



DATAWEAVER: Authoring Data-Driven Narratives through the Integrated Composition of Visualization and Text

Y. Fu¹ , D. Bromley²  and V. Setlur² ¹Georgia Institute of Technology, Atlanta, GA, USA²Tableau Research, Seattle, WA and Palo Alto, CA, USA

Abstract

Data-driven storytelling has gained prominence in journalism and other data reporting fields. However, the process of creating these stories remains challenging, often requiring the integration of effective visualizations with compelling narratives to form a cohesive, interactive presentation. To help streamline this process, we present an integrated authoring framework and system, DATAWEAVER, that supports both visualization-to-text and text-to-visualization composition. DATAWEAVER enables users to create data narratives anchored to data facts derived from “call-out” interactions, i.e., user-initiated highlights of visualization elements that prompt relevant narrative content. In addition to this “vis-to-text” composition, DATAWEAVER also supports a “text-initiated” approach, generating relevant interactive visualizations from existing narratives. Key findings from an evaluation with 13 participants highlighted the utility and usability of DATAWEAVER and the effectiveness of its integrated authoring framework. The evaluation also revealed opportunities to enhance the framework by refining filtering mechanisms and visualization recommendations and better support authoring creativity by introducing advanced customization options.

CCS Concepts

• **Human-centered computing** → *Visualization systems and tools; Interactive systems and tools;*

1. Introduction

Data-driven storytelling, a blend of visualization, interactivity, and narrative techniques, serves as a powerful medium for effectively communicating complex data insights, engaging the public, and delivering clear and actionable information [RHDC18, LRIC15, FS23]. By transforming data into compelling and often interactive narrative visualizations [SH10, HD11], data journalists make complex insights both understandable and impactful.

Despite its potential, authoring data-driven stories poses an interdisciplinary challenge, requiring a combination of skills in data analysis, visualization, and narrative writing [RHDC18, LRIC15]. Prior work has introduced tools to streamline the process, supporting chart creation and annotation, identifying data facts, and storytelling. Yet, the diversity in authors’ skillsets and workflows, coupled with the complexity of the process [LRIC15], poses challenges to achieving an efficient and integrated authoring experience. First, composing data-driven narratives is often a tedious and error-prone process, requiring manual transcription of data into textual facts and, in turn, transforming these facts into coherent narratives [CX22, FGB*24]. Second, current authoring tools typically follow a unidirectional ‘data-first’ workflow, offering limited entry points for traditional authors who often begin with a narrative perspective and then seek data to refine and expand their story [FS23]. Third, the storytelling process involves context-switching between

data analysis, visualization creation, and narrative writing, making it challenging to achieve a cohesive and integrated experience.

Drawing from a formative study with domain experts and existing work, our work proposes a framework (Figure 1) for an integrated visualization-text composition for data-driven storytelling, implemented in a system, DATAWEAVER. The system supports *vis-to-text* composition by incorporating deictic referencing through chart interactions, enabling users to create charts, highlight visual elements, and anchor Large Language Model (LLM)-based narratives to selected data facts, ensuring the narratives are aligned with the authors’ intent. Complementing this approach, DATAWEAVER also enables *text-to-vis* composition by recommending and generating interactive visualizations based on selected text, helping authors jumpstart relevant data exploration. This dual approach addresses distinct but complementary needs; *vis-to-text* helps ground narratives in accurate, data-driven insights, reducing manual transcription and error, while *text-to-vis* supports narrative-first workflows, particularly valuable for authors with limited data expertise by providing tailored visualization recommendations driven by text. We evaluated DATAWEAVER’s utility and usability with 13 participants and conducted a week-long diary study with 5 of them. Participants provided valuable feedback for refining the system, particularly regarding the data facts filtering mechanism, chart recommendations, and customization options.

2. Related Work

Our work draws on research in data-driven storytelling, chart and dashboard authoring, and natural language interfaces, contributing to the broader goal of effectively communicating data insights.

2.1. Data-driven Storytelling and Narrative Visualization

Data-driven storytelling has gained prominence as a visualization and HCI research focus, with prior work addressing key aspects such as insight discovery [FS22, WSZ*20, SXS*21], streamlining the authoring process [SH14, CX22, SBL21, SS23b, CVTH21, LZK*22], diversifying presentation mediums [LKS13, ARL*16, BRCP17, CECX23, SZZW24, SKH*23], exploring the design space [YXL*22, LWS*22, LSW*21], and developing taxonomies and frameworks [LWQ24, CLA*20, ZE23, CTL*18]. With the recent rise in LLMs, a key focus has been on integrating AI technologies into storytelling tools, harnessing the complementary strengths of both humans and AI [LWQ24]. Our work aligns with this direction, placing particular emphasis on facilitating a more efficient, fluid, and accurate authoring experience.

Several tools focus on bridging the gap between complex data analysis and accessible storytelling. For example, Ellipsis [CRS*14] facilitates narrative visualization through user-specified transitions and annotations, while Idyll Studio [CVTH21] provides a structured environment for authoring interactive articles that integrate text and visualizations. These tools, however, largely focus on pre-defined workflows, limiting flexibility for iterative or exploratory storytelling. Kori [LZK*22] offers an interactive platform for synthesizing text and charts, enabling users to create coherent narratives by effectively combining these elements. Cross-Data [CX22] enhances this capability by leveraging connections between text and data, streamlining the authoring process of data documents. Similarly, VizFlow [SBL21] supports dynamical layouts by leveraging text-chart links. However, bidirectional authoring support in these tools is limited to basic linkages between existing text and charts, lacking mechanisms to dynamically generate context-specific insights by leveraging their inherent connections.

