




StreetWeave: A Declarative Grammar for Street-Overlaid Visualization of Multivariate Data

Sanjana Srabanti , G. Elisabeta Marai , and Fabio Miranda 

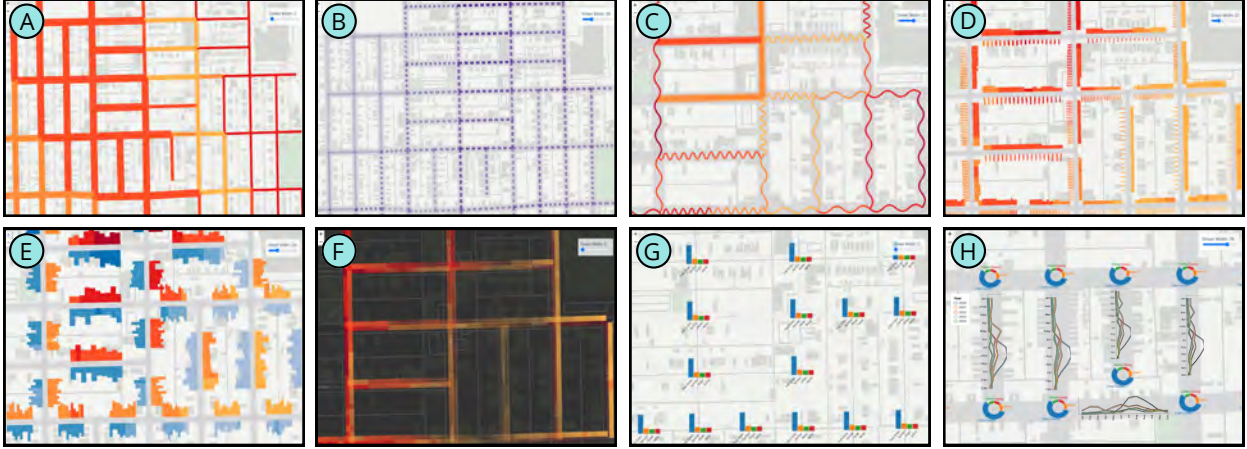


Fig. 1: Street and pedestrian network visualizations with StreetWeave. (A) A multivariate line map, with attributes like color and line width. (B) & (C) A pattern-based encoding approach for street overlays, utilizing dashed lines, variable opacities, and wave patterns to represent data attributes. (D) & (E) A bristle map and a line map that communicate data through variations in width, height, color, and opacity, with an additional alignment dimension integrating multiple encodings. (F) A multi-dimensional grid, where each cell encodes a distinct data value, and the background map color can be customized. (G) & (H) Charts applied directly to street intersections and segments, with an orientation dimension that allows for adjustments of plot orientation.

Abstract—The visualization and analysis of street and pedestrian networks are important to various domain experts, including urban planners, climate researchers, and health experts. This has led to the development of new techniques for street and pedestrian network visualization, expanding possibilities for effective data presentation and interpretation. Despite their increasing adoption, there is no established design framework to guide the creation of these visualizations while addressing the diverse requirements of various domains. When exploring a feature of interest, domain experts often need to transform, integrate, and visualize a combination of thematic data (e.g., demographic, socioeconomic, pollution) and physical data (e.g., zip codes, street networks), often spanning multiple spatial and temporal scales. This not only complicates the process of visual data exploration and system implementation for developers but also creates significant entry barriers for experts who lack a background in programming. With this in mind, in this paper, we reviewed 45 studies utilizing street-overlaid visualizations to understand how they are applied in practice. Through qualitative coding of these visualizations, we analyzed three key aspects of street and pedestrian network visualization usage: their analytical purposes, the visualization approaches employed, and the data sources used in their creation. Building on this design space, we introduce StreetWeave, a declarative grammar for designing custom visualizations of multivariate spatial network data across multiple resolutions. We demonstrate how StreetWeave can be used to create various street-overlaid visualizations, enabling effective exploration and analysis of spatial data. StreetWeave is available at urbank.org/streetweave.

Index Terms—Urban visual analytics, design space, street-overlaid visualization, visualization grammar.

1 INTRODUCTION

Street-overlaid visualization has a rich history that dates back to 1854, when John Snow created one of the most iconic examples: a map of the cholera outbreak in London. By overlaying cholera cases on a street map of London, he could communicate the source of the outbreak [31]. Since then, street-overlaid visualizations have become popular across various domains, including climate and environmental sciences [14, 17, 21, 32, 35, 54], urban planning [42, 53], and urban accessibility studies [18, 45], as they provide in-depth insights into

multiple factors. At the same time, visualizing street and pedestrian networks is challenging due to their complex structures, vast datasets, and diverse analytical requirements. Meaningful visualization outcomes often depend on interdisciplinary collaborations involving experts from computer science, urban planning, transportation engineering, architecture, geography, and public health [26]. These collaborations typically produce specialized workflows for analyzing and visualizing data spanning various spatial and temporal scales. While visualizations of street and pedestrian networks have supported diverse analytical tasks across domains, including identifying pollution hotspots [42] or highlighting accessibility concerns [42, 53], the research community remains fragmented [55]. Many projects arise from collaborations between computer scientists and domain experts focused on specific problems, resulting in one-off solutions that are difficult to adapt or reuse across contexts or regions. Creating them demands skills in visualization, graphics programming, and data integration, raising entry barriers for those without technical expertise. While recent progress in visualization toolkits and grammars has lowered some technical barriers [28],

• Sanjana Srabanti, G. Elisabeta Marai, and Fabio Miranda are with the University of Illinois Chicago. E-mail: {ssraba2, gmarai, fabiom}@uic.edu

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

there is still a lack of a structured and dedicated approach tailored specifically for visualizing street and pedestrian networks.

Given the significance of street and pedestrian network visualizations, understanding their design space is necessary. Prior work has explored design spaces for networks, geospatial networks, and map visualizations. For instance, Schöttler et al.'s survey reviewed visualization and interaction techniques for geospatial networks, particularly in contexts like social networks, and trade and migration [40], while Högreför et al. surveyed map-like visualization techniques [20]. However, there has not yet been a dedicated effort to create a design space specifically for street and pedestrian network visualizations. Visualizing data on streets poses several challenges, including overplotting, limited space for information encoding, and the risk of obscuring important spatial details [43, 44]. Additionally, aggregating data by geographic regions leads to a loss of spatial specificity, while simultaneously encoding multiple attributes on streets can result in visual clutter [23]. This limitation highlights the need for more structured guidance on effectively embedding multivariate data directly onto streets, ensuring a clear, integrated, and meaningful understanding of spatial information. To address this need, we conducted a qualitative study analyzing 45 existing visualizations overlaid on street maps. Through a systematic review of these works, we examined the analytical objectives that motivated the use of street-overlaid visualizations, the visual components employed to construct them, and the data types used to drive these visual encodings. We abstracted common tasks such as identifying specific segments of interest (e.g., high-crime or congested streets), comparing spatial or temporal conditions (e.g., weekday vs. weekend traffic, neighborhood accessibility), and analyzing patterns (e.g., uncertainty in travel, shifting volumes, regional movement, or congestion spread). Additionally, we analyzed how designers used visual channels, marks, and compositions to represent information directly on street networks. Finally, we studied the forms of data, ranging from raw observations to derived metrics, that informed these visualizations. This analysis provides a broad understanding of how street-overlaid visualizations are constructed and used across domains. Our design space serves as a systematic and actionable framework for developing visualizations tailored to diverse analytical needs for street and pedestrian networks. By leveraging this structure, StreetWeave supports task-specific visualization design, ensuring that users can effectively explore and interpret spatial data within their areas of interest.

Building on these insights, we introduce StreetWeave, a declarative grammar designed to simplify the creation of street and pedestrian network visualizations. StreetWeave lowers entry barriers by reducing the need for deep technical knowledge. Specifically, it supports easy prototyping, clear evaluation of visualization designs, and flexible adaptation across various applications and geographic contexts. The system is reusable and extensible, supporting visualization development across a range of scenarios and domains. StreetWeave's grammar-based approach eliminates the need for low-level code manipulation, enabling users to focus on what to visualize rather than how to implement it. This makes the system especially beneficial for our primary target users, urban planners, transportation engineers, public health analysts, climate researchers, and other domain experts, who often work with spatial and thematic data but may not have extensive programming experience. StreetWeave aims to bridge the gap between complex visualization techniques and practical usability. It supports reuse, cross-domain applications, and offers a simplified approach to embedding information directly within street networks. Unlike abstract network analysis tools, StreetWeave treats streets as structural scaffolds for organizing data while preserving geographic context. By lowering entry barriers, StreetWeave allows users to integrate diverse datasets directly onto streets and pedestrian pathways, enhancing the way we analyze and interact with urban spatial data. We demonstrate how StreetWeave enables the creation of diverse street and pedestrian network visualizations, facilitating clear exploration and in-depth analysis of spatial data across various applications. Our key contributions are:

- We propose a design space for street and pedestrian network visualizations, based on a review of 45 papers, organizing insights by analytical tasks, visualization techniques, and data types.

