identify and relate events across space, time, and variables. Our case studies on flight delays, urban traffic and crime, and Amazon land use illustrate how ManyWaves supports domains ranging from urban planning to environmental monitoring. Expert feedback highlighted the usefulness of the linked interactions, while also noting the learning curve associated with interpreting wavelet-based representations.

Future work includes improving scalability with respect to the number of variables, locations, and timestamps. We also plan to explore automatic methods for suggesting relevant subsets of variables, regions, and graph constructions, as well as techniques for directly detecting multivariate events. A more extensive user study and lightweight onboarding materials are promising directions for reducing barriers to adopting wavelet-based visual analytics in practice.

Use of Generative AI During the preparation of this manuscript, AI tools (Gemini and Grammarly) were used for text polishing.

Acknowledgments

This work was partially supported by the Brazilian National Council for Scientific and Technological Development (CNPq, grants #311425/2023-2 and #313454/2026-4), the São Paulo Research Foundation (FAPESP grants #2022/09091-8 and #2021/07012-0), the Carlos Chagas Filho Foundation for Research Support of the State of Rio de Janeiro (FAPERJ, grant #E-26/210.585/2025), CONCYTEC-PROCIENCIA E041-2024-01 (grant #PE501087483-2024), and Fundação Getulio Vargas (FGV).

