

From Instruction to Insight: Exploring the Functional and Semantic Roles of Text in Interactive Dashboards

Nicole Sultanum  and Vidya Setlur 




Fig. 1: Examples of interactive dashboards that demonstrate how descriptive text elements can guide users through interactive data exploration and communicate key insights. Left (by *Kevin Wee, #105*): An analytic-style dashboard that provides descriptive and instructional text for the geographical distribution and profiles of U.S. higher education institutions with data analytics programs. Users can interact with the map to discover information about each institute, with supplementary bar charts and histograms providing details on the degrees offered, campus settings, student enrollment sizes, average net prices, and graduation rates. Right (by *Infocepts, #173*): An infographic-style dashboard showing changes in appointments, total revenue, and sales, with qualitative assessments to narrate the broader effects on shopping behavior and wedding planning. Visual elements, along with quantitative metrics and descriptive narratives, are employed to draw attention to data trends and provide an overview of the industry's adaptation to virtual platforms.

Abstract— There is increased interest in understanding the interplay between text and visuals in the field of data visualization. However, this attention has predominantly been on the use of text in standalone visualizations (such as text annotation overlays) or augmenting text stories supported by a series of independent views. In this paper, we shift from the traditional focus on single-chart annotations to characterize the nuanced but crucial communication role of text in the complex environment of interactive dashboards. Through a survey and analysis of 190 dashboards in the wild, plus 13 expert interview sessions with experienced dashboard authors, we highlight the distinctive nature of text as an integral component of the dashboard experience, while delving into the categories, semantic levels, and functional roles of text, and exploring how these text elements are coalesced by dashboard authors to guide and inform dashboard users. Our contributions are threefold. First, we distill qualitative and quantitative findings from our studies to characterize *current practices* of text use in dashboards, including a categorization of text-based components and design patterns. Second, we leverage current practices and existing literature to propose, discuss, and validate *recommended practices* for text in dashboards, embodied as a set of 12 heuristics that underscore the semantic and functional role of text in offering navigational cues, contextualizing data insights, supporting reading order, among other concerns. Third, we reflect on our findings to identify gaps and propose *opportunities* for data visualization researchers to push the boundaries on text usage for dashboards, from authoring support and interactivity to text generation and content personalization. Our research underscores the significance of elevating text as a first-class citizen in data visualization, and the need to support the inclusion of textual components and their interactive affordances in dashboard design.

Index Terms—Text, dashboards, semantic levels, metadata, interactivity, instruction, description, takeaways, conversational heuristics.

1 INTRODUCTION

Dashboards have become a prevalent analytical artifact for visualizing and interacting with data across a variety of domains. From their original conception as glanceable visual monitoring systems for time-critical decision making [13,48], dashboards have evolved into versatile *data communication tools* that support both analytical and storytelling purposes [2,45]. This shift in utility has led to an extensive body of research aimed at cataloging and categorizing dashboards, focusing on their communication goals [45], dashboard intent [42], their interaction mechanisms [62], and conversational design heuristics [48],

- Nicole Sultanum and Vidya Setlur are with Tableau Research. E-mail: {nsultanum,vsetlur}@tableau.com

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

all aspects that underscore the growing importance of dashboards to support effective communication of data to their intended audience.

Among the diverse components that make up a dashboard, text elements such as titles, labels, captions, and data descriptions play a pivotal role in this communication [13,25,47,64]. An assessment of public dashboards in the wild found a majority (55%) contained blocks of text [52], stressing the prevalence of textual elements in dashboard design. There has been increased interest in studying relevant uses of text in the context of individual charts [20,24,58], but text use in dashboards has received relatively little attention so far. And arguably, it is a timely concern. The current repertoire of dashboard generation tools focuses on automating chart and layout creation, placing less emphasis on text generation and placement [16,30]. With the rise of large language models, this gap can be more easily filled, but without a solid understanding of practices and expectations for text, we are ill-equipped to design appropriate guardrails that can steer dashboard content generation in useful, meaningful directions.

Acknowledging the growing relevance of text in data visualization

and the current landscape of generative AI for text, our work aims to understand and formalize the use of text in dashboards, and to provide guidance for dashboard authors and dashboard tool designers alike. Our approach was informed by (a) a survey and qualitative analysis of 190 publicly available dashboards to catalog features of text and characterize their impact on dashboard design and communication, and (b) design review sessions with 13 dashboard creators to reflect and learn from on their personal practices for text use (§3.2). Our contributions are threefold. First, we (i) outline *current practices* for text in dashboards based on dashboard analysis and expert feedback (§4), which helped inform a typology of text elements in dashboards and a compilation of text use patterns. Second, we discuss (ii) *recommended practices* for text use in dashboards, based on an adaptation of Setlur et al.’s dashboard heuristics [48] and validated by experts (§5). Third, we outline (iii) *future opportunities* for the usage of text in dashboards, reflecting on emerging challenges for text use and potential solutions based on the state of the art in data visualization research (§6). Our work sheds light onto the functional, structural, and semantic aspects of text and reflects on its vital role in dashboards for analytical and navigation support. Our investigation shows that text is much more than a mere accompaniment to visuals but rather a fundamental enabler for narrative, explanation, and insight discovery.

2 RELATED WORK

Our work builds on several lines of work: dashboard characteristics & guidelines, text & chart integration, and dashboard authoring & search.

2.1 Dashboard Characteristics and Guidelines

The role of dashboards has evolved to support the increasing complexity of user needs and the recognition that dashboards serve not just as a means for data consumption but as platforms for data conversation. Tory et al.’s [62] concept of *data conversations* highlights the diverse communication goals of dashboard users, emphasizing tasks such as summarizing, monitoring, and predicting insights.

Dhanoa et al. [10] and Sarikaya et al. [45] identify key elements that contribute to an effective dashboard, including the importance of guiding users through the data exploration process with clear narratives and structured interaction pathways. Bach et al. [2] extend this understanding by identifying design patterns and dashboard genres that cater to varied user intents and interaction styles, where text plays a key role in data narration and guiding user exploration. Setlur et al. [48] present a set of 39 dashboard design heuristics grounded in the Gricean maxims of cooperative conversation, a framework that views dashboard interaction as a form of *analytical conversation*. The heuristics consider the role of text in facilitating analytical conversation, ranging from legibility and descriptiveness, effective use of iconography, and the provision of navigational cues and breadcrumbs. Aspects of communication that dashboard text contributes to are acknowledged and studied in these works, but not how text specifically influences this communication.

Our work is also informed by a long history of design heuristics for interfaces and visualizations that include Nielsen’s ten usability guidelines [38, 39] and Shneiderman et al.’s eight golden rules for interface design and evaluation [51], which provide directives for user-centered design. The examination of usability heuristics for evaluating visualizations by Tory & Möller [63] emphasizes the benefit of incorporating both usability and visualization expertise in creating effective analytical interfaces. Subsequent work has aimed to develop and evaluate visualization-specific heuristics [8, 11, 14, 61, 65, 75], expanding upon the principles of user-centered design to cater specifically to the unique challenges and opportunities presented by visual data representation. Resources by Few [13], Yigitbasioğlu [71], and Wexler [68] introduce a user-centered design perspective to dashboard design, complemented by frameworks that focus on user goals and intents [27, 31]. Together, these works guide visualization design towards more meaningful user experiences and support users’ cognitive processes [67].

We build upon this body of research by examining the role of text in dashboards along with its categories, semantic levels, and functional roles in guiding and informing users. Through an extensive dashboard

coding exercise, we uncover the interplay between text and dashboard visualizations to support analytical and narrative communication goals.

2.2 Text and Chart Integration

Recent research has explored the integration of text and charts in visualizations through various lenses, offering insights into how text can influence the perception and understanding of visualized data.

Kim et al. [20] conducted a crowdsourced study to understand how readers synthesize information from both charts and captions, finding that the emphasis on high-prominence features within both elements led readers to identify those features as key takeaways. Their research underscores the importance of coherence between visual and textual elements and how external context can enhance the reader’s comprehension of the chart’s message. Building on these insights, Lundgard and Satyanarayan [34] proposed a four-level semantic categorization of text content designed to enhance visualization accessibility. Their framework distinguishes between perceiver-independent descriptions, such as objective chart specifications, and perceiver-dependent insights, including interpreted observations and contextual knowledge. Recent work broadens the scope of accessible and inclusive data representation by championing the parity of visualization, textual description, and sonification in multimodal data analysis [74].

Further exploring the role of text in visualizations, Stokes [58] observed that study participants favored heavily annotated charts over less annotated charts or text alone. This preference highlights the added value of text in aiding data interpretation, with emphasis on how different types of semantic content — statistical, relational, elemental, or encoded — impact the takeaways drawn by readers. This finding aligns with Stokes and Hearst’s advocacy for treating text as co-equal to visualization, urging researchers to focus on the readability and integration between these two modalities [57]. Their call to action reflects a growing body of work, including Otley et al. [41], exploring the critical role of text in visual analysis and its importance for conveying key messages to the reader. Brath provides for a more integrated and text-centric perspective on visualization, suggesting new avenues for enhancing data communication through textual integration [5].

