

An Analysis of Text Functions in Information Visualization

Chase Stokes , Anjana Arunkumar , Marti A. Hearst , and Lace Padilla 




Fig. 1: These designs demonstrate REPLACE MAPPINGS: the use of text elements to replace conventional mapping components in a chart, such as *Axes* and *Legends*. This function showcases the integration of text and visual design choices. Within this paper, we propose and analyze ten distinct text functions, including REPLACE MAPPINGS, and show how they can be organized into four factors.

Abstract—Text is an integral but understudied component of visualization design. Although recent studies have examined how text elements (e.g., titles and annotations) influence comprehension, preferences, and predictions, many questions remain about textual design and use in practice. This paper introduces a framework for understanding text functions in information visualizations, building on and filling gaps in prior classifications and taxonomies. Through an analysis of 120 real-world visualizations and 804 text elements, we identified ten distinct text functions, ranging from identifying data mappings to presenting valenced subtext. We further identify patterns in text usage and conduct a factor analysis, revealing four overarching text-informed design strategies: Attribution and Variables, Annotation-Centric Design, Visual Embellishments, and Narrative Framing. In addition to these factors, we explore features of title rhetoric and text multifunctionality, while also uncovering previously unexamined text functions, such as text replacing visual elements. Our findings highlight the flexibility of text, demonstrating how different text elements in a given design can combine to communicate, synthesize, and frame visual information. This framework adds important nuance and detail to existing frameworks that analyze the diverse roles of text in visualization.

Index Terms—Visualization, text, language, text function, factor analysis, design patterns.

1 INTRODUCTION

Visualization research has produced important empirical and theoretical work that provides recommendations for using graphical elements (e.g., [8, 15, 69, 72]) while historically overlooking the impact of text. Designers use text in almost every visualization, but unlike visual design standards that optimize data encodings for precision and clarity, text-related decisions often rely on designer intuition and conventions. Text plays a critical role in visualizations; titles, captions, and annotations provide context, guide attention and frame understanding. Language is a highly flexible tool, offering countless ways to convey an idea, making it a promising subject for deep analysis.

Recent studies have begun to reveal how variations in text placement, function, and style influence reader experience and comprehension

(e.g., [1, 32, 59–62, 75]), and prior work has built useful foundations for visualization-based text classifications [3, 11, 42, 43, 52]. However, there remain many open questions, including a gap in understanding the specific functions that textual elements perform. Previous efforts often focus on broad categories or isolated uses of text and may not clearly define the specific functions of text elements. The lack of a clear, functional framework creates a barrier to understanding how text works in concert with visual elements to support data communication.

We define **text function** as the specific role a text element plays in supporting the design and data communication of a visualization, focusing on the actions it performs and the target it operates on. The “action” terms borrow verbs directly from prior taxonomies in visualizations [11, 52]. The “target” terms were compiled through the exploratory analysis described in Sec. 3 as well as close readings of taxonomic definition in prior work [43, 52]. A *function* is distinct from the broader notion of the role of text, as studied in prior works [52, 59, 62], which describes the general purpose or effect of text in guiding reader interpretations.

In an effort to further clarify this distinction, we investigated our first research question: **RQ1: What are the functions of text in visualization designs?** This question seeks to identify and characterize the detailed functions individual text components fulfill. We develop this text function framework via a detailed qualitative analysis of 804 individual text components drawn from 120 visualizations from several

- Chase Stokes and Marti Hearst are with University of California, Berkeley.
E-mail: cstokes@ischool.berkeley.edu; hearst@berkeley.edu
- Anjana Arunkumar and Lace Padilla are with Northeastern University.
E-mail: xxx

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

domains, such as blogs, government reports, and news sources. In our framework, functions describe how visualizations use language at the level of an individual text component, analyzed independently from but in conjunction with its type (e.g., *Titles*, *Legends*, *Annotations*).

Through the qualitative analysis, we uncover several novel insights into the roles text plays in visualization design. For instance, we find that text usage in visualizations is often multifunctional; *Titles* can state the chart’s variables, synthesize data, compare values, and hint at emotional valence all at once. Another insight is shown in Fig. 1. REPLACE MAPPINGS is a previously unexamined function in which a text element replaces a conventional component of a visualization, as when data values on the chart replace the y-axis.

The detailed understanding of text function can be used to create stable characterizations of text in visualization design. Based on the foundational understanding of text functions distilled in **RQ1**, we explore **RQ2: What text design patterns emerge across visualizations, and which functions are engaged in each pattern?**

We conducted a factor analysis to investigate combinations of text functions within text elements and entire designs, aiming to uncover common design patterns. This analysis revealed four characterizing factors of how text is used across designs, which we named: Attribution and Variables, Annotation-Centric Design, Visual Embellishments, and Narrative Framing. These factors not only highlight trends in text usage but also help clarify how specific text choices can align with broader design goals.

Contributions: In this work, we build on existing frameworks for interpreting text in visualizations, focusing on the function of each text component with respect to its content and how it interacts with other elements in the design. We contribute:

- **A Set of Text Functions:** We develop a structured set of text functions, based on prior work and our own detailed analysis. These functions fill existing gaps in text taxonomies.
- **Application of Text Functions:** We apply these function definitions to more than 800 text components using a detailed codebook, achieving high inter-annotator agreement ($Mean = 0.97$). We show how text functions occur across individual text types (e.g., titles), how often a single text element exhibits multiple functions, and other notable uses distinct from overarching design patterns.
- **Design Patterns Based on Text Function:** We examine how text functions are used across text components within a given design by performing a factor analysis on the coded data. We extract four factors that represent the main trends in text design across the corpus. We conduct computational validation of these factors to ensure the robustness and stability of item associations.

Clearer guidance on text’s varied roles may help designers use language more effectively. By establishing a structured framework for text functions, we provide a foundation for *future* improvements to both human and automated design decisions.

2 RELATED WORK

This paper proposes a new framework for classifying the text used in information visualizations. This section summarizes related work on text in visualization, prior work on taxonomies in the information visualization space, and remaining open questions.

2.1 Text in Visualizations

A growing body of work focuses on the combination of text and visual information. Text elements of a visualization have a substantial role in directing and attracting attention when viewing a visualization. Text paragraphs, titles, legends, labels, and axes are fixated upon within the first three seconds of viewing a chart, typically before the reader examines the data displayed [14]. The content viewers recall from visualizations is often based on information or content they extract from the title [9]. Different titles lead to different information recalled from the same visualizations [36, 37].

Annotations and captions also affect conclusions drawn from visualizations, depending on the content. When specific features are

highlighted in captions, readers are more likely to mention those features in their takeaways, particularly if they align with visually salient elements [33, 75]. Annotations similarly guide interpretation, with participants more likely to conclude statistical or external information if it is explicitly written in an annotation [62].

Text can affect interpretations beyond conclusions; using annotations to highlight one outcome over another for a prediction task led to higher ratings of designer bias compared to neutral text conditions [59]. It is likely that text elements have differing impacts depending on the task or context. In the same study, text had only a minimal effect on the prediction task itself, despite influencing other perceptions of the chart.

Visualizations that use text to ‘focus’ the view of the data tend to be perceived by readers as clearer and more visually appealing [1, 35]. Assessment of probability representations revealed that important information was easier to identify with visuals but easier to comprehend with text [49]. The impact of multimedia communication may also depend in part on the audience. In one study examining decision-making under uncertainty, participants with lower working memory capacity benefited more from the combination of text and visual information than those with a higher working memory capacity [7].

Research endeavors have primarily focused on the impact of text elements on perceptions and understanding of visualized data, but there are also many open questions about designer *practice* regarding text elements. Prior work examining the role of text in interactive dashboards [63] demonstrates how text serves both navigational and interpretative roles. This work additionally considered the use of text formatting (e.g., italics) as a tool for designers to distinguish between different text elements or purposes. Text and natural language are increasingly incorporated into design tools, including question-answering pipelines [19, 31], automatically generated summaries [30, 46, 65], and assistance in data storytelling [44, 64].

Visualization researchers have demonstrated a growing interest in leveraging text effectively in visualization design. However, the concept of a text’s role or function across these studies is often ambiguously defined and context- or task-specific. We need to develop a structured and thorough understanding of text function in visualizations to develop better systems and guidelines for designers. Initial efforts in this area include several variations on taxonomies of text.

