

Reducing Ambiguities in Line-based Density Plots by Image-space Colorization

Yumeng Xue, Patrick Paetzold, Rebecca Kehlbeck, Bin Chen, Kin Chung Kwan, Yunhai Wang, and Oliver Deussen

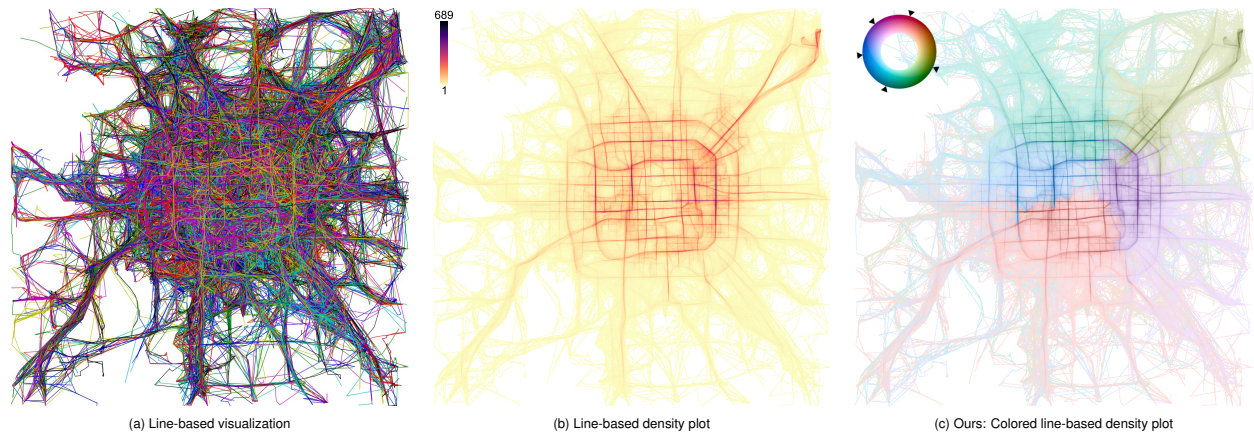


Fig. 1: Trajectories of taxi rides in Beijing: (a) A line-based visualization of the trajectories is cluttered and convoluted; (b) Using a density plot eliminates clutter, but the continuation of the trends is ambiguous, e.g., it looks like taxis follow the prominent, circular route around the city center; (c) Pixel-based colorization reveals clusters, we see that taxis mostly stay in one part of the city.

Abstract—Line-based density plots are used to reduce visual clutter in line charts with a multitude of individual lines. However, these traditional density plots are often perceived ambiguously, which obstructs the user’s identification of underlying trends in complex datasets. Thus, we propose a novel image space coloring method for line-based density plots that enhances their interpretability. Our method employs color not only to visually communicate data density but also to highlight similar regions in the plot, allowing users to identify and distinguish trends easily. We achieve this by performing hierarchical clustering based on the lines passing through each region and mapping the identified clusters to the hue circle using circular MDS. Additionally, we propose a heuristic approach to assign each line to the most probable cluster, enabling users to analyze density and individual lines. We motivate our method by conducting a small-scale user study, demonstrating the effectiveness of our method using synthetic and real-world datasets, and providing an interactive online tool for generating colored line-based density plots.

Index Terms—Trajectory data, times series, density-based visualization, clustering, coloring

1 INTRODUCTION

Line-based plots are a popular way to visualize time series data [20] and trajectory data [10]. These plots are widely used across domains, such as finance, healthcare, and navigation, to highlight evolving trends in complex datasets. To emphasize patterns and gain insights, two effective techniques are line coloring and density plots. By using line coloring, different lines in a plot can be assigned distinct colors, making it easier to distinguish between different groups or categories. This technique is particularly useful when dealing with multiple time series or trajectories that are plotted together. Alternatively, density plots show the distribution of data at each point in time or space. These plots use color or shading to indicate areas of high or low data density,

providing a more detailed picture of the underlying patterns in the data.

Line coloring assigns a color to each line. However, even in small datasets, line intersections can cause overplotting, where lines are drawn on top of each other, making it difficult to interpret the data. This problem becomes more pronounced in complex datasets with a larger number of lines, resulting in visual clutter that can obscure underlying trends [30], see Fig. 1a. Mainly, this has been addressed by replacing individual lines through visual abstractions [17, 42, 73], which use abstract glyphs to display statistical information about the lines.

Using density plots [34, 44] shifts the focus from visualizing individual data items to aggregating data attributes in image space, communicating the structure of the data better. By using a continuous color map to visually communicate the density at a given position, density plots can effectively reduce visual clutter. However, the individual line information is lost. Such plots do not reveal which and how lines intersect at a specific position. To alleviate this issue, interaction techniques have to be used (e.g., timebox [23]) with representative lines (such as kd-box [81]) being used to explore patterns of interest.

In general, due to the continuous nature of time series and trajectory data, line-based density plots are more challenging to interpret than, e.g., point-based density plots. Visual ambiguity within the perceived patterns often leads to false conclusions, as the Gestalt principles of good continuation and similarity strongly influence the perception of interacting lines. Similarly, the Principle of Unambiguous Data Depiction

• Y. Xue is with University of Konstanz and Shandong University. E-mail: yumeng.xue@uni-konstanz.de.

• P. Paetzold, R. Kehlbeck, B. Chen, and O. Deussen are with University of Konstanz. E-mail: firstname.lastname@uni-konstanz.de.

• K.C. Kwan is with California State University Sacramento. E-mail: kwan@csus.edu

• Y. Wang is with Shandong University. E-mail: cloudseawang@gmail.com.

• O. Deussen and Y. Wang are joint corresponding authors.

introduced by Kindlmann and Scheidegger [33] describes visualizations that fail to be unambiguous as having so-called “confusers”. This means different data inputs can result in visualizations that are not distinguishable. For example, in Fig. 1b, it looks like most taxis drive on the circular route around the city center. The question is if this observed trend exists in the data or if other combinations of patterns could result in a similar visualization. Both line coloring and traditional density plots offer no solution to help users to identify the existence, continuation, and intertwinement of trends. Clustering the data using our method shows that there is no main circular trend around the city center; rather, taxis mostly drive within local segments of the city.

Our idea is to use an image-space coloring method to analyze line-based density plots. In comparison to traditional density plots, we use color not just as an indicator for density but also to visually highlight similar regions in the plot. Each region is characterized by the lines passing through it. As a result, regions with similar lines are colored similarly, allowing users to identify and distinguish trends. To obtain the colors, we perform hierarchical clustering based on the lines passing through regions and subsequently map the identified cluster to the hue circle using circular MDS. To provide additional line-based analysis, once the bins in the density plot are clustered, we suggest a heuristic approach to allocate each line to the most probable cluster. This enables the user to explore not only the density plot but also individual lines. We demonstrate the capabilities of our method using selected synthetic and real-world datasets. Our main contributions of this paper are:

- A novel image space coloring scheme to enrich line-based density plots with similarity information;
- A heuristic method to assign lines to their most likely cluster;
- A motivational user study to demonstrate the ambiguous perception of line-based density plots; and
- An interactive online tool to generate colored density plots.

In the subsequent sections, we illustrate ambiguities in the perception of trends in line-based density plots using various examples, propose our novel image-based coloring method, and perform a user study to identify ambiguities in line-based density plots.

2 RELATED WORK

In our discussion of related work, we focus on density plot generation, visual ambiguities and illusions, line-based analysis, and color mapping techniques.

2.1 Density-based Visualizations

Density plots are a popular method to generate uncluttered scatter plots [8]. There are several techniques to visualize density, including opacity blending [39], kernel density estimation (KDE) [16, 59], and binning followed by summing for color mapping [74]. To reduce visual clutter, density-based adaptations of various types of visualizations for different data types have been proposed. For example, Artero *et al.* [3] used density-based filtering for parallel coordinates, while Zinsmaier *et al.* [83] proposed an interactive rendering method for large-scale graphs using KDE-based node aggregation. Scheepens *et al.* [58] used variable KDE kernel radii for user-customizable trajectory exploration with density maps and extended their technique to combine density fields of multiple attributes in a single visualization [57]. Wickham [74] employed binning, summarizing, and smoothing to abstract large datasets and emphasize patterns, while Jerding and Stasko [31] introduced a reduced representation of line charts by using gray-scale values based on the level of overlap. Lampe *et al.* [34] extended KDE to curves and called their method curve density estimates (CDE). Recently, Moritz and Fisher [44] proposed DenseLines, a discretized CDE variation that allows parallel computing on the GPU. However, all these works focus on calculating pixel density and do not consider the relationships between pixels.

2.2 Visual Ambiguities and Illusions

Much research has been done to tackle visual ambiguities or illusions that lead humans to draw wrong conclusions. Feng *et al.* [16] visually

encoded uncertainty present in the data in scatter plots and parallel coordinates to prevent users from drawing false conclusions about the data. Pomerence *et al.* [51] explored the relationship between slope and the perceived prominence of lines in ghost clusters and proposed a slope-dependent density correction method to reduce visual errors. For scatter plots, Liu *et al.* [37] refine the orientation of marks to guide the users in estimating trends to prevent imprecise estimations biased by human vision. Hong *et al.* [26] studied how the size and lightness in scatterplots affect the perceived mean. Some other methods use additional visual coding to reduce visual ambiguities and illusions. Novotný and Hauser [46] proposed an outlier preserving method for parallel coordinates. They extract outliers by binning in the adjacent dimensions and detect outlier bins. Trautner *et al.* [64] highlight outliers by overlaying additional visual coding on the scatter-based density map. Unfortunately, there is no current work that focuses on visual ambiguities and illusions in line-based density plots.

