

DASH: A Bimodal Data Exploration Tool for Interactive Text and Visualizations

Dennis Bromley*

Vidya Setlur†

Tableau Research

DASH: Data Analysis with Semantic Hierarchies

Fluid bidirectional data exploration with text and charts

In the quest to find a suitable neighborhood for a family of four, several factors must be considered, including average number of bedrooms, bathrooms, house size, lot size, and average house price. For a family of four, neighborhoods like **98105**, **98112**, and **98117** emerge as top choices due to their higher average number of bedrooms (figure 1) and bathrooms, accommodating the needs of a larger family unit. However, **98112** not only offers a larger average house size but also features properties with sizable lots, providing ample outdoor space for family activities. However, considering the average house price is also crucial. **98178** presents a more affordable option (figure 2) with decent house size and lot sizes, making it an attractive choice for families on a budget. In contrast, **98101** and **98121** are less suitable due to their lower average number of bedrooms and bathrooms, along with smaller house sizes, which may not meet the needs of a four-person family. To succeed in the Seattle real estate market, families should consider not only the size and price of the home but also the neighborhood's access to quality schools, parks, and community services, which are vital for a thriving family life. Therefore, while **98105**, **98112**, and **98117** offer spacious and well-appointed homes, **98178** provides a balance of affordability and livability, making it a compelling option for families seeking value.

The average house and lot sizes across different zip codes are crucial indicators for families seeking a new home, especially for a family of four. Larger average house sizes suggest more spacious living conditions, which can accommodate the needs of a growing family. Similarly, larger lot sizes not only offer more outdoor space for recreational activities but also indicate a potential for future expansions or additions to the property.

The data points for zip code 98112, with an average of 3.54 bedrooms and 2.77 bathrooms per home, indicate it as a highly suitable area for families seeking spacious living accommodations.

Tell Me More

The average house and lot sizes across different zip codes are crucial indicators for families seeking a new home, especially for a family of four. Larger average house sizes suggest more spacious living conditions, which can accommodate the needs of a growing family. Similarly, larger lot sizes not only offer more outdoor space for recreational activities but also indicate a potential for future expansions or additions to the property.

The data points for zip code 98112, with an average of 3.54 bedrooms and 2.77 bathrooms per home, indicate it as a highly suitable area for families seeking spacious living accommodations.

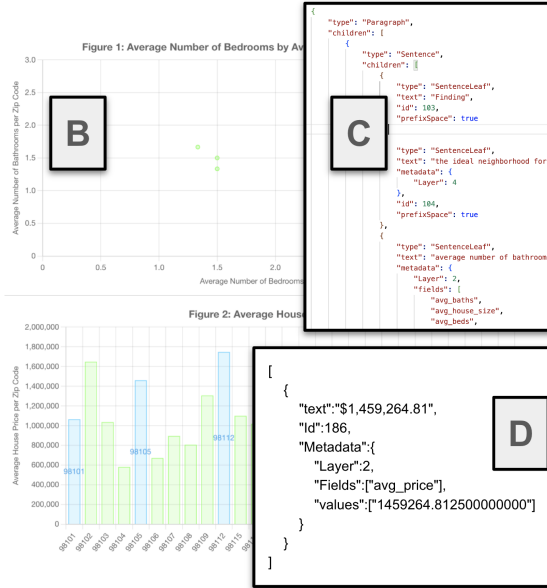


Figure 1: DASH tool's interface for supporting textual and visual data analysis. (A, B) Text and charts encoded with DASH semantic metadata. (C) A JSON representation of semantic metadata, including the semantic level, the data field, and the data value. (D) The DASH data-exchange JSON packet comprises the interactive text, its metadata, and identifiers that link the textual narrative to specific data points. (E) DASH semantic level assignment using the Lundgard *et al.* [20] color encoding.

ABSTRACT

Integrating textual content, such as titles, annotations, and captions, with visualizations facilitates comprehension and takeaways during data exploration. Yet current tools often lack mechanisms for integrating meaningful long-form prose with visual data. This paper introduces DASH, a bimodal data exploration tool that supports integrating semantic levels into the interactive process of visualization and text-based analysis. DASH operationalizes a modified version of Lundgard *et al.*'s semantic hierarchy model that categorizes data descriptions into four levels ranging from basic encodings to high-level insights. By leveraging this structured semantic level framework and a large language model's text generation capabilities, DASH enables the creation of data-driven narratives via drag-and-drop user interaction. Through a preliminary user evaluation, we discuss the utility of DASH's text and chart integration capabilities when participants perform data exploration with the tool.