- We introduce StreetWeave, a high-level visualization grammar designed specifically for street and pedestrian network visualizations, enabling users to design interactive visualizations with minimal technical effort. StreetWeave is available at urbantk.org/streetweave.
- We report on a set of use cases inspired by previous works demonstrating StreetWeave's ease of use, flexibility, and reproducibility.

2 RELATED WORK

We begin by reviewing prior work on the design space of geospatial visualizations, then examine street-overlaid visualizations, and finally explore related work on visualization grammars and authoring tools.

2.1 Design spaces for geospatial visualizations

Data has always been important for understanding and navigating spaces [50]. From early hand-drawn maps to modern visual analytics systems [13, 27], advancements in data collection and analysis have enabled the creation of highly accurate and insightful maps [58]. This progress has highlighted the need for structured design spaces for map and geospatial network visualizations. Högreför et al. [20] proposed a design space for map-like visualizations by categorizing existing techniques from two perspectives: imitation and schematization. They focused on understanding how abstract data could resemble traditional maps (imitation) or how maps could be made more abstract for specific purposes (schematization). Each of these perspectives is divided into four categories based on the affected visual elements: point, line, area, and field. Their work provides a comprehensive taxonomy for map-like visualizations by grouping similar techniques within each category. Similarly, Zhang et al. [58] analyzed over 100 geo-infographics to establish a design space for geo-infographic creation, focusing on four main dimensions: basic map representations, data encoding channels, label design and placement, and highlighting techniques for key information. They even created a web-based authoring tool, MapCraft, to make geo-infographic creation accessible to a wider audience. Additionally, Schöttler et al. [40] conducted a survey exploring visualization and interaction techniques for geospatial networks, covering multiple fields to create a design space. This work categorizes visualization techniques based on geography representation, network depiction, composition, and interactivity. Despite the existence of design spaces for various geospatial visualizations, there is still no dedicated design space for street and pedestrian network visualizations that specifically guide the creation of these types of visual representations.

2.2 Street-overlaid visualizations

Street-overlaid visualizations have emerged as an approach for exploring and analyzing multivariate spatiotemporal data, effectively integrating geographical, temporal, and attribute-based information [23, 43, 44]. These visualizations display data directly on the geometry of street or pedestrian networks, using streets as the visual anchor to encode additional data attributes. As the volume and complexity of spatial data continue to grow, these visualizations enable users to uncover patterns, trends, and anomalies that are otherwise hard to detect while preserving the spatial structure of the network. Street-overlaid visualizations play a key role in several types of analyses, such as public transportation routes, route choice behavior, trajectory analysis, and overall traffic patterns [1–3, 11, 19, 25, 33, 46–49, 56, 57]. Moreover, several new methods have been developed to further enhance street-overlaid visualizations. Techniques such as occlusion-free route zooming [44], street enlargement [43], and aggregation methods for multilevel data analysis [9] have all contributed to making these visualizations more effective and easier to interpret. In addition to transportation analysis, street-level visualizations play a significant role in crime analysis [15, 16]. By overlaying crime incident data on street maps, analysts can quickly identify crime hotspots, monitor changes over time, and understand spatial crime patterns. This helps in developing targeted interventions and policies for improving public safety. Similarly, in urban planning, visualizing urban data directly at the street level provides critical insights into infrastructure needs, population density, and accessibility, facilitating better planning and decision-making processes [12, 22, 34, 45, 51].

Street-overlaid visualizations allow urban planners to evaluate scenarios, identify areas needing improvement, and create plans that enhance connectivity, accessibility, and quality of life.

2.3 Visualization grammars and authoring tools

Visualization grammars and authoring tools have become widely adopted approaches for simplifying the creation of interactive visualizations. Generally, these tools are grouped into three main categories: template-based tools, low-level visualization libraries, and high-level visualization grammars [13]. The template-based visualization tools typically utilize graphical interfaces, allowing users to create visualizations through drag-and-drop interactions without requiring programming knowledge [52]. Despite their ease of use, these tools offer limited flexibility and customization, making it difficult to handle the complex analytical tasks associated with street and pedestrian network visualizations. On the other hand, low-level visualization libraries, such as D3 [6], provide considerable freedom in creating custom visualizations. However, their use demands programming expertise, limiting accessibility to non-technical users. Additionally, these libraries often lack built-in capabilities for easily integrating multiple data layers and interactions. High-level visualization grammars represent a middle ground, offering a flexible yet accessible way to create visualizations through declarative specifications. Tools like Vega [39], Vega-Lite [38], the Urban Toolkit (UTK) [28], and similar grammar-based frameworks allow users to define visualizations through concise and declarative specifications. Users can quickly experiment with different visualization types, encoding strategies, and interaction methods without writing extensive code. Although powerful, these grammars typically lack specialized support for directly integrating and interacting with multiple thematic and physical layers specifically associated with street networks.

Inspired by these existing grammars and authoring tools, our work on StreetWeave aims to overcome these limitations. StreetWeave introduces a high-level visualization grammar specifically designed for creating street and pedestrian network visualizations. It allows users to create complex visualizations using simple, declarative rules, eliminating the need for low-level programming. StreetWeave also supports the integration of multiple thematic and spatial data layers directly onto street networks, offering both flexibility and usability for visualization and domain experts alike.

3 STREETWEAVE'S DESIGN GOALS

The design of StreetWeave follows a structured approach informed by prior research on street and pedestrian network visualizations [3, 23, 43, 44, 48], a review of 45 existing studies, and our analysis of key challenges. Our goal was to streamline the process of designing, prototyping, and sharing interactive visualizations of street and pedestrian networks. This comprehensive review shaped the development of StreetWeave, which unfolded in three main phases. In the first phase, we systematically analyzed the existing literature to identify limitations and tradeoffs in previous visualization tools and practices, with the results presented in Section 4. Next, we reflected on our personal experiences as visualization researchers and tool builders, focusing on ways in which a unified framework could enhance current methodologies. This reflection allowed us to define clear, detailed design objectives and capabilities for StreetWeave. Finally, during the framework's implementation phase, we continuously evaluated whether our established design goals were effectively met. Next, we present StreetWeave's design goals, inspired by our own UTK [28]. We refer to these goals throughout the subsequent sections, particularly when discussing StreetWeave's grammar (Section 5) and usage scenarios (Section 6). We also reflect on how well these goals were achieved in Section 7.3.

Ease of use and minimal technical barriers. Creating street and pedestrian network visualizations often requires programming skill, restricting their accessibility to non-technical domain experts. To address this issue, StreetWeave must provide a straightforward declarative grammar that simplifies complex technical tasks such as data loading, rendering, and interaction design. By reducing technical complexity,

StreetWeave must allow users to focus on visual analysis and exploration, instead of low-level implementation details, making advanced visualization accessible even to those without specialized programming experience.

Rapid prototyping and flexibility. Visual analytics workflows are inherently exploratory, often requiring users to iteratively test and refine visualization ideas. Traditional tools can hinder this process due to the extensive programming effort involved in adapting visualizations for new questions, datasets, or geographical contexts. StreetWeave must support rapid prototyping by enabling users to easily experiment with different visualization techniques, data encodings, and interactions within the same declarative framework. Additionally, StreetWeave must allow flexible adaptation across regions, scales, and datasets without requiring extensive reimplementations, supporting broader use by domain experts beyond the tool's original development context.

Extensibility. Beyond rapid iteration, street and pedestrian network analysis requires applying custom visual designs across diverse spatial scales, aggregation levels, and thematic layers. Insights often develop through shifting analytical perspectives and evolving hypotheses, which can be hindered by rigid tooling. StreetWeave must enable users to define and extend custom visualizations that align with varied analytical workflows, including advanced operations like data aggregation and comparative spatial analysis. Additionally, StreetWeave must support embedding Vega-Lite specifications within its grammar, leveraging a wealth of easy-to-use visualizations from the Vega community to extend its capabilities.

Reproducibility. Reproducing street and pedestrian network visualizations remains challenging. Many tools are closed-source and depend on complex, custom features and specialized libraries, making replication difficult even when data is available [13]. StreetWeave must address reproducibility through two strategies. First, it must use standard input data formats, allowing users to easily package and share datasets used in visualizations. Second, it must treat visualization configurations as structured, declarative specifications, so reproducing a visualization requires only sharing the high-level specification and corresponding datasets, enabling straightforward verification and reuse.

