

Visual Stenography: Feature Recreation and Preservation in Sketches of Noisy Line Charts

Rifat Ara Proma, Michael Correll, Ghulam Jilani Quadri, Paul Rosen

Abstract—Line charts surface many features in time series data, from trends to periodicity to peaks and valleys. However, not every potentially important feature in the data may correspond to a *visual feature* which readers can detect or prioritize. In this study, we conducted a *visual stenography* task, where participants re-drew line charts to solicit information about the visual features they believed to be important. We systematically varied noise levels (SNR $\approx$ 5–30 dB) across line charts to observe how visual clutter influences which features people prioritize in their sketches. We identified three key strategies that correlated with the noise present in the stimuli: the *Replicator* attempted to retain all major features of the line chart including noise; the *Trend Keeper* prioritized trends disregarding periodicity and peaks; and the *De-noiser* filtered out noise while preserving other features. Further, we found that participants tended to *faithfully* retain trends and peaks and valleys when these features were present, while periodicity and noise were represented in more qualitative or gestural ways: *semantically* rather than accurately. These results suggest a need to consider more flexible and human-centric ways of presenting, summarizing, pre-processing, or clustering time series data.

Index Terms—Line charts, graphical perception, time series.

I. INTRODUCTION

LINE charts are widely used for visualizing time-series data, enabling viewers to observe key patterns such as trends (e.g., in stock prices [1]), periodicity or cyclical patterns (e.g., in environmental data analysis [2]), or the peaks and valleys (e.g., in COVID-19 pandemic data [3]). However, there can be mismatches between the *statistical* patterns in time series data and the *visual* patterns in line charts. While statistical patterns may be present, how viewers interpret and internalize them remains unclear.

Prior work has explored how viewers match statistical patterns to visual features in charts, examining tasks like mean estimation [4], trend detection [5], and similarity evaluation [6]. Research also shows contextual cues like captions [7] and anchoring effects [8] influence interpretation. However, little attention has been given to which features (e.g., trends, periodic patterns, or peaks and valleys) naturally draw viewers' attention, especially under varying noise levels, which can impact cognitive load [9] and judgement [10]. Knowing how people attend to these features is critical to designing visualizations that reflect human understanding of data.

While many prior works measure people's accuracy for graphical-perception tasks, our goal is instead to reveal how

people choose which visual features to preserve when sketching noisy line charts and to characterize the distinct strategies that emerge. To address this, we conducted a human subjects evaluation of 20 participants to identify how individuals represent trends, periodic patterns, and peaks and valleys in line charts given a free-form sketching medium, a valuable method for capturing how individuals perceive visual features [11], [12] as it directly captures how users naturally structure and prioritize visual features without imposing explicit constraints on their perceptual judgments [13]. Each participant saw nine line charts and completed what we term a *visual stenography* task—re-drawing the time series data to capture the noteworthy visual features of the chart (see Sect. IV). The nine datasets were selected to include a variety of trends, periodic patterns, and notable peaks and valleys (see Sect. IV-A). We performed both qualitative (see Sect. V-A) and quantitative (see Sect. V-B) analyses of the resulting sketches to assess the retention of visual features under various levels of noise. We explicitly focused on noise as our independent variable because real-world line charts commonly include random fluctuations that often cause visual clutter. By systematically controlling signal-to-noise ratio (SNR), we reveal how increasing visual complexity influences viewers' mental prioritization and representation of visual features.

We then conducted a validation experiment (Sect. VII-B), where participants selected pre-generated sketches that best-represented reference charts. This allowed us to validate findings from our primary study, particularly given the subjectivity of sketching [14], [15], and gain insights into their selection rationale.

Our analysis revealed three distinct patterns of behavior surrounding the inclusion and exclusion of feature: (1) *Replicator* (see Fig. 1a), preserve all major features, including noise. (2) *Trend Keeper* (see Fig. 1b) focused primarily on retaining trends in the data while disregarding finer details such as periodic patterns and peaks (3) *De-noiser* (see Fig. 1c) filtered out noise while retaining other major features. Noise significantly influenced behavior, with participants more likely to adopt *Trend Keeper* strategies as noise increased. We also noted similar behavior in how accurately subjects represented features. While trends and peaks and valleys were often represented *faithfully* (accurately with respect to the stimuli observed), periodicity and noise were represented *semantically*: in ways that qualitatively denote the presence of a feature but may not accurately reflect exact statistical properties of the feature like scale or location. Our results suggest the use of not only *mathematical* but also *human-centric* representations of time series data for goals like designing and annotating visualizations, selecting algorithms for pre-

Rifat Ara Proma is with the University of Utah. E-mail: rifat.proma@utah.edu.

Michael Correll is with Northeastern University. E-mail: m.correll@northeastern.edu.

Ghulam Jilani Quadri is with the University of Oklahoma. E-mail: quadri@ou.edu.

Paul Rosen is with the University of Utah. E-mail: prosen@sci.utah.edu.




Fig. 1. The illustration on the left showcases the process of visual stenography, where the participants were shown line charts and were tasked to re-draw them on an iPad. The stimuli are in blue, and the participant sketches are in red. (a-c) We compared the sketches with the shown stimuli and identified three behavior pattern clusters—Replicator, Trend Keeper, and De-noiser. (d-f) Further, participants showed general robustness to noise in terms of preserving trends, periodicity, peaks, and valleys in their sketches across different datasets, and (g-i) they tended to represent the periodicity and noisiness of the stimuli semantically (i.e., conceptually) rather than faithful (i.e., accurately) replication in their sketches.

processing data (e.g., smoothing [16] or down sampling [17]), and performing data-related tasks (e.g., visual querying [18] or summarization [19]).

In summary, the key contributions of this work are:

- 1) Introduction of visual stenography as a novel sketch-based method for externalizing viewers' subjective prioritization of chart features under noise.
- 2) Identification of three behavior groups based on feature inclusion and exclusion patterns;
- 3) A study that verified the robustness of visual feature extraction of trends, periodicity, and peaks and valleys under various noise levels; and
- 4) Insights into human-centric feature preservation, showing that trends and peaks are represented more faithfully, while periodicity and noise are represented semantically.

II. PRIOR WORK

In this section, we discuss prior work in graphical perception and the utility of sketching. In Sect. III, we provide background on visual features in line charts.

A. Graphical Perception

An extensive body of work on perceptual and cognitive processes underlies human estimations in visualizations from seminal papers like Cleveland & McGill [20]. More recent work has examined not just extraction or comparison of individual values but graphical perception for more holistic statistical tasks [21], [22]. For instance, Gleicher et al. [23] found that viewers can reliably estimate average values in scatterplots, even with noise. Other studies confirm that viewers can identify scatterplot trends, even in the presence of noise and outliers [24]–[27]. Furthermore, investigations into the reliability of assumptions individuals rely on when estimating trends from various plots have shed light on viewers'

capabilities to accurately assess trends in line charts and other bivariate visualizations [5]. However, noise's impact on visual estimation differs from its effect on statistical analysis, suggesting that visual estimations may not align with statistical summaries [28]. The estimation of summary statistics from visualizations is influenced by the attributes of the visualization design and underlying data properties [29]–[31].

Most of this research focuses on scatterplots using synthetic data. In contrast, our study examines whether these findings hold when line charts depict real-world time series augmented with controlled levels of synthetic noise, focusing on how viewers preserve key features as visual clutter increases.

B. Utility of Sketching

Sketching is an unbounded medium for generating visual content [32], allowing individuals to highlight areas of interest [18]. Therefore, various studies have adopted pen-and-touch interfaces for chart sketching [11], [12], [33], recognizing their adaptability and effectiveness in generating more expressive and customized visual representations. Unlike selection-based methods, free-form sketching avoids priming participants toward predefined visual features, capturing more authentic perceptual priorities [11], [12], [18]. Although explicitly querying participants about their intentions could clarify perceptual priorities, we intentionally avoided this to preserve the spontaneity and natural authenticity of their sketching behavior. Nonetheless, inherent subjectivity in sketching, especially among non-experts, remains a methodological challenge [32].

While visual query systems leverage sketching to represent queries visually, these systems face challenges, including the inherent ambiguity in sketches [18] and the frequent inability or reluctance of users to depict their intentions accurately [15]. Prior work has shown how user input through free-hand sketches [13] or example-based labeling [34] can effectively tune a system to human perceptual priorities. Our work

complements these efforts by systematically studying which visual features people naturally prioritize for preservation, followed by a validation evaluation where participants chose sketches and explained their reasons to reduce the subjectivity associated with sketching.

Collectively, these approaches provide valuable insights into feature prioritization and contribute to the development of improved visual query systems.

III. BACKGROUND ON VISUAL FEATURES

Existing literature presents different ways of defining and identifying visual features such as trends, periodicity, peaks and valleys, and noise within time series data. While not exhaustive, these features align with the principles of scagnostics [35], [36]. Adapting scagnostics for line charts allows these features to be interpreted as measurable properties. Additionally, while scagnostics measures do not explicitly define noise, it remains a critical feature to explore due to its impact on perception [9]. This section establishes a foundation for the upcoming analysis by reviewing established methods for estimating and quantifying these visual features and explaining their role in the current study.

A. Defining Trend

Trends are widely used to interpret time series data [37], but their definition varies across fields. In mathematics and statistics, they are associated with monotonic patterns, event rates, and timing [38], while in climate science, they represent systematic variations in parameters like temperature and extreme weather events [39], [40].