The introduction of block-based approaches, such as DataParticles [CECX23], adds another dimension to storytelling by enabling dynamic and animated data representation. The exploration of LLMs in tools, such as DataTales [SS23b] further pushes the boundaries of narrative generation, though it presents challenges with scalability and accuracy with respect to the generated summaries. Charagraph [MMCV23] facilitates real-time annotation and narrative generation by allowing users to create interactive charts within data-rich paragraphs, enhancing the narrative process. Meanwhile, ChartAccent [RBL*17] provides tools for annotating data visualizations, enabling users to highlight key insights and make data stories more comprehensible. Automation also plays a crucial role in advancing data-driven storytelling. Caliope [SXS*21] automates the generation of visual data stories from spreadsheets, translating raw data into narrative presentations. Notable [LYZ*23] streamlines the storytelling process within computational notebooks, allowing for the seamless integration of narratives with analytical workflows. Lastly, Erato [SCC*23] introduces a collaborative aspect by enabling cooperative editing of data sto-

ries through fact interpolation, facilitating a shared narrative construction process among multiple users. Although these tools advance the capabilities of data-driven storytelling, they often focus on the linear conversion of data into stories. Our work further enhances the storytelling process by employing a novel framework that integrates both vis-to-text and text-to-vis workflows.

2.2. Chart and Dashboard Authoring

In data visualization and dashboard authoring, several tools aim to enhance the expressiveness of visual presentations. Medley [PSS23] provides intent-based recommendations to support the composition of dashboards. Data Illustrator [LTW*18] and Data Animator [TLS21] augment traditional tools by integrating data binding functionalities, enabling users to directly link datasets to graphical elements for dynamic visualization design. Beyond traditional visualization tools, innovations in data content authoring further expand the scope of data storytelling. Epigraphics [ZHC24] introduces a message-driven approach to infographics authoring, focusing on the narrative elements that drive the creation of infographics and conveying data insights in a more expressive way. DataInk [XHRC*18] provides a creative platform for direct data-oriented drawing, allowing users to create visualizations through an intuitive drawing interface. By focusing on creativity and direct manipulation of data, DataInk enables users to explore and present data in innovative and personalized ways. These authoring tools often focus on either visualization or narrative as the starting point; however, DATAWEAVER enables dynamic transitions between visualization creation and narrative development, anchoring narratives to data facts and visualization recommendations.

2.3. Natural Language Interfaces

The development of natural language interfaces (NLIs) [SBT*16, DMN*17, GDA*15, SS18] has enabled users to engage with data through conversational interactions, simplifying complex visual analysis by enabling natural language queries to explore and interpret data. NL4DV [NSS21] maps natural language queries to JSON-based analytic specifications with data attributes, tasks, and Vega-Lite visualizations. Voder [SDES19] enhances data interpretation by making textual data facts interactive for exploring alternative visualizations. FlowSense [YS20] integrates natural language processing with visual data exploration, allowing users to construct and modify visualizations via conversational input within a dataflow system. BOLT [SS23a] enhances dashboard authoring by allowing users to create and modify dashboards using conversational interactions, minimizing the need for manual configurations. While existing tools for data-driven storytelling, chart authoring, and natural language interfaces offer various capabilities, they often lack a fully integrated approach that supports fluid transitions between data and narrative. Many focus on either visualization or narrative as the primary entry point, limiting the flexibility of authors in switching between different modes of creation. DATAWEAVER addresses these gaps by providing a bidirectional storytelling environment that supports both *data-first* and *text-first* authoring modes. The system provides a novel, flow-based, zoomable interface that enables navigation between visualizations and narratives, supporting the integration of data and text during the storytelling process.

3. Understanding the Challenges in Data-driven Storytelling

Prior research highlights the complexity of data storytelling, which often involves iterative data exploration, the composition of text and visuals, story arrangement, and final presentation [LRIC15]. Our work aims to help streamline this process by integrating visualization and text authoring into a cohesive workflow. To uncover key obstacles in achieving this integrated experience, we conducted a formative study involving a data story authoring exercise and post-study interviews with five professional visualization practitioners (P'_1 – P'_5) recruited via our organization's mailing list. All participants self-identified as experts with extensive experience in creating visualizations and data-driven stories. In the authoring exercise, we asked participants to explore data and uncover insights using multiple interactive visualizations, preferably their own creations; for those who did not have their own creations, we provided a curated list of visualizations sourced from Tableau Public [tab] and The New York Times [NYT] (included in supplementary material). They presented their findings by capturing screenshots, drafting textual narratives, and adding annotations to create a cohesive presentation. Throughout the process, we encouraged participants to think aloud and share their reasoning. Each session lasted 30–40 minutes, followed by a 10–15 minute semi-structured interview. The study allowed us to observe authoring processes, generate artifacts for analysis, and gather feedback on interaction methods, challenges, and perspectives on AI integration. Sessions were recorded and transcribed, and the generated stories were collected using Google Slides. From our observations, transcripts, and collected stories, we identified the following key findings:

F1. Pre-computed data facts can be helpful. Three participants expressed a desire for pre-calculated data facts describing statistics computed from relevant data points as they interacted with visualizations. P'_4 noted, “Having some words already available to be plugged in without writing the calculation, like the percentage change, would be fantastic.” P'_3 concurred, expressing a desire for the calculation of the “absolute value” of an increasing trend.

F2. Challenges in manual data transcription from visualizations. Four participants highlighted the challenges of manually transcribing numerical data from visualizations into text. P'_4 described this as “[The hardest part] was getting some of the numbers from the visual into the text...so I had to go back and forth.” This issue is exacerbated when values are not directly visible, requiring interactions like hovering to reveal them. For instance, in a chart showing score differences, P'_1 needed the exact value (e.g., 25-point lead) to write an accurate narrative.