Eye-tracking studies by Borkin et al. [4] have shown that participants were more likely to fixate on and recall textual content surrounding visualizations, such as titles and labels. Similarly, Kong et al. [22, 23] investigated the influence of titles on the perceived message of visualizations, discovering that slanted framings (e.g., emphasizing only part of the chart’s message) significantly impacted recall and interpretation. Zhi et al. [73] found that explicit text-chart integration in data-driven stories may increase recall and engagement. Several tools build on these synergies, e.g., FacetNotes [3], a novel concept for integrating text annotations directly with data points on interactive dashboards; Kori [29], a tool to support identification of text-chart links; and VizFlow [60], a tool that leverages text-chart links to generate dynamic data-driven articles.

However, the preference for integrating text and visual elements is not always universal. Hearst & Tory [15] investigated visualization preferences in conversational interfaces with chatbots and found a notable division in participant preferences. Nearly half of the participants expressed a preference against viewing charts, while those who did prefer charts also favored the inclusion of additional contextual data within the visual representation. Based on these foundational studies, we specifically focus on the dynamic interplay between text and chart features in the context of interactive dashboards. Unlike previous research that primarily considers static visualizations or singular aspects of text integration, our work investigates how authors specifically incorporate text elements in dashboards, exploring both the role of text and their various forms.

2.3 Dashboard Authoring and Search

Understanding the analytical intent for authoring dashboards is crucial for effectively integrating text and charts for enhancing the overall communicative effectiveness of dashboards [6, 46]. Recent work by Pandey et al. [42] proposed a set of dashboard intents and objectives accompanied by specific objectives like summarizing measures, comparing

categories across dimensions, and displaying univariate summaries, to support dashboard composition through auto-generated variations of intent and objectives. Srinivasan and Setlur [55] extend the capability for dashboard composition through a natural language interface, enabling authors to generate and arrange dashboard views based on user-provided linguistic utterances. Machine learning techniques have also been explored for dashboard authoring and search. Ma et al. [35] and Wu et al. [69] propose deep learning-based methods to assist in dashboard generation and chart view suggestions, respectively, demonstrating the potential of ML to expedite dashboard prototyping and enhance the generation process. Similarly, Deng et al. [9] introduce DashBot, leveraging reinforcement learning to automate dashboard generation, focusing on the creation of valid dashboards that facilitate data insight discovery and support direct manipulation of recommendations. However, this prior work tends to focus less on the systematic integration of text beyond basic labels and titles in dashboard generation, presenting an opportunity to explore the role of text to further enhance the dashboard’s communicative effectiveness.

Oppermann et al. [40] explore leveraging text-based content within dashboards to aid in dashboard search and recommendation for specific analytical questions. Their work highlights the challenge of sparse text presence in dashboards and the necessity of including non-visible text for context, thus offering a pathway toward more meaningful recommendations based on textual content. EmphasisChecker [19] explores the integration of text and charts more directly in the authoring experience by identifying mismatches between the visual emphasis of charts and the textual emphasis of captions. Our work extends beyond prior research that primarily focused on dashboard intents, objectives, and composition; rather, we explore the specific interplay between text and visualizations, wherein text is not merely an accessory to visual elements but a primary medium for communication within dashboards, conveying key insights to the intended audience.

3 METHODOLOGY

Our efforts to characterize text use in dashboards are threefold: to identify relevant features of current text use, to provide recommendations for text use, and to identify opportunities for improvement. To support these goals, we performed two elicitation studies. First, we surveyed and analyzed a large collection of publicly available dashboards, which informed a typology of dashboard text elements and provided samples to illustrate best practices. Second, we conducted interviews with 13 dashboard designers, which helped contextualize and validate findings from our dashboard survey and inform challenges of text use in dashboards, which in turn helped inform research opportunities for text use in dashboards. We detail our study methodology in the following subsections, and our cross-cut findings are presented throughout our investigation goals to describe *current practices* (§4), *recommended practices* (§5), and *opportunities* (§6).

3.1 Dashboard Corpus

We compiled a large and diverse collection of 190 publicly available dashboards (Tableau (116), Power BI (50), and other web visualization tools (24)) sourced four ways:

1. **Dashboard corpora from prior work** (52). We included dashboards listed in prior surveys [2, 45, 48] that were both in English and publicly available, yielding 52 dashboards (from a total of 173). While these dashboards cover a relatively diverse range of platforms and genres, they have been opportunistically sampled, which could suggest a skewed selection of dashboards from limited sources and topics.

2. **Curated sample from Tableau Public** (74). To broaden diversity, we gathered dashboards from a large dataset of Tableau Public [1] dashboards featured in Srinivasan et al. [52] (with assistance from the authors). From an initial random sample of dashboards, we filtered for dashboards in English featuring five or more dashboard zones (i.e., charts, filters, legends, text, multimedia blocks), with at least one text zone and one chart zone. This led to a diverse set of 74 dashboards, spanning a broad range of genres, topics, communication goals, analytical complexity, and levels of refinement.

3. **Popular Power BI dashboards** (43). We also included a third set of 43 Power BI dashboards, particularly those highly regarded in the “Top Kudos” category (i.e., high community value) of the Power BI Data Stories Galleries [44] to ensure platform diversity and to analyze best practices in dashboard design across different software ecosystems.

4. **Dashboards from expert interview participants** (21). We included a fourth set of 21 dashboards from interviews (§3.3) that were authored by expert interview participants and handpicked by them as meaningful examples of text use in their dashboard practice. While most (19) are Tableau dashboards, they present a balanced coverage of genres and topics, from business dashboards to infographics for the general public.

3.2 Dashboard Analysis

The first author conducted an initial round of inductive analysis on a subset of dashboards to flesh out a preliminary typology of text-based components. Both authors were then involved in a follow-up deductive stage to refine these codes and incorporate other text-relevant features from past literature that encode dashboard semantics, including *dashboard genres* [2], *dashboard goals* [45] and *semantic levels of data-driven content* [34]. This resulted in a code book for independent coding covering a typology of text components and relevant dashboard-level features. A detailed breakdown of our derived typology is presented in §4.

We also compiled a set of heuristics for text usage in dashboards. Based on the work by Setlur et al. [48] that proposes 39 heuristics for dashboard design to foster analytical conversation, we reframed their heuristics from a perspective of text use, leading to a set of 12 heuristics for text use. These heuristics were appended to the code book alongside a detailed rubric listing 4 levels of application for each heuristic: *strong application*, *weak application*, *weak violation*, *strong violation*. A detailed discussion on the derived heuristics is presented in §5.

Using the code book, both authors conducted an independent assessment of dashboards in our corpus to extract meaningful text features and assess the application of the heuristics. Authors met regularly during the coding process to resolve disagreements and refine code book descriptions and rubric. In parallel, recurrent and interesting applications of text use were also cataloged. Unanimous consensus was reached on all features for the 190 dashboards in our collection.

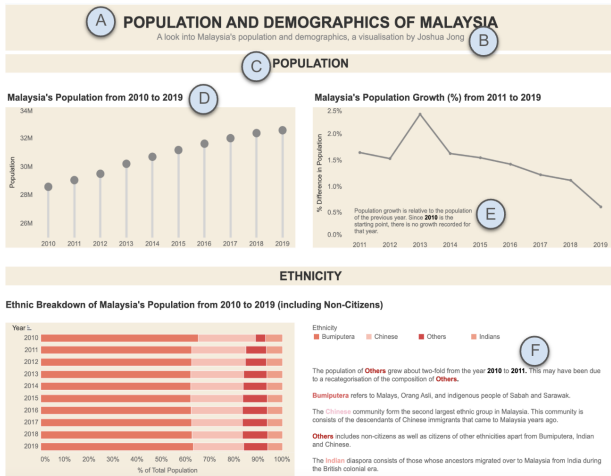
Outcomes of this coding effort showcased the diversity of our corpus, featuring a broad range of dashboard *genres* [2] (*analytic* (59%), *infographic* (25%), *static* (10%), *repository* (8%) and *magazine* (4%)) and dashboard *goals* [45] (*decision making* (40%), *awareness* (22%), *motivation & learning* (38%)). It also contributed a more systematic look over patterns of text use (§4), meaningful examples to showcase patterns (§5), and inspiration for future opportunities for text use in dashboards (§6). Our coded dashboard corpus and code book can be found in Supplemental Materials. We reference specific dashboard instances by their corpus (*#ID*) (with embedded access links).

3.3 Expert Interviews

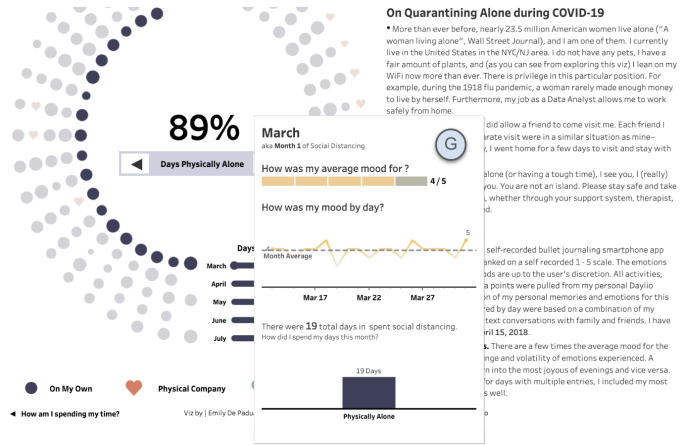
To complement our dashboard analysis efforts, we conducted interview sessions with experienced dashboard creators. Our goal was to validate observations of text use from the dashboard analysis, elicit rationale for design patterns, document common practices for text use, and validate our proposed set of heuristics.