In the visualization community, trends are generally seen as upward or downward movements over time [41]–[43], often depicted using trend lines or regression lines [5], [44]–[47]. However, real-world data often lacks clear trends [37]. In such scenarios, locally estimated scatterplot smoothing (LOESS) regression [48] is useful for capturing the main patterns [37]. Likewise, in signal processing, the Fast Fourier Transform (FFT) is often used to analyze time series [49] and detect trends by removing higher frequencies [50]–[52], based on Duda et al. [51]’s definition of trend as the non-periodic component of a time series. Our analysis employed LOESS regression and FFT with a low pass filter to capture the non-linear trends in our real-world data.

B. Detecting Periodicity

A time series is considered periodic if values repeat at (almost) regular intervals over time [53]. Identifying periodicity in time series for a number of tasks, including aiding interpretation [54], detecting anomalies [55], and forecasting [56]. Several approaches exist for detecting periodicity including FFTs [50], [57], [58], wavelet transformation [59]–[61], and dynamic time warping [62]–[64]. In this study, we used FFTs to detect prominent periodic patterns by identifying the most prominent peak in the frequency domain as described by Musbah et al. [57].




Fig. 2. Line charts for all nine datasets used in this study.

C. Detecting Peaks and Valleys

Peaks are local maxima [16], [65] and valleys local minima [66] in time series data. Multiple studies explore methods for peak and valley detection, such as entropy-based methods [67], wavelet transformation [68], [69], and Topological Data Analysis (TDA) [65]. This paper adopts the TDA approach due to its robustness to noise and its effectiveness in preserving prominent peaks and valleys [16].

D. Detecting Noise

Time series data often contain random fluctuations that mask the underlying signal, referred to as noise [70]. Various techniques, such as filtering or smoothing [16] and downsampling [17], are employed to mitigate noise. Additionally, frequency domain low-pass filters, including the Butterworth [71] and Chebyshev [72] filters, are effective in noise reduction. This study adopts the approach of employing an FFT with a high-pass filter to isolate noise components within the signal [57], [73], [74]. Given that noise can complicate the interpretation of visualizations [75], this research also incorporates the use of Pixel Approximate Entropy (PAE) [76] as a metric to estimate the noise level in line charts.

IV. PRIMARY STUDY DESIGN

Understanding how individuals attend to and prioritize visual features in noisy line charts offers valuable insights for designing compelling visualizations. To this end, we conducted a laboratory study with 20 participants to examine which features of line charts they would recreate under varying noise levels. Participants first completed a stenography task where they were asked to re-draw (see Sect. IV-C) line charts of nine real-world datasets (see Sect. IV-A), each with one of five levels of Gaussian noise (see Sect. IV-B). We then assessed feature retention in their sketches compared to the stimuli, both quantitatively and qualitatively (see Sect. V). Participants were later invited back for a follow-up study to select preferred sketches and provide a verbal rationale for their choices (see Sect. VII).

A. Datasets and Properties

We selected nine datasets (see Fig. 2), including: Daily Apple Stock Closing Price [77]; Radio Astronomy Signal [78]; Monthly Homicide Rate in Chicago [79]; Daily High Temperature at JFK Airport [80]; Daily Doge Coin Closing Price [81]; Single Channel of EEG Data [82]; Weekly US Domestic Flights [83]; Monthly Tourists in New Zealand [84]; and US Monthly Unemployment in Agriculture [85]. Each of the datasets demonstrated properties (see Table I) often encountered in real-world data visualization scenarios, including long-term stock data with overall directional trends (e.g., Fig. 2a and Fig. 2d), cyclical or periodic patterns seen in climate data (e.g., Fig. 2g), prominent peaks and valleys (e.g., Fig. 2b and Fig. 2f). Importantly, most datasets demonstrated multiple properties simultaneously.

B. Introducing Noise

One aspect of our experimental design was examining how noise impacts visual features in participants' sketches. While each input dataset contained some inherent noise, we initially treated them as noise-free (i.e., signal-only). Afterwards, Gaussian noise, chosen for its resemblance to real-world noise [86], was introduced across datasets at five distinct signal-to-noise ratio (SNR) levels:

- **None:** The dataset remained unchanged, acting as the control group without additional noise;
- **Low:** SNR of ≈ 30 dB after adding noise;
- **Medium:** SNR of ≈ 20 dB after adding noise;
- **High:** SNR of ≈ 10 dB after adding noise; and
- **Max:** SNR of ≈ 5 dB after adding noise representing a scenario with substantial noise interference.

One drawback of using real-world series is that each contains some native noise, so it is important to note that the same noise level had different effects on different datasets (see Fig. 3) due to data density and structure variations. For example, some were denser (e.g., Fig. 2h / Fig. 3h) and others sparser (see Fig. 2i / Fig. 3i), while some contained smaller peaks (e.g., Fig. 2a / Fig. 3a) and others contained more prominent peaks (e.g., Fig. 2b / Fig. 3b). Figures of each dataset with all noise levels can be found in the supplemental material. While we could have used fully synthetic signals to control noise, working with real data allows preserving the




Fig. 3. Line charts for all nine datasets with max noise (SNR ≈ 5 dB).

natural trends, periodicity, peaks, and valleys that give the results high ecological validity, while still varying only the noise level as the experimental factor.

C. Experiment

Stimuli: We created 45 stimuli with 9 datasets (within-subject) crossed with 5 noise levels (between-subject). Participants saw each dataset once, in the same order, at a randomly selected noise level. The noise levels were shuffled such that all 45 stimuli appeared for every five participants. We used Vega-Lite [87] to generate the line charts. To focus solely on features observed in the chart and to avoid copying the data, none of the generated line charts included axes, gridlines, or ticks. Each dataset's minimum and maximum values were used to scale the y-axis to minimize the white space.

Task: Participants viewed each of the nine line charts sequentially on a monitor at a resolution of 950 x 375. They used an Apple iPad Air (5th Gen) running a web application that captured their sketch via an Apple Pencil (2nd Gen), with both the monitor and iPad canvases set to the same 2.53:1 aspect ratio. Participants were tasked with re-drawing the line chart without any prior context about the data used to ensure that their sketches reflected purely visual interpretations rather than domain-driven bias. During the sketching process, the stimulus line chart remained visible at all times. They were encouraged to emphasize elements in their sketches that they deemed significant, with the understanding that an exact replication of the shown stimuli was not required. The web application would record a single stroke only (i.e., the pen always had to stay on the canvas). They were informed that lifting the pen would prompt them to accept or reset their sketch, with unlimited resets allowed. Although no time limit was imposed, the experiment generally took 20 minutes, with one participant taking approximately 40 minutes.

Participants: We recruited 20 participants (5 females and 15 males) who were engineering students (1 undergraduate and 19 graduate). Each participant had experience conducting research in the visualization domain and affirmed their familiarity with line charts. Participants did not receive any compensation. The study was IRB-approved.

TABLE I

A DESCRIPTION OF THE DATASETS USED AND THE PROPERTIES WE OBSERVED, INCLUDING OVERALL TREND ($\nearrow$ UP, $\searrow$ DOWN, OR $\rightarrow$ CONSTANT), PERIODIC PATTERNS, PROMINENT PEAKS AND/OR VALLEYS.

Dataset	Description	Overall Trend	Periodic Pattern	Peaks and/or Valleys
Apple	Daily Apple Stock Closing Price [77]	$\nearrow$		$\checkmark$
Astronomy	Radio Astronomy Signal [78]	$\rightarrow$		$\checkmark$
Chicago	Monthly Homicide Rate in Chicago [79]	$\rightarrow$	$\checkmark$	$\checkmark$
Temperature	Daily High Temp. at JFK Airport [80]	$\rightarrow$	$\checkmark$	
Doge	Daily Doge Coin Closing Price [81]	$\searrow$		$\checkmark$
EEG	Single Channel of EEG Data [82]	$\nearrow$		$\checkmark$
Flights	Weekly US Domestic Flights [83]	$\nearrow$		$\checkmark$
Tourists	Monthly Tourists in New Zealand [84]	$\nearrow$		$\checkmark$
Unemployment	US Monthly Unemploy. in Agriculture [85]	$\rightarrow$	$\checkmark$	$\checkmark$

V. ANALYSIS

We employ a mixed-methods approach to assess the behaviors observed in the participant-sketched line charts. We performed a qualitative evaluation (see Sect. V-A) to identify which visual features participants retained, while we performed a quantitative analysis (see Sect. V-B) to determine the accuracy of this feature replication in the resulting sketches. This combination captured both numerical similarity and semantically significant representations.

A. Qualitative Analysis

We adopted a thematic coding process for qualitative analysis to systematically examine key aspects of participant sketches, which involved addressing questions about four factors of the sketches: the overall trend, periodicity, prominent peaks and/or valleys, and the noise level (for datasets where these features were present, see Table I). Each sketch was shown juxtaposed over the stimuli as a reference, as shown in Fig. 4. The coders were given the following questions with three options each:

Feature	Coding Questions	Options
Overall Trend	How well does the trend in the sketch (Red) match the trend in the stimuli (Blue)?	Very well / Somewhat / Not at all
Periodicity Match	How well does the periodicity in the sketch (Red) match the periodicity in the stimuli (Blue)?	Very well / Somewhat / Not at all
Peak & Valley Preservation	How many of the peaks and/or valleys of the stimuli (Blue) are preserved in the sketch (Red)?	Most / Some / None of them
Noise Match	How well does the noise in the sketch (Red) match the noise in the stimuli (Blue)?	Very well / Somewhat / Not at all

Two authors of the paper independently coded the sketches. Before coding, they aligned their understanding of trends, periodicity, peaks, valleys, and noise to establish a consensus code. The independent coding yielded a Cohen’s kappa score of 0.733 with 445 matches and 115 mismatches, indicating a “substantial” agreement [88]. Moreover, all 115 mismatches indicated only minor interpretive differences (e.g., one coder classified a sketch as matching the original trend *Very well* while another classified it as *Somewhat*) rather than fundamental disagreement. Therefore, results from the first coder




Fig. 4. Participant sketches (red) are overlaid on the stimuli (blue) of the Chicago dataset at four different noise levels. Coders evaluated the sketches and assigned them to categories based on retained features.

were used for all analyses. The complete coding is included in the supplemental material.

