Information Systems 139 (2026) 102712

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Enhancing urban data exploration: Layer Toggling and Visibility-Preserving Lenses for multi-attribute spatial analysis

Karelia Salinas [a],*, Luis Gustavo Nonato [a], Jean-Daniel Fekete [b], Fernanda Bartolo dos Santos Saran [a]

[a] Institute of Mathematics and Computer Sciences, University of São Paulo (ICMC-USP), Avenida Trabalhador São-carlense, 400, São Carlos, 13566-590, São Paulo, Brazil

[b] Université Paris-Saclay, CNRS, Inria, LISN, Bâtiment 660, Rue Noetzlin, Gif-sur-Yvette, 91190, Île-de-France, France

ARTICLE INFO

Keywords:
Multi-attribute spatial data
Toggling interaction technique
Spatial data exploration

ABSTRACT

This manuscript proposes two novel interaction techniques for visualization-assisted exploration of urban data, namely, Layer Toggling and Visibility-Preserving Lenses. The former mitigates visual overload by organizing information into distinct layers while enabling multi-layer comparisons through controlled overlays. The technique supports focused analyses without sacrificing spatial context and enables users to quickly switch between layers through a dedicated physical button interface. Visibility-Preserving Lenses, on the other hand, dynamically adapt their size and transparency so that users can effectively examine dense spatial regions and temporal attributes in detail. Both techniques support urban data exploration and improve prediction.

Exploring urban data is essential for understanding complex phenomena related to crime, mobility, and residents' behavior and equally important is the ability to predict and explain how they evolve over time, supporting informed urban planning and policymaking. However, navigating urban data in all their complexity is challenging, often resulting in cognitive overload, loss of spatial context, and excessive visual clutter due to the many layers that must be examined simultaneously. Although layered visualizations aim to mitigate those challenges, they face limitations with occlusion and effortless comparisons across data layers. Additionally, interaction methods are typically confined to mouse-based controls, limiting the fluidity of dynamic exploration.

The visualization tool was validated through a comprehensive user study that measured user performance, cognitive load, and interaction efficiency across multiple devices. Using real-world data from São Paulo, including mobility patterns, climate conditions, and crime statistics, the way the approach enhances both exploratory and analytical tasks is demonstrated. The results also show how users perform when playing with different interactive devices, providing guidelines for future developments and improvements.

Contents

* Corresponding author.
E-mail address: kareliavs17@gmail.com (K. Salinas).

https://doi.org/10.1016/j.is.2026.102712
Received 15 June 2025; Received in revised form 17 January 2026; Accepted 24 February 2026

1. Introduction

Urban data visualization plays a crucial role for the understanding and predicting various phenomena, including accidents, crime, and mobility patterns. As cities become increasingly complex, the ability to visualize multivariate and multimodal data is essential for informed decision-making and effective resource allocation. In São Paulo, Brazil, the dynamics of taxi trips can intersect with criminal activities involving passengers, drivers, and external parties. Numerous predictive models have been developed towards anticipating such behaviors [1–6] and can be evaluated along three critical dimensions, namely, success rate (How Much?), underlying reasons for their success (Why?), and specific locations where they excel (Where?). While machine learning techniques address the "How Much" aspect through accuracy metrics, explainability methods tackle the "Why" question. This manuscript proposes leveraging 2D spatial data visualization that integrates multiple attributes and datasets for effectively addressing the "Where" component.

Recent advancements in predictive modeling have shifted focus from solely relying on historical data (e.g., criminal incidents) to incorporating a broader range of urban information, including socioeconomic factors, points of interest, and climate data. Such an integration gives rise to a comprehensive amount of data serving as feature vectors for machine learning models, which, from a visualization perspective, can be represented as a series of overlapping layers, each corresponding to a specific data type (e.g., origin–destination graphs for traffic and heatmaps for crime and real estate values).

Approaches to those data visualization often oscillate between two extremes, namely, space multiplexing and utilization of the entire screen space. The former enables multiple views, but limits spatial analysis, while full-screen views reveal correlations at the cost of visual clutter. Aggregation and filtering help, but may obscure important patterns; a detailed discussion is provided in Section 2.

This manuscript introduces a novel methodology according to which each data layer occupies the entire screen space, thereby avoiding spatial multiplexing and maintaining high spatial resolution and correlation between layers. Layer visibility and order are efficiently managed with the use of physical buttons or keys and dynamic lenses for filtering towards ensuring the layers remain readable. The approach incorporates *layer toggling*, in which each layer is assigned to a key on an auxiliary device, enabling users to maintain their view focus while easily comparing features, serving as a lightweight alternative to coordinated multiple views and interactive dashboards that fragment spatial context across separate panels. Additionally, visibility-preserving lenses are proposed as a central contribution of this study, including adaptive, density-based dynamic lenses for temporal and spatial data, in which the lens size and filtering behavior automatically adjust according to data density and temporal distribution. Those lenses explicitly preserve underlying spatial structures and correlations while selectively revealing relevant information, addressing visual clutter without sacrificing contextual awareness within a single unified view.

The development of a novel method for visual analysis is essential in this context, particularly in light of the prevalent tendency for division of screens into multiple views, which may hinder the visualization of the whole data and miss correlations between information layers. Moreover, when the output of predictive models is part of the analysis, identifying areas with poor predictions is vital for analyses of causes and improvements in model performance.

An extensive dataset focused on the city of São Paulo (denoted São Paulo for now on), the largest city in South America and which offers a rich diversity of open data, was used for the validation of the methodology, enabling the incorporation of varied information into the analytical framework.

The methodology consists of several key steps. The process begins by aggregating data with the use of CityHub library [7] and a machine-learning classification algorithm then predicts whether a taxi trip would involve a crime occurrence.

`kepler.gl` software [8] was used for a visual presentation of the datasets and prediction results, incorporating customized filtering and navigation functions, along with a "stream deck" button box for efficient layer toggling when compared to a baseline condition with no auxiliary device. Finally, the effectiveness of the approach was validated through an in-depth analysis of São Paulo, yielding valuable insights.

The approach was tested with users divided into several groups for evaluations of whether the use of layers aided urban analyses and assessment of differences in performance and efficiency with three different auxiliary devices for layer toggling, namely, mouse, keyboard shortcuts, and button box, as well as a no-device baseline condition. In a second phase, static versus dynamic filters were compared.

This user study demonstrated layered data visualization significantly enhances exploration and understanding; however, no significant differences were found among the various devices used when compared to each other. Moreover, dynamic filters facilitated tasks compared to static filters in terms of execution time.

The contributions of this study include:

- Proposal of a novel analysis method that organizes data into layers.
- Introduction of the Layer Toggling technique for navigation.
- Implementation of density-based dynamic filtering for spatial and temporal data.
- A visual evaluation of predictive system performance, highlighting areas of success and failure.
- User tests for assessments of the effectiveness of layered visualization in facilitating exploration and evaluating the impact of different devices on layer toggling and comparing the efficiency of static and dynamic filters through statistical analysis.

2. Related work

Urban data visualization often involves complex datasets that can lead to visual clutter, obscuring meaningful insights and effectively addressing that issue requires exploration of various taxonomies of urban data representation. Boyandin et al. [9] and Gu et al. [10] provided a comprehensive classification of data representation approaches, which can be broadly categorized into cartographic contexts and other diagrammatic forms.

Origin–destination, in which several representations and visualizations stand out, is a widely investigated type of urban data. Flow Maps, for instance, successfully encode multiple data components using flow symbols, enabling an efficient visualization of both directed and undirected movements [11]. In contrast, Thematic Maps do not employ line symbols to indicate flow; instead, they utilize alternative visual variables to represent the flow between origins and destinations [12]. Other diagrammatic approaches include Matrices, in which columns and rows correspond to origins and destinations, with entries indicating flows [10]. Node-Link Diagrams use node symbols for origins and destinations, while linear symbols represent flows, employing various layouts such as Arc Diagrams [13], Alluvial Diagrams, Circos [14], and Hive Plots [15].

Despite the advances, the cluttering problem remains a significant challenge in urban data visualization [16], encompassing several facets, including the cluttering problem itself, the Modifiable Area Unit Problem (MAUP), in which different aggregations may reveal distinct patterns, and the Normalization (or Size-Difference) Problem, which refers to the challenge of accurately representing data when the sizes of visual elements do not correspond proportionally to the underlying values they represent. Such discrepancy may lead to misinterpretations and skewed understandings of the data, thus hampering users from drawing accurate conclusions from the visualizations. Among the several techniques that have emerged towards combating clutter are filtering of Origin-Destination (OD) data [17] and optimization of symbols and layout through location aggregation [18], which utilizes spatial clustering [19], graph partitioning [17], and discretization in high-level administrative units. Additional methods include layout adjustment techniques [20], edge bundling algorithms [21], and interactive triangular irregular networks (TIN) modification [22]. On the other hand, those techniques can obscure significant data, smooth patterns, and lead to misinterpretations. Maintaining data integrity without altering the underlying information is crucial; however, it may can lead to visual overload, which is addressed in this work through the use of layers and dynamic filtering. The previous study [23] introduced a preliminary implementation of layer toggling and visibility-preserving lenses and the present one extends the discussion by providing a detailed explanation of data sources, preprocessing methods, the structural modeling process, and information on implementation. Additionally, the visualizations have been improved through enhanced tooltips, optimization of color contrast, refinement of user feedback within the button box, and a more noticeable distinction between activated and deactivated buttons. The previous Study reported on an experiment conducted with ten users, who tested three different devices for layer toggling and two filtering approaches, namely, dynamic and static, ultimately favoring the button box and the dynamic filter. The present study expands the number of participants, segments users into different groups, quantitatively measures their performance (including task completion time and click count), and refines the result analysis phase for better justifying design decisions. Moreover, it addresses the cluttering and spatial correlation problems with two interactive techniques, namely, layer toggling and visibility-preserving lenses.