2.3 Line-based Analysis

Line-based data analysis methods use lines as fundamental units of data and enable various analyses, including clustering, abstraction, and interactive exploration. Line clustering uses different metrics [80] to measure the distance between 2D or 3D lines, such as the Euclidean distance [11, 55], curvature or torsion [41, 75], and user-specified stream-line predicates [56]. Dynamic time warping (DTW) [45] is another widely used metric to measure similarity between time series data. The clustering of lines can be achieved by applying these similarity metrics. To obtain cluster centers for clustering time series data, Petitjean *et al.* [49] proposed a global averaging method based on DTW, while Gaffney and Smyth [21] presented a framework for clustering curves based on probabilistic curve alignment models. Some methods use Euclidean space vectors to represent the line curves and then apply Euclidean distance-based clustering methods [9, 55].

However, even after obtaining clustering information, further analysis remains a challenge because lines can still be cluttered, and clustering methods are sensitive to noise. Uncertainty visualizations therefore use glyphs to visualize statistical information and enable overviews. Mirzargar *et al.* [42] generalized the curve boxplot for ensembles of curves by introducing a functional band depth. Ferstl *et al.* [17] applied clustering to construct curve-boxplot-like abstractions of multiple line bundles. Palmas *et al.* [48] abstracted clusters in parallel coordinates through edge bundling, facilitating further interaction with clusters. These visualization methods are effective only for regularly distributed lines and are not suitable for noisy datasets.

In addition to the previously introduced clustering approaches that either visually cluster lines by color or abstract them, edge bundling techniques [24, 25, 66] alter the course of individual lines to combine them into bundles to reduce visual clutter. To incorporate geographical restrictions like roads, specially tailored edge bundling techniques [62, 79] were proposed. Recently Wallinger *et al.* [68, 69] proposed an edge-path bundling method that reduced the ambiguities introduced by edge bundling. Our method is focused on reducing the ambiguity while not altering the original data, as edge bundling does.

Interaction methods are also important for analyzing lines. For example, QuerySketch [71] allows users to sketch freely and query the data lines that match the shape of the sketched line. Hurter *et al.* [28] introduced interactive paradigms to extract trajectories of interest from large-scale airline trajectory data. Hochheiser and Shneiderman [23] presented Timebox to query data that passes through a box, representing a range of positions within a certain time. Recently, Zhao *et al.* [81] proposed KD-Box by using a KD-tree to speed up the timebox query and introduce representative lines to assist analysts in exploring the local pattern details of interest. However, these methods require significant user interaction and learning costs, and the small number of representative lines may not fully reflect the overall composition of patterns.

2.4 Color Mapping

Color maps are an important aspect of data visualization, as they have a significant impact on how effectively the viewer perceives and un-

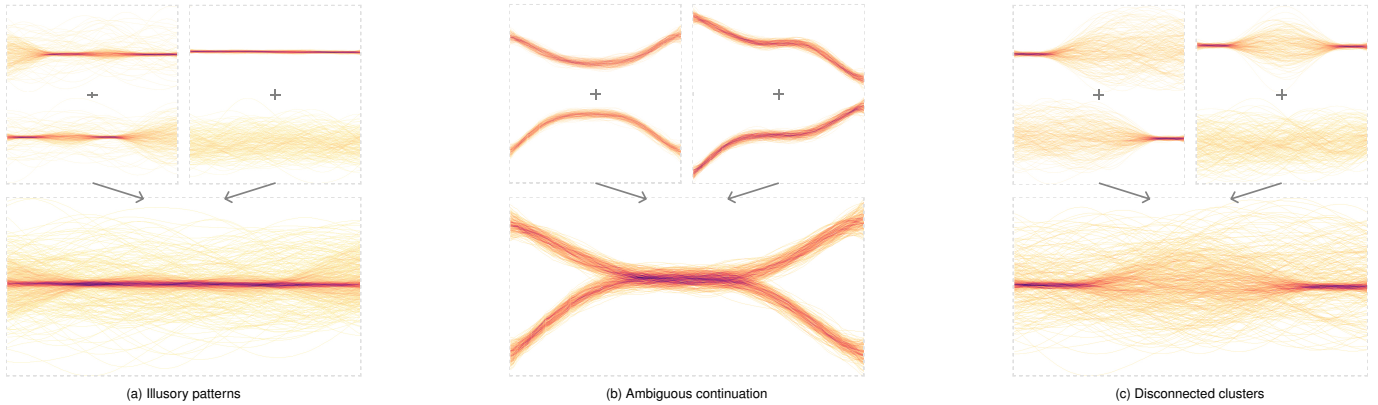


Fig. 2: Three examples of ambiguities in line-based density plots. Each density plot in the lower row could be a superposition of one of the two trends in the upper row. (a) The density could come from two trends with varying densities or from a clear trend and added noise. (b) The pattern could be two touching trends or two crossing ones. (c) The trend could be a single, continuous trend, or a combination of two trends that diverge.

derstands the underlying data. Zhou and Hansen [82] conducted a comprehensive survey of color map generation methods and classified them into four categories: procedural-based [43, 50, 53, 65], user-study-based [22, 32, 54], rule-based [4, 52, 63], and data-driven [35, 60, 61, 70]. In the context of density plots, sequential color maps are typically used to display the gradual changes in density values. The Brewer color palettes, established by Brewer *et al.* [5–7], provide a set of commonly used sequential, diverging, and qualitative color maps and guidelines for selecting appropriate colors for different types of visualizations.

Color maps may introduce visual artifacts or lack perceptual consistency. Perceptually consistent color maps are designed to match the human perception of color and are constructed using color spaces such as HCL [27]. Zeileis *et al.* [78] proposed a perceptually consistent single-hue color map for density visualizations. Lu *et al.* [38] proposed Palettator to optimize categorical colorization. It can be used to obtain perceptually consistent colors while maintaining the discriminability of positionally close categories. However, coloring methods for categorical data (e.g., multiclass scatterplots) are not applicable to our problem, as the colors they assign do not carry density information. Our goal is twofold: we want to visualize a cluster membership while retaining density information. Therefore, multiple single-hue colormaps based on HCL color space are well suited for this problem. They convey clustering information using hues while maintaining visual consistency for the same density. In addition, the hue harmonic method proposed by Cohen-Or *et al.* [12] can be used to further optimize the hue selection.

3 DENSITY PLOTS OF LINES

We follow the process proposed by Moritz and Fisher [44] to create the line-based density plot. Firstly, the visual area is divided into bins (in the smallest case, a bin is one pixel). For each bin, the number of entities — points in scatter plots and lines in line charts — are counted, and the density is mapped to a sequential color palette. This is the application of Wickham’s bin-summarize-smooth [74] paradigm to line data. In the subsequent paragraphs, we provide a detailed overview of sources of ambiguity in the interpretation of line-based density plots. In Section 6, we use the three defined sources of ambiguities to show in a small-scale motivational user study how participants perceive ambiguous line-based density plots.

3.1 Time series and Trajectory Definition

Mirzargar *et al.* [42] define trajectories mathematically as parametric curves mapping from a domain \mathbf{D} to a potentially higher dimensional co-domain \mathbf{R} . As we aim to generate two-dimensional density plots, we restrict our consideration of trajectories to one and two-dimensional co-domains. For one-dimensional co-domains, the trajectory is a time series representation. Two-dimensional co-domains can represent, for example, geospatial data, like traffic data or human movements, or more abstract spaces, like eye-tracking trajectories.

3.2 Ambiguities in Density Plots

While phenomena like visual ambiguities and visual artifacts introduced by scatter plots [40] and line-based plots have been studied [51], little research has been done on density plots. Compared to density representations of scatter plots, which are straightforward to interpret, the interplay of multiple lines in a line-based density plot might add visual ambiguities hampering their clarity. Investigating line-based density plots of real-world datasets, we found patterns with ambiguities and further compiled a list of three abstracted types of patterns (Fig. 2). We focus ourselves on these three recurring patterns in our investigation and do not claim the list to be complete. In Section 6, we show the results of our user study, underpinning that users perceive line-based density plots ambiguously. The upper rows of each subfigure show possibilities of individual trends that, when combined, form a similar pattern. Even if only two trends are superimposed, it is challenging to decide which individual trends the density plot is composed of.

Illusory patterns. The combined density plot in the lower row of Fig. 2a is visually dominated by the central horizontal line. An observer could assume this dominant region is superimposed on a less dense but noisy region surrounding it. This perceived continuation is also supported by the Gestalt principle of good continuity [72]. Our example illustrates that this perception is deceptive, as the observed pattern is actually a composition of the two patterns on the top left. This ambiguity is often present in real-world datasets, such as ship trajectories before and after passing through a strait (see Fig. 6). We refer to patterns that appear to have a high-density trend only in the combined line-based density plot that is not present in the individual underlying patterns as “illusory patterns”.

Ambiguous continuation. If multiple independent, well-separated trends partly overlap with a similar orientation in the overlapping region, their continuation is ambiguous. In the lower row of Fig. 2b, it is not clear how the individual trends continue after intersecting. Both combinations on the top could create such a pattern. Thus, their continuation remains unclear. This ambiguity frequently appears in data we have observed, such as time series data with crossovers and overlaps between trends (see Fig. 4).

Disconnected clusters. Dense regions in plots attract attention. If a line-based density plot contains multiple dense regions, it is ambiguous if they belong to the same trend. A line-based density plot with two such dense regions is shown in the lower row of Fig. 2c. Due to their visual saliency and similar density, they appear to be related. Perceptual psychology explains this with the Gestalt principle of similarity. However, this impression is deceptive since the pattern is actually composed of the two completely independent patterns on the left of the upper row of Fig. 2c. Thus it is ambiguous if the combined density plot consists of two line bundles fanning out in the center of the plot or two independent trends. It is common to see disconnected clusters in real-world datasets, an example can be observed in Fig. 5. Even these straightforward examples show that interpreting line-based density plots is challenging,




Fig. 3: Processing pipeline: per region (bin) sets of features (line IDs) are obtained by checking which lines pass through a bin. By sampling, we obtain a subset of the bins that allow us to cluster them efficiently. Hierarchical clustering is then applied to find similar bins. Using circular MDS we map cluster centers to the HCL color space to create the rendition.

and identifying the underlying trends is often ambiguous. Patterns in real-world data are even more complex than our examples. We acknowledge the possibility of the existence of other types of ambiguities, which is an interesting topic for further research. Already just the three types which we identified show that there is a need to visually better convey trends and patterns. This is the motivation for us to propose our density plot coloring method that helps to identify individual patterns within line-based density plots. Our method assigns similar colors to regions with similar lines passing through (Figs. 8a, 8c, and 8e), which helps users intuitively understand the true composition of the related line patterns and their relationships.