We recruited 13 participants (*P1-P13*) from various channels, including social media, specialized community forums, targeted recruitment, and word of mouth. Most participants create dashboards professionally; listed occupations include business intelligence (BI) managers, data analysts, BI consultants, and students in data-relevant fields. All participants had at least one year of experience producing dashboards, with 10 reporting at least five years of professional experience on the topic. Tableau (11/13) and Power BI (5/13) were the most commonly cited dashboard platforms, but several participants also used other tools (5/13), and many stated being conversant in more than one tool (6/13).

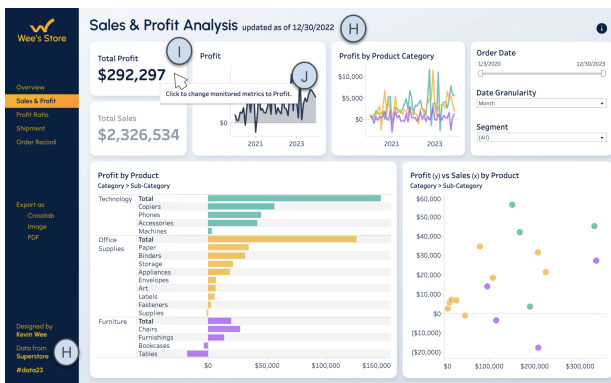
Sessions were one hour long and consisted of a semi-structured interview on text practices for dashboards followed by a feedback segment on the heuristics. Prior to the session, participants were asked to provide links to dashboards they authored that made for interesting



(a) A dashboard visualizing the population and demographic trends in Malaysia.



(b) A personal dashboard tracking the experience of social distancing during the COVID-19 pandemic.



(c) A business dashboard for sales and profit analysis.



(d) A business dashboard for sales monitoring.

Fig. 2: Text components in various dashboard snippets. Figure 2(a) (by *Joshua Jong, #87*): (A) The title provides an overarching theme; (B) The subheading offers additional context and credits the creator; (C) Section headers categorize the dashboard into thematic areas; (D) chart titles denote the focus of each visualization; (E) chart annotations explain specific data points or trends for clarity; and (F) Content blocks supply narrative insights, interpreting the ethnic composition changes over time and providing definitions for terms such as 'Bumiputera' and the 'Indian' diaspora. Figure 2(b) (by *Emily De Padua, #136*): Features a prominent percentage highlighting days spent alone, a narrative text block reflecting personal reflections on quarantine, and a (G) custom tooltip displayed on hover, showing an interactive timeline charting average daily mood. Figure 2(c) (by *Kevin Wee, #59*): (H) Metadata blocks, indicating author, data source, and when the dashboard was last updated; (I) text-data summary blocks featuring total profit and sales; and (J) interaction guidance on tooltip hover. Figure 2(d) (by *Fowmy Abdulmuttalib, #156*): (K) A strip of text-data summaries featuring icons; and (L) a data table featuring embedded graphics.

discussion on the use of text (and were later included in our fourth sample set, §3.1). At the session, we asked them to share design rationale on text use for their dashboards, including semantics, level of detail, text formatting choices, layout and reading order, user guidance, dynamic text, and accessibility. We also inquired about practices and challenges they generally face with text. We then presented our list of heuristics and asked them to reflect on how representative they were of their current practices and what aspects of text practices we should further consider. Sessions took place over video conference and were screen-recorded and auto-transcribed. Findings from transcripts were derived via thematic analysis and contributed to contextualize patterns of dashboard text use (§4), validate our proposed heuristics (§5), and inform opportunities for text use in dashboards (§6).

4 CHARACTERIZING CURRENT TEXT PRACTICES

Our examination of current practices for dashboard text is based on findings from our coding exercise (§3.2) and interviews with expert dashboard creators (§3.3). First, we provide a characterization of text components in dashboards from a *functional* lens by outlining different embodiments and roles of text use (§4.1) and from a *semantic* lens via an adaptation of Lundgard and Satyanarayan's four-level model of semantic content [34] (§4.2). We then discuss prominent text usage

patterns that emerged from the coding exercise and expert interviews, presented alongside implications for design (§4.3).

4.1 Text Components

We identified two groups of dashboard components where text plays an integral role: (a) *text blocks*, i.e., where text appears as a standalone element; and (b) *text-based components*, where text is intrinsically combined with visual and interactive features but is still prominent. Text decorators to interactive widgets (e.g., filter labels, drop-down labels) and primarily visual components (e.g., chart axis and ticks and color legend labels) are not considered in this assessment, as their content and functionality are dictated by design requirements of corresponding charts and legends, rather than offering independent or additional narrative or analytical insights. We also did not explicitly code for text visualization instances, as their use of text is highly bespoke.

The list of *text blocks* we identified include:

- **Title** (present in 92% of dashboards in our corpus). A prominent, discernible label (or short sentence) displayed with a font size larger than other text blocks in the dashboard (Fig. 2(A)).
- **Subheading** (36%). Text that goes below (or beside) the dashboard title and adds more context to the title (Fig. 2(B)).

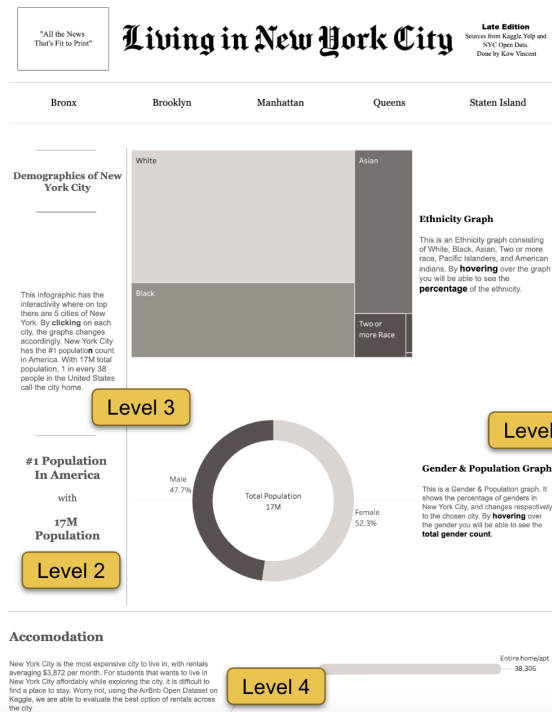


Fig. 3: A snippet of a dashboard displaying text components at varying levels of semantic detail (by *Kow Vincent, #89*). Level 1: basic chart descriptions and titles, indicating visible data dimensions like “Gender & Population Graph.” Level 2: statistical concepts shown as BANs providing numerical summaries such as “17M Total Population.” Level 3: perceptual insights on population segments elucidate more complex data features. Level 4: contextual and domain-specific insights offer additional information regarding accommodation affordability in New York City.

- **Section Header (29%)**. A textual element that labels a partition of a dashboard containing multiple components (e.g., a text block alongside one or more views) and provides a high-level overview of what the ensuing section contains (Fig. 2(C)).
- **Chart Title (84%)**. Text labeling an individual chart or view (Fig. 2(D)). Even though they play a supportive role to charts, they allow for more authoring freedom than text decorators and are thus included in our assessment as a standalone component.
- **Interaction Guidance (58%)**. Text that describes possible interactions with the dashboard, often written in imperative language, e.g., “click here to filter” (Fig. 2(J)).
- **Metadata (76%)**. A feature borrowed from Bach et al. [2], it designates text that describes data sources (53%), author (43%), last update (18%), and data-related disclaimers (26%) (Fig. 2(H)).
- **Content Block (67%)**. Sentence or multiline text that (a) supports the analytical and narrative goals of the dashboard (such as explaining what the data is, the visual encodings, salient points, and relevant context), and (b) is hierarchically perceived as “body” level text (Fig. 2(F)).

The list of *text-based components* that we consider are:

- ⇒ **Text-data summaries (58%)**. Also known as BANs (i.e., “big-a** numbers”), they feature key measures (i.e., a numerical value and a corresponding label) in a prominent way (Fig. 2(I) and (K)).
- ⇒ **Chart Annotation (19%)**. Text callouts overlaid on charts to bring attention to particular data elements, e.g., salient points, trends, data-specific context, and interaction guidance (Fig. 2(E)).
- ⇒ **Data Table (35%)**. A (still very prominent) text-based alternative to showcase data (Fig. 2(L)).
- ⇒ **Tooltip**. A dynamic text callout that appears when hovering data marks (Fig. 2(G) and (J)). Dashboards created in Tableau and

Power BI implement basic default tooltips that show the numerical values for the hovered data points, and so they are nearly ubiquitous (i.e., appearing in nearly 100% of our dashboards). However, a substantial subset of them (42%) are customized to display text content in a richer way (e.g., in the form of sentences), making tooltips a relevant content authoring resource.