Index Terms: Semantic levels, LLMs, text generation.

*dbromley@tableau.com

†vsetlur@tableau.com

1 INTRODUCTION

The interplay between text and visualizations is an important aspect of enhancing comprehension during data exploration [24, 17]. Research underscores the critical role of text in the reader's interpretation of visualizations, where it serves to summarize statistical attributes, explain insights, and offer contextual insights [14]. Effective textual descriptions not only reinforce the visual elements of a chart but also improve reader engagement and trust [28, 16]. However, existing tools often fall short of providing robust support for authoring text alongside visualizations, typically offering limited solutions for titles and chart/text alignment [19, 13]. To explore new ways of supporting interactive data exploration using *both* text and visualizations, we introduce DASH (Data Analysis using Semantic Hierarchies), a tool that leverages the integration of semantic levels in text for data analysis [20]. The tool provides large language model (LLM)-generated textual content through direct interactions with the visualization and vice-versa, supporting a *bimodal* composition process that is reflective of the underlying data semantics. In Figure 1, (A) shows LLM-generated interactive text that contextualizes data points with semantic metadata, while (B) displays the corresponding visual data through charts. (C) outlines the JSON metadata representation that describes the semantic levels and data fields, (D) illustrates a JSON packet that facilitates bimodal/bidirectional interactive data exploration and manipulation of the narrative in real-time, and (E) shows the semantic level assignment using the Lundgard and Satyanarayan [20] color encoding. A preliminary user evaluation of DASH indicates utility in in-

tegrating automated and manual authoring capabilities to enhance the overall data exploration process.

2 RELATED WORK

Our work is informed by research areas that explore the interplay of text and visualizations, organized into the following themes:

2.1 Typologies and taxonomies for analytical tasks

In the research space of visual analytics and data visualization, efforts have been made to understand and map users' data-exploration interactions to their underlying motivations and analytical goals. Brehmer and Munzner's typology of visualization tasks provides a foundational framework for this approach, offering a way to classify and interpret user actions within data visualizations [8].

Complementing this work, Amar and Stasko developed a taxonomy categorizing the different types of questions users ask when engaging with visualizations, providing a user-centered perspective on the design and evaluation of visualization tools [6]. Shneiderman offers a practical sequence for interactive visualization design, emphasizing the user's path from an overview to specifics [25]. Thomas and Cook expand upon these ideas, proposing a broad framework for the visual analytics process that aligns with analytical reasoning and supports the development of interactive systems [32]. Further refining these taxonomies, Tory and Moller address the task-based design of visualization tools, promoting an understanding of users' tasks as a central design principle [33]. Gotz and Zhou examine how visual analytics tools support users in forming insights and making decisions, contributing to the understanding of how people derive value from visual data exploration [10].

Recent work has sought to apply these theoretical constructs more directly to user interaction data. Lundgard and Satyanarayan's taxonomy [20] further refine this approach by breaking down the semantic content conveyed by visualizations into a hierarchy of comprehensible levels. This model starts at the elemental and encoded properties of the data and extends to higher-order cognitive and contextual insights, providing a structured methodology to assess and design visualization and textual description interplay. Taking inspiration from the aforementioned taxonomies, DASH operationalizes the semantic hierarchy proposed by Lundgard and Satyanarayan, encoding textual and visual content with metadata that represents different levels of data analysis, from basic data representation to integrated domain knowledge. Users can then engage in interactive data exploration through both the generated text and visualizations.

2.2 The role of text in visual analytics

There is a growing consensus within the research community that text should be treated as co-equal to visualization and can significantly enhance data comprehension. Kim et al. [14] conducted a study on how readers interpret line charts with accompanying captions, finding that coherence between textual and visual elements enables readers to glean insights more effectively. This coherence is particularly useful when both text and visuals emphasize the same key features, enhancing the overall understanding of the data presented. Research by Stokes et al. [28] highlights that users show a preference for visualizations that are heavily annotated as opposed to those with sparse or no text, suggesting that detailed annotations can reduce cognitive load and clarify complex datasets. This premise is further elaborated by Quadri et al. [23] and Fan et al. [9], who investigate how textual details affect comprehension and the perception of thematic maps, revealing that text can significantly influence how users interpret spatial and statistical data.

