TopoText: Context-Preserving Text Data Exploration Across Multiple Spatial Scales

Jiawei Zhang,[1] Chittayong Surakitbanharn,[1] Niklas Elmqvist,[2]
Ross Maciejewski,[3] Zhenyu Qian,[1] and David S. Ebert[1]

[1]Purdue University
{zhan1486|csurakit|qianz|ebertd}@purdue.edu

[2]University of Maryland
elm@umd.edu

[3]Arizona State University
rmacieje@asu.edu

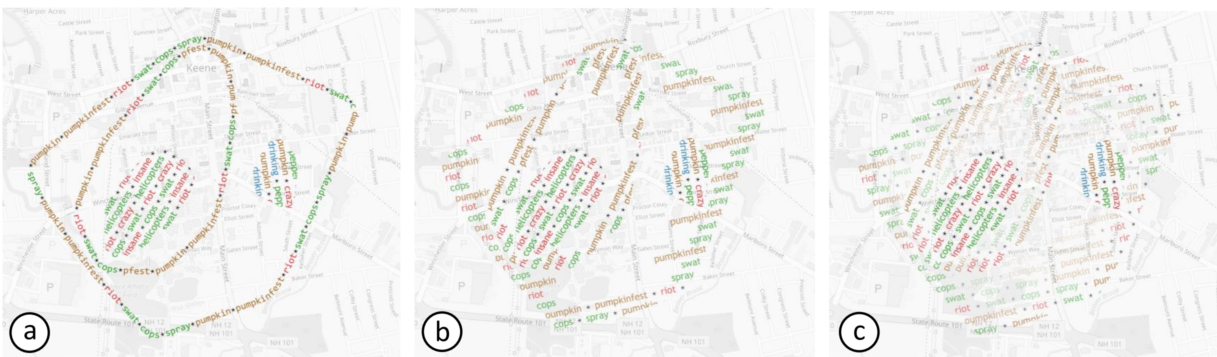

Figure 1. TopoText showing the prominent topics (encoded by color) at different spatial scales on social media around the city of Keene in the state of New Hampshire, during the Pumpkin Festival riots in 2014. TopoText creates novel text-based visualizations to couple the multi-level textual information in the same visual display for context preservation. (a): The multi-scale boundary-dominant visualization; (b): The multi-scale boundary-space hybrid visualization; (c): The multi-scale space-dominant visualization.

ABSTRACT

TopoText is a context-preserving technique for visualizing text data for multi-scale spatial aggregates to gain insight into spatial phenomena. Conventional exploration requires users to navigate across multiple scales but only presents the information related to the current scale. This limitation potentially adds more steps of interaction and cognitive overload to the users. TopoText renders multi-scale aggregates into a single visual display combining novel text-based encoding and layout methods that draw labels along the boundary or filled within the aggregates. The text itself not only summarizes the semantics at each individual scale, but also indicates the spatial coverage of the aggregates and their underlying hierarchical relationships. We validate TopoText with both a user study as well as several application examples.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

CHI 2018, April 21–26, 2018, Montréal, QC, Canada.
Copyright © 2018 ACM ISBN 978-1-4503-5620-6/18/04 ...$15.00.
http://dx.doi.org/10.1145/3173574.3173611

Author Keywords

Geospatial visualization; text visualization; typographic map; context preservation; multi-scale analysis.

INTRODUCTION

Spatial data aggregated at different scales often yield different distribution patterns. For example, the dominant criminal offense types in an entire city may be theft and robbery. However, local regions such as a university campus may be dominated by different crimes, i.e., liquor law violation and noise. Thus, there are multiple interpretations of the data depending on scale [36, 52], and acquiring an accurate picture of the city requires understanding them all. Typical approaches to support multi-scale navigation either juxtapose multiple views of each scale for easy comparison [19, 24, 31], combine the analysis results at different scales into a holistic display [26, 51, 57], or focus on optimization to improve navigation efficiency [30, 38, 56]. However, all these approaches either require additional, often significant, interaction to navigate through multiple spatial scales, or yield high visual complexity and clutter. Furthermore, while much spatial data is now textual in nature, such as geotagged social media, most existing techniques are designed for scalar or tensor data.

In this paper, we propose TopoText, a technique designed for the visual analysis of multi-scale textual information associated with spatial data, such as geo-tagged social media, crime incident reports, and demographic data. Based on a recent context-preserving visualization of multi-scale spatial aggregates [57], TopoText maximizes screen space utilization and

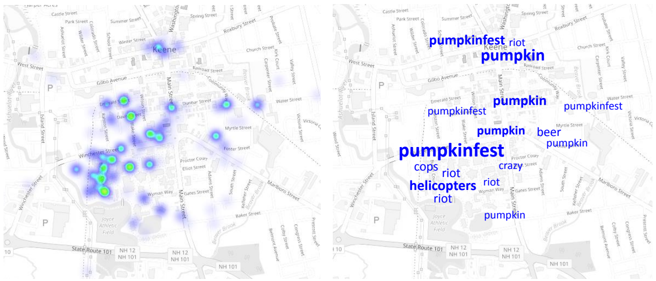

Figure 2. A heat map (left) reflects the spatial data distribution but does not support exploring the textual information. A tag map (right) depicts the major keywords at different regions at the current spatial scale, but does not indicate the variation of the text data across multiple scales. (The same data are visualized by TopoText and shown in Figure 1.)

minimizes visual complexity by using the shape of each geographic aggregate (cluster) as the major visualization element. Specifically, TopoText eschews graphical lines or filled areas to convey these shapes, and instead renders them using the textual labels themselves as graphical features. This is done by either drawing the labels along the boundary or by filling the interior of the shape while appropriately adjusting perceptual channels including translucency, color, orientation, etc. In this way, the text itself indicates both the textual information as well as the spatial extents it was aggregated from (Figure 1).

We have conducted two user studies, one on semantic contents and one on the spatial hierarchy, to evaluate the efficacy of different designs alternatives in TopoText. Each study asked participants to perform two tasks using a specific variant of the technique. Results suggest appropriate design choices for TopoText to enable users to more effectively and efficiently navigate in a multi-scale space. Not surprisingly, our findings indicate that there is hardly a one-size-fits-all solution and the optimal design to choose depends on the problem, task, and user requirements. We also present practical applications for TopoText through two examples on social media analysis to demonstrate the efficacy of the technique in both a static authoring process and an interactive exploratory environment.

BACKGROUND

Here we review prior work related to visual analysis of multiple spatial scales, text exploration of spatial data, and text-based visualization techniques with spatial constraints.

Visual Analysis of Multiple Spatial Scales

Conventional multi-scale exploration requires heavy interactions and can increase cognitive overload [6,30,32,36]. Visual analytics research has explored juxtaposing the visual results at different scales for direct and quick comparison [19,24,31], or optimizing navigation operations such as zooming and panning [30,38,56]. Although these methods support intuitive exploration across scales, they require the users to navigate to the specific scales of interest, thus easily causing interaction overload. To this end, several techniques have been proposed to reduce interactions and maintain the context by combining multi-scale results in the same visual display. Turkay et al. [51] utilize a single chart to summarize the multi-scale statistical results in a static visualization. Goodwin et al. [26] propose a compact glyph design called Scale Mosaic, which consists a set of concentric rings to encode statistical correlations from the global to the local scales. Zhang et al. [57] adopt this idea for a technique called TopoGroups that visualizes the spatial clusters at multiple scales by distorting their boundaries. In large-display visualization, a similar idea is applied that blends multiple visual representations into the same visual space to accommodate users of different viewing distances [29].

As an extension of TopoGroups [57], TopoText adopts a similar idea by creating a visual summarization of multi-scale spatial clusters in a single display. However, unlike TopoGroups, which visualizes the hierarchical and statistical information related to individual aggregates, TopoText focuses on the visual analysis of the textual information, such as terms, phrases, and topics associated with the spatial data.

Text Exploration of Spatial Data

Typical approaches to visualizing textual information extracted from spatial data visualize them in a view that is physically separate from, but linked to, the geographical space [20,23,56]. However, they usually require the users to switch between multiple views and perform additional interactions in order to correlate the spatial and textual dimensions, potentially adding to the cognitive load of the user. Research has explored combining text within the geographical space in order to reduce the overload. One common technique is Tag Maps [5,7,12,49,54], a variant of tag clouds that appropriately positions the words on a map to indicate their geographical distribution and prominence. Other work also utilizes the spatial dimension for visualization, where the position of the textual features do not necessarily represent their geographic locations. For example, Nguyen et al. [46,47] sort words based on the user-defined order and position the text on the map along the vertical skeleton of the geographical boundary. Brath and Banissi [9,10] extend common set visualization techniques [3] to coupling textual attributes.