B. Quantitative Analysis

We conducted a quantitative analysis of participant sketches through first pre-processing to correct artifacts and normalize sketches (see Sect. V-B1) and then establishing metrics for measuring the visual features (see Sect. V-B2). All the figures from this analysis are included in the supplemental material.

1) Pre-processing Participant Sketches:

Artifact Removal: Due to the free-hand nature of the task, the participant sketches could have artifacts that violated the requirements of time series data, temporal ordering in particular (i.e., data must be in chronological order). This would arise in sketches as strokes that would go backward and/or contain loops (see Fig. 5a). To address this issue, we implemented a correction process. Since the sketch was done in a single stroke, given two temporally consecutive points, x_i and x_{i+1} , if the x-coordinate of the first data point was larger than the second (i.e., $x_i > x_{i+1}$), the x-coordinate of the second data point was set to be 0.1 larger than the first (i.e., $x_{i+1} = x_i + 0.1$). Here, 0.1 was chosen as a minimal increment on our 950 px-wide canvas, small enough to correct ordering without noticeably altering the sketch’s shape. This approach resolved these issues, such as the loops as shown in Fig. 5b.

Normalization: The next step involved normalizing the data, such that the participant sketch and the stimuli can be compared point-by-point. First, in the x-direction the time series participant sketches were stretched or compressed to fit the canvas. Next, both the stimuli and participant sketched time series were upsampled to 9,500 samples (10x the original resolution), using linear interpolation (see Fig. 5c). In the y-direction, both sketches and the stimuli were aligned on center by subtracting their average value.

2) *Estimating Features and Metrics for Similarity:* After converting participant sketches into time series with equal data points to the stimuli, we quantified how well they preserved features discussed in Sect. III, including trends, periodic patterns, prominent peaks and valleys, and noise. However, measuring time series similarity is complex, as prior work shows the human perception of similarity is subjective [31], [89]. Ding et al. [90] found no universally superior similarity measure, as accuracy is often dataset and domain-dependent. Since selecting metrics tailored to data and domain was beyond our scope, we opted for generalizable metrics.

Estimating and Measuring Trends: The ground truth trend in real-world datasets is often unavailable [37], which is why it is estimated. We pursued two methods for trend estimation: frequency domain analysis and LOESS curve fitting. We used Fast Fourier Transformation (FFT) for frequency domain analysis on each line chart (see Fig. 6b), applying a low pass filter to approximate the trend. Specifically, we considered only the first one-third of the log scale frequency components, yielding an estimated trend as depicted in Fig. 6c. As an additional validation, we utilized the LOESS curve fitting technique (see Fig. 6a), inspired by the work of Wang et




Fig. 5. (a) A participant sketch of the *Apple* dataset with max noise (SNR $\approx$ 5dB), which is not a valid time series, as pointed out by the orange enclosure. (b) After artifact removal, the data is now a valid time series. (c) Data is normalized in order to compare it to the stimuli data.

al. [37] to captures the deterministic component of variation in a time series. Consistent with their methodology, we employed a span value of 0.4 to capture large-scale patterns.

We evaluated several common metrics (L^∞ , L^2 , L^1 -norms, and Dynamic Time Warping (DTW)), observing consistent results across all (see supplemental material). We report L^2 -norm due to its widespread use in assessing overall differences in continuous data and its sensitivity to cumulative variation. L^2 -norm, $\ell_2(X, Y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$ quantifies the overall Euclidean distance between the trends of the stimuli and participant sketches. In Fig. 8a, it is shown as the cumulative squared differences across all points in the lines.

Estimating and Measuring Periodicity: In datasets featuring periodic patterns, researchers often utilize frequency domain analysis, as indicated by prior studies [50], [57]. According to Musbah et al. [57], significant peaks in the FFT represent frequencies corresponding to the periodicity in the signal. Following this methodology, we identified the maximum frequency component in the FFT for line charts with periodic characteristics. We transformed it back to the time domain, as depicted in Fig. 6d. We calculated the amplitude and period difference between the estimated periodicity from the stimuli and participant sketches as illustrated in Fig. 8b. The estimated periodicity in the stimuli is represented by blue line and the participant sketch by a red line. If the number of period in the blue line is T_1 and T_2 in the red line, the difference in period is $\delta T = |T_1 - T_2|$. Similarly, If the amplitude of the blue line is A_1 and A_2 in the red line, the difference in amplitude is $\delta A = |A_1 - A_2|$.

Estimating and Measuring Peaks and Valleys: To identify

peaks and valleys, we utilize techniques from Topological Data Analysis. Peaks and valleys can be defined as the local minima and maxima [16], [67]. The local minima and maxima from a line chart are paired using a process described in detail in [65] and placed into an object called a persistence diagram [91]. Fig. 8c shows examples of such peaks. To compare line charts, the persistence diagrams from the two line charts, $\bar{X}$ and $\bar{Y}$ are compared using bottleneck distance, which identifies the peaks with the maximum variance. Let η be a bijection between the two sets, The bottleneck distance is $W_\infty(\bar{X}, \bar{Y}) = \inf_{\eta: \bar{X} \rightarrow \bar{Y}} \sup_{\bar{x} \in \bar{X}} \|\bar{x} - \eta(\bar{x})\|_\infty$.

Estimating and Measuring Noise: Noise can be considered as the high-frequency component in the data. Therefore, we employed FFT to identify and isolate the noise components as suggested by previous research [50], [57]. To ensure a consistent and generalized interpretation of noise, we applied a high-pass filter to the FFT. Specifically, we focused on the last one-third of the frequency components and transformed them back to the time domain, as illustrated in Fig. 6e.

The area preserved was calculated for the estimated noise (see Fig. 6e). The change in area, $\delta a(X, Y) = |\sum x_i - \sum y_i|$, was computed by taking the difference between the stimuli (X) and the participant sketch (Y). The process is shown in Fig. 8d. The sum of the blue bars represents the area covered by the estimated noise in the stimuli, and the sum of the red bars indicates the participant sketches. In addition, a second metric was used to estimate the noise, namely Pixel Approximate Entropy (PAE). PAE is a measure of the perceptual complexity in line charts [76], and served as an indirect measure of noise level, with higher entropy values indicating greater noise.

VI. RESULTS

We present three main findings. First, the qualitative analysis revealed that individual **participants did not consistently exhibit the same behavior across all stimuli** but instead clustered into three behavior types: **Replicator**, when they replicated all features of the stimuli including noise in sketches; **Trend Keeper**, when they retained only the general trend in their sketches with or without noise; and the **De-noiser**, when they removed the noise from the shown stimuli and retained all the other features in sketches.

Second, we found that viewers **retained most features in line charts across varying noise levels**. However, the same noise level introduced different amounts of visual clutter depending on the dataset. Despite this, most participants effectively retained features, showing robustness to noise.




Fig. 6. Illustration of feature extraction. (a) shows the *Chicago* dataset in a line chart in blue, with its LOESS fit curve trend in green. (b) The Fast Fourier Transform (FFT) decomposes the data into frequency components. (c) The inversion of low-frequency components back to the time domain provides the estimated trend. (d) The most prominent FFT peak is transformed to the time domain to estimate the periodic pattern. (e) The high-frequency components are converted to the time domain to estimate noise.




Fig. 7. Participant sketches overlaid on stimuli for four datasets, categorized into clusters based on feature inclusion or exclusion in the sketches.

Finally, while participants retained features, they did so either *faithfully* (with low quantitative error) or *semantically* (with higher quantitative error but capturing the presence of features qualitatively). For example, in some sketches, the *amplitude* and *period* of the periodic pattern differed from the stimuli, yet they still conveyed the *presence* of these features.

A. Clustering Participant Behavior

Due to the free-form nature of participant sketches, there is no reliable numeric measure to group them by shared drawing strategies, and the artifact removal and normalization required for any quantitative analysis would erase the nuances of sketches we aimed to capture. Consequently, we applied qualitative thematic analysis in Sect. V-A to identify behavior

patterns directly from the sketches. Based on the analysis, we formed three clusters of participant behavior based on the inclusion or exclusion of visual features in sketches (see Fig. 9a):

- **Replicator**: These sketches (e.g., Fig. 7a-Fig. 7d) generally maintained all features in the stimuli. Trends and periodicity (if present) were preserved either *very well* or *somewhat*, along with *most* or *some* peaks and valleys (if present). Noise was also replicated *very well* or *somewhat*.
- **Trend Keeper**: These sketches (e.g., Fig. 7e-Fig. 7h) predominantly maintained trends *very well* or *somewhat*, but did not preserve periodicity, peaks, or valleys, even when these features were present in the stimuli. Noise was optionally preserved in these sketches.
- **De-noiser**: These sketches (e.g., Fig. 7i-Fig. 7l) maintained all the features of the stimuli like trends and periodicity (if present in the stimuli) *very well* or *somewhat* and *most* or *some* of the peaks and valleys (if present in the stimuli) but removed the noise present in the stimuli. The only difference between a Replicator and a De-noiser is the presence of noise.