2.1. Layer navigation

Although urban data visualization has achieved significant advancements in recent years, many methodologies have revealed critical limitations that adversely impact user experience. Techniques for layered navigation often fragment the display into multiple modules, resulting in cognitive overload and impeding a comprehensive understanding of the presented data. Kraak and Ormeling [24], Gao et al. [25], García et al. [2], and Salinas et al. [7] collectively emphasized subdividing the visual space restricts the overall analysis of attributes or datasets, hampering the identification of spatial correlations.

Some studies have focused on unifying heterogeneous urban data sources for adressing integration challenges. Wang et al. [26] introduced UrbanDataLayer, which fuses urban data with spatio-temporal base layers through a unified pipeline, improving dataset cohesion, but still facing issues of visual clarity at scale. Similarly, Sideris et al. [27] employed graph-based convolutional neural networks to model complex relationships in urban data, advancing analytical capabilities at the cost of computational complexity.

Other studies have prioritized reductions in visual clutter and improvements in user navigation. Bentlin [28] advocated for graphical reduction and selective removal of unnecessary layers towards highlighting meaningful relationships, despite risking oversimplifying data and omitting critical details. Zhou and Hsu [29] introduced interactive layering techniques for enhancing engagement; however, they may inadvertently increase fragmentation and distract from core analysis tasks.

The complexity of visualizing three-dimensional urban data adds further challenges. Miranda et al. [30] discussed the importance of direct associations between data layers in 3D environments for boosting user understanding, but warned of potential cognitive overload inherent to such representations, and Wagner et al. [31] explored immersive analytics via a space–time cube metaphor, enhancing spatial comprehension, yet hampering navigation due to its immersive design.

Finally, foundational perspectives of Dodge and Kitchin [32] highlight visualization's essential role in geographic information science, but do not offer concrete solutions for seamless multi-layer navigation.

Although it may introduce visual clutter and challenges related to data aggregation, leveraging the entire screen for a single visualization enhances the clarity of both local and global relationships among features.

Towards addressing those issues, a visualization system that facilitates *eye fixing*, eliminating the need to subdivide the display into multiple modules has been implemented and a unified screen with layered visualizations that provide a more coherent and less distracting navigation experience is proposed in this study. Furthermore, an *external button box* equipped with intuitive controls for layer manipulation has been incorporated, thus enabling users to interact with the data seamlessly without losing eye focus on their primary objectives. Such an innovative design significantly enhances user engagement with relevant information while reducing the cognitive burden typically associated with split displays.

2.2. Lenses

As discussed by Tominski et al. [33], lenses provide dynamic filtering capabilities so that users can focus on relevant data while minimizing clutter. Their application as a filtering mechanism in data visualization is prevalent; however, many contemporary approaches have deficiencies in adaptability and usability. Traditional filtering methods frequently fail to dynamically adjust to user needs, leading to information saturation. Bach and Carpendale [34] explored the potential of lenses to enhance data exploration, yet they do not sufficiently address the way the lack of adaptability can detrimentally affect user experience. Similarly, Keim and Ward [35] highlighted various filtering applications, but overlooked the necessity of maintaining a balanced data representation on-screen for avoiding overwhelming users.

Integrating lenses into urban data visualization can significantly enhance user experience by enabling dynamic filtering tailored to individual needs. Fischer and Hegarty [36] investigated the use of temporal lenses, contributing with valuable insights into time-dependent data visualization. However, a comprehensive solution that adapts to the evolving requirements of users remains elusive.

Adaptive filters that adjust to user preferences while ensuring a balanced presentation of data, thus preventing information overload were implemented in this study. The approach not only enhances usability, but also optimizes the comprehension of the visualized information. By merging advanced layered navigation techniques with adaptive lens functionality, the study represents a substantial advancement in urban data visualization, addressing the shortcomings of existing methodologies and incorporating user-centered design principles, and, therefore, improving data interpretation and fostering a more effective and engaging user experience.

3. Methodology

3.1. Design requirements and tasks

The development of the tool was rooted in extensive experience and work with several entities in the management of urban data, particularly in the manipulation of origin–destination datasets. Several pressing needs were identified during the presentation of information to transportation analysts who may lack expertise to navigate complex software systems. Towards addressing such challenges, the tool and its enhancements were guided by the following key requirements:

- **R1.** Maintenance of resolution by avoiding fragmentation of the screen into modules, ensuring all information is displayed in a single, cohesive visualization.
- **R2.** Use of cartographic maps for data representation, since such an approach is more intuitive for transportation professionals and facilitates a better understanding of spatial relationships.
- **R3.** Enhancement of user interaction with the tool for supporting eye fixation, thereby minimizing distractions from keyboard manipulation and manual movements.
- **R4.** Incorporation of filters necessary for the analysis of spatio-temporal data, enabling users to delve deeper into the insights provided by the data.
- **R5.** Spatial representation of the results of predictive models towards identification of failure zones, uncovering of correlations between layers, and highlight of the significance of features that contribute to predictive outcomes.

The tool has been designed to fulfill the following primary two primary tasks, each tailored to specific user types, as illustrated in Fig. 2.

- **T1. Exploration:**, which focuses on inexperienced users, such as regular citizens, who may wish to explore the conditions of specific neighborhoods for assessing their suitability for living. The functionality enables users to gain valuable insights into various urban factors that impact their quality of life; and
- **T2. Analysis:**, aimed at analyst users, including data analysts and specialists in urbanism, transportation, and law enforcement. These users examine the relationships among diverse datasets to identify patterns of abnormal behavior, such as crime and traffic issues. Additionally, the tool serves as a critical resource for evaluations of the effectiveness of predictive models.

In summary, the methodology presented not only addresses the technical requirements for an effective data visualization, but also considers the diverse users' needs, ensuring both regular users and analysts can derive meaningful insights from urban data. Table 1 shows the requirements and tasks to be observed and attended when data layers and the different functionalities to be implemented are managed.

3.2. Predictive modeling pipeline

The methodology integrates a predictive modeling pipeline that estimates the likelihood of crime occurrence associated with taxi trips for supporting the analysis task (T2) and fulfilling requirement R5. Due to the heterogeneous, temporal, and highly imbalanced nature of the data, the modeling pipeline was designed to ensure robustness, interpretability, and suitability for expert users, including transportation analysts and law enforcement professionals.

Prior to model induction, categorical attributes are transformed with the use of one-hot encoding, while ordinal attributes are mapped to numerical values. All features are normalized to the [0,1] range so as to ensure comparability across attributes with different scales. Given the temporal dependency of ride-hailing data, the dataset is split chronologically into training, testing, and validation sets for avoiding information leakage and better reflecting real-world deployment conditions.

Multiple learning algorithms with different inductive biases, including Logistic Regression, Decision Trees, Random Forests, Support Vector Machines, Multilayer Perceptrons, and Extreme Gradient Boosting (XGBoost) were systematically evaluated towards selection of an appropriate classifier and algorithm-level strategies (class weighting) and preprocessing techniques (e.g., Random Undersampling of the majority class) were considered so that class imbalance could be addressed. Model selection was performed through a grid search over combinations of preprocessing steps, classifiers, and hyperparameters, resulting in a large set of candidate models evaluated with the use of Gmean metric, which balances performance across majority and minority classes.

XGBoost combined with Random Undersampling achieved the best overall performance among all evaluated candidates and was, therefore, selected due to its ability to capture non-linear relationships, handle heterogeneous feature spaces, and maintain robust performance under severe class imbalance. Simpler baselines, such as Logistic Regression, were also tested; however, they showed inferior performance, particularly in capturing complex interactions among spatial, temporal, and contextual variables.

The classifier tends to underperform in spatial regions where the class distribution is close to be balanced (i.e., near a 1:1 ratio). In such regions, the lack of strong discriminative patterns limits the model's ability to confidently separate classes, leading to higher uncertainty in predictions. That behavior motivates the development of visual analysis tools that enable experts to inspect prediction outcomes spatially, identify failure zones, and relate model performance to underlying urban context, thereby closing the loop between predictive modeling and interactive visualization within a unified analytical workflow.

4. Kepler.gl tool

Kepler.gl project [37], an advanced data visualization and analysis platform that enables explorations of large-scale spatial and temporal

Fig. 1. Multisource urban databases are rendered as spatially aligned layers overlaid towards creating a comprehensive view so that each layer visualizes a specific type of spatial data using a dense representation. Each of the nine layers on the left is associated with a button in the button box, enabling users to interactively toggle its visibility on the screen and explore spatial correlations between layers, leveraging retinal persistence. Access can also be achieved through keyboard shortcuts corresponding to each layer number or via mouse interactions. The Visibility-Preserving Lens on the top-right dynamically adjusts the brush radius according to the graph link density at the mouse position, thereby enhancing graph legibility and serving as a spatial filtering mechanism. The size of the range sliders adapts dynamically to the local distribution for temporal data and animation can be employed to sweep across the entire range while maintaining controlled legibility. The layers interact and overlap, facilitating analyses of relationships among several types of data.

Table 1
Relationship between layers and resources with the requirements (R) they fulfill and the tasks (T) they address. See more details in Section 3.1.