Unlike existing research that only focus on a single scale in one display, TopoText adopts the geographic space to show the multi-scale textual information in the same visual space, aiming to reduce the interaction and cognitive overload that exists in the previous work (Figure 2).

Text Visualization with Spatial Constraints

Text-based design space involves a rich set of the visual attributes. Among them, position is probably the most critical aspect to consider as it potentially indicates the latent relationships among different text entities and can reflect other information dimensions when properly encoded. When positioning text, Spatial constraints commonly exist in various text-based visualizations. The most well-known technique, tag clouds [43], and its descendants [13,17,18,33,37,39,54], typically generate a compact and occlusion-free word layout in which the feasible position of the individual words are constrained by the existing words in the visual space. Other spatial constraints are defined based on the additional information dimensions associated with text, such as the geometric elements in either 2D or 3D space, where text labels provide supporting

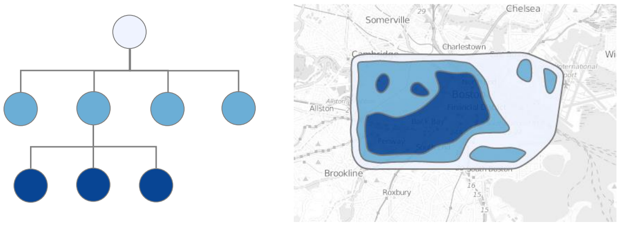

Figure 3. The occlusion-free and context-preserving visualization generated in the TopoGroups technique [57]. Left: The multi-scale aggregate hierarchy; Right: The corresponding geospatial representation.

information. Wong et al. [53] combine text and visual elements (e.g., nodes and edges) in a graph in order to recycle the space resource and avoid visual clutter among multiple elements. Maharik et al. [41] propose digital micrograms that creates calligrams (text arranged to form a shape that illustrates its semantic meaning, which has been crafted by artists and poets even before the emergence of computer graphics) by calculating the vector fields for the graphical elements in the image in order to guide the text layout. Xu and Kaplan [55] introduce Calligraphic Packing, a technique that divides an image into segments and warps and fills letters into each segment. Afzal et al. [1] automate typographic maps [42], in which the text layout is constrained by the underlying geographical elements. Similarly, Godwin et al. [25] apply the typographic map to visualizing semantic topics extracted from social media [25]. Moreover, the spatial constraints commonly exist in various map design applications, where the label placement is carefully executed in order to indicate the feature locations and avoid potential ambiguity or contradiction [11, 28, 40, 50].

The spatial constraint in TopoText is the boundaries generated from the multi-scale clusters. Inspired by the typographic map technique [1], TopoText explores different design alternatives that embed text within the individual shapes to fully utilize the visual resources and maintain the semantic context across scales, and further evaluates the effectiveness of these design choices depending on the problems and tasks.

TOPOTEXT: MULTI-SCALE TEXT DATA EXPLORATION

The TopoText technique is designed to support effective visualization and interactive exploration of textual data aggregated at multiple scales. In this section, we first motivate our work from the TopoGroups technique [57]. Then we formulate the design goals for multi-scale text visualization, followed by the detailed approach of the TopoText technique.

Motivation

As Figure 3 shows, TopoGroups [57] is a technique for preserving context across multi-scale spatial aggregation. TopoGroups performs hierarchical clustering on the spatial data points [22], combines multi-scale clusters in the same visual display and distorts their boundaries to avoid visual clutter. This occlusion-free representation guarantees proper visual space between the adjacent boundaries to potentially encode different information dimensions. TopoGroups consists of visual encoding methods to indicate the statistical or categorical information associated with individual aggregates, enabling the users to compare and correlate them within a multi-scale space. Nevertheless, exploring textual information aggregated at multiple scales in TopoGroups and other visual analytics techniques (e.g., tagmap in Figure 2) is inefficient because the displayed text changes at different spatial scales, requiring the users to switch between scales and mentally remember the multi-scale results.

TopoText extends TopoGroups to tackle the challenges of visualizing text at multiple scales. Inspired by the typographic maps [1, 42], TopoText utilizes the occlusion-free property and employs textual labels as its primary visualization entity to reduce visual complexity. Although the design space of the text-based visualization is broad and consists of multiple perceptual channels (color, size, density, position, shading, etc.), employing too many attributes may easily increase visual complexity and overload readers. Thus, we tailored a hierarchical aggregation and visualization model [22] to develop design goals for a multi-scale text exploration technique. Then we identified a subset of perceptual channels for text rendering that meet the design goals (consideration space), proposed appropriate design choices and rejected bad ones at the design time (proposal space). Finally, we evaluated the efficacy of the design candidates in a user study setup (selected solution) [44].

Design Goals

Our primary design principle is consistent with the concept of the geographical mashup [54] by regarding textual information as a secondary information dimension overlaid on the primary geographical dimension. Doing so visually indicates the correlation between these two dimensions and provides contextual information of the spatial patterns. This can also reduce the overload caused by switching to separate views with textual information. Below, we detail our design goals regarding text data exploration of multi-scale spatial data. They are mainly extended from a hierarchical aggregation model for information visualization [22].

Entity Budget (G1): The entity budget for the text data is proportional to the budget for the aggregate since the text labels are visually associated with the geographic location of the corresponding aggregate. Hence, an aggregate with larger/smaller spatial coverage (not necessarily a larger/smaller number of data points) has a correspondingly larger/smaller visualization budget for the text. Aggregates that are too tiny (e.g., occupy less than 5∗5 pixels in the screen space) or outside the viewport are not visualized.

Visual Summary (G2): The accuracy of the textual information presented at a single scale should be compromised or at least not prioritized. Hence, for each scale we should show a coarse-grained summary instead of details. One should not expect that a visualization design shows the entire multi-scale hierarchy while being able to depict the fine-grained information at each individual scale (e.g., tag map).

Visual Simplicity (G3): The amount of the textual information presented at a single scale should be limited. In order words, the representation should be simple and clean in order to avoid generating visual complexity. G2 and G3 constraint the textual information at a single scale in order to express the textual information at multiple scales succinctly

Figure 4. Design alternatives for visualizing the text data on a single aggregate. (a): The text labels are placed along the boundary; (b): The text labels are filled within the area of the aggregate; (c): The space-filling visualization is enhanced by applying a transparency gradient on the text labels; (d): The text labels that are close to the boundary are placed inside the aggregate.

and avoid overloading the users. These two principles are especially critical within the domain of the text-based visualization since the design space of text is complicated and can easily involve design choices that confuse or overload readers. Hence, a reasonable design should identify a small and optimal set of orthogonal visual channels and establish a reasonable mapping between them and the data dimensions that need to be conveyed.

Discriminability (G4): This refers to the capability of visually distinguishing between the aggregate and data items. The data items (e.g., the geospatial data points) are usually represented as simple dots or more complex glyphs on the geographical map. Therefore, the visual entity of the aggregate—the text label—is easily distinguishable from the data item and does not require additional decoration as suggested in the model [22] to facilitate visual discrimination.

Fidelity (G5): The fidelity issue is often involved in visualizing aggregates. Since only a summary of the entire textual features is visualized (G2), the readers may have a biased interpretation on the textual information associated with the aggregate. Furthermore, this issue also exists in text visualization due to inappropriate encoding choices. For example, when the font size is fixed, a longer length keyword typically occupies more screen space, which can visually mislead its importance value and introduce perceptual bias [2]. Hence, text encodings should be carefully executed in order to prevent visual confusion. In essence, a trade-off between these potentially contradictory principles needs to be considered.

Interpretability (G6): Inappropriate text visualization methods may also hamper its interpretability. For example, a radial layout of a set of text should make necessary adjustment to avoid rendering the text upside-down [16, 56]; Rendering a word sequence along a curve with sharp angles may result in distorted letters and inconsistent spacing between them [1, 41, 53]. A reasonable design should avoid such flaws related to the low-level visual attributes.