4 METHOD

Our goal is to reduce ambiguities in line-based density plots, identifying and highlighting similar regions. Previously introduced data-space clustering methods cluster individual data items, like lines or scatter points. In contrast, we apply clustering in the image-space, grouping regions based on the lines passing through them.

Generally, creating density plots involves three main steps: discretizing the canvas into 2D bins, counting the data items that touch each bin, calculating the bin’s density, and optionally applying a smoothing method, e.g., a kernel method. This bin-summarize-smooth paradigm was introduced by Wickham [74]. Our method assigns colors to bins in the image space based on their similarity. Similar bins are grouped into clusters with hue values assigned to them.

Our pipeline (see Fig. 3) consists of four main steps: bin-based feature set extraction, reduction of the number of bins by sampling and thresholding, hierarchical clustering, and finally, the assignment of colors to bins. We implemented our method as a web-based analysis tool that offers various user-controllable parameters to analyze the data and steer the colorization.

4.1 Feature Set Extraction

As a first step, we discretize the plot’s canvas into bins. We see a bin as an atomic unit of the density plot, which can contain multiple pixels, but each bin can only have one color, i.e., the number of bins can be considered as the logical resolution. For simplicity, in the following, a bin is equivalent to a pixel. We define a feature set for each bin that contains the identifiers of all lines touching it. This feature set expresses the relationship between a bin p and every line $L = \{L_1, L_2, \dots, L_N\}$ in the dataset. It encodes which lines of L are spatially close to the center of p . For each bin, we obtain a feature set $S(p)$:

$$S(p) = \{i \mid L_i \in L \wedge \text{distance}(L_i, p) < T\}. \quad (1)$$

where $\text{distance}(L_i, p)$ is the closest distance between line L_i and center of bin p . The parameter T is conceptually the radius of a disc positioned at the center of p , defining the extraction area for that bin. As default, we set $T = 1$. Due to the circular distance measure, the disc created by T may cover a slightly larger area than the bin. The feature set $S(p)$ then contains the identifiers of all lines running through the disc at position p . Increasing T expands the extraction area of bins, which will result in more similar feature sets for adjacent bins.

4.2 Bin Sampling

If the bins correspond to individual pixels, the number of feature sets $S(p)$ is equal to the number of pixels in the density plot. Even a medium-resolution density plot with a resolution of 1000×500 pixels

contains 500,000 bins. In the subsequent step of our pipeline, we use agglomerative hierarchical clustering [29] to group bins based on their similarity. It has a time complexity of $O(n^2 \log(n))$, which makes it challenging to use it to cluster such a large number of bins. One potential way to reduce the number of bins to increase the bin size. However, this inevitably sacrifices details and results in a loss of resolution. Alternatively, we can either sample bins or exclude bins of low-density regions. Using sampling has the advantage that we can exploit the continuous nature of the line-based datasets that underlay the visualization - feature sets of neighboring bins are typically similar due to the spatial continuity of lines. Thus sampling does roughly retain the distribution of the feature sets. In addition, users are typically interested in sufficiently dense regions. Low-density bins (e.g., those touching fewer than ten lines) typically contain a lot of noise and are less noticeable to the user. To filter out such bins, we provide a user-customizable minimum density threshold.

4.3 Hierarchical Clustering

We group the sampled bins using agglomerative hierarchical clustering. To cluster the bins, we have to use a distance metric that is appropriate for set-type data. Commonly employed are the Jaccard index $J(A, B) = \frac{|A \cap B|}{|A \cup B|}$ and Sørensen-Dice coefficient $DSC(A, B) = \frac{2|A \cap B|}{|A| + |B|}$ for sets A and B . A ratio close to 1 indicates a high degree of similarity. As densities across line plots can vary greatly, the sizes of the feature sets also differ heavily. Because the Jaccard index and Sørensen-Dice coefficient are susceptible to variations in the set size, their values tend to be lower when comparing sets of vastly different sizes, even if one is a subset of the other. Therefore, for line-based density plots the overlap coefficient [67] seems to be more suitable:

$$\text{overlap}(A, B) = \frac{|A \cap B|}{\min(|A|, |B|)} \quad (2)$$

It is insensitive to differences in set sizes as the denominator is solely determined by the smaller set. In a line-based density plot, the feature set of a bin in a low-density region (with dispersed lines) is typically a subset of the feature set of a bin in high-density areas. Because lines passing through a lower-density bin are likely to also pass through a higher-density bin, where the lines are more aggregated. In this case, we expect the two sets to have a high degree of similarity. The overlap coefficient accurately reflects this phenomenon. A more detailed comparison is available in Section 2 of the supplementary material.

We use average linkage hierarchical clustering [15] as it can be used with arbitrary distance metrics, and the number of clusters does not have to be specified in advance. Additionally, the user can interactively divide the clusters to analyze the data further. We did not use well-established partitioning clustering methods such as PAM to identify the clusters since, here, the number of clusters has to be known in advance.

4.4 Cluster Assignment

In the previous step, we grouped the sampled bins into clusters based on their similarity. However, the bins that were not part of the sample have not been attributed to a cluster as of now. To assign clusters to these bins, we have to calculate the similarity between clusters and bins, which is not easily possible using the above-mentioned overlap measure. Since only the similarity between bins can be calculated, we would have to calculate the overlap coefficient similarity of an unassigned bin

to all bins of each cluster, which is time-consuming. To solve this, we instead compute a mean feature vector $M(C)$ for each cluster C . The feature vector $V(p)$ of the feature set $S(p)$ of a bin is a binary vector with entries set to 1 if the identifier i of the line L_i is contained in $S(p)$ and 0 otherwise:

$$V(p) = (v_0^p, v_1^p, \dots, v_n^p), \quad v_i^p = \begin{cases} 1, & i \in S(p) \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

The mean feature vector of a cluster C is defined as the element-wise sum of the binary feature vectors $V(p)$ of all its bins p , divided by the number of bins in C .

$$M(C) = \frac{1}{|C|} \sum_{p \in C} V(p) \quad (4)$$

To assign a cluster to each bin, we identify the most similar cluster C for this bin p . To calculate the distance between the bin's feature vector $V(p)$ and the mean feature vector $M(C)$ of the cluster C , we use the Euclidean distance as a similarity measure. Note that when calculating the distance, we only consider the position where v_i^p is 1:

$$D(p, C) = \sum_{i \in S(p)} (1 - M(C)_i)^2 \quad (5)$$

This is similar to our rationale for choosing the overlap coefficient because if we consider all v_i^p , it would also count the lines not touching the bin, which leads to an unreasonable increase in the distance.

4.5 Cluster Colorization

In the next step, we assign colors to clusters so that these colors visually represent the similarities among them – bins of similar clusters should be colored similarly, while dissimilar clusters should be colored differently. Thus, a bin's color solely depends on the cluster it is assigned to and does not, e.g., reflect the orientation of the lines passing through it. A bin's color should visually communicate both the bin's assignment to a cluster and the density of the bin. Changing a bin's cluster assignment should not drastically alter its perceived density. Therefore, we use the perceptually uniform color space HCL (Hue-Chroma-Luminance) [27] that uses hue, chroma, and luminance as its dimensions. Since it is derived from perception science [78], it is increasingly used for visualization. Luminance can be altered independently of chroma and hue [78]. Different clusters with the same density should ideally be perceptually comparable to enable the user to compare densities across multiple clusters colored with different hues. Mapping density only to luminance would allow for an easy comparison of densities, but the resulting color map would not be rich in contrast. Similar to [78], we map the density of a bin simultaneously to chroma and luminance. It seems more desirable to map densities to a contrast-rich color map that matches the typical characteristics of density plot colorizations, such as darker colors for high density and light colors lower density.

Hue values can be represented as values on a hue circle in the range $[0, 2\pi]$. Therefore we need a mapping from the high dimensional mean feature vector $M(C)$ of a cluster to the hue circle that preserves the similarity between feature vectors.

Circular multidimensional scaling (MDS) is a nonlinear, non-metric dimensionality reduction method that was developed for exactly this application. It differs from the commonly used MDS, which uses the Euclidean metric. The stress measure in this variant of MDS [13] is defined as follows:

$$S = \sqrt{\frac{\sum_{i < j} (\delta(x_{i,j}) - d_{i,j})^2}{\sum_{i < j} d_{i,j}^2}} \quad (6)$$

Here, $\delta(x_{i,j})$ is the linear scaling of the distance between points i and j to the distance on the circle, and $d_{i,j}$ is the distance between points i and j after dimensionality reduction to the unit circle. In contrast to the

approximate distances used by TF Cox et. al. [14], the $d_{i,j}$ in our case is the accurate distances along the circular arc:

$$d_{i,j} = \begin{cases} |\theta_i - \theta_j|, & |\theta_i - \theta_j| \leq \pi \\ 2\pi - |\theta_i - \theta_j|, & |\theta_i - \theta_j| > \pi \end{cases} \quad (7)$$

Therefore, for the unit circle, $\delta(x_{i,j}) \in [0, \pi]$. We use the partial derivative $\frac{\partial S}{\partial \theta_k}$ (details see the Section 1 of the supplementary materials) of the stress S with respect to the angle θ_k of a point k for allowing gradient descent. By doing so we optimize the angle of each point and this way determine the hue value corresponding to each cluster center.