4 STREETWEAVE'S DESIGN SPACE

To develop a clear understanding of how street and pedestrian network visualizations are created and used, we first conducted a systematic survey of existing works. We compiled a diverse collection of 45 papers, focused on street-overlaid visualizations across various domains to identify common usage patterns and practices.

4.1 Selection process

To gain a comprehensive understanding of street and pedestrian network visualizations, we began by searching for papers that featured street-overlaid visualizations. Initially, we focused on proceedings and collections from major venues where visualization research is typically published, such as IEEE Transactions on Visualization and Computer Graphics, Computer Graphics Forum, ACM Conference on Human Factors in Computing Systems, and IEEE PacificVis. For each of these venues, we manually scanned the proceedings and retrieved papers based on their titles. If a title contained keywords such as 'street', 'trajectory', 'route', or 'road', we included those papers in our collection. We then manually examined each paper to confirm if it contained street-overlaid visualizations, discarding any that did not. Through this process, we initially gathered 25 papers from these venues. To expand our collection, we then conducted additional searches on Google Scholar using the same keywords. Through this method, we identified a substantial number of additional papers from various venues. We followed the same manual scanning process to verify whether these papers included street-overlaid visualizations, ultimately adding 20 more relevant papers to our collection.

4.2 Data coding

After compiling the 45 papers featuring street-overlaid visualizations, we conducted a structured thematic coding process to analyze how these visualizations were constructed and for what purposes. The first

author individually reviewed each paper to verify inclusion criteria and then systematically coded each figure using a framework informed by prior geospatial visualization literature [20,40,58]. The coding focused on three key dimensions: (1) analytical tasks (e.g., identify, compare, pattern analysis), (2) visual encodings (e.g., color, opacity, width, orientation), and (3) data types (e.g., direct observations vs. derived metrics). Each paper could include multiple visualization instances, and all were carefully annotated to record how these dimensions were used. Although coding was performed by a single researcher, detailed notes and criteria were maintained to ensure consistency across the dataset. The frequency of appearance for each visual encoding, task category, and data type was quantified and is reported in the corresponding sections of the paper (Sections 4.3.1, 4.3.2, and 4.3.3), where we break down the number of papers exhibiting each element. We acknowledge that including a coding protocol with multiple coders would further strengthen reliability and plan to incorporate that in future extensions of this work.

4.3 Design framework

To understand how these visualizations are constructed and used in practice, we structured the design space around common analytical tasks, visualization techniques, and data types observed across a broad set of research papers.

4.3.1 Analytical task abstraction

Understanding the analytical goals behind street-overlaid visualizations is essential for designing tools that effectively support user tasks. To identify these goals, we reviewed 45 research papers spanning domains such as urban planning, transportation, crime, and climate studies, extracting common analytical tasks supported by street-overlaid visualizations. In some cases, this extraction was challenging, particularly in papers centered on visualization techniques or algorithms without explicitly stated analytical objectives. Nevertheless, our review revealed several recurring analytical tasks across the surveyed studies, specifically **identify**, **compare**, and **pattern analysis** tasks. While these high-level tasks serve as an initial abstraction, StreetWeave explicitly targets concrete, realistic analytical scenarios informed by our qualitative analysis of existing street-overlaid visualizations.

Identify tasks focus on locating specific spatial areas, such as street segments, intersections, or sidewalk regions, that satisfy a known condition. These goal-driven tasks appeared in 93% of the reviewed papers. Examples include identifying roads with the highest nighttime crime rates [15,16], locating areas with high bicycle usage [19], pinpointing congested street segments [33,48], or highlighting sidewalks that are inaccessible to people with disabilities [45]. In urban planning, a targeted example is detecting sidewalk segments with significant accessibility issues (e.g., missing or inadequate curb ramps) to help planners prioritize interventions [24]. **Compare** tasks involve evaluating differences or similarities across spatial regions, such as streets, intersections, or neighborhoods. These tasks facilitate the assessment of spatial variation across the street network and were present in 88% of the reviewed papers. Examples include comparing traffic flow between streets [3,33,47] or assessing accessibility scores across different road segments [45]. A specific example is comparing sidewalk conditions between adjacent neighborhoods to guide equitable resource allocation [24,45]. **Pattern analysis** tasks center on detecting broader spatial or spatiotemporal trends and relationships across street networks. This category appeared in 62% of the reviewed papers. Examples include analyzing uncertainty in traffic behavior, monitoring changes in traffic volume [33,48], identifying mobility or population flow patterns [56], evaluating connectivity within street networks [2], or crime hotspots and clusters of service requests along streets to inform targeted safety or maintenance initiatives [15,16].

4.3.2 Visual component abstraction

This dimension examines the visual elements used in street-overlaid visualizations to align encoding choices with analytical tasks. We categorize these elements into **visual channels** (e.g., color, size, opacity),

marks (e.g., lines, points, glyphs), and **compositions** (e.g., juxtaposition, superimposition). Color hue was the most frequently used channel (78% of the papers), allowing clear differentiation of features [8,9,25,45,51], while color intensity (33%) conveyed variations in data values [15,23,49] supporting identify, compare, and pattern analysis tasks. Opacity (7%) indicated data density or confidence through transparency [23]. Geometric properties such as length, width, and height further supported multivariate spatial data representation, while additional channels like size, density, dashed lines, and blur encoded complementary attributes. For marks, lines (51%) were commonly used, often combined with color and opacity for multivariate encoding. Points and circles conveyed data at specific locations, while waves and curved lines occasionally depicted uncertainty or directional flows. Glyphs provided compact representations of multivariate and spatiotemporal data, supporting complex pattern analysis and comparison [5,22]. Labels in text or symbol form aided direct identification and interpretation [15,22]. In terms of composition, we observed that juxtaposition and superimposition were the dominant strategies in street-overlaid visualizations. Juxtaposed views (38% of the papers) placed visualizations side-by-side for coordinated comparison [43,44,49], while superimposed views (used in all papers) layered visual elements directly on the street network [9,23,33], supporting interpretation while risking clutter when encoding multiple variables.

4.3.3 Data abstraction

Our review shows that street-overlaid visualizations primarily use data that can be categorized into two broad types: **direct observation** and **derived metrics**, both requiring geographic references (e.g., latitude/longitude or segment identifiers) to align data with street networks. **Direct observations** refer to raw data collected without further transformation or spatial computation. Examples include GPS traces capturing pedestrian or vehicle trajectories [19,33], sensor readings recording traffic speed [11], and manual audits of sidewalk conditions [45]. These values can be mapped directly onto the street network based on their geographic coordinates or segment associations. **Derived metrics** are generated through processing, aggregation, or spatial analysis over raw data. These include summary statistics such as the average traffic volume on a street segment over a week [33,47], the flood vulnerability recorded at intersections [8], or aggregated counts of crime incidents along streets [15,16]. Derived metrics can also involve complex spatial computations, such as calculating accessibility scores [45,51].

5 STREETWEAVE’S GRAMMAR

Based on the design goals and insights from our systematic review, we introduce StreetWeave, a declarative visualization grammar tailored for street and pedestrian network visualizations. StreetWeave provides clear, structured rules for specifying visualizations, covering spatial units (e.g., area, segment, node), multiple data layers (thematic and physical data integrated via spatial relations), and diverse visual encoding options. Spatial units serve as the grammar’s foundation, with all other specifications defined relative to each visualization layer. In the following sections, we detail the grammar, organizing it into three parts: (1) general grammar specification, (2) data layer definitions (including thematic and physical data, spatial relations, and analytical operations), and (3) data plotting specifications. For formal presentation, we adopt the notation of Ren et al. [36], where $:=$ denotes assignment, $|$ indicates “or”, $?$ marks optional elements, and $+$ represents one or more elements.

5.1 General specification

To structure the StreetWeave visualization grammar, we begin by defining a clear general specification. Each layer in the grammar comprises three key components: a map, a unit (Section 5.2), and data layer aggregations (Section 5.3). All properties include recommended default settings to guide users through the visualization process while minimizing the need for extensive manual adjustments, particularly for those with limited technical expertise.




Fig. 2: Integrating thematic and physical layers through spatial relations. StreetWeave integrates thematic data with the physical street network by applying spatial relations at each street segment or intersection. The nearest neighbor relation selects the 50 closest data points to a segment or intersection, while the contains relation includes only those within a segment's or intersection's boundaries. The buffer relation gathers points within a user-defined radius from a segment's or intersection's midpoint. Aggregation operations like mean or sum can be applied to these points for analysis.

```
StreetWeave := (map+, unit+, data+, relation?, query?)
map := streetColor?, streetWidth?, background?
background := 'light' | 'dark'
```

The main configuration defines the foundational visual and interactive settings for presenting street and pedestrian network visualizations. It specifies **map**, **unit**, **data**, **relation**, and **query**. Additional map attributes allow customization of the visualization's appearance, such as selecting the **background** style as light or dark, setting **streetColor** to any desired color, and specifying **streetWidth** to control line thickness.