Ottley et al. [22] explore how text annotations can affect perceptions of bias and credibility of the presented data. Their findings emphasize that text not only serves to explain and summarize data but also shapes trust and reliability attributed to the visualizations by the viewers. Stokes et al. [27] further contribute to this topic by

examining how different types of semantic content impact reader takeaways, underscoring the nuanced role that text plays in data visualization environments. Our work continues this line of research by exploring how interactive and dynamically generated text can complement the visual analysis process. DASH allows for the real-time adaptation of both text and visualizations as users interactively navigate across various semantic levels of text description.

2.3 Visualization and text tools

Research focusing on tools that support the integration of visualization and text has explored techniques for the creation, interpretation, and enhancement of visualizations with textual elements. He et al. [11] surveyed the utilization of LLMs for crafting narrative visualizations, underscoring the potential of AI in bolstering the narrative aspects of data visualization. Other tools, such as AutoTitle [19], an interactive title generator, and Vistext [31], a benchmarking tool for semantically rich chart captioning, demonstrate novel capabilities in automatic text generation.

VizFlow [30] has shown the effectiveness of facilitating interaction between authors and readers by dynamically connecting text segments to corresponding visual elements, enriching the storytelling experience. Similarly, systems such as Kori [18] provide an interactive environment for synthesizing text and charts within data documents, emphasizing seamless integration for enhanced communication. Tools such as EmphasisChecker [13], DataDive [15], and FigurAlly [26] focus on specific aspects such as guiding chart and caption creation, aiding readers in contextualizing statistical statements and supporting the crafting of accessible scientific text, respectively. Recent advances have also been made with systems like SciCarpenter [12], which supports the composition of scientific figure captions using AI-generated content, reflecting a broader trend toward integrating generative AI into the visualization process. Basole and Major [7] discuss how generative AI tools offer creative assistance and automation throughout the visualization workflow, marking a shift towards a "human-led, AI-assisted" paradigm. This premise includes everything from user and design requirements to data preparation and insight augmentation.

DASH contributes to this research by adopting a mixed-initiative approach that leverages an LLM to enhance interactivity and semantic coherence between text and visualizations. The tool employs a semantic framework, allowing for interaction and bidirectional manipulation of both text and visual elements in real time.

3 DASH TOOL

3.1 Motivation

Text and visual modalities each excel at different aspects of data analysis; visual charts compress large amounts of data into data-rich images, while news agencies, blogs, and even this article are *text-first* documents that emphasize higher level data-analysis concepts such as inter-data relationships, conceptual discussions, and speculative narratives. This chart/text split between low-level and high-level semantic analysis inspired the design of the DASH tool.

3.2 Applying semantic levels for data exploration

Lundgard and Satyanarayan [20] devised a four-level data analysis model where analysis levels increase in knowledge integration and semantic granularity from basic chart information such as axis-ranges and encodings to domain-specific insights and broader multi-domain narratives. These four semantic levels provided additional inspiration around *when* to use *which* modality and how to craft a compelling bimodal data exploration narrative. Therefore, we propose a similar four-level (Semantic Levels 1-4, hereafter referred to as *L1-L4*) data-analysis semantic hierarchy specifically designed to be operationalizable in real-world tools (Table 1). The DASH tool utilizes this semantic hierarchy to help create a fluid

data analysis experience where text and charts are both first-class concepts, each leveraging their own strengths.

Level	Role	Modality	Example
4	Insights and Integration of Domain Knowledge	Text	“The steady increase in CEO pay may be playing a role in the recent uptick in investor activism.”
3	Relationships among data & statistics	Text, Chart	Correlations, clustering, outliers, and trends.
2	Statistics	Text, Chart	STDEV(profit)
1	Base Data	Text, Chart	Rows and columns of a SQL database

Table 1: Semantic levels for data analysis, with preferred modalities and examples at each level. Higher levels refer to higher-level semantic abstraction and knowledge integration. The *Text* on Level 1 modality indicates that text is often used to present individual data values but is typically not used for data presentation.