2.2 Visualization Taxonomies

2.2.1 Taxonomies of Visualization Tasks

Scholars have developed structured frameworks to help designers navigate the multitude of choices available in visual design. Taxonomies in visualization research categorize the elementary tasks that readers complete with charts [2, 39, 57], high-level comprehension tasks [3, 13, 68], abstract combinations of these tasks [11]. These frameworks help designers make informed choices on key aspects such as visual encodings, interactions, and spatial organization.

Low-level task taxonomies [2, 39, 57] guide designers on how visualizations might support analytical activities. These taxonomies break down visualization tasks into core components, including “filter,” “sort,” and “compare,” enabling designers to tailor their visualizations to specific analytic needs. The tasks have also been evaluated or examined according to different charts or data types, providing more specificity in how different design choices can support different tasks.

In contrast, high-level task taxonomies are typically more focused on the broader contexts of design or use. For example, Tory and Möller define a taxonomy based on the design models of visualization algorithms, distinguishing between continuous and discrete assumptions about data representation [68]. Burns et al. apply Bloom’s taxonomy to visualizations to map and understand different levels of user understanding [13]. This application of the educational taxonomy has also been used to measure graphical literacy [4, 50].

The abstract task taxonomy proposed by Brehmer and Munzner offers a bridge between low- and high-level task classifications by introducing a multi-level abstract framework structured around the questions of why, how, and what tasks are performed [11]. This work suggests a set of verbs that describe the perspective and goals of users (e.g., discover), search actions they may perform (e.g., lookup), and

elementary tasks they could perform with the data points (e.g., identify). This encompasses the *why*. The *what* refers to the possible inputs and outputs of the particular task, such as values, structures, or other visualization features. *How* refers to specific actions or interactions from the user (e.g., select, filter).

2.2.2 Taxonomies of Text In Visualization

Emerging research has sought to develop frameworks and taxonomies for textual elements in visualizations. Rahman et al. [52] built on the verbs from the abstract task taxonomy (identify, compare, summarize, and present) [11], conducting an extensive thematic coding of real-world annotations (both visual and textual). This work presents a design space of annotation types, which includes enclosure, connector, text, glyph, color, indicator, and geometric. The visual elements could be combined with text information as part of a single annotation grouping.

Using a “how, why, and what” structure, Rahman et al. outline a design space for the analytic purposes of the annotations, the strategies available, and the types of data needed to generate the annotations. Although this work incorporates text annotations, it does not account for specific communicative functions performed by text; we build on the verb structure from this taxonomy, and others [11] to explicitly assess the functionality of text.

Liu et al. [42] developed a structured design space for visualization titles based on a survey of existing practices. Their framework identified two key dimensions: generic information (i.e., variables and encoding schemes) and data features (i.e., mathematical operations such as trends and aggregations). For example, a title might simply state the variables shown (“Staff Counts Across Healthcare Systems”) or describe a specific trend in the data (“NHS Has Fewer Staff Than Some Counterparts”). Their work focused specifically on titles that could be generated automatically from visualization inputs, excluding external context (e.g., current events). We use this taxonomy to inform how we define generic and descriptive functionality.

Moving beyond single chart images, Hao et al. examined taxonomies of data-driven news articles [24], including how text appeared in visualizations. Titles often mirrored headlines, stating issues or emphasizing data features; captions tended to cite sources and scope, occasionally describing the visual content. Annotations typically labeled data values but sometimes conveyed thresholds or summaries.

Prior work on dashboard design [6, 55, 58, 63] also incorporates studies of how text ‘blocks’ are used to structure and support visualizations. Text in dashboards is often added to provide additional information beyond basic data facts. Rather than focus on visualization structure, our work examines the content *inside* the text blocks themselves.

A related but distinct line of work focuses on the use of text to improve accessibility, particularly through alternative (alt) text. To support alt text generation, Lundgard and Satyanarayan proposed a conceptual model with four semantic levels: *encoded* (e.g., chart type), *statistical* (e.g., extrema), *perceptual* (e.g., pattern synthesis), and *contextual* (e.g., current events). This model primarily provides guidance for writing alt text but has also been used for evaluating text content in and around visualizations [59, 62, 63, 75].

2.2.3 Gaps in Existing Taxonomies

After careful review of prior work, we identify several important gaps in the existing formalizations with respect to text in visualization:

1. Text is characterized only at a broad level.
2. One kind of text use is considered (e.g., titles, alt text), rather than classifying across all uses within a visualization.
3. A given text element is assigned only a single function.
4. Interactions between text elements in a design are not considered.

If these gaps remain unaddressed, designers and researchers risk relying on incomplete or oversimplified models that fail to capture the complexity of text use in visualizations, leading to suboptimal design decisions and understanding. To address these issues, we pose two guiding research questions. First, RQ1: **what are the functions of text in visualization designs?** This question addresses Gaps 1-3 by clarifying the *specific* roles text can play across a range of elements, including

cases where a single element serves multiple functions. Building on existing frameworks, we propose a detailed set of text functions.

We use this set of functions to then investigate RQ2: **what text design patterns emerge across visualizations, and which functions are engaged in each pattern?** Addressing Gap 4, this question considers how features of text elements, including their functions, interact and combine within a given design. Using a factor analysis on a set of text design variables to systematically group these functions, we provide a thorough assessment of the text design space, including possible design goals. Together, these questions establish a foundation for a structured and nuanced representation of text in visualization.

3 METHODOLOGY FOR FUNCTION FRAMEWORK CREATION

3.1 Corpus Creation

We used two collections of visualizations: the first for open coding of functions and the second to analyze the resulting set of functions¹. We examined a subset of real-world visualizations sourced from the Image-to-Information corpus [5], which itself draws stimuli from MASSVIS [9, 10]. The original Image-to-Information corpus consisted of 500 images curated by sampling from various sources equivalently and supplementing collected infographics with more standard visualization types. We selected this dataset since it represents a broad spectrum of “in the wild” visualization designs from diverse sources such as news media, government reports, and scientific literature.

For this paper, we examined a subset ($n = 120$) of the original corpus, focusing on visualizations that conveyed information through relatively common chart types and structures. To do so, we excluded a total of 380 designs. First, we filtered out 273 designs classified as “image” rather than “information” in the original study [5]. This filtering ensured that our focus remained on visualizations meant to inform readers.

We further refined this subcorpus by removing infographics (54), diagrams (25), tables (10), and interactive designs (2). Additionally, we omitted any designs with unreadable text due to poor image resolution. These criteria resulted in a set of 120 visualizations, each offering rich text examples for analysis. Compared to the original Image-to-Information corpus, our subset leans more heavily on visualizations from news media ($n = 46$) and government reports (19). While our corpus focuses on a narrower range of sources, it captures key visualizations aimed at general audiences, and our text function methodology can be extended to genre-specific studies, as discussed in Sec. 7.

3.2 Code Development

With insights from our review of prior taxonomies, we conducted open coding of 18 existing designs, including designs encountered in the real world and drawn from MASSVIS and other chart corpora [10, 43, 52]. Because the open coding was conducted on a small subset of designs in preparation for closed coding on the primary corpus, we drew from diverse sources to avoid missing potential text functions.

This open coding took place at the sentence or phrase level and consisted of writing short phrases to describe what functions the text was performing (e.g., highlighting an outlier in the data). We then performed axial coding, grouping together similar phrases to distill sets of candidate functions. We continued to review and discuss charts containing text that did not clearly fit within our initial function set, refining our categories accordingly. Through repeated cycles of open and axial coding, we developed a structured set of distinct text functions to capture the range of ways text contributes to visualization design.

Building on these coded insights, we returned to foundational taxonomies, such as the abstract task [11] and annotation taxonomies [52], to further refine the functions. From these frameworks, we adopted a verb-focused specification, borrowing the specific verbs (e.g., “Compare”) used in both taxonomies.

3.3 Function Coding

We coded the corpus of 120 selected visualizations and 804 distinct text components according to the set of text functions. We manually

¹Three images overlapped between the two collections.

extracted text from each visualization², recording either the exact content or summarizing repetitive elements such as categorical headers or date labels. For example, we used “set of categories” to describe repeated labels on a categorical axis. Multi-line text was treated as a single element to preserve context.

We coded additional features of each text element, as shown in Tab. 1. These codes drew from prior definitions when possible [63] and were derived via open coding and discussion. Specific definitions are available in supplementary materials. A second author reviewed all extracted text and independently coded this metadata, with discrepancies resolved through discussion.

Metadata	Description	Codes
Text Type	Category or position of text element used in visualization design	Titles, Subtitles, Annotations, Captions, Axes, Legends, Paragraphs
Non-Data Visual Elements	Visual elements associated with the text, if any	Arrows, Circles, Logos, Icons, Lines, Rectangles
Color Use	Role of color in the text element, if any	Encoding, Highlight, Style

Table 1: Coding types other than text functions, which are defined in Tab. 2. Not all text elements contain visual elements or color.