5.2 Unit and visual encodings specification

In StreetWeave's grammar, the concept of a unit defines both the spatial granularity for data aggregation and analysis, and the corresponding visualization specification.

```
unit := type, density?, method, opacity?, color?, dash?,
squiggle?, width?, height?, chart?, rows?, columns?,
orientation?, alignment?, zoom?
type := 'segment' | 'node' | 'point'
method := 'line' | 'rect' | 'matrix'
chart := vega_spec
orientation := 'parallel' | 'perpendicular'
alignment := 'left' | 'center' | 'right'
zoom := [ minZoom, maxZoom ]
```

The StreetWeave grammar supports flexible spatial resolutions through the **type** property, which specifies units as *segment*, *node*, or *point*. A *segment* represents individual street segments and can be subdivided using the **density** property, which accepts a fixed value or a thematic variable for finer granularity, such as visualizing sidewalk severity (Figure 5). This subdivision enables users to control the spatial resolution of their analysis, allowing denser sampling in areas of interest while maintaining coarser aggregation elsewhere. By mapping **density** to a data attribute, users can simultaneously encode an additional variable, supporting multivariate analysis through combinations of **density** and other visual properties (e.g., color, width). A *node* corresponds to intersections, enabling aggregation and visualization of data at street junctions (Figure 6). A *point* allows the placement of point-based thematic data, such as event or sensor locations, directly onto the network. StreetWeave also supports scale-dependent rendering through the **zoom** property, which specifies the zoom levels at which a unit will be displayed. The **zoom** property is defined as a range, [minZoom, maxZoom], so that users can control the visibility of units depending on the current map zoom level. This feature enables the

creation of multi-scale visualizations, where detailed information can be shown at higher zoom levels and clutter is reduced when viewing the network at broader scales.

StreetWeave's grammar enables embedding visual encodings directly on street networks using two approaches. The first utilizes graphical methods specified by **method** (e.g., *line*, *rect*, *matrix*), which can be customized with attributes such as color, opacity, width, and height. Additional stylistic properties like **dash** and **squiggle** allow visual differentiation, aiding in the representation of uncertainty or categorical distinctions (Figure 1 (B, C)). Importantly, these properties can either be assigned fixed values or mapped dynamically to data attributes, supporting the representation of quantitative or categorical information directly within the visualization. The second approach leverages the **chart** property, allowing users to embed Vega-Lite specifications directly into StreetWeave visualizations. This integration provides access to a wide range of pre-defined visualization types, facilitating the inclusion of complex, multivariate encodings while maintaining alignment and orientation on the street network. The **orientation** property further extends expressive capacity by allowing users to render visual elements either *parallel* or *perpendicular* to the street network. For example, rendering a perpendicular bar chart along a street segment while encoding color parallel to the segment can help users interpret traffic volume alongside congestion levels without visual overlap.

5.3 Data and spatial relations specification

StreetWeave's grammar enables users to load, visualize, integrate, and apply analytical operations directly on data layers, classified into two distinct categories: thematic layers and physical layers. This separation, inspired by our prior work on UTK [28], facilitates clear organization of spatial data for analysis and visualization. Thematic layers represent data attributes measured at specific locations, such as crime incidents, pollution levels, or noise complaints. Typically, thematic data consists of discrete points, each with spatial coordinates and associated values. In contrast, physical layers depict structural or environmental features, including road networks, pedestrian networks, and intersections. They define the geometric reference on which visualizations are overlaid, giving spatial context for interpreting thematic data.

```
data := physical?, thematic?
physical := path
thematic := path, latColumn?, lonColumn?
relation := spatial, value, typerelation
spatial := 'buffer' | 'nn' | 'contains'
typerelation := 'sum' | 'mean' | 'min' | 'max'
```

In addition, StreetWeave computes the bearing angle of street segments to capture their orientation, aiding in the analysis of movement

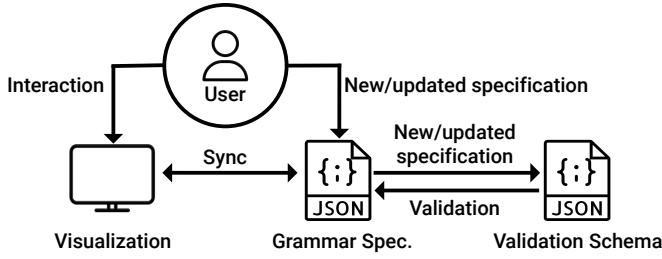


Fig. 3: Overview of StreetWeave’s system architecture and execution pipeline. User specifications are processed, interpreted, and rendered to support interactive network visualizations.

patterns and directional dependencies within the street network. The grammar supports spatial operations to integrate thematic and physical layers effectively (Figure 2). By defining a spatial relation (e.g., buffer with a specified radius, contains, or nearest neighbor), thematic data can be linked to physical geometries, enabling spatially grounded analyses. When multiple thematic data points correspond to a single physical geometry, the grammar supports aggregation operations (e.g., mean, sum, max, min) to summarize these values. For example, pedestrian counts or noise levels recorded at discrete locations can be aggregated onto street segments or intersections based on the defined spatial relation. In doing so, StreetWeave allows analyses such as evaluating accessibility around intersections or summarizing pollution levels along street segments within a consistent spatial framework.

5.4 Query specification

With StreetWeave, users can selectively modify styling (e.g., width) of segments using two optional query parameters: a geographic address and a radius around that address.

query:= address[?] , radius[?]

By specifying an **address** and a **radius**, users can define geographic areas within which width modifications will apply, supporting localized adjustments without affecting the entire network. For example, users can emphasize streets within a specific neighborhood to highlight areas with high pedestrian activity, event impact zones, or regions undergoing construction.

5.5 Implementation

StreetWeave is implemented as a grammar-based toolkit using JavaScript. Users specify visualizations by authoring a JSON-based specification that follows the grammar previously introduced. This specification is parsed and validated against StreetWeave’s JSON schema to ensure syntactic correctness and to guide users during the authoring process. Users explicitly define the layers, visual encodings, and operations required for their visualization, which are then rendered directly on the street or pedestrian network within the interface. Input data includes physical layers provided in JSON format and thematic layers supplied as CSV files. StreetWeave’s interface consists of two primary components: a text editor and a map view panel. The text editor allows users to input and edit their grammar specifications, which are validated using a JSON schema to ensure correctness and provide guidance. Users reference their physical and thematic data layers directly within the specification. Once specified, visualizations are rendered in the map view panel, allowing users to explore and refine designs iteratively. Figure 3 presents a high-level overview of StreetWeave’s system architecture. StreetWeave uses React.js for interface management, Leaflet.js for interactive map rendering, and D3.js for custom SVG-based overlays and data-driven visual encodings. By integrating these libraries, StreetWeave supports smooth interaction capabilities, including panning and zooming, while maintaining efficient rendering performance even with complex, multivariate visualizations. To

support extensibility and user experimentation, StreetWeave allows the embedding of Vega-Lite specifications within its grammar. An interpreter reads the user’s JSON-based visualization specification and dynamically generates corresponding SVG elements, which are then appended to the map. Because StreetWeave operates on a layer-based map architecture, it can leverage any mapping platform that supports layer management.

6 USAGE SCENARIOS

To demonstrate StreetWeave’s capabilities, flexibility, and ease of use, we present a series of scenarios grounded in real-world data challenges. These examples illustrate how StreetWeave can integrate diverse datasets to generate detailed street-overlaid visualizations that support in-depth analysis and interpretation.

6.1 Scenario 1: Analyzing sidewalk accessibility data

To illustrate StreetWeave’s capabilities and flexibility, we present an analysis scenario using the Project Sidewalk [37] dataset for Chicago, which contains geo-located data on sidewalk accessibility issues, such as severity of crosswalk and curb ramp problems, missing curb ramps, absent sidewalks, obstacles, and surface issues. In this example, we visualize the distribution and severity of these accessibility issues across Chicago’s street network, enabling simultaneous exploration of multiple concerns. Using StreetWeave’s grammar, we define a visualization where each line segment encodes different thematic attributes through width, allowing users to analyze various conditions in a single integrated view. We begin by loading the physical layer data (geometric information for street segments) and the thematic layer data (sidewalk severity ratings in CSV format). Using StreetWeave’s grammar within the text editor, we specify spatial relations between thematic and physical layers, applying a *buffer* relation with a 10-meter radius and a *mean* aggregation type to summarize sidewalk severity onto street segments. For visualization, we select *segment* as the unit type to display data at the street level. The center-aligned line encodes curb ramp severity through its width. A left-aligned purple line encodes the severity of the surface problem using width. Additionally, a right-aligned line represents missing sidewalks, also using width for encoding severity.