3.3 Implementation

The DASH tool is implemented using Typescript [5] and React [3] and consists of four primary pieces: a DASH metadata definition comprising information on the data field, specific data values, semantic level, and the user-facing text (Figure 1D), a generative LLM (OpenAI API model *gpt-4-turbo-preview*), Slate.JS [4], a text presentation library that maintains text-associated DASH metadata and renders the text appropriately, and the chart library Chart.JS [1].

The DASH tool is initialized with 1) a dataset (our prototype used Seattle real estate data [2]) 2) a natural language (NL) description of the dataset, 3) an NL description of the analytical end goal (in our case, to find a home for a family of four), and 4) a description of the DASH metadata format and instructions on how to assign the different semantic levels to the LLM’s response text. The LLM produces the final narrative text formatted as a tree-like hierarchy where each node in the tree contains metadata concerning semantic level (‘Layer’ in the Figure 1D JSON code) and associated data fields and values (Figure 1C). This tree structure maintains the text organization (*Paragraph* → *Sentence* → *SentenceLeaf*) and is necessary for consumption by the text rendering component.

The text rendering component provides rendering callbacks to format the text based on the text’s metadata. LLM response is shown in Figure 1E formatted with color-coded semantic levels (*L1* (pink), *L2* (green), *L3* (yellow), and *L4* (blue).) When text is dragged from either the text or chart components of the DASH interface, we recover the text’s underlying metadata from the component. While text is assigned a semantic level by the LLM, chart data in this prototype is assigned to L2 because chart data is semantically higher than L1 but lower than L4. More expressive charts with analytical overlays, such as clustering and correlation, could comprise L3 data. From these data, whether sourced from text or chart, we construct the drag-and-drop JSON object shown in Figure 1D and store the object in browser memory. The object can then be dropped anywhere within DASH, independent of where it came from, a crucial aspect of DASH’s bidirectional data flow.

Tell Me More, *Show Me More*, and charts (Figure 2) are all capable of receiving the JSON object above via mouse drag-and-drop. *Show Me More* creates a new chart with the metadata from the dropped JSON object; this metadata details the data fields and values to be charted. Inspired by tools like Polaris [29] and ShowMe [21], one field produces a bar chart; two fields produce a

scatter plot, specific data values produce reference lines, and specific zip codes highlight the indicated mark. If the JSON object is dropped directly onto an existing chart, the chart is updated in the same manner, possibly upgrading a bar chart to a scatter plot if a new field is added. If the drag-and-drop JSON object is dropped on *Tell Me More*, DASH re-queries the LLM for further discussion about the fields and values in the object. To produce a narrative style, DASH produces semantically complementary responses by offering high-level (*L3* & *L4*) analytical responses to low-level (*L1* & *L2*) data observations and low-level data-centric responses to high-level observations. As shown in Figure 2, this semantic-level-aware, source-agnostic data flow facilitates data exploration between DASH’s components, including text-to-text, text-to-chart, chart-to-text, and chart-to-chart interactivity. While the DASH interface employs drag-and-drop, the framework supports any affordance that allows the user to interactively direct data from one data-rendering modality to another.

3.4 Interface and Interactivity Behavior

Figure 2 shows an example of DASH interactivity with generated text paragraphs (Px) and interactions (Ix). DASH presents an initial paragraph (P1) discussing a Seattle real estate dataset through the lens of a real estate agent looking for a good neighborhood for a family of four. The user drags the words “house and lot sizes” to the *Tell Me More* button (I1). DASH’s LLM then generates a new high-level explanation of *why* those attributes are important for their analytical goal (P2). The user then drags the phrase “affordable option” (whose metadata contains the ‘avg_price’ field) to the *Show Me More* button (I2). DASH generates the chart “Average House Price” and a chart reference is inserted into the text. The user then drags the words “98101” and “98105, 98112, and 98117” into the “Average House Price” chart (I3 and I4), highlighting those zip codes. The text “bedrooms and bathrooms” is then dragged to *Show Me More* (I5), triggering DASH to create a new scatter plot “Average Number of Bedrooms by Average Number of Bathrooms”. Chart references are re-numbered by order of appearance in the text. Looking at the house price bar chart, the user wonders why zip code 98112 is so expensive, so they drag the “98112” bar from the bar chart to the scatter plot (I6), highlighting “98112” in the upper right of the scatter plot. Looking for a higher-level analysis of what this position means, the user drags the “98112” mark to *Tell Me More* (I7), whereupon the DASH LLM generates text (P3) explaining that that zip code has high-end but expensive homes. The end result of this interaction is a custom bimodal dashboard tailored for a specific analytical end goal, where high-level analysis and low-level data presentation work together.