We then coded the presence of functions in each text element. To maintain rigor, the coding process was carried out iteratively, with a focus on achieving high reliability across all function codes [27, 47]. Initially, two coders independently analyzed a random subset of 20 designs, and interrater reliability (IRR) was calculated [21]. If the IRR did not surpass a threshold of 75% of codes with $\kappa > 0.8$, the coders discussed discrepancies and refined the codebook. This iterative process continued until the threshold was reached.

After two rounds of subset coding, we surpassed the reliability threshold, achieving an average kappa value of 0.97 across the codes. More than 75% of codes showed complete agreement between coders, although emotional engagement proved slightly more subjective, with a kappa value of 0.62. Each coder then coded the entire corpus independently, with discrepancies resolved through in-depth discussions among all four authors. The final codebook reflects these deliberations, providing a reliable framework for future analysis.

4 TEXT FUNCTIONS

To clarify the relationship between text and visual elements and identify recurring design patterns across visualizations, we developed a framework comprising ten distinct text functions. This framework was informed by existing research on visualization tasks and text taxonomies and was further refined through open coding of a representative set of existing visualizations. Tab. 2 summarizes the resulting functions, and Fig. 2 illustrates them. Frequency information for each function can be found in Tab. 3; text elements could serve multiple functions.

4.1 Text Function Framework

Text elements that IDENTIFY MAPPINGS communicate the conventions for interpreting the chart. *Axes* and *Legends* frequently serve this function by mapping data to visual channels such as position or hue. However, this function is not always handled by traditional components. On the right-hand side of Fig. 2, no formal y-axis exists. Instead, data labels (the number *Annotations* on the bars) take on the IDENTIFY MAPPINGS function and show how height relates to numeric value.

These same labels also serve a second function: IDENTIFY VALUES. By providing a complete set of numeric values for each bar in the ‘Spain’ category, they allow readers to retrieve exact values without estimating.

²In initial tests, large-language models (LLMs) extracted text fairly well with some errors (e.g., capturing axes). At the time of analysis, LLMs did not support automatic coding of text with our labels, so we did this work manually.

Crucially, IDENTIFY VALUES only applies when all relevant values are labeled; selective emphasis, such as tagging only an outlier, would fall under COMPARE VALUES.

Some text provides context beyond the chart’s encodings or data. The function PRESENT METADATA captures when text includes details like data sources, collection methods, or the intended audience. This function often appears in *Captions*, such as the note in Fig. 2 that reads, “Data from Eurostat and Spain’s Budget Ministry,” supporting transparency and credibility.

During open coding, we identified another distinct function: REPLACE MAPPINGS. Unlike text that complements conventional encodings, this function substitutes for them entirely. In the right of Fig. 2, for example, the y-axis is omitted. Instead, text on each bar communicates the encoded values, replacing what the axis would have shown. Similarly, *Annotations* such as ‘Spain’ and ‘Euro-Zone Average’ take the place of *Legends*. Even on the left side of Fig. 2, the legend title “Government deficit or surplus as a share of GDP” doubles as a replacement for a y-axis label. These examples illustrate how text can step in when conventional elements are minimized or removed.

Other text elements work to interpret or rephrase chart encodings for the viewer. The COMPARE MAPPINGS function translates visual encodings into more conceptual or accessible terms. For instance, the phrase, “Greater budget surplus,” on the right-hand side of Fig. 2 offers a qualitative interpretation of the vertical position. Unlike IDENTIFY MAPPINGS, which describe standard mappings, COMPARE MAPPINGS represents the *meaning* of the encoding in the context of the data.

When text highlights relationships between specific data points, it performs the COMPARE VALUES function. This can be as part of a longer phrase or as a single data label. This includes callouts to outliers, extremes, or other noteworthy comparisons. For instance, an *Annotation* identifying the minimum value (as on the right of Fig. 2) emphasizes how it stands apart from the rest, even if other points are left unlabeled.

Text can also summarize broader patterns. The function SUMMARIZE VALUES applies when text aggregates or calculates across multiple points, such as when describing an average or trend. On the left of Fig. 2, the *Subtitle*, “Spain is struggling to narrow its budget deficit amid a deep economic contraction,” reflects the full set of data for Spain and conveys a high-level trend. Both SUMMARIZE VALUES and COMPARE VALUES have a set of more precise subfunctions drawn from a low-level task taxonomy [2].

The same *Subtitle* also serves another function: SUMMARIZE CONCEPTS. This function captures broader conceptual framing or synthesis, often seen in *Titles* or *Subtitles*. Through close review of title-specific studies and taxonomies [36, 37, 42], we refined this category and identified two subtypes. SUMMARIZE CONCEPTS: SYNTHESIS involves interpretation or synthesis, distilling the chart into a takeaway (e.g., left of Fig. 2: “Spain is struggling...”). SUMMARIZE CONCEPTS: VARIABLES, by contrast, lists the variables or chart contents *without* added interpretation (e.g., right of Fig. 2: “Economic Targets...”). We also examined rhetorical strategies used in SUMMARIZE CONCEPTS, such as puns, associations, exaggerations, repetitions, etc. [24, 41].

Finally, text can introduce information not visible in the chart. PRESENT CONTEXT includes external background, such as domain knowledge or current events. In the left-hand chart in Fig. 2, the phrase “deep economic contraction” situates the data in a broader economic narrative. Even brief text, like the “Targets” *Annotation*, can signal contextual knowledge that these are not just values, but policy goals.

The PRESENT VALENCE SUBTEXT function promotes emotionally charged or value-laden language. For instance, describing Spain as “struggling” introduces an affective tone not inherent in the data presentation otherwise. To inform our definition and use of this function, we reviewed studies on affect in visualization design [38, 40]. Since we did not have access to designer intents, we focused instead on the text’s *potential* to elicit an emotional or affective response in viewers. A useful heuristic is whether the text could be reasonably rewritten in a more neutral form; if so, the original likely carries affective weight.

Table 2: Text functions in visualizations, constructed with a verb to describe the text’s action and a noun to capture what it acts upon. “Identify,” “Present,” “Compare,” and “Summarize” are verbs borrowed from prior work [11, 52]. Additional examples can be found in supplementary materials.

Verb	Noun	Definition
Q IDENTIFY	📍 MAPPINGS	Communicates conventions for reading the charts, i.e., the relationship between data point(s) and corresponding visual elements or structures, such as position, color, shape, or other visual channels.
Q IDENTIFY	📊 VALUES	Directly identifies and labels all data points with the relevant value, category, or other point-specific information. Serves as a reference for data values, rather than selective emphasis or comparison.
🖥 PRESENT	📄 METADATA	Provides information about the source of the data, the transformations applied to the data, the visualization elements, its purpose, the people involved in its creation, and its intended audience. Includes definitions of a data category, variable, or other terms used.
🔄 REPLACE	📍 MAPPINGS	Provides an unconventional representation for a conventionally presented element used to establish the basic structure of data encoding when such an element (e.g., axis, legend) is omitted from the display.
⚖ COMPARE	📍 MAPPINGS	Translates a data mapping into more understandable or contextually relevant terms by rephrasing or interpreting technical, complex, or abstract data mappings into more relational language.
⚖ COMPARE	📊 VALUES	Describes relationships among and between data points through direct comparison of points or groups. Text may also highlight one or multiple data points in comparison to the overall dataset.
📊 SUMMARIZE	📊 VALUES	Describes relationships among and between data points through aggregation (e.g., average) or mathematical function (e.g., addition). Text may also group or filter points along a given dimension.
📊 SUMMARIZE	🧠 CONCEPTS	Provides a high-level summary of some or all aspects of the chart; text can either provide information synthesis (subcode: SYNTHESIS) or describe the variables displayed (subcode: VARIABLES). Augmented by an additional taxonomy of rhetorical strategies [24, 41].
🖥 PRESENT	🌐 CONTEXT	Integrates contextual information, including background knowledge about the world (such as geographic, cultural, and political relationships), knowledge about current events, and domain-specific information stemming from expertise in a particular field of research or scholarship.
🖥 PRESENT	💡 VALENCE SUBTEXT	Promotes an emotional response, appealing to emotions, values, and personal experience. Conveys an emotional tone that would be diminished or lost in a more neutral phrasing.