4.6 Cluster-based Line Filtering

Our coloring method identifies the main trends and highlights them through colorization. While being able to analyze individual trends in isolation and identifying the lines contributing to them is an important task, traditional density plots and our colorization aggregate lines in bins and thus do not directly provide a correspondence between trends in the density plot and individual lines. Therefore, we have to assign each line to a cluster of the density plot. To determine the cluster of a line, we focus on two aspects: grouping the bins of each cluster and summing up their weights. Counting the bins of each cluster that a line passes through already gives us a good intuition about which cluster may describe the line best. However, the perceived trends mainly depend on high-density regions, so bins of these regions should be more influential. Inspired by density-based edge clustering [36] and representative line selection [81], we, therefore, also consider the density of the bins a line passes through. Bins are expressed as:

$$C_k^i = \{p_j \mid p_j \in C_i \wedge i \in S(p_j)\} \quad (8)$$

where C_k^i represents the set of bins of the k -th cluster C_k that line i passes through. Then, we compute the sum of densities of the bins in the set C_k^i as follows:

$$W_k^i = \sum_{p_j \in C_k^i} D(p_j) \quad (9)$$

where $D(p_j)$ is the density of bin p_j . Thus, we assign line i to the cluster C_k with the greatest summed-up weight W_k^i . This allows users to analyze the details of the lines matched to each cluster.

5 RESULTS

To demonstrate the effectiveness of our approach and compare it to traditional coloring schemes of density plots. We selected four datasets: two common time-series datasets that have been used in [44, 81] and two trajectory datasets: Hellenic Trench AIS data [18] and taxi rides in Beijing [76, 77] (featured in the Fig. 1).

Temperature data. This dataset contains 293,175 weekly maximum temperature values of 6187 time series obtained from stations across the United States [1]. The line plot in Fig. 4a is highly convoluted and offers no insight into the underlying trends in the data. The conventional density plot in Fig. 4b visually separates high from low-density regions, but there are ambiguities about how individual trends continue after crossing, e.g., in the region marked "A" on the right. In addition, the parts of the trends circled as "B" show zigzag features and align with the ambiguous continuation problem shown in Fig. 2b. Our coloring method in Fig. 4c helps to reduce these ambiguities and reveals patterns that are not easily visible otherwise. In particular, the blue cluster is visually separated from other trends, indicating areas with consistently high temperatures, which makes region "B" more clearly stand out. Meanwhile, the purple and green trends intersect at two positions (marked as "A" in Fig. 4c), revealing areas that experience sudden temperature changes, making them more anomalous at certain times. The line counts for these clusters were 2180 (blue, Fig. 4d), 793 (green, Fig. 4e), and 742 (purple, Fig. 4d). This demonstrates that there are indeed many lines that follow the trend of these clusters. The lines in the blue cluster remain at higher temperatures, while lines of the green

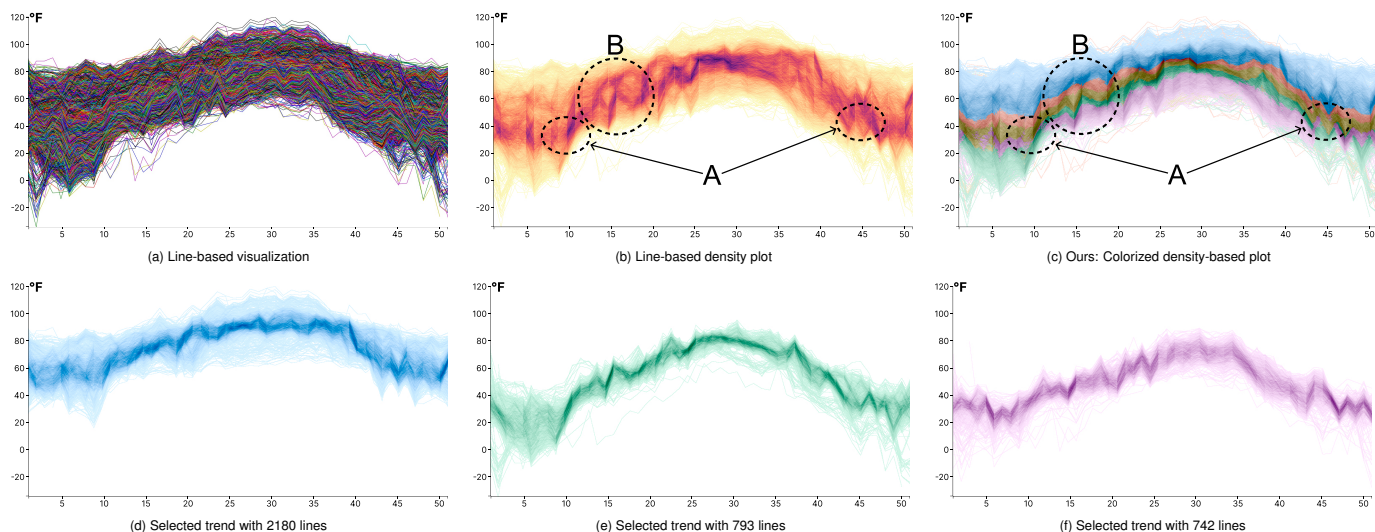


Fig. 4: Time series of temperature values (6187 lines) visualized using (a) line-based visualization, (b) line-based density plot, and (c) our line-based density plot colorization scheme. (d) - (f): Individual trends selected using our interactive tool.

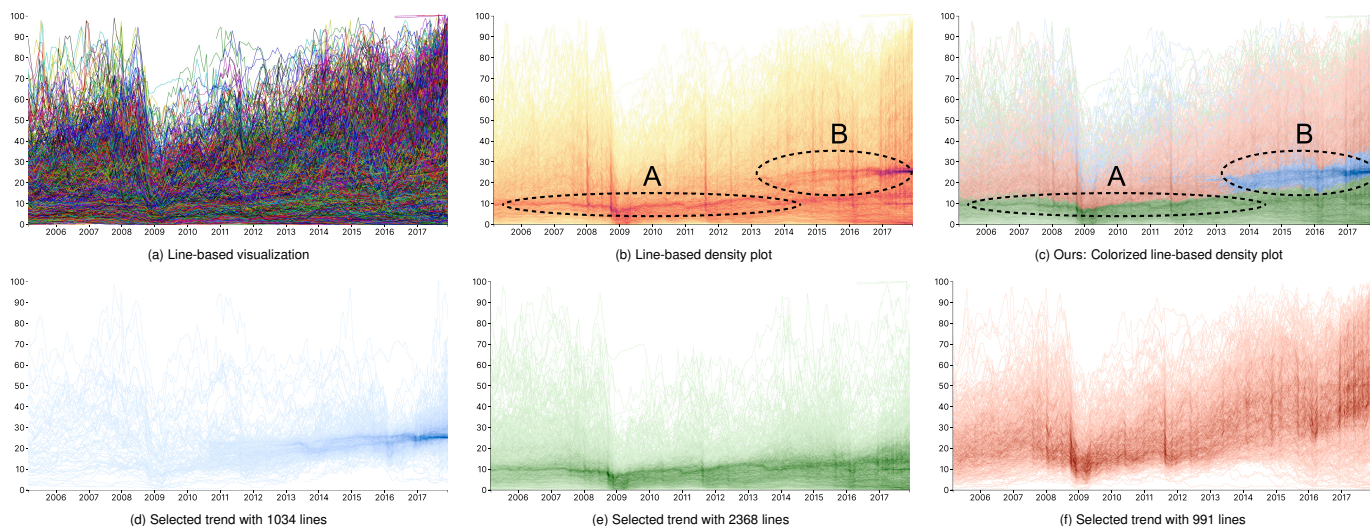


Fig. 5: Daily closing prices of 4393 stock at the New York Stock Exchange from 2005 - 2017 visualized using (a) line-based visualization, (b) line-based density plot, and (c) our line-based density plot colorization scheme. (d) - (f): Individual trends selected using our interactive tool.

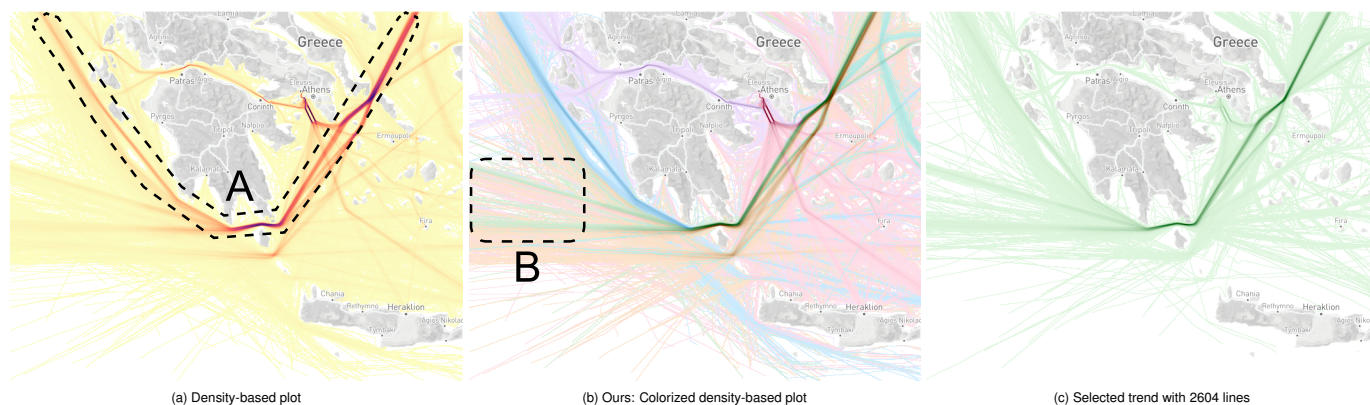


Fig. 6: Hellenic Trench AIS data [18]: (a) density-based visualization of randomly selected 10,000 trajectories with 1,048,575 points. A “U”-shaped trend seems to dominate the plot; (b) The colored density-based plot instead highlights several high-density trends that intersect; (c) The lines associated with the green cluster show a converging pattern of the cluster after vessels passed the channel (from west to east) between two islands.

and purple clusters cross at certain positions, supporting our previous conclusions drawn from the colored line-based density plot.

