Eurographics Conference on Visualization (EuroVis) 2019
M. Gleicher, H. Leitte, and I. Viola
(Guest Editors)

Volume 38 (2019), Number 3

Bridging the Data Analysis Communication Gap Utilizing a Three-Component Summarized Line Graph

C. Yau,[1] M. Karimzadeh,[1] C. Surakitbanharn,[2] N. Elmqvist,[3] and D. S. Ebert[1]

[1]Purdue University, [2]Stanford University, [3]University of Maryland, College Park

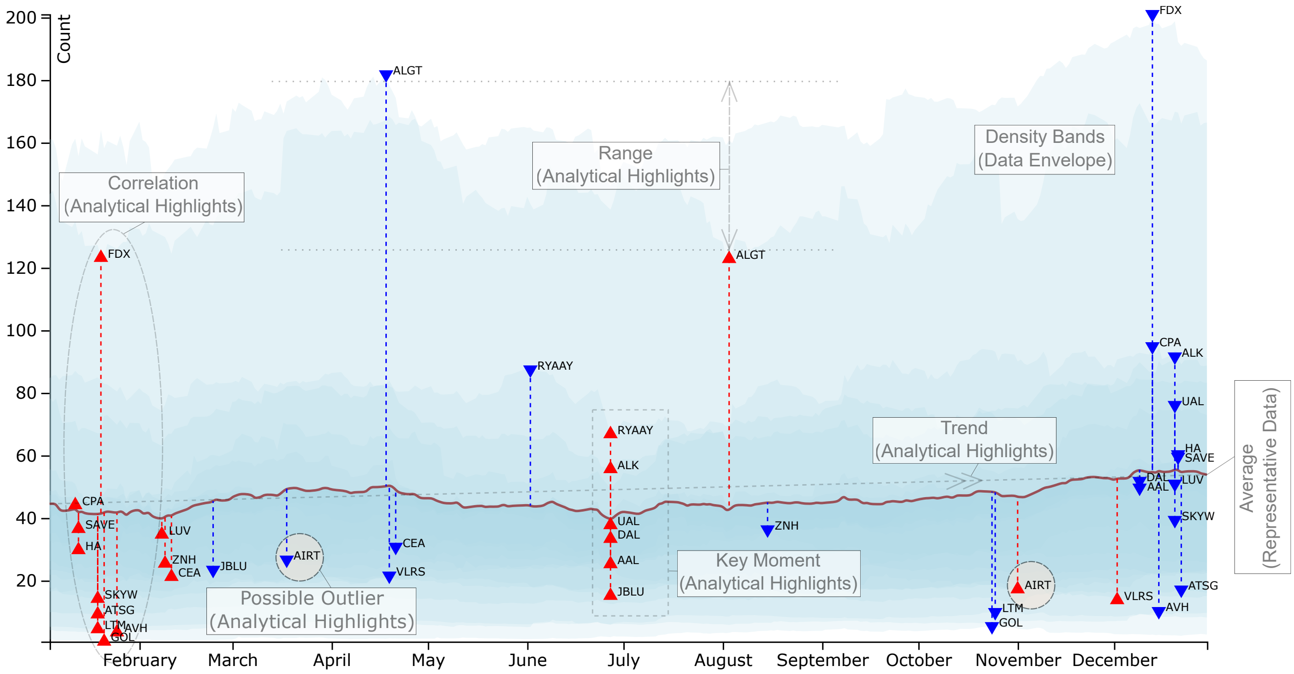

Figure 1: *A three-component summarized line graph showing Nasdaq stock prices of the transportation industry in the air freight/delivery service during 2016 (21 stocks over a year). The tan line is the representative data: an average curve providing the mean value for the entire summarized dataset. Along the time axis, analytical highlights are shown as ranges, trends, correlations, outliers, and key moments called out using dotted lines and triangles; red triangles represent the absolute minimums of each line, and blue triangles the absolute maximums. Finally, the light blue bands in the background provide the data envelope representing the data distribution over the entire time.*

Abstract

Communication-minded visualizations are designed to provide their audience—managers, decision-makers, and the public—with new knowledge. Authoring such visualizations effectively is challenging because the audience often lacks the expertise, context, and time that professional analysts have at their disposal to explore and understand datasets. We present a novel summarized line graph visualization technique designed specifically for data analysts to communicate data to decision-makers more effectively and efficiently. Our summarized line graph reduces a large and detailed dataset of multiple quantitative time-series into (1) representative data that provides a quick takeaway of the full dataset; (2) analytical highlights that distinguish specific insights of interest; and (3) a data envelope that summarizes the remaining aggregated data. Our summarized line graph achieved the best overall results when evaluated against line graphs, band graphs, stream graphs, and horizon graphs on four representative tasks.

CCS Concepts

• ***Human-centered computing*** → ***Visualization techniques; Information visualization;***

1. Introduction

Understanding and analyzing complex datasets using visual analytics typically involves a long sequence of interactive steps, such as attaining an overview, zooming, and filtering [Shn96]. This kind of visual exploration, while proven to be effective for analysts to gain first-hand understanding and insights into data, does not capture this full analysis history to individuals who did not participate in the exploration. For example, a manager may be making a decision based on the information prepared and presented by her data analysts. Without the context of the full analysis session, including hypotheses tested and discarded, the manager has limited understanding of the overall findings, no way to detect potential biases, and no recourse to check the work herself. This gap in the understanding of the data could prevent the manager from making optimal decisions.

In this paper, we draw on the concept of communication-minded visualization [VW06] to bridge this gap between analysts and their audience (e.g., managers, decision-makers, the public) by incorporating contextual knowledge into visual summaries. To achieve this, we first have to understand how decision-makers require data analysts to present this knowledge. We surveyed decision-makers from the first responder community (e.g., public safety, police, rescue) on how data is prepared and presented to them by data analysts. We used this survey to compile a list of requirements for communicating data to increase efficiency, improve understanding, and reduce bias:

R1 **Comprehensibility:** Understanding the overall data behavior.
R2 **Accuracy:** Identifying insights that are important to a decision.
R3 **Fidelity:** Faithfulness to the original raw data.
R4 **Precision:** Obtaining actual data values.
R5 **Comparison:** Comparing significant aspects of the data.

Based on these requirements, we designed a visualization approach utilizing three visual components: (1) representative data, (2) analytical highlights, and (3) data envelope. The representative data provides a takeaway of the entire dataset, the analytical highlights allow the audience to identify insights of interest with ease, and the data envelope summarizes the remaining aggregated data to allow exploration of the raw statistics. Since our target audience often deals with multiple quantitative time-series, we focus our efforts on temporal visualization. Thus, our design approach is based on the line graph as it is one of the most frequently used visualization techniques [Cle93] and can be compared against many other time-series visualization techniques. The result is our *summarized line graph*.

As annotated in Figure 1, our summarized line graph features the mean as representative data; ranges, trends, correlations, outliers and key moments as analytical highlights; and extent, density, and traces as the data envelope. We have compared the summarized line graph against traditional line graphs, band graphs, stream graphs, and horizon graphs in a user study on both their complexity and ability to meet the requirements listed above. Our study measured performance in terms of both accuracy and completion time for four representative tasks drawn from the requirements: identifying the original data, the overall trend, the outliers, and the key moments in the time-series. While our results indicate that our summarized line graph does not outperform other techniques in every task, it achieves the best overall result when all four tasks are considered. We also see potential in bringing this three-component summarization design approach to visualizing other data types in future work.

2. Survey: Communicating Insights to Stakeholders

Our work was inspired by regular interactions with decision-makers. Based on their feedback, we designed a survey to understand the problems inherent in communicating insights from data analysis.