Layer/Resource	R1	R2	R3	R4	R5	T1	T2
Crime Layer	✓	✓				✓	
Taxi Trips Layer	✓	✓				✓	
Weather Layer	✓	✓				✓	
Public Transportation Layer	✓	✓				✓	
Favelas Layer	✓	✓				✓	
Socioeconomic Layer	✓	✓				✓	
Graph Layer	✓	✓					✓
Hotspots Layer	✓	✓					✓
Prediction Layer	✓	✓			✓		✓
Layer Toggling	✓	✓	✓			✓	✓
Visibility-Preserving Lenses				✓		✓	✓
Correlation Matrix					✓		✓
Shapley Values					✓		✓

datasets, was the basis for the creation of the visualization. It emphasizes interactive tools that empower users to intuitively identify patterns and relationships within their data [37].

Kepler is constructed from modern web technologies, incorporation of frameworks such as React for the user interface and deck.gl, a data visualization library developed by Uber's Data Visualization team. It also leverages WebGL and react-mapbox-gl for rendering dynamic, high-performance visualizations. Such a robust architecture ensures Kepler can handle complex datasets with efficiency and responsiveness. Kepler.gl v2.5.0, an actively maintained version that continues to receive updates and enhancements from its creators, was used for the implementation. Several of its modules were customized, taking full advantage of the tool's open-source nature.

The key features of Kepler include:

- **Spatial Data Visualization:** Kepler enables the display of geospatial data on interactive maps, facilitating the identification of trends, anomalies, and geographic patterns.
- **Interactivity:** The platform enables users to manipulate visualizations in real-time, adjusting parameters and filters for exploring the data from multiple perspectives.
- **Integration of Multiple Data Sources:** Kepler supports the integration of diverse datasets, facilitating more comprehensive and contextualized analyses.
- **Ease of Use:** Designed for accessibility, Kepler's user interface is intuitive enough for both technical and non-technical users.
- **Modular Architecture:** Kepler's modular design facilitates both extension and integration of new features. Its backend is powered by Node.js, ensuring scalability and high performance when large datasets are being handled.
- **Open Source:** As an open-source project, Kepler encourages contributions from developers and enables customization to suiting specific needs.

The system is particularly valuable in fields such as urban planning, environmental management, and scientific research, where powerful data visualization tools can lead to better decision-making and groundbreaking discoveries.

4.1. Data

São Paulo data were collected from various sources and organized into distinct layers within the developed analytical tool. Each database is tailored to a specific data type and representation — as an example, geolocated data is depicted as points, a feature already integrated into Kepler. The data were categorized into two primary types, namely, dynamic and static. Dynamic data, such as crime statistics, taxi trips, and weather conditions, frequently change over time, Fig. 2 provides a comprehensive summary of data types, their respective representations across each layer, the way they are aggregated, the user expected to handle them, and the way the respective layer can be accessed in the different devices.

The following subsections provide a detailed description of each database, elucidating the source, visual representation, and methodology that integrates the data into the graph nodes for predictive tasks. The construction of the graph is based on the city's street map, where nodes represent street corners and edges represent streets connecting them.

4.1.1. Crime

Crime data were obtained from the Open Data Portal of the Secretary of Public Security of the State of São Paulo (www.ssp.sp.gov.br) and each crime record includes information such as geolocation (latitude and longitude), date and time of occurrence, and type of crime. This study focused specifically on crimes categorized as vehicle and mobile phone thefts that occurred in 2020. In the visualization, crime data are represented as geolocated points, with red points indicating vehicle thefts and blue points representing mobile phone thefts. Furthermore, crime records were assigned to the nodes of the street graph by the nearest corner approach, as proposed by Garcia-Zanabria et al. [2].

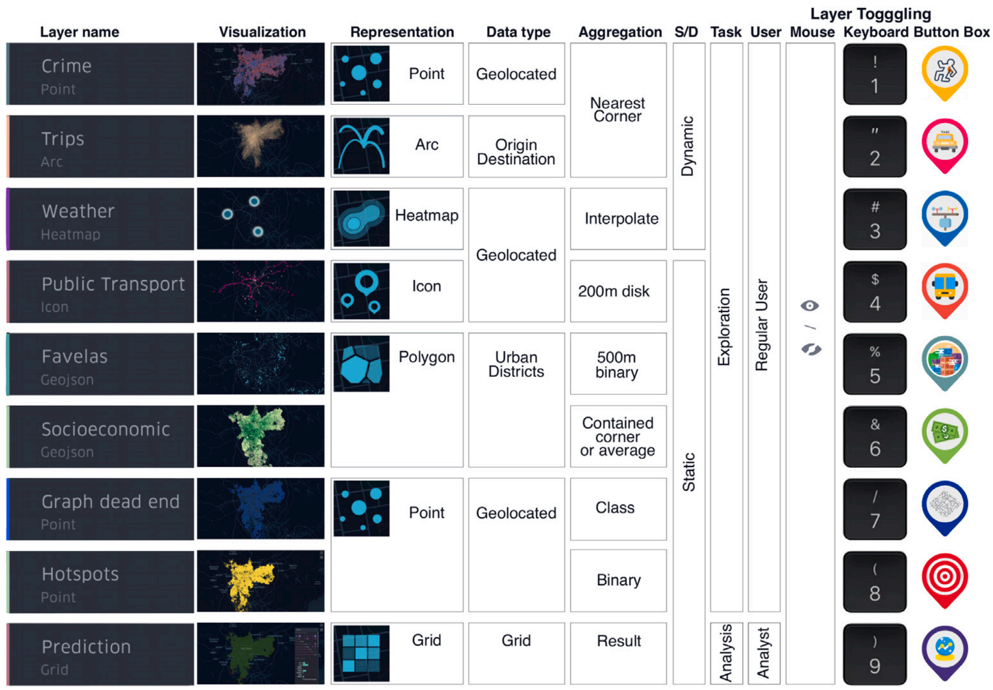

Fig. 2. The layers under consideration — each layer corresponds to a distinct database. The first column lists the buttons as they appear in the interface and the second displays the associated visual representation. The third column illustrates the data representations based on the available types in Kepler (e.g., points and arcs) and the fourth denotes data type within each database, including geolocated data and origin–destination. The fifth column shows the way data are integrated into the spatial domain representation, which is the street a graph in this case. The integration gives rise to feature vectors at each vertex of the graph (street corner) used in the predictive model. The sixth column classifies the data as dynamic or static, according to update frequency and temporal variability, and the seventh and the eighth columns outline the tasks associated with each layer and the corresponding potential user, respectively. The last three columns describe how the interaction with different devices for toggling is performed, which can be accomplished using a mouse, keyboard shortcuts, or the button box, with each layer represented by an icon.

4.1.2. Taxi trips

According to population density in census tracts, a synthetic ride-hailing dataset that emulates taxi data and comprises 87,000 records labeled as "regular" or "occurrence" was generated. The former represents trips with no criminal incidents (e.g., robbery or theft) and the latter corresponds to trips in which such incidents occurred. Normalized population density is used as probabilities for randomly sampling census track as origin and destination — distance to crime hotspots were used for labeling trips as "occurrence" or "regular" in an approximately 1:90 ratio when the distance from the origin or destination to the hotspot is smaller than 500 m. The procedure generated a highly imbalanced dataset. Each trip record includes information such as the closest street corner (graph node) to the origin and destination computed from the location of the census track, period of day–morning [6 AM–12 PM], afternoon [12 PM–6 PM], night [6 PM–12 AM], or dawn [12 AM–6 AM]–(more samples are drawn in the working hours), day of the week (fewer samples are drawn on weekends), and trip month (approximately the same number of samples in each month). Origin–destination data are represented as arcs connecting starting and ending points of the trips.

The predictive task focuses on that dataset for forecasting whether a trip would involve a criminal occurrence. A feature vector was, therefore, constructed for both origin and destination, concatenated them, and processed the combined vector through a binary response predictive model for determining the likelihood of crime. This database not only enriches the graph, but also leverages it for predictive purposes. Additionally, taxi pick-up and drop-off count can be recorded at each node, following the same nearest corner criterion.

4.1.3. Weather

Temperature and rainfall data were collected from three weather stations around the city in 2020, sourced from the Brazilian National Meteorology Institute. Those data are spatially sparse, but temporally rich. The exact locations of the three weather stations in São Paulo, located at the districts of Barueri, Interlagos, and Mirante, are pinpointed in the map visualization. The climate value at each graph node is obtained by linearly interpolating information from the three stations and using as weight coefficients the inverse of the distance between the location and the three stations.

4.1.4. Public transportation

Location and types of public transportation facilities (bus stops, terminals, subway, and train stations) were obtained from Geosampa portal (). The data are geolocated and distinct icon representations are utilized for each category on the map. The count of each type of transportation facility within a 200-meter radius centered on each node is used as data.

4.1.5. Favelas — Subnormal agglomerates

In the Brazilian urban context, *subnormal agglomerates*, commonly known as *favelas*, refer to densely populated informal settlements characterized by irregular land occupation, limited access to public services, and higher levels of social and economic vulnerability. Such areas often display distinct mobility patterns and are frequently associated with increased exposure to urban risks, which makes them particularly relevant for crime-related analyses.

Geolocation data for subnormal agglomerates were obtained from the IBGE repository (www.ibge.gov.br) for 2019 and are visually represent as polygons — GeoJSON data are used for those urban districts. A binary variable is assigned to each street graph node, indicating whether it is located within 500 m of a subnormal agglomerate.