4.2 Semantic Levels

Beyond a functional understanding of text components, our coding also sought to capture how text content shapes dashboard communication. To this end, we adapted Lundgard and Satyanarayan’s four-level model for semantic content [34], which was originally created in the context of accessibility to categorize analytical levels of detail of chart caption paragraphs. Given that the scope of our coding encompasses shorter, label-style text as opposed to fully formed sentences, we made a few adaptations when applying the four-level model to dashboard text components, described below and exemplified in Fig. 3:

- **Level 1 (LV1): Elementar and encoded properties (97%)**. Refers to the description of charts and visible chart data, e.g., adjacent content blocks, and chart titles.
- **Level 2 (LV2): Statistical concepts and relations (74%)**. Refers to mentions of statistical features of the data formulated at a “data facts” [53] level, e.g., descriptive statistics, outliers, point-wise comparisons. Text-data summaries generally fit in this category as they often depict sums, averages, and extrema in a similar granularity as LV2 statements.
- **Level 3 (LV3): Perceptual and cognitive phenomena (27%)**. Refers to more complex data statements, such as higher-level trends and pattern synthesis, while using more “natural-sounding” language than LV2 statements. Most often found in content blocks, and occasionally chart annotations and subheadings.
- **Level 4 (LV4): Contextual and domain-specific insights (50%)**. Text that provides additional context to the observed patterns and trends, e.g., historical, social, and political factors. Largely present in content blocks, and sometimes tooltips.

4.3 Outlining Dashboard Text Practices

Building on the functional and semantic categorizations above, plus dashboard-level features [2, 45] and expert interviews, we discuss various usage patterns for dashboard text alongside implications for design.

4.3.1 Dashboard Text Practices are Based on Experience

When asked about their text use more generally, most participants (8/13) stated that they are not explicitly formalized (P1, P4, P6, P8) nor taught (P6), but rather based on experience, exposure, and honed senses, e.g., learning from others (P6, P7, P13), learning on the job (P5, P12), and developing a feel for “what looks good” to them (P3, P7), which feeds into an implicit set of shared practices: “even if it hasn’t been formalized, there is kind of this collective understanding (...) because you’ve seen it repeated in other forms of conception” (P4). The more explicit guidance emerging from expert sessions was on structure and formatting guidelines borrowed from other established fields, such as publishing & typography (e.g., magazines) (P1, P4, P7), and graphic & web design (P5, P7, P13).

Implications: While text practices are not formalized, shared practices do exist and borrow inspiration from other communication and UX fields. Explicitly formalizing text practices is valuable and viable to some extent, but the most sustainable path forward is likely learning from examples (e.g., machine-learning approaches).

4.3.2 Text Components and Semantic Levels

Our coding exercise helped inform how text components tend to cover semantic levels. Chart titles were the most frequent contributors of LV1 content, and text-data summaries of LV2 content. Content blocks encompassed a lot of the higher level, fine-grained detail and context in dashboards, appearing in 96% of dashboards featuring LV3 content,




Fig. 4: Tooltip featuring LV2-LV3 content (by *Andy Cotgreave, #77*).

and 90% of dashboards featuring LV4 content (versus 67% of all dashboards); and so did chart annotations (present in 42% of dashboards with LV3 and 28% with LV4 content, versus 19% of all dashboards).

That said, text components also showcase a lot of diversity in terms of the range of semantic levels they may support. Tooltip instances in our corpus expanded far beyond their default role of displaying data values under hovered marks (LV1), to also correlate marks with other data dimensions (Fig. 2b (#136)), list associated data facts (#137), add in historical context (#55), and showcase disclaimers (#34). Chart annotation content was found to be equally diverse, covering the whole range of LV1 (e.g., overlaying associated data values, (#126)), LV2 (e.g., emphasizing salient points and trends, (#111) (#118)), LV3 (e.g., explaining complex trends, (#87) (#37)) and LV4 content (e.g., incorporating external context, (#79) (#58)). We also noted several instances of titles and subheadings being used to convey LV3 content, such as summarizing higher-level insights from the data and framing them as takeaways, e.g., (#178), Fig. 1 (right) (#173).

Related to this last observation, it is worth noting that the hierarchical structure of text blocks (titles, subheading, section headers, chart titles) also tend to mirror a hierarchy of semantic concepts, going from titles and subheadings that allude to higher level analytical goals down to chart titles which are closer to the data.

Implications: There is a strong relationship between structural levels of text and their semantic level coverage. This underscores that the visual design and formatting of text is more than an aesthetic concern. From a content generation perspective, this means not only looking at semantic hierarchy, but also how that hierarchy may be scaffolded by an explicit, visual text structure.

4.3.3 Text Use Across Dashboard Genres & Goals

The strongest signal we found to discriminate text use, based on the coding exercise and on expert feedback, is if the dashboard is more narrative-driven (e.g., *infographic* and *magazine* genres, often *motivation & learning* goal, for general audiences, casual use), or more exploratory (e.g., largely *analytic* genre, and *decision making* goals, for expert audiences, recurrent use).

By definition, narrative-driven dashboards make far more liberal use of explicit salient points and takeaways. This emphasis is clear when we look at the prevalence of perceiver-dependent content (LV3-LV4) and associated text components against *infographic* (more explanatory) and *analytic* (more exploratory) dashboards. Infographics feature more content blocks (91% vs. 53%), chart annotations (29% vs. 14%), and LV3-LV4 content (57% and 76% vs. 14% and 33%, respectively) than analytic dashboards. Differences between dashboard goals (*motivation & learning* vs. *decision making*) are even more pronounced, for content blocks (93% vs. 49%), chart annotations (32% vs. 12%) and LV3-LV4 semantic levels (58% and 75% vs. 8% and 29%).

Many experts (5/13) report on a similar divide on the design of enterprise dashboards as “*meant to be exploratory and analytical rather than narrative-driven*” (P9), and they acknowledge actively avoiding “*prescribed analytics*” (P4) in these cases: “*the last thing I want to do is to feed conclusions to the user*” (P12). That said, there is still a place for LV4 content in corporate dashboards, in the form of metric definitions and explanations (P4, P7, P9, P12), metadata (e.g., when last updated) (P7), data alerts (P4), and interaction guidance (P1, P5): “*You’re supposed to guide them on how to use the dashboard but not guiding them to the insights*” (P5).

Implications: Context of use is a key factor guiding the choice of text elements and semantic level coverage. Dashboard genres, dashboard goals, dashboard audience, and frequency of use are potentially useful signals to discriminate text use at scale.

4.3.4 Text Formatting for Analytical Support and Emphasis

In our coding exercise, we found that text formatting and placement is frequently used to support analytical reasoning. From interviews, we found that text formatting was the aspect of text use that experts were most intentional about, and the majority (7/13) reported following some form of text formatting guideline. The practices we compiled are:

- Leveraging font sizes to convey *hierarchy of text components* (e.g., titles vs. chart captions) and support reading order (P3, P4).
- Text color used as *color legends* (P6, P7, P9) (Figs. 2a and 2d).
- Boldface and text color for *emphasis* (e.g., data values, data labels, takeaways) (P7, P8); (Figs. 2b and 2a).
- Italics and bold face to *distinguish* interaction guidance from content blocks (P13) (Fig. 2d).
- *Avoid small fonts*, for legibility and accessibility reasons (5/13): “*screen readers can read best anything that’s 12 pt and up*” (P4)
- *Consistent use of fonts* (4/13): “*I don’t like mixing fonts*” (P4).

While seemingly obvious, these text practices require deliberate action from dashboard designers. There is also a general consensus that text formatting affordances in dashboard tools are minimal, and scriptable tasks like aligning text color to match visual marks require much manual intervention. Unsurprisingly, several experts reported the use of *style guides* in their corporate practice (P4, P7, P11) to help inform and align formatting practices across larger design teams.

Implications: Text formatting plays a significant role in dashboard communication and accessibility, but realizing its potential presently requires designer knowledge and experience. Given reasonable consensus around formatting practices, rule-based approaches to facilitate enforcement of best practices might be feasible as a first step. More importantly, explicit text-data linking would enable useful formatting affordances.

4.3.5 Managing Content Overload with Details on Demand

Regarding overall dashboard content, many experts expressed a desire to minimize “*big walls of text*” (P1, P11, P12), following a general understanding that users “*tend to glance over*” text content (P1, P4, P8, P13): “*my approach is to shorten everything that you can shorten*” (P11). These considerations are particularly relevant in the context of explanatory dashboards, where experts report being very deliberate in choosing what content to include and trying to “*keep it as relevant as possible*” (P7) by focusing on the key takeaways (P4, P6, P7): “*I’m constantly questioning myself. Why am I not including this? what is the ulterior motive of why you’re doing this?*” (P4).

As for minimizing visible text content, the most cited strategy is providing details on demand via *tooltips* (P1, P7, P9, P10, P13). While dynamic views and tabs were occasionally used for this same purpose, tooltips were prevalent and featured a broad range of implementations, including embedded visual reports (e.g., Fig. 2b (#136)). Another notable tooltip adaptation is having content formatted as a templated sentence (seen in 9% of dashboards, e.g., Fig. 4), aimed at simplifying the interpretation of data facts (P8, P13): “*if it’s a less technical audience, then I’ll try and turn it into a sentence*” (P13).

Implications: Choosing what content goes into dashboards is a nuanced task that leans on minimizing visible text. Content summarization, text templating tools, and more content nesting affordances would be helpful for authors. From a content generation perspective, ranking strategies should also be considered.