F3. Concerns regarding direct input of raw data to AI. While three participants showed a strong willingness to use AI, particularly LLMs for data narrative composition, four participants voiced significant concerns about feeding raw datasets directly to AI. They cited concerns over accuracy of generated narratives and data security, emphasizing their reluctance to expose sensitive data to AI.

F4. Narrative-guided exploration as an authoring strategy. We observed that authors working with familiar datasets or topics often adopt a narrative-driven approach—rather than engaging in open-ended data analysis, they begin with a specific storyline in mind, actively seeking data points or visual facts to substantiate or refine their narrative. For instance, P'_1 , while exploring sports data,

already had a clear narrative direction based on their personal experience of watching a specific game; they intentionally identified data facts to enrich and help validate their pre-existing narrative.

Design Goals: Drawing from these key findings and prior research [LRIC15, CX22, FS23], we distill five design goals for devising an integrated authoring framework to help streamline both visualization and text composition in data-driven narratives:

- **DG1.** Provide accurate data insights/facts that align with user intentions.
- **DG2.** Streamline data transcription and narrative composition.
- **DG3.** Prevent the direct input of raw datasets into LLMs.
- **DG4.** Enable narrative-driven visual data exploration.
- **DG5.** Foster a seamless and cohesive authoring experience.

4. DATAWEAVER System

4.1. An Integrated Framework for Visual-Text Composition

To achieve our design goals, we propose an integrated framework (Figure 1) for data story authoring, which combines two compositional approaches: *vis-to-text* and *text-to-vis*.

Vis-to-Text Composition addresses design goals **DG1** and **DG2** by facilitating narrative generation anchored to authors' *intended* insights and *accurate* data facts. Specifically, we propose the following approach. First, we allow users to *specify visual elements of interest through chart interactions*. These interactions enable deictic referencing by highlighting elements to indicate that the elements form the focus of the discussion. Second, we rely on algorithms to *compute descriptive statistics and data facts for the highlighted data points* through an intermediate data fact layer, ensuring accurate and efficient computation by delegating tasks to algorithms instead of LLMs. The feature also enables users to select relevant data facts to enhance narrative relevance and anchors generated narratives to a defined set of data facts to improve accuracy. Finally, we *leverage LLMs to generate narratives anchored to the selected data facts*. By weaving pre-computed data facts into narrative text, we address challenges such as restricted context length, limited numerical precision, and difficulties with complex data operations, avoiding the need to feed raw data directly into LLMs, fulfilling **DG3**, and offering an additional layer of privacy control.

Text-to-Vis Composition addresses **DG4** by introducing text as an alternative starting point for data story creation. A common challenge in storytelling is the need to create visualizations to support further analysis or to expand existing narratives. By simplifying this process, we aim to streamline authoring and reduce the burden on users with limited data expertise. To achieve this, the generation of relevant visualizations from text requires a degree of inter-modal semantic mapping. However, the free-text nature of DATAWEAVER was not amenable to a strictly structured parsing approach. Given these design considerations, we leverage LLMs as a pragmatic solution, offering a reasonable balance between semantic interpretation and free-text comprehension. While not without limitations, we found that LLMs provide an effective means of suggesting relevant analyses and generating chart specifications under the constraints of unstructured textual input. This approach also complements *vis-to-text* composition, as the resulting interactive visualizations can serve as new starting points for further data exploration and narrative development. Together, these bidirectional workflows

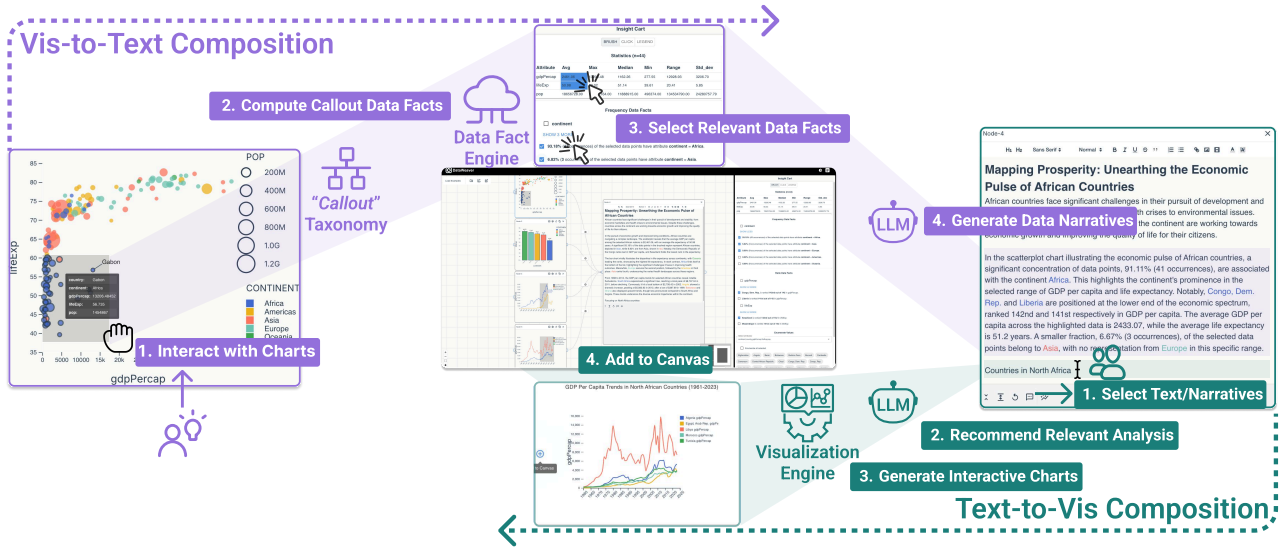


Figure 1: The integrated framework underlying DATAWEAVER. The interface is detailed in Figure 2. The purple-colored flow shows the Vis-to-Text composition (Figure 3). The teal-colored flow demonstrates Text-to-Vis composition (Figure 4)

create a ‘boomerang-like’ loop, enabling a seamless and iterative authoring experience that addresses DG5.