Three **Anomaly** sketches did not fit any cluster and were excluded from further analysis. Feature inclusion/exclusion for these sketches is shown in Fig. 9a, with additional figures in the supplemental materials.

Observations of these behavior patterns varied both across and within participants. Fig. 9b shows the clusters assigned to all participant sketches, sorted by the proportion of **Replicator**, **Trend Keeper**, and **De-noiser** behavior. The majority of sketches, 64% (115/180), were labeled **Replicator**, 18% (33/180) **Trend Keeper**, and 16% (29/180) **De-noiser**. Four participants [P01-P04] consistently exhibited **Replicator** behavior




Fig. 8. Illustration of similarity measurements used in our approach. (a) shows measuring trend similarity with a total variation using ℓ_2 . (b) illustrates comparing the period and amplitude, where δT denotes the difference in the number of periods, and δA denotes the difference in amplitude between the blue and red signals. (c) illustrates matching the peaks & valleys in the data, where W_∞ measures the largest variation pointed with arrow shows. (d) shows measuring noise levels with area preservation, where δa measures the difference between the two shaded areas.




Fig. 9. The figure illustrates the distribution of different behavioral patterns—Replicator, Trend Keeper, De-noiser, and Anomaly—across features, participants, noise levels, and datasets. (a) An Upset plot [92] of the visual features preserved by cluster type is shown. (b) The color-coded 2D grid depicts each participant’s behaviors in their 9 sketches. It was organized in decreasing order to highlight the prevalence of Replicator behavior and subsequently Trend Keeper, De-noiser, and Anomaly. (c) The distribution of these clusters at different noise levels is shown, where each bar signifies a distinct noise level. (d) The cluster distribution across the 9 datasets.




Fig. 10. Participant sketches superimposed on their corresponding stimuli, with an SNR $\approx$ 5dB. In these datasets, participants demonstrated robustness to noise by capturing features, such as trends, periodicity, peaks, and valleys, in their sketches despite the presence of a significant level of noise. These sketches were classified as Replicator

across all stimuli, while others showed a mix of different behaviors in reaction to different stimuli.

Cluster Relationship to Noise and Dataset: To understand variation in participant behaviors, we examined cluster distribution across noise levels (see Fig. 9c) and datasets (see Fig. 9d). The analysis revealed a dependency of behaviors on both the dataset type and the noise level. As noise increases, Replicator and De-noiser behaviors decrease, while Trend Keeper behavior increases (Fig. 9c).

Among all the datasets, Temperature had the most number of De-noiser sketches, while Unemployment had the most Replicator sketches (see Fig. 9d). One possible reason behind this could be the nature of the datasets. Temperature dataset’s periodic pattern (see Fig. 10b) remained visible even at higher noise levels, making it less susceptible to noise. In contrast, the Unemployment (see Fig. 10d) dataset remained less cluttered under high noise, making replication easier. The Tourists dataset had no Trend Keeper sketches, as it lacked periodicity or prominent peaks and valleys (Table I). As such, sketches preserving trends without noise fell into Replicator

or De-noiser clusters.

B. Robustness to Noise

We observed varying participant robustness to noise, with feature preservation differing across datasets. We calculated a regression line per dataset and quantitative metric to measure error increases in participant sketches as noise increases (e.g., red lines in Fig. 11). The slope between no noise and max noise (SNR $\approx$ 5dB) represents the percentage error increase. Results for all datasets, features, and metrics are in Table II. To illustrate distinct slope patterns (upward, downward, or constant), we included plots from specific datasets in the paper, with additional plots available in the supplemental materials.

Trend Robustness to Noise: Fig. 11a and Fig. 11b show downward regression lines, indicating reduced error (difference between the estimated trend in stimuli and participant sketches) as noise increases in the Tourists dataset. Conversely, Fig. 11c and Fig. 11d show upward trends, indicating increased error for the Unemployment as noise increased.

As seen in Table II, apart from the Temperature and Unemployment datasets, most datasets show flat or even lower trend preservation errors as noise increases. This suggests that trend preservation depends on both dataset characteristics and noise levels. However, given the general error reduction across most datasets for substantial increases in noise, we conclude that participants’ trend preservation was generally robust to noise.

Periodicity Robustness to Noise: Participants identified and preserved periodicity in their sketches despite varying noise levels. Fig. 12 shows the amplitude and period difference between the stimuli and all the participant sketches for the Temperature dataset. The regression line in the amplitude plot (Fig. 12a) indicates a decrease in error with higher noise levels, while the period plot (Fig. 12b) shows a consistent error rate. Table II also shows that for both Chicago and Unemployment datasets, error decreased for amplitude and period differences, suggesting a general insensitivity to noise

levels when estimating amplitude and periodicity for these datasets.

Quantitative analysis confirms participants identified periodicity across noise levels, though the accuracy of period count or amplitude varied. As shown in Fig. 9d, 53 of 60 sketches from periodic datasets (Chicago, Temperature, Unemployment) were classified as **Replicator** or **De-noiser**, indicating participants' robustness in identifying and reproducing periodicity despite noise.

Peak and Valley Robustness to Noise: Preserving peaks and valleys showed different patterns than trends and periodicity as noise increased. Fig. 13 shows the bottleneck distance between the stimuli and participant sketches for two datasets. Fig. 13(Chicago) shows an increasing error with higher noise levels in the Chicago dataset, a tendency also seen in Apple, Flights, and Unemployment (Table II). In contrast, Fig. 13(EEG) shows decreasing error in the EEG dataset, a pattern also seen in Doge and Astronomy (Ta-

ble II). As noise can create spurious peaks and valleys, only very prominent peaks remained visible with increased noise, making these features harder to preserve in some datasets.

In conclusion, feature preservation amidst noise varies by dataset. Datasets like Astronomy, Chicago, Doge, EEG, and Tourists show robustness, enabling accurate feature replication even under high noise. This may result from fewer data points or more prominent features. In contrast, Apple, Temperature, Flights, and Unemployment are more susceptible to noise, leading to greater errors. Peaks and valleys were harder to preserve in Apple and Unemployment, while trend preservation was more challenging in Temperature and Unemployment. The higher errors in Apple and Flights may be due to more data points, increasing visual clutter at lower SNR levels and blurring peaks and valleys (Fig. 3a, Fig. 3f).

C. Semantic and Faithful Representation of Features

The analyses in Sect. VI-A show that participants preserved many salient features in their sketches. Where features were preserved accurately in sketches, we deem this representation *faithful*. Where features were represented visually, but in ways that diverge from their quantitative properties, we deem this representation *semantic*. Quantitative analysis indicated that periodicity and noise were mainly represented semantically, while trends, peaks, and valleys were more faithfully preserved. However, we acknowledge the ambiguity and subjectiveness between the *semantic* and *faithful* distinction.

1) More Faithful Features: Trends, Peaks and Valleys

Trends: The qualitative analysis showed that participants effectively preserved the trends, with all clusters coded as either *Very well* or *Somewhat* preserving trends. Fig. 11 shows the




Fig. 11. Scatterplot of the L^2 -norm of estimated trends via FFT and LOESS across Tourists and Unemployment dataset. Each point represents the L^2 -norm between the participant sketch and the stimuli. The point colors correspond to the qualitative coding for the associated participant sketch.

error in the estimated trends across noise levels for *Tourists* and *Unemployment*. As discussed in Sect. VI-B, participants exhibited robustness to noise, consistently preserving trends across all datasets and noise levels.

Peaks and Valleys: As noted in Sect. VI-B, participants were not consistently robust to noise when preserving peaks and valleys. However, Fig. 13 shows that sketches retaining most peaks and valleys exhibited a lower Bottleneck distance, indicating a closer resemblance to the stimuli with lower error. Despite phase shifts in sketches where peak and valley locations differed, prominent peaks and valleys were faithfully represented, as seen in Fig. 10a and Fig. 10c. This tendency was also observed in other datasets. This indicates that the qualitative code of participant sketches regarding preserving prominent peaks and valleys matched the quantitative measurement, further supporting a faithful representation.

2) More Semantic Features: Periodicity and Noise

Periodicity: Analysis in Sect. VI-B showed that participants' sketches did not precisely replicate the amplitude and period of the stimuli. Yet, they consistently conveyed its periodic characteristics. Fig. 12a and Fig. 12b show the difference in amplitude and number of periods preserved, respectively, for the *Temperature* dataset. For this dataset, the number of periods remained the same across all the stimuli despite varying noise levels. This shows an apparent discrepancy: despite the uniformity in the number of periods across all stimuli, the sketches predominantly did not capture the precise period of the stimuli. However, most sketches were coded as preserving periodicity *Very well* or *Somewhat*, with only one




Fig. 12. Scatterplot displaying (a) amplitude difference and (b) period difference between all participant sketches and the stimuli across *Temperature* dataset. The data points are color-coded based on their assigned code in the qualitative analysis.




Fig. 13. Scatterplot displaying the Bottleneck distance between prominent peaks and valleys for *Chicago* and *EEG* dataset. Data point colors correspond to the qualitative coding for the associated participant sketch. The regression line in red is derived from participant sketches.

exception marked as *Not at all*. This suggests participants grasped the underlying periodicity of the data but did not reproduce it precisely, as demonstrated in Fig. 10b. This semantic representation pattern without exact replication was also observed in the *Chicago* and *Unemployment* datasets, which also exhibit periodic trends.