Visualizing the Text Data: A Single Aggregate

An aggregate is the basic element in the multi-scale aggregation hierarchy, which can be visually represented as a node in the dendrogram representation (Figure 3). TopoText creates four primary design alternatives for showing text of a single aggregate as described below (Figure 4). These approaches solely rely on appropriate text encoding and layout to indicate both the textual information and the geographic characteristic of the aggregate.

S-bd **S**ingle-scale **b**oun**d**ary-based visualization: The text labels are placed along the boundary (Figure 4(a)). Particularly, TopoText identifies sharp angles along the boundary and divides it into segments accordingly. This ensures that the segments have low curvature without any sharp change in direction. The text labels are then placed within the segment with potential distortion avoided [41] (G6).

S-sp **S**ingle-scale **sp**ace-filling visualization: The text labels are filled within the area of the aggregate based on the sweep line approach in order to fully utilize the inner space (Figure 4(b)). The text labels are clipped based on the boundary to visually indicate the shape of the aggregate. The direction of the text layout is determined by the direction of the diameter (the longest axis) of the polygon instead of a fixed direction (e.g., a horizontal layout) to avoid generating short and fragmented text lines that are hard to interpret [41] (G6). The vertical and horizontal spacing between adjacent text labels within one aggregate is set as a constant value in order to provide a simple and clean visual effect.

S-tsp **S**ingle-scale **t**ranslucent **sp**ace-filling visualization: When the aggregate occupies a relatively large space on the screen, directly applying the space-filling method (S-sp) can result in a large number of the text labels visible, potentially overloading the users. To this end, we apply a transparency gradient to the visualization such that the labels close to the boundary have a higher opacity value while those close to the aggregate's center have a lower opacity value (Figure 4(c)). A cubic function is used in the transparency gradient to enhance the visual perception of the boundary.

S-bh **S**ingle-scale **b**oundary-space **h**ybrid visualization: The text layout strategy in this design is similar to the space-filling visualization (S-sp), except that only the text labels that are close to the boundary are visualized to visually indicate the boundary shape (Figure 4(d)). The distance measure is based on the Euclidean distance between the center point of the text labels and the edge of the polygon that is the closest to the center point.

We note that in the four proposed design choices the position of the text is determined based on the available space resource of the aggregate and does not reflect the spatial distribution of the keywords within the aggregate (G2 and G3). Furthermore, the text labels in an aggregate have a fixed font size. The rationale behind these designs is that we aim to provide a visual semantic summary that doesn't cause potential information overload by employing too many visual channels (G2). The color of the text can be used to encode information such as topics, sentiment, etc., and TopoText allows the users to change the setting interactively.

Visualizing the Text Data: Multi-Scale Aggregates

As the multi-scale aggregates introduce more complexity to the visualization space, an effective visual representation should be free of visual occlusion and constrain the number of the visual elements presented to the user. We enumerates a set of potential design candidates by extending the single-scale design choices (Section 3.3) to the multi-scale aggregates. Then we identify the potential limitations in each design, perform appropriate refinement and propose the satisfying solutions that are listed below [44]. The visualization results of these solutions are shown in Figure 1.

M-bd **M**ulti-scale **b**oundary-**d**ominant visualization: The boundary-based technique (S-bd) is applied to the multi-scale aggregates that are not at the lowest aggregation level (Figure 1(a)). Since the lowest-level aggregates do not have children in their inner space, the space-filling visualization (S-sp) is applied to them in order to improve the space resource utilization (Figure 5(a)). This approach generates an occlusion-free visual result by taking advantage of the proper spacing between the boundaries [57].

M-bh **M**ulti-scale **b**oundary-space **h**ybrid visualization: The hybrid visualization (S-bh) is applied to the multi-scale aggregates that are not at the lowest aggregation level (Figure 1(b)). To avoid visual clutter generated by the overlapping aggregates (typically the aggregates that have a parent-children relationship), the child aggregate is visualized on top of the parent aggregate such that the parent's area that is covered by the child is invisible to the user. Because the direction of text is dependent on the diameter of the aggregate, this variation in the direction makes it easier for users to distinguish the adjacent aggregates [1]. Furthermore, the text labels at a higher (abstract) level are more transparent and sparse while those at a lower (detail) level are more opaque and dense [45]. Similar to M-bd, the space-filling technique is applied to the lowest-level aggregates.

M-sp **M**ulti-scale **sp**ace-dominant visualization: The translucent space-filling visualization (S-tsp) is applied to the multi-scale aggregates that are not at the lowest aggregation level (Figure 1(c)). Simply applying the space-filling technique (S-sp) may produce significant information overload (Figure 5(b)). Similar to M-bh, the opacity and density of the text increases from the higher-level aggregates to the lower-level ones. Similarly, the space-filling technique is applied to the lowest-level aggregates.

We have conducted a user study to evaluate the efficacy of the aforementioned design choices in conveying the textual information of the multi-scale aggregates while retaining the

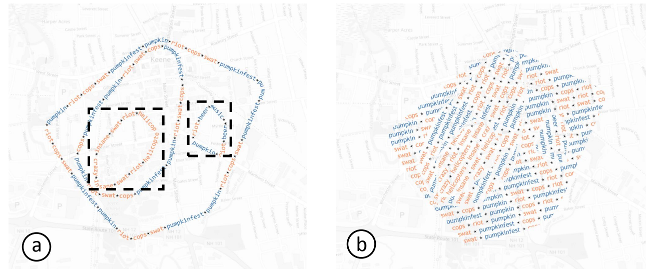

Figure 5. (a): Applying the boundary-based visualization (S-bd) to the multi-scale aggregates. The space utilization of this design can be improved by filling the text labels in the lowest-level aggregates (shown in the black rectangles). (b): Applying the space-filling visualization (S-sp) to the multi-scale aggregates. Since the number of the text labels in the visualization can potentially be large, this design may add significant visual overload to the user.

geographical and hierarchical relationships of these aggregates. The results are reported in the evaluation section.

We also note that additional visual attributes besides the textual features can be integrated to encode different information dimensions. For example, the background color of an aggregate can be used to encode the data density or the aggregation level [57] (Figure 7). A blue-red scheme is applied in TopoText by default. But more color schemes are supported to account for personal preferences and accommodate color blindness. When the aggregate's background is rendered, TopoText chooses a color scheme that has high contrast with the text color for the purpose of better readability. TopoText also applies the halo effect on the boundary of the aggregate in order to produce a visual effect that the child aggregates stack on top of their parents, thus enhancing the perception of the aggregates' hierarchy [57] (Figure 7 and Figure 8). The users can toggle the halo on or off in the interface of TopoText.

INTERACTION AND INTERFACE IN TopoText

The interface of TopoText mainly consists of a geographic map view that visualizes the multi-scale text data (Figure 7(b)) and a tree view that overviews the multi-scale hierarchy (Figure 7(a)). The map view visualizes the aggregates that intersect with the current viewport and occupy a reasonable amount of screen space, e.g., more than 100 pixels (G1). As the user navigates to different regions and scales on the map, the nodes (aggregates) that are visible in the viewport are highlighted in the tree view accordingly. When a node of interest in the tree view is selected, the map smoothly zooms and pans to center the corresponding aggregate in the viewport. The two coordinated views enable the users to navigate to different scales and details on the map while being able to maintain the context of the entire analysis space.

When the text-based techniques are applied to the aggregates that have a limited visual budget (e.g., the aggregate occupies a relatively small region), the text labels may be partially visible to the users and thus hamper information fidelity (G5) or interpretability (G6). In these cases, TopoText utilizes a set of boundary-based encoding strategies from TopoGroups [57] that typically visualize a sequence of colored segments or colored dashes on the boundary to summarize relevant information, such as the volume of the messages corresponding to the different topics (the aggregate B in Figure 7(b)).

Given a limited spatial visualization budget for an aggregate, TopoText provides common methods to determine the top K representative keywords to visualize, which includes term and inverse document frequency (TF-IDF), latent Dirichlet allocation (LDA), and lexicon-based matching, and supports the users to toggle between different options and adjust the value of K. Furthermore, as the user hovers over a specific textual feature in the aggregate or searches for a keyword in the control panel, the aggregates that contain the same feature highlight accordingly (Figure 7(e)).

As TopoText visualizes a summary of the textual information, a detail-on-demand interaction design is supported to enable quick access to detailed information that is not presented in the current visualization. When the user specifies an aggregate, the child aggregates inside it fade out and the space-filling technique (S-sp) is applied to the aggregate for the purpose of fully utilizing the inner space to present textual features. Moreover, when the user performs a scrolling operation on the aggregate, the textual labels dynamically move up or down depending on the scrolling direction, thus presenting the previously invisible text to the user [47].