2.1. Method

To better understand how data is communicated between data analysts and decision-makers, we surveyed six decision-makers from the first responder community. Our participants, who represent our primary target user community, are decision-makers at different levels in police and public safety departments. Of course, with such a limited sample from one specific group, the result may not represent all practice; we have, however, found similar needs in a survey of financial analysts [KCA*16]: providing context, supporting analyses, allowing comparisons of the details, etc. From our surveys, we identified a few key points in the current practices that identify the limitations of this communication process between data analysts and decision-makers and derived the requirements of our design in Section 1. These results support the feedback we have received from many decision-makers in the past.

2.2. Findings

First, common among all respondents was that *data analysts are often given only limited time to present their findings*, often limited to five minutes or less. As a result, the data analysts can only communicate a limited amount of data. This ties into the requirements of our visual design **R1** (comprehensibility) and **R2** (accuracy), where the decision-makers must be able to quickly and efficiently reach an understanding of the dataset and its highlights. This time limitation also forces the presentation to focus on the dataset and the insights instead of allowing a full walkthrough of the exploration process.

We also found that *decision-makers only have limited exposure to the data*. Due to the limited amount of time, analysts are not likely to walk the decision-maker through the entire dataset. From our surveys, none of the decision-makers are always presented with raw statistics, and one-third of their data analysts almost never include raw statistics in their presentation. On the other hand, all of the decision-makers acknowledge the positive impact on their decision-making of seeing raw statistics and like to see the raw statistics at one point or another. Two of the six decision-makers actually wanted to see raw statistics at all times. It is, of course, impractical to present all the data with the limited amount of time, but as one of the decision-makers stated: "[The data analysts] have [the data] in volume. I need it in highlights with the ability to ask for more." The decision-makers should be given the ability to understand the raw statistics better, which ties into our requirement **R3** (fidelity).

Data analysts and the decision-makers understand data differently. One of the decision-makers stated that "[data analysts] tend to focus on the manner in which data is captured whereas [decision-makers] tend to focus on the story the data is telling." While this may not hold true for all data analysts, it is not surprising that important details may be lost during the filtering process because the data analysts have a different focus in mind while exploring and preparing the data for presentation. For example, when presenting a dataset with just the average and standard deviation, anomalous spikes in a specific data source that lasted only a short amount of time can easily be overlooked and not presented. This ties into our requirements **R2** (accuracy) and **R3** (fidelity) in which decision-makers must be able to identify insights that are important to their decision and explore the data to a certain extent, rather than leaving decisions on the level of importance entirely up to the analysts.

Additionally, *information bias can exist in the data as presented.* Five of the six decision-makers have experienced situations where presented information appeared to be biased toward a decision. One decision-maker noted the biased information could sometimes point toward what the data analysts viewed as the preferred courses of action based on their exploration of the data. However, as stated in the work of Ajzen et al. [ABR96], personal relevance could affect what is viewed as the preferred outcome. This leads us back to our requirement **R1** (comprehensibility), **R2** (accuracy), and **R3** (fidelity); the decision-makers must be able to obtain an understanding of the dataset that can minimize the information bias from the information presented from the data analysts.

Also, *data influences real-world decision making*. All survey participants acknowledged that the data sometimes affect their decision outcome. It is important for the data to be measurable and comparable to allow the decision-makers to link the data to additional real-world variables for decision-making. This ties into our requirements **R4** (precision) and **R5** (comparison) where even in a summarized visualization, decision-makers must be able to measure and compare significant factors relevant to their decisions.

Finally, *the presentation can be limited by its medium*. With limited time and different settings, analysts are sometimes limited to presenting data using static images. Presenting such static charts means the final display should be self-contained, i.e., it should incorporate all insights without the need for interaction or animation. By summarizing the dataset visually and including noteworthy insights, we can reduce the cluttering on screen and the dependency on interaction or animation, allowing the decision-makers to retrieve the same information more efficiently even with just static images.

2.3. Corollary: The Communication Gap

Findings from our survey aptly confirm prior art on what has come to be informally known as the “communication gap” in visualization and visual analytics [TC05, VW06]: that while Pirolli and Card’s sensemaking loop [PC05] includes presentation as an integrated part of analysis, current tools often distinguish between the phases (e.g., stories vs. worksheets and dashboards in Tableau). There is clearly a need for visualization techniques intended for casual experts that bridge this gap by both summarizing large-scale data as well as providing in-depth analytical insights for representative tasks.

3. Related Work

Here we review relevant prior art on the topic of this paper.

3.1. Data Communication

Our work falls in between the traditional use of visual analytics and that of casual information visualization [PSM07] for the purposes of communicating data between data analysts and their final audiences. Our work shares some of the characteristics of casual information visualization such as targeting audiences who are not experts of visualizations and the design challenge of modifying the design for different users, data, and insights needed. However, our work also targets specific work tasks rather than everyday tasks. Our audiences, closer to what are referred to as casual experts [MEV*14], may not be trained in understanding visualization, but should have the domain expertise in the data presented and the problem in question.

Viegas and Wattenberg [VW06] introduced communication-minded visualization to support communication and collaborative analysis through visualization designs. Inspired by their work, we focus on solving the communication gap between data analysts and their audiences through novel visualization techniques.

Segel and Heer [SH10] suggested design strategies for narrative visualization to tell data stories. Their work discussed the importance of balancing author-driven and reader-driven stories; our work shares characteristics of both. Segal and Heer also suggested that storytelling of data is most effective when there is constrained interaction. However, as described in Section 2, the scope of this work focuses on communicating data when interaction is not available.

Hullman and Diakopoulos [HD11] presented a narrative visualization framework based on rhetoric visualizations to tell data stories more effectively using a combination of visual representations, annotations, and interactivity layers in the design. This work follows a similar storytelling strategy by overlaying the overview and analysis result on top of the remaining aggregated data.

Some visualization techniques also present multiple data types to complement the story telling. A common combination used to tell stories with more context utilizes numeric and textual data. Textual data are often overlaid on traditional numeric visualization techniques as (interactive) annotations [BMW17, HDA13] to give reasons behind the numeric data behavior. The approach is more effective when displaying a comparatively smaller amount of contextual data, as it can suffer from scalability issues.

3.2. Data Abstraction

Descriptive statistics are common and well-developed measures for summarizing scalar data, quantitatively describing features of a collection of information [Man10]. Mean and standard deviation can describe a snapshot of the entire dataset providing a representative value and a basis for interpreting the data through distribution characteristics. In contrast, our work provides a more detailed view of the subsets of the data as well as the change over time.

Sampling is also often applied to summarize large data [BC94]. However, when used in time-series data, sampling typically focuses on summarizing the time-axis rather than the multiple time-series, which is different from the focus of this work. Another issue with data sampling, similar to segmentation, dimensionality reduction, and clustering [DOL03], is that it treats all data equally and does not evaluate what information is removed in the process, whereas we include data that is often overlooked.

Visual analytics is often used for big data analysis [WT04], typically aided by visual summarization [CAFG12]. Aside from specific visualization designs, there also exist techniques that allow the audience to better perceive the data such as hierarchical aggregation [EF10] and aspect ratio adjustment [Cle93, HA06].

3.3. Time-Series Visualization

To first test our design approach, we apply our three-component visual summary to the most frequently used graphic design: the time-series plot [Tuf83]. In this subsection, we compare and contrast several time-series visualization techniques relevant to our design.

Familiar to most users, the line graph designed by Playfair [Pla01] is one of the most common statistical graphics [Cle93] for time-series data. It is measurable and easily comparable when the number of lines is small and the range of their values are close. However, as the number of lines and the range of their values increase the graph becomes more complex, precise tasks (R2, R5) become difficult and the users start to experience cognitive overload [APM*11].