4.3.6 Guiding Users (or Not)

As discussed earlier, different dashboard genres entail different degrees of user guidance, from *analytical support* that highlights salient points




Fig. 5: Dashboard depicted in Fig. 2c (#59) with active instruction overlay.

and articulates takeaways to *usability support* via visualization interpretation instructions and interaction guidance. A lot of the text practices already covered pertain to user guidance, including the significant prevalence of interaction guidance components (58%), text formatting practices for guidance text, and simplifying analytical interpretation with sentence-style tooltips.

Despite the differences in analytical guidance needs across dashboard genres, the presence of interaction guidance in our corpus is fairly balanced across all non-*static* genres (between 53% and 65%). We found interaction guidance text to be most often displayed as collocated text blocks (either adjacent to or embedded in corresponding views) or as instructions in interactive widgets (e.g., filters). Other creative but less common implementations include: using icons in lieu of text (#82); *instruction overlays* that, when active, reveals text instructions next to relevant dashboard components, e.g., (#22), Fig. 5 (#59); and *staged presentation* (#136), where portions of the dashboard are shown or highlighted one at a time, e.g., (#172) (#181).

While useful, guidance may not always be needed. Relevant use cases for guidance, according to experts, include audiences with lower visualization literacy (P3, P13) and bespoke or unusual interactions (P6, P7, P13): “I like to assume most people know how general filters work, so I don’t tend to explain those unless there is some sort of nuance that needs to be explained” (P13). One participant argues that explicit interaction guidance may not even be always beneficial: “They’re really useful, but they’re only useful once, right? (...) An area I continually struggle with is, do you make that explicit and make the experience for first-time users better, or do you not make that explicit and make it better for continuous users?” (P9). They also commented on cases where discovering hidden interactions makes for a “*delightful*” experience: “That’s always fun to get the element of surprise. If it can be done well, you can spark joy” (P9).

Implications: Determining appropriate levels of analytical and interaction guidance is a key design consideration that requires understanding the analytical needs and visualization literacy levels of the target audience. While guidance is not always necessary or desirable, there is value in providing better support for authors to create more and better guidance and encourage thinking about end users and their information needs.

4.3.7 Dynamic Text for Analytical Support and Breadcrumbs

One of the aspects considered in our coding exercise was the presence of dynamic text, i.e., understanding if and how text components respond to interaction or embody animations. Dynamic text is present in 49% of all dashboards in our corpus and 62% of *analytic* dashboards, underscoring its role in more exploratory-driven dashboards. Common versions of dynamic text include text-data summaries that update with dashboard filters and interaction breadcrumbs for titles and chart titles that update to reflect newly selected data dimensions. These embodiments play an important awareness role in communicating what data slices are included or excluded, e.g., “I’m putting on [filter indicators] on the

dashboard so that they can see very clearly, ‘Oh, I have filters applied, I don’t really mean to do that’ ”(P6).

Less common but interesting uses of dynamic text are on dynamic takeaways, i.e., pre-computed text findings that change depending on the filter selections (e.g., (#53) (#149) (#190)); and a circular “news ticker” strips with data highlights (e.g., (#22) (#189)). On the latter, some experts commented on the value of animated features for user engagement (P3, P7, P9, P10): “Maybe you hover on top of a tiny text and it’s just goes bigger, little things that just keep you engaged” (P7); and some also mentioned practical uses of dynamic formatting and alerts in their corporate practice (P4, P6): “if it’s an overall data issue, it’s at the very top and it gets color-coded by brand” (P4). But despite the excitement and the meaningful use cases for dynamic text and dynamic formatting, adding such dynamic behaviors with mainstream dashboard tools is, again, reportedly challenging (P2, P5, P6, P13).

Implications: Dynamic text is a compelling resource to enhance comprehension and delight dashboard users but is difficult to implement with current tools. There are opportunities to provide more user-friendly ways to link dashboard text and data properties, and more flexible strategies for templated content generation.

4.3.8 Communication Elements Beyond Text

Finally, we mention a few text-adjacent topics that have an impact on dashboard communication and that emerged in the coding and discussions. First, we noted a prevalence of icons for communication (33%) either as space-efficient alternatives to text, as redundant visual elements, or as text decorators (e.g., Fig. 2d). When well designed, icons and symbols enhance dashboard comprehension, but again, practices are not formalized. There is room to facilitate more principled icon usage via standard icon libraries and dedicated design recommendations.

Second, text plays an important role in dashboard *accessibility*, as both an object (e.g., with formatting concerns) and enabler (e.g., via alt text) of accessibility. Experts agree that this is an important consideration, but also not a top-of-mind concern for most (7/13), and something that entails design trade-offs (P9) (e.g., limited font options).

Third, an accessibility-adjacent topic that has received relatively little attention is language translation support. Our English corpus features two distinct uses for multilingual dashboards, including one where all text is duplicated to feature both English and Japanese (#58); and one featuring protest chants in Polish (with English translations) for affect (#78). Advocating for the expansion of text visualization techniques and fostering multilingual accessibility is crucial for engaging a global audience and ensuring dashboard usability.

Implications: While text is a major player in dashboard communication, many other communication aspects can affect or be affected by text. Keeping a broader outlook on the dashboard communication landscape may help inform additional areas of research.

5 RECOMMENDATIONS FOR TEXT USE

Reflecting on the empirical nature of text practices, we argue it is useful to materialize such tacit knowledge in the form of explicit design guidance, not only for dashboard authors but also for dashboard tool designers. As described in §3.2, we turned to the work of Setlur et al. [48] and their heuristics for analytical conversation in dashboards. We reframed their 39 heuristics from a text-centered perspective to derive a set of 12 heuristics (Table 1) specifically on the use of text in dashboards to support analytical conversation. We present our derived heuristics (§5.1), discuss findings from applying these heuristics to our dashboard corpus (§5.2), and share expert impressions (§5.3).

5.1 Heuristics for Use of Text in Dashboards

Our construction process started by assessing each of the 39 heuristics in Setlur et al. [48] (H1-H39) and transforming relevant ones to be centered around text use. Some heuristics made explicit mentions of the appropriate use of text (e.g., H6 “the text in the dashboard is legible,

Analytical Support	HA1	Text in the dashboard supports specific analytical questions or tasks.
	HA2	The text in the dashboard should emphasize the most salient points of what the visuals in the dashboard convey.
	HA3	The text in the dashboard supports readers in deriving clear takeaways from the dashboard.
Semantics	HS1	The text provides sufficient contextual information to describe what the dashboard is about.
	HS2	Text content within the dashboard is at the appropriate level of detail to convey the intended message.
	HS3	The text in the dashboard plays a role in disclosing sources, disclaimers, and biases.
Presentation	HP1	The text in the dashboard is legible, easy to read, and useful. Different parts of the dashboard are well-described.
	HP2	Text supports a clear reading order within the dashboard, and it is logical.
	HP3	Text formatting and placement provide a clear and consistent visual style, mood, and guidance in understanding the analysis.
	HP4	Icons and symbols are used with text to help communicate patterns in data.
Interaction	HN1	Text makes it clear to the user where they need to start interacting with the dashboard.
	HN2	The text indicates a clear path and breadcrumbs for performed user actions within the dashboard.

Table 1: Our set of 12 heuristics for the use of text in dashboards, based on the work of Setlur et al. [48].

easy to read, and useful”, here mapping to HP1) and are used almost (if not) verbatim. Others covered communication aspects where text plays an implied or supporting role, which we then redrafted to explicitly mention text as an actor, e.g., H15 “*Charts within the dashboard are at the appropriate level of detail to convey the intended message*”, here mapping to HS2, “*Text content within the dashboard are at the appropriate level of detail to convey the intended message*”. Heuristics with similar meaning after these redraftings were merged (e.g., H9 thru H12, now HP3). While we iteratively fine-tuned descriptions and rubrics, the essence and scope of the heuristics remained.

Instead of the five conversation states used to organize the heuristics in Setlur et al. [48], we grouped ours into four guidance categories tailored to more directly address properties and roles of text use in dashboards: **Analytical support**, referring to how text aids users in understanding the data and analytical questions they can derive from the dashboard; **Semantics**, referring to the role of context and granularity of information to support analytical tasks; **Presentation**, referring to formatting and layout considerations to support navigation; and **Interaction**, highlighting the role of interaction guidance and awareness.

5.2 Applying the Heuristics

We wanted to understand how heuristics may be used in practice to assess dashboards in supporting analytical conversation and what features of text use tend to contribute to each heuristic. As mentioned in §3.2, both authors independently assessed dashboards in the corpus using the heuristics via a rubric that spans four assessment levels (*strong application*, *weak application*, *weak violation* and *strong violation*), and engaged in in-depth discussions to reach consensus on what aspects of dashboard text contributed to the assessment of each heuristic. We made our best efforts to account for differences in topic, genre, and dashboard goals in each case. This process allowed us to shift from a creator’s perspective to a user’s perspective and to reflect on how text components in dashboards may best serve their intended audience. In the following segments, we share our process findings from applying the heuristics to our dashboard corpus and how our functional (§4.1) and semantic (§4.2) categorizations may be leveraged to guide a heuristics-based dashboard assessment. For those seeking to perform similar heuristic assessments, we encourage the use of our assessment rubric in our code book (Supplemental Materials).