4.2. Callout Taxonomy

Given the extensive range of statistics and data facts that can be derived from a dataset [SXS*21], a key challenge for our proposed framework lies in striking a balance between providing users with granular control to define ‘what is salient’, while minimizing their cognitive effort in identifying and selecting relevant data facts.

Achieving this balance requires an effective mechanism for computing and organizing data facts—one that clarifies user intent and generates a focused list of insights based on interactions. User intent behind interactions has been studied across various contexts [YKSJ07, SH24], such as analysts’ selection behaviors in scatterplots [GGC*21] and annotations in grouped bar charts [RQSR24]. We extend this understanding by examining intent behind ‘calling out’ visual elements, i.e., ‘‘When users highlight data points in a chart, what insights or facts are they likely to incorporate into their narrative or discussion?’’ We developed a preliminary *Callout Taxonomy* (Table 1) as a theoretical guidance for more relevant data facts generation and their effective organization. Drawing from prior literature on visualization interaction [AES05, YKSJ07], data facts and insights [WSZ*20, SXS*21], and analytical intent [GGC*21, SH24], we posit that the intended data facts are associated with the chart type and interaction type that users engage with. In this context, we define *Callout Interactions* as user actions that highlight specific data points in a visualization, *Callout Intent* as the motivation behind such highlights, and *Callout Data Facts* as the insights derived from these interactions. We selected a set of exemplary chart types, both standard and advanced, and curated a list of *callout interactions* for each, drawing from popular visualization tools/libraries, and prior research [HAW08]. Two authors collaboratively developed the cor-

responding *callout data facts*: one created the initial coding, while the other independently generated a separate set. The lists were iteratively refined into a final taxonomy, extendable to additional chart types, interactions, or even compounded and chained interactions.

Chart	Interaction	Callout Data Facts
Scatterplot	Area Selection (2-D Brush)	Summary statistics*, Frequency*, Group vs. Global (stats), Rank*, Values*
	Discrete Selection (Click)	Values*, Outliers*, Rank*, Summary statistics*, Group vs. Global, Frequency*
	Group Selection (Legend Click)	Summary statistics*, Frequency*, Rank*, Group vs. Global (one group selected), Group vs. Group (multiple groups), Outliers*, Values*
	Add Trendline	Correlation between variables, Trendline
Bar Chart	Discrete Selection (Click)	Values*, Rank/Extreme*, Difference*, Summary statistics*
	Category Selection (Legend Click)	Summary statistics*, Frequency*, Rank*, Group vs. Global Difference, Group vs. Group Difference, Outliers*, Values*
	Area Selection (1-D Brush)	Summary statistics*, Frequency*, Rank*, Group vs. Global (one group), Values*
	Line Chart	Trend*, Start/end time*, Extreme points*, Range*, Difference*
Stacked Bar Chart	Timeframe Selection (Brush)	Trend*, Compare selected lines*, Compare selected to others
	Line Selection (Legend Click)	Value/Rank for individuals, Compare selected date, Analyze position relative to trend
	Temporal Point Selection (Click)	Comparison/relation/joint contribution, Joint relationship to others, Comparison to each other, Compare individuals to group
	Subcategories Selection (Legend Click)	Proportions, Compare individual proportions
Donut/Pie Chart	Segments Selection	Proportional values, Compare selected proportions, Compare sum proportion to total
	Discrete Selection (Click)	Proportional values, Compare relative proportions, Compare/aggregate hierarchical proportions
Sunburst	Discrete Selection (Click)	Chained proportion
	Chained Selection (Double-click)	

Table 1: *Callout Taxonomy* categorizes chart types, their corresponding callout interaction types, and the associated callout data facts. A * denotes a calculation implemented in DATAWEAVER.