IMPLEMENTATION DETAILS

TopoText is implemented based on a two-layered SVG canvas using the *D3* toolkit [8]. A map layer (OpenStreetMap) provides a gray-scale geographic context at the bottom of the canvas. The visualization layer stays on top of the map and renders text labels, aggregate boundaries and halos.

To position text labels along the boundary (Figure 4(a)), TopoText divides aggregate boundaries into segments of low curvature and renders the text using the `<textPath>` element. When the labels are visualized on the path iteratively, the `<startOffset>` attribute is used to define the position of the label and updated accordingly that guides the layout of the label to be rendered next. To fill the text inside an aggregate (Figure 4(b)), TopoText identifies the diameter of the polygon and calculates the bounding box in parallel with the diameter. TopoText then positions the text labels inside the bounding box using the `<transform>` attribute such that the orientation of text is in parallel with the diameter. An `<clipPath>` element is initialized based on the aggregate boundary and forces the rendering to be masked against the boundary.

TopoText implements the transparency gradient (Figure 4(c)) using the `<linearGradient>` element. The gradient vector is calculated based on the relative position of the text labels and the center of the polygon and is specified using the `<x>` and `<y>` attributes associated with the `<linearGradient>`. The `<stop>` element and its `<offset>` attribute are used to define the ramp of the opacity value to use on a gradient. As only a linear gradient is supported in the SVG, we sample multiple points along the gradient vector to approximate a higher-order gradient such as a quadratic or cubic function. We found 5 points to produce a visually appealing effect given the fact that the text labels have relatively short lengths.

EVALUATION

To evaluate TopoText, we focus on the two major aspects that are typically involved in the multi-scale analysis and text analysis tasks. (1) How effective does the technique express the textual information related to the multi-scale aggregates? (2) How effective does the technique reveal the geographic characteristics of the multi-scale aggregates and their relationships in the hierarchy?

Participants, Apparatus and Procedure

16 participants (4 female, 12 male, age range of 24 to 30) were recruited in the first study, and 14 participants (7 female, 7 male, age range of 22 to 64) were recruited in the second study. Most of the participants were students and staff from an engineering college and had some basic understanding of geographic applications, data clustering and data visualization. The entire study lasted around 30 minutes and each participant was paid $5 for participation in one study. We used a Dell monitor with a 1920×1080 resolution to present the system interface and the task description. The major visualization occupies an area of 1024×1024 within the screen space.

The procedures for the two studies were similar but were conducted independently. The investigator introduced the participant to the research background as well as the visualization techniques that were being tested. Then a training session was conducted to allow the participants to get familiar with the designs and the tasks. Special characters or symbols that appeared in the text-based visualization were also explained at this stage to avoid causing potential confusion to the participants (e.g., the hash ("#") or the AT ("@") symbol in a tweet). For the study that tests the textual information, the investigator also presented the participants a list of keywords that would be shown in the tasks, which familiarized the participants with the text content. The participants were asked to raise any questions during the training session. The main study included a set of multiple-choice questions which were answered afterwards. The accuracy and the completion time were recorded for each trial. The participants ended the study by finishing all the trials and filling in a post-experiment survey.

Techniques and Task Design

The techniques being evaluated in the two studies included the three multi-scale techniques: the boundary-dominant visualization (M-bd), the boundary-space hybrid visualization (M-bh), and the space-dominant visualization (M-sp). In order to focus on typography-based design choices (i.e., involving only text in the visualization and varying the visual attributes associated with text labels to generate design alternatives) and reduce the complexity of the evaluation process, the user studies did not involve additional visual channels such as the background color of the aggregate or the halo effect along the boundary. For the same reason, we did not design and involve a baseline technique (i.e., the TopoGroups [57] technique combined with a word cloud visualization to show semantic content) for comparison. We note that these are potential limitations of the evaluation and we leave them as future work.

Inspired by previous research on representative analytical tasks regarding geospatial exploratory analysis [4, 35, 48], we in-

Table 1. The analytical task design in the user studies.

User Study	Highlighted entity in the visualization	Analytical task	Task taxonomy
Study 1: Semantics (text)	Multiple aggregates $\{A,B,C\}$	Identify the one in $\{A,B,C\}$ that contains a target keyword	Locate, search
	One aggregate A and multiple aggregates $\{X,Y,Z\}$	Identify the one in $\{X,Y,Z\}$ that has one or more keywords in common with A	Compare, correlate
Study 2: Hierarchy	Multiple aggregates $\{A,B,C\}$	Identify the one in $\{A,B,C\}$ that is at a higher (lower) aggregation level	Rank, compare
	One aggregate A and multiple aggregates $\{X,Y,Z\}$	Identify the one in $\{X,Y,Z\}$ that is a child of A	Locate, compare

volved four types of tasks in the two studies (Table 1). The first two tasks evaluated the capability of the technique to convey textual information (Study 1). The last two investigated the effectiveness of the technique to convey geographic and hierarchical relationships among the aggregates (Study 2). For each trial, a static image was shown to the participant, in which one of the techniques being tested was applied to visualize the textual information of the multi-scale aggregates (similar to Figure 1). As Table 1 shows, specific aggregates related to the task were highlighted in the image using a black arrow and an upper case letter, such as X, Y, Z. The participants were asked to read the image, perform the corresponding task, and choose an answer from a list of options. For example, in the first task, three aggregates labeled as A, B and C are highlighted to the participant, and they are asked to find the aggregate that contains a target keyword W in the image. As a control variable, the color of the text labels in the image remains a constant value. We use synthetic data in the two studies.

For Study 1, we controlled the difficulty level D of each trial based on the complexity of the textual information, which can be quantified based on the number of distinct keywords for each aggregate shown in the visualization (we use 2, 4 and 8 in the study). The three techniques were presented in a counter-balanced order to prevent potential bias. The entire study consists of 3 (technique) $\times$ 2 (task type) $\times$ 3 (difficulty level) $\times$ 2 (repetition) = 36 trials. For Study 2, we controlled the difficulty level D of each trial based on the complexity of the hierarchy, which is quantified based on the number of scales in the hierarchy, or the depth of the hierarchy (we use 2, 4 and 6 in the study) Similarly, the three techniques were presented in a counter-balanced order to prevent potential bias. The entire study consists of 3 (technique) $\times$ 2 (task type) $\times$ 3 (difficulty level) $\times$ 2 (repetition) = 36 trials.

Study 1: Results and Observations

The accuracy across the three techniques ranges from 90.6% to 94.1% (92.5% on average). This is because the visualization provides the necessary information—the keywords to search among aggregates—for the participants to identify the correct answer and there was no time limit for the tasks.

The participants spent the most time on the boundary-dominant technique (M-bd) (24.38 seconds), followed by the hybrid technique (M-bh) (16.83 seconds) and the space-dominant technique (M-sp) (14.72 seconds) (Figure 6(a)). Visualization technique V had a significant main effect on completion time ($F(2, 26) = 17.14$, $p < .0001$). Pairwise comparison between visualization techniques using a Tukey HSD showed that the pairs (M-bd, M-sp) and (M-bd, M-bh) have statistical significance ($p < .0001$). Difficulty level D also had a significant main effect on completion time as well ($F(2, 26) = 10.59$, $p < .0001$). These results indicate that the boundary-dominant technique (M-bd) was inefficient in conveying the textual information. This can be explained by the fact that placing text on the boundary may potentially distort the letters and hamper readability. In order to read text along the boundary, the participants had to visually cover a larger distance in the screen space, thus requiring a longer time. In contrast, the space-dominant technique (M-sp) and the hybrid technique (M-bh) rendered text in a fixed direction without distortion, enabling an easier visual perception. Moreover, as the space-dominant technique filled text entirely, the amount of information presented within the unit of screen space was maximal. This enabled the users to focus on a smaller region to search or match keywords, thus reducing the overhead to switch visual focus across distant areas on the screen.