Stack zooming [JE10] allows users to examine and compare focus points while retaining the overview context and provides visual clues to connect the two. Our work also provides the ability to examine the details while keeping the overview in context, but as a visual display rather than an interaction function.

Various systems already allow communicating large scale time-series data effectively through interactive exploration [HS03, KL06, LKL*04, MMKN08, VWVS99]. Our work supports exploration of time-series data when interactive functions are not available.

Treemaps [Shn92] are often used to display financial data [Wat99], which is the primary type of time-series data we examine in our use case. While treemaps are powerful in displaying the hierarchical structure and the trend of both the combined group and the individual commodities, they are not capable of displaying detailed changes over time (R3, R4) which our work also aims to summarize.

The simple design of the band graph enables it to be a powerful tool in describing the overview of a dataset whose audiences have no prior knowledge of the said dataset [Mun11]. However, the band graph does not support examining the individual time-series (R4, R5) for further explorations. We adopted its design and applied the boundaries and the central value components to our visualization technique to communicate the overview and the aggregated summary. Similar to the band graph, Fua et al. [FWR99] introduced multi-dimensional graduated bands that encode the extent and the mean of polyline clusters in hierarchical parallel coordinates. However, their work focuses on multivariate datasets which are beyond the scope of this paper.

The stream graph [BW08] utilizes the ThemeRiver [HHN00] layout to visualize the overall theme and its changes over time by moving all values around a varying central line and preserving limited measurability on the individual lines. Though sharing a similar appearance, our work plots the lines using their true y-axis values, providing easier measurement and comparison (R5).

The horizon graph [Rei08] utilizes two-tone pseudo coloring [SMY*05] and separate charts for each time-series data to provide efficient comparison across a larger visual span [JME10] while preserving the movement of the individual commodities. Cloudlines [KBK11] shares a similar design strategy to the horizon graph utilizing separate and normalized space-saving design with the additional lens magnification interactive function to support a closer examination of the details. However, neither graphs' visual design provides value measurement (R4) which is important to applications such as analyzing stock market data. Our work, on the other hand, provides a simplified comparison between individual commodities on the important factors while retaining enough measurability.

Many charts precisely communicate one aspect of data but leave out the context that casual experts need to identify potential biases. For example, while treemaps communicate price and trend effectively, a user cannot determine how the comparison between different stocks change over time and whether the trend is likely to continue by treemaps alone. Commercial tools, such as Tableau, allow trained analysts to explore datasets effectively by providing multiple instances of such charts, but are not designed to communicate the knowledge gained throughout exploration to casual experts efficiently. For example, online trading platforms often utilize visualizations such as line graph (moving average, advance/decline indexes, etc.) and candlestick chart (high, low, open and closing prices) to allow its users to examine stocks and market indicators closely, generally, one at a time using separate views. The sector or market summary is primarily visualized using graphs (treemaps, candlestick graphs, etc.) where the users cannot identify detailed information for individual stock under the group. This requires the users to obtain and compare the information between different views during different steps of interactive exploration to retrieve the insights desired with context and explanation. Our work focuses on efficiently communicating that knowledge to the casual experts by highlighting the important insights while preserving and linking the analyses to the context in one single visualization. In our example of stock data, for instance, our compact three-component visualization enables decision-makers to gain quick insight on the long term trends/highs/lows, indicators for short-term investment (e.g., sector increasing, but one or two stocks at low over a six month period) and for long-term investment (e.g., multiple sector stocks reaching all time high but showing downward trend indicating time to divest).

Aside from treemaps which do not encode time and techniques that rely heavily on interactive functions for exploration, we will compare visualization techniques utilizing the linear time structure [AMM*07] alongside our summarized line graph in their capability to complete different tasks in Section 5 and Section 6.

4. Summarized Line Graph

After studying our audience and the design requirements, we apply the findings to the design of a summarized line graph. We selected the line graph because it is one of the most commonly used and understood visualization techniques for representing simple time-series data [Cle93] that has supported decision-making in multiple domains. For example, knowledge gained from financial data can aid investors and financial fraud investigators in decision-making, and knowledge gained from public safety data can aid decisions in resource allocation, analyzing crime analytics, search and rescue, etc. Since there are many visualization techniques that focus on the efficient analysis of time-series data [AMST11], we can use these visualization techniques to evaluate the effectiveness of our summarized line graph. Our proposed visualization technique targets casual experts, e.g., decision-makers with strong domain knowledge but limited time or training for advanced visualizations. The proposed technique is able to effectively and efficiently communicate multiple quantitative time-series data and their correlations. We first present the overall design direction and discuss our design choices for the three components in our summarized line graph.

4.1. Design Approach

To summarize the data for efficient knowledge retrieval without losing important details, our summarization approach is driven by the following characteristics derived from our design requirements:

First, a summary visualization must represent the full dataset [EF10]. Since summarizing the main takeaway can result in losing perspective on parts of the dataset, it is important for users to be able to obtain a basic understanding of the range and the distribution of the actual data to satisfy **R3** (fidelity) and **R5** (comparison) even with the visualization focusing primarily on the summarized components. Additionally, to avoid information bias and satisfy **R2** (accuracy), the visualization should not only enable quick extraction of important insights but also allow its audience to easily understand how these insights are retrieved from the data. Finally, to satisfy **R4** (precision), the above need to be measurable.

To include these characteristics, the final visual design separates the data into three components, each with a different focus and priority in encoding the data and combine them to provide a balanced summary presentation. The three components are: *Representative Data*, *Analytical Highlights*, and the *Data Envelope*.

The Representative Data provides the audience with a simple but precise description of the dataset (**R1**). It should be clear and easy for casual users to understand without additional training and should be easy to communicated quickly. Visually, the representative data should be the most prominent element in the visual summary.

Analytical Highlights are added to the visual summary as the second visual component to reduce the time needed to extract useful insights from the dataset, to ensure it is clear how the insights are extracted, and to minimize the loss of important discoveries during the exploration (**R1**). This component should be designed to address the specific insights of interest to the decision-makers. In the visual design, the analytical highlights should be easy to identify and should also provide a connection to the raw data indicating how the insight is extracted (**R3**). Visually the analytical highlights should not outshine the representative data but remain easily recognizable.

The Data Envelope summarizes the remaining aggregated data to put the first two components into context (**R3**) and, therefore, aids the users in identifying possible information bias. It should provide simple yet specific (**R3**) details (e.g., boundary values) of the raw data that are not included in the representative data, and possibly allow basic comparisons between different data points (**R5**). Visually, the data envelope should be less prominent compared to the representative data and the analytical highlights, so that it provides overall context but does not distract.

To accommodate the casual experts' lack of training in data analysis, we utilize familiar visualizations and statistics (e.g., time-series graphs, bands, average, etc.) to create the summarized line graph.

4.2. Representative Data

Our summarized line graph plots the mean as the representative data. The mean provides measurable values (**R4**) and changes over time that represent the central tendency of the entire dataset (**R1**) and is easily understood by casual users. Summarizing the entire dataset with one line on the graph creates the initial focus of the visualization in a simple yet effective manner.

We chose mean over median to focus on values rather than order. While mean is susceptible to the influence of outliers and can be misleading when extreme values exist in the dataset, the data envelope is designed to counteract this problem.