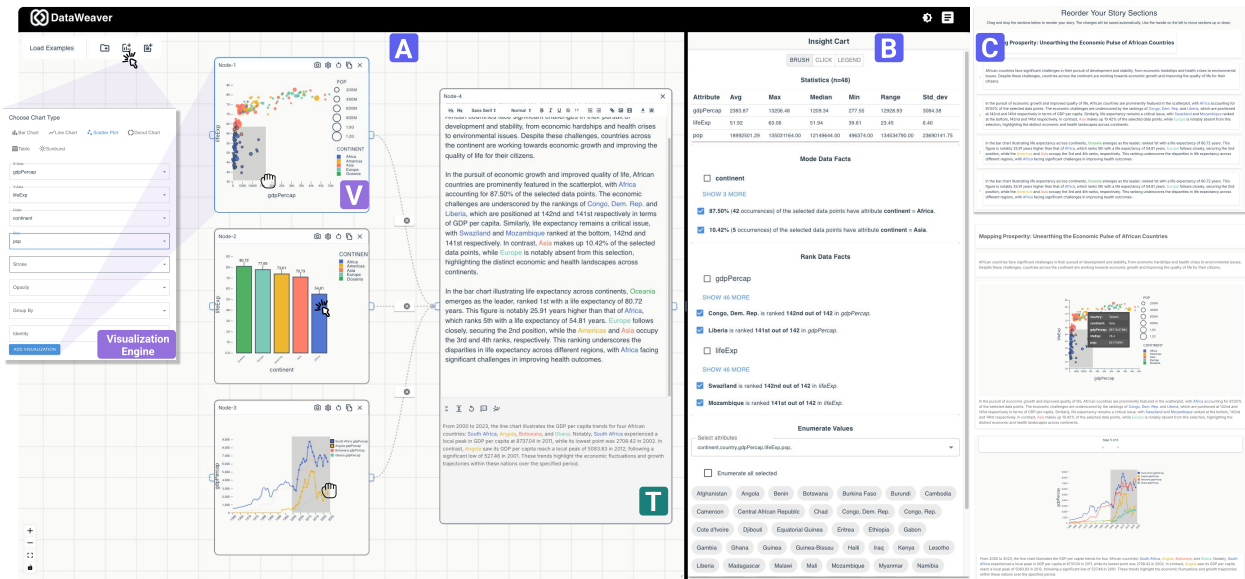


Figure 2: An overview of DATAWEAVER’s interface. The **Authoring Canvas** (A) is a flow-based, zoomable interface where users can add vis-nodes (V) using a visualization engine and text-nodes (T), as well as edges to connect them. The **Insight Cart** (B) is used for insight management while the **Review Page** (C) allows users to reorder the story and convert the story pieces into different presentation formats.

4.3. User Interface

Building on this taxonomy, we integrated its principles into DATAWEAVER to inform the design and functionality of the system’s interface (Figure 2) and core features. The frontend is built with React and uses D3.js for rendering interactive visualizations. DATAWEAVER’s backend, built with Python Flask, handles server-side logic and API requests, utilizing GPT-4o as its LLM along with other libraries for text generation, semantic parsing, and various computational tasks. The interface comprises three major components: a flow-based *Authoring Canvas* (Figure 2-A) for composing visualizations and text, an *Insight Cart* (Figure 2-B) for managing insights, and a *Review Page* (Figure 2-C) for refining the story and previewing the final visual data narrative in multiple formats.

Authoring Canvas (Figure 2-A): The authoring canvas allows users to add two types of nodes: visualization nodes (*vis-nodes*) and text editor nodes (*text-nodes*). Edges connect these nodes, defining their relationships and flows.

- **Nodes:** Users can add various types of interactive visualizations as *vis-nodes*. DATAWEAVER fundamentally supports interactive visualization creation by leveraging a visualization engine powered by D3.js. Users can upload and manipulate datasets (e.g., filtering), select the chart type, and specify the attributes that determine the chart’s visual encodings. The chart types supported are congruent with Table 1. The embedded visualizations are interactive, featuring tooltips and corresponding callout interactions. *Text-nodes* provide a straightforward, rich text editor for narrative composition, along with widgets for additional text manipulation. Users can also move, zoom, resize, duplicate, and remove nodes from the canvas.
- **Edges:** Edges help specify the flow of data facts generated by chart interactions. For example, if *vis-node-1* and *vis-node-2* both

have edges directed to *text-node-3*, the data facts provided by these two visualization nodes will collectively contribute to the text-node. Users can reuse the data facts from the same nodes to drive different text-nodes. This mechanism aims to assist users in managing story content and flow, especially when composing a data story that involves multiple sets of data facts and visualizations. Datasets bound to *vis-nodes* will also automatically serve as source data for visualization recommendations. Similarly, users can serve the datasets as data tables without specifying the chart.

Insight Cart (Figure 2-B): The insight cart, like a “shopping cart,” acts as a temporary repository to store facts for “checkout,” i.e., integration into a data story. Each node has its own insight cart.

- **Vis-nodes’ Insight Carts:** Insight carts associated with *vis-nodes* contain insights (e.g., descriptive statistics or data facts) derived from users’ callout interactions. These insights are displayed as clickable tables and grouped checkboxes, allowing users to make selections. The selected data facts automatically flow into the insight carts of downstream *text-nodes*.
- **Text-nodes’ Insight Carts:** Insight carts associated with *text-nodes* consist of two segments: a data-facts segment, which displays grouped data facts from upstream *vis-nodes*, and a visual insights segment, which suggests recommended visualizations.

Review Page (Figure 2-C): During or after composing the data story, authors can utilize the *review page* to adjust the order of the content and preview the final output in various forms. The *adjusting mode* automatically processes text from different nodes into a nested structure, enabling efficient restructuring of both the narrative and the linked visualizations. The *previewing mode* allows for a preview of the final product in common narrative visualization formats: *continuous page*, *scrollytelling*, and *stepper*.