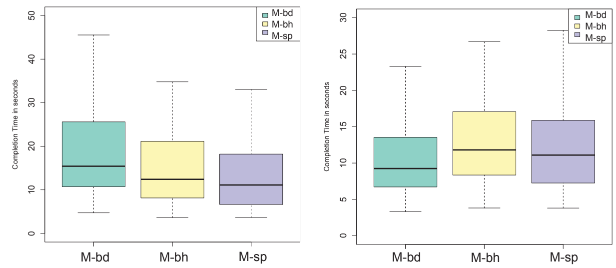

(a) Completion Time (Study 1) (b) Completion Time (Study 2)

Figure 6. Completion time for the two studies. Left: The space-dominant technique (M-sp) was the most effective for understanding the textual information visually. Right: The participants spent the least time identifying the aggregates' hierarchy based on the boundary-dominant technique (M-bd).

The subjective feedback was consistent with the analysis on the completion time. 9 participants (64%) agreed that the space-dominant technique (M-sp) was the most efficient. One participant noted that *filling text compactly helped the finding of keywords in the cluster. The transparency distinguishes the boundary of clusters.* Most participants (86%) disliked the boundary-dominant technique (M-bd). One participant mentioned that *I have to keep "relocating" my eye focus in order to read the text.* Another participant noted that *words are written in different orientation, so I had to twist my head to read words*. 3 participants (21%) preferred the hybrid technique (M-bh) over the space-dominant technique (M-sp). They seemed to have been distracted by the transparency effect: *I had to squint a lot to read the fading out effect.*

Study 2: Results and Observations

The accuracy across the three techniques ranges from 90.9% to 95.2% (93.2% on average). Similarly, the participants were able to successfully understand the hierarchical relationships within the visualization presented and the time spent on each trial was not constrained.

In terms of the completion time, the participants spent the most time on the hybrid technique (M-bh) (14.44 seconds), followed by the space-dominant technique (M-sp) (13.57 seconds), followed by the boundary-dominant technique (M-bd) (11.66 seconds) (Figure 6(b)). Visualization technique V had a significant main effect on completion time ($F(2, 30) = 5.37$, $p < .005$). Pairwise comparison between visualization techniques using a Tukey HSD showed that the pairs (M-bd, M-sp) and (M-bd, M-bh) have statistical significance ($p < .05$). Difficulty level D had a significant main effect on completion time as well ($F(2, 30) = 22.97$, $p < .0001$). The results indicate that the boundary-dominant technique (M-bd) was the most effective design for visualizing the hierarchical structure of the multi-scale aggregates. In the perspective of visual perception, this technique utilized the minimum space resource that was required to convey the aggregate hierarchy, thus reducing the cognitive overload to the readers. The space-filling-based approaches (M-bh and M-sp) were less effective, especially when a parent had too many children and the children were located near the boundary of the parent. In these cases, the visual space between the adjacent boundaries were filled with text labels and made it challenging for the readers to understand the shape of the aggregates. In the experiment, the participants spent less time on average on the space-dominant technique (M-sp) than on the hybrid technique (M-bh). One explanation may be the fact that the visual perception of the boundary was enhanced by the higher-order transparency gradient. In contrast, the hybrid technique had labels of varying sizes near the aggregate's boundary, adding potential visual confusion to the readers. However, we note that we did not find statistical significance between the two techniques.

In the post-experiment survey, all of the participants agreed that the boundary-dominant technique (M-bd) was the most effective in terms of conveying the hierarchical relationships among aggregates. One participant noted that *the boundaries of the clusters were clear and distinct and helped me identify the children easily. I had to put more efforts in the other designs.* Another participant noted that *it's clear even for a deeply nested structure*. A majority of the participants (75%) disliked the hybrid technique (M-bh). The major limitation commented by them was its inefficiency at distinguishing between the parent and children visually. One participant who disliked the hybrid design said that *the words along the boundary had different lengths and looked messy*. One participant mentioned that *the transparency change helped to better recognize the boundary shape compared to the one without it*.

Takeaways: The two user studies show that visualizing the text on the boundary (M-bd) more effectively depicts the aggregates' hierarchy while filling text inside the space (M-sp, M-bh) more effectively convey the textual information. These results essentially reflect the fact that when the visualization budget is limited, a trade-off exists between retaining an effective overview of the multi-scale hierarchy and providing detailed information related to individual aggregates. In the typical multi-scale exploration process, the boundary-dominant approach might be suitable for the initial or pilot stage that requires the analysts to obtain a coarse-grained understanding of the analysis space and identify potential exploration directions. With the analysis narrowed down to small-scale subspaces, the space-dominant approach can present more detailed information and support a fine-grained investigation. However, designing an optimal solution is challenging, and requires taking into account different perspectives such as the problem, task and user requirement.

PRACTICAL APPLICATIONS

We present two use cases to demonstrate the capability of TopoText for visualizing the textual information and maintaining the semantic context in the multi-scale aggregate space.

Keene Pumpkin Festival Riot

We analyzed the location-based social media (Twitter, 1507 tweets) generated during the 2014 riots in the city of Keene in the state of New Hampshire during its annual pumpkin festival. We started the exploration by extracting the trending topics during the event using LDA topic modeling. The top five topics related to jobs ("hiring", "job", "career", "retail"), festival ("#pumpfest", "pumpkin", "#pfest"), entertainment ("drinking", "beer", "music"), riot ("riot", "crazy", "injured") and police ("cop", "helicopter", "police"). We filtered out the job-related topic since most of the relevant posts were online advertisements, and visualized the textual information related to the other three topics in TopoText as shown in Figure 7(b). The festival-related topics were prominent at the abstract level in this region, as the majority of the keywords associated with the outward aggregate were rendered in yellow. Some keywords related to the riot and law enforcement (e.g., swat, riot, crazy) also appeared in the outward aggregate, indicating that quite a few social media users discussed about the riot. We also noticed that at the lower levels, a large aggregate was generated around the Keene State College that mainly contained riot-related (e.g., crazy, insane) and police-related (e.g., helicopter) keywords (aggregate A). This reflects the fact that the riot mainly originated from the college. In contrast, the northern (aggregate B) and eastern (aggregate C) regions had more tweets related to the festival and entertainment. Since the aggregate B occupied a relatively small screen space, instead of the text-based visualization, the yellow dashed lines [57] were rendered on the boundary of the aggregate to indicate the major topics were festival-related.

We zoomed further into the college region to explore the event in details (Figure 7(d)). Two large aggregates (bottom left) on the campus were identified that contained keywords including "#injured", "#flipped", "#dumpster", "pepper", "spray", etc., which indicated that the celebration spun out of control and the law enforcement had to use pepper spray to subdue the rioters. We hovered on a keyword (e.g, "#injured"), with the interface automatically highlighting similar keywords in other aggregates (Figure 7(e)). We then navigated and zoomed into the northern region in the city (Figure 7(c)) and identified that this region was the downtown of the city (Central Square) where there were a lot of bars and clubs. Since the riot did not spread to this region, the riot-related keywords rarely appeared.

The topic summaries provided by TopoText allow the users to not only capture what has happened (e.g., the chaos), but also understand to what spatial extent the event has spread (e.g.,