References

- [AA20] ANDRIENKO N., ANDRIENKO G.: Spatio-temporal visual analytics: a vision for 2020s. *Journal of Spatial Information Science*, 20 (2020), 87–95. 2
- [AAFW16] ANDRIENKO G., ANDRIENKO N., FUCHS G., WOOD J.: Revealing patterns and trends of mass mobility through spatial and temporal abstraction of origin-destination movement data. *IEEE transactions on visualization and computer graphics* 23, 9 (2016), 2120–2136. 2
- [AAG03] ANDRIENKO N., ANDRIENKO G., GATalsky P.: Exploratory spatio-temporal visualization: an analytical review. *Journal of Visual Languages & Computing* 14, 6 (2003), 503–541. Visual Data Mining. doi:10.1016/S1045-926X(03)00046-6. 1, 2, 3
- [ACF*23] ALBERT J. S., CARNAVAL A. C., FLANTUA S. G. A., LOHMANN L. G., RIBAS C. C., RIFF D., CARRILLO J. D., FAN Y., FIGUEIREDO J. J. P., GUAYASAMIN J. M., HOORN C., DE MELO G. H., NASCIMENTO N., QUESADA C. A., ULLOA C. U., VAL P., ARIEIRA J., ENCALADA A. C., NOBRE C. A.: Human impacts outpace natural processes in the amazon. *Science* 379, 6630 (2023), eabo5003. doi:10.1126/science.abo5003. 9
- [Ami20] AMIGO I.: When will the amazon hit a tipping point? *Nature* 578, 7796 (2020), 505–508. 9
- [AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: Visualization of time-oriented data. In *Human-Computer Interaction Series* (2011). URL: <https://api.semanticscholar.org/CorpusID:42472812>. 1, 2
- [BDA*17] BACH B., DRAGICEVIC P., ARCHAMBAULT D., HURTER C., CARPENDALE S.: A descriptive framework for temporal data visualizations based on generalized space-time cubes. In *Computer graphics forum* (2017), vol. 36, Wiley Online Library, pp. 36–61. 2
- [BP00] BRUZZONE L., PRIETO D. F.: Automatic analysis of the difference image for unsupervised change detection. *IEEE Transactions on Geoscience and Remote sensing* 38, 3 (2000), 1171–1182. 2
- [BSC*21] BUCHMÜLLER J. F., SCHLEGEL U., CAKMAK E., KEIM D. A., DIMARA E.: Spatialrugs: A compact visualization of space and time for analyzing collective movement data. *Computers & Graphics* 101 (2021), 23–34. 5
- [CHW*17] CHEN W., HUANG Z., WU F., ZHU M., GUAN H., MACIEJEWSKI R.: Vaud: A visual analysis approach for exploring spatio-temporal urban data. *IEEE transactions on visualization and computer graphics* 24, 9 (2017), 2636–2648. 2
- [Cyb22] CYBULSKI P.: An empirical study on the effects of temporal trends in spatial patterns on animated choropleth maps. *ISPRS International Journal of Geo-Information* 11, 5 (2022). doi:10.3390/ijgi11050273. 2
- [DFD*14] DORAISWAMY H., FERREIRA N., DAMOULAS T., FREIRE J., SILVA C.: Using topological analysis to support event-guided exploration in urban data. *IEEE Trans. on Visualization and Computer Graphics* 20, 12 (2014), 2634–2643. URL: <http://ieeexplore.ieee.org/document/6876004/>. 1, 2, 3
- [DMPP17] DEFFERRARD M., MARTIN L., PENA R., PERRAUDIN N.: Pygsp: Graph signal processing in python, Oct. 2017. doi:10.5281/zenodo.1003158. 6
- [DVP*17] DAL COL A., VALDIVIA P., PETRONETTO F., DIAS F., SILVA C. T., NONATO L. G.: Wavelet-based visual analysis for data exploration. *Computing in Science & Engineering* 19, 5 (2017), 85–91. 3, 4
- [DVP*18] DAL COL A., VALDIVIA P., PETRONETTO F., DIAS F., SILVA C. T., NONATO L. G.: Wavelet-based visual analysis of dynamic networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 8 (2018), 2456–2469. doi:10.1109/TVCG.2017.2746080. 5