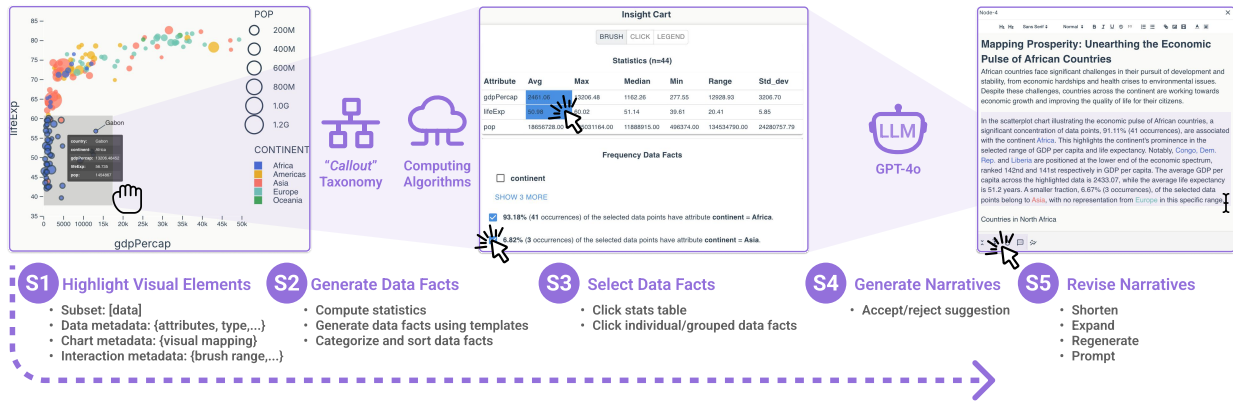


Figure 3: Demonstration of Vis-to-Text composition workflow. After users apply callout interaction to visualization (S1), DATAWEAVER computes the data facts and presents (S2) them in the insight cart. Users then select desired data facts (S3). An LLM then generates data narratives (S4) based on the selected data facts and metadata. Users can revise the generated narratives using the buttons (S5).

4.3.1. Vis-to-Text Composition Workflow

The vis-to-text composition starts with a created visualization and follows an *interact-compute-select-generate-revise* workflow with users, algorithms, and LLM collaboratively handle respective steps:

S1. User interacts with charts to highlight visual elements: In each visualization node, users initiate *callout interactions* to select visual elements, prompting the system to retrieve and dispatch a package with the data subset, metadata, and interaction details. This metadata is divided into three categories: *data metadata*, containing attribute names and value types; *chart metadata*, which captures visual encodings like color, size, and tooltip content (e.g., mapped variables or identity variables like a country’s name); and *interaction metadata*, which varies depending on the chart and interaction type. For instance, brushing a scatterplot captures a 2D spatial selection with coordinate ranges $[x_1, y_1]$ to $[x_2, y_2]$ and value ranges $[xValue_1, yValue_1]$ to $[xValue_2, yValue_2]$.

S2. Algorithms compute and organize data facts: Upon receiving the ‘callout package,’ DATAWEAVER’s backend, following Table 1, first computes the relevant statistics following the taxonomy and converts the computed values into a statistical table or template-based data facts, such as “95.24% of the selected data points have the attribute continent = Africa”, or “Gabon is an outlier in lifeExp.” These data facts are presented in the *vis-nodes’ insight cart*. To manage this potentially extensive list of data facts and reduce users’ cognitive load, DATAWEAVER categorizes them based on derived callout intents and applies sorting mechanisms to the list. The data facts are organized hierarchically in a nested structure: *Fact Types* \triangleright *Attributes* \triangleright *Data Facts*. The *Fact Type* layer ordering is guided by both our taxonomy and the goal of maintaining UI consistency. For example, the statistical table is always placed at the top to ensure uniformity across the interface, and the *Frequency* category is positioned above the *Rank* category when brushing a scatterplot. The *Attribute* layer does not adhere to a specific ordering scheme. However, to reduce the computational load from handling numerous attributes and ease users’ cognitive load when reviewing data facts, DATAWEAVER processes only the *attributes of interest*, i.e., variables that are visually mapped and explicitly se-

lected by the user. Within each list, *Data Facts* are sorted based on their ‘significance’ or ‘interestingness,’ a method frequently used by recent studies [SDES19, SXS*21]. DATAWEAVER utilizes bespoke sorting algorithms tailored to specific fact types, aiming to promote the most ‘significant’ data facts within each category. For example, when brushing over the lower quadrant of a scatterplot visualizing life expectancy against GDP per capita, DATAWEAVER surfaces data facts, such as “93.18% of them are African countries” but “none of them are European countries” in the *Frequency* \times *Continent* category. DATAWEAVER then sorts these *frequency* data facts by calculating weighted scores, taking into account normalized differences, ratios, and entropy, while prioritizing significant deviations and prominence within the subset. The detailed sorting mechanism and explanations are in the supplementary materials.

S3. User selects data facts: Users click on the statistical table cells or checkboxes to select individual or grouped data facts they want to include from each visualization node. The selected data facts are then streamed to the subsequent text-nodes, where they appear in the corresponding insight cart in a nested format.

S4. LLM generates data narratives: The narrative generation can be triggered by the user pressing the Tab key. The selected data facts are then fed to the LLM, along with the preceding context and metadata, organized into a prompt template. The prompt template encapsulates both a generic template and chart-specific templates. The generic template offers a structured approach to understanding the visualization and its context. The specification includes steps such as recognizing the visualization type, examining the article context, and synthesizing data facts into a coherent narrative. The type-specific templates provide tailored guidance and context. For example, brushing interactions provide the range of the brushed area versus the full axes range. A pseudo prompt is as follows:

- Understand the visualization: *chart type, chart metadata*
- Consider the context: “*preceding text content*”
- Consider the callout interaction: *interaction metadata*
- Focus on these data facts: [“*data fact 1*”...].
- Synthesize information and write a narrative based on the data facts.

Users can decide to accept or reject the generated narrative. Prompt templates are available in the supplementary materials.