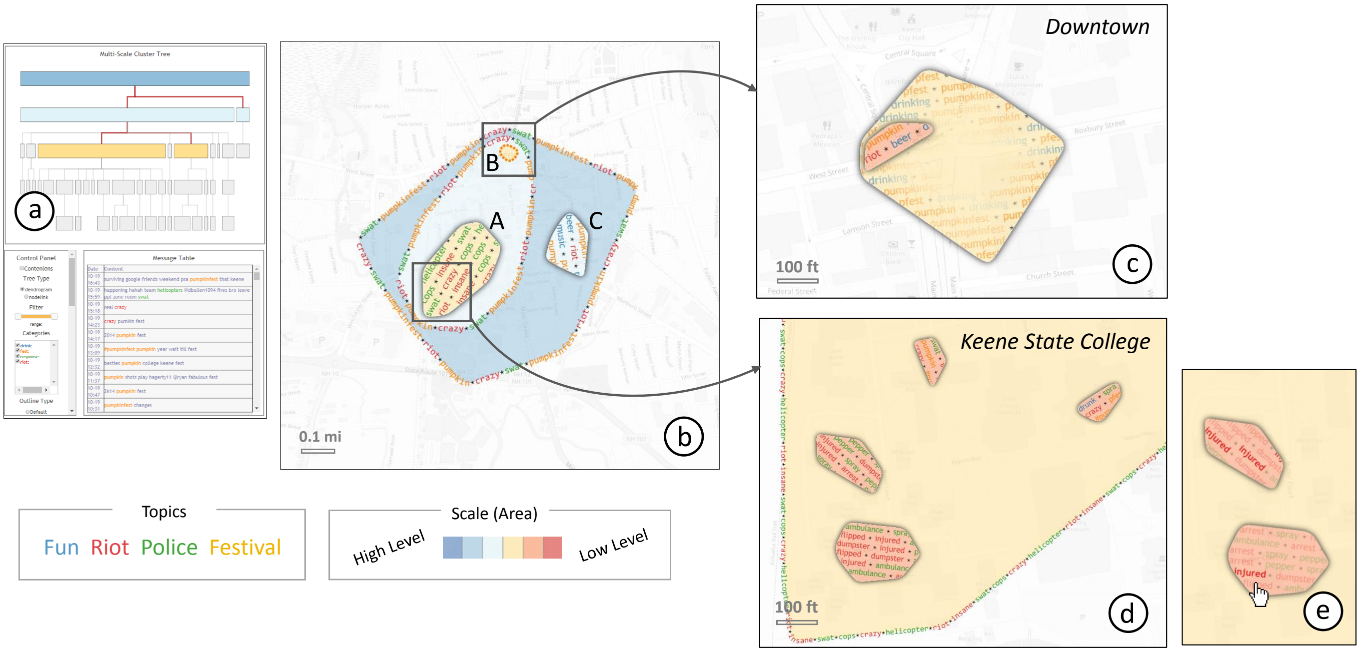

Figure 7. The interface of TopoText consists of a geographic map view (b) for visualizing the multi-scale aggregates and their textual information and a tree view that provides an overview of the multi-scale hierarchy (a). TopoText utilizes a blue-red color scheme to render the inner space of the aggregates based on their aggregation levels. TopoText also allows for text-oriented interactions: e.g., hovering on a specific keyword highlights similar keywords in other aggregates (e).

the campus area) and identify locally concentrated actionable information (e.g., "flipped", "dumpsters") that were overshadowed by more general discussions (e.g., "crazy", "insane"). By utilizing the visual outcome from TopoText, an emergency manager is able to further evaluate the scale and impact of the event and perform effective resource allocation (e.g., city police or college police); A journalist who hear a series of reports from the witnesses at the incident is able to corroborate the first-hand accounts to determine whether each story fits with the overall trends of what was happening at the time.

Republican National Convention

We investigated the social media posts (8839 tweets) collected during the 2016 Republican National Convention (RNC) in the region around the city of Cleveland, OH. Similarly, we filtered out job-related posts and identified four major topics: RNC-related ("gop", "#rnc", "convention"), traffic-related ("vehicle", "blocked", "accident"), protest-related ("#protest", "police", "#rally") and drinking-related ("drinking", "wine").

As Figure 8 shows, the region was dominated by the RNC-related topic since the outward aggregate mainly contained keywords such as "#rncincle", "trump" and "convention". As we continued to examine the lower levels, the multi-scale text-based visualization clearly revealed different topical patterns at the city level. A large aggregate around the city of Cleveland (highlighted in the figure) showed a high frequency of RNC-related and protest-related topics, potentially enhancing situational awareness for public safety personnel. In contrast, the nearby cities surrounding Cleveland contained more posts relevant to drinking and traffic. By further investigating the details associated with the individual aggregates, we found that the delegates and attendees were accommodated in the hotels in the nearby cities and suburbs and there were traffic restrictions near the convention center, causing some congestion and accidents. Therefore, the visual outcome generated by TopoText effectively preserves the semantic context and highlights the variance of spatial patterns at multiple scales.

DISCUSSION

TopoText implements a hierarchical aggregation and visualization model [22] by effectively allocating screen space to the multi-scale aggregates and visualizing the semantic summary accordingly. Unlike the original model [22] that “treats” the aggregates at different levels equally, TopoText visually highlights the ones at the lowest level in the current viewport. This is achieved by showing the lowest-level aggregates (focus) on top of other levels (context) and increasing their opacity value [14]. The rationale behind this design is that in various analytical tasks within a hierarchical space, the users are required to navigate from the top level (abstract) to the bottom level (detail). Since the visual representation typically consists of a sub-space of the hierarchy, highlighting the lowest-level aggregates in the current sub-space can visually indicate the entry to the deeper levels and effectively guide the users to navigate within the multi-scale hierarchy.

The text labels rendered within the aggregate or on its boundary may potentially be truncated to visually indicate the shape of the aggregate. We note that this truncation issue is an inherent visual output in TopoText. Essentially, this is an NP-hard packing problem [15] that aims to arrange bins of different sizes into a container in order to minimize the empty space within the container. While applying advanced layout algo-

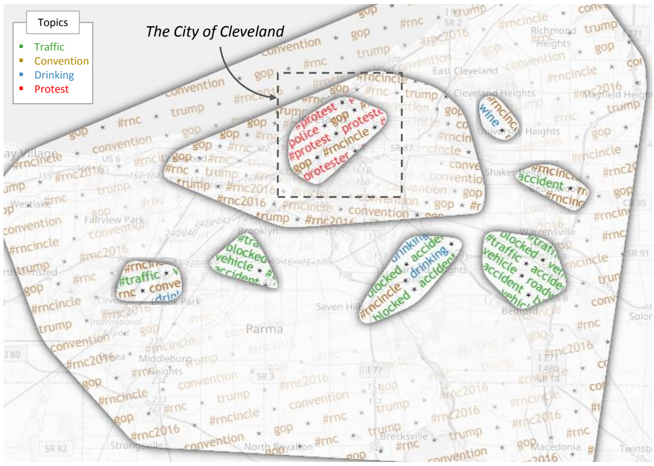

Figure 8. Applying the TopoText technique to visualizing the social media data around the city of Cleveland, OH, during the 2016 Republican National Convention (RNC). The halo effect is enabled to highlight the aggregates' hierarchy. While the region shows a high frequency of RNC-related topics, the area of Cleveland also contains topics related protest. In contrast, suburban areas have more posts relevant to traffic and drinking.

rithms may reduce the truncation issue, it is beyond the scope of this work. TopoText summarizes the semantic content of multi-scale aggregates by rendering the top K (i.e. less than 10) representative words for each one (G2 and G3). Although the number of words associated with an aggregate could potentially be large, visualizing too many distinct words within the multi-scale context can easily overwhelm the user. When the user is interested in a specific aggregate and narrows down (i.e., perform the zooming operation) to that region, the visual space for that aggregate is enlarged accordingly to accommodate more distinct keywords.

TopoText utilizes the boundary or the inner space of the aggregate for the visual budget to present the relevant textual information. Intuitively, the amount of visual budget is proportional to the area occupied by the aggregate in the geographic space. Considering that the complexity of the semantic for an aggregate is usually proportional to the volume of the data within it, instead of the occupied area, this approach may introduce a potential inefficiency of the visual space utilization. For example, the tweets posted around the college stadium during a major football game may be more complex than those posted on a common day across the campus, although the aggregate around the stadium occupies a much smaller area than the entire campus. Overcoming this inconsistency between the semantic dimension and the geographic dimension requires additional distortion or transformation to the visual representation. A potential solution for this might be to design a cartogram [21] that distorts the shape of the aggregates such that the area is proportional to the complexity of the semantics. However, since the aggregates are associated with a geographic context, this distortion can easily cause fidelity issues (G5) and add visual confusion to the users. In general, a trade-off exists between the semantic expressiveness and the geographic accuracy. Which factor to focus on depends on the problems and user requirements.

TopoText generates high-quality and resolution-independent SVG imaging that supports efficient interaction handling as the SVG elements are organized as nodes in the browser DOM. However, the rendering performance may degrade when a large number of graphical elements are added in the DOM. While the current interface supports nearly interactive response (the latency is usually less than 2 seconds), the rendering performance can further be improved by precomputing the visual results at different spatial scales and organizing them as hierarchical map tiles in order to improve the interactivity and alleviate the rendering overload in the browser side.

TopoText uses a density-based clustering algorithm (DBSCAN) to establish a static multi-scale hierarchy. The clustering process is performed in the back-end server for performance efficiency. Future work includes extending TopoText to analyzing streaming datasets in real-time scenarios. This can be achieved by managing a moving time window and visualizing the dynamic aggregate hierarchy within the window. Additional perceptual support, such as animated transitions [56], can provide a smooth transition between real-time visual changes in adjacent frames.