4.1.6. Socioeconomic data

Socioeconomic data were collected from the Brazilian Census database (www.ibge.gov.br) for 2010 and are originally aggregated in census tracts represented by a polygonal curve. Seven socioeconomic indicators, namely, average household income, average householder income, unemployment rate among householders, literacy rate for children aged 7 to 15 years, percentage of residents under 18 years, percentage of residents aged 18 to 65 years, and percentage of residents over 65 years were considered. Towards aggregation of those data in the graph, each corner captures all data from the census sector to which it belongs — if located on the boundary between multiple sectors, it stores the average of all data.

4.1.7. Graph classification

Each node of the street graph is classified as *dead end*, *near dead end*, or *regular*. A node is a *dead end* if it is connected to a node of degree one, indicating a street with no exit and a dead-end node within a 100-meter path from a node is a *near dead end*. Otherwise, the node is classified as *regular*. Those classifications are represented as points in the visualization and each node preserves its value in the graph.

4.1.8. Identification of crime hotspot

Identifying crime hotspots is a significant task in crime mapping [38–40] and hotspots identification techniques detect locations with a high risk of crime. A Markov Model was employed in the present analysis for detecting hotspots [41,42], also represented as points in the map. The graph nodes preserve their value.

4.1.9. Prediction

Effective classification models depend on robust feature vectors. Since a critical challenge in our context is the highly imbalanced nature of the ride-hailing dataset a rigorous experiment was run for the selection of an appropriate classification model that accommodates that imbalance. Among the classification algorithms considered for imbalanced datasets, XGBoost [43] combined with Random Undersampling [44] yielded the best results, achieving a 0.89 G-mean performance.

Each aforementioned database corresponds to a data layer, from which a feature vector was generated for each node of the street graph. The predictive task relies on origin–destination trips, in which the feature vector for the origin is concatenated with that of the destination and input into the classification model for predicting whether a crime occurrence would take place during a taxi trip. The prediction is then compared with the labels in our trip database — correct predictions are counted as positive points and incorrect ones are marked as failures. Subsequently, a regular grid where the counts of successes and failures are aggregated with each grid cell is used as domain discretization. The result is visualized in a three-dimensional view, where the height corresponds to the count within each cell, with green color indicating success and red denoting failure. That layer fulfills requirement R5 and supports task T2 related to the analysis.

Fig. 4 illustrates the prediction grid, where the central area shows a higher concentration of correct predictions. Conversely, the peripheral regions display a greater number of failures. The height of the grid is particularly significant, for it correlates with the larger volume of trips occurring in the central area. In contrast, the peripheral areas, where the model tends to underperform, reflect a lower density of trips.

5. Visualization design

The proposed solution has been built upon Kepler.gl [37] open-source project described in Section 4, although the solutions can also be adaptable to other platforms. The implementation adheres to the requirements outlined in Section 3.1 and effectively addresses exploration and analysis tasks. In what follows is the implementation of the solutions proposed for tackling data layer-based entire screen visualization and adaptive lens filtering problems.

5.1. Layer toggling

Implemented to satisfy requirement R1, this feature accommodates diverse databases and deals with dynamic and static data, properly handling different data types and visual representations. Each database is treated as a layer, as illustrated in Fig. 1. The layers occupy the entire screen viewport and can overlap during visualization. A technique known as Layer Toggling has been developed for facilitating seamless transitions between layers while maintaining retinal persistence. It enables instantaneous layer changes without losing spatial focus, delegating the task to manual control rather than visual attention. Three versions were implemented for accommodating its operation in different devices, namely, mouse, keyboard, and button box, as detailed in what follows.

- Mouse Interaction: An icon is added to each layer button, positioned at the rightmost part adjacent to the map. The setup enables the layer to be displayed when user clicks with mouse on symbol and turned off when they click on .
- Keyboard Interaction: A key on the computer keyboard is assigned to each layer. Users have two options, namely, assign a number to each layer for an easier correspondence or assign customized keys. The latter is useful when users have keyboards with tactile reliefs for intuitive navigation, facilitating navigation of layers without diverting their gaze from the screen.
- Button Box: Whereas the aforementioned options are suitable for traditional computers, an additional device would significantly enhance navigation capabilities. Therefore, "Elgato Stream Deck Classic" [45], a live production controller that features 15 customizable LCD keys and an adjustable stand, was selected. Unlike other versions, it enables users to calibrate the tilt angle and is priced under 200 euros.
  The Stream Deck Classic is organized as a 5×3 key layout, as depicted on the top-left in Fig. 1. For ergonomic purposes, only the four horizontal buttons from each of the three columns, each assigned to index, middle, ring, and little fingers were used. Such a configuration enables fingers to navigate across the nine activation buttons, of which each corresponds to a layer. An icon representing the content alongside the layer name, which is visible from all tilt angles, was included for aiding memory retention.

All three aforementioned interactive resources aim to address requirement R3, although the use of keyboard without tactile reliefs diverges a bit from this goal.

5.2. Visibility-preserving lenses

The following two types of lenses for dynamic filtering, whose adjustment is based on an initial parameter that can be either spatial, or temporal have been developed and implemented towards addressing requirement R4.

- Spatial Lens: The first type is a brush that enables exploration of a specified number of points, such as one hundred crime incidents. When the cursor is moved over the map, the brush highlights only groups of one hundred points; in areas with higher point concentrations, its size decreases and increases in areas with fewer points.
  The implementation of the feature was inspired in [46], which combines Quad Tree data structure [47] with K-Nearest Neighbor (KNN) clustering algorithm [48]. Fig. 3 illustrates the integration of both methods.
  Data points are inserted into a Quad Tree, a process applicable only to layers with geolocated or origin–destination data types, for they have specific latitude and longitude coordinates on the map.

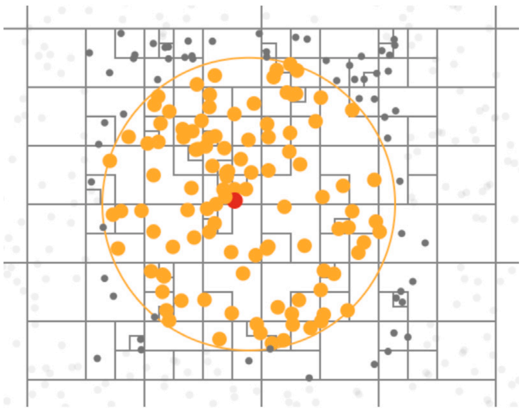

Fig. 3. Division of space using a Quad Tree, represented by gray quadrants. The application of KNN to find the "n" points closest to the cursor is highlighted in red. Explored points appear in gray, while found points are shown in orange.

A slider included in the interface specifies the number of points to be kept in the visualization and, combined with the cursor's position, serves as input for KNN algorithm, enabling searches for "k" closest points. The location of the last point found, e.g., the most distant from the cursor is then saved and the distance between both points is measured for the establishment of the brush radius.

- Temporal Lens: Temporal data, such as distribution of crimes by month, is depicted as a histogram within a folding window. The implemented lens enables users to select a range of bars to display, for instance, crimes in the months of February and March. Additionally, Kepler offers an animation feature that begins with the initial range input (e.g., two months) and progresses to displaying subsequent months until the entire distribution has been covered. Upon completion, the animation seamlessly loops back to the beginning.
Several improvements have also been incorporated into the tool. As an example, let us suppose we want to compare crime rates in February and March to other periods of the year. The dynamic lens then accepts inputs in terms of occurrence count or density. Consequently, the animation's progression is no longer fixed at two months, but dynamically adjusts to representing time intervals equivalent to the occurrences in February and March. The cumulative sum algorithm [49] facilitates the process. As the animation progresses, each subsequent bar is added until the initially set amount has been reached, with the first bar being subtracted. The approach prevents screen overload by consistently displaying the same quantity of data in each frame. Furthermore, the animation reverses direction upon reaching the end of the histogram, rather than looping back to the beginning so that continuity in the visualization is maintained.

Several other customized tools have been implemented for improving computational performance:

- Efficient Handling of Large Parquet Files: A database reader and manager has been integrated with the use of DuckDB library [50], significantly enhancing the speed of reading and querying, addressing challenges associated with big data.
- Correlation Matrix: A correlation matrix of the layers has been incorporated into the map legend Fig. 4 for simplifying the selection of layers, thus enabling expert analysts to seamlessly switch between the most correlated layers to elucidate hypothesis, thereby addressing requirement R5. The correlation matrix is constructed by first computing the pairwise correlations (with the use of Pearson's coefficient) of all 52 features that contribute to the predictor. The features are grouped into eight thematic layers. The average of the correlation values within each layer is computed towards synthesizing the matrix at the layer level. The resulting reduced matrix, shown in the legend, summarizes the average intra-layer correlations.
- Shapley Values: A bar chart displaying Shapley values [51] has been included in the legend for indicating characteristics that most significantly contribute to the classification model decison, which highlights important layers that should be prioritized in the analysis. Only absolute values are considered, since the direction of contribution (positive or negative) is less relevant than the percentage of contribution or importance. The tool also contributes to the fulfillment of requirement R5.

6. Illustrative case study

This section presents a scenario in which a user utilizes the tool to fulfill both exploration and analysis tasks. Initially, the user engages with the tool as a regular user for exploratory purposes, followed by an analytical approach as a data analyst. Such an illustrative case study demonstrates how the proposed methodology can support urban data exploration and provide valuable insights for users.

Let us consider a scenario where an individual is searching for an apartment to rent in São Paulo, specifically around Paulista Avenue. He is interested in assessing the neighborhood's security, transportation facilities, and overall cost of living. He uses the proposed approach to *explore* the city, first inserting an address in the neighborhood of interest, resulting in an icon placed at that location.