4.3. Analytical Highlights

In this paper, we create a generic design that utilizes absolute/global extrema to extract simple analytical highlights that are relevant to multiple domains: ranges, trends, correlations, outliers, and key moments (time steps when external events may have influenced multiple time-series). We plot the absolute maximum as downward-pointing triangles in blue and the absolute minimum as upward-pointing triangles in red for each of the time-series. We label the triangles and align them to the mean curve using vertical dotted lines for better time point measurement and comparison. With the absolute extrema triangles, viewers can extract the global and individual ranges using the y-values of the extrema and the approximate global and individual trends by comparing the time stamps and orders of the extrema. These characteristics provide a sufficient overview for each of the time-series data with little visual clutter (two data points each); aligning and comparing these overviews alongside the representative data allows users to identify possible correlations, outliers, and key moments and compare different subsets of the data (**R5**). Additionally, by comparing the ratio of growing trends to decaying trends, users can perform analyses similar to the market indicator of advance-decline issues. If the design instead highlights local extrema, users can also examine the local extrema to perform analyses of new highs-new lows. Finally, plotting the extrema provides measurable values of the ranges and the key moments (**R4**).

By analyzing the values and time points of the extrema, we reduce the challenges of analyzing values of overlapping lines and lines that suffer from adjusting to the overall scale of the dataset. By extracting these insights from the highlighted extrema, the summarized line graph also allows users to understand how the insights are extracted from the dataset (**R3**). This design assumes the fluctuation of the lines has a smaller vertical impact than the actual trend over time, which from our observation is the case for most real-life data.

For example, in Figure 1 where Nasdaq stock prices of the transportation industry in the air freight/delivery service subsector during the year of 2016 are displayed, the stock price of FedEx (FDX) ranges between \$125 and \$195 and has an overall growth trend. By comparing the extrema and the density bands, the graph suggests the subsector mostly shares a positive correlation with no obvious single outlier. It also suggests that the end of June is worth further exploration as six of the airline stock prices reached their absolute minimum on the same day, indicated by the aligned dotted lines. While these highlights do not target a specific scenario, these insights can be useful to traders looking to invest, managers trying to understand the performance of their company against its competitors, or security advisers searching for attacks and insider trading.

4.4. Data Envelope

Since the data envelope summarizes the remaining data, it provides important information not presented in the representative data and analytical highlights in a simplified and contextual manner. To minimize the potential misleading information from extreme values, the data envelope adds density bands between the mean and each of the original lines to visualize the distribution of the time-series. The transparency of the band is defined in Equation 1, where C_o is the user-chosen opacity, normally between 1 and 2. The equation is designed to incorporate the standard deviation, the total range, and the line count to provide better separation of the different densities.

$$Opacity = C_o \times \frac{log(\frac{10 \times std\ dev}{max - min})}{line\ count} \tag{1}$$

With multiple overlapping layers, the final opacity will inform the audience of the distributions of the lines, allowing them to better understand the original dataset and its effect on the representative data (**R3**). Density Bands also aid the audience in connecting the extrema of a line in further exploring the original dataset (**R3**). We chose not to use a conventional confidence band to preserve more information on the individual time-series. Similar to Novotný et al.'s focus+context design [NH06], the semi-transparency design allows the audience to focus more on the other two components. By placing the transparent bands within the 2D plane, it provides the audience with enough measurability for the data envelope component (**R4**). Additionally, by examining the density bands and the mean in the same graph, users can examine the distribution of time-series above and below the mean, similar to how market indicators examine the percentage of stocks above and below key moving averages.

4.5. Generalizability

To demonstrate the generalizability of our summarized line graph, we present an alternative design using Pearson's correlation coefficient as the analytical highlights for comparing point-wise trends of individual time-series against the trend of aggregated time-series [CGK*07] for crime analytics [MME*12]. As shown in Figure 2, the graph highlights the representative crime (car prowl, the time-series with a strong positive correlation; blue, with an up arrow following the label), the crime with the most opposite trend to the overall trend (narcotic, moderate negative correlation; yellow, with a down arrow following the label), and the the crime most indepedent to the overall trend (street robbery, very weak correlation close to zero; red, with a dash following the label) alongside the the average number of crime reports per month for the city of Seattle from 2008 to 2018. The line-width reflects the strength of the correlation.

5. Evaluation

We conducted a user evaluation with 30 minute sessions for our summarized line graph design. The evaluation examined the complexity and the ability of the summarized line graph to extract insights on

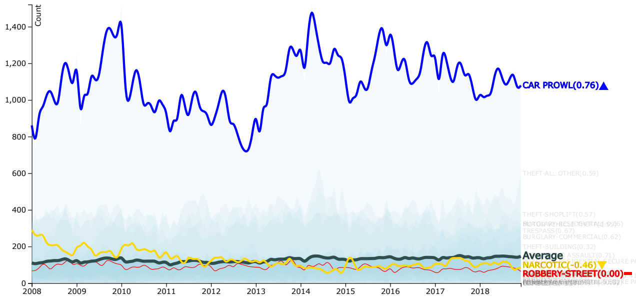

Figure 2: *An alternative summarized line graph design using correlation analytical highlights, showing the number of reports for 30 crime subcategories from the city of Seattle between 2008 and 2018.*

the overall trend, outliers, and key moments against four visualization techniques capable of communicating time-series data through static presentation: the traditional line graph, the band graph, the stream graph, and the horizon graph. We designed the tasks to be as suitable as possible to all of the visualization techniques tested. The training material, sample task images, the raw data, and the analysis results can be found in the supplemental material.

5.1. Hypotheses

Our ultimate design goal is for the summarized line graph to satisfy all the design requirements. We test whether this design is capable of providing a balanced and effective analysis on *all* of the tasks, and compare to other visualization techniques. Based on the visual designs, we hypothesize that

H1 *The summarized line graph will perform better in identifying outliers (accuracy and time) and locating key moments (time) compared to the line graph.* They will perform similarly in identifying overall trends (accuracy and time).

H2 *The summarized line graph will perform similarly to a band graph (accuracy and time) in identifying the original graph and the overall trend.* Because a band graph does not support examining individual time-series, the summarized line graph will perform better in identifying outliers and key moments.

H3 *The summarized line graph will perform better in identifying the original graph and outliers, and locating key moments (accuracy and time) compared to the stream graph.* They will perform similarly in identifying overall trends (accuracy and time).

H4 *The summarized line graph will perform better in identifying the original graph (accuracy and time) compared to the horizon graph.* The two techniques will perform similarly in identifying the overall trend and outliers and locating key moments (accuracy and time). The summarized line graph also supports measuring actual values which the horizon graph does not.

5.2. Participants

For this evaluation, we recruited 22 university student volunteers (13 male, 9 female) ranging from 18 to 32 years of age (average age of 25) with backgrounds in Computer Science, Electrical and Computer Engineering, Industrial Engineering, Aerospace Engineering,