Fig. 2: Two versions of a Wall Street Journal (WSJ) visualization demonstrate how changing text elements and the functions they perform can affect the overall information provided by the design. **Left:** Original WSJ design [26]. Most of the text elements contribute to the narrative framing of the data. This design falls under F4:Narrative Framing (see Sec. 5.6). **Right:** Visualization redesigned by this paper’s authors as a more annotation-centric (F2:Annotation-Centric Design; see Sec. 5.4) design. Data synthesis was also replaced with a neutral summary. This design replaces the legend and y-axis with annotations providing direct values. Many redundant functions are not labeled due to space considerations.











4.2 Comparison to Other Frameworks

In comparison to existing frameworks, our approach aligns with prior taxonomies while also offering meaningful extensions. Several of our functions, such as IDENTIFY MAPPINGS, SUMMARIZE VALUES, and COMPARE VALUES, align with the action-oriented verbs in Rahman et al.’s annotation taxonomy [52]. Similarly, Lundgard and Satyanarayan’s four-level semantic model [43] captures broader categories like contextual or encoded information, which correspond to our PRESENT CONTEXT and IDENTIFY MAPPINGS functions. Our framework builds on these foundations by introducing finer-grained, context-sensitive functions that interpret text in relation to its role within the overall chart

design. This added specificity enables more detailed analysis of design choices and more precisely reflects how text operates as an integrated part of visualization structure, rather than as isolated content.

Our schema also diverges from other taxonomies in its inclusion of functions such as REPLACE MAPPINGS and PRESENT VALENCE SUBTEXT, which capture nuanced roles of text that are largely absent from existing taxonomies. These functions were developed through the open coding procedures; existing taxonomies were unable to capture the role of the respective text. For example, while the semantic levels address text that describes visual encodings, they do not account for cases where text replaces traditional components, such as axes or

Table 3: Percent of designs ($n = 120$) containing each function, faceted by text type. Designs could contain multiple types of text serving the same function. Cells are shaded based on values *within* each column, with overall occurrence percentages highlighted in blue and type-specific percentages in gray. Annotations and subtitles were the most functionally diverse types of text.

Function	Percent of Corpus	Titles	Subtitles	Annotations	Captions	Axes	Legends	Paragraphs
 Identify Mappings	100%	11%	9%	58%	3%	75%	44%	2%
 Identify Values	56%	0%	2%	52%	0%	4%	1%	1%
 Present Metadata	68%	6%	12%	7%	53%	0%	4%	2%
 Replace Mappings	56%	8%	8%	38%	1%	4%	2%	2%
 Compare Mappings	19%	0%	1%	12%	0%	2%	2%	2%
 Compare Values	34%	5%	6%	24%	1%	1%	1%	3%
 Summarize Values	35%	8%	12%	13%	0%	0%	1%	4%
 Summarize Concepts	88%	78%	30%	13%	7%	1%	2%	3%
 Present Context	32%	3%	12%	18%	5%	2%	2%	8%
 Present Valenced Subtext	25%	21%	8%	2%	0%	0%	0%	2%

legends, as captured by REPLACE MAPPINGS. As another example, COMPARE MAPPINGS bridges L1 (elemental and encoded) and L2 (statistical and relational) levels [43] by connecting abstract encodings to relational terms. The potential for text to evoke emotion, represented by PRESENT VALENCED SUBTEXT, goes beyond the scope of most taxonomies. Additionally, our framework captures the fact that even seemingly simple text elements can perform *multiple* functions, something not commonly labeled in other schemes.

4.3 Exploratory Analysis Findings

We analyzed the prevalence and representation of different functions; further details can be found in supplementary materials. This analysis addresses our first research question: **What are the functions of text in visualization designs?** We examine how function frequencies and usage patterns align with established conventions in text design, and to what extent the function framework reveals additional or previously undercharacterized aspects of how text is used in visualizations.

4.3.1 Characterizing Titles and Subtitles

Prior research indicates that titles provide a strong initial impression of the visualization’s content and meaning [9, 36, 37], but the majority of prior work does not provide detailed analysis of the construction or interplay of titles and subtitles. Consider the two examples in Fig. 2:

- T1: “Economic Targets for Spain for 2012-2014: Government deficit or surplus as a % share of GDP”
T2: “Belt-Tightening: Spain is struggling to narrow its budget deficit amid a deep economic contraction.”

These examples exhibit text function structure that is common within our collection. Type T1 is labeled in our framework as SUMMARIZE CONCEPTS: VARIABLES and we find that in almost every case (94%), this type contains noun phrases but no verbs. In contrast, type T2 is labeled as SUMMARIZE CONCEPTS: SYNTHESIS, and these often contain verbs (68%).

T2 demonstrates another pattern common to some news visualizations: a short title exhibiting a rhetorical move, in this case a metaphor (“Belt-tightening”), as a bid to capture the reader’s attention. Using a taxonomy of headline strategies from advertising [41] and data-driven journalism [24], we analyzed rhetorical devices such as puns, associations, exaggerations, word omissions, and repetition, alongside broader distinctions between formal and informal registers [29]. We observed much higher rates of puns and metaphors in short *Titles* (three or fewer words; 62%) compared to *Subtitles* or longer *Titles* (3%). For example, the *Title* “Falling Rains” (Image 243) encapsulates both a variable (rain) and a data trend (declining rainfall), leveraging wordplay in the form of a pun to reinforce meaning and engage the reader. In contrast, *Subtitles* rarely used distinct headline rhetorical strategies, suggesting

that *Subtitles* serve a different role, often providing additional context rather than relying on wordplay or rhetorical techniques.

Titles that used rhetorical strategies often also PRESENT VALENCED SUBTEXT through associations or wordplay, which is a part of the “Narrative” factor. These short *Titles* are typically paired with longer *Subtitles*. Additionally, by using the double meaning in a word, *Titles* can simultaneously perform SUMMARIZE CONCEPTS: SYNTHESIS and SUMMARIZE CONCEPTS: VARIABLES.

Although *Titles* are a common element in visualizations, 26 images in our corpus omitted them entirely. In these cases, text information was primarily conveyed through *Annotations*, with an average of 3.65 *Annotations* per image ($SD = 2.56$). These *Annotations* were most commonly used to IDENTIFY VALUES, though some engaged in data synthesis, summarization, or comparison. *Annotations* were the predominant location for COMPARE VALUES, allowing designers to position comparative statements directly alongside relevant data. This pattern reflects an overlap between the Annotation-Centric Design and Narrative Framing factors described in Sec. 5. *Annotations* can both serve as substitutes for conventional chart elements and contribute to storytelling by emphasizing comparisons.

4.3.2 Replacing Axes and Legends

Across our corpus, 56% of designs omitted at least one traditional mapping element, instead using REPLACE MAPPINGS. This function was most common in *Annotations* but also appeared in *Titles* and *Subtitles*, where text elements assumed roles typically associated with axis labels. *Axes* were replaced slightly more often than *Legends*, likely influenced by their more frequent presence in visualizations; 31 images replaced *Legend*, and 41 images replaced *Axes*. Six designs replaced both a *Legend* and an *Axis* with other text elements. We observed three primary representations of REPLACE MAPPINGS (see Fig. 1).

The most common use of REPLACE MAPPINGS replaced *Legends* with *Annotations* ($n = 28$), sometimes incorporating the use of color directly into the text ($n = 14$). For example, the right image in Fig. 1 uses color in the “Democrat” and “Republican” *Annotations*, allowing the designer to avoid finding suitable space for a *Legend*. This method is a form of direct labeling [1, 20, 35], reducing the need for users to reference a *Legend*. In cases where color was not used, the data points were directly labeled, or the color encoding was simply described in the text, as in the case of Image 403 in Fig. 3. An *Annotation* explicitly states that red shading indicates uncertainty, eliminating the need for a separate legend box. This method may be particularly useful when space is limited, or the color encodings require more explanation than a *Legend* affords.

Titles or *Subtitles* could also be used to present *Axis* information ($n = 17$), eliminating the need for a separate axis label. This approach could reduce redundant text while maintaining a neutral framing of the data. For example, the center image in Fig. 1 uses the *Subtitle*

“Height in metres, (*year of completion*),” which provides the same information as an axis label but in a more prominent location. This method aligns with minimalist design practices, though it requires viewers to synthesize information across different text elements rather than relying on a localized *Axis* label.