Noise: We observed a similar pattern in the participants' sketches for preserving noise. Using Pixel Approximate Entropy (PAE), we compared noise levels in the stimuli to those in sketches. Fig. 14a shows PAE values for stimuli with various noise levels and corresponding participant sketches, with each data point for sketch color-coded based on qualitative coding. The figure reveals that captured noise in sketches was generally lower than in the stimuli, even when qualitatively coded as matching noise *Very well*. This indicates participants tended to underrepresent the noise in their sketches, but semantically, they captured the noisiness as shown in Fig. 14c and Fig. 14d. However, noise levels in sketches were closer to stimuli noise at SNR ≈ 5 dB (Fig. 14b). One sketch Fig. 14e captured the noise in the shown stimuli both qualitatively and quantitatively, indicating a more faithful representation. However, semantic representation was more common across all datasets. Additional PAE values and area difference plots are available in the supplemental material.

VII. VALIDATION STUDY DESIGN

Our primary visual stenography experiment had some limitations, such as a lack of verbal feedback from participants regarding their thought processes and imposed sketching constraints, such as using a single continuous stroke. Moreover,




Fig. 14. The (a) Pixel Approximate Entropy (PAE) and (b) Area Preserved by the estimated noise of all participant drawn sketches for the *Tourists* dataset. Data point colors denote assigned qualitative code. The blue line in (a) indicates PAE for stimuli at different noise levels. The annotated data points as ① and ② highlight two sketches (c-d), coded as closely matching the noise level in the stimuli. However, (a) reveals these sketches have lower noise levels than their stimuli. In (b), the annotated data points confirm a discrepancy in noise preservation, suggesting a semantic interpretation of noise in these sketches. Whereas the annotated datapoint as ③ highlights a sketch (e) which shows more faithful representation of noise.

most people are not faithful artists [32] and struggle to depict what they want to convey accurately [14]. To address these limitations, we designed a follow-up evaluation to validate the findings from the visual stenography experiment and gain a deeper insight into feature prioritization and reasoning. Participants were presented with reference stimuli and pre-generated sketches from the primary experiment. This approach allowed participants to select a well-drawn sketch that they believed best represented the important aspects of the reference stimuli, regardless of their sketching abilities.

A. Experiment

Stimuli: Participants viewed nine reference stimuli (one per dataset) in a fixed order, followed by corresponding sketches we selected from the initial experiment. For each dataset, we selected sketches from the noise level with the most diversity in behavior cluster, excluding *Anomaly*, which we skipped due to their rarity and lack of consistent patterns. We prioritized higher noise levels to explore noise perception. An exception was the *Flights* dataset: we selected SNR ≈ 20 dB instead of SNR ≈ 5 dB, as it uniquely produced *Trend Keeper* and *De-noiser* behavior, with no *Replicator* behavior. For each task, one sketch per cluster type was shown. When multiple sketches were available for a cluster and noise level, sketches were selected at random to reduce bias. A figure illustrating the selected noise level for each dataset is provided in the supplemental material.

Task: For this task, each reference stimulus was displayed for 10 seconds before disappearing, after which participants were presented with multiple sketches. Participants were asked the following questions:

Q1. *Based on your understanding of the reference image shown earlier, which one would you prefer to sketch to capture its important aspects? If none of them capture what you find important, feel free to say so.*

Q2. *What is the reason behind your choice?*

Like the visual stenography task experiment described in Sect. IV-C, participants were encouraged to select a sketch that preserved the significant features. It was emphasized that this was not a matching task but a choice based on their interpretation of important aspects. No specific criteria or guidelines for assessing similarity were provided to avoid influencing their judgment. Instead, we aimed to capture their intuitive perceptions of visual features, an approach aligned with methods used in prior studies [18], [31].

After completing the sketch selection for all nine reference stimuli, participants were directed to a follow-up questions page. This page had a draggable list of features, including trend, periodicity, peaks and valleys, and noise. They were asked the following additional questions:

Q3. *When seeing a line chart, how do you prioritize its features? Drag and drop the following features to rank them from highest to lowest priority.*

Q4. *What other features do you think are important when viewing a line chart?*

We asked participants for a general ranking of features, without providing dataset examples or task context, to assess their intuitive, context-free prioritization of visual features.

The study was conducted online via Zoom. Consent was obtained before each session, and interviews were recorded. The experiment took around 20 minutes.

Participants: 12 of the 20 participants from the first study (2 females and 10 males) participated, with no compensation. All were graduate students. To minimize recall bias, we ran this validation study six months after the initial study. For clarity, participants in the validation study were re-anonymized using alphabetical labels.

B. Analysis

We used a mixed-methods approach to analyze participant responses. Sketch choices were categorized into **Replicator**, **Trend Keeper**, and **De-noiser** across participants, datasets, and noise levels. The feature priority rankings provided quantitative data, while participants' explanations required qualitative analysis to understand their thought processes.

Data Preprocessing: Responses were captured via Zoom transcripts, with one missing transcript generated using Descript. Transcripts were anonymized and organized by stimulus. An error in the EEG dataset selection (SNR $\approx$ 30dB instead of SNR $\approx$ 5dB) was discovered. As a result, it was excluded from the analysis, reducing it to eight datasets. All plots, including EEG data, are in the supplemental material.

Thematic Coding: Participants' explanations for sketch choices were coded by one of the authors to identify alignment with behavior patterns (i.e., **Replicator**, **Trend Keeper**, or **De-noiser**). Mixed behaviors or deviations from predefined clusters were noted. Feature ranking responses were analyzed to find common patterns. The anonymized transcripts and

coding are in the supplemental material.

C. Results

The follow-up study confirmed that participants' sketch choices varied across clusters, consistent with the initial visual stenography experiment. Fig. 15a shows that **Replicator** sketches were chosen most frequently (35/96), followed by **Trend Keeper** (31/96) and **De-noiser** (30/96). We did not ask participants whether they were prioritizing *faithful* versus *semantic* representations of features. We avoided asking this directly to prevent framing effects that might bias their selections, and so we cannot definitively know their underlying intentions. Nonetheless, the selections indicate that for periodicity and noise, participants favored sketches that conveyed those features *semantically* (i.e., clearly signaled their presence), whereas for trends, peaks, and valleys, they leaned more toward *faithful* representations. Sketches from all clusters were selected across noise levels (Fig. 15b) and datasets (Fig. 15c). However, the Apple dataset differed, with participants consistently choosing **Replicator** or **Trend Keeper** sketches despite being shown sketches from all clusters.

The alignment between chosen sketches and participants' rationale was analyzed qualitatively. Fig. 16 shows that most sketch selections (76/96) matched participants' rationale, while 20 instances showed mismatches, particularly in datasets like Chicago, Temperature, Flights, and Unemployment. The Temperature dataset had the most mismatches (8/12). In these cases, participants often provided rationale aligned with **Replicator** behavior but chose **Trend Keeper** sketches. Most participants selected sketches through elimination, rejecting those they found too noisy or insufficiently detailed. A similar elimination pattern was seen in Chicago, Flights, and Unemployment datasets. For Chicago, Participant C mentioned that their preference fluctuated between the simpler **Trend Keeper** and more detailed **Replicator** sketches. Regarding Flights, Participant H noted that both sketches captured key features, but the **De-noiser** sketch was clearer. Two instances did not align with predefined behavior patterns (Fig. 16). For Unemployment, Participant A focused on valley position over trends or noise, leading to a **De-noiser** selection. For Chicago, Participant K saw no




Fig. 15. Chosen sketches from different behavioral patterns—**Replicator**, **Trend Keeper**, **De-noiser**—across participants, noise levels, and datasets. (a) A color-coded 2D grid depicting the assigned behavior pattern of the sketch chosen by each participant for all 8 stimuli. The data was arranged in descending order **Replicator**, **Trend Keeper**, and **De-noiser** behaviors. (b) The distribution of these clusters at different noise levels is shown, where each bar signifies a distinct noise level, with colors indicating the behavioral clusters. (c) The cluster distribution across the 8 datasets. The color boxes below each dataset label represent stimuli sketch types associated with each dataset.

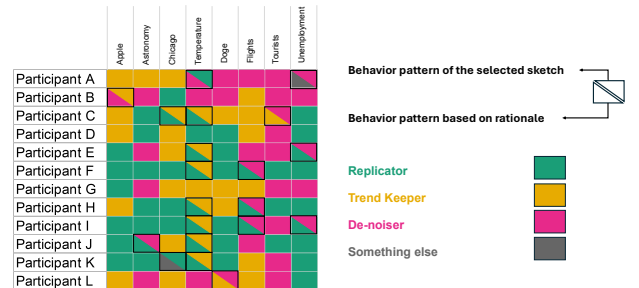


Fig. 16. The colored matrix compares the selected sketch's behavior pattern and the rationale for choosing them. Single-colored cells indicate a match between the behavior pattern between the selected sketch and rationale, while two-colored cells indicate a mismatch. The bottom triangle represents the behavior pattern demonstrated through the rationale, and the top triangle represents the assigned behavior pattern of the selected sketch during the primary study. Two Gray triangles indicate instances where the participant rationale did not align with predefined behavior patterns.

clear pattern in the reference image and chose the **Replicator** sketch’s recognizable peaks and valleys.

These findings reveal two key insights: first, the alignment between sketch clusters and rationale clusters supports the validity of the visual stenography task, indicating that the sketches generally preserved the key features participants aimed to capture. Second, line chart interpretation is more nuanced than initially assumed. Participants’ intentions often diverged from their final choices, suggesting that what they *intended* to preserve did not always align with what they *perceived* in the available sketches.