He activates the "Crime" layer to evaluate the security aspect and observes a notable concentration of crimes in the area. Knowing that, he decides to go deeper and analyze if crimes are, in fact, concentrated near the location chosen. He then uses the spatial lenses, hovering the mouse along the streets and avenues. The brush contracts particularly nearby his specified address, indicating a greater concentration of crime, thus providing valuable insights into both distribution and density of crime occurrences along the street. He wants to go even further towards knowing the most problematic hours so as to avoid them. He uses the temporal filter and visualizes the distribution of crimes by hours, with a bigger concentration at night. He realizes the number of crimes at 8 pm is equivalent to all dawn [12am-6am] occurrences, thanks to the temporal dynamic lens. Now he knows the areas and the times to be avoided to be safer.

Regarding mobility, the user investigates public transportation facilities near his address, swapping instantaneously to the "Public transportation" layer — he does not lose the view of the particular address, for he has the auxiliary keyboard in his left hand with fingers placed in the layers of interest. He sees many nearby bus stops and three subway stations. Although the place is well served in terms of public transportation, what would happen if he decided to take a taxi? Is it safe in that zone? He then switches to the "Taxi trips" layer, where he sees all the taxi trips with some criminal occurrences highlighted. He raises the hypothesis the economic level of people living in that zone might be attracting criminals to the area.

He decides to confirm that hypothesis and toggles to the "Socioeconomic" layer. The darker green color of the neighborhood reveals a high income rate, corroborating his suspicions. Despite the higher rents and expensive cost of life in that region, he also noticed a high rate of literate children, which might ensure a good community for his family.

Since he is also an expert *analyst*, he decides to go further in explaining crime in taxi trips. He then seeks to identify additional factors contributing to crime and analyzes whether a standard prediction algorithm would face challenges in that area.

The crime prediction literature reports several hypothetical factors that might lead to a good crime prediction — many criminological

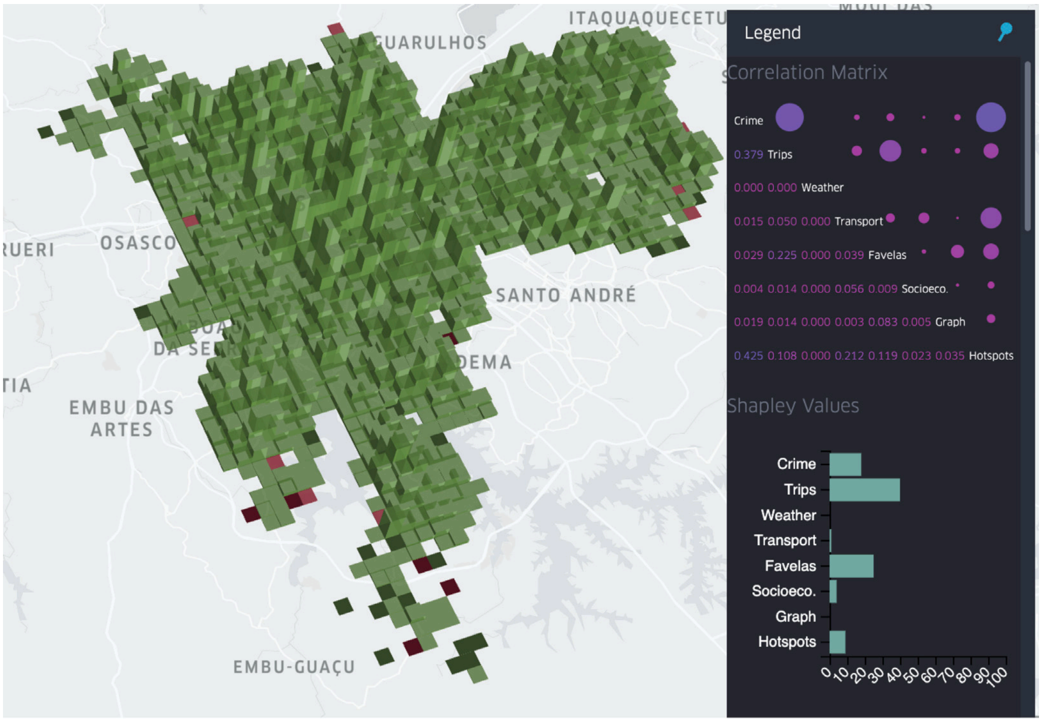

Fig. 4. Resources for Analytical Tasks: The figure illustrates the spatial representation of the classification model's behavior, addressing the initial question of where the model underperforms, specifically in the peripheral areas. Other resources for analysts, including the correlation matrix and Shapley values are presented on the right side.

studies suggest a strong correlation of climate and crime [52], of which our user is aware. He moves to the "Weather" layer, which contains maximum and minimum temperature, and total precipitation. Using the filters, the user verifies the highest incidence of crime occurs in the months of February and March, which also coincides with the high temperature in São Paulo during those months.

Several studies have performed crime prediction with a basis on the detection of hotspots [53], and he learned about that. He then switches to the "Hotspots" layer and sees hotspots are present in most corners, which is the granularity of the data, but the difference is the number of occurrences accumulated in the hotspots and using the filter the user is able to realize that right in the indicated location there are 4 hotspots, which is not favorable for their security.

Towards better understanding the data modeling, the user explores the "Graph" layer and verifies each corner is associated with a feature vector holding information from the underlying layers. He notices the graph itself provides meaningful information, including presence of dead ends, which can be an indicator of insecurity.

São Paulo encompasses some very densely populated settlements or shantytowns, the well-known "Favelas", whose residents are usually low-income individuals who face economic challenges. Such as communities can be associated with issues that include poverty, crime, and inadequate access to essential services. The user observes the favelas are mostly on the outskirts of the city and no favelas are near his location. However, he also finds a correlation between "Favelas" and "Trips" layers, since a large number of trips with criminal occurrence depart from favelas, or have them as their destination.

Finally, the user explores the "Prediction" layer and notices the predictive model fails more frequently in the peripheries, especially in the southwestern part, an area that also concentrates a large number of taxi trips with criminal occurrences, and unlike the rest of the city, the number of trips with and without occurrences is balanced. He concludes the predictive model that generally works well for unbalanced data fails in that balanced region.

7. User study

The user study was designed to evaluate the effectiveness of layer toggling and dynamic filtering features through a rigorous experimental framework. 23 participants were recruited for the study and completed both phases of the experiment.

They were predominantly aged between 18–34 years (15 participants, approximately 68%) and 35–54 years (7 participants, approximately 32%), with a gender distribution of 5 women, 17 men, and 1 who preferred not to disclose. Most of them (22 individuals) were postgraduate students, primarily pursuing PhDs in fields of Computer Science, Mathematics, Statistics, and Engineering, indicating a robust foundation in quantitative analysis and research methodologies. Three participants work professionally in the area of crime prediction and one is involved in public transportation policy, contributing with domain-specific expertise to the application context. Additionally, a significant number of the participants were international students with limited prior knowledge of the city of São Paulo, which aligns with the study scenario of new residents seeking information to support housing-related decisions.

Four participants identified themselves as developers, which enhances their technical engagement with data visualization tools, and one served as an assistant professor, offering an academic perspective. Three had prior experience in urban data visualization, providing valuable insights and skills relevant to the study, whereas the remaining ones lacked experience in the area, potentially influencing their interactions with the visualization tools used in the research.

An instruction phase was conducted, during which a demonstration of the tool was provided, including an analysis of a neighborhood, adhering to specific criteria that mirrored the tasks participants would later undertake. They were then tasked with completing similar activities while taking notes on a provided form and, at the end, they filled out a questionnaire for the capture of their feedback and experiences.

7.1. Study 1: Layer toggling

Objectives:

- To assess whether enabling and disabling layers reduced visual clutter when compared to displaying all layers concurrently, and
- To evaluate whether the Button Box device promoted a more efficient layer switching.

In the first study, participants were randomly assigned to one out of four groups for evaluating the functionality of toggling between layers, i.e., activating or deactivating specific data layers during visual exploration. The groups were organized as follows: Group A — with no toggling functionality and button box; Group B — use of the mouse to toggle; Group C — use of the keyboard to toggle; and Group D — use of the Button Box.

The primary goal was to determine whether the ability to enable and disable layers significantly improved the user experience. Participants were tasked with gathering specific data from active layers related to transportation, socioeconomic information, and favelas in the vicinity of Avenida Paulista, in São Paulo. They responded to quantitative questions on number of subway stations, average household income, and presence of favelas within a three-block radius. Their confidence in the answers was measured on a Likert scale from 1 to 5, where 1 indicated low confidence and 5 indicated high confidence. It was anticipated that Group A would require more time to complete the tasks and showed signs of discomfort, whereas Groups B, C, and D would perform more efficiently — Group D was expected to achieve the best performance.

7.1.1. Study 2: Dynamic filtering

Objectives:

- To determine whether dynamic filters (spatial and temporal) were more efficient than static filters in performing specific tasks.

In the second study, all 23 participants engaged in a follow-up experiment aimed at evaluating the effectiveness of dynamic filters in comparison to static ones. The objective was to ascertain whether dynamic filters (spatial and temporal) offered greater efficiency than static filters of the original application. Participants were divided into two groups, namely, Group X, which utilized static filters, and Group Y, which employed the proposed dynamic filters, specifically the Dynamic Brush tool and the Temporal Data Dynamic Histogram.