Stock market data. In the upper row of Fig. 5, we show 552,559 daily closing prices of 4393 stocks (with prices between 0 and 100; we removed stocks with only one close price) at the New York Stock Exchange [2] from 2005 to 2017. Although the raw line-based visualization in Fig. 5a is cluttered, it is still possible to observe that there are many lines at the bottom of the graph. The pattern (highlighted as “A”) at the bottom of the density plot in Fig. 5b aligns with this observation. However, there is another high-density pattern, which is shorter and concentrated towards the end of the graph (highlighted as “B”), that only appears in the density plot. These two disconnected high-density patterns show the issue of disconnected clusters we presented in Fig. 2c in a real-world dataset. Our colorization in Fig. 5c illustrates that the clusters marked “A” and “B” are actually disconnected. The cluster “B” does not start from the beginning of the time period but instead starts from in between. This is because many stocks were newly listed after 2010 and did not have corresponding records before that time (see Fig. 5d). After matching lines to the three clusters in the colored line-based density plot, the lines of the individual clusters are shown in the lower row of Fig. 5. The green cluster represents the main trends and contains more than half of the total lines (Fig. 5e). It represents a large number of stocks fluctuating at lower prices. The red cluster contains fewer lines, but the prices of those stocks are relatively high, although they are more scattered (Fig. 5f).

Ship trajectories. The Hellenic Trench AIS dataset [18], containing more than 170,000 vessel trajectories, was originally collected to analyze the impact of ship routes on the survival of Mediterranean sperm whales [19]. At first, we cleaned the dataset by removing trajectories crossing land, likely introduced by GPS recording errors. Due to the memory limitation of the browser, we randomly selected 10,000 trajectories with 1,048,575 points to generate the density plot and our colorized version. The density plot in Fig. 6a highlights several intersecting high-density trends. Visually, a dominant “U”-shaped trend (marked as “A” in Fig. 6a) begins west of the Peloponnese and continues eastwards to Athens. Similar to Fig. 2a, based on the density map alone, it is unclear if the trend is illusory or not. After coloring the density plot by our method (Fig. 6b), the western and eastern parts of the “U”-shaped trend are colored blue and green. Thus they are likely two separate patterns. Through our line assignment method, 2604 lines were assigned to the green cluster, as shown in Fig. 6c. The lines of the green cluster are almost on the same route after entering the channel but beforehand come from various directions. Although there are some trajectories that match the visual “U”-shaped trend, such lines are not as dense as indicated in the original density map. Our method succeeds in finding lines that match other clusters and is able to decompose them into several clusters of line bundles. Another point worth noting is that the pixels in the lower left corner of Fig. 6b (marked as “B”) look like the lines are headed in one direction, but they have been grouped into different clusters. This reflects a limitation of our approach: it is not able to, e.g., measure parallelism or other geometric features of lines.

Taxi trajectories. The Beijing taxi trajectory dataset [76, 77] contains GPS trajectories of 10,357 taxis on February 2-8, 2008. As this dataset is very large, loading it in our tool exceeded the browser’s memory limitations. Thus, we sampled 6502 trajectories of 500 taxis with 734,967 individual time points to generate Fig. 1. The line plot in Fig. 1(a) shows severe overplotting, hiding most of the dataset’s structure. However, although the line-based density plot in Fig. 1(b) reveals the road network of Beijing, it fails to display frequent taxi routes, e.g., it is ambiguous if the taxis circle around the city center or if they drive outwards or inwards from there. Our method allows to identify different clusters. The lines associated with the green and purple clusters in Fig. 1 are shown in Fig. 7. While the green cluster represents taxis driving along the airport highway, where many taxis exit at the toll booths to reach different destinations, the purple cluster represents taxis traveling within the eastern part of Beijing. This can be explained by the presence of a large residential area east of Beijing. Similar patterns can be observed for the other clusters. Taxis are unlikely to circle the city center but are likely to drive in one part of the city.



Fig. 7: Filtered lines from two clusters of Fig. 1.

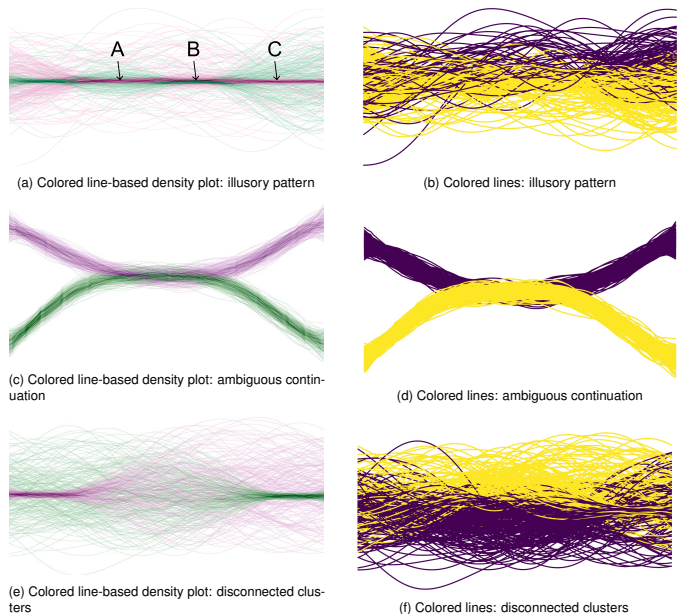


Fig. 8: Comparison between colored line-based density plots and line-based clustering for the scenarios of Fig. 2.

5.1 Comparison to Line Clustering Approaches

We use the three synthetic datasets shown in Fig. 2 to illustrate the results of coloring lines by clusters. The dataset used for Fig. 2a contains 400 lines representing the actual pattern and, additionally, 100 lines of background noise. Figs. 2b and 2c contain 200 lines per pattern.

Existing methods treat lines as basic data units and focus on line-based analysis to find trends in the data [17] using clustering. To assign colors, the lines are clustered using a similarity measure tailored to line data. While these approaches work well for lines with distinct line bundles, they do not work well for noisy data. To highlight the limitations of such methods, we use the state-of-the-art line-based method proposed by Ferstl *et al.* [17] and compare their results to ours. Their approach involves embedding lines into a higher-dimensional vector space, followed by clustering in a lower-dimensional Euclidean space after using PCA.

To compare our approach to the line-based method, we apply both methods to the examples shown in Fig. 2. While the line-based method fails to distinguish the line bundles in two of the three cases due to a less separable data distribution in low-dimensional space, our approach effectively reveals the underlying structure. For instance, in Fig. 8a, our method colors the pixels with two colors, showing that the pattern is not a continuous trend (marked as “A”, “B”, “C”) but is composed of different individual patterns. In contrast, the result of the line-based approach in Fig. 8b cannot reveal the actual components of the trend. Similarly, in Fig. 8c, our method provides a clear view of the trends of the two line bundles, achieving results that are similar to the line-based method (Fig. 8d). Finally, our approach successfully divides the pixels of the disconnected clusters (Fig. 8e), whereas the line-based method (Fig. 8f) fails to provide valuable information in this case. Section 5 of

the supplementary material, contains further comparative results based on real-world datasets.

5.2 Implementation and Interaction

Implementation. Results were obtained using an Apple Silicon M1 processor with 16GB RAM. Our interactive system (<https://color-line-density-plot.github.io/>) is written in Javascript and does not utilize the GPU. We used Firefox for the runtime tests. The time necessary to calculate the overlap coefficients is dependent on the size of the dataset, resulting in longer computation times for larger sets. For 10,000 lines, our hierarchical tree construction needs about 10s. Classifying all bins per division of clusters took less than 4s. The processing speed improves significantly with decreasing cluster sizes. In summary, our method can complete pre-processing (including threshold filtering, sampling, and hierarchical tree construction) for all datasets presented in this paper within 20s. Interactions to separate clusters took less than 4s, mostly between 1s and 2s.

Density threshold and sampling rate. Section 4.2 outlines the sampling strategy that requires setting density threshold and sampling rate. To enhance the intuitiveness of the minimum threshold, bins below the threshold are hidden using a corresponding slider. The user is informed when the selected bins exceed the processing capacity.

Cluster operations. Users can set the number of clusters as well as interactively separate clusters. After setting a number, clusters will be divided accordingly, based on the distance between the nodes of the hierarchical tree. Clicking on a cluster allows the user to split it into two sub-clusters. The feature vectors of these sub-clusters will then be mapped using circular MDS to find appropriate hue values. Users can select cluster colors that are fixed in the hue circle. Then, during gradient descent, only unfixed points are moved.

Hue adjusting and harmonization. In Section 4.5, we describe our use of multidimensional scaling (MDS) to ensure that the distance relationships between hues within each cluster correspond to their distance relationship in high-dimensional data space. Nonetheless, users can modify the hues assigned to the clusters themselves. Additionally, we provide users with eight harmonic templates, as defined by Cohen-Or *et al.* [12]. Once a template is selected, the hues for each cluster will be mapped to the corresponding area, with the movement of the dots being constrained accordingly.

6 MOTIVATIONAL USER STUDY

Here, we present the findings of a small motivational user study we performed to understand the ability of users to identify and trace trends in line-based density plots. We explored whether users have issues separating trends of lines. Additionally, we investigated if users are even aware of potential ambiguities in given plots. We did not include a comparison between our proposed colorization method and traditional density plots. We expected participants to perceive different trends overall, but as long as the line-based density plot remains visually similar, even if it is generated based on different underlying data, their perceptions should remain unchanged. We posed the following two hypotheses:

H1: There are different understandings when participants perceive trends in ambiguous line-based density plots.

H2: The participants are mostly unaware of the ambiguities in line-based density plots.

6.1 Experimental Design

We conducted a browser-based online user study with university students and Ph.D. candidates. A total of ten line datasets were included, all presented in the form of line-based density plots. Participants had to identify trends in line-based density plots and complete three tasks for each dataset. All plots used in the user study and the trends they are composed of are available in Section 3 of the supplementary material.