After specifying these configurations, we apply the specification, and the map view panel instantly updates to display the multi-encoded visualization (Figure 4 (left)). Users can interact with the map by zooming and panning to identify problematic street segments and prioritize areas for infrastructure improvements. Because sidewalk accessibility data is highly multivariate, StreetWeave’s ability to handle multiple data attributes within a single view is particularly valuable, enabling users to analyze patterns and identify critical areas based on several factors.

6.2 Scenario 2: Detailed sidewalk accessibility analysis

In the previous example, we demonstrated how StreetWeave can identify problematic street segments across Chicago by visualizing sidewalk accessibility data directly on the streets. However, marking an entire street segment as problematic may not provide the granularity needed for targeted urban improvements. To pinpoint which portions of a segment exhibit accessibility challenges, a finer-grained analysis is necessary. For this example, we again utilize the Project Sidewalk dataset but we now focus on sub-segment granularity by subdividing each street segment into smaller sections. This enables spatial aggregation and visualization on each sub-segment, providing more accurate and detailed insights into sidewalk accessibility conditions. To visualize this enhanced detail, we first create a line map that encodes multiple accessibility conditions using **color** and **width**, with the visual encoding oriented *parallel* to the street (Figure 5 (left)). By simply changing the **orientation** to *perpendicular*, **method** to *line*, and adding the **height** dimension, we produce a bristle-style visualization inspired by Bristle Maps [23] (Figure 5 (middle)). In this variant, each perpendicular “bristle” extending from the street represents a sub-segment, with its visual attributes directly reflecting sidewalk accessibility data. Users can encode multiple variables simultaneously in different combinations, such as color and height or color and width, providing clear visual cues



Fig. 4: Visualizing the sidewalk accessibility severity. Left: A multi-line visualization maps sidewalk accessibility attributes onto street segments, using line width to encode data values across three overlaid layers: a center-aligned green line encodes curb ramp severity, a right-aligned brown line represents missing sidewalks, and a left-aligned purple line shows surface problem severity. Right: A multivariate encoding combines color and width to represent multiple data attributes simultaneously. The corresponding StreetWeave specifications demonstrate the grammar-based approach that configures and adapts these multivariate encodings.

that highlight different aspects of the data. This configuration allows users to observe detailed variations in accessibility conditions within a single street segment, which may be too subtle to detect in a parallel line map.

This rapid reconfiguration is achieved by modifying a few attributes in StreetWeave’s grammar. Users can explore and compare multiple visualization setups without additional effort. StreetWeave also supports the simultaneous display of multiple visualizations on the same map, such as displaying different bristle maps on either side of the street, enabling direct comparisons and deeper spatial analysis. In Figure 5 (right), two layers of bristle maps are shown side by side, illustrating how different accessibility conditions can vary along the same street segment, with some conditions showing low severity while others exhibit higher severity.

6.3 Scenario 3: Exploring crime and 311 service data

In addition to sidewalk accessibility, understanding patterns in urban crime and city service requests is essential for creating safer and more responsive urban environments. For this scenario, we demonstrate how StreetWeave can simultaneously visualize and analyze two distinct yet interrelated datasets: crime incidents and 311 service requests in Chicago, showcasing StreetWeave’s flexibility and support for iterative visualization design. We begin by visualizing crime incident data aggregated at street intersections (nodes). Leveraging StreetWeave’s integration with Vega-Lite, we embed compact charts directly within the visualization, using *contains* as the spatial relation and *mean* as the aggregation type to summarize crime data at each intersection. This approach highlights the distribution of crime across intersections while demonstrating StreetWeave’s extensibility through seamless incorporation of external visualization libraries. Next, we analyze 311 service requests from 2020 to 2024, plotting them along street segments. We again use the *contains* spatial relation, now with the *sum* aggregation type, to represent the frequency of service requests along each segment. Initially, we applied a *perpendicular* orientation for these encodings (Figure 6 (top)). However, upon reviewing the visualization, we found that switching the orientation to *parallel* (Figure 6 (bottom)) improved readability by aligning the encodings with street paths, facilitating easier trend tracing. As shown in Figure 6 (bottom), clear differences in service request distributions are visible even on adjacent streets (highlighted by the rectangle), revealing varying service needs within small geographic areas and underscoring StreetWeave’s capacity for detailed spatial analysis. The visualizations in this scenario serve as

illustrative examples of StreetWeave’s grammar, demonstrating that it is not limited to these encodings. The underlying grammar remains extensible, supporting the incorporation of additional types or custom visual components as needed for diverse analytical tasks.

This iterative design process highlights StreetWeave’s strengths in rapid prototyping, enabling users to experiment with different visualization parameters and immediately assess their effectiveness. By facilitating the simultaneous visualization of crime and service request data, StreetWeave helps users explore potential relationships or divergences between crime hotspots and areas with high service demand. These insights can assist urban planners and city officials in prioritizing safety and service improvements effectively.

7 DISCUSSION

In this section, we first discuss feedback gathered from domain experts, then compare our contributions with existing visualization tools, and finally reflect on how StreetWeave’s design goals were met.

7.1 Expert feedback

To evaluate StreetWeave’s utility and alignment with real-world needs, we conducted semi-structured interviews with five experts working extensively in the urban domain. None of the participants are authors of this paper. Participants included one expert with a PhD in Urban Systems [E1](#), one with a PhD in Computer Science [E2](#), one with a PhD in Civil Engineering [E3](#), one PhD student in Computer Science [E4](#), and one PhD student in Urban Planning [E5](#). Quotes from participants have been lightly edited for clarity.

Ease of use and workflow simplification emerged as a consistent theme, with experts appreciating how StreetWeave can reduce the complexity of creating multivariate, street-overlaid visualizations compared to D3, GeoPandas, or multi-step pipelines. [E3](#) noted that she “has used D3, and it is complicated,” adding that StreetWeave “is easy (...) it is good that you do not need to think about the code beneath it.” Similarly, [E4](#) noted that she “never felt comfortable with D3.” While [E3](#) acknowledged that “multilayers can be done in GeoPandas,” she explained that it required “creating multiple maps,” whereas “StreetWeave can easily do this.” [E2](#) emphasized that “while D3 is valuable for heavy customization, it requires effort to create visualizations.” She highlighted StreetWeave’s usefulness for analyzing climate data, including temperature, rain, and wind, and visualizing heat waves and urban heat islands at street level. [E4](#) highlighted that “it normally takes a very long time to create street-overlaid visualizations, requiring multiple

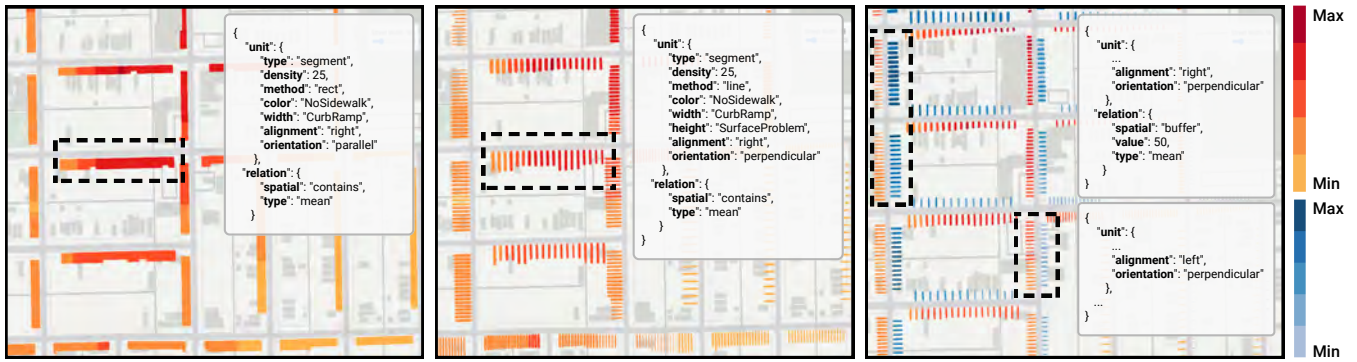


Fig. 5: Fine-grained street segment analysis. Left: A line map displays sidewalk accessibility severity using variations in color and width. Middle: A bristle map incorporates an additional height dimension to more accurately highlight problematic areas, as seen in the highlighted regions. Right: Dual bristle maps, aligned to the left and right, visualize multivariate data side-by-side, providing another perspective on localized accessibility issues.

software packages,” while StreetWeave “provides a simple domain-specific language that even non-coders can use.” E5, who had not used Vega-Lite before, found StreetWeave very interesting and expressed that he would have used it in previous projects had he known about it earlier. He appreciated the ability to “set up a map once and then change the visualization quickly without extra coding,” contrasting this with his usual workflow in GeoPandas, where “you need to know what you want beforehand,” whereas StreetWeave “lets you quickly try out different design ideas to see what works.” E5 noted that while Python felt intuitive for spatial work, StreetWeave’s declarative commands offered faster iteration on visual styles, reducing the need to repeatedly adjust styles manually as he currently does in GeoPandas.