Participants utilized the spatial brush filter, which adjusts its size according to number of crimes, for identifying the block with highest concentration of crimes. A smaller brush indicates a denser region of crime, faciliating the identification of hotspots. Participants should analyze crime data along Avenida Paulista, identifying the block with highest concentration of crimes and the peak time for those incidents. They were again asked to rate their confidence in their answers on a Likert scale, providing qualitative feedback on their experience.

After the experiments, participants completed a survey for assessment of their acceptance of the tool. The questionnaire addressed various aspects, including ergonomics, novelty, learning curve, and overall user experience. Additionally, participants who used a device were queried about their cognitive memory of layer associations and any distractions caused by the auxiliary device.

Table 2 shows the distribution of the groups and details of the two user studies, including the forms, can be found in the supplementary material.

7.2. Availability of data and materials

Datasets, modified software, experimental materials, including design documents, videos, forms, and results from the current study are available at https://github.com/Kareliavs/Toggler/.

Table 2
User Study Group Distribution: In the first study, participants were assigned to Group A (no device), Group B (mouse), Group C (keyboard), and Group D (Button Box). In the second study, Groups X and Y used static and dynamic filters, respectively.

Participants	Group	Functionality	Device
23	A (5)	With no functionality	None
	B (5)	Toggle between layers	Mouse
	C (5)	Toggle between layers	Keyboard
	D (8)	Toggle between layers	Button Box
23	X (10)	Static filters	–
	Y (13)	Dynamic filters	–

Table 3
Kruskal–Wallis (K-W) statistics, p-values and effect sizes for differences in the following variables among groups A, B, C and D: number of correct answers in the questionnaire (Hits), confidence level in responses (Confidence), time in seconds (Times), and the number of clicks (Clicks) required to complete the task.

Variable	K-W statistics	P-value	Effect size
Hits	2.0099	0.5703	0.0000
Confidence	2.0594	0.5602	0.0000
Times	11.921	0.0077	0.4880
Clicks	1.9356	0.5859	0.0000

8. Results

In Study 1, Kruskal–Wallis test assessed the differences in observations for each of the following variables among groups A, B, C, and D: number of correct answers in the questionnaire (hits), confidence level in responses (confidence), time in seconds (times) and the number of clicks (clicks) required to complete the task. Three outliers were identified for time measurements, according to the boxplot criterion — two from group C and one from group D. They were removed, except for one in group C, which, although classified as an outlier, was not considered an extreme value. Regarding number of clicks, one extreme outlier from group D was also removed. Table 3 shows the statistics, p-values and effect sizes from the tests performed. The effect sizes were taken according to the eta squared measure and negative results were reported as zero.

The results indicate no significant differences (at 5% significance level) in number of correct answers, confidence levels, or number of clicks among the groups, but a significant one in task completion time. The average times (in seconds) for each group were 107 (A), 31.4 (B), 38 (C) and 24.6 (D) and a bootstrap study on the differences in means between groups was conducted for analyses of multiple comparisons. Bootstrap sample means were compared for each pair of groups and the mean of those differences was estimated, along with the 2.5% and 97.5% percentiles as confidence intervals for this parameter. Fig. 5 shows the confidence intervals for the difference in means between groups.

It has been concluded that task completion times are significantly longer for the group that did not use any device (“Button Box - No”, “Mouse - No”, and “No - Keyboard” comparisons in Fig. 5). For the remaining groups, completion times are considered statistically similar, although their confidence interval limits are close to zero, making the result borderline (“Button Box - Keyboard”, “Mouse - Button Box”, and “Mouse - Keyboard” comparisons in Fig. 5). Since no difference in completion time across devices, it can be concluded that the Button Box did not improve the user experience in this regard.

In study 2, Kruskal–Wallis test also assessed differences in observations for each of the previously cited variables between groups X and Y. Outliers and extreme values were observed only for number of clicks in group Y and were removed. Table 4 provides the statistics and p-values from the tests.

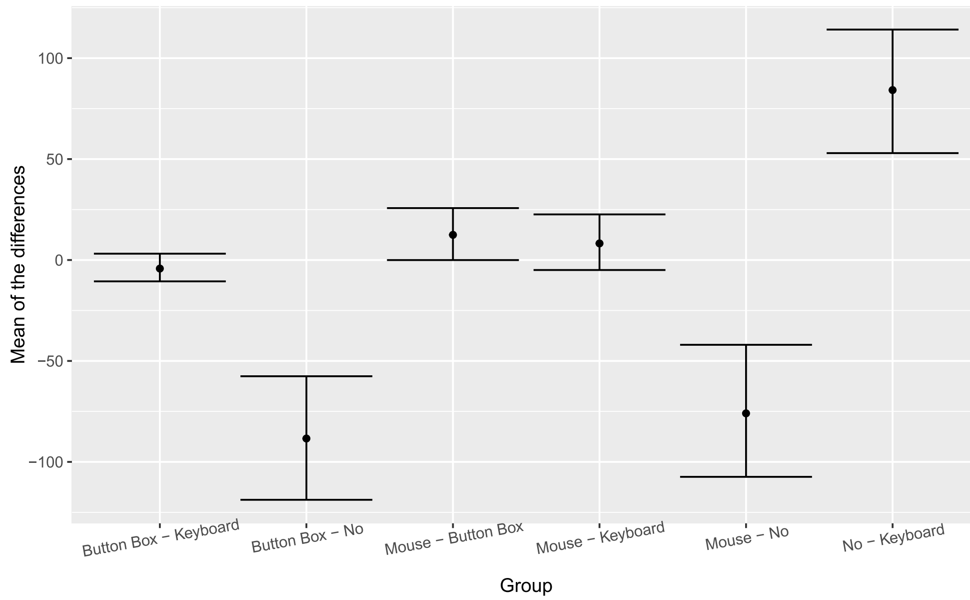

Fig. 5. 95% Bootstrap Confidence interval for the difference in means considering all comparisons between groups (e.g., the "Button Box - Keyboard" comparison represents difference calculated as the mean of the "Button Box" group minus the mean of the "Keyboard" group).

Table 4
Kruskal–Wallis (K-W) statistics, p-values and effect sizes for differences in the following variables between groups X and Y: number of correct answers in the questionnaire (Hits), confidence level in responses (Confidence), time in seconds (Times), and number of clicks (Clicks) required to complete the task.

Variable	K-W statistics	P-value	Effect size
Hits	0.0491	0.8247	0.0000
Confidence	3.6151	0.0573	0.1250
Times	9.4287	0.0021	0.4010
Clicks	1.9003	0.1680	0.0474

A borderline difference in confidence levels was observed, with group Y (Dynamic filters) showing higher confidence in their answers, as illustrated in Fig. 6 and in the 95% bootstrap confidence interval for the differences in means, shown in Fig. 7. Although it includes zero, the upper limit is close to it. No significant difference was found in number of correct answers or clicks between groups X and Y. However, task completion time was significantly different between the groups at a 5% significance level, with group Y having the shortest times (in seconds), i.e., 131 (X) and 65.7 (Y). The 95% bootstrap confidence interval for the difference between groups are shown in Fig. 8. This confirms that dynamic filters effectively reduce the time required to complete a task.

Regarding additional questionnaire evaluations, the relationship between number of cases in which a user experience visual information overload and their assigned filter-related group was analyzed. At a 5% significance level, the hypothesis that number of overloads does not differ between groups X and Y was not rejected, leading to the conclusion dynamic filters do not increase or decrease the cases of overcharge compared to static filters. The study also assessed the relationship between distraction level and ease of use of filters among groups that received different devices. These variables were not found to be significantly different between groups (see Table 5 for the results).

According to the statistical tests, task completion time was the only variable for which the difference between groups was significant, indicating although the use (or lack thereof) of devices (e.g., mouse, keyboard, and button box) does not affect users' accuracy in their responses, their confidence levels, or number of clicks required to complete the task, the use of these devices is associated with a shorter task completion time. However, no difference was observed in completion times among participants who used different devices.

Table 5
Kruskal–Wallis test statistics (K-W Stats), p-values and effect size for differences in the following variables between groups: overload, distraction level, and easy of use.

Variable	Groups	K-W Stats	P-value	Effect size
Overload	X and Y	0.7591	0.3836	0.0000
Distraction level	A, B, C and D	3.5333	0.1709	0.1020
Easy of use	A, B, C and D	1.0081	0.6041	0.0000

Additionally, a statistical difference was detected in task completion times between participants who were provided with dynamic and static filters — the group that used dynamic filters completed the tasks in less time.