Tasks:

- *Tracing task.* The participants were asked to trace each dominant trend they identified.
- *Counting task.* Participants were asked to count the number of trends they perceived for each line-based density plot.

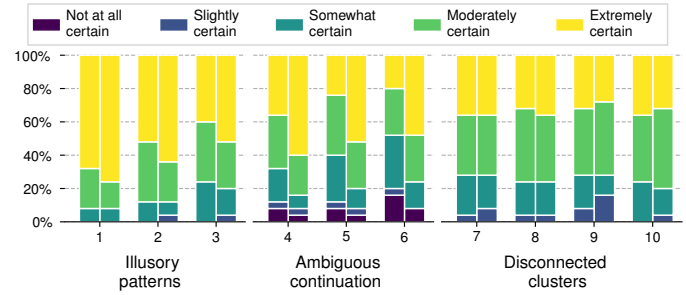


Fig. 9: Responses to the Likert scale of 25 participants for 10 ambiguous line-based density plots. The left bars represent the participant’s certainty about the number of trends they identified. The right bars represent the participant’s certainty about the trends they traced.

- *Certainty task.* We let participants rate how certain they were about their judgments on the number and tracing of trends.

Ambiguous pattern generation. We generated ten line datasets in total based on the three types of ambiguities introduced in Section 3.2. For each kind of ambiguous pattern, we provided multiple examples with varying degrees of ambiguity or different numbers of trends. In the first three datasets, we used illusory patterns, similar to Fig. 2a. The line-based density plot of the first dataset was composed of a horizontal trend superimposed on noise, which is similar to the right trends in the figure. For the following two datasets, we combined trends like the ones on the left in the top row of the figure. We increased the fan out in the center of each individual trend. Thus the plots were getting progressively less ambiguous, as seen in Fig. 10b. As the second kind of ambiguous pattern, we added ambiguous continuations to datasets 4 to 6. The combined patterns from Fig. 2b, both the upper left and the upper right composition, are used for datasets 4 and 5. In addition, we added another line dataset whose density plot was composed of three individual trends (dataset 6). The third kind of pattern showed disconnected clusters. In the last four datasets, the density plots contained potentially disconnected clusters. datasets 7 and 8 were composed of distinct patterns. On one hand, dataset 7 consisted of a continuous trend with a central fan-out region superimposed on low-density noise, as shown on the right side of the upper row in Fig. 2c. On the other hand, dataset 8 comprised two independent trends, which are illustrated in Fig. 10d. To decrease the ambiguity of the pattern, we increased the horizontal length of the high-density regions of the two independent trends to create datasets 9 and 10. This reduced the size of the central overlapping region as shown in Fig. 10e.

Experiment organization. We displayed the line-based density plots of the datasets colored using a multi-hue color map, one dataset per page, on our survey website. All participants had to fulfill the three tasks for each of the ten datasets in the same order. For the datasets with illusory patterns (datasets 1 to 3), because we were interested in whether participants could tell if multiple trends were combined, we asked them to trace only the main trends. For the datasets with ambiguous continuation (datasets 4 to 6), we asked participants to trace the individual trends they perceived with different color. For the disconnected clusters (datasets 7 to 10), we asked participants to mark the trends they thought were connected with the same color of strokes, and if they thought they were disconnected, use different colors.

Participants. In total, we collected inputs from 25 participants who completed our user study effectively. Three other participants were excluded as they either did not trace any trends or entered numbers of perceived trends that vastly differed from the trends they actually drew. We gathered the traces for all ten tasks drawn by each participant. The participants’ age was, on average, 26 years with a standard deviation of 4. The gender distribution among our participants was 56% male, 36% female, and 8 % preferring not to tell their gender.

Procedure. Each participant went through the following procedure: (1) completing a short questionnaire with questions about age, gender, and familiarity with line charts and density plots; (2) viewing an introduction about the tasks based on an already marked density plot; and (3) performing the tasks based on each dataset.

6.2 Experimental Results

Fig. 9 illustrates the participant’s responses using a Likert scale to the questions about their certainty regarding the number of trends and the tracings of the trends. Fig. 10 shows the distribution of the number of trends the participants perceived for each line-based density plot.

Illusory patterns. Most participants perceived a central horizontal trend in the density plots of the illusory patterns shown in Fig. 10a. The first dataset, consisting of only a horizontal trend superimposed by noise, was correctly identified by all participants and traced using a horizontal line. For all datasets with illusory patterns, more than 70% participants stated that they were “Moderately certain” or “Extremely certain” about the number of trends and their tracings. However, increasing the fan-out region of the combined trends in the illusory pattern, as shown in Fig. 10b, decreased the participant’s certainty about how many trends they perceived and how they traced them. Fig. 9 shows a decrease in the portion of participants being “Extremely certain” about the number of trends they identified and how they traced them from the first dataset to the third. Additionally, the portion of participants perceiving at least two trends increased from 0% to 16%, as shown in Fig. 10. This indicates that participants were increasingly unsure about the trends in the density plots. Still, even though datasets 2 and 3 were created from data containing multiple trends, participants mostly perceived only a single trend for all datasets (92% and 84%, respectively), which supports **H2**.

Ambiguous continuation. All participants correctly identified, with moderate to high certainty, that the combined line-based density plots of the individual ambiguous pattern shown in Fig. 2b are composed of two trends. For dataset 6, consisting of three trends, all participants except for one correctly identified the number of trends. In the tracing task, 88% of the participants identified the same trends in both datasets 4 and 5. 60% of the participants traced the combination of two “U”-shaped trends in both plots. While 28% of the participants traced crossing trends for both plots in an “X” shape. For datasets 4 and 5, whose density plots are almost identical, we expect participants to perceive them ambiguously. For dataset 4, 16 participants identified a “U”-shaped pattern and 9 an “X”-shaped. For dataset 5, 17 participants traced a “U” and 8 an “X”. For dataset 6, which contains a combination of three trends, the participants’ understandings are more diverse. We only offered the ability to trace a single trend in the data. Participants were aware that there might be multiple possible trends in the plot but drew only the most dominant one. Therefore, the results do not show whether participants potentially perceived multiple patterns. The distribution of participants among both identified trends indicates that there was no consensus on which trend was dominantly perceived. We conclude that this supports **H1** with respect to the group of participants, but not for an individual participant.

Disconnected clusters. While participants largely agreed on the number of perceived trends in the previous patterns, they were torn about whether the disconnected clusters in Fig. 2c contained one or two trends. For dataset 7 56% and for dataset 10 64% of the participants indicated that the trends were independent, like in the left example of Fig. 2c. The remaining portion of participants highlighted just one continuous trend, as shown on the right. Thus participants did not agree, while over 80% of them were at least moderately certain about the number of patterns they identified and their tracings. Fig. 10c shows the disagreement of the participants regarding the number of trends for datasets 7-10. As the density plots of datasets 7 and 8 are almost identical but contain different trends, we expect participants to give similar answers. For dataset 7, 44% of the participants answered one trend, and 56% of the participants answered two. While for dataset 8, it is 40% and 60%, respectively. For datasets 9 and 10, the participation response percentages are still similar, although they are less ambiguous. This further supports **H1**.

6.3 Experimental Findings

While our user study was meant to exemplarily show that the user’s understanding of line-based density plots is ambiguous, it is not a large-scale study with extensive possibilities for statistical evaluations. Nevertheless, it supports our hypothesis **H1** and **H2** and shows trends in

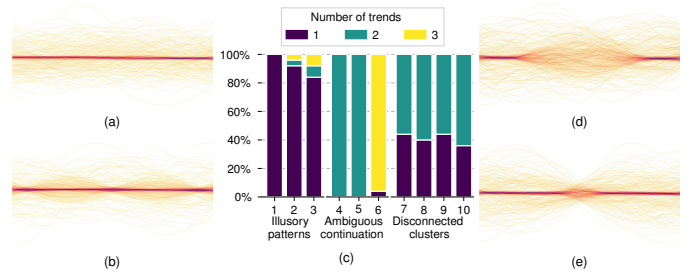


Fig. 10: Line-based density plots with illusory trends with varying fan-out of lines from (a) to (b) and disconnected clusters with varying separation from (d) to (e). (c) Distribution of the number of perceived trends in 10 line-based density plots.

line-based density plots are perceived ambiguously. While participants can correctly identify the number of trends for the patterns that showed ambiguous continuation, they disagreed about the flow of the trends. Our results indicate that although there is a great deal of disagreement among the participants regarding the underlying trends, the individual participants seemed to be fairly certain about the number of trends and their tracing. Thus it is important to support users in identifying individual trends and avoiding wrong perceptions.

7 DISCUSSION AND CONCLUSION

In this paper, we presented a novel coloring method for line-based density plots. Although our approach is an image-based analysis of line data, it could be applied to density plots of parallel coordinates or even other types of data. In contrast to the in Section 2 introduced line-based methods, which group similar lines by colorization of the lines, abstract them with statistical information of ensembles, or combine them by altering their paths using edge-bundling, we aggregate the individual line into line-based density plots and enrich them by clustering similar regions in the image space.

Future works and limitations. The basic unit of data for our presented method is pixels, which means the higher the resolution of the visualization, the more computationally demanding it is. However, it is unnecessary to perform binning for each pixel. Different bin sizes can be used in different areas depending on the distribution of the data. For example, a larger bin size can be used for unimportant areas. So the plot could be divided into differently-sized bins to reduce the number of bins. Therefore, the next step would be to explore automatic methods for finding the best binning strategy. In subsequent works, we also want to reduce the number of clustering errors by investigating other sampling strategies.

Edge bundling is a widely used approach to address the clutter in line-based visualizations. Recent methods address the issue of ambiguities introduced by edge bundling. Our method could also be applied to disambiguate further the results generated by edge bundling methods.