E1 highlighted a promising new application for StreetWeave in the uncertainty analysis of computer vision models used for built environment audits, noting that StreetWeave “could enable per-segment evaluation rather than more coarse block-level visualizations.” However, E1 also pointed out a potential limitation, explaining that many models do not work at the segment level, but rather at a polygon level, such as models that extract sidewalk networks from satellite imagery. She noted that “such cases would require additional computations to map the uncertainty to edges” to fully leverage StreetWeave in these contexts.

Reproducibility was also noted as a major strength, with experts recognizing the value of sharing JSON specifications for replicable research workflows. E3 noted that “you can share the JSON files and that definitely ensures the reproducibility, which is a big concern when you’re doing research.” E3 also emphasized that reproducibility extends beyond sharing static specifications, suggesting that coupling StreetWeave’s outputs with clear data provenance could further strengthen transparency and facilitate validation across research projects.

The experts also suggested further extensions. E4 mentioned the option to embed functions within the grammar, “giving advanced users more customization capabilities,” as well as the ability to “auto-detect coordinate reference systems to reduce friction for non-GIS users.” Furthermore, experts noted the value of enabling StreetWeave to interoperate with standard tools like QGIS and ArcGIS. E5 expressed interest in additional usability features, including the ability to add titles, legends, and axis labels, as well as options for exporting maps as PDF or image files for inclusion in reports and presentations.

7.2 Comparison with existing tools

In this section, we compare StreetWeave with widely used GIS tools like ArcGIS and QGIS. While powerful, these tools often overwhelm users with extensive operations and toolboxes requiring specialized knowledge [59]. They also lack direct support for effective street-overlaid visualizations, pushing users towards complex workflows or custom scripting. In contrast, StreetWeave offers a concise, grammar-based approach tailored for street and pedestrian network visualizations. Its clear specification reduces complexity, enabling rapid prototyping

and iteration without extensive technical expertise. Traditional GIS tools also struggle with reproducibility and easy sharing due to licensing constraints and complex dependencies. StreetWeave addresses this by storing configurations as grammar specifications paired with JSON data, supporting collaboration and replication. Unlike existing GIS tools, it enables flexible, multivariate visualizations directly on street networks using diverse encodings, enhancing contextual interpretation and interactive analysis.

In addition to GIS platforms, StreetWeave also differs significantly from conventional network visualization tools such as Gephi [4], Cytoscape [41], and Graphviz [10], which are primarily designed to visualize abstract relational networks, such as social graphs, biological pathways, or citation networks, using algorithmic layouts like force-directed or radial positioning. These tools prioritize structural features like node centrality, clustering, and connectivity, often abstracting away geographic space entirely. In contrast, StreetWeave focuses on spatially grounded networks where the geometry of streets and intersections is essential to the analytical task. Rather than exploring topological relationships, StreetWeave supports multivariate spatial analysis by aligning thematic data directly with physical street geometry and enabling flexible visual encodings at multiple spatial resolutions. This makes StreetWeave more suitable for urban domains where tasks revolve around comparing infrastructure quality, accessibility, or socio-environmental indicators across real-world locations, rather than navigating purely relational graph structures.

7.3 Reflection on design goals

StreetWeave was designed to address the critical challenges of street and pedestrian network visualizations, emphasizing accessibility, flexibility, extensibility, and reproducibility. Through the examples demonstrated here, we have shown how StreetWeave meets each of these previous objectives. In terms of **ease of use and minimal technical barriers**, Examples 1 and 2 illustrate the accessible design of the StreetWeave grammar. By simplifying the process of creating custom visualizations through an easy-to-understand, JSON-based grammar, StreetWeave lowers the entry barriers for non-technical users. The ability to define complex encodings, such as multivariate lines with varying orientations, not only enables nuanced and expressive visualization designs but also empowers visualization practitioners and domain experts to construct detailed analyses without requiring extensive programming knowledge, allowing them to focus directly on analytical insight rather than implementation.

The goal of **rapid prototyping and flexibility** is highlighted in Examples 2 and 3. By allowing quick alterations, such as changing the visualization orientations (from perpendicular to parallel) and method (from rect to line), StreetWeave facilitates iterative experimentation. Users can rapidly refine visualizations, testing various dimensions to identify the best presentations that effectively communicate complex data relationships. This capability shortens the iterative cycles common in exploratory visual analytics.

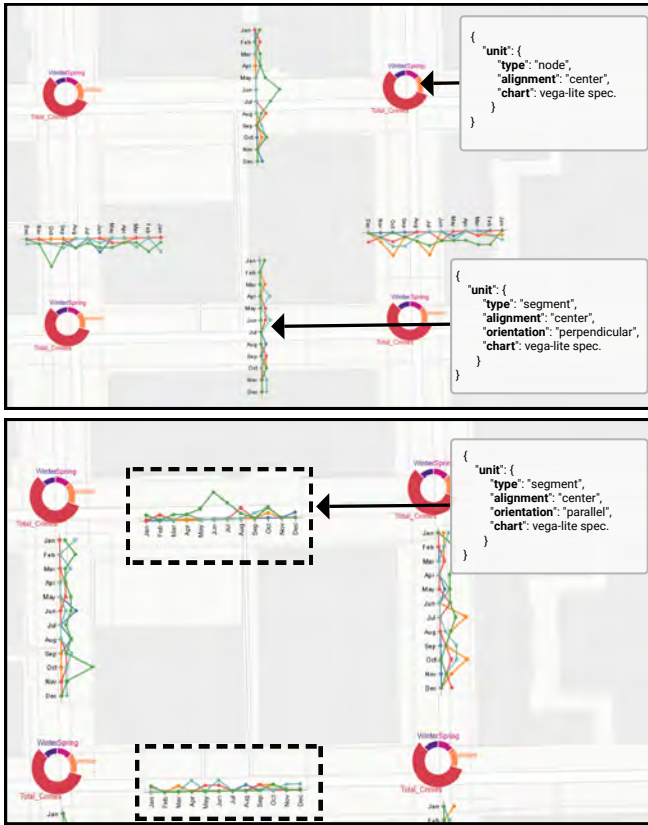


Fig. 6: Visualizing crime patterns and 311 service requests. Top: Charts are applied to street intersections (radial charts) and street segments (line charts) using perpendicular orientation, enabling multi-layered representation. Bottom: Charts are rendered parallel to the streets for clearer presentation. This reveals noticeable contrasts in service call distributions, even between adjacent streets (as highlighted by the rectangle).

Extensibility is primarily addressed through the integration of external visualization capabilities, namely Vega-Lite, as shown in Example 3. This showcases how StreetWeave can incorporate an additional framework and embed external visual encodings without requiring fundamental changes to its underlying architecture. Users can directly extend the grammar to include diverse visualization types, providing the flexibility to adapt to evolving analytical demands.

Regarding **reproducibility**, StreetWeave emphasizes standardized data formats and declarative grammar specifications. Users can share their visual analyses by simply exchanging grammar specifications and datasets in JSON format, significantly enhancing the reproducibility and transparency of their workflows. This approach facilitates broader verification and adoption within both the visualization research community and among domain practitioners.

8 LIMITATIONS, CONCLUSIONS, AND FUTURE WORK

Limitations. While StreetWeave demonstrates the potential of grammar-based street network visualization, it has a few limitations. First, we have not yet conducted a formal user study to evaluate usability and effectiveness across users with diverse knowledge levels, domain expertise, and disciplinary backgrounds. Second, we have not quantitatively assessed how comprehensively StreetWeave can reproduce visualization designs within the broader design space of urban street visualization, limiting our understanding of coverage and expressiveness. Third, StreetWeave currently lacks a modular plugin system for custom visualization components outside of Vega-Lite, restricting more advanced and bespoke views. Fourth, StreetWeave relies on SVG-based rendering, which can become slow when visualizing large street networks or several data layers, restricting performance in very complex scenarios. Furthermore, the system’s default values for encod-

ings, colors, and sizing are based on practical experience rather than empirical studies, which may not align optimally with user preferences or perceptual best practices. Finally, while the grammar simplifies reproducibility, integrating StreetWeave workflows with existing GIS or analysis pipelines still requires manual coordination, limiting end-to-end workflows.