4 PRELIMINARY EVALUATION

To explore the utility of DASH in supporting interactive data exploration, we conducted a preliminary user study involving 12 participants (*S1* - *S12*) with diverse levels of experience in data analysis (six novices, two intermediate, and four experts). The primary goals were to: 1) gather feedback on the integration of semantic levels in textual and visual data exploration and 2) assess DASH’s current design and identify areas for improvement.

Each session was conducted remotely, lasting about 30 minutes, with participants recruited via professional networks and social media platforms dedicated to data analytics communities. To mitigate any bias, participants were given only a brief introduction to the tool’s capabilities at the beginning of the session. They were then asked to engage with DASH using the Seattle real estate data [2], focusing on using the tool to generate data-driven narratives through the manipulation of text and visual elements. Participants performed the task of identifying a suitable neighborhood for a family of four by dragging and dropping data references across different semantic levels. This manipulation allowed participants to observe

DASH: Data Analysis with Semantic Hierarchies

Fluid bidirectional data exploration with text and charts

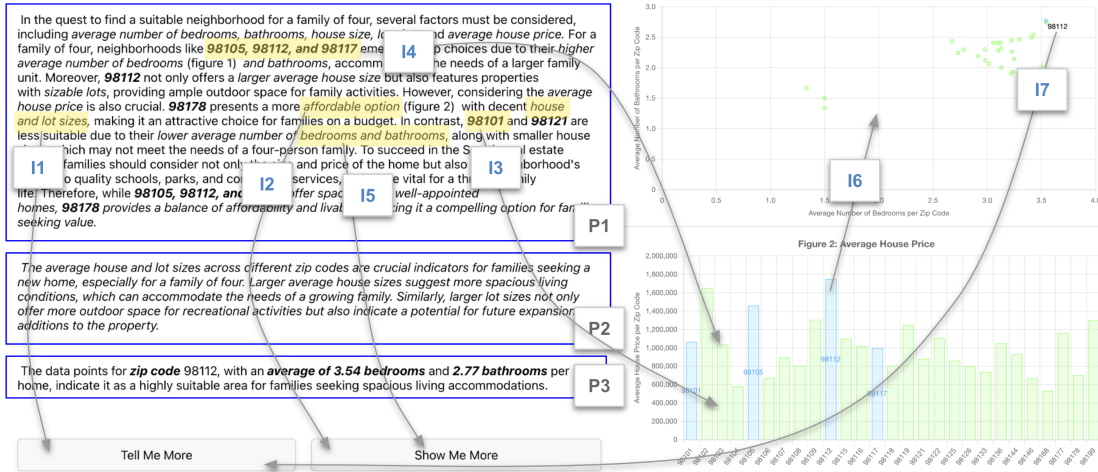


Figure 2: The DASH interface and an associated interactivity example. Generated text paragraphs are labeled as Px, and interactive mouse-drag gestures are labeled as Ix. The DASH interface comprises two primary sections: the text area on the left and the chart area on the right. The *Tell Me More* and *Show Me More* buttons at the bottom left trigger additional text and chart rendering, respectively. Text-to-text (I1), text-to-chart (I2-I5), chart-to-chart (I6), and chart-to-text (I7) interactions are illustrated.

how actions changed the narrative and visualization context, providing insights into how well DASH facilitates exploration toward the task goal. Feedback was collected through active observation and a semi-structured interview after the interactive session.

4.1 Results and Discussion

Participants found DASH to be a valuable tool for interactive data exploration in the context of the study task provided.