With an increasing number of clusters, the angle between assigned colors on the hue wheel decreases, which decreases their distinguishability. Already for 7 different colors (see Fig. 6b), the colors start to be barely distinguished. Furthermore, perception depends on cluster sizes, spacing of hues, and positional relationships between clusters. Conducting a user study could shed light on these aspects.

Our small-scale motivational study was only concerned with showing whether and to what extent the perception of trends in line-based density plots is ambiguous. It encompassed only university students and Ph.D. candidates. But it underpinned our hypothesis that line-based density plots are indeed perceived ambiguously by users. We investigate only whether the user’s understanding of line-based density plots is ambiguous. Therefore, we did not include a comparison of line-based coloring methods and our method in the user study. A further study should have a larger and more diverse population, including perception and understanding of our results.

We currently summarize three types of ambiguities in line-based density plots, but more types may exist. In future work, we plan to propose a more comprehensive taxonomy to classify the ambiguities in line-based density plots.

ACKNOWLEDGMENTS

This work was supported in part by Deutsche Forschungsgemeinschaft (DFG) Project 410883423, Project 251654672 – TRR 161 “Quantitative methods for visual computing”, KE 740/17-2 of the FOR2111 “Questions at the Interface”, the National Key R&D Program of China (2022ZD0160805), NSF China (No.62132017, 62141217), and the Shandong Provincial Natural Science Foundation (No. ZQ2022JQ32).

REFERENCES

- [1] Acis web services. http://www.rcc-acis.org/docs_webservices.html, 2017. Accessed: 2023-06-30. 5
- [2] New york stock exchange. <https://www.nyse.com/market-data/historical>, 2023. Accessed: 2023-06-30. 7
- [3] A. Artero, M. de Oliveira, and H. Levkowitz. Uncovering clusters in crowded parallel coordinates visualizations. In *IEEE Symposium on Information Visualization*, pp. 81–88, 2004. doi: 10.1109/INFVIS.2004.682
- [4] L. Bergman, B. Rogowitz, and L. Treinish. A rule-based tool for assisting colormap selection. In *Proceedings Visualization '95*, pp. 118–125, 1995. doi: 10.1109/VISUAL.1995.480803 3
- [5] C. A. BREWER. Chapter 7 - color use guidelines for mapping and visualization. In A. M. MACEACHREN and D. F. TAYLOR, eds., *Visualization in Modern Cartography*, vol. 2 of *Modern Cartography Series*, pp. 123–147. Academic Press, 1994. doi: 10.1016/B978-0-08-042415-6.50014-4 3
- [6] C. A. Brewer. Guidelines for use of the perceptual dimensions of color for mapping and visualization. In J. Bares, ed., *Color Hard Copy and Graphic Arts III*, vol. 2171, pp. 54 – 63. International Society for Optics and Photonics, SPIE, 1994. doi: 10.1117/12.175328 3
- [7] C. A. Brewer et al. Color use guidelines for data representation. In *Proceedings of the Section on Statistical Graphics, American Statistical Association*, pp. 55–60, 1999. 3
- [8] D. B. Carr, R. J. Littlefield, W. L. Nicholson, and J. S. Littlefield. Scatterplot matrix techniques for large n. *Journal of the American Statistical Association*, 82(398):424–436, 1987. doi: 10.1080/01621459.1987.10478445 2
- [9] C.-K. Chen, S. Yan, H. Yu, N. Max, and K.-L. Ma. An illustrative visualization framework for 3d vector fields. *Computer Graphics Forum*, 30(7):1941–1951, 2011. doi: 10.1111/j.1467-8659.2011.02064.x 2
- [10] W. Chen, F. Guo, and F.-Y. Wang. A survey of traffic data visualization. *IEEE Transactions on Intelligent Transportation Systems*, 16(6):2970–2984, 2015. doi: 10.1109/ITITS.2015.2436897 1
- [11] Y. Chen, J. Cohen, and J. Krolik. Similarity-guided streamline placement with error evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1448–1455, 2007. doi: 10.1109/TVCG.2007.70595 2
- [12] D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y.-Q. Xu. Color harmonization. *ACM Trans. Graph.*, 25(3):624–630, 2006. doi: 10.1145/1141911.1141933 3, 8
- [13] M. A. A. Cox and T. F. Cox. *Multidimensional Scaling*, pp. 315–347. Springer, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-33037-0_14 5
- [14] T. F. Cox and M. A. Cox. Multidimensional scaling on a sphere. *Communications in Statistics - Theory and Methods*, 20(9):2943–2953, 1991. doi: 10.1080/03610929108830679 5
- [15] W. H. Day and H. Edelsbrunner. Efficient algorithms for agglomerative hierarchical clustering methods. *Journal of classification*, 1(1):7–24, 1984. doi: 10.1007/BF01890115 4
- [16] D. Feng, L. Kwock, Y. Lee, and R. Taylor. Matching visual saliency to confidence in plots of uncertain data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):980–989, 2010. doi: 10.1109/TVCG.2010.176 2
- [17] F. Ferstl, K. Bürger, and R. Westermann. Streamline variability plots for characterizing the uncertainty in vector field ensembles. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):767–776, 2016. doi: 10.1109/TVCG.2015.2467204 1, 2, 7
- [18] A. Frantzis, R. Leaper, P. Alexiadou, A. Prospathopoulos, and D. Lekkas. Hellenic trench ais data. 2018. doi: 10.17882/57040 5, 6, 7
- [19] A. Frantzis, R. Leaper, P. Alexiadou, A. Prospathopoulos, and D. Lekkas. Shipping routes through core habitat of endangered sperm whales along the hellenic trench, greece: Can we reduce collision risks? *PLOS ONE*, 14(2):1–21, 2019. doi: 10.1371/journal.pone.0212016 7
- [20] B. D. Fulcher, M. A. Little, and N. S. Jones. Highly comparative time-series analysis: the empirical structure of time series and their methods. *Journal of the Royal Society Interface*, 10(83):20130048, 2013. doi: 10.1098/rsif.2013.0048 1
- [21] S. Gaffney and P. Smyth. Joint probabilistic curve clustering and alignment. In *Proceedings of the 17th International Conference on Neural Information Processing Systems, NIPS'04*, p. 473–480. MIT Press, Cambridge, MA, USA, 2004. 2
- [22] C. Healey. Choosing effective colours for data visualization. In *Proceedings of Seventh Annual IEEE Visualization '96*, pp. 263–270, 1996. doi: 10.1109/VISUAL.1996.568118 3
- [23] H. Hochheiser and B. Shneiderman. Dynamic query tools for time series data sets: Timebox widgets for interactive exploration. *Information Visualization*, 3(1):1–18, 2004. doi: 10.1057/palgrave.ivs.9500061 1, 2
- [24] D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):741–748, 2006. doi: 10.1109/TVCG.2006.147 2
- [25] D. Holten and J. J. Van Wijk. Force-directed edge bundling for graph visualization. *Computer Graphics Forum*, 28(3):983–990, 2009. doi: 10.1111/j.1467-8659.2009.01450.x 2
- [26] M.-H. Hong, J. K. Witt, and D. A. Szafir. The weighted average illusion: Biases in perceived mean position in scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):987–997, 2022. doi: 10.1109/TVCG.2021.3114783 2
- [27] K. Hornik, F. Leisch, and A. Zeileis. Colour for presentation graphics. In *Proceedings of DSC*, p. 2, 2003. 3, 5
- [28] C. Hurter, B. Tissoires, and S. Conversy. Fromdady: Spreading aircraft trajectories across views to support iterative queries. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1017–1024, 2009. doi: 10.1109/TVCG.2009.145 2
- [29] A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: A review. *ACM Comput. Surv.*, 31(3):264–323, 1999. doi: 10.1145/331499.331504 4
- [30] W. Javed, B. McDonnell, and N. Elmqvist. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):927–934, 2010. doi: 10.1109/TVCG.2010.162 1
- [31] D. F. Jerding and J. T. Stasko. The information mural: a technique for displaying and navigating large information spaces. *IEEE Transactions on Visualization and Computer Graphics*, 4(3):257–271, 1998. doi: 10.1109/2945.722299 2
- [32] G. Kindlmann, E. Reinhard, and S. Creem. Face-based luminance matching for perceptual colormap generation. In *IEEE Visualization*, pp. 299–306, 2002. doi: 10.1109/VISUAL.2002.1183788 3
- [33] G. Kindlmann and C. Scheidegger. An algebraic process for visualization design. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2181–2190, 2014. doi: 10.1109/tvcg.2014.2346325 2
- [34] O. D. Lampe and H. Hauser. Curve Density Estimates. *Computer Graphics Forum*, 2011. doi: 10.1111/j.1467-8659.2011.01912.x 1, 2
- [35] S. Lee, M. Sips, and H.-P. Seidel. Perceptually driven visibility optimization for categorical data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(10):1746–1757, 2013. doi: 10.1109/TVCG.2012.315 3
- [36] A. Lhuillier, C. Hurter, and A. Telea. State of the art in edge and trail bundling techniques. *Computer Graphics Forum*, 36(3):619–645, 2017. doi: 10.1111/egf.13213 5
- [37] T. Liu, X. Li, C. Bao, M. Correll, C. Tu, O. Deussen, and Y. Wang. Data-driven mark orientation for trend estimation in scatterplots. In *Proceedings of the CHI Conference on Human Factors in Computing Systems, CHI '21*. ACM, New York, 2021. doi: 10.1145/3411764.3445751 2
- [38] K. Lu, M. Feng, X. Chen, M. Sedlmair, O. Deussen, D. Lischinski, Z. Cheng, and Y. Wang. Palettailor: Discriminable colorization for categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):475–484, 2021. doi: 10.1109/TVCG.2020.3030406 3
- [39] J. Matejka, F. Anderson, and G. Fitzmaurice. Dynamic opacity optimization for scatter plots. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, p. 2707–2710. ACM, New York, 2015. doi: 10.1145/2702123.2702585 2
- [40] A. Mayorga and M. Gleicher. Splatterplots: Overcoming overdraw in scatter plots. *IEEE Transactions on Visualization and Computer Graphics*, 19(9):1526–1538, 2013. doi: 10.1109/TVCG.2013.65 3
- [41] T. McLoughlin, M. W. Jones, R. S. Laramée, R. Malki, I. Masters, and C. D. Hansen. Similarity measures for enhancing interactive streamline seeding. *IEEE Transactions on Visualization and Computer Graphics*, 19(8):1342–1353, 2013. doi: 10.1109/TVCG.2012.150 2