Conclusions. In conclusion, we introduced StreetWeave, a declarative visualization grammar designed to streamline the creation, prototyping, and sharing of street and pedestrian network visualizations. Through a systematic review of 45 prior studies, we analyzed the analytical purposes, visualization techniques, and data types employed in street-overlaid visualizations. This analysis informed a structured design space that captures these insights, providing clear guidance to visualization designers across diverse domains. To bring the design space into practice, we developed StreetWeave, a grammar-based specification specifically tailored for street and pedestrian network visualizations. StreetWeave reduces technical complexity, enabling users without extensive programming experience to design custom visualizations. The use cases presented demonstrate StreetWeave’s ease of use, flexibility, and reproducibility in practical scenarios, including analyzing sidewalk accessibility and exploring crime and service request patterns. We see StreetWeave as a step towards bringing declarative grammars to urban-specific visualization challenges. By capturing urban domain requirements with a grammar, we hope to pave the way for broader adoption of visualization grammars in urban analytics.

Future work. StreetWeave opens new research opportunities for advancing street network visualization. Much like how our evaluation study in Mota et al. [30] on 3D urban visualization catalyzed a series of papers expanding methods and applications, we hope StreetWeave will similarly spark future research in street-overlaid visualization. Building on our prototype, we will conduct formal user studies across diverse domains and user expertise levels to evaluate the effectiveness of different street-overlaid visualizations. These studies will also help ground StreetWeave’s default design and encoding choices in empirical evidence, aligning them with yet-to-be-uncovered perceptual best practices, guidelines, and user preferences to ensure effective communication for diverse urban analysis tasks. We also plan to quantitatively assess how comprehensively StreetWeave can recreate visualization designs within the broader urban street visualization design space, clarifying its coverage, expressiveness, and identifying gaps for further extension. We will situate this work within our broader efforts on domain-specific visualization grammars, including the ongoing development of the UTK [28], positioning StreetWeave as a reusable component in creating reproducible, urban-specific visualization tools. We also envision extending StreetWeave to enable “what-if” analyses, allowing users to modify street structures, such as simulating widened roads or added pedestrian paths, and observe impacts on walkability, accessibility, and traffic flow. Achieving this will require deeper integration with simulation and predictive models, expanding StreetWeave’s utility for urban planning, policy evaluation, and participatory design contexts. Building on this, we see promising opportunities to develop grammars for broader scenario planning tasks, enabling experts to define and compare development or policy scenarios more systematically through linked visualization specifications. From a technical standpoint, to address current performance limitations, we plan to port StreetWeave from SVG to other rendering systems such as WebGL and WebGPU, better providing scalable, interactive rendering for large networks, and supporting physical layer deformation, similar to UrbanRama [7]. Additionally, we plan to integrate StreetWeave with the UTK and Curio frameworks [28, 29], enabling users to incorporate StreetWeave visualizations as modular components within broader urban data analysis workflows. From a community standpoint, we also plan to create a “StreetWeave Gallery” for visualization specifications to be shared for reuse in other cities and projects.

ACKNOWLEDGMENTS

We thank the reviewers for their constructive feedback, the U.S. National Science Foundation (#2320261, #2330565, #2411223), and the U.S. National Institutes of Health (R01CA258827, UG3TR004501).

REFERENCES

- [1] S. Al-Dohuki, Y. Wu, F. Kamw, J. Yang, X. Li, Y. Zhao, X. Ye, W. Chen, C. Ma, and F. Wang. Semantictraj: A new approach to interacting with massive taxi trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):11–20, 2016. [2](#)
- [2] N. Andrienko, G. Andrienko, F. Patterson, and H. Stange. Visual analysis of place connectedness by public transport. *IEEE Transactions on Intelligent Transportation Systems*, 21(8):3196–3208, 2019. [2, 4](#)
- [3] N. Andrienko, G. Andrienko, and S. Rinzivillo. Leveraging spatial abstraction in traffic analysis and forecasting with visual analytics. *Information Systems*, 57:172–194, 2016. [2, 3, 4](#)
- [4] M. Bastian, S. Heymann, and M. Jacomy. Gephi: An open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 3, pp. 361–362, 2009. [8](#)
- [5] R. Borgo, J. Kehrler, D. H. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics*, pp. 39–63, 2013. [4](#)
- [6] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011. [3](#)
- [7] S. Chen, F. Miranda, N. Ferreira, M. Lage, H. Doraiswamy, C. Brenner, C. Defanti, M. Koutsoubis, L. Wilson, K. Perlin, and C. Silva. UrbanRama: Navigating cities in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4685–4699, 2022. [9](#)
- [8] D. Cornel, A. Konev, B. Sadransky, Z. Horváth, E. Gröller, and J. Waser. Visualization of object-centered vulnerability to possible flood hazards. *Computer Graphics Forum*, 34(3):331–340, 2015. [4](#)
- [9] Z. Deng, S. Chen, X. Xie, G. Sun, M. Xu, D. Weng, and Y. Wu. Multi-level visual analysis of aggregate geo-networks. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):3135–3150, 2022. [2, 4](#)
- [10] J. Ellson, E. Gansner, L. Koutsofios, S. C. North, and G. Woodhull. Graphviz—open source graph drawing tools. In *International Symposium on Graph Drawing*, pp. 483–484. Springer, Berlin, Heidelberg, 2002. [8](#)
- [11] J. Fawcett and P. Robinson. Adaptive routing for road traffic. *IEEE Computer Graphics and Applications*, 20(3):46–53, 2000. [2, 4](#)
- [12] Z. Feng, H. Li, W. Zeng, S.-H. Yang, and H. Qu. Topology density map for urban data visualization and analysis. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):828–838, 2020. [2](#)
- [13] L. Ferreira, G. Moreira, M. Hosseini, M. Lage, N. Ferreira, and F. Miranda. Assessing the landscape of toolkits, frameworks, and authoring tools for urban visual analytics systems. *Computers & Graphics*, 123:104013, 2024. [2, 3](#)
- [14] M. Fu, K. Ban, L. Jin, and D. Wu. How urban street spatial composition affects land surface temperature in areas with different population densities: A case study of Zhengzhou, China. *Sustainability*, 16(22):9938, 2024. [1](#)
- [15] G. Garcia-Zanabria, E. Gomez-Nieto, J. Silveira, J. Poco, M. Nery, S. Adorno, and L. G. Nonato. Mirante: A visualization tool for analyzing urban crimes. In *SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 148–155. IEEE, 2020. [2, 4](#)
- [16] G. Garcia-Zanabria, M. M. Raimundo, J. Poco, M. B. Nery, C. T. Silva, S. Adorno, and L. G. Nonato. Cripav: Street-level crime patterns analysis and visualization. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4000–4015, 2021. [2, 4](#)
- [17] F. Guo, M. Luo, C. Zhang, J. Cai, X. Zhang, H. Zhang, and J. Dong. The mechanism of street spatial form on thermal comfort from urban morphology and human-centered perspectives: A study based on multi-source data. *Buildings*, 14(10):3253, 2024. [1](#)
- [18] K. Hara. Scalable methods to collect and visualize sidewalk accessibility data for people with mobility impairments: An overview. *SIGACCESS Accessibility and Computing*, (111):34–37, 2015. [1](#)
- [19] T. He, J. Bao, S. Ruan, R. Li, Y. Li, H. He, and Y. Zheng. Interactive bike lane planning using sharing bikes’ trajectories. *IEEE Transactions on Knowledge and Data Engineering*, 32(8):1529–1542, 2019. [2, 4](#)
- [20] M. Hografer, M. Heitzler, and H.-J. Schulz. The state of the art in map-like visualization. *Computer Graphics Forum*, 39(3):647–674, 2020. [2, 4](#)
- [21] Y. Hu, F. Qian, H. Yan, A. Middel, R. Wu, M. Zhu, Q. Han, K. Zhao, H. Wang, F. Shao, et al. Which street is hotter? Street morphology may hold clues-thermal environment mapping based on street view imagery. *Building and Environment*, 262:111838, 2024. [1](#)
- [22] X. Huang, Y. Zhao, C. Ma, J. Yang, X. Ye, and C. Zhang. TrajGraph: A graph-based visual analytics approach to studying urban network centralities using taxi trajectory data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):160–169, 2015. [2, 4](#)
- [23] S. Kim, R. Maciejewski, A. Malik, Y. Jang, D. S. Ebert, and T. Isenberg. Bristle maps: A multivariate abstraction technique for geovisualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(9):1438–1454, 2013. [2, 3, 4, 6](#)
- [24] C. Li, K. O. Y. Ma, M. Saugstad, K. Fujii, M. Delaney, Y. Eisenberg, D. Labbé, J. L. Shanley, D. Snyder, F. P. Thomas, et al. “I never realized sidewalks were a big deal”: A case study of a community-driven sidewalk accessibility assessment using Project Sidewalk. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI ’24, pp. 1–18. ACM, 2024. [4](#)
- [25] M. Lu, C. Lai, T. Ye, J. Liang, and X. Yuan. Visual analysis of route choice behaviour based on GPS trajectories. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 203–204. IEEE, 2015. [2, 4](#)
- [26] G. C. Mariano, V. Adnani, I. Kewalramani, B. Wang, M. J. Roorda, J. Bowes, and S. Diamond. Designing a dashboard visualization tool for urban planners to assess the completeness of streets. In *Human Interface and the Management of Information. Designing Information: Thematic Area, HIMI 2020*, pp. 85–103. Springer, Cham, 2020. [1](#)
- [27] F. Miranda, T. Ortner, G. Moreira, M. Hosseini, M. Vuckovic, F. Biljecki, C. T. Silva, M. Lage, and N. Ferreira. The state of the art in visual analytics for 3D urban data. *Computer Graphics Forum*, 43(3):e15112, 2024. [2](#)
- [28] G. Moreira, M. Hosseini, M. N. A. Nipu, M. Lage, N. Ferreira, and F. Miranda. The Urban Toolkit: A grammar-based framework for urban visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):1402–1412, 2023. [1, 3, 5, 9](#)
- [29] G. Moreira, M. Hosseini, C. Veiga, L. Alexandre, N. Colaninno, D. de Oliveira, N. Ferreira, M. Lage, and F. Miranda. Curio: A dataflow-based framework for collaborative urban visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):1224–1234, 2025. [9](#)
- [30] R. Mota, N. Ferreira, J. D. Silva, M. Horga, M. Lage, L. Ceferino, U. Alim, E. Sharlin, and F. Miranda. A comparison of spatiotemporal visualizations for 3D urban analytics. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1277–1287, 2023. [9](#)
- [31] M. Nöllenburg. *Geographic Visualization*, pp. 257–294. Springer, Berlin, Heidelberg, 2007. [1](#)
- [32] A. Peters, D. Hernández, M.-A. Kioumourtzoglou, M. A. Johnson, S. N. Chillrud, and M. Hilpert. Assessing neighborhood-scale traffic from crowd-sensed traffic data: Findings from an environmental justice community in New York City. *Environmental Science & Policy*, 133:155–163, 2022. [1](#)
- [33] M. Pi, H. Yeon, H. Son, and Y. Jang. Visual cause analytics for traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2186–2201, 2019. [2, 4](#)
- [34] E. Prelovskaya and A. Levashev. Modern approach of street space design. *Transportation Research Procedia*, 20:523–528, 2017. [2](#)
- [35] M. S. Rahman, M. Meenar, S. Labib, T. Howell, D. Adlakha, and B. Woodward. Unveiling environmental justice in two US cities through greenspace accessibility and visible greenness exposure. *Urban Forestry & Urban Greening*, 101:128493, 2024. [1](#)
- [36] D. Ren, B. Lee, and M. Brehmer. Chartulator: Interactive construction of bespoke chart layouts. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):789–799, 2019. [4](#)
- [37] M. Saha, M. Saugstad, H. T. Maddali, A. Zeng, R. Holland, S. Bower, A. Dash, S. Chen, A. Li, K. Hara, and J. Froehlich. Project Sidewalk: A web-based crowdsourcing tool for collecting sidewalk accessibility data at scale. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, 14 pages, p. 1–14. ACM, 2019. [6](#)
- [38] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-Lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):341–350, 2017. [3](#)
- [39] A. Satyanarayan, R. Russell, J. Hoffswell, and J. Heer. Reactive Vega: A streaming dataflow architecture for declarative interactive visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):659–668, 2016. [3](#)
- [40] S. Schöttler, Y. Yang, H. Pfister, and B. Bach. Visualizing and interacting with geospatial networks: A survey and design space. *Computer Graphics Forum*, 40(6):5–33, 2021. [2, 4](#)
- [41] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage,