Analytical support. We found that having a combination of complementary text components that work well together is key for good analytical support. Header-level text blocks like *title*, *subheading*, *section headers*, and *chart titles* lend structure to connect data-level chart communication goals (HA1) to higher-level takeaways (HA3). Some components make direct and obvious contributions to heuristics, such as *chart annotations*, *text-data summaries*, and *content blocks* to salient points (HA2), and subheading to takeaways (HA3), but as discussed earlier in §4.3.3, prescriptive analytics are not always needed or desired, nor more content equate better communication. Several dashboards in our corpus showcased good application of these heuristics even with minimal use of text (#53) (#98) (#111).

Semantics. Assessing sufficient context (HS1) is a subjective measure that heavily depends on the use case and intended audience [17]. We

assessed this metric based on the extent we were able to interpret and analyze dashboard data given available information (e.g., in *content blocks*, *tooltips*, and *chart annotations*). Dashboards for expert users or general knowledge topics (like sports or the pandemic) might be sufficiently grounded without explicit context (LV4), but we found that having explicit context was generally helpful. On the other hand, the assessment of semantic levels of detail (HS2) is more objective and can be characterized in terms of the coverage of the LV1-LV4 content spectrum to support analytical reasoning. A strong application here means that there are clear reasoning paths from LV1 and LV2 findings to overarching conclusions (LV3 and LV4). All text components collectively play a role in supporting this aspect. Finally, sources, disclaimers, and biases (HS3) are present in *metadata blocks*; in our assessment, applications featured easily retrievable sources (45%), and strong applications featured notes on data collection and processing (19%).

Presentation. Presentation factors like formatting and placement play both functional and aesthetic roles in dashboard communication, and are addressed by HP2 and HP3. Good applications for HP2 refer to the extent that dashboard text supports users with navigation: it is about conveying a strong visual sense of content hierarchy via font size and header placement for content navigation, as well as strategic placement of text blocks close to relevant visual components. HP3, on the other hand, focuses on appropriate and consistent use of text formatting more holistically, including fonts, color, mood, and branding. HP1 pertained specifically to the quantity and quality of text at appropriate levels to fulfill dashboard goals; dashboards we assessed lower in this criteria featured little to no content beyond *chart titles*. Finally, HP4 is about ensuring that icons (when used) have a synergistic role with text, clearly communicating their intent and matching dashboard aesthetics.

Interaction. The final two heuristics pertain to the role of text in scaffolding interaction. The first is about aiding discovery (HN1), which we assessed as the presence of explicit text instructions (i.e., *interaction guidance blocks*) to point new users to key interactive features of the dashboard; strong violations happened when meaningful interactive features were missed due to lack of explicit guidance. The second refers to text communicating interaction outcomes (HN2), which was most frequently embodied as dynamic labels updating to match filters, e.g., a *chart title* updating alongside an updated chart. Strong violations here map to cases where no traces of user action are provided in text or visuals, e.g., filtered views misrepresented as whole dataset views.

5.3 Expert Feedback on Heuristics

Expert feedback on heuristics was largely positive. The empirical nature of text practices is again acknowledged (P4, P6, P12): e.g., “*I teach data visualization frequently (...) [and] we don’t really talk about text that much, or how to use text appropriately*” (P6); and that having these practices formalized as heuristics is useful (10/13): “*For those who are not familiar with the norms or may come from an experience where the norms are different, [this formalization] is a unifier*” (P4). Cited benefits and uses entailed: fostering intentional and “*thoughtful usage of text*” (P5, P6, P13); guidance for novice designers (P6, P8) and sanity checks for experienced designers (P5, P9); and providing a base for more bespoke guidance in applied scenarios (P7). Participants also

stated that our set of heuristics felt complete (8/13) and representative of their practices (5/13), and they also appreciated the structure with four guidance categories (P3, P4, P7).

While the heuristics were found to provide “a good template” (P7) for recommended text practices, many participants argued more specificity is needed to support practical hands-on support (P3, P4, P8, P9), e.g., “If it’s just guidance to point people in the right direction, great. But if you want people to have a standard and follow, I need a little more functional direction and just something to grasp and hold on to” (P4). A participant also remarked that the heuristics on their own will not ensure dashboard quality, e.g., “you tick the boxes altogether and then maybe you might get an ideal dashboard, but sometimes you might not” (P12). Topic suggestions for more specific guidance entailed best practices in text formatting (P4, P5, P7), guidance for annotations (P3), and an example catalog to showcase both good and bad applications of individual heuristics (P3, P5, P11, P13). More generally, some also underscored the importance of “context and perspective of the user” (P10, P12), and better communicate the extent to which each heuristic matters on different use cases (P10).

6 OPPORTUNITIES FOR TEXT USE IN DASHBOARDS

We identified several areas for future research that build on our dashboard findings, discussions with experts, and application of heuristics.

Dashboard linters. Based on the strongly empirical nature of dashboard text practices (§4.3.1) and informal dashboard text practices (§4.3.4), novice dashboard designers would greatly benefit from direct design guidance. Akin to the goals of visualization linters [7, 36], dashboard linters could provide real-time feedback on text use during authoring, ensuring adherence to best practices within the scope of our 12 heuristics (§5), e.g., typography, readability, and information hierarchy. These tools could also potentially offer suggestions for improving text clarity, as well as accessibility [54].

Example-based style transfer and recommendations. In view of the prevalent use of style guides in large design teams (§4.3.4), we also see value in more easily transferring text design features from one dashboard to another, similar to the notion of style transfer in image rendering [18]. By analyzing exemplary dashboards within various domains and genres, algorithms can also provide smart defaults and recommendations for effective text formatting styles (e.g., font sizes, colors, hierarchy) that can help lower the barriers for text design that can effectively support navigation and reading order (§4.3.1, §4.3.2).

Automatic explanations. Reflecting on user guidance needs for exploratory and explanatory dashboards alike (§4.3.3, §4.3.6), we see an important opportunity for the development of automated and semi-automated content generation tools for context-specific explanations on underlying data. In the context of corporate dashboards, for example, explanations on visible metrics are widely used but onerous to create and keep consistent with changes in the data processing pipeline. Being able to generate accessible, domain-specific explanations that can be vetted by dashboard authors would significantly ease the burden on the author to produce comprehensive explanations.

Provisions for text+charts links in dashboard design. Following research on text-chart links for visual storytelling [26, 28, 37, 60, 72, 73], facilitating the linking of dashboard text and charts would allow for exciting opportunities to expand the capabilities of dynamic text for analytical communication and storytelling (§4.3.4, §4.3.7). Automating the linking of text and charts in this context would involve developing algorithms that not only detect and establish connections between textual descriptions and corresponding data points for guidance [29, 37, 43], but also suggest contextually relevant emphasis [59], such as bolding key takeaways, coloring data references, or italicizing instructions (§4.3.4, §4.3.6). Future research could focus on creating intelligent storytelling tools that adapt text formatting based on the data narrative and integrate these capabilities into existing dashboard authoring platforms.

Guidance overlays and personalization. Incorporating instruction overlays into dashboards can guide users through complex datasets or analysis features, enhancing learnability (§4.3.6). Drawing from just-

in-time learning and in-situ instructions [56], research could explore the inclusion of overlays that adapt to user preferences and interaction patterns, offering better scaffolding for dashboard users. By analyzing user behavior, along with analytical and accessibility needs (§4.3.3, §4.3.8), the dashboard could then take user assistance a step further by tailoring personalized walkthrough versions for specific audiences, featuring data insights relevant to the user’s specific needs or roles.

Reactive documents meet dashboards. Reflecting on next steps for dynamic text (§4.3.7) and the benefits of prose text for analytical understanding (§4.3.5), we are inspired by past work on explorable documents [12, 66]. Applying these concepts to interactive dashboards presents a fertile ground for future research where interactive narratives and user-driven data exploration coalesce. For example, revealing textual explanations based on user interaction; breadcrumbs that provide dynamic summaries of the analytical journey; and contextual information surfacing on demand. Future work could also explore how text can play an active role in the storytelling experience through interactive text widgets that allow users to manipulate data directly from the narrative.

7 LIMITATIONS

While we made deliberate efforts to increase diversity and representation in our dashboard sample, our research primarily analyzes publicly available dashboards, a limitation given the vast universe of private corporate dashboards that remain beyond our purview. Business dashboards in public platforms (e.g., Tableau Public) are often portfolio versions depicting mock data, whereas real business dashboards are embedded in intricate data ecosystems, feature complex visualizations, and tend to be tailored for specific organizational contexts and decision-making processes, potentially incorporating proprietary design elements and interaction paradigms. This exclusion could mean that our analysis misses crucial insights into dashboard design and use in professional environments and might raise concerns regarding the generalizability of our findings. While our interview experts helped us uncover various text-related concerns relevant to their professional practice, we argue that collaborations with professional and media organizations could help identify novel applications of text in domain-specific use cases, thus broadening the applicability of this work.

We also acknowledge the inherent subjectivity that accompanies our evaluative framework of the heuristics. Arriving at a unanimous consensus on the application or violation of heuristics was a laborious and nuanced exercise, but also an informative one, and we stand by the value of these heuristics as tools for design guidance. That said, we agree that further assessment with wider usage would be valuable, e.g., evaluating the heuristics with novices as per original work [48].

Finally, while the heuristics proposed in our work provide a structured approach to evaluating and designing text for dashboards, they may not offer the specificity required for direct application in all authoring scenarios. This limitation arises partly from the diverse contexts in which dashboards operate, making it challenging to formulate universally applicable rules without sacrificing relevance to specific analytical use cases and target audiences. However, the broad nature of our guidelines also presents an opportunity for designers to interpret these guidelines in ways that best suit their unique objectives and user needs.