- [42] M. Mirzargar, R. T. Whitaker, and R. M. Kirby. Curve boxplot: Generalization of boxplot for ensembles of curves. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2654–2663, 2014. doi: 10.1109/TVCG.2014.2346455 1, 2, 3
- [43] K. Moreland. Diverging color maps for scientific visualization. In *International Symposium on Visual Computing*, pp. 92–103. Springer, 2009. doi: 10.1007/978-3-642-10520-3_9 3
- [44] D. Moritz and D. Fisher. Visualizing a million time series with the density line chart, 2018. doi: 10.48550/arXiv.1808.06019 1, 2, 3, 5
- [45] M. Müller. *Dynamic Time Warping*, pp. 69–84. Springer, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-74048-3_4 2
- [46] M. Novotny and H. Hauser. Outlier-preserving focus+context visualization in parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):893–900, 2006. doi: 10.1109/TVCG.2006.170 2
- [47] OpenStreetMap contributors. Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org>, 2017. 11
- [48] G. Palmas, M. Bachynskiy, A. Oulasvirta, H. P. Seidel, and T. Weinkauff. An edge-bundling layout for interactive parallel coordinates. In *2014 IEEE Pacific Visualization Symposium*, pp. 57–64, 2014. doi: 10.1109/PacificVis.2014.40 2
- [49] F. Petitjean, A. Ketterlin, and P. Gançarski. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recogn.*, 44(3):678–693, 2011. doi: 10.1016/j.patcog.2010.09.013 2
- [50] B. Pham. Spline-based color sequences for univariate, bivariate and trivariate mapping. In *Proceedings of the First IEEE Conference on Visualization: Visualization '90*, pp. 202–208, 1990. doi: 10.1109/VISUAL.1990.146383 3
- [51] D. Pomeranke, F. L. Dennig, D. A. Keim, J. Fuchs, and M. Blumenschein. Slope-dependent rendering of parallel coordinates to reduce density distortion and ghost clusters. In *2019 IEEE Visualization Conference (VIS)*, pp. 86–90, 2019. doi: 10.1109/VISUAL.2019.8933706 2, 3
- [52] P. L. Rheingans. Task-based color scale design. In W. R. Oliver, ed., *28th AIPR Workshop: 3D Visualization for Data Exploration and Decision Making*, vol. 3905, pp. 35–43. International Society for Optics and Photonics, SPIE, 2000. doi: 10.1117/12.384882 3
- [53] P. K. Robertson and J. F. O’Callaghan. The generation of color sequences for univariate and bivariate mapping. *IEEE Computer Graphics and Applications*, 6(2):24–32, 1986. doi: 10.1109/MCG.1986.276688 3
- [54] B. Rogowitz and A. Kalvin. The "which blair project": a quick visual method for evaluating perceptual color maps. In *Proceedings Visualization, 2001. VIS '01.*, pp. 183–556, 2001. doi: 10.1109/VISUAL.2001.964510 3
- [55] C. Rossel and H. Theisel. Streamline embedding for 3d vector field exploration. *IEEE Transactions on Visualization and Computer Graphics*, 18(3):407–420, 2012. doi: 10.1109/TVCG.2011.78 2
- [56] T. Salzbrunn and G. Scheuermann. Streamline predicates. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1601–1612, 2006. doi: 10.1109/TVCG.2006.104 2
- [57] R. Scheepens, N. Willems, H. van de Wetering, G. Andrienko, N. Andrienko, and J. J. van Wijk. Composite density maps for multivariate trajectories. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2518–2527, 2011. doi: 10.1109/TVCG.2011.181 2
- [58] R. Scheepens, N. Willems, H. van de Wetering, and J. J. van Wijk. Interactive visualization of multivariate trajectory data with density maps. In *2011 IEEE Pacific Visualization Symposium*, pp. 147–154, 2011. doi: 10.1109/PACIFICVIS.2011.5742384 2
- [59] B. W. Silverman. *Density estimation for statistics and data analysis*, vol. 26. CRC press, 1986. 2
- [60] M. Tennekes and E. de Jonge. Tree colors: Color schemes for tree-structured data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2072–2081, 2014. doi: 10.1109/TVCG.2014.2346277 3
- [61] D. Thompson, J. Bennett, C. Seshadhri, and A. Pinar. A provably-robust sampling method for generating colormaps of large data. In *2013 IEEE Symposium on Large-Scale Data Analysis and Visualization (LDAV)*, pp. 77–84, 2013. doi: 10.1109/LDAV.2013.6675161 3
- [62] M. Thöny and R. Pajarola. Vector map constrained path bundling in 3d environments. In *Proceedings of the 6th ACM SIGSPATIAL International Workshop on GeoStreaming, IWGS '15*, pp. 33–42. ACM, New York, 2015. doi: 10.1145/2833165.2833168 2
- [63] C. Tominski, G. Fuchs, and H. Schumann. Task-driven color coding. In *2008 12th International Conference Information Visualisation*, pp. 373–380, 2008. doi: 10.1109/IV.2008.24 3
- [64] T. Trautner, F. Bolte, S. Stoppel, and S. Bruckner. Sunspot Plots: Model-based Structure Enhancement for Dense Scatter Plots. *Computer Graphics Forum*, 2020. doi: 10.1111/cgf.14001 2
- [65] B. E. Trumbo. A theory for coloring bivariate statistical maps. *The American Statistician*, 35(4):220–226, 1981. 3
- [66] M. van der Zwan, V. Codreanu, and A. Telea. Cubu: Universal real-time bundling for large graphs. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2550–2563, 2016. doi: 10.1109/TVCG.2016.2515611 2
- [67] M. Vijaymeena and K. Kavitha. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal*, 3:19–28, 03 2016. doi: 10.5121/mlaij.2016.3103 4
- [68] M. Wallinger, D. Archambault, D. Auber, M. Nöllenburg, and J. Peltonen. Edge-path bundling: A less ambiguous edge bundling approach. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):313–323, 2022. doi: 10.1109/TVCG.2021.3114795 2
- [69] M. Wallinger, D. Archambault, D. Auber, M. Nöllenburg, and J. Peltonen. Faster edge-path bundling through graph spanners. *Computer Graphics Forum*, 2023. doi: 10.1111/cgf.14789 2
- [70] L. Wang and A. Kaufman. Importance driven automatic color design for direct volume rendering. *Computer Graphics Forum*, 31(3):1305–1314, 2012. doi: 10.1111/j.1467-8659.2012.03123.x 3
- [71] M. Wattenberg. Sketching a graph to query a time-series database. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '01, p. 381–382. ACM, New York, 2001. doi: 10.1145/634067.634292 2
- [72] M. Wertheimer. Laws of organization in perceptual forms. *A source book of Gestalt psychology*, pp. 71–88, 1938. doi: 10.1037/11496-005 3
- [73] R. T. Whitaker, M. Mirzargar, and R. M. Kirby. Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2713–2722, 2013. doi: 10.1109/TVCG.2013.143 1
- [74] H. Wickham. Bin-summarise-smooth: a framework for visualising large data. *had. co. nz, Tech. Rep*, 2013. 2, 3, 4
- [75] H. Yu, C. Wang, C.-K. Shene, and J. H. Chen. Hierarchical streamline bundles. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1353–1367, 2012. doi: 10.1109/TVCG.2011.155 2
- [76] J. Yuan, Y. Zheng, X. Xie, and G. Sun. Driving with knowledge from the physical world. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11*, p. 316–324. ACM, New York, 2011. doi: 10.1145/2020408.2020462 5, 7
- [77] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang. T-drive: Driving directions based on taxi trajectories. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS '10*, p. 99–108. ACM, New York, 2010. doi: 10.1145/1869790.1869807 5, 7
- [78] A. Zeileis, K. Hornik, and P. Murrell. Escaping rgbland: Selecting colors for statistical graphics. *Computational Statistics & Data Analysis*, 53(9):3259–3270, 2009. doi: 10.1016/j.csda.2008.11.033 3, 5
- [79] W. Zeng, Q. Shen, Y. Jiang, and A. Telea. Route-aware edge bundling for visualizing origin-destination trails in urban traffic. *Computer Graphics Forum*, 38(3):581–593, 2019. doi: 10.1111/cgf.13712 2
- [80] Z. Zhang, K. Huang, and T. Tan. Comparison of similarity measures for trajectory clustering in outdoor surveillance scenes. In *18th International Conference on Pattern Recognition (ICPR'06)*, vol. 3, pp. 1135–1138, 2006. doi: 10.1109/ICPR.2006.392 2
- [81] Y. Zhao, Y. Wang, J. Zhang, C.-W. Fu, M. Xu, and D. Moritz. Kd-box: Line-segment-based kd-tree for interactive exploration of large-scale time-series data. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):890–900, 2022. doi: 10.1109/TVCG.2021.3114865 1, 2, 5
- [82] L. Zhou and C. D. Hansen. A survey of colormaps in visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(8):2051–2069, 2016. doi: 10.1109/TVCG.2015.2489649 3
- [83] M. Zinsmaier, U. Brandes, O. Deussen, and H. Strobel. Interactive level-of-detail rendering of large graphs. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2486–2495, 2012. doi: 10.1109/TVCG.2012.238 2

A FIGURE CREDITS

For Fig. 6, the gray background image of Greece was taken from OpenStreetMaps [47] at location <https://www.openstreetmap.org/#map=7/37.055/22.313&layers=T>.