Visual Features in Line charts:

Participants ranked four features—trend, periodicity, peaks and valleys, and noise—by priority when viewing line charts. Most participants ranked trend as the highest priority, followed by peaks and valleys, periodicity, and noise (Fig. 17). This aligns with findings from the primary experiment, where sketches consistently preserved trends across clusters. Noise was consistently ranked as the lowest priority feature.




Fig. 17. The ranking distribution of features—trend, periodicity, peaks and valleys, and noise by participants. A rank of 1 indicates the highest priority, while 4 represents the lowest priority. The figure shows that the trend was ranked first most of the time, while noise was ranked last consistently.

While trends were prioritized without context, participants emphasized that priorities shift based on the data’s purpose. In specific contexts, noise may become a top priority. Participant C shared a personal example involving microscopy data, where analyzing noise was essential to evaluate the performance and sensitivity of the microscope. In other situations, noisy line charts can obscure critical features like peaks and valleys, underscoring the importance of filtering noise or adding annotations to enhance visibility. Participant F highlighted the difficulty in distinguishing between peaks and valleys and noise without data context. Participant C commented: *“I think these may depend a little bit on the task and the data you are dealing with. For instance, if you have a line chart showing a patient’s heart rate over time, extreme values are actually very, very important.”*

Some participants suggested additional relevant features such as contextual markers, axes, change points, average lines, and uncertainty boundaries. Although opinions varied on whether noise added value or was distracting, those who mentioned additional features preferred context-driven enhancements for transparent and accurate visualizations.

VIII. DISCUSSION

Given the findings of our primary study (Sect. VI) and follow-up validation study (Sect. VII-C), we identified three behavior groups while re-creating line charts. Further analysis showed that some people treated the same stimuli differently

regarding feature inclusion or exclusion. Also, a single individual would follow different patterns for line charts of different datasets or noise levels. This shows that what people see and re-create in line charts is subjective and depends on specific data properties, noise levels, and the person. Therefore, predicting what people will notice in a line chart is tough. However, generally speaking, people showed robustness to different noise levels by capturing trends and periodicity in their sketches. However, they were not quite as robust to noise in re-creating the peaks and valleys.

Another significant finding of our research highlights the tendency of participants to interpret and depict periodicity and noise in line charts through a semantic lens rather than with exact replication. However, the representation of trends, peaks, and valleys is more faithful to the data. These findings suggest the following implications for design:

Smoothing May Not Always Be Necessary to Show Trend and Periodicity:

Both studies indicate that individuals can effectively recreate trends and periodic patterns within datasets. Therefore, smoothing or de-noising may be unnecessary if the noise level is still reasonable and the primary objective is highlighting trends or periodicity. However, what counts as *reasonable* noise is often data and context dependent, as reflected in our analysis (Sect. VI-B). Since noise can also adversely affect judgments [24] or trust [93] in data, determining reasonable thresholds and their effects remains an open question that requires further investigation.

Highlight Peaks and Valleys in Noisy Line Charts:

Findings from the primary study reveal that noise significantly impacts the re-creation of peaks and valleys compared to periodicity and trends in line chart sketches. The validation study also reveals that people struggle to distinguish between true extrema in the data and fluctuations caused by noise. Some ways to ensure viewers notice peaks and valleys in noisy line charts include smoothing or using annotations to highlight significant peaks and/or valleys [43], [94]. Depending on the dataset, task, and features to be preserved, selecting an appropriate level of smoothing can make peaks and valleys more prominent. Such levels can be explored using tools like LineSmooth [16], which allow designers to adjust smoothing levels and compare multiple algorithms, including TopoLines [65], designed to reduce noise while preserving extrema.

Annotate Important Features: Both studies highlight the difficulty in predicting which aspects of a line chart will capture attention due to variability in behavior. Participants in the validation study noted the importance of annotations [95], making it a reliable approach to emphasize key features and ensure effective communication.

Make Visual Query Systems More Adaptable: Participants often abstracted periodic patterns, sketching varying periods and amplitudes that did not match the stimuli. Instead, they aimed to semantically convey the cycles in the data. Similarly, participants represented noise semantically rather than precisely replicating it. This observation sheds light on the potential challenges and considerations in visual querying, particularly in sketch-based query interfaces where diverse

algorithms may be necessary to interpret ambiguous or vague human input. These findings resonate with Correll et al.'s discussion on the potential irrelevance of specific data attributes, such as amplitude and noise, in sketch-based queries, calling such attributes invariants [18]. Our study suggests that periodicity could also be one such invariant.

Representation of Line Charts May Require Less Detail Than You Think: Individuals often express periodicity and noisiness in data more semantically, suggesting that detailed information on these aspects might not significantly enhance line chart representation. Unless the quantity of periods or the amplitude holds particular importance, incorporating such details into visual representation like sparklines [19] could be considered unnecessary. Conversely, trends, peaks, and valleys are typically more precisely or faithfully represented. Therefore, emphasizing these elements in representation could provide more value, offering a clearer understanding of the data's key features.

IX. LIMITATIONS

Our studies provide insights into user strategies for preserving visual features in line charts. Nevertheless, these findings must be considered within the scope of the study design choices we made. Notably, our primary study did not explicitly consider participants' sketching abilities, domain knowledge, dataset familiarity, or task framing, factors that could influence the interpretation and representation of visual features [8], [15]. To partially mitigate these concerns, we conducted a second validation study designed to assess user priorities and intentions independently of their sketching abilities. However, since this validation relied on sketches produced by participants from the first study, the influence of individual differences remains partially embedded within the evaluated materials.

We utilized participant-generated sketches, drawn in a single continuous stroke, to systematically capture user strategies for preserving visual features in line charts. This design simplified qualitative analysis by reducing variability and allowed direct comparisons across different conditions. However, the single-stroke limitation potentially restricted participants' ability to depict complex visual features accurately. The alternative approach of enabling multi-stroke sketches could offer greater representational flexibility, but at the expense of introducing analytical complexity, losing the clarity of comparison achieved by our chosen constraint.

Another limitation of our study is occasional mismatch between the qualitative codes and the quantitative metrics for sketches. For example, a sketch that we coded qualitatively as "Very well" at preserving a feature may show a relatively large L^2 -norm value, while a sketch coded as "Not at all" may score low as shown in Fig. 11. Such deviations can stem from multiple factors: the pre-processing steps (i.e., artifact removal, normalization, interpolation) that reshape sketches, the methods used to estimate features, or the choice of similarity metric itself. These factors can produce artificially high or low quantitative fidelity that does not fully reflect participants' original sketching choices.

In our second study, we intentionally limited the sketches presented per noise level (four per condition) to reduce participant fatigue and enhance focus, thus ensuring high-quality responses. This facilitated clear insights into user feature preferences and priorities. However, the small set of sketches might have restricted the representational breadth, thereby affecting the generalizability and robustness of our validation findings. Although a larger sketch pool would provide broader insights, it would also necessitate careful design to avoid cognitive overload and reduced participant attention.

X. CONCLUSION

In conclusion, we explored how individuals recreate features of line charts, focusing on trends, periodicity, peaks, and valleys. Our findings reveal three significant insights. First, participants exhibited diverse strategies in recreating visual features, balancing fidelity with abstraction based on perceived importance. Second, individuals demonstrated resilience in retaining underlying data patterns despite varying noise levels. Finally, our research indicates that people may represent periodicity and noise semantically or conceptually rather than precisely. In contrast, their depictions of trends, peaks, and valleys are more faithful to the actual data. Based on these findings, we propose guidelines for designing line charts that better align with human perceptual priorities, providing valuable direction for the visualization community to create more effective line chart designs. All the supplemental materials are available at: https://osf.io/4p3ze/?view_only=1dbd44ff26c24728b21a65f2907ccabb

REFERENCES

- [1] T. Mussweiler and K. Schneller, "What goes up must come down-how charts influence decisions to buy and sell stocks," *Journal of Behavioral Finance*, vol. 4, no. 3, pp. 121–130, 2003.
- [2] J. Assfalg, T. Bernecker, H.-P. Kriegel, P. Kröger, and M. Renz, "Periodic pattern analysis in time series databases," in *Database Systems for Advanced Applications (DASFAA)*, 2009, pp. 354–368.
- [3] O. Stein, A. Jacobson, and F. Chevalier, "The effect of smoothing on the interpretation of time series data: A covid-19 case study," *arXiv preprint*, 2023.
- [4] M. Correll, D. Albers, S. Franconeri, and M. Gleicher, "Comparing averages in time series data," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2012, pp. 1095–1104.
- [5] M. Correll and J. Heer, "Regression by eye: Estimating trends in bivariate visualizations," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2017, pp. 1387–1396.
- [6] G. E. Batista, E. J. Keogh, O. M. Tatav, and V. M. De Souza, "Cid: an efficient complexity-invariant distance for time series," *Data Mining and Knowledge Discovery*, vol. 28, pp. 634–669, 2014.
- [7] D. H. Kim, V. Setlur, and M. Agrawala, "Towards understanding how readers integrate charts and captions: A case study with line charts," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–11.
- [8] C. Xiong, L. Van Weelden, and S. Franconeri, "The curse of knowledge in visual data communication," *IEEE Trans. on Visualization and Computer Graphics*, vol. 26, no. 10, pp. 3051–3062, 2019.
- [9] A. Vredeveldt, G. J. Hitch, and A. D. Baddeley, "Eyeclosure helps memory by reducing cognitive load and enhancing visualisation," *Memory & cognition*, vol. 39, pp. 1253–1263, 2011.
- [10] J. Hu, F. Jia, and W. Liu, "Application of fast fourier transform," *Highlights in Science, Engineering and Technology*, vol. 38, pp. 590–597, 2023.
- [11] B. Lee, R. H. Kazi, and G. Smith, "Sketchstory: Telling more engaging stories with data through freeform sketching," *IEEE Trans. on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2416–2425, 2013.
- [12] H. Xia, N. Henry Riche, F. Chevalier, B. De Araujo, and D. Wigdor, "Dataink: Direct and creative data-oriented drawing," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–13.