Flexibility of interface for data exploration. Participants generally found the interface of DASH to be intuitive. They noted the ease with which they could manipulate data by dragging and dropping references across different semantic levels. S1 remarked, “*This feels very fluid, almost like moving pieces on a chessboard, but each piece reveals more information about the neighborhood.*” The flexibility to dynamically alter the data narrative and visualization by adjusting semantic levels was appreciated by all the participants.

Use of semantic layers to adapt the narrative. DASH was noted for its effective use of semantic layers to facilitate deeper insights into the data. Participants were impressed with the system’s responsiveness and its ability to adapt the narrative and visualization based on their interactions. Challenges were noted when semantic layers were not immediately clear or when transitions between layers seemed abrupt, but overall, the system managed well. S5 noted, “*Seeing how my actions change the narrative and visual output in real time helps me understand the data at a much deeper level.*”

Relevance of data-driven text. The relevance of the data-driven narratives to the task of finding a suitable neighborhood was consistently highlighted. Participants valued the tool’s capability to sift through data and present nuanced insights, such as identifying areas with good schools and low crime rates. For example, S3 excitedly shared, “*I dragged the average income level to the chart, and it immediately showed me areas that fit my family’s lifestyle expectations.*” S11 appreciated how DASH could tailor the data presentation to reflect the task, stating, “*It’s impressive how it pulls together various data points to paint a clear picture of each neighborhood.*” Our evaluation also highlighted several areas for improvement:

Distinction across higher-level semantic levels. Participants pointed out some limitations in the granularity expressed through the semantic layers. The interface did not always capture subtle

distinctions within the data, which sometimes led to a generalized view rather than detailed insights. As S7 noted, “*When I tried to delve into more contextual data about local amenities, it felt like the system didn’t quite capture every detail I was interested in.*”

Need for fine-tuning of generated content. There was also feedback on the need for robust mechanisms to correct and refine data interactions. Participants expressed a desire for features in DASH that allowed them to modify their interactions or fine-tune the semantic layers post-initial manipulation to better refine their exploration. S6 suggested, “*It would be helpful if I could tweak or even backtrack on the semantic layers after adjusting them and then feed it back to the system to fine-tune the information displayed.*” Additionally, while the use of LLMs simplifies the implementation of DASH’s bimodal flow, they are not without challenges. For example, during this project the LLM would sometimes mislabel text or inaccurately reference data. Taken together with the performance and memory constraints of LLMs, where metadata labeling can take between 20–40 seconds, there are clearly areas for LLM improvement before they meet production tool standards.

5 FUTURE DIRECTIONS FOR BIMODAL DATA ANALYSIS

This preliminary work aims to operationalize the notion of semantic levels and demonstrate the feasibility of a fluid text/chart bidirectional analytical workflow. We identify several interesting directions for further enhancing the integration of text and charts in data exploration. One area for research is exploring techniques to enhance the customization of the semantic layers, providing users more flexibility to define and adjust these layers according to their domain-specific analysis needs. Future tools could incorporate machine learning models that not only interpret data but also predict trends and suggest new areas of exploration based on the user’s interaction history and data patterns. This proactive suggestion system could guide users toward hidden insights and enhance their understanding of the data. Finally, integrating such tools with real-time data is important for wider adoption. Dynamically updating visualizations and text based on live data changes could be beneficial in environments where data is continuously evolving, such as in financial markets or social media analytics.