- [FAM23] FILIPOV V., ARLEO A., MIKSCH S.: Are we there yet? a roadmap of network visualization from surveys to task taxonomies. In *Computer Graphics Forum* (2023), vol. 42, Wiley Online Library, p. e14794. 2
- [FPV*13] FERREIRA N., POCO J., VO H. T., FREIRE J., SILVA C. T.: Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2149–2158. 2
- [FSN20] FILHO J. A. W., STUERZLINGER W., NEDEL L.: Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524. doi:10.1109/TVCG.2019.2934415. 2
- [FSR*21] FILIPOV V., SCHETINGER V., RAMINGER K., SOURSOS N., ZAPKE S., MIKSCH S.: Gone full circle: A radial approach to visualize event-based networks in digital humanities. *Visual Informatics* 5, 1 (2021), 45–60. URL: <https://www.sciencedirect.com/science/article/pii/S2468502X21000012>, doi:10.1016/j.visinf.2021.01.001. 2
- [GGZZ22] GENG R., GAO Y., ZHANG H., ZU J.: Analysis of the spatio-temporal dynamics of covid-19 in massachusetts via spectral graph wavelet theory. *IEEE Transactions on Signal and Information Processing over Networks* 8 (2022), 670–683. doi:10.1109/TSPN.2022.3193252. 1, 3, 4, 8, 12
- [GO92] GETIS A., ORD J. K.: The analysis of spatial association by use of distance statistics. *Geographical Analysis* 24, 3 (1992), 189–206. doi:10.1111/j.1538-4632.1992.tb00261.x. 2, 4
- [GRM*15] GIBBS H. K., RAUSCH L., MUNGER J., SCHELLY I., MORTON D. C., NOOJPADY P., SOARES-FILHO B., BARRETO P., MICOL L., WALKER N. F.: Brazil's soy moratorium. *Science* 347, 6220 (2015), 377–378. doi:10.1126/science.aaa0181. 10
- [GS99] GURALNIK V., SRIVASTAVA J.: Event detection from time series data. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining* (1999), pp. 33–42. 3
- [GZRP*21] GARCIA-ZANABRIA G., RAIMUNDO M. M., POCO J., NERY M. B., SILVA C. T., ADORNO S., NONATO L. G.: Cripav: Street-level crime patterns analysis and visualization. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2021), 4000–4015. 2, 5
- [HSS15] HADLAK S., SCHUMANN H., SCHULZ H.-J.: A survey of multi-faceted graph visualization. In *EuroVis (STARs)* (2015), pp. 1–20. 2
- [HVG11] HAMMOND D. K., VANDERGHEYNST P., GRIBONVAL R.: Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* 30, 2 (2011), 129–150. doi:10.1016/j.acha.2010.04.005. 2, 4
- [KAA*21] KIRANYAZ S., AVCI O., ABDELJABER O., INCE T., GABBOUJ M., INMAN D. J.: 1d convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing* 151 (2021), 107398. doi:10.1016/j.ymssp.2020.107398. 5
- [KCK17] KIM K., CARLIS J. V., KEEFE D. F.: Comparison techniques utilized in spatial 3d and 4d data visualizations: A survey and future directions. *Computers & Graphics* 67 (2017), 138–147. 2
- [KH13] KEHRER J., HAUSER H.: Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Transactions on Visualization and Computer Graphics* 19, 3 (2013), 495–513. doi:10.1109/TVCG.2012.110. 2
- [KMM*13] KIM S., MACIEJEWSKI R., MALIK A., JANG Y., EBERT D. S., ISENBERG T.: Bristle maps: A multivariate abstraction technique for geovisualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 9 (2013), 1438–1454. 2
- [Kra03] KRAAK M.: The space-time cube revisited from a geovisualization perspective. In *ICC 2003 : Proceedings of the 21st International Cartographic Conference* (New Zealand, 2003), International Cartographic Association, pp. 1988–1996. 2
- [LB] LANDSEA C., BEVEN J.: The revised Atlantic hurricane database (HURDAT2). accessed: 2025-12-01. URL: <https://www.aoml.noaa.gov/hrd/hurdat>. 8