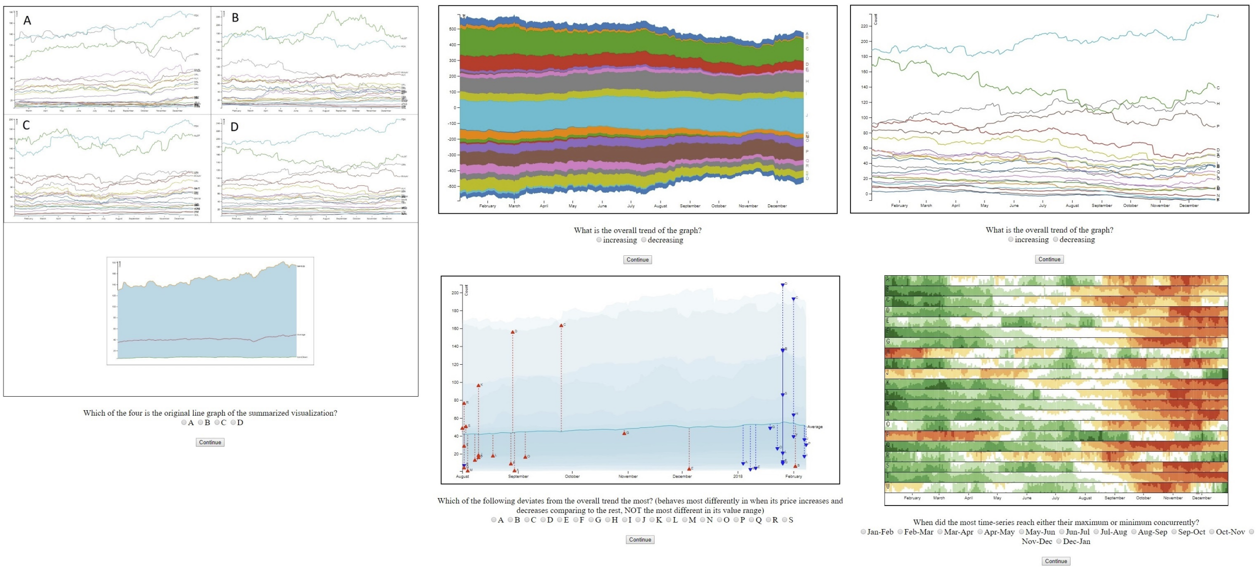

Figure 3: *Examples of the five visualization techniques and the four tasks used in the study: identifying the original graph using a band graph (left), identifying the overall trend using a stream graph (top center), identifying the overall trend using a traditional line graph (top right), identifying the outlier using a summarized line graph (bottom center), and locating the key moment using a horizon graph (bottom right).*

Medicinal Chemistry, and Linguistics through the university public email lists and campus billboards. The majority of the participants were familiar with basic Excel-level visualization techniques, including the traditional line graph. There were no color-blind participants (self-reported). The participants were compensated at a \$10 hourly rate. All participants were fluent in English.

5.3. Apparatus

The evaluation was conducted on standard desktop computer equipped with a mouse, a keyboard, and a 30" monitor set to 2560×1600 resolution. The evaluation was performed on a web browser page maximized on the screen. Each image was displayed at a 960×500 resolution. Only the mouse was used for the tasks.

5.4. Tasks

During the evaluation, the participants were given four types of tasks to evaluate the complexity of the visualization techniques and how they support our design requirements R1, R2, R3, and R5. Design requirement R4 was not included in the evaluation as it is straightforward from the design of the visualization techniques. The analytical tasks are inspired by Amar et al.'s taxonomy tasks [AES05], which explore the characteristics of an entire dataset, are not easily achievable by the majority of visualization techniques, and are reasonable for scenarios working with time-series data. We used two years of historical Nasdaq stock market data from the airline industries and four years of historical Nasdaq stock market data from the technology industries. We altered the time range and the stocks used in each question, typically a year's worth of data for 20 to 30 stocks, to prevent participants from memorizing the answer. As a result, each question was given a "unique dataset." The correct response for each task was pre-calculated using the raw data.

All considered visualization techniques were used to complete each task in a random order. Figure 3 shows samples of the tasks and the visualizations techniques used in the study. Each task was evaluated based on completion time and correctness. We evaluate the performance of our summarized line graph on these tasks against four representative time-series visualization techniques: the traditional line graph, the band graph, the stream graph, and the horizon graph. The traditional/simple line graph, the stream/stacked graph, and the horizon graph are representative visualization techniques for displaying multiple time-series data [JME10], and the band graph provides a simple yet effective overview of multiple time-series data while sacrificing the ability to explore the individual series.

The traditional line graph shares a similar visual appearance and attributes with braided graph and scattered plot [WHZ*18] and is one of the most commonly used and understood visualization techniques [Tuf83]. The band graph shares similar appearance, functionality, and limitations as the river plot [BPS*07] and the functional boxplot [SG11] with a more simple and direct presentation. The stream graph is a good representation of stacked graphs that highlights the overall dynamics and the individual contributors. Finally, the horizon graph is a good representative visualization technique that utilizes small multiples to save space and explore both the individual and the overall dataset. We chose horizon graph over conventional small multiples because conventional small multiples can take up noticeably more vertical space, which may not be available during the knowledge transfer between the data analysts and the decision-maker. The normalization and the different binning in the horizon graph also allow easier trend identification and comparison. The band graph, the stream graph, the horizon graph, and our summarized line graph can all be derived from the line graph. Each visualization technique was given two questions for each tasks. We excluded the choice “undeterminable” from the answers, forcing the participants to make their best guess when the answer is not obvious; this option complicates the calculation of accuracy and can influence the decision time measurement since the participants may give up at different levels of frustration.

5.4.1. Identifying the Original Graph

For each visualization technique (excluding the traditional line graph), we gave two questions per task: one for identifying the original graph composed of 20 time-series, and a similar question with 30 time-series in the graph. For possible answers, the participants were given a choice of four line graphs (each with 20 or 30 time-series, respectively) to identify as the one from which the given visualization (i.e. summarized line graph, band graph, stream band, and horizon graph) was derived. By analyzing the time and accuracy required to identify the original graph, we can better understand the complexity of each visualization technique and the user’s ability to create the mental image of the raw form of the data through such techniques. The result reflects the techniques’ ability to meet design requirement R3 (fidelity).

5.4.2. Identifying the Overall Trend

The participants were asked to identify the overall trend for the dataset using the five visualization techniques. For each question, the participants were asked to identify whether a given graph had an overall increasing or decreasing trend. Each visualization technique was given two questions, one with an overall growth or fall of five percent, and the other thirty percent. The result of the task reflects the techniques’ ability to meet requirement R1 (comprehensibility).

5.4.3. Identifying the Outlier

The participants were asked to identify the outlier using the five visualization techniques excluding the band graph. For each question, the participants were asked to select one time-series in the given graph that deviated from the overall trend the most. This task focuses on anomalous behavior, meaning a data source’s value is increasing or decreasing in the opposite direction of the rest of the group, rather than a data source having values significantly higher or lower than the rest of the group. Note that we removed the band graph starting with this task as the individual time-series are not identifiable with this visualization technique. The result of the task reflects the visualization techniques’ ability to meet the design requirement R1 (comprehensibility) and R2 (accuracy).

5.4.4. Locating the Key Moment

Finally, the participants were asked to identify the time step when a key moment occurred (multiple time-series reaching their extrema concurrently) using the five visualization techniques excluding the band graph. For each question, the participants were asked to identify the month when the most time-series reached either their maximum or minimum concurrently in the given graph. The result of the task reflects the visualization techniques’ ability to meet design requirements R2 (accuracy) and R5 (comparison).

Table 1: *T-test on the effects of difficulties.*

Task	Correctness	Completion Time
Original Graph	p-value = 0.80	p-value = 0.19
Overall Trend	p-value = 0.64	p-value = 0.45

5.5. Procedure

After each participant provided informed consent, we provided a 10-minute training session describing how our summarized line graph and other visualization techniques used in the evaluation were derived from the traditional line graph. We then administered three sample questions for participants to test their understanding of the visualization techniques and the tasks to complete.

During the evaluation, the participants answered multiple-choice questions for the tasks. The evaluation question order was randomized and updated for each participant using a Latin Square randomization order [CC50] to ensure an even distribution of the question types throughout the evaluation trials to minimize the learning effect in the results. After the evaluation, the participants were surveyed about their demographic, self-reported skill level, and thoughts on the tasks and the visualization techniques.