- N. Amin, B. Schwikowski, and T. Ideker. Cytoscape: A software environment for integrated models of biomolecular interaction networks. *Genome Research*, 13(11):2498–2504, 2003. [8](#)
- [42] A. Steed, S. Spinello, B. Croxford, and C. Greenhalgh. e-Science in the streets: Urban pollution monitoring. In *Proceedings of the UK e-Science All Hands Meeting*. Citeseer, 2003. [1](#)
- [43] G. Sun, R. Liang, H. Qu, and Y. Wu. Embedding spatio-temporal information into maps by route-zooming. *IEEE Transactions on Visualization and Computer Graphics*, 23(5):1506–1519, 2016. [2](#), [3](#), [4](#)
- [44] G. Sun, Y. Liu, W. Wu, R. Liang, and H. Qu. Embedding temporal display into maps for occlusion-free visualization of spatio-temporal data. In *IEEE Pacific Visualization Symposium*, pp. 185–192. IEEE, 2014. [2](#), [3](#), [4](#)
- [45] M. Tajgardoon and H. A. Karimi. Simulating and visualizing sidewalk accessibility for wayfinding of people with disabilities. *International Journal of Cartography*, 1(1):79–93, 2015. [1](#), [2](#), [4](#)
- [46] F. Wang, W. Chen, F. Wu, Y. Zhao, H. Hong, T. Gu, L. Wang, R. Liang, and H. Bao. A visual reasoning approach for data-driven transport assessment on urban roads. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 103–112. IEEE, 2014. [2](#)
- [47] Z. Wang, M. Lu, X. Yuan, J. Zhang, and H. Van De Wetering. Visual traffic jam analysis based on trajectory data. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2159–2168, 2013. [2](#), [4](#)
- [48] Z. Wang, T. Ye, M. Lu, X. Yuan, H. Qu, J. Yuan, and Q. Wu. Visual exploration of sparse traffic trajectory data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1813–1822, 2014. [2](#), [3](#), [4](#)
- [49] D. Weng, C. Zheng, Z. Deng, M. Ma, J. Bao, Y. Zheng, M. Xu, and Y. Wu. Towards better bus networks: A visual analytics approach. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):817–827, 2020. [2](#), [4](#)
- [50] W. F. Wiecek and A. M. Delmerico. Geographic information systems. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(2):167–186, 2009. [2](#)
- [51] L. Wilson, J. Danforth, D. Harvey, and N. LiCalzi. Quantifying the urban experience: Establishing criteria for performance based zoning. In *Proceedings of the Symposium on Simulation for Architecture and Urban Design*, pp. 1–8, 2018. [2](#), [4](#)
- [52] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pp. 2648–2659. ACM, 2017. [3](#)
- [53] J.-H. Woo, S. B. Lim, K. Jeong, H. Kim, J.-J. Kim, J. Lee, J. Kim, T. Lee, R. Ryoo, S. Kim, et al. Aircscope: A micro-scale urban air quality management system. In *IEEE Fourth International Conference on eScience*, pp. 378–379. IEEE, 2008. [1](#)
- [54] S. Yang, A. Chong, P. Liu, and F. Biljecki. Thermal comfort in sight: Thermal affordance and its visual assessment for sustainable streetscape design. *Building and Environment*, 271:112569, 2025. [1](#)
- [55] W. Yap, P. Janssen, and F. Biljecki. Free and open source urbanism: Software for urban planning practice. *Computers, Environment and Urban Systems*, 96:101825, 2022. [1](#)
- [56] W. Zeng, C.-W. Fu, S. M. Arisona, A. Erath, and H. Qu. Visualizing mobility of public transportation system. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1833–1842, 2014. [2](#), [4](#)
- [57] W. Zeng, C.-W. Fu, S. Müller Arisona, A. Erath, and H. Qu. Visualizing waypoints-constrained origin-destination patterns for massive transportation data. *Computer Graphics Forum*, 35(8):95–107, 2016. [2](#)
- [58] X. Zhang, Y. Xu, K. Li, L. Yu, and Y. Liu. ChinaVis24.MapCraft: Dissecting and designing custom geo-infographics. *Journal of Visualization*, 28(4):837–857, 2025. [2](#), [4](#)
- [59] P. Ziegler and S. E. Chasins. A need-finding study with users of geospatial data. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pp. 1–16. ACM, 2023. [8](#)