The multi-scale hierarchy established in TopoText represents the spatial proximity of data points at different scales. The application of TopoText is not limited to geographic datasets and includes various types of spatial datasets. It can also be applied to non-spatial datasets that can be spatialized into the 2D space such that the pair-wise distance of the 2D points represents the proximity of certain data dimensions. Typical examples include low-dimensional representations generated from high-dimensional data based on dimension reduction (e.g., SOM, MDS, t-SNE) [27, 34]. As the projection often preserves the pairwise distance of data points, summarizing the multi-scale aggregation hierarchy at the low dimension can potentially provide the insight into the characteristics of the data patterns in the original high-dimensional space.

CONCLUSION

We have presented a text-based visualization technique called TopoText for maintaining the semantic context in the multi-scale aggregation space. Our primary contribution includes a set of visual encoding and layout strategies that spatialize visual text labels on the boundary or in the inner space of the aggregates. We have explored and evaluated several design choices that utilize different visual attributes of text labels including color, opacity, density and orientation for multi-scale text exploration tasks.

Our future work includes optimizing the rendering performance of TopoText by precomputing the visualizations and organizing them as map tiles. We also plan on extending TopoText to supporting other types of spatial data, or non-spatial data that can be spatialized in a meaningful way. Finally, we plan on extending TopoText to exploring the multi-scale aggregation dynamics in real-time applications.

ACKNOWLEDGMENTS

This work was funded by the U.S. Department of Homeland Security VACCINE Center under Award Number 2009-ST-061-CI0003.

REFERENCES

1. Shehzad Afzal, Ross Maciejewski, Yun Jang, Niklas Elmqvist, and David S Ebert. 2012. Spatial text visualization using automatic typographic maps. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2556–2564. DOI: **http://dx.doi.org/10.1109/TVCG.2012.264**
2. Eric Carlson Alexander, Chih-Ching Chang, Mariana Shimabukuro, Steve Franconeri, Christopher Collins, and Michael Gleicher. 2017. Perceptual Biases in Font Size as a Data Encoding. *IEEE transactions on visualization and computer graphics* PP, 99 (2017), 1–1. DOI: **http://dx.doi.org/10.1109/tvcg.2017.2723397**
3. Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. 2014. Visualizing sets and set-typed data: State-of-the-art and future challenges. In *Proceedings of the IEEE VGTC/EuroGraphics Conference on Visualization*. 1–21. DOI: **http://dx.doi.org/10.2312/eurovisstar.20141170**
4. Robert Amar, James Eagan, and John Stasko. 2005. Low-level components of analytic activity in information visualization. In *Proceedings of the IEEE Symposium on Information Visualization*. IEEE, 111–117. DOI: **http://dx.doi.org/10.1109/infovis.2005.24**
5. Gennady Andrienko, Natalia Andrienko, Harald Bosch, Thomas Ertl, Georg Fuchs, Piotr Jankowski, and Dennis Thom. 2013. Thematic patterns in georeferenced tweets through space-time visual analytics. *Computing in Science & Engineering* 15, 3 (2013), 72–82. DOI: **http://dx.doi.org/10.1109/mcse.2013.70**
6. Caroline Appert and Jean-Daniel Fekete. 2006. OrthoZoom scroller: 1D multi-scale navigation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, 21–30. DOI: **http://dx.doi.org/10.1145/1124772.1124776**
7. Harald Bosch, Dennis Thom, Florian Heimerl, Edwin Puttmann, Stephan Koch, Robert Kruger, Michael Worner, and Thomas Ertl. 2013. Scatterblogs2: Real-time monitoring of microblog messages through user-guided filtering. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2022–2031. DOI: **http://dx.doi.org/10.1109/TVCG.2013.186**
8. Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D^3: Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. DOI: **http://dx.doi.org/10.1109/TVCG.2011.185**
9. Richard Brath and Ebad Banissi. 2016a. Typographic sets: Labeled set elements with font attributes. In *International Workshop on Set Visualization and Reasoning*. 29–43.
10. Richard Brath and Ebad Banissi. 2016b. Using Typography to Expand the Design Space of Data Visualization. *She Ji: The Journal of Design, Economics, and Innovation* 2, 1 (2016), 59–87. DOI: **http://dx.doi.org/10.1016/j.sheji.2016.05.003**
11. Cynthia Brewer. 2004. *Designing Better Maps: A Guide for GIS Users*. Environmental Systems Research.
12. Junghoon Chae, Yuchen Cui, Yun Jang, Guizhen Wang, Abish Malik, and David S. Ebert. 2015. Trajectory-based Visual Analytics for Anomalous Human Movement Analysis using Social Media. In *Proceedings of the EuroVis Workshop on Visual Analytics*. The Eurographics Association. DOI: **http://dx.doi.org/10.2312/eurova.20151102**
13. Ming-Te Chi, Shih-Syun Lin, Shiang-Yi Chen, Chao-Hung Lin, and Tong-Yee Lee. 2015. Morphable word clouds for time-varying text data visualization. *IEEE Transactions on Visualization and Computer Graphics* 21, 12 (2015), 1415–1426. DOI: **http://dx.doi.org/10.1109/TVCG.2015.2440241**
14. Andy Cockburn, Amy Karlson, and Benjamin B Bederson. 2009. A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys (CSUR)* 41, 1 (2009), 2. DOI: **http://dx.doi.org/10.1145/1456650.1456652**
15. Edward G Coffman Jr, Michael R Garey, and David S Johnson. 1984. Approximation algorithms for bin-packing: an updated survey. In *Algorithm design for computer system design*. Springer, 49–106. DOI: **http://dx.doi.org/10.1007/978-3-7091-4338-4_3**
16. Christopher Collins, Sheelagh Carpendale, and Gerald Penn. 2009a. DocuBurst: Visualizing Document Content Using Language Structure. *Computer Graphics Forum* 28, 3 (2009), 1039–1046. DOI: **http://dx.doi.org/10.1111/j.1467-8659.2009.01439.x**
17. Christopher Collins, Fernanda B Viegas, and Martin Wattenberg. 2009b. Parallel tag clouds to explore and analyze faceted text corpora. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. IEEE, 91–98. DOI: **http://dx.doi.org/10.1109/vast.2009.5333443**
18. Weiwei Cui, Yingcai Wu, Shixia Liu, Furu Wei, Michelle X. Zhou, and Huamin Qu. 2010. Context preserving dynamic word cloud visualization. In *Proceedings of the IEEE Pacific Symposium on Visualization*. 121–128. DOI: **http://dx.doi.org/10.1109/PACIFICVIS.2010.5429600**
19. Jean-Yves Delort. 2010. Vizualizing large spatial datasets in interactive maps. In *Proceedings of the International Conference on Advanced Geographic Information Systems, Applications, and Services*. IEEE, 33–38. DOI: **http://dx.doi.org/10.1109/GEOProcessing.2010.13**
20. Marian Dörk, Sheelagh Carpendale, Christopher Collins, and Carey Williamson. 2008. VisGets: Coordinated visualizations for web-based information exploration and discovery. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1205–1212. DOI: **http://dx.doi.org/10.1109/TVCG.2008.175**

21. Herbert Edelsbrunner and Roman Waupotitsch. 1997. A combinatorial approach to cartograms. *Computational Geometry* 7, 5-6 (1997), 343–360. DOI: http://dx.doi.org/10.1016/S0925-7721(96)00006-5

22. Niklas Elmqvist and Jean-Daniel Fekete. 2010. Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Transactions on Visualization and Computer Graphics* 16, 3 (2010), 439–454. DOI: http://dx.doi.org/10.1109/TVCG.2009.84

23. Cristian Felix, Anshul Vikram Pandey, and Enrico Bertini. 2017. TextTile: An Interactive Visualization Tool for Seamless Exploratory Analysis of Structured Data and Unstructured Text. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 161–170.