9. Discussion, implications, and limitations

This paper has introduced a novel approach to urban data visualization that addresses the challenges of analyzing multivariate and multimodal datasets, particularly in the context of predictive modeling for crime and mobility in São Paulo, and its key contributions include:

1. A Layered Visualization Methodology: A method proposed organizes urban data into distinct layers so that each layer to occupies the entire screen space. The approach enhances spatial resolution and facilitates the perception of both local and global correlations among features.
2. Dynamic Layer Toggling and Filtering Techniques: The innovative layer toggling mechanism, combined with dynamic filtering capabilities, enables users to navigate complex datasets efficiently. Allowing users to maintain focus on the screen while comparing features enhances the overall user experience.
3. A Comprehensive Evaluation of Predictive Systems: A framework introduced for visual evaluations of predictive models enables the identification of areas with successful predictions as well as those with poor performance, providing valuable insights for further refinements of predictive algorithms.
4. The statistical methods employed showed the proposed tool for the use of dynamic filters to improve data visualization enhanced user's experience in terms of time required to complete a task within the platform and the use of the Button Box device does

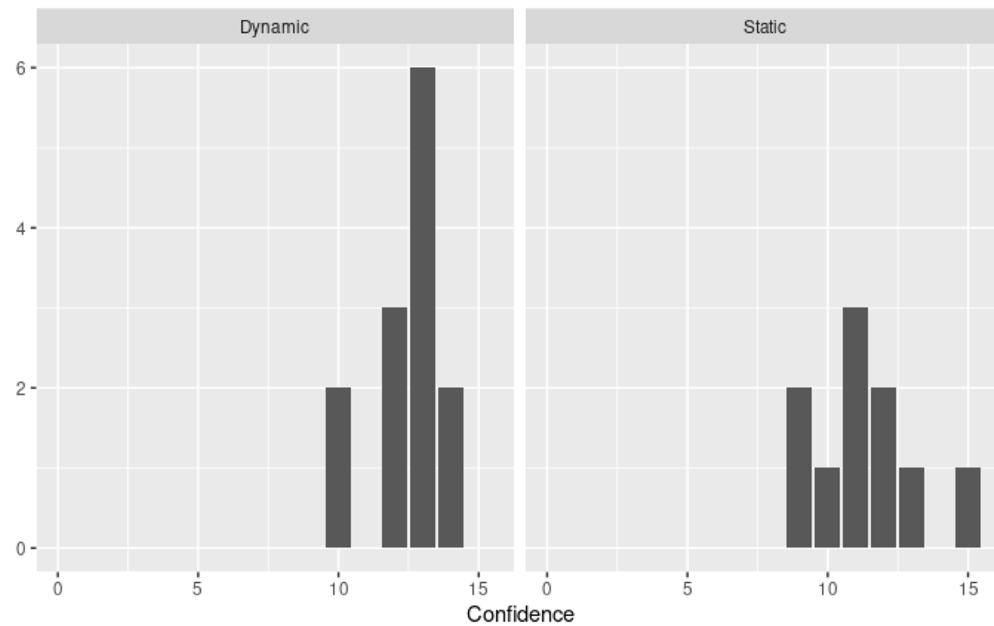

Fig. 6. Bar plot of the confidence level demonstrated by the participants in answering questions about the filters.

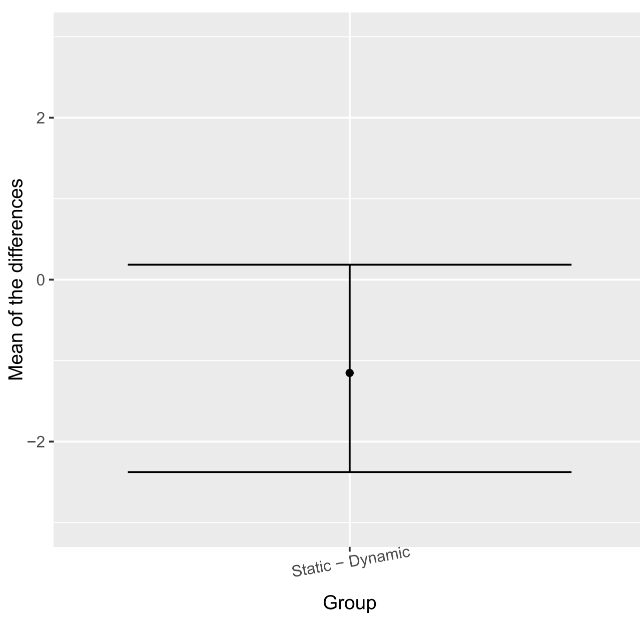

Fig. 7. 95% Bootstrap confidence interval for the difference in means (variable Confidence) considering the comparison between groups (X and Y). The "Static - Dynamic" comparison represents the difference calculated as the mean of the "Static" group minus the mean of the "Dynamic" group.

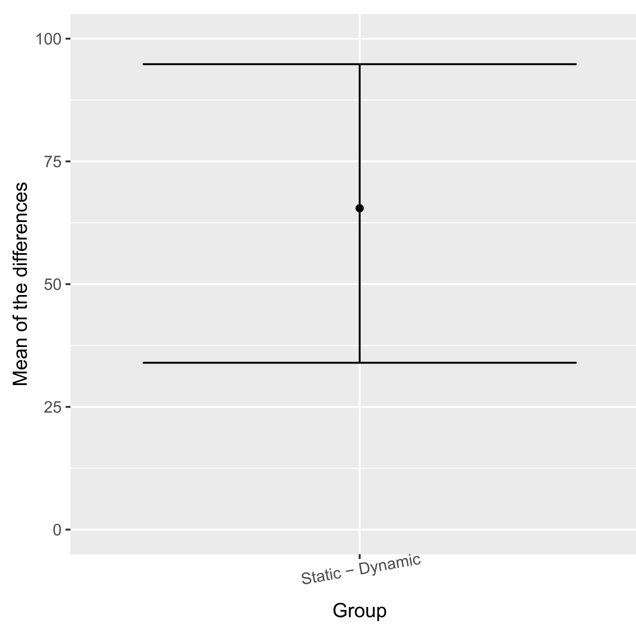

Fig. 8. 95% Bootstrap confidence interval for the difference in means (variable Times) considering the comparison between groups (X and Y). The "Static - Dynamic" comparison represents the difference calculated as the mean of the "Static" group minus the mean of the "Dynamic" group.

not reduce task completion time compared to other devices, such as mouse and keyboard. However, the completion times remain equivalent and significantly better than the absence of a device. Although dynamic filters reduced task execution time, the results indicate a borderline effect on increasing users' confidence in their responses and no evidence that the change improved response accuracy or reduced the number of clicks required to complete the task was found. The same was observed for the use of devices, since the inclusion of the Button Box did not appear to enhance confidence and accuracy, or reduce the number of clicks. Although the button box may not be better than other devices regarding times of task executions, the user's experience was impacted in terms of time required to complete a task with implementation of dynamic filters and the use of any device – whether a keyboard, mouse, or button box – leading to faster execution.

Despite the promising results of the proposed approach, several limitations should be considered in the interpretation of the findings. The user study involved a relatively small sample size and was conducted in a controlled experimental setting, which may restrict the generalizability of the results to broader populations or real-world decision-making contexts. In addition, the evaluation focused on a single metropolitan area and outcomes may differ when applied to cities with distinct urban structures or levels of data availability.

From an applied standpoint, the results suggest the proposed visualization framework can support analysts, policymakers, and citizens in exploring complex predictive data. Nevertheless, its adoption in operational environments would require integration with existing decision-support workflows and adaptation to domain-specific requirements.

There are several potential avenues for future research. Exploring the integration of real-time data streams might enhance the responsiveness of the proposed visualization system, enabling more dynamic and timely analyses, and expanding the methodology towards incorporating additional data types (e.g., social media interactions or sensor data) might provide a more comprehensive view of urban phenomena.

investigations on user interactions with the proposed visualization tools through user studies can yield insights into usability and effectiveness, leading to further refinements in design and the application of the approach to other urban contexts would validate its adaptability and effectiveness across different environments and datasets. An important direction for future work involves extending the evaluation to additional cities equipped with richer and more diverse multi-attribute datasets. Such multi-city analyses would enable a deeper examination of generalizability and comparative behavior across distinct urban morphologies. As more comprehensive open datasets become available, broader studies can be conducted for assessments of scalability and transferability of the proposed techniques.

In summary, this study can contribute to the field of urban data visualization by providing a robust framework for analyses of complex datasets, while also highlighting opportunities for future exploration and innovation.

10. Conclusions

This manuscript presented a novel layered visualization framework for the exploration and evaluation of multivariate and multimodal urban data, with a focus on predictive modeling for crime and mobility. By combining full-screen layered views with dynamic filtering and interaction mechanisms, the proposed approach enhances spatial perception and supports more efficient exploratory analyses.

Overall, the findings showed dynamic visualization techniques can significantly improve task efficiency, while also revealing important constraints related to user confidence, accuracy, and device effectiveness. The results contribute to the field of urban data visualization by offering both methodological advances and empirical evidence, while laying the groundwork for future studies aimed at improving generalizability, scalability, and real-world applicability.

CRediT authorship contribution statement

Karelia Salinas: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Luis Gustavo Nonato:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Jean-Daniel Fekete:** Writing – review & editing, Supervision, Methodology. **Fernanda Bartolo dos Santos Saran:** Writing – original draft, Validation, Formal analysis.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Karelia Salinas reports financial support was provided by State of Sao Paulo Research Foundation. Luis Gustavo Nonato reports financial support was provided by State of Sao Paulo Research Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was financially supported by FAPESP (#2020/07012-8, #2022/09091-8). The opinions, hypotheses, conclusions, and recommendations expressed in this material are the responsibility of the authors and do not necessarily reflect the views of FAPESP.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.is.2026.102712.

Data availability

Public datasets were used. The integrated kepler.gl JSON file generated for this study is available on request.

References

[1] G. Garcıa, J. Silveira, J. Poco, A. Paiva, M.B. Nery, C.T. Silva, S. Adorno, L.G. Nonato, CrimAnalyzer: Understanding crime patterns in São Paulo, IEEE Trans. Vis. Comput. Graphics 27 (4) (2019) 2313–2328.

[2] G. Garcia-Zanabria, M.M. Raimundo, J. Poco, M.B. Nery, C.T. Silva, S.F.A. de Abreu, L.G. Nonato, CriPAV: Street-level crime patterns analysis and visualization, IEEE Trans. Vis. Comput. Graphics (2021).