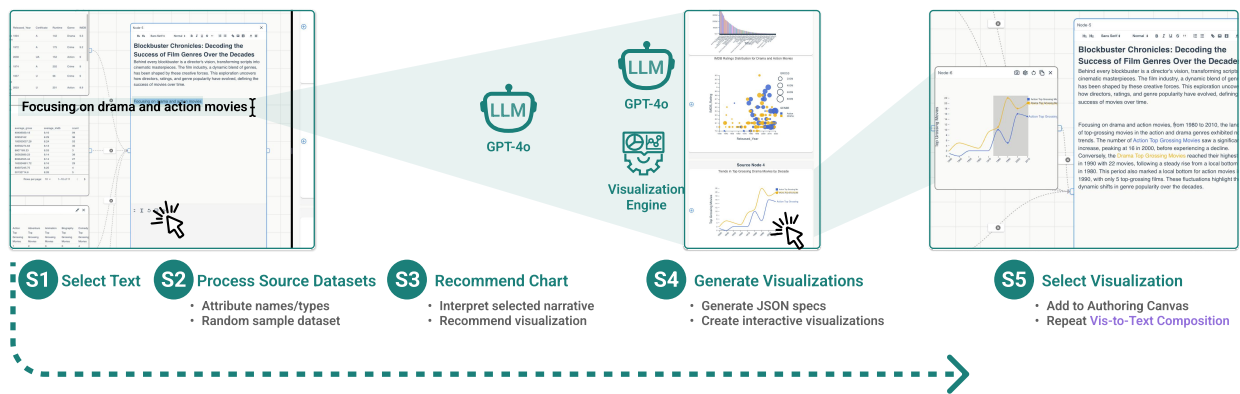


Figure 4: Demonstration of Text-to-Vis composition workflow. Users first type or select text to focus (S1), and DATAWEAVER retrieves and processes the datasets from upstream nodes (S2). An LLM then interprets the text and metadata and recommends relevant charts (S3). Based on the charts’ types, LLM then generates JSON specifications that contain both data operation and visualization schemas used to create interactive charts (S4). Users finally review the generated charts and add desired ones as new vis-nodes (S5).

S5. User and LLM revise the generated data narrative. After accepting the initial narrative, users can either manually revise it or use the LLM for further refinement. They can select any part of the generated text and prompt the LLM to revise it. DATAWEAVER offers three shortcut buttons (*shorten*, *expand*, *regenerate*) and a text instruction window for direct prompting. The pre-defined or user-specified instructions will be integrated into a new prompt alongside the generated content and the original prompt to ensure that the new content accurately aligns with the data facts.

4.3.2. Text-to-Vis Composition Workflow

DATAWEAVER utilizes LLMs to understand and infer contextually relevant information, enabling authors to expand their narratives using relevant information from the data. Additionally, DATAWEAVER leverages LLMs’ semantic parsing capabilities to provide a well-structured JSON specification that can be used for data operations and visualization generation [FGB*24].

S1. User selects or types text to focus. In a text node, users first select any portion of the existing text of interest, or they can type new sentences. For example, an author might focus on a specific region of Africa and type “Countries in North Africa...” They then select this segment and click the *Recommend Visualization* button.

S2. Algorithm processes the datasets. DATAWEAVER simultaneously retrieves all the underlying datasets from upstream nodes as reference data and dispatches them for further processing, e.g., extracting attributes, aiming to provide an overview of the available datasets without uploading the entire dataset.

S3. LLM interprets the narrative and recommends relevant charts for analysis. The selected text from S1 and attribute names from S2 are fed into the LLM. The prompt guides the LLM to first interpret the available datasets using processed information from S2 and then to recommend analyses and visualizations for the selected text. For example, providing the narrative, “Women’s participation in the Olympics has increased over time” alongside a dataset of athlete counts by gender across Olympic history prompts

the LLM to suggest analyses, such as the percentage of genders over the years, and a line chart to support the analysis.

S4. LLM generates JSON specifications for data operation and visualization generation. Based on different chart types, DATAWEAVER prompts the LLM to generate a JSON specification for the data operations (e.g., filtering, aggregation) and visualization creation. Under the hood, DATAWEAVER uses a JSON schema tailored for D3 to render different visualizations. For example, a scatterplot requires numerical `xAttr` and `yAttr` as mandatory inputs, while other attributes (e.g., `colorAttr`) are optional. The LLM leverages these schemas to generate corresponding JSON specifications, allowing DATAWEAVER to create visualizations with integrated callout interactions, which are then displayed in the visual insight section of the text-node’s insight cart.

S5. User selects generated visualizations. Eventually, users review the recommended charts and, by clicking an ‘add’ button, select those they wish to incorporate into the canvas as new visualization nodes. These new nodes are now available for interaction, data-fact generation, and narrative inclusion in the same manner as the existing visualization nodes. The new visualization nodes have completed the bidirectional loop, allowing users to generate new data facts and narratives that are relevant to the focused text.

To summarize, DATAWEAVER allows users to initiate the authoring process using either *Vis-to-Text* or *Text-to-Vis* workflows. In the former (Figure 3), users begin by creating interactive visualizations using DATAWEAVER’s foundational chart authoring feature, then compose data narratives by interacting with the charts and the automatically generated data facts (Figure 2-B). In the latter (Figure 4), users select the textual narratives they want to focus on, prompting DATAWEAVER to recommend and generate relevant interactive charts, thus completing the loop for data-driven narrative composition. The narrative can be further refined manually or through the system’s LLM-based revision feature. Finally, the visual-to-text connections automatically established enable users to rapidly adjust the sequence of content and review it in dynamic presentation

formats (Figure 2-C). A video showcasing DATAWEAVER’s interactivity and features is included in the supplementary materials.

5. User Evaluation

To evaluate the effectiveness and usability of DATAWEAVER and gather valuable feedback, we conducted a two-part evaluation: a focused user study, followed by a longitudinal week-long diary study.