REFERENCES

- [1] Chart.js. <http://chartjs.org>. Accessed on 2024-07-26. 3
- [2] Houses Price Prediction. <https://www.kaggle.com/datasets/ahmedessammukhtar/houses-price-prediction?select=train-data.csv>. Accessed on 2027-07-26. 3
- [3] React. <http://react.dev>. Accessed on 2024-07-26. 3
- [4] Slate.js. <http://slatejs.org>. Accessed on 2024-07-26. 3
- [5] Typescript language. <http://www.typescriptlang.org>. Accessed on 2024-07-26. 3
- [6] 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.2004.10 2
- [7] R. C. Basole and T. Major. Generative AI for Visualization: Opportunities and Challenges. *IEEE Computer Graphics and Applications*, 44(2):55–64, 2024. doi: 10.1109/MCG.2024.3362168 2
- [8] 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
- [9] A. Fan, F. Lei, M. Mancenido, A. M. Maceachren, and R. Maciejewski. Understanding Reader Takeaways in Thematic Maps Under Varying Text, Detail, and Spatial Autocorrelation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3613904.3642132 2
- [10] D. Gotz and M. X. Zhou. Characterizing Users' Visual Analytic Activity for Insight Provenance. *Information Visualization*, 8(1):42–55, Jan 2009. doi: 10.1057/ivs.2008.31 2
- [11] Y. He, S. Cao, Y. Shi, Q. Chen, K. Xu, and N. Cao. Leveraging Large Models for Crafting Narrative Visualization: A Survey. *arXiv preprint arXiv:2401.14010*, 2024. doi: 10.48550/arXiv.2401.14010 2
- [12] T.-Y. Hsu, C.-Y. Huang, S.-H. Huang, R. Rossi, S. Kim, T. Yu, C. L. Giles, and T.-H. K. Huang. SciCapenter: Supporting Caption Composition for Scientific Figures with Machine-Generated Captions and Ratings. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI EA '24. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3613905.3650738 2
- [13] 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*, 2023. doi: 10.1109/TVCG.2023.3327150 1, 2
- [14] 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. doi: 10.1145/3411764.3445443 1, 2
- [15] H. Kim, K. D. Le, G. Lim, D. H. Kim, Y. J. Hong, and J. Kim. DataDive: Supporting Readers' Contextualization of Statistical Statements with Data Exploration. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, IUI '24, p. 623–639. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3640543.3645155 2
- [16] 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*, CHI '19, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3290605.3300576 1
- [17] R. Kosara and J. Mackinlay. Storytelling: The Next Step for Visualization. *Computer*, 46(5):44–50, 2013. doi: 10.1109/MC.2013.36 1
- [18] 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, Jan 2022. doi: 10.1109/TVCG.2021.3114802 2
- [19] C. Liu, Y. Guo, and X. Yuan. AutoTitle: An Interactive Title Generator for Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5276–5288, 2024. doi: 10.1109/TVCG.2023.3290241 1, 2
- [20] 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, Jan 2022. doi: 10.1109/TVCG.2021.3114770 1, 2
- [21] J. Mackinlay, P. Hanrahan, and C. Stolte. Show Me: Automatic Presentation for Visual Analysis. *IEEE Transactions on Visualization & Computer Graphics*, 13(06):1137–1144, Nov 2007. doi: 10.1109/TVCG.2007.70594 3
- [22] A. Ottley, A. Kaszowska, R. J. Crouser, and E. M. Peck. The Curious Case of Combining Text and Visualization. In J. Johansson, F. Sadlo, and G. E. Marai, eds., *EuroVis 2019 - Short Papers*. The Eurographics Association, 2019. doi: 10.2312/evs.20191181 2
- [23] G. J. Quadri, A. Z. Wang, Z. Wang, J. Adorno, P. Rosen, and D. A. Szafir. Do You See What I See? A Qualitative Study Eliciting High-Level Visualization Comprehension. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3613904.3642813 2
- [24] 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
- [25] B. Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*, pp. 336–343, 1996. doi: 10.1109/VL.1996.545307 2
- [26] N. Singh, L. L. Wang, and J. Bragg. FigurA11y: AI Assistance for Writing Scientific Alt Text. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, IUI '24, p. 886–906. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3640543.3645212 2
- [27] C. Stokes, C. X. Bearfield, and M. A. Hearst. The Role of Text in Visualizations: How Annotations Shape Perceptions of Bias and Influence Predictions. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–12, 2023. doi: 10.1109/TVCG.2023.3338451 2
- [28] 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, 2023. doi: 10.1109/TVCG.2022.3209383 1, 2
- [29] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A System for Query, Analysis, and Visualization of Multidimensional Relational Databases. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):52–65, 2002. doi: 10.1109/2945.981851 3
- [30] 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*, CHI '21. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3411764.3445354 2
- [31] B. J. Tang, A. Boggust, and A. Satyanarayan. VisText: A Benchmark for Semantically Rich Chart Captioning. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 2023. doi: 10.18653/v1/2023.acl-long.401 2
- [32] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE, 2005. 2
- [33] M. Tory and T. Moller. Human Factors in Visualization Research. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):72–84, 2004. doi: 10.1109/TVCG.2004.1260759 2