[3] N. Ferreira, J. Poco, H.T. Vo, J. Freire, C.T. Silva, Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips, IEEE Trans. Vis. Comput. Graphics 19 (12) (2013) 2149–2158.

[4] J. Freire, C.T. Silva, H.T. Vo, H. Doraiswamy, N. Ferreira, J. Poco, Riding from urban data to insight using new york city taxis, IEEE Data Eng. Bull. 37 (4) (2014) 43–55.

[5] A. Bogomolov, B. Lepri, J. Staiano, N. Oliver, F. Pianesi, A. Pentland, Once upon a crime: towards crime prediction from demographics and mobile data, in: Proceedings of the 16th International Conference on Multimodal Interaction, 2014, pp. 427–434.

[6] S.F. Tekin, S.S. Kozat, Crime prediction with graph neural networks and multivariate normal distributions, 2021, arXiv preprint arXiv:2111.14733.

[7] K. Salinas, T. Gonçalves, V. Barella, T. Vieira, L.G. Nonato, Cityhub: A library for urban data integration, SIBGRAPI, in: 2022 35th SIBGRAPI Conference on Graphics, Patterns and Images, vol. 1, IEEE, 2022, pp. 43–48.

[8] S. He, From beautiful maps to actionable insights: Introducing Kepler. gl, Uber's open source geospatial toolbox, Uber Eng. 29 (2018).

[9] I. Boyandin, Visualization of Temporal Origin-Destination Data (Ph.D. thesis), Faculty of Science of the University of Fribourg, 2013, URL https://ilya.boyandin.me/assets/thesis.pdf.

[10] Y. Gu, M.-J. Kraak, Y. Engelhardt, F.-B. Mocnik, A classification scheme for static origin–destination data visualizations, Int. J. Geogr. Inf. Sci. (2023) 1–28.

[11] B.D. Dent, Cartography: Thematic map design, (No Title) (1999).

[12] J. Bertin, Semiology of Graphics, University of Wisconsin Press, 1983.

[13] M. Wattenberg, Arc diagrams: Visualizing structure in strings, in: IEEE Symposium on Information Visualization, 2002, INFOVIS 2002, IEEE, 2002, pp. 110–116.

[14] M. Krzywinski, J. Schein, I. Birol, J. Connors, R. Gascoyne, D. Horsman, S.J. Jones, M.A. Marra, Circos: an information aesthetic for comparative genomics, Genome Res. 19 (9) (2009) 1639–1645.

[15] M. Krzywinski, I. Birol, S.J. Jones, M.A. Marra, Hive plots—rational approach to visualizing networks, Brief. Bioinform. 13 (5) (2012) 627–644.

[16] D. Guo, X. Zhu, Origin-destination flow data smoothing and mapping, IEEE Trans. Vis. Comput. Graphics 20 (12) (2014) 2043–2052.

[17] D. Guo, Flow mapping and multivariate visualization of large spatial interaction data, IEEE Trans. Vis. Comput. Graphics 15 (6) (2009) 1041–1048.

[18] T. Von Landesberger, F. Brodkorb, P. Roskosch, N. Andrienko, G. Andrienko, A. Kerren, MobilityGraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering, IEEE Trans. Vis. Comput. Graphics 22 (1) (2015) 11–20.

[19] N. Adrienko, G. Adrienko, Spatial generalization and aggregation of massive movement data, IEEE Trans. Vis. Comput. Graphics 17 (2) (2010) 205–219.

[20] K. Buchin, B. Speckmann, K. Verbeek, Flow map layout via spiral trees, IEEE Trans. Vis. Comput. Graphics 17 (12) (2011) 2536–2544.

[21] W. Cui, H. Zhou, H. Qu, P.C. Wong, X. Li, Geometry-based edge clustering for graph visualization, IEEE Trans. Vis. Comput. Graphics 14 (6) (2008) 1277–1284.

[22] M. Dakowicz, C. Gold, Interactive tin modification with a cutting tool, in: Proc. of 4th ISPRS Workshop on Dynamic and Multi-Dimensional GIS, Pontypridd, Wales, Citeseer, 2005, pp. 5–9.

[23] K. Salinas, J.-D. Fekete, L.G. Nonato, Navigating multi-attribute spatial data through layer toggling and visibility-preserving lenses, in: Posters of VIS 2024-IEEE Visualization, 2024.
[24] M.-J. Kraak, F. Ormeling, Cartography: Visualization of spatial data, Geogr. Inf. Sci. 25 (1) (2010) 1–10.
[25] R. Gao, H. Tao, H. Chen, W. Wang, J. Zhang, Multi-view display coordinated visualization design for crime solving analysis: Vast challenge 2014: Honorable mention for effective use of coordinated visualizations, in: 2014 IEEE Conference on Visual Analytics Science and Technology, VAST, IEEE, 2014, pp. 321–322.
[26] Y. Wang, T. Wang, Y.Y. Zhang, H. Zhang, UrbanDataLayer: A unified data pipeline for urban science, in: Proceedings of the Neural Information Processing Systems, NeurIPS, 2024, URL https://proceedings.neurips.cc.
[27] N. Sideris, G. Bardis, A. Voulodimos, G. Miaoulis, Enhancing urban data analysis: Leveraging graph-based convolutional neural networks for a visual semantic decision support system, Sensors 24 (3) (2024) 567–580, URL https://www.mdpi.com.
[28] F. Bentlin, The urban layer analysis, in: Handbook of Qualitative and Visual Methods in Spatial Research, Springer, 2024, URL https://library.oapen.org.
[29] X. Zhou, W. Hsu, Layered visualization for enhanced user engagement, Inf. Vis. 17 (3) (2018) 189–204.
[30] F. Miranda, T. Ortner, G. Moreira, The state of the art in visual analytics for 3D urban data, Comput. Graph. Forum 43 (2) (2024) 123–145, URL https://www.wiley.com.
[31] J. Wagner, C.T. Silva, W. Stuerzlinger, Reimagining TaxiVis through an immersive space-time cube metaphor and reflecting on potential benefits of immersive analytics for urban data exploration, in: Proceedings of the 2024 IEEE Conference on Immersive Analytics, 2024, URL https://ieeexplore.ieee.org.
[32] M. Dodge, R. Kitchin, Outlines of geographical information science, Geogr. Rev. 103 (2) (2013) 145–162.
[33] C. Tominski, S. Gladisch, U. Kister, R. Dachselt, H. Schumann, A survey on interactive lenses in visualization, in: EuroVis (STARs), 2014, pp. 1–20.
[34] B. Bach, S. Carpendale, The role of lenses in data exploration, IEEE Trans. Vis. Comput. Graphics 22 (1) (2016) 1–10.
[35] D. Keim, M. Ward, Visual data mining: Techniques and applications, Comput. Graph. 33 (4) (2009) 431–438.
[36] F. Fischer, M. Hegarty, Temporal lenses in data visualization, J. Vis. 18 (2) (2015) 1–12.
[37] Uber Technologies, Inc., Kepler.gl, 2024, https://kepler.gl/. (Accessed 23 February 2024).
[38] S. Chainey, J. Ratcliffe, GIS and Crime Mapping, John Wiley & Sons, 2013.
[39] J. Eck, S. Chainey, J. Cameron, R. Wilson, Mapping crime: Understanding hotspots, 2005.
[40] E. Eftelioglu, S. Shekhar, X. Tang, Crime hotspot detection: A computational perspective, in: Improving the Safety and Efficiency of Emergency Services: Emerging Tools and Technologies for First Responders, IGI Global, 2020, pp. 209–238.
[41] D.L. Robertson, W.S. Goodridge, Predicting density of serious crime incidents using a multiple-input hidden Markov maximization a posteriori model, Mach. Learn. Appl. 7 (2022) 100231.
[42] S. Mondal, D. Singh, R. Kumar, Crime hotspot detection using statistical and geospatial methods: a case study of Pune City, Maharashtra, India, GeoJournal 87 (6) (2022) 5287–5303.
[43] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.
[44] H. He, E.A. Garcia, Learning from imbalanced data, IEEE Trans. Knowl. Data Eng. 21 (9) (2009) 1263–1284.
[45] Elgato Systems, Welcome to stream deck, 2024, https://www.elgato.com/uk/en/s/welcome-to-stream-deck. (Accessed 23 February 2024).
[46] Bea, K-nearest-neighbor search using D3 quadtrees, 2024, https://observablehq.com/@llb4ll/k-nearest-neighbor-search-using-d3-quadtrees. (Accessed 23 February 2024).
[47] R.A. Finkel, J.L. Bentley, Quad trees a data structure for retrieval on composite keys, Acta Inform. 4 (1974) 1–9.
[48] T. Cover, P. Hart, Nearest neighbor pattern classification, IEEE Trans. Inform. Theory 13 (1) (1967) 21–27.
[49] D.E. Knuth, The Art of Computer Programming, vol. 3, Pearson Education, 1997.
[50] DuckDB Foundation, DuckDB is a fast in-process analytical database, 2024, https://duckdb.org/. (Accessed 23 February 2024).
[51] A.E. Roth, Introduction to the Shapley value, Shapley Value 1 (1988).
[52] C.A. Anderson, K.B. Anderson, N. Dorr, K.M. DeNeve, M. Flanagan, Temperature and aggression, in: Advances in Experimental Social Psychology, vol. 32, Elsevier, 2000, pp. 63–133.
[53] D. Yang, T. Heaney, A. Tonon, L. Wang, P. Cudré-Mauroux, CrimeTelescope: crime hotspot prediction based on urban and social media data fusion, World Wide Web 21 (2018) 1323–1347.