Axes could also be replaced with direct data labels in the form of *Annotations* ($n = 14$). Instead of requiring viewers to estimate values along a scale, designs used numeric values positioned directly within or beside data points. This strategy can be seen in the left image of Fig. 1 where the percent values are directly labeled and the 0 to 100 y-axis is only implied. In this case, the design prioritizes precision and reduces the need for external reference points. However, it can also increase text density, which may become overwhelming in data-dense visualizations. For example, in Fig. 2, the decision to label Spain’s data directly required widening the design.

4.3.3 Single Text, Multiple Functions

Within our corpus, 39% of text components served multiple functions, with *Annotations* being the most frequently multifunctional (66%), followed by *Subtitles* (58%) and *Titles* (51%). These text elements often worked to construct a visualization’s narrative structure rather than just supporting visual encodings.

Several common function pairings emerged from our analysis. SUMMARIZE VALUES and COMPARE VALUES frequently co-occurred, with 58% of SUMMARIZE VALUES instances also involving COMPARE VALUES. The combination of these functions suggests that data synthesis often involved both aggregation and direct comparisons. This kind of combination may allow for complex data operations and calculations to be summarized in a single sentence or phrase.

Similarly, PRESENT VALENCE SUBTEXT almost always (90%) appeared in combination with SUMMARIZE CONCEPTS and, to a lesser extent (28%), PRESENT CONTEXT. This pattern highlights the role of some *Titles* and *Subtitles* in shaping emotional framing and narrative context rather than presenting data in a purely neutral manner, as we explored in Sec. 4.3.1. The frequency of this function combination highlights the strategies used by designs to engage with readers.

Text that expressed the COMPARE MAPPINGS functions also frequently expressed SUMMARIZE VALUES (22%). It was also common for COMPARE MAPPINGS to co-occur (34%) with REPLACE MAPPINGS; in these cases, the design would typically omit one or both *Axes* in favor of *Annotations*.

Some functions tended to operate independently of others. PRESENT METADATA was largely independent of other text functions, with its highest co-occurrence being 11% with IDENTIFY MAPPINGS, and even lower with all other functions. These findings suggest that although some text functions may combine to form layered interpretative structures, others remain more specialized or independent, serving distinct roles within the visualization.

5 DESIGN PATTERNS BASED ON TEXT FUNCTIONS

To investigate how text functions and design features interact to form latent constructs — underlying patterns that emerge from correlated text and design elements in visualizations — we conducted an exploratory factor analysis (EFA) [22, 66] with varimax rotation in R Studio using the psych package [51, 54]. This analysis addresses our second research question: **What text design patterns emerge across visualizations, and which functions are engaged in each pattern?** Similar clustering analyses have been performed in prior work [6, 55, 58].

We included binary features based on: our proposed text **functions**, text **type** (e.g., *Title*), image **domain** (e.g., News), **color** use (e.g., encoding), **visual elements** associated with the text (e.g., arrows), and normalized **word count** measures (e.g., total word count). See Tab. 1 for more information on the text metadata categories. IDENTIFY MAPPINGS was not included in this analysis since it appeared in every design. Further details can be found in the supplementary materials.

5.1 Factor Selection

We examined the feature correlation matrix to identify relationships between text properties. As expected, there were strong correlations

among certain variables; PRESENT METADATA and *Captions* had the highest correlation ($r = 0.74$). At the same time, many features were only weakly correlated, reinforcing the need for a multifactor solution. To determine the number of factors, we used Bayesian information criterion (BIC), model complexity, and eigenvalues [25]. These measures collectively supported a 4-factor model, which best captured variation in text use while preserving factor independence and interpretability.

5.2 Factor Analysis Evaluation

To assess the reliability of the factor analysis in identifying latent textual constructs within our corpus, we deployed several validation approaches, including blind expert categorization, factor reliability assessment through resampling, and permutation testing.

As an initial evaluation, three coauthors — blinded to the factor analysis results — independently categorized a subset of visualizations from the corpus. Using the four identified factors as a reference, they predicted which factor each visualization would be most associated with. This blind categorization was performed for five images, and their classifications showed strong alignment with the factor loadings derived from the analysis, providing initial support for factor validity.

We further evaluated factor reliability through resampling and permutation testing. For resampling, we randomly selected 75% of images in each iteration and re-ran the 4-factor analysis, extracting factor loadings. The 75% subset size balanced sample representativeness and computational feasibility. We repeated this process for 50 iterations, determined empirically through stabilization of intraclass correlation (ICC) values. ICC values showed minimal change beyond 50 iterations, indicating convergence. After 50 iterations, ICC values for all four factors exceeded 0.93, confirming high factor reliability.

Additionally, we tested robustness by removing the top 10% of images based on their factor loadings to check if a small subset of highly weighted images disproportionately influenced the factor structure. The resulting factors remained comparable to those from the full dataset, supporting the validity of our extracted factors. These findings reinforce that the identified factors reflect underlying text relationships in visualizations rather than being artifacts of specific samples or outliers.

Together, these validation steps ensure that the factor analysis reliably captures meaningful text-based constructs within the visualization corpus, reinforcing the robustness of our findings.

Our exploratory factor analysis identified four distinct approaches to text design; examples of each factor along with its loadings can be seen in Fig. 3. We assigned names to these factors to make them easier to reference and understand [55], while acknowledging that the names may not capture every specific variable perfectly. Variables with loadings above 0.6 were considered strongly associated; variables with loadings below 0.3 were considered weakly associated. We supplemented the factor analysis with additional exploratory analyses to identify cross-factor findings and practices.

The right-hand version of the visualization shown in Fig. 2 illustrates the use of these factors. The original version on the left aligns with Narrative Framing. We converted this design to align with the variables that load highly on Annotation-Centric Design, by shifting the emphasis toward IDENTIFY VALUES and *Annotation*.

5.3 Factor 1: Attribution and Variables

We refer to F1 as Attribution and Variables, since this factor has relatively strong loadings from PRESENT METADATA and SUMMARIZE CONCEPTS: VARIABLES. Overall, loadings indicate that text primarily plays a neutral, factual role.

The highest loading item in this factor is *Captions*, which overwhelmingly served to PRESENT METADATA (93%), following standard journalistic conventions for citing sources and providing attribution [24]. *Titles* were moderately associated with this factor, most of performed SUMMARIZE CONCEPTS: VARIABLES, providing no additional framing to the data (87%). *Subtitles* also had a moderate loading on this factor. Overall, Attribution and Variables appears to capture common text design practices in journalism visualizations, with a moderate loading from news sources. Conversely, blog and media-based visualizations were moderately negatively associated with F1.

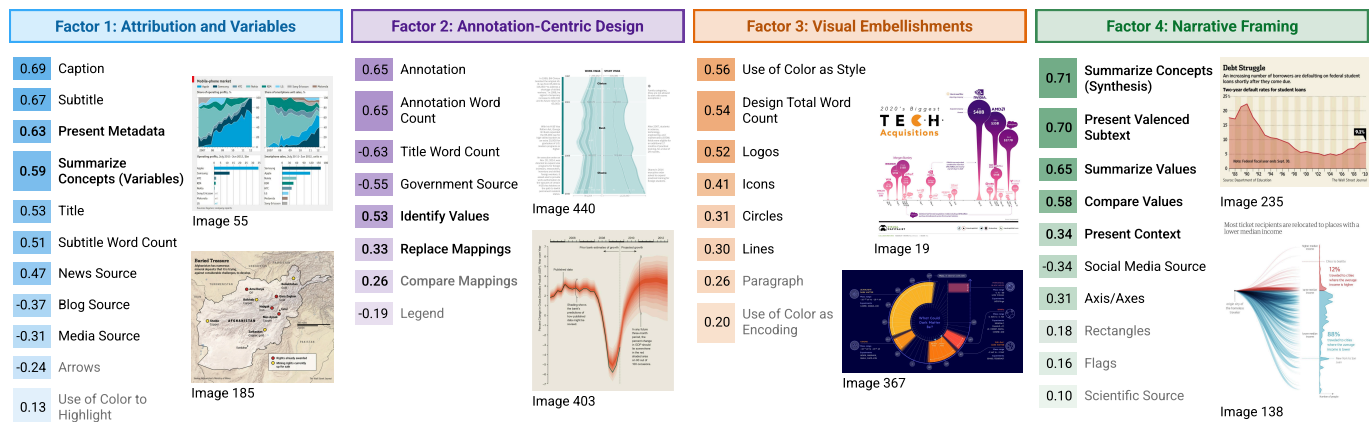


Fig. 3: Four factors identified in exploratory factor analysis, including the relevant items, their loadings, and example images for each factor. These factors represent four patterns of text design practices in visualizations: Attribution and Variables, Annotation-Centric Design, Visual Embellishments, and Narrative Framing. Loading scores indicate the strength of each variable's association with the factors. Example images shown are also available in supplementary materials for reference.