24. Nivan Ferreira, Lauro Lins, Daniel Fink, Steve Kelling, Chris Wood, Juliana Freire, and Claudio Silva. 2011. BirdVis: Visualizing and understanding bird populations. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2374–2383. DOI: http://dx.doi.org/10.1109/TVCG.2011.176

25. Alex Godwin, Yongxin Wang, and John Stasko. 2017. TypoTweet Maps: Characterizing Urban Areas through Typographic Social Media Visualization. In *Proceedings of the IEEE VGTC/EuroGraphics Conference on Visualization*. 25–29. DOI: http://dx.doi.org/10.2312/eurovisshort.20171128

26. Sarah Goodwin, Jason Dykes, Aidan Slingsby, and Cagatay Turkay. 2016. Visualizing multiple variables across scale and geography. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 599–608. DOI: http://dx.doi.org/10.1109/TVCG.2015.2467199

27. Florian Heimerl, Markus John, Qi Han, Steffen Koch, and Thomas Ertl. 2016. DocuCompass: Effective exploration of document landscapes. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*. IEEE, 11–20. DOI: http://dx.doi.org/10.1109/VAST.2016.7883507

28. Eduard Imhof. 1975. Positioning names on maps. *The American Cartographer* 2, 2 (1975), 128–144. DOI: http://dx.doi.org/10.1559/152304075784313304

29. Petra Isenberg, Pierre Dragicevic, Wesley Willett, Anastasia Bezerianos, and Jean-Daniel Fekete. 2013. Hybrid-image visualization for large viewing environments. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2346–2355. DOI: http://dx.doi.org/10.1109/tvcg.2013.163

30. Waqas Javed, Sohaib Ghani, and Niklas Elmqvist. 2012a. GravNav: Using a gravity model for multi-scale navigation. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. ACM, 217–224. DOI: http://dx.doi.org/10.1145/2254556.2254597

31. Waqas Javed, Sohaib Ghani, and Niklas Elmqvist. 2012b. PolyZoom: multiscale and multifocus exploration in 2D visual spaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 287–296. DOI: http://dx.doi.org/10.1145/2207676.2207716

32. Susanne Jul and George W. Furnas. 1998. Critical Zones in Desert Fog: Aids to Multiscale Navigation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. 97–106. DOI: http://dx.doi.org/10.1145/288392.288578

33. KyungTae Kim, Sungahn Ko, Niklas Elmqvist, and David S Ebert. 2011. WordBridge: Using composite tag clouds in node-link diagrams for visualizing content and relations in text corpora. In *Proceedings of the Hawaii International Conference on System Sciences*. 1–8. DOI: http://dx.doi.org/10.1109/HICSS.2011.499

34. Minjeong Kim, Kyeongpil Kang, Deokgun Park, Jaegul Choo, and Niklas Elmqvist. 2017. TopicLens: Efficient multi-level visual topic exploration of large-scale document collections. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 151–160. DOI: http://dx.doi.org/10.1109/TVCG.2016.2598445

35. Etien L. Koua, Alan MacEachren, and M.-J. Kraak. 2006. Evaluating the usability of visualization methods in an exploratory geovisualization environment. *International Journal of Geographical Information Science* 20, 4 (2006), 425–448. DOI: http://dx.doi.org/10.1080/13658810600607550

36. Nina Siu-Ngan Lam and Dale A. Quattrochi. 1992. On the issues of scale, resolution, and fractal analysis in the mapping sciences. *The Professional Geographer* 44, 1 (1992), 88–98. DOI:http://dx.doi.org/10.1111/j.0033-0124.1992.00088.x

37. Shixia Liu, Jialun Yin, Xiting Wang, Weiwei Cui, Kelei Cao, and Jian Pei. 2016. Online visual analytics of text streams. *IEEE transactions on visualization and computer graphics* 22, 11 (2016), 2451–2466. DOI: http://dx.doi.org/10.1109/TVCG.2015.2509990

38. Wanyu Liu, Rafael Lucas DÓliveira, Michel Beaudouin-Lafon, and Olivier Rioul. 2017. BIGnav: Bayesian Information Gain for Guiding Multiscale Navigation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 5869–5880. DOI: http://dx.doi.org/10.1145/3025453.3025524

39. Steffen Lohmann, Florian Heimerl, Fabian Bopp, Michael Burch, and Thomas Ertl. 2015. ConcentriCloud: Word cloud visualization for multiple text documents. In *Proceedings of the International Conference on Information Visualization*. 114–120. DOI: http://dx.doi.org/10.1109/iV.2015.30

40. Alan M. MacEachren. 1995. *How Maps Work: Representation, Visualization, and Design*. Guilford Press.

41. Ron Maharik, Mikhail Bessmeltsev, Alla Sheffer, Ariel Shamir, and Nathan Carr. 2011. Digital Micrography. In *ACM Transactions on Graphics*. 100:1–100:12. DOI: http://dx.doi.org/10.1145/2010324.1964995
42. Axis Maps. 2017. Typographic maps. http://store.axismaps.com/. (2017). Accessed: 2017-08-15.
43. Stanley Milgram. 1976. Psychological maps of Paris. *Environmental psychology: People and their physical settings* (1976), 104–124.
44. Tamara Munzner. 2014. *Visualization Analysis and Design*. CRC Press.
45. Miguel Nacenta, Uta Hinrichs, and Sheelagh Carpendale. 2012. FatFonts: combining the symbolic and visual aspects of numbers. In *Proceedings of the ACM Conference on Advanced Visual Interfaces*. ACM, 407–414. DOI: http://dx.doi.org/10.1145/2254556.2254636
46. Dinh-Quyen Nguyen and Heidrun Schumann. 2010. Taggram: Exploring geo-data on maps through a tag cloud-based visualization. In *Proceedings of the International Conference on Information Visualization*. 322–328. DOI: http://dx.doi.org/10.1109/IV.2010.52
47. Dinh Quyen Nguyen, Christian Tominski, Heidrun Schumann, and Tuan Anh Ta. 2011. Visualizing tags with spatiotemporal references. In *Proceedings of the International Conference on Information Visualization*. 32–39. DOI: http://dx.doi.org/10.1109/IV.2011.43
48. Victoria Rautenbach, Serena Coetzee, and Arzu Çöltekin. 2016. Development and evaluation of a specialized task taxonomy for spatial planning — A map literacy experiment with topographic maps. *ISPRS Journal of Photogrammetry and Remote Sensing* 127 (2016), 16–26. DOI: http://dx.doi.org/10.1016/j.isprsjprs.2016.06.013
49. Aidan Slingsby, Jason Dykes, Jo Wood, and Keith Clarke. 2007. Interactive tag maps and tag clouds for the multiscale exploration of large spatio-temporal datasets. In *Proceedings of the International Conference on Information Visualization*. 497–504. DOI: http://dx.doi.org/10.1109/IV.2007.71
50. Jonathan Stott, Peter Rodgers, Juan Carlos Martinez-Ovando, and Stephen G Walker. 2011. Automatic metro map layout using multicriteria optimization. *IEEE Transactions on Visualization and Computer Graphics* 17, 1 (2011), 101–114. DOI: http://dx.doi.org/10.1109/tvcg.2010.24
51. Cagatay Turkay, Aidan Slingsby, Helwig Hauser, Jo Wood, and Jason Dykes. 2014. Attribute signatures: Dynamic visual summaries for analyzing multivariate geographical data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2033–2042. DOI: http://dx.doi.org/10.1109/TVCG.2014.2346265
52. David Wong. 2009. The modifiable areal unit problem (MAUP). *The SAGE Handbook of Spatial Analysis* (2009), 105–123.
53. Pak Chung Wong, Patrick Mackey, Ken Perrine, James Eagan, Harlan Foote, and Jim Thomas. 2005. Dynamic visualization of graphs with extended labels. In *Proceedings of the IEEE Symposium on Information Visualization*. 73–80. DOI: http://dx.doi.org/10.1109/infvis.2005.1532131
54. Jo Wood, Jason Dykes, Aidan Slingsby, and Keith Clarke. 2007. Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1176–1183. DOI: http://dx.doi.org/10.1109/TVCG.2007.70570
55. Jie Xu and Craig S. Kaplan. 2007. Calligraphic Packing. In *Proceedings of Graphics Interface*. 43–50. DOI: http://dx.doi.org/10.1145/1268517.1268527
56. Jiawei Zhang, Benjamin Ahlbrand, Abish Malik, Junghoon Chae, Zhiyu Min, Sungahn Ko, and David S. Ebert. 2016. A Visual Analytics Framework for Microblog Data Analysis at Multiple Scales of Aggregation. *Computer Graphics Forum* 35, 3 (2016), 441–450. DOI: http://dx.doi.org/10.1111/cgf.12920
57. Jiawei Zhang, Abish Malik, Benjamin Ahlbrand, Niklas Elmqvist, Ross Maciejewski, and David S. Ebert. 2017. TopoGroups: Context-Preserving Visual Illustration of Multi-Scale Spatial Aggregates. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 2940–2951. DOI: http://dx.doi.org/10.1145/3025453.3025801