- [13] P. K. Muthumanickam, K. Vrotsou, M. Cooper, and J. Johansson, "Shape grammar extraction for efficient query-by-sketch pattern matching in long time series," in *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2016, pp. 121–130.
- [14] D. J. Cohen and S. Bennett, "Why can't most people draw what they see?" *Journal of Experimental Psychology: Human Perception and Performance*, vol. 23, no. 3, p. 609, 1997.
- [15] D. J.-L. Lee, J. Lee, T. Siddiqui, J. Kim, K. Karahalios, and A. Parameswaran, "You can't always sketch what you want: Understanding sensemaking in visual query systems," *IEEE Trans. on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 1267–1277, 2019.
- [16] P. Rosen and G. J. Quadri, "Linesmooth: An analytical framework for evaluating the effectiveness of smoothing techniques on line charts," *IEEE Trans. on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1536–1546, 2020.
- [17] S. Steinarrsson, "Downsampling time series for visual representation," Ph.D. dissertation, University of Iceland, 2013.
- [18] M. Correll and M. Gleicher, "The semantics of sketch: Flexibility in visual query systems for time series data," in *IEEE Visual Analytics Science and Technology (VAST)*, 2016, pp. 131–140.
- [19] E. R. Tufte, *Beautiful evidence*. Graphics Press Cheshire, CT, 2006.
- [20] W. S. Cleveland and R. McGill, "Graphical perception: Theory, experimentation, and application to the development of graphical methods," *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 531–554, 1984.
- [21] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman, "The science of visual data communication: What works," *Psychological Science in the Public Interest*, vol. 22, no. 3, pp. 110–161, 2021.
- [22] D. A. Szafrir, S. Haroz, M. Gleicher, and S. Franconeri, "Four types of ensemble coding in data visualizations," *Journal of Vision*, vol. 16, no. 5, pp. 11–11, 2016.
- [23] M. Gleicher, M. Correll, C. Nothelfer, and S. Franconeri, "Perception of average value in multiclass scatterplots," *IEEE Trans. on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2316–2325, 2013.
- [24] N. Harvey Teresa Ewart Robert West, "Effects of data noise on statistical judgement," *Thinking & Reasoning*, vol. 3, no. 2, pp. 111–132, 1997.
- [25] L. Ciccione and S. Dehaene, "Can humans perform mental regression on a graph? accuracy and bias in the perception of scatterplots," *Cognitive Psychology*, vol. 128, p. 101406, 2021.
- [26] M.-H. Hong, J. K. Witt, and D. A. Szafrir, "The weighted average illusion: Biases in perceived mean position in scatterplots," *IEEE Trans. on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 987–997, 2021.
- [27] L. Ciccione, M. Sablé-Meyer, and S. Dehaene, "Analyzing the misperception of exponential growth in graphs," *Cognition*, vol. 225, p. 105112, 2022.
- [28] D. Reimann, C. Blech, N. Ram, and R. Gaschler, "Visual model fit estimation in scatterplots: Influence of amount and decentering of noise," *IEEE Trans. on Visualization and Computer Graphics*, vol. 27, no. 9, pp. 3834–3838, 2021.
- [29] Y. Kim and J. Heer, "Assessing effects of task and data distribution on the effectiveness of visual encodings," *Computer Graphics Forum*, vol. 37, no. 3, pp. 157–167, 2018.
- [30] L. Ciccione, G. Dehaene, and S. Dehaene, "Outlier detection and rejection in scatterplots: Do outliers influence intuitive statistical judgments?" *Journal of Experimental Psychology: Human Perception and Performance*, vol. 49, no. 1, p. 129, 2023.
- [31] A. Gogolou, T. Tsandilas, T. Palpanas, and A. Bezerianos, "Comparing similarity perception in time series visualizations," *IEEE Trans. on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 523–533, 2019.
- [32] M. Eitz, J. Hays, and M. Alexa, "How do humans sketch objects?" *ACM Trans. on Graphics (TOG)*, vol. 31, no. 4, pp. 1–10, 2012.
- [33] M. Mannino and A. Abouzied, "Expressive time series querying with hand-drawn scale-free sketches," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–13.
- [34] F. Lekschas, B. Peterson, D. Haehn, E. Ma, N. Gehlenborg, and H. Pfister, "Peax: Interactive visual pattern search in sequential data using unsupervised deep representation learning," in *Computer Graphics Forum*, vol. 39, no. 3. Wiley Online Library, 2020, pp. 167–179.
- [35] L. Wilkinson, A. Anand, and R. Grossman, "Graph-theoretic scagnostics," in *IEEE Symposium on Information Visualization*, 2005, pp. 21–21.
- [36] J. H. Friedman and J. W. Tukey, "A projection pursuit algorithm for exploratory data analysis," *IEEE Trans. on computers*, vol. 100, no. 9, pp. 881–890, 1974.
- [37] Y. Wang, F. Han, L. Zhu, O. Deussen, and B. Chen, "Line graph or scatter plot? automatic selection of methods for visualizing trends in time series," *IEEE Trans. on Visualization and Computer Graphics*, vol. 24, no. 2, pp. 1141–1154, 2017.
- [38] J. Lawless, C. Çiğsar, and R. Cook, "Testing for monotone trend in recurrent event processes," *Technometrics*, vol. 54, no. 2, pp. 147–158, 2012.
- [39] L. V. Alexander, X. Zhang, T. C. Peterson, J. Caesar, B. Gleason, A. Klein Tank, M. Haylock, D. Collins, B. Trewin, F. Rahimzadeh et al., "Global observed changes in daily climate extremes of temperature and precipitation," *Journal of Geophysical Research: Atmospheres*, vol. 111, no. D5, 2006.
- [40] L. S. Seregina, A. H. Fink, R. van der Linden, N. A. Elagib, and J. G. Pinto, "A new and flexible rainy season definition: Validation for the greater horn of africa and application to rainfall trends," *International Journal of Climatology*, vol. 39, no. 2, pp. 989–1012, 2019.
- [41] T. D. Wang, C. Plaisant, B. Shneiderman, N. Spring, D. Roseman, G. Marchand, V. Mukherjee, and M. Smith, "Temporal summaries: Supporting temporal categorical searching, aggregation and comparison," *IEEE Trans. on Visualization and Computer Graphics*, vol. 15, no. 6, pp. 1049–1056, 2009.
- [42] G. Robertson, R. Fernandez, D. Fisher, B. Lee, and J. Stasko, "Effectiveness of animation in trend visualization," *IEEE Trans. on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1325–1332, 2008.
- [43] A. Fan, Y. Ma, M. Mancenido, and R. Maciejewski, "Annotating line charts for addressing deception," in *ACM SIGCHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–12.
- [44] S. Ahmed, M. J. Islam, and H. Rajan, "Semantics and anomaly preserving sampling strategy for large-scale time series data," *ACM/IMS Trans. on Data Science (TDS)*, vol. 2, no. 4, pp. 1–25, 2022.
- [45] S. Rahman, M. Aliakbarpour, H. K. Kong, E. Blais, K. Karahalios, A. Parameswaran, and R. Rubinfield, "I've seen" enough" incrementally improving visualizations to support rapid decision making," *Proceedings of the VLDB Endowment*, vol. 10, no. 11, pp. 1262–1273, 2017.
- [46] A. Filipowicz, S. Carter, N. Bravo, R. Iliev, S. Hakimi, D. A. Shamma, K. Lyons, C. Hogan, and C. Wu, "Visual elements and cognitive biases influence interpretations of trends in scatter plots," *arXiv preprint*, 2023.
- [47] M. Johnson and K. Ottenbacher, "Trend line influence on visual analysis of single-subject data in rehabilitation research," *International Disability Studies*, vol. 13, no. 2, pp. 55–59, 1991.
- [48] W. S. Cleveland, "Robust locally weighted regression and smoothing scatterplots," *Journal of the American Statistical Association*, vol. 74, no. 368, pp. 829–836, 1979.
- [49] I. Popivanov and R. J. Miller, "Similarity search over time-series data using wavelets," in *International Conference on Data Engineering*, 2002, pp. 212–221.
- [50] H. Musbah, H. H. Aly, and T. A. Little, "A novel approach for seasonality and trend detection using fast fourier transform in box-jenkins algorithm," in *IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, 2020, pp. 1–5.
- [51] J. T. Duda and T. Pelech-Pilichowski, "Detection of periodic components from seasonal time series with moving trend method and low pass filtering," in *Advanced Solutions in Diagnostics and Fault Tolerant Control*, 2018, pp. 192–202.
- [52] J. Denholm-Price and J. Rees, "A practical example of low-frequency trend removal," *Boundary-layer Meteorology*, vol. 86, pp. 181–187, 1998.
- [53] M. G. Elfeky, W. G. Aref, and A. K. Elmagarmid, "Periodicity detection in time series databases," *IEEE Trans. on Knowledge and Data Engineering*, vol. 17, no. 7, pp. 875–887, 2005.
- [54] S. Lazcano and C. D. Johnson, "Improved data interpretation through identification of time series periodicity changes," Pacific Northwest National Laboratory, Richland, WA, USA, Tech. Rep., 2022.
- [55] Y. Shehu and R. Harper, "Efficient periodicity analysis for real-time anomaly detection," in *IEEE/IFIP Network Operations and Management Symposium*, 2023, pp. 1–6.
- [56] Q. Wen, K. He, L. Sun, Y. Zhang, M. Ke, and H. Xu, "Robustperiod: Robust time-frequency mining for multiple periodicity detection," in *International Conference on Management of Data*, 2021, pp. 2328–2337.
- [57] H. Musbah, M. El-Hawary, and H. Aly, "Identifying seasonality in time series by applying fast fourier transform," in *IEEE Electrical Power and Energy Conference (EPEC)*, 2019, pp. 1–4.
- [58] Q. Wen, L. Yang, and L. Sun, "Robust dominant periodicity detection for time series with missing data," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [59] M. Y. Y. Mohammed and M. Celik, "Developing fast techniques for periodicity analysis of time series," in *International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)*, 2019, pp. 1–5.
- [60] D. Stratimirović, D. Sarvan, V. Miljković, and S. Blesić, "Analysis of cyclical behavior in time series of stock market returns," *Communications in Nonlinear Science and Numerical Simulation*, vol. 54, pp. 21–33, 2018.