One common design pattern in this factor is highly neutral text, where text functions almost entirely PRESENT METADATA and SUMMARIZE CONCEPTS: VARIABLES. For example, in Image 55 in Fig. 3, nearly all non-axis text is dedicated to these neutral functions, resulting in a design where text serves as an organizational and attributional tool rather than contributing to interpretation.

A second pattern within this factor uses long *Subtitles* to provide additional context; *Subtitle* word length also had a moderate loading on F1. For instance, in Image 239 in Fig. 3, the *Title* is the elliptical 'Buried Treasure,' and the *Subtitle* states, "Afghanistan has numerous mineral deposits that it is trying, against considerable challenges, to develop." The *Title* offers only minimal insight into the data, but the *Subtitle* elaborates on the content displayed.

Note that this *Subtitle* introduces contextual information; most examples within this factor contain text that is largely but not 100% neutral. These sets of loadings suggest that F1 captures a broader "fact-based" focus to charts. Designs in the Attribution and Variables factor tended to have an emphasis on neutral summaries and statements of attribution but could still incorporate additional context or data synthesis through the use of longer *Subtitles*.

5.4 Factor 2: Annotation-Centric Design

We refer to F2 as Annotation-Centric Design, since it encompasses a design approach characterized by a high density of *Annotations*, the use of text to IDENTIFY VALUES, and substitutions in the form of REPLACE MAPPINGS. Designs in this factor tended to embed details directly within the chart itself rather than using longer *Titles* to describe data features.

This pattern reveals a possible design strategy that prioritizes in situ explanation, with positive loadings for both *Annotations* (0.65) and *Annotation* word count (0.65), alongside a negative loading for *Title* word count (-0.63). By distributing a possibly lengthy title within the chart, the text density of the design may not be greatly increased by added *Annotations*. Based on the moderate loading of IDENTIFY VALUES (0.53) also within F2, these annotations likely also provide additional data precision in many cases.

REPLACE MAPPINGS had a moderate loading (0.33) in this factor (see 4.3.2 for more detail on this function), likely corresponding in part with the use of IDENTIFY VALUES and the weak negative loading of *Legends*. Government-produced visualizations also had a negative loading, which suggests that government designs may rely more on conventional data presentation methods, with fewer *Annotations* replacing traditional mapping elements.

Although not a major component of F2, COMPARE MAPPINGS had a weakly positive association with the factor. *Annotations* tended to be the primary location for the COMPARE MAPPINGS function (Tab. 3).

COMPARE MAPPINGS had a relatively equal but negative loading on Attribution and Variables, indicating that F1 designs tend to rely on more conventional visual mappings rather than creative uses of text, such as COMPARE MAPPINGS or REPLACE MAPPINGS.

5.5 Factor 3: Visual Embellishments

F3 is referred to as Visual Embellishments, since it primarily captures how text interacts with stylistic elements in a design, highlighting the use of color, icons, and other graphical elements to add emphasis. Unlike other factors, which have strong loadings from text functions, variables in this factor primarily reflect aesthetic decisions. In addition to using color to style the text, designs in this factor incorporate embellishments, often connecting the text to a specific data point (with a circle or line) or to a higher-level concept (using logos or icons).

The strongest-loading item in this factor is the use of color in text for style purposes. This includes cases where text is colored for reasons beyond standard data encoding, such as matching a designer's branding. For example, in Image 19 in Fig. 3, part of the *Caption* is rendered in green to align with the logo of the designer group (Visual Capitalist), reinforcing the visual identity of the source.

Word count also strongly loaded onto this factor, indicating that these designs tended to feature longer text components. On average, designs in F3 contained 1.5 more text elements per visualization ($Mean = 9.4, SD = 3.1$) than those in other factors ($Mean = 7.9, SD = 3.6$). This may help explain the higher word counts observed in these designs since more elements contribute to the overall textual density.

Beyond color and word count, the presence of additional visual elements (e.g., logos, icons, circles, and lines) was another key feature of Visual Embellishments. Most visual elements (87%) appeared alongside *Annotations*, with only four appearing near *Captions*. Visual embellishments were not used alongside *Titles* or *Subtitles* in our corpus, reinforcing the idea that these stylistic choices may be primarily applied to supporting text elements, such as annotations, rather than typically high-level text elements, such as titles.

Although text functions are essential for understanding how text supports visualization design, the combination of variables in Visual Embellishments highlights that visual styling can serve a distinct role. The presence of F3 underscores the need for future research into how visual styling may interact with or operate separately from text function in visualization design.

5.6 Factor 4: Narrative Framing

We refer to F4 as Narrative Framing, due to text usage patterns that align with data storytelling practices [53,56]. SUMMARIZE CONCEPTS: SYNTHESIS and PRESENT VALENCE SUBTEXT had the strongest loadings (> 0.7), followed by SUMMARIZE VALUES and COMPARE

VALUES (> 0.5). PRESENT CONTEXT also had a moderate loading on this factor (0.34). Visualizations associated with this factor could emphasize specific takeaways by integrating text that framed the data in relation to broader events or real-world implications.

Variables in this factor represent a narrative approach to visualization design. The strong association with SUMMARIZE CONCEPTS: SYNTHESIS indicates these visualizations frequently contain text that distills complex information into concise summaries, and the nearly equal loading for PRESENT VALENCE SUBTEXT suggests these texts often incorporate emotional language, framing, or use of associations/puns.

The moderately strong loadings for SUMMARIZE VALUES and COMPARE VALUES further characterize this factor through text elements that highlight specific data points and their relationships, sometimes including PRESENT CONTEXT, which had a moderately weak loading on F4. Some of the association between these factors at the image level may have its roots in the specific text elements themselves; several function combinations that were characteristic of this factor also frequently combined within individual text elements as well (see Sec. 4.3.3).

Social media visualizations had a moderate negative association with this factor, suggesting that charts shared on social platforms may rely less on in-chart data synthesis. Although our corpus did not have a high degree of visualizations that were collected from social media, this suggests an interesting future direction of research to better understand the prevalence of these factors across contexts.

6 DISCUSSION

We addressed two major research questions: **What are the functions of text in visualization designs and which functions are engaged with specific emergent design patterns?** By developing and applying a functional framework grounded in real-world examples, we aim to establish a baseline for how textual elements operate in practice: what they do, where they appear, and how they co-occur. This work augments and improves on previous taxonomies by providing more specific characterization of functions, examining multiple types of text, considering interactions between text and visual elements, and capturing multifunctionality within a single text element.

One striking result is the prevalence of multifunctional text elements. These elements challenge assumptions that each component of a visualization serves a single, fixed role, opening up new design questions: when is multifunctionality effective, and when might it obscure meaning? Our framework offers one way to answer that by supporting structured audits of existing designs through a combination of functional analyses and reader assessments or feedback. For example, a designer or evaluator might identify areas where additional text could clarify ambiguous encodings, or where overloaded annotations might be better split into separate components with distinct functions.

Beyond its research contributions, this framework also has the potential to support visualization designers who seek greater structure and intentionality in how they use text. Having a clear set of functional roles offers a way to systematically plan, revise, and evaluate text components. Designers can more effectively align their text with specific communicative goals or contexts.

Within this paper, we demonstrate the use of factor-motivated redesigns (Fig. 2). The factors offer a structured way to modify text elements based on communicative goals. For example, visualizations associated with Visual Embellishments may be well suited for public-facing reports and marketing materials, where color and graphical elements can enhance engagement. In contrast, when presenting the same data to stakeholders, visualizations that score higher on Narrative Framing may be more appropriate, in which designers emphasize key takeaways and persuasive messaging. Designers could intentionally shift between text design strategies depending on audience, purpose, and available space.

These patterns could also serve as starting points or templates. By providing a structured framework of standard text functions and their co-occurrence patterns, our framework offers a foundation for assessing data visualization designs to be sure the text is meaningful and contextually appropriate; our set of codes provides a start to a corpus that could also be used to evaluate visualizations in the future.

This work contributes to the foundation for a deeper theoretical understanding of text in visualization, moving beyond broad categorizations of text content to a functional approach that accounts for the diverse roles text plays in structuring and shaping visual information. The ability to differentiate designs based solely on text design practices further underscores its role as an active design element. By systematically defining and analyzing text functions, we contribute to a growing body of research that situates text as a fundamental design element rather than a secondary consideration. Future studies can build on and adapt this framework to continue expanding our understanding and recommendations for text design in visualizations.