5.6. Results

To analyze the results, we examined the 95% confidence intervals calculated utilizing the bootstrapping method [Efr92] with 1,000 iterations to alleviate the small sample size. Figure 4 presents the accuracy and completion time of each visualization under each task, and Figure 5 presents an overall comparison between the techniques.

In this section we compare the performance of the techniques using the overlap-test [SG01] and the t-test [Stu08]. Table 1 lists the difference in the difficulties of the tasks to identify the original graph and the overall trend, which is not significant for the correctness and the completion time. Therefore, the following analysis treats the results from the different difficulties equally.

Figure 4 shows the accuracy of the summarized line graph to be consistently above 80% correct. For the task to identify the original graph, summarized line graphs (μ=91%) perform significantly stronger in correctness compared to stream graphs (μ=62%) and horizon graphs (μ=33%), and similarly to band graphs (μ=93%). For the task to identify the overall trend, summarized line graphs (μ=100%) perform significantly stronger in correctness compared to stream graphs (μ=83%) and horizon graphs (μ=81%), and similarly to band graphs (μ=98%) and line graphs (μ=98%). For the task to identify the outlier, summarized line graphs (μ = 93%) perform significantly stronger in correctness compared to line graphs (μ=67%) and stream graphs (μ=22%), and similarly to horizon graphs (μ=91%). Finally, for the task to locate the key moment, summarized line graphs (μ=93%) perform significantly stronger in correctness compared to stream graphs (μ=43%) and horizon graphs (μ=71%), and similarly to line graphs (μ=79%, p-value=0.08). Figure 5 shows that over the scope of this experiment, which was designed to reflect the visualization technique’s ability to satisfy the requirements listed in the introduction, the summarized line graph

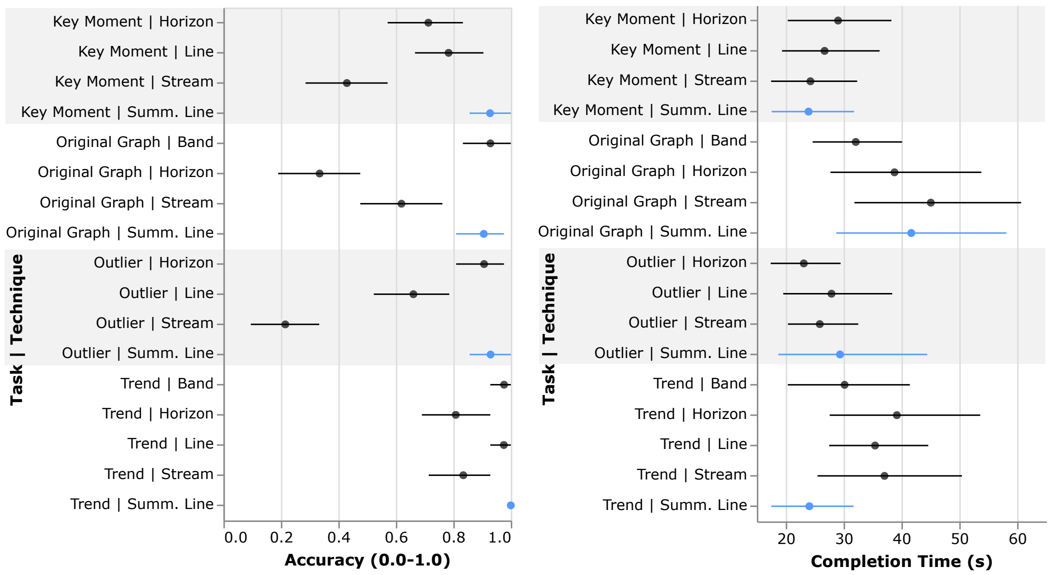

Figure 4: *95% confidence interval plots of the study results in accuracy (left) and completion time (right) separated by task and technique.*

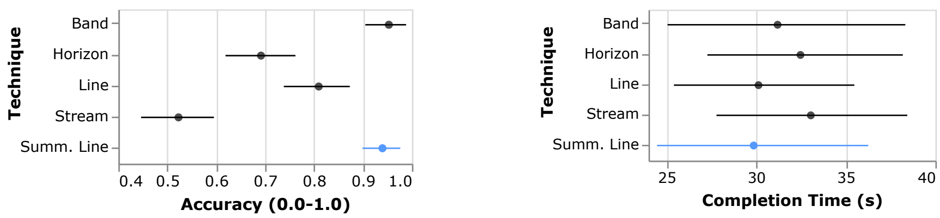

Figure 5: *95% confidence interval plots for the different techniques in accuracy (left) and completion time (right).*

(μ=94%) performs significantly stronger in correctness compared to the line graph (μ=81%), the stream graph (μ=52%) and the horizon graph (μ=69%), and similarly to the band graph (μ=95%).

Figure 4 and Figure 5 show that the summarized line graph has the shortest average completion time in overall comparison, trend identification and key moment locating. However, the majority of the differences in the average completion time are not statistically significant. The only exception lies in identifying the overall trend, where the summarized line graphs (μ=23.94s) perform significantly more efficiently compared to line graphs (μ=26.98s, p-value=0.05).

The summarized line graph received positive feedback from the participants in the post-experiment survey. The participants appreciated its cleaner aesthetic and found its resemblance to the more familiar traditional line graph helpful. The participants also found the average curve and the removal of original lines useful when examining the data. Finally, the participants agreed that the summarized line graph is easy to interpret and helps them identify the insights required for the tasks. However, a few participants also expressed minor frustration about the additional time it took to find the labels in the summarized line graph.

6. Discussion

Here we first explain the results of our evaluation. We then discuss aspects of the summarized line graph and our design approach.

6.1. Evaluation Results

Comparing between the summarized line graph and the traditional line graph, we can conclude from our study result that the summarized line graph is more accurate in identifying the outliers and more efficient in identifying the overall trend, confirming and exceeding the predictions in H1. The summarized line graph and the band graph performed similarly regarding accuracy and efficiency, which confirms H2. The simple design of a band graph allows it to be a powerful tool in communicating an overview of the data, but is not capable of the more detailed tasks a summarized line graph can handle. The summarized line graph performs significantly stronger than the stream graph in the correctness of every task, which exceeds the accuracy predictions in H3, but does not meet the efficiency predictions in the hypothesis. Finally, the summarized line graph outperforms the horizon graph in the accuracy to identify the original graph, to identify the overall trend, and to locate the key moment which exceeds the accuracy predictions in H4 but does not meet the predicted efficiency in identifying the original line graph. These results matches our overall expectation as our visualization technique is designed specifically for tasks common to our targeted audiences, similar to the ones used in the evaluation. With the lack of visualization techniques designed specifically for the same goal, we compared against designs that are capable of extracting the same information and proved that our design provides a more accurate result for such tasks than currently existing tools.

Unfortunately, due to the large variance in the completion time between the different participants, most of our comparisons between the efficiency are inconclusive. However, while the differences are not statistically significant, the experiment shows the summarized line graph to be at least as effective as the existing techniques we tested against. The wide range of completion time may result from some participants giving up and moving forward with random guesses at different points of time on tasks that require more effort. We will consider rewarding the participants with a bonus for results above a certain level of correctness in future studies to stress the importance of getting the correct answer.

6.2. Limitations

Our summarized line graph is not an intuitive visualization design and will require some training before one can use it. Based on our user study results, however, a 10-minute training is sufficient. Using the global maximum and minimum may be effective for examining data across a long period of time, but the audience may be confused by fluctuations when examining data that span a shorter period of time or have a stable global trend. Also, placing the labels next to the extrema makes the design less suitable for searching for specific time-series of interest without prior knowledge of their behavior. Finaly, The semi-transparent density bands can also be misleading to audiences familiar with stream graphs as the two techniques share a similar appearance but are read differently.