5.1. Participants and Study Protocol

Participant Recruitment: We recruited 13 participants (P1–P13) to evaluate DATAWEAVER. Each received a \$20 Amazon gift card for a one-hour virtual session conducted via Google Meet. Participants (6 female, 7 male) self-reported high familiarity with data-driven stories (8 “highly familiar” and 8 “somewhat familiar”), diverse educational backgrounds (6 PhDs, 2 master’s, 4 bachelor’s, 1 other), and job roles (6 data analysts, 4 students, 1 postdoc, 1 data coach, 1 software manager). Five participants took part in a subsequent diary study, each receiving a \$60 Amazon gift card.

User Study: The study comprised four main phases. Participants began with an *Introduction* that included a verbal overview of the study’s background and objectives and a video walkthrough to ensure familiarity with DATAWEAVER before proceeding to the tasks. This was followed by the *Reproduction Walkthrough*, where participants explored DATAWEAVER’s core functionalities by recreating a visual data story similar to the one shown in the introductory video. To streamline this process, we provided three preconfigured visualization nodes and a text node with leading narratives, and participants were encouraged to ask questions as needed. In the *Open Authoring Exercise*, participants created their own visual data stories using an IMDb movie dataset [IMD24], with the option to start from four preconfigured charts or build from scratch using text linked to the datasets. We encouraged them to focus on core tasks, such as data fact selection, narrative generation, and chart recommendations, enabling us to evaluate DATAWEAVER’s integrated framework and both composition workflows. Throughout this task, participants were encouraged to think aloud, especially when making decisions about visual elements and data facts. Finally, in the *Questionnaire and Post-study Interview*, participants provided feedback on DATAWEAVER through a Likert-scale questionnaire that included nine utility questions and ten System Usability Scale (SUS) questions [Bro96]. Afterward, a follow-up interview was conducted to gather deeper insights and qualitative feedback on specific aspects of DATAWEAVER. Detailed materials, including the questionnaire, procedure, scripts, and results, are available in the supplementary materials.

Diary Study: In the diary study, participants used DATAWEAVER independently over a week, spending 20–30 minutes each day creating visual data stories. They took screenshots of their work, documented their experiences, challenges, and observations in daily diary entries, and provided feedback on DATAWEAVER’s features, usability, and workflows. At the end of the week, participants shared overarching reflections and recommendations for enhancing DATAWEAVER’s effectiveness and overall experience.

5.2. Results and Discussion

We analyzed responses using quantitative SUS scores and thematic coding of study transcripts. Two authors independently coded the transcripts using an open-coding approach to identify recurring themes and insights on the authoring experience, core functionalities, and areas for improvement. A third author examined the diary study documents, identified new themes, and synthesized findings from both studies to form the qualitative results.




Figure 5: The charts depict questionnaire response distributions on a 5-point Likert scale. The top chart presents results of nine utility questions, while the bottom two display usability questions. White lines and numerical values represent adjusted scores (out of 100), calculated using the SUS algorithm [Bro96].

DATAWEAVER received an average SUS score of 85.77, placing the tool in the “excellent” usability range, with a system utility score of 83.33 [Bro96, BKM09]. Figure 5 shows the distribution of responses to the system utility and usability questions. Qualitative feedback further supports this rating. DATAWEAVER’s ease of use, efficiency in data narrative composition, and integrated framework were repeatedly commended. P11 described DATAWEAVER as “something that I could definitely see myself using in my job.” P8 expressed excitement about DATAWEAVER’s potential to address a “major hurdle” and “open the doors to a lot of [non-data-experts] getting a lot closer to data.” P7 noted that generating text based on the selected data in the charts helped streamline the authoring process: “It’s so time-consuming to write insights manually, and the system does it for you, which is a huge time-saver” while P9 found the overall workflow “buttery smooth”.

Beyond feedback on the overall effectiveness of DATAWEAVER, the evaluation results revealed key insights into the tool’s framework and core features, along with areas for improvement:

The integrated framework is powerful, intuitive, and unique but can benefit from flexible interface layout options. Participants appreciated the intuitive integrated, bidirectional framework of DATAWEAVER, emphasizing that this functionality allowed for a more seamless transition between data and narrative composition,

as P7 remarked, “The ability to jump from text to visualization and back again without losing context was really powerful.” P9 added, “I loved how the system responded when I selected text—it immediately gave me a chart that matched the narrative, and that felt very smooth.” P8 described this feature as “super powerful” and “really unique.” Participants also noted that the flexibility to begin with either direction significantly enhanced their creative process. For instance, P4 mentioned, “I love that I can start with visualization and create text, or start with text and get a chart. It’s really versatile and fits the way I think.” This flexibility enables users to follow their workflow naturally, whether starting with data or using the narrative to guide analysis. Participants also emphasized DATAWEAVER’s ability to integrate the fragmented authoring process; P12 noted, “This is so much faster than what my current flow is,” referring to the frequent switching between different tools. P2 concurred, “I don’t need to switch tools for visualization and text—I can accomplish everything within a single tool.” While the integrated framework showed promise, participants expressed diverse preferences for interface layouts that aligned with their mental models. P7 expressed a strong preference for our flow-based canvas, stating, “I really love this kind of node and link connection way of seeing it... this is so fun.” P2 remarked that the combination of text and visualization nodes allowed her to “iterate and refine the relationship between texts and visualizations.” Meanwhile, three participants (P1, P9, P12) preferred a “side-by-side” or “left-to-right” layout to distinctly separate charts and text, preferably with the insight cart “in between these two columns” [P9]. P5 advocated for a “notebook-style” layout, allowing users to choose their preferred option.

Vis-to-text composition supports efficiency and reduces cognitive workload but demands a more advanced mechanism for filtering and selecting data facts. All participants appreciated the relevance and utility of the generated facts. P1