- [LCZ*19] LI J., CHEN S., ZHANG K., ANDRIENKO G., ANDRIENKO N.: Cope: Interactive exploration of co-occurrence patterns in spatial time series. *IEEE Transactions on Visualization and Computer Graphics* 25, 8 (2019), 2554–2567. doi:10.1109/TVCG.2018.2851227. 2, 3
- [Lew95] LEWIS J. P.: Fast normalized cross-correlation. In *Vision interface* (1995), vol. 10, pp. 120–123. 5
- [LLW*20] LI Q., LIN H., WEI X., HUANG Y., FAN L., DU J., MA X., CHEN T.: Maravis: Representation and coordinated intervention of medical encounters in urban marathon. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–12. 2
- [LMW*17] LIU S., MALJOVEC D., WANG B., BREMER P.-T., PASCUCCI V.: Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics* 23, 3 (2017), 1249–1268. doi:10.1109/TVCG.2016.2640960. 2
- [LXR19] LIU D., XU P., REN L.: Tpflo: Progressive partition and multidimensional pattern extraction for large-scale spatio-temporal data analysis. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 1–11. doi:10.1109/TVCG.2018.2865018. 1, 2, 3
- [Map22] MAPBIOMAS: *MapBiomass General “Handbook” Algorithm Theoretical Basis Document*, 2022. 10
- [MFS*23] MOTA R., FERREIRA N., SILVA J., HORGA M., LAGE M., CEFERINO L., ALIM U., SHARLIN E., MIRANDA F.: A comparison of spatiotemporal visualizations for 3D urban analytics. *IEEE Trans. on Visualization and Computer Graphics* 29, 1 (2023), 1277–1287. URL: <https://ieeexplore.ieee.org/document/9904453/>. 2
- [MME*12] MALIK A., MACIEJEWSKI R., ELMQVIST N., JANG Y., EBERT D. S., HUANG W.: A correlative analysis process in a visual analytics environment. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2012), IEEE, pp. 33–42. 2
- [Mor50] MORAN P. A. P.: Notes on continuous stochastic phenomena. *Biometrika* 37, 1/2 (1950), 17–23. 5
- [MRH*09] MACIEJEWSKI R., RUDOLPH S., HAFEN R., ABUSALAH A., YAKOUT M., OUZZANI M., CLEVELAND W. S., GRANNIS S. J., EBERT D. S.: A visual analytics approach to understanding spatiotemporal hotspots. *IEEE Transactions on Visualization and Computer Graphics* 16, 2 (2009), 205–220. 2
- [NNN11] NAGARAJ S., NATARAJAN V., NANJUNDIAH R. S.: A gradient-based comparison measure for visual analysis of multifield data. *Computer Graphics Forum* 30, 3 (2011), 1101–1110. doi:10.1111/j.1467-8659.2011.01959.x. 12
- [NSA06] NEPSTAD D. C., STICKLER C. M., ALMEIDA O. T.: Globalization of the amazon soy and beef industries: Opportunities for conservation. *Conservation Biology* 20, 6 (2006), 1595–1603. 10
- [PFC19] PAHINS C. A., FERREIRA N., COMBA J. L.: Real-time exploration of large spatiotemporal datasets based on order statistics. *IEEE Transactions on Visualization and Computer Graphics* 26, 11 (2019), 3314–3326. 8
- [PLC*19] PARK J., LEE M., CHANG H. J., LEE K., CHOI J. Y.: Symmetric graph convolutional autoencoder for unsupervised graph representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision* (2019), pp. 6519–6528. 4
- [PRI] PRISM CLIMATE GROUP, OREGON STATE UNIVERSITY: PRISM climate data. accessed: 2025-12-01. URL: <https://prism.oregonstate.edu>. 8
- [PWMB24] PICCOLOTTO N., WALLINGER M., MIKSCH S., BÖGL M.: Underground tubes: Exploring spatial data with multidimensional projections and set visualization. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2024), 196–206. 2
- [RAAKR05] RADKE R., ANDRA S., AL-KOFAHI O., ROYSAM B.: Image change detection algorithms: a systematic survey. *IEEE Transactions on Image Processing* 14, 3 (2005), 294–307. doi:10.1109/TIP.2004.838698. 2, 4
- [Rod05] RODGERS P.: Graph drawing techniques for geographic visualization. In *Exploring geovisualization*. Elsevier, 2005, pp. 143–158. 2
- [RSJE12] ROSA I. M., SOUZA JR C., EWERS R. M.: Changes in size of deforested patches in the brazilian amazon. *Conservation biology* 26, 5 (2012), 932–937. 11
- [Sec] SECRETARIA DA SEGURANÇA PÚBLICA DO ESTADO DE SÃO PAULO: Estatísticas – consultas. accessed: 2025-12-01. URL: <https://www.ssp.sp.gov.br/estatistica/consultas>. 9
- [SHKW20] SAMET H., HAN Y., KASTNER J., WEI H.: Using animation to visualize spatio-temporal varying covid-19 data. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Modeling and Understanding the Spread of COVID-19* (New York, NY, USA, 2020), COVID-19, Association for Computing Machinery, p. 53–62. doi:10.1145/3423459.3430761. 2
- [Shn03] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*. Elsevier, 2003, pp. 364–371. 3
- [SJPC*21] SILVA JUNIOR C. H., PESSÔA A. C., CARVALHO N. S., REIS J. B., ANDERSON L. O., ARAGÃO L. E.: The brazilian amazon deforestation rate in 2020 is the greatest of the decade. *Nature Ecology & Evolution* 5, 2 (2021), 144–145. 9, 10, 11
- [SL18] SILVA C., LIMA M.: Soy moratorium in mato grosso: Deforestation undermines the agreement. *Land Use Policy* 71 (2018), 540–542. doi:10.1016/j.landusepol.2017.11.011. 11
- [SYPB21] SCHÖTTLER S., YANG Y., PFISTER H., BACH B.: Visualizing and interacting with geospatial networks: A survey and design space. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 5–33. 2
- [TMF*12] TATU A., MAASS F., FÄRBER I., BERTINI E., SCHRECK T., SEIDL T., KEIM D.: Subspace search and visualization to make sense of alternative clusterings in high-dimensional data. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2012), pp. 63–72. doi:10.1109/VAST.2012.6400488. 11
- [TSH*14] TURKAY C., SLINGSBY A., HAUSER H., WOOD J., DYKES J.: Attribute signatures: Dynamic visual summaries for analyzing multivariate geographical data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2033–2042. 2
- [TSL01] TUCKER A., SWIFT S., LIU X.: Variable grouping in multivariate time series via correlation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 31, 2 (2001), 235–245. doi:10.1109/3477.915346. 11
- [U.S] U.S. DEPARTMENT OF TRANSPORTATION: On-time performance dataset. accessed: 2025-12-01. URL: https://transtats.bts.gov/DatabaseInfo.asp?QO_VQ=EFD&DB_URL=8. 8
- [VDP*15] VALDIVIA P., DIAS F., PETRONETTO F., SILVA C., NONATO L.: Wavelet-based visualization of time-varying data on graphs. In *Proc. of the 2015 IEEE Conf. on Visual Analytics Science and Technology* (Chicago, USA, 2015), VAST '15, pp. 1–8. URL: <http://ieeexplore.ieee.org/document/7347624/>. 1, 2, 3, 4, 8, 12
- [VLBR*15] VON LANDESBERGER T., BRODKORB F., ROSKOSCH P., ANDRIENKO N., ANDRIENKO G., KERREN A.: Mobilitygraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 11–20. 2
- [VLKS*11] VON LANDESBERGER T., KUIJPER A., SCHRECK T., KOHLHAMMER J., VAN WIJK J. J., FEKETE J.-D., FELLNER D. W.: Visual analysis of large graphs: state-of-the-art and future research challenges. In *Computer graphics forum* (2011), vol. 30, Wiley Online Library, pp. 1719–1749. 2
- [Waz] WAZE: Waze for cities data. accessed: 2025-12-01. URL: <https://www.waze.com/wazeforcities>. 9