8 CONCLUSION

Our work highlights the fundamental role of text in the context of dashboard design and communication, wherein textual elements do much more than accompany visualizations; they enrich and clarify, providing essential context and direction. We hope this paper provides useful insights and future directions for not only dashboard authoring tools to better leverage the utility of text with visualizations, but also to guide effective analytical conversation and text generation in dashboard-adjacent text-data spaces, such as data stories [49, 50], data videos [70], computational notebooks [32], visualization thumbnails [21], and chart captioning [33, 54]. This synergy between text and visuals reaffirms the adage, “A picture is worth a thousand words, but a word can paint a thousand pictures”, underscoring the dual power of text to distill complex data into information understanding.

REFERENCES

- [1] Tableau Public. <https://public.tableau.com/>. Accessed: 2024. 3
- [2] B. Bach, E. Freeman, A. Abdul-Rahman, C. Turkay, S. Khan, Y. Fan, and M. Chen. Dashboard Design Patterns. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):342–352, 2022. doi: 10.1109/TVCG.2022.3209448 1, 2, 3, 5
- [3] S. K. Badam, S. Chandrasegaran, and N. Elmqvist. Integrating Annotations into Multidimensional Visual Dashboards. *Information Visualization*, 21(3):270–284, 2022. doi: 10.1177/14738716221079591 2
- [4] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, H. Pfister, and A. Oliva. Beyond Memorability: Visualization Recognition and Recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):519–528, 2015. doi: 10.1109/TVCG.2015.2467732 2
- [5] R. Brath. *Visualizing with Text*. CRC Press, 2020. doi: 10.1201/9780429290565 2
- [6] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013. doi: 10.1109/TVCG.2013.124 2
- [7] Q. Chen, F. Sun, X. Xu, Z. Chen, J. Wang, and N. Cao. VizLinter: A Linter and Fixer Framework for Data Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 28(01):206–216, Jan 2022. doi: 10.1109/TVCG.2021.3114804 9
- [8] A. Cuttone, M. K. Petersen, and J. E. Larsen. Four Data Visualization Heuristics to Facilitate Reflection in Personal Informatics. In *International Conference on Universal Access in Human-Computer Interaction*, pp. 541–552. Springer, 2014. doi: 10.1007/978-3-319-07509-9_51 2
- [9] D. Deng, A. Wu, H. Qu, and Y. Wu. Dashbot: Insight-driven Dashboard Generation Based on Deep Reinforcement Learning. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):690–700, 2022. doi: 10.1109/TVCG.2022.3209468 3
- [10] V. Dhanoa, C. Walchshofer, A. Hinterreiter, H. Stitz, E. Groeller, and M. Streit. A Process Model for Dashboard Onboarding. In *Computer Graphics Forum*, vol. 41, pp. 501–513. Wiley Online Library, 2022. doi: 10.1111/cgf.14558 2
- [11] D. Dowding and J. A. Merrill. The Development of Heuristics for Evaluation of Dashboard Visualizations. *Applied Clinical Informatics*, 9(03):511–518, 2018. doi: 10.1055/s-0038-1666842 2
- [12] P. Dragicevic, Y. Jansen, A. Sarma, M. Kay, and F. Chevalier. Increasing the Transparency of Research Papers with Explorable Multiverse Analyses. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 15 pages, 2019. doi: 10.1145/3290605.3300295 9
- [13] S. Few. *Information Dashboard Design: The Effective Visual Communication of Data*. O’Reilly Media, Inc., 2006. 1, 2
- [14] C. Forsell and J. Johansson. A Heuristic Set for Evaluation in Information Visualization. In *Proceedings of the International Conference on Advanced Visual Interfaces*, pp. 199–206, 2010. doi: 10.1145/1842993.1843029 2
- [15] M. Hearst and M. Tory. Would You Like a Chart with That? Incorporating Visualizations into Conversational Interfaces. In *2019 IEEE Visualization Conference (VIS)*, 5 pages. IEEE, 2019. doi: 10.1109/VISUAL.2019.8933766 2
- [16] M. A. Hearst. Show It or Tell It? Text, Visualization, and Their Combination. *Communications ACM*, 66(10):68–75, 8 pages, Sep 2023. doi: 10.1145/3593580 1
- [17] J. Hullman and N. Diakopoulos. Visualization Rhetoric: Framing Effects in Narrative Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2231–2240, 2011. doi: 10.1109/TVCG.2011.255 8
- [18] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song. Neural Style Transfer: A Review. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3365–3385, 2019. doi: 10.1109/TVCG.2019.2921336 9
- [19] D. H. Kim, S. Choi, J. Kim, V. Setlur, and M. Agrawala. EMPHASIS-CHECKER: A Tool for Guiding Chart and Caption Emphasis. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):120–130, Jan. 2023. doi: 10.1109/tvcg.2023.3327150 3
- [20] D. H. Kim, V. Setlur, and M. Agrawala. Towards Understanding How Readers Integrate Charts and Captions: A Case Study with Line Charts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, article no. 610, 11 pages. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3411764.3445443 1, 2
- [21] H. Kim, J. Kim, Y. Han, H. Hong, O.-S. Kwon, Y.-W. Park, N. Elmqvist, S. Ko, and B. C. Kwon. Towards Visualization Thumbnail Designs That Entice Reading Data-Driven Articles. *IEEE Transactions on Visualization and Computer Graphics (In press)*, 16 pages, 2023. doi: 10.1109/TVCG.2023.3278304 9
- [22] H.-K. Kong, Z. Liu, and K. Karahalios. Frames and Slants in Titles of Visualizations on Controversial Topics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 12 pages, 2018. doi: 10.1145/3173574.3174012 2
- [23] H.-K. Kong, Z. Liu, and K. Karahalios. Trust and Recall of Information Across Varying Degrees of Title-visualization Misalignment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 13 pages, 2019. doi: 10.1145/3290605.3300576 2
- [24] N. Kong, M. A. Hearst, and M. Agrawala. Extracting References between Text and Charts via Crowdsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’14, 10 pages, p. 31–40. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2556288.2557241 1
- [25] R. Kosara and J. D. Mackinlay. Storytelling: The Next Step for Visualization. *Computer*, 46(5):44–50, 2013. doi: 10.1109/MC.2013.36 1
- [26] B. C. Kwon, F. Stoffel, D. Jäckle, B. Lee, and D. Keim. VisJockey: Enriching Data Stories Through Orchestrated Interactive Visualization. In *Poster Compendium of the Computation+Journalism Symposium*, vol. 3, p. 3, 2014. 9
- [27] H. Lam, M. Tory, and T. Munzner. Bridging from Goals to Tasks with Design Study Analysis Reports. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):435–445, 2017. doi: 10.1109/TVCG.2017.2744319 2
- [28] S. Latif, S. Chen, and F. Beck. A Deeper Understanding of Visualization-Text Interplay in Geographic Data-driven Stories. *Computer Graphics Forum*, 40(3):311–322, 2021. doi: 10.1111/cgf.14309 9
- [29] S. Latif, Z. Zhou, Y. Kim, F. Beck, and N. W. Kim. Kori: Interactive Synthesis of Text and Charts in Data Documents. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):184–194, 2021. doi: 10.1109/TVCG.2021.3114802 2, 9
- [30] B. Lee, N. H. Riche, P. Isenberg, and S. Carpendale. More Than Telling a Story: Transforming Data into Visually Shared Stories. *IEEE Computer Graphics and Applications*, 35(5):84–90, 2015. doi: 10.1109/MCG.2015.99 1
- [31] E. Lee-Robbins and E. Adar. Affective Learning Objectives for Communicative Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1), 11 pages, 2022. doi: 10.1109/TVCG.2022.3209500 2
- [32] H. Li, L. Ying, H. Zhang, Y. Wu, H. Qu, and Y. Wang. Notable: On-the-fly Assistant for Data Storytelling in Computational Notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 16 pages, 2023. doi: 10.1145/3544548.3580965 9
- [33] C. Liu, L. Xie, Y. Han, D. Wei, and X. Yuan. AutoCaption: An Approach to Generate Natural Language Description from Visualization Automatically. In *2020 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 191–195. IEEE, 2020. doi: 10.1109/PacificVis48177.2020.1043 9
- [34] A. Lundgard and A. Satyanarayan. Accessible Visualization via Natural Language Descriptions: A Four-level Model of Semantic Content. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1073–1083, 2021. doi: 10.1109/TVCG.2021.3114770 2, 3, 4, 5
- [35] R. Ma, H. Mei, H. Guan, W. Huang, F. Zhang, C. Xin, W. Dai, X. Wen, and W. Chen. LADV: Deep Learning Assisted Authoring of Dashboard Visualizations from Images and Sketches. *IEEE Transactions on Visualization and Computer Graphics*, 27(9):3717–3732, 2020. doi: 10.1109/TVCG.2020.2980227 3
- [36] A. McNutt and G. Kindlmann. Linting for Visualization: Towards a Practical Automated Visualization Guidance System. In *VisGuides: 2nd Workshop on the Creation, Curation, Critique and Conditioning of Principles and Guidelines in Visualization*, vol. 1, p. 2, 2018. 9
- [37] R. Metoyer, Q. Zhi, B. Janczuk, and W. Scheirer. Coupling Story to Visualization: Using Textual Analysis as a Bridge Between Data and Interpretation. In *23rd International Conference on Intelligent User Interfaces*, pp. 503–507, 2018. doi: 10.1145/3172944.3173007 9
- [38] J. Nielsen. Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of the 1994 CHI Conference on Human Factors in Computing Systems*, pp. 152–158, 1994. doi: 10.1145/191666.191729 2
- [39] J. Nielsen. Ten Usability Heuristics. <http://www.nngroup.com/articles/ten-usability-heuristics/>, 2005. 2
- [40] M. Oppermann, R. Kincaid, and T. Munzner. VizCommender: Computing