- [61] X. Otazu, M. Ribó, M. Peracaula, J. Paredes, and J. Nunez, "Detection of superimposed periodic signals using wavelets," *Monthly Notices of the Royal Astronomical Society*, vol. 333, no. 2, pp. 365–372, 2002.
- [62] M. G. Elfeky, W. G. Aref, and A. K. Elmagarmid, "Warp: time warping for periodicity detection," in *IEEE International Conference on Data Mining (ICDM)*, 2005, pp. 8–pp.
- [63] I. Boulnemour and B. Boucheham, "Qp-dtw: upgrading dynamic time warping to handle quasi periodic time series alignment," *Journal of Information Processing Systems*, vol. 14, no. 4, pp. 851–876, 2018.
- [64] X. Song, Q. Wen, Y. Li, and L. Sun, "Robust time series dissimilarity measure for outlier detection and periodicity detection," in *ACM International Conference on Information & Knowledge Management*, 2022, pp. 4510–4514.
- [65] P. Rosen, A. Suh, C. Salgado, and M. Hajji, "TopoLines: Topological Smoothing for Line Charts," in *EuroVis (Short Papers)*, 2020.
- [66] R. A. Rensink and G. Baldrige, "The perception of correlation in scatterplots," *Computer Graphics Forum*, vol. 29, no. 3, pp. 1203–1210, 2010.
- [67] G. Palshikar *et al.*, "Simple algorithms for peak detection in time-series," in *International Conference on Advanced Data Analysis, Business Analytics and Intelligence*, vol. 122, 2009.
- [68] Y. Zhou, J. Ma, F. Li, B. Chen, T. Xian, and X. Wei, "An improved algorithm for peak detection based on weighted continuous wavelet transform," *IEEE Access*, vol. 10, pp. 118 779–118 788, 2022.
- [69] P. Du, W. A. Kibbe, and S. M. Lin, "Improved peak detection in mass spectrum by incorporating continuous wavelet transform-based pattern matching," *Bioinformatics*, vol. 22, no. 17, pp. 2059–2065, 2006.
- [70] J. Price and T. Goble, "Signals and noise," in *Telecommunications Engineer's Reference Book*. Elsevier, 1993, pp. 10–1.
- [71] S. Butterworth *et al.*, "On the theory of filter amplifiers," *Wireless Engineer*, vol. 7, no. 6, 1930.
- [72] J. D. Rhodes and S. Alseyab, "The generalized chebyshev low-pass prototype filter," *International Journal of Circuit Theory and Applications*, vol. 8, no. 2, pp. 113–125, 1980.
- [73] B. Raj, L. Turicchia, B. Schmidt-Nielsen, and R. Sarpeshkar, "An ftf-based companding front end for noise-robust automatic speech recognition," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2007, pp. 1–13, 2007.
- [74] A. Karar, R. Tanaka, and J. C. Vanel, "Apd's excess noise measurements using spectral analysis (fft)," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 387, no. 1-2, pp. 205–210, 1997.
- [75] V. Parekh, M. Bilalpur, C. Jawahar, S. Kumar, S. Winkler, and R. Subramanian, "Investigating the generalizability of eeg-based cognitive load estimation across visualizations," in *International Conference on Multimodal Interaction: Adjunct*, 2018, pp. 1–5.
- [76] G. Ryan, A. Mosca, R. Chang, and E. Wu, "At a glance: Pixel approximate entropy as a measure of line chart complexity," *IEEE Trans. on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 872–881, 2019.
- [77] M. Nagadia, "Apple stock price from 1980-2021," 2022. [Online]. Available: <https://www.kaggle.com/datasets/meetnagadia/apple-stock-price-from-19802021>
- [78] "Alma Science Archive," <http://almascience.nrao.edu/aq/>, 2020.
- [79] Chicago Police Department, "Crimes - 2001 to present: City of Chicago: Data Portal," <https://data.cityofchicago.org/Public-Safety/Crimes-2001-to-present/ijzp-q8t2>, 2020.
- [80] National Centers for Environmental Information, "Climate Data Online: Web Services Documentation," <https://www.ncdc.noaa.gov/cdo-web/webservices/v2>, 2020.
- [81] D. Dave, "Dogecoin historical data," 2021. [Online]. Available: <https://www.kaggle.com/ds/1179216>
- [82] A. Delorme, "EEG / ERP Public Database," https://scn.ucsd.edu/~arno/fam2data/publicly_available_EEG_data.html, 2020.
- [83] M. Bostack, "Observable: Zoomable Area Chart," <https://observablehq.com/@d3/zoomable-area-chart>, 2020.
- [84] "New Zealand Tourist Arrivals," <https://tradingeconomics.com/new-zealand/tourist-arrivals>, 2020.
- [85] "U.S. Bureau of Labor Statistics," <https://www.bls.gov/>, 2020.
- [86] E. Kim, J. Kim, H. Lee, and S. Kim, "Adaptive data augmentation to achieve noise robustness and overcome data deficiency for deep learning," *Applied Sciences*, vol. 11, no. 12, p. 5586, 2021.
- [87] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-lite: A grammar of interactive graphics," *IEEE Trans. on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 341–350, 2016.
- [88] M. L. McHugh, "Interrater reliability: the kappa statistic," *Biochemia medica*, vol. 22, no. 3, pp. 276–282, 2012.
- [89] P. Eichmann and E. Zraggen, "Evaluating subjective accuracy in time series pattern-matching using human-annotated rankings," in *International Conference on Intelligent User Interfaces (IUI)*, 2015, pp. 28–37.
- [90] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh, "Querying and mining of time series data: experimental comparison of representations and distance measures," *Proceedings of the VLDB Endowment*, vol. 1, no. 2, pp. 1542–1552, 2008.
- [91] H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction*. Providence, RI, USA: American Mathematical Society, 2010.
- [92] A. Lex, N. Gehlenborg, H. Strobel, R. Vuillemot, and H. Pfister, "Upset: visualization of intersecting sets," *IEEE Trans. on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 1983–1992, 2014.
- [93] H. Elhamedi, A. Gaba, Y.-S. Kim, and C. Xiong, "How do we measure trust in visual data communication?" in *Evaluation and Beyond-Methodological Approaches for Visualization (BELIV)*, 2022, pp. 85–92.
- [94] N. Kong and M. Agrawala, "Perceptual interpretation of ink annotations on line charts," in *ACM Symposium on User Interface Software and Technology (UIST)*, 2009, pp. 233–236.
- [95] M. D. Rahman, G. J. Quadri, B. Doppalapudi, D. A. Szafr, and P. Rosen, "A qualitative analysis of common practices in annotations: A taxonomy and design space," *IEEE Trans. on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 360–370, 2025.



Rifat Ara Proma is a Ph.D. Candidate in the School of Computing at the University of Utah. Her dissertation focuses on developing strategies to mitigate visual clutter in line charts. She holds a Master's in Computer Science from the University of South Florida. Her research interests include information visualization, human-centric visualizations and human-computer interaction.



Michael Correll is a Research Associate Professor at Northeastern University's Roux Institute. He received the Ph.D. degree from the University of Wisconsin-Madison. His research interests include data ethics, communicating statistics to mass audiences, and investigating biased or misleading data visualizations.



Ghulam Jilani Quadri is an Assistant Professor in the School of Computer Science at the University of Oklahoma. Prior to that, Dr. Quadri was a CIFellow postdoc at the University of North Carolina at Chapel Hill. He received his Ph.D. from the University of South Florida. Quadri received the 2021 Computing Innovation Fellow award. His research interests include creating human-centered frameworks to optimize visualization design and improve decision-making quality.



Paul Rosen is an Associate Professor at the University of Utah. He received his Ph.D. from Purdue University. His research interests include applying geometry- and topology-based approaches to problems in information visualization. Along with his collaborators, he has received best paper awards or honorable mentions at IEEE VIS, IEEE PacificVis, CG&A, IVAPP, and SIBGRAPI. Dr. Rosen received a National Science Foundation CAREER Award in 2019.