[WLY*13] WANG Z., LU M., YUAN X., ZHANG J., VAN DE WETERING H.: Visual traffic jam analysis based on trajectory data. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2159–2168. doi:10.1109/TVCG.2013.228. 2

Appendix A: Implementation Details of Wavelet Encoding

As described in Sec. 5.3, we compute the *Wavelet Encoding* from the high-frequency wavelet coefficients of all variables. For each location i , we construct a sequence $\{\vec{w}_{i,1}, \dots, \vec{w}_{i,T}\}$, where $\vec{w}_{i,t} \in \mathbb{R}^d$ contains the highest-frequency coefficient of each variable at time t . These sequences are used as input to a 1D convolutional autoencoder.

The encoder consists of two 1D convolutional layers with kernel size 3 and padding 1. Each convolutional layer is followed by ReLU activation and 1D max pooling, which halves the temporal resolution. The first convolution produces 128 channels, and the second produces 64 channels. The resulting representation is flattened and passed through a fully connected layer that produces a 64-dimensional embedding for each location. The decoder mirrors the encoder structure and reconstructs the input sequence from the embedding. These embeddings are then projected to two dimensions with t-SNE for visualization in the *Wavelet Encoding* view.

Appendix B: Implementation Details of Quantitative Evaluation

The synthetic dataset used in Sec. 4.3 contains 200 georeferenced time series of length 20 and is displayed in Fig. S1. For a node v_{it} representing location i at time t , the base signal is defined as

$$f(v_{it}) = 2 \sin\left(\frac{\pi}{2} \cdot \frac{t}{20}\right).$$

We randomly sample 20 events. Each event has a start time $t_0 \in [0, 15]$, a duration $d \in [3, 8]$, and affects a randomly selected group of 1 to 15 spatially close nodes. For each event, we sample a random phase $k \in [0, \pi]$ and define the event signal as

$$g(v_{it}) = 0.7 \sin\left(k + \frac{(t - t_0)\pi}{d}\right).$$

The final signal is $s(v_{it}) = f(v_{it}) + g(v_{it}) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, 0.05)$.

The graph autoencoder is based on Park et al. [P19] and follows the implementation from <https://github.com/daniel-unyi-42/GALA>. The encoder consists of a graph convolutional layer with dimensions 1×8 , followed by a ReLU activation and a second graph convolutional layer with dimensions 8×16 . The latent dimension is 16, and the decoder follows the reverse architecture. The graph autoencoder is trained for 1000 epochs or until the decrease in loss is lower than 10^{-4} . The anomaly score is computed as $(x - x')^2$, where x is the signal value and x' is its reconstruction.

Reference

[P19] PARK J., LEE M., CHANG H. J., LEE K., CHOI J. Y.: Symmetric graph convolutional autoencoder for unsupervised graph representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 6519–6528.




Figure S1: First five timestamps of the synthetic dataset used in the GWT comparison. Circles encode signal values at each node; squares mark nodes that participate in an event. While all nodes' signals increase over time, event nodes exhibit atypical behavior within their local neighborhoods, which GWT is designed to capture.