Several advantages outweigh these drawbacks. As a shared-space technique [JME10], the summarized line graph's display size is independent of the number of time-series it displays, unlike techniques that create a separate chart for each series. While shared-space techniques are traditionally more efficient with fewer lines, the summarized line graph's design should reduce the impact of overlap and clutter better than traditional shared-space time-series plots as it aggregates the original lines into polygonal visual elements. Furthermore, an important advantage of our summarized line graph is that it is simple to read direct values. In comparison, horizon graphs make reading values difficult, and reading the stream graph requires estimating the width of a band. While this can be supported by interaction, such interactions are not always available.

The use of extrema and the automatic selection were chosen as "shortcuts" of typical analysis tasks, and we demonstrate the benefit of this simple design using the three-component visual summary approach. None of our participants explicitly requested additional forms of analytical highlights, and based on their performance, appeared to perform well for the specific tasks in our evaluation. However, we leave surveying domain experts on effective indicator-task combinations to future work. More complex highlights and semi-automated selection of features can be added to address the needs of other scenarios. Similarly, the scalability of the technique depends on the data and the chosen analytical highlight. From our study, the design was able to clearly visualize at least 31 time-series. However, we leave a formal scalability study also to future work.

6.3. Three-Component Visual Summary

Differing from summarization techniques that focus on the aggregation of data and present the data item equally, the three-component visual summary uses components with different priorities that are reflected in their corresponding visual design. The intention is to provide audiences with an order to explore the visual summary in a manner similar to the Shneiderman mantra [Shn96]. The flexibility to choose the variables for the three components also allows the visual summary to be a more customizable and focused experience, which works toward reducing the communication gap.

However, selecting which variables to visualize can be challenging when applying this methodology to any visualization technique, as there are no prescribed steps to follow. Rather, the task depends on the designer's knowledge of the variables' roles, the strength of the visual components, and the insights important to the scenario.

Based on the result of the summarized line graph evaluation, we believe that the approach can enable the audience to extract a more accurate and extensive understanding of the dataset within the same amount of time compared to existing visualization techniques. While this paper focuses on the line graph use case, we believe the approach can be applied to more visualization techniques as long as the designer can identify components that satisfy the characteristics described in Section 4.1. We plan to explore this in future work.

7. Conclusion and Future Work

In this paper, we present a novel line graph summarization design to better bridge the data analysis communication gap. The summarized line graph presents a dataset through three visual components: mean as the representative data; trends, correlations and key moments as the analytical highlights; and data density as the data envelope. We evaluated the summarized line graph in its ability to capture the scope of the dataset and support identifying trends and outliers, and locating key moments. We compared the results in both task completion accuracy and efficiency against traditional line graphs, band graphs, stream graphs, and horizon graphs.

In future research, we will explore methods to improve the scalability of our summarization technique, procedures of choosing the optimal variables for the three components in a visual summary design, and the ability to display multiple summaries within the same visual display to support a two-level hierarchical data structure. We will also explore the effectiveness of this three-component visual summary approach when applied to other visualization techniques, data types, and multi-view interactive systems.

8. Acknowledgements

The authors would like to thank José Florencio de Queiroz Neto and Christina Stober for proofreading the paper as well as the participants who volunteered to evaluate the design.

References

[ABR96] AJZEN I., BROWN T. C., ROSENTHAL L. H.: Information bias in contingent valuation: effects of personal relevance, quality of information, and motivational orientation. *Journal of Environmental Economics and Management 30*, 1 (1996), 43–57. doi:10.1006/jeem.1996.0004. 3

[AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *Proceedings of the IEEE Symposium on Information Visualization* (2005), IEEE, pp. 111–117. doi:10.1109/INFVIS.2005.1532136. 7

[AMM*07] AIGNER W., MIKSCH S., MÜLLER W., SCHUMANN H., TOMINSKI C.: Visualizing time–oriented data–a systematic view. *Computers & Graphics 31*, 3 (2007), 401–409. doi:10.1016/j.cag.2007.01.030. 4

[AMST11] AIGNER W., MIKSCH S., SCHUMANN H., TOMINSKI C.: Survey of visualization techniques. In *Visualization of Time-Oriented Data*. Springer, 2011, pp. 147–254. doi:10.1007/978-0-85729-079-3_7. 4

[APM*11] ANDERSON E. W., POTTER K. C., MATZEN L. E., SHEPHERD J. F., PRESTON G. A., SILVA C. T.: A user study of visualization effectiveness using EEG and cognitive load. *Computer Graphics Forum 30*, 3 (2011), 791–800. doi:10.1111/j.1467-8659.2011.01928.x. 4

[BC94] BRYMAN A., CRAMER D.: *Quantitative Data Analysis for Social Scientists*. Routledge, New York, NY, 10001, 1994. 3

[BMW17] BRYAN C., MA K.-L., WOODRING J.: Temporal summary images: An approach to narrative visualization via interactive annotation generation and placement. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 511–520. doi:10.1109/TVCG.2016.2598876. 3

[BPS*07] BUONO P., PLAISANT C., SIMEONE A., ARIS A., SHNEIDERMAN B., SHMUELI G., JANK W.: Similarity-based forecasting with simultaneous previews: A river plot interface for time series forecasting. In *Proceedings of the International Conference on Information Visualization* (2007), IEEE, pp. 191–196. doi:10.1109/IV.2007.101. 7

[BW08] BYRON L., WATTENBERG M.: Stacked graphs—geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (Nov 2008), 1245–1252. doi:10.1109/TVCG.2008.166. 4

[CAFG12] CORRELL M., ALBERS D., FRANCONERI S., GLEICHER M.: Comparing averages in time series data. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (2012), ACM, pp. 1095–1104. doi:10.1145/2207676.2208556. 3

[CC50] COCHRAN W. G., COX G. M.: *Experimental Designs*. Wiley, 1950. 8

[CGK*07] CHANG R., GHONIEM M., KOSARA R., RIBARSKY W., YANG J., SUMA E., ZIEMKIEWICZ C., KERN D., SUDJIANTO A.: WireVis: Visualization of categorical, time-varying data from financial transactions. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2007), IEEE, pp. 155–162. doi:10.1109/VAST.2007.4389009. 6

[Cle93] CLEVELAND W. S.: *Visualizing Data*. Hobart Press, 1993. 2, 4

[DOL03] DE OLIVEIRA M. F., LEVKOWITZ H.: From visual data exploration to visual data mining: a survey. *IEEE Transactions on Visualization and Computer Graphics 9*, 3 (July 2003), 378–394. doi:10.1109/TVCG.2003.1207445. 3

[EF10] ELMQVIST N., FEKETE J.-D.: Hierarchical aggregation for information visualization: Overview, techniques and design guidelines. *IEEE Transactions on Visualization and Computer Graphics 16*, 3 (May 2010), 439–454. doi:10.1109/TVCG.2009.84. 4, 5

[Efr92] EFRON B.: Bootstrap methods: another look at the jackknife. In *Breakthroughs in Statistics*. Springer, 1992, pp. 569–593. doi:10.1007/978-1-4612-4380-9_41. 8

[FWR99] FUA Y.-H., WARD M. O., RUNDENSTEINER E. A.: Hierarchical parallel coordinates for exploration of large datasets. In *Proceedings of the IEEE Conference on Visualization* (Oct 1999), IEEE Computer Society Press, pp. 43–50. doi:10.1109/VISUAL.1999.809866. 4

[HA06] HEER J., AGRAWALA M.: Multi-scale banking to 45 degrees. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (2006), 701–708. doi:10.1109/TVCG.2006.163. 4