- Text-based Similarity in Visualization Repositories for Content-based Recommendations. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):495–505, 2020. doi: 10.1109/TVCG.2020.3030387 3
- [41] A. Ottley, A. Kaszowska, R. J. Crouser, and E. M. Peck. The Curious Case of Combining Text and Visualization. In *EuroVis (Short Papers)*, pp. 121–125, 2019. doi: 10.2312/evs.20191181 2
- [42] A. Pandey, A. Srinivasan, and V. Setlur. MEDLEY: Intent-based Recommendations to Support Dashboard Composition. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1135–1145, 2023. doi: 10.1109/TVCG.2022.3209421 1, 2
- [43] J. Pinheiro and J. Poco. ChartText: Linking Text with Charts in Documents. *arXiv preprint arXiv:2201.05043*, 2022. doi: 10.48550/arXiv.2201.05043 9
- [44] Power BI. Power BI Data Stories Gallery. <https://community.fabric.microsoft.com/t5/Data-Stories-Gallery/bd/p/DataStoriesGallery>, 2023. Accessed: August 2023. 3
- [45] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics*, 25(1):682–692, 2018. doi: 10.1109/TVCG.2018.2864903 1, 2, 3, 5
- [46] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A Design Space of Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2366–2375, 2013. doi: 10.1109/TVCG.2013.120 2
- [47] E. Segel and J. Heer. Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, 2010. doi: 10.1109/TVCG.2010.179 1
- [48] V. Setlur, M. Correll, A. Satyanarayan, and M. Tory. Heuristics for Supporting Cooperative Dashboard Design. *IEEE Transactions on Visualization and Computer Graphics*, 30(1), 2023. doi: 10.1109/TVCG.2023.3327158 1, 2, 3, 7, 8, 9
- [49] D. Shi, X. Xu, F. Sun, Y. Shi, and N. Cao. Calliope: Automatic Visual Data Story Generation from a Spreadsheet. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):453–463, 2020. doi: 10.1109/TVCG.2020.3030403 9
- [50] M. Shin, J. Kim, Y. Han, L. Xie, M. Whitelaw, B. C. Kwon, S. Ko, and N. Elmquist. Roslingifier: Semi-automated Storytelling for Animated Scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 2022. doi: 10.1109/TVCG.2022.3146329 9
- [51] B. Shneiderman, C. Plaisant, M. S. Cohen, S. Jacobs, N. Elmquist, and N. Diakopoulos. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Pearson, 2016. 2
- [52] A. Srinivasan, M. Correll, L. Battle, V. Setlur, A. Crisan, et al. Toward a Scalable Census of Dashboard Designs in the Wild: A Case Study with Tableau Public. *arXiv preprint arXiv:2306.16513*, 2023. doi: 10.48550/arXiv.2306.16513 1, 3
- [53] A. Srinivasan, S. Drucker, A. Endert, and J. Stasko. Augmenting Visualizations with Interactive Data Facts to Facilitate Interpretation and Communication. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):672–681, 08 2018. doi: 10.1109/TVCG.2018.2865145 5
- [54] A. Srinivasan, T. Harshbarger, D. Hilliker, and J. Mankoff. Azimuth: Designing Accessible Dashboards for Screen Reader Users. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*, 16 pages, 2023. doi: 10.1145/3597638.3608405 9
- [55] A. Srinivasan and V. Setlur. BOLT: A Natural Language Interface for Dashboard Authoring. In *EuroVis 2023 - Short Papers*. The Eurographics Association, 2023. doi: 10.2312/evs.20231035 3
- [56] C. Stoiber, D. Ceneda, M. Wagner, V. Schetinger, T. Gschwandtner, M. Streit, S. Miksch, and W. Aigner. Perspectives of Visualization Onboarding and Guidance in VA. *Visual Informatics*, 6(1):68–83, 2022. doi: 10.1016/j.visinf.2022.02.005 9
- [57] C. Stokes and M. A. Hearst. Give Text A Chance: Advocating for Equal Consideration for Language and Visualization. In *NL VIZ: Workshop on Exploring Opportunities and Challenges for Natural Language Techniques to Support Visual Analysis*. IEEE, 2021. doi: 10.48550/arXiv.2404.00131 2
- [58] C. Stokes, V. Setlur, B. Cogley, A. Satyanarayan, and M. A. Hearst. Striking a Balance: Reader Takeaways and Preferences when Integrating Text and Charts. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1233–1243, 2022. doi: 10.1109/TVCG.2022.3209383 1, 2
- [59] H. Strobel, D. Oelke, B. C. Kwon, T. Schreck, and H. Pfister. Guidelines for Effective Usage of Text Highlighting Techniques. *IEEE transactions on visualization and computer graphics*, 22(1):489–498, 2015. doi: 10.1109/TVCG.2015.2467759 9
- [60] N. Sultanum, F. Chevalier, Z. Bylinskii, and Z. Liu. Leveraging Text-chart Links to Support Authoring of Data-driven Articles with VizFlow. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 17 pages, 2021. doi: 10.1145/3411764.3445354 2, 9
- [61] A. Tarrell, A. Fruhling, R. Borgo, C. Forsell, G. Grinstein, and J. Scholtz. Toward Visualization-Specific Heuristic Evaluation. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pp. 110–117, 2014. doi: 10.1145/2669557.2669580 2
- [62] M. Tory, L. Bartram, B. Fiore-Gartland, and A. Crisan. Finding their Data Voice: Practices and Challenges of Dashboard Users. *IEEE Computer Graphics and Applications*, 2021. doi: 10.31219/osf.io/pjwrx 1, 2
- [63] M. Tory and T. Moller. Evaluating Visualizations: Do Expert Reviews Work? *IEEE Computer Graphics & Applications*, 25(5):8–11, 2005. doi: 10.1109/MCG.2005.102 2
- [64] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2001. doi: 10.1097/01445442-198507000-00012 1
- [65] H. Vätätjää, J. Varsaluoma, T. Heimonen, K. Tiitinen, J. Hakulinen, M. Turunen, H. Nieminen, and P. Ihantola. Information Visualization Heuristics in Practical Expert Evaluation. In *BELIV Workshop*, pp. 36–43, 2016. doi: 10.1145/2993901.2993918 2
- [66] B. Victor. Scientific Communication as Sequential Art. *Online*. <http://worrydream.com/ScientificCommunicationAsSequentialArt>, 2011. 9
- [67] C. Ware. *Visual Thinking for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008. 2
- [68] S. Wexler, J. Shaffer, and A. Cotgreave. *The Big Book of Dashboards: Visualizing your Data using Real-world Business Scenarios*. John Wiley & Sons, 2017. doi: 10.1002/9781119283089 2
- [69] A. Wu, Y. Wang, M. Zhou, X. He, H. Zhang, H. Qu, and D. Zhang. MultiVision: Designing Analytical Dashboards with Deep Learning Based Recommendation. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):162–172, 2021. doi: 10.1109/TVCG.2021.3114826 3
- [70] X. Xu, A. Wu, L. Yang, Z. Wei, R. Huang, D. Yip, and H. Qu. Is It the End? Guidelines for Cinematic Endings in Data Videos. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 16 pages, 2023. doi: 10.1145/3544548.3580701 9
- [71] O. M. Yigitbasioglu and O. Velcu. A Review of Dashboards in Performance Management: Implications for Design and Research. *International Journal of Accounting Information Systems*, 13(1):41–59, 2012. doi: 10.1016/j.accinf.2011.08.002 2
- [72] C. Zheng and X. Ma. Evaluating the Effect of Enhanced Text-visualization Integration on Combating Misinformation in Data Story. In *2022 IEEE 15th Pacific Visualization Symposium (PacificVis)*, pp. 141–150. IEEE, 2022. doi: 10.48550/arXiv.2205.11911 9
- [73] Q. Zhi, A. Ottley, and R. Metoyer. Linking and Layout: Exploring the Integration of Text and Visualization in Storytelling. *Computer Graphics Forum*, 38(3):675–685, 2019. doi: 10.1111/cgf.13719 2, 9
- [74] J. Zong, I. Pedraza Pineros, M. K. Chen, D. Hajas, and A. Satyanarayan. Umwelt: Accessible Structured Editing of Multimodal Data Representations. In *Proceedings of the 2024 ACM CHI Conference on Human Factors in Computing Systems*, 20 pages. ACM, 2024. doi: 10.1145/3613904.3641996 2
- [75] T. Zuk, L. Schlesier, P. Neumann, M. S. Hancock, and S. Carpendale. Heuristics for Information Visualization Evaluation. In *Proceedings of the 2006 AVI Workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization*, 6 pages, 2006. doi: 10.1145/1168149.1168162 2