7 LIMITATIONS AND FUTURE WORK

This study provides a broad analysis of text functions in visualization design, but there are several limitations to consider. One constraint of this work is that our corpus is primarily composed of visualizations drawn from prior research, meaning it may not fully reflect the most recent trends in text design. Additionally, the corpus is largely made up of news visualizations, which may limit the applicability of these findings to other contexts. Expanding future analyses to include a broader range of domains could provide additional insights into how text functions vary depending on audience, purpose, and platform.

Building on these findings, several directions for future research could further refine our understanding of text in visualization design. A targeted, domain-specific analysis of text function would provide valuable insights into how text is used in specific contexts, such as social media posts or scientific reports. For example, visualizations shared on social media may rely more heavily on captions and post text to provide context rather than integrating text directly into the chart itself. Interactive and multi-visualization designs (i.e., dashboards) also warrant further investigation. A deeper understanding of text in visualization design also has the potential to provide more useful explanations of data for audiences with lower graphical literacy.

This work characterizes the current use of text in visualizations but does not assess the *usability* of these choices to further validate the framework. Future work should also consider user studies to evaluate the impact of different text design choices on reader comprehension, trust, and engagement. Specifically, examining the role of text function on audience interpretation would help to shape guidelines or educational materials around text design for visualizations. Although this work explores the practices in existing designs, there is an additional need for researchers to establish empirical recommendations.

Another area of future work could use automation for text classification and generation. Current AI systems frequently omit or make errors in essential text elements such as titles, creating notable gaps in automated chart design [16, 67]. Training an LLM on these functions and datasets could enable more efficient extraction and classification of text in visualizations. This would provide a faster way to track evolving text design practices in public charts. Additionally, AI-assisted visualization systems could also provide suggestions or guidance to designers, helping them make more informed text design choices.

8 CONCLUSION

This work provides a detailed framework for understanding text functions in visualization, expanding beyond broad classifications. Through an analysis of 120 designs and 804 text elements, we identify ten distinct functions. Our factor analysis further uncovers four key factors of text design, demonstrating how designers combine different textual features to support visual communication. By combining fine-grained text functions with broader design factors, this framework highlights practices and opportunities for intentional text integration in visualizations. Future work should further validate this framework and explore how text functions influence reader interpretations, how function-based frameworks might inform AI-assisted design tools, and how visualization designers can use this framework to support more intentional text design. This work lays the foundation for both critique and creation: it provides researchers with additional tools to study text systematically, and gives designers a vocabulary for using it more purposefully.

SUPPLEMENTAL MATERIALS

All supplemental materials are available on OSF ([link](#)), released under a CC BY 4.0 license. In particular, they include (1) data and analysis files, (2) images from the corpus analyzed, (3) figures used in this paper.

FIGURE CREDITS

A. Figure 1 image credits: Img 131: Alyson Hurt / NPR [original [28]] and Flowing Data [reprint [73]], September 2015 ; Img 70: Graphic Detail / Economist, October 2011 [17] ; Img 448: Neil King Jr. / Wall Street Journal, September 2012 [34]

B. Figure 2 image credit (Img 241): Jonathan House, David Román / Wall Street Journal, December 2011 [26]

C. Figure 3 image credits: Img 55: Graphic Detail / Economist, September 2012 [18] ; Img 185: Dion Nissenbaum / Wall Street Journal, June 2012 [48] ; Img 440: Samuel Granados / Washington Post, February 2017 [23] ; Img 403: Jen Christiansen / Flowing Data [reprint [74]], September 2019 ; Img 19: Omri Wallach / Visual Capitalist, December 2020 [71] ; Img 367: Samuel Velasco / Quanta Magazine, November 2020 [70] ; Img 235: Josh Mitchell and Rachel Louise Ensign / Wall Street Journal, September 2012 [45] ; Img 138: Nadiyah Bremer / Visual Cinnamon [reprint [12]], December 2017

ACKNOWLEDGMENTS

The authors wish to thank Michael Correll, Will Wang, Md Dilshadur Rahman, and Simone Laszuk for early comments on the manuscript, as well as the anonymous reviewers for helpful feedback which improved the paper. This work was supported in part by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 2146752.

REFERENCES

- [1] K. Ajani, E. Lee, C. Xiong, C. N. Knaflitz, W. Kemper, and S. Franconeri. Declutter and focus: Empirically evaluating design guidelines for effective data communication. *IEEE Transactions on Visualization and Computer Graphics*, 28(10):3351–3364, 2021. 1, 2, 6
- [2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 111–117. IEEE, 2005. 2, 4
- [3] R. Amar and J. Stasko. A knowledge task-based framework for design and evaluation of information visualizations. In *IEEE Symposium on Information Visualization*, pp. 143–150, 2004. doi: 10.1109/INFVIS.2004.101, 2
- [4] J. B. Arneson and E. G. Offerdahl. Visual literacy in bloom: Using bloom’s taxonomy to support visual learning skills. *CBE—Life Sciences Education*, 17(1):ar7, 2018. 2
- [5] A. Arunkumar, L. Padilla, G.-Y. Bae, and C. Bryan. Image or information? examining the nature and impact of visualization perceptual classification. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 3
- [6] 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, 2023. doi: 10.1109/TVCG.2022.3209448 3, 7
- [7] M. Bancilhon, A. Wright, S. Ha, R. J. Crouser, and A. Ottley. Why combining text and visualization could improve bayesian reasoning: A cognitive load perspective. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2023. 2
- [8] J. Bertin. *Semiology of graphics*. University of Wisconsin Press, 1983. 1
- [9] 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. 2, 3, 6
- [10] M. A. Borkin, A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2306–2315, 2013. doi: 10.1109/TVCG.2013.234 3
- [11] 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. 1, 2, 3, 5
- [12] N. Bremer. Techniques for data visualization on both mobile & desktop. <https://www.visualcinnamon.com/2019/04/mobile-vs-desktop-dataviz/>, 2019. (accessed March 26, 2025). 10
- [13] A. Burns, C. Xiong, S. Franconeri, A. Cairo, and N. Mahyar. How to Evaluate Data Visualizations across Different Levels of Understanding. In *2020 IEEE Workshop on Evaluation and Beyond-Methodological Approaches to Visualization (BELIV)*, pp. 19–28. IEEE, 2020. 2
- [14] Z. Bylinskii, M. A. Borkin, N. W. Kim, H. Pfister, and A. Oliva. Eye fixation metrics for large scale evaluation and comparison of information visualizations. In *Workshop on Eye Tracking and Visualization*, pp. 235–255. Springer, 2015. 2
- [15] S. K. Card, J. D. Mackinlay, and B. Shneiderman. Using vision to think. *Readings in information visualization: using vision to think*, pp. 579–581, 1999. 1
- [16] N. Chen, Y. Zhang, J. Xu, K. Ren, and Y. Yang. Viseval: A benchmark for data visualization in the era of large language models. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2024. doi: 10.1109/TVCG.2024.3456320 9
- [17] G. Detail. Tottering towers. <https://www.economist.com/graphic-detail/2011/10/24/tottering-towers>, 2011. (accessed March 26, 2025). 10
- [18] G. Detail. Mega phone. <https://www.economist.com/graphic-detail/2012/09/12/mega-phone>, 2012. (accessed March 26, 2025). 10
- [19] K. Dhamdhere, K. S. McCurley, R. Nahmias, M. Sundararajan, and Q. Yan. Analyza: Exploring data with conversation. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*, pp. 493–504, 2017. 2
- [20] 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*, 22(3):110–161, 2021. 6
- [21] M. Gamer, J. Lemon, I. Fellows, and P. Singh. Various coefficients of interrater reliability and agreement (r package version 0.84.1)[computer software], 2012. 4
- [22] R. L. Gorsuch. Exploratory factor analysis. In *Handbook of multivariate experimental psychology*, pp. 231–258. Springer, 1988. 7
- [23] S. Granados. How today’s visa restrictions might impact tomorrow’s america. <https://www.washingtonpost.com/graphics/national/visas-impact/>, 2017. (accessed March 26, 2025). 10
- [24] S. Hao, Z. Wang, B. Bach, and L. Pschetz. Design patterns for data-driven news articles. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, article no. 231, 16 pages. Association for Computing Machinery, 2024. doi: 10.1145/3613904.3641916 3, 4, 5, 6, 7
- [25] J. C. Hayton, D. G. Allen, and V. Scarpello. Factor retention decisions in exploratory factor analysis: A tutorial on parallel analysis. *Organizational research methods*, 7(2):191–205, 2004. 7
- [26] J. House and D. Román. Spain unveils new set of overhauls. <https://www.wsj.com/articles/SB10000872396390443916104578021692765950384>, 2012. (accessed March 26, 2025). 5, 10
- [27] D. J. Hruschka, D. Schwartz, D. C. St. John, E. Picone-Decaro, R. A. Jenkins, and J. W. Carey. Reliability in coding open-ended data: Lessons learned from hiv behavioral research. *Field methods*, 16(3):307–331, 2004. 4
- [28] A. Hurt. Coal, gas, nuclear, hydro? how your state generates power. <https://www.npr.org/2015/09/10/319535020/coal-gas-nuclear-hydro-how-your-state-generates-power>, 2015. (accessed March 26, 2025). 10
- [29] A typology of language registers. Standard, International Organization for Standardization, Geneva, CH, 2018. 6
- [30] S. Kantharaj, R. T. Leong, X. Lin, A. Masry, M. Thakkar, E. Hoque, and S. Joty. Chart-to-text: A large-scale benchmark for chart summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4005–4023, May 2022. doi: 10.18653/v1/2022.acl-long.277 2
- [31] D. H. Kim, E. Hoque, and M. Agrawala. Answering questions about charts and generating visual explanations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. 2
- [32] 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*, pp. 1–11, 2021. 1
- [33] Y.-S. Kim, K. Reinecke, and J. Hullman. Explaining the gap: Visualizing one’s predictions improves recall and comprehension of data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 12 pages, p. 1375–1386, 2017. doi: 10.1145/3025453.3025592 2
- [34] N. King Jr. Lost in translation: Gop struggles with hispanics. <https://www.wsj.com/articles/SB10000872396390444082904577605>