[HD11] HULLMAN J., DIAKOPOULOS N.: Visualization rhetoric: Framing effects in narrative visualization. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2231–2240. doi:10.1109/TVCG.2011.255. 3

[HDA13] HULLMAN J., DIAKOPOULOS N., ADAR E.: Contextifier: automatic generation of annotated stock visualizations. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (2013), ACM, pp. 2707–2716. doi:10.1145/2470654.2481374. 3

[HHN00] HAVRE S., HETZLER B., NOWELL L.: ThemeRiver: visualizing theme changes over time. In *Proceedings of the IEEE Symposium on Information Visualization* (2000), pp. 115–123. doi:10.1109/INFVIS.2000.885098. 4

[HS03] HOCHHEISER H., SHNEIDERMAN B.: Interactive exploration of time series data. In *The Craft of Information Visualization*. Elsevier, 2003, pp. 313–315. doi:10.1016/B978-155860915-0/50039-1. 4

[JE10] JAVED W., ELMQVIST N.: Stack zooming for multi-focus interaction in time-series data visualization. In *Proceedings of the IEEE Pacific Symposium on Visualization* (2010), pp. 33–40. doi:10.1109/PACIFICVIS.2010.5429613. 4

[JME10] JAVED W., MCDONNEL B., ELMQVIST N.: Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (Nov 2010), 927–934. doi:10.1109/TVCG.2010.162. 4, 7, 10

[KBK11] KRSTAJIC M., BERTINI E., KEIM D.: Cloudlines: Compact display of event episodes in multiple time-series. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (2011), 2432–2439. doi:10.1109/TVCG.2011.179. 4

[KCA*16] KO S., CHO I., AFZAL S., YAU C., CHAE J., MALIK A., BECK K., JANG Y., RIBARSKY W., EBERT D. S.: A survey on visual analysis approaches for financial data. *Computer Graphics Forum 35*, 3 (2016), 599–617. doi:10.1111/cgf.12931. 2

[KL06] KINCAID R., LAM H.: Line graph explorer: scalable display of line graphs using focus+ context. In *Proceedings of the ACM Conference on Advanced Visual Interfaces* (2006), ACM, pp. 404–411. doi:10.1145/1133265.1133348. 4

[LKL*04] LIN J., KEOGH E., LONARDI S., LANKFORD J. P., NYSTROM D. M.: Visually mining and monitoring massive time series. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining* (2004), ACM, pp. 460–469. doi:10.1145/1014052.1014104. 4

[Man10] MANN P.: *Introductory Statistics*. John Wiley & Sons, 2010. URL: https://books.google.com/books?id=N_mEBiCYaqkC. 3

[MEV*14] MADHAVAN K., ELMQVIST N., VORVOREANU M., CHEN X., WONG Y., XIAN H., DONG Z., JOHRI A.: DIA2: Web-based cyberinfrastructure for visual analysis of funding portfolios. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (Dec 2014), 1823–1832. doi:10.1109/TVCG.2014.2346747. 3

[MME*12] MALIK A., MACIEJEWSKI R., ELMQVIST N., JANG Y., EBERT D. S., HUANG W.: A correlative analysis process in a visual analytics environment. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (Oct 2012), pp. 33–42. doi:10.1109/VAST.2012.6400491. 6

[MMKN08] MCLACHLAN P., MUNZNER T., KOUTSOFIOS E., NORTH S.: LiveRAC: interactive visual exploration of system management time-series data. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (2008), ACM, pp. 1483–1492. doi:10.1145/1357054.1357286. 4

[Mun11] MUNDIGL R.: An underrated chart type: The band chart. http://www.clearlyandsimply.com/clearly_and_simply/2011/04/an-underrated-chart-type-the-band-chart.html, 2011. 4

[NH06] NOVOTNÝ M., HAUSER H.: Outlier-preserving focus+ context visualization in parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (2006), 893–900. 6

[PC05] PIROLLI P., CARD S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of the International Conference on Intelligence Analysis* (Bedford, MA, USA, 2005), vol. 5, MITRE Corporation, pp. 2–4. 3

[Pla01] PLAYFAIR W.: *The commercial and political atlas: representing, by means of stained copper-plate charts, the progress of the commerce, revenues, expenditure and debts of England during the whole of the eighteenth century*. T. Burton, 1801. 4

[PSM07] POUSMAN Z., STASKO J., MATEAS M.: Casual information visualization: Depictions of data in everyday life. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (Nov 2007), 1145–1152. doi:10.1109/TVCG.2007.70541. 3

[Rei08] REIJNER H.: The development of the horizon graph. In *Proceedings of the IEEE VIS Workshop on From Theory to Practice: Design, Vision and Visualization* (2008). 4

[SG01] SCHENKER N., GENTLEMAN J. F.: On judging the significance of differences by examining the overlap between confidence intervals. *The American Statistician 55*, 3 (2001), 182–186. doi:10.1198/000313001317097960. 8

[SG11] SUN Y., GENTON M. G.: Functional boxplots. *Journal of Computational and Graphical Statistics 20*, 2 (2011), 316–334. doi:10.1198/jcgs.2011.09224. 7

[SH10] SEGEL E., HEER J.: Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (Nov 2010), 1139–1148. doi:10.1109/TVCG.2010.179. 3

[Shn92] SHNEIDERMAN B.: Tree visualization with tree-maps: 2-D space-filling approach. *ACM Transactions on Graphics 11*, 1 (Jan. 1992), 92–99. doi:10.1145/102377.115768. 4

[Shn96] SHNEIDERMAN B.: The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages* (1996), pp. 336–343. doi:10.1109/VL.1996.545307. 2, 10

[SMY*05] SAITO T., MIYAMURA H. N., YAMAMOTO M., SAITO H., HOSHIYA Y., KASEDA T.: Two-tone pseudo coloring: Compact visualization for one-dimensional data. In *Proceedings of the IEEE Symposium on Information Visualization* (Oct 2005), IEEE, pp. 173–180. doi:10.1109/INFVIS.2005.1532144. 4

[Stu08] STUDENT: The probable error of a mean. *Biometrika* (1908), 1–25. doi:10.2307/2331554. 8

[TC05] THOMAS J. J., COOK K. A. (Eds.): *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society Press, Piscataway, NJ, USA, 2005. 3

[Tuf83] TUFTE E.: *The Visual Display of Quantitative Information*. Graphics Press, 1983. 4, 7

[VW06] VIEGAS F. B., WATTENBERG M.: Communication-minded visualization: A call to action. *IBM Systems Journal 45*, 4 (2006), 801–812. doi:10.1147/sj.454.0801. 2, 3

[VWVS99] VAN WIJK J. J., VAN SELOW E. R.: Cluster and calendar based visualization of time series data. In *Proceedings of the IEEE Symposium on Information Visualization* (1999), IEEE, pp. 4–9. doi:10.1109/INFVIS.1999.801851. 4

[Wat99] WATTENBERG M.: Visualizing the stock market. In *Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 1999), ACM, pp. 188–189. doi:10.1145/632716.632834. 4

[WHZ*18] WANG Y., HAN F., ZHU L., DEUSSEN O., CHEN B.: Line graph or scatter plot? automatic selection of methods for visualizing trends in time series. *IEEE Transactions on Visualization and Computer Graphics 24*, 2 (2018), 1141–1154. doi:10.1109/TVCG.2017.2653106. 7

[WT04] WONG P. C., THOMAS J.: Visual analytics. *IEEE Computer Graphics and Applications 24*, 5 (Sept 2004), 20–21. doi:10.1109/MCG.2004.39. 3