- 824074935132, 2012. (accessed March 26, 2025). 10
- [35] C. N. Knaflitz. *Storytelling with data: A data visualization guide for business professionals*. John Wiley & Sons, 2015. 2, 6
 - [36] 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*, pp. 1–12, 2018. 2, 4, 6
 - [37] 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*, pp. 1–13, 2019. 2, 4, 6
 - [38] X. Lan, Y. Wu, and N. Cao. Affective visualization design: Leveraging the emotional impact of data. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 4
 - [39] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry. Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI Workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization*, BELIV '06, 5 pages, p. 1–5. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1168149.1168168 2
 - [40] E. Lee-Robbins and E. Adar. Affective learning objectives for communicative visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1–11, 2022. 4
 - [41] J. H. Leigh. The use of figures of speech in print ad headlines. *Journal of advertising*, 23(2):17–33, 1994. 4, 5, 6
 - [42] C. Liu, Y. Guo, and X. Yuan. Autotitle: An interactive title generator for visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 1, 3, 4
 - [43] 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 1, 3, 5, 6
 - [44] R. Metoyer, Q. Zhi, B. Janczuk, and W. Scheirer. Coupling story to visualization: Using textual analysis as a bridge between data and interpretation. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces*, IUI '18, 5 pages, p. 503–507. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3172944.3173007 2
 - [45] J. Mitchell and R. L. Ensign. Student-loan defaults mount again. <https://www.wsj.com/articles/SB10000872396390444712904578024412786219802>, 2012. (accessed March 26, 2025). 10
 - [46] V. O. Mittal, J. D. Moore, G. Carenini, and S. Roth. Describing complex charts in natural language: A caption generation system. *Computational Linguistics*, 24(3):431–467, 1998. 2
 - [47] A. Nili, M. Tate, A. Barros, and D. Johnstone. An approach for selecting and using a method of inter-coder reliability in information management research. *International Journal of Information Management*, 54:102154, 2020. 4
 - [48] D. Nissenbaum. Delays imperil mining riches afghans need after pullout. <https://www.wsj.com/articles/SB100014240527023040919404577418363707647028>, 2012. (accessed March 26, 2025). 10
 - [49] 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. 2
 - [50] I. Peng, E. E. Firat, R. S. Laramée, A. Joshi, J. Bourdin, and E. Paquette. Evaluating bloom’s taxonomy-based learning modules for parallel coordinates literacy. In *Eurographics*, 2022. 2
 - [51] Posit team. *RStudio: Integrated Development Environment for R*. Posit Software, PBC, Boston, MA, 2025. 7
 - [52] M. D. Rahman, G. J. Quadri, B. Doppalapudi, D. A. Szafir, and P. Rosen. A qualitative analysis of common practices in annotations: A taxonomy and design space. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 1, 3, 5
 - [53] D. Ren, M. Brehmer, B. Lee, T. Höllerer, and E. K. Choe. Chartaccent: Annotation for data-driven storytelling. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 230–239, 2017. doi: 10.1109/PACIFICVIS.2017.8031599 8
 - [54] W. Revelle. psych: Procedures for psychological, psychometric, and personality research (r package version 2.412)[computer software], 2024. 7
 - [55] 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, 2019. doi: 10.1109/TVCG.2018.2864903 3, 7
 - [56] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, 2010. 8
 - [57] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*, pp. 364–371. Elsevier, 2003. 2
 - [58] S. Srinivasan, J. Purich, M. Correll, L. Battle, V. Setlur, and A. Crisan. From dashboard zoo to census: A case study with tableau public. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–15, 2024. doi: 10.1109/TVCG.2024.3490259 3, 7
 - [59] C. Stokes, C. X. Bearfield, and M. A. Hearst. The role of text in visualizations: How annotations shape perceptions of bias and influence predictions. *under review*, 2023. 1, 2, 3
 - [60] C. Stokes and M. A. Hearst. Why more text is (often) better: Themes from reader preferences for integration of charts and text. In *NL VIZ: Workshop on Exploring Opportunities and Challenges for Natural Language Techniques to Support Visual Analysis*. IEEE, 2022. 1
 - [61] C. Stokes, C. Sanker, B. Cogley, and V. Setlur. From Delays to Densities: Exploring Data Uncertainty through Speech, Text, and Visualization. In *Computer Graphics Forum*, vol. 43, p. e15100. Wiley Online Library, 2024. doi: 10.1111/cgf.15100 1
 - [62] 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. 1, 2, 3
 - [63] N. Sultanum and V. Setlur. From instruction to insight: Exploring the functional and semantic roles of text in interactive dashboards. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2, 3, 4
 - [64] N. Sultanum and A. Srinivasan. Datatales: Investigating the use of large language models for authoring data-driven articles. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
 - [65] Tableau. Tableau Summary Card. Tableau Software help pages, 2022. https://help.tableau.com/current/pro/desktop/en-us/insp ectdata_summary.htm. 2
 - [66] B. Thompson. Exploratory and confirmatory factor analysis: Understanding concepts and applications. *Washington, DC*, 10694(000):3, 2004. 7
 - [67] Y. Tian, W. Cui, D. Deng, X. Yi, Y. Yang, H. Zhang, and Y. Wu. Chartgpt: Leveraging llms to generate charts from abstract natural language. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–15, 2024. doi: 10.1109/TVCG.2024.3368621 9
 - [68] M. Tory and T. Moller. Rethinking visualization: A high-level taxonomy. In *IEEE Symposium on Information Visualization*, pp. 151–158, 2004. doi: 10.1109/INFVIS.2004.59 2
 - [69] E. R. Tufte. *The visual display of quantitative information*. Graphics Press, 1983. 1
 - [70] S. Velasco. The search for dark matter is dramatically expanding. <https://www.quantamagazine.org/physicists-are-expanding-the-search-for-dark-matter-20201123/>, 2020. (accessed March 26, 2025). 10
 - [71] O. Wallach. Visualizing the biggest tech mergers and acquisitions of 2020. <https://www.visualcapitalist.com/visualizing-biggest-tech-mergers-and-acquisitions-of-2020/>, 2020. (accessed March 26, 2025). 10
 - [72] C. Ware. *Information visualization: perception for design*. Morgan Kaufmann, 2019. 1
 - [73] N. Yau. Power sources in each state. <https://flowingdata.com/2015/09/18/power-sources-in-each-state/>, 2015. (accessed March 26, 2025). 10
 - [74] N. Yau. Gallery of uncertainty visualization methods. <https://flowingdata.com/2019/09/03/gallery-of-uncertainty-visualization-methods/>, 2019. (accessed March 26, 2025). 10
 - [75] H. H. Zhu, S. S. Cheng, and E. Wu. How do captions affect visualization reading? In *Workshop on Visualization for Communication (VisComm)*. IEEE, 2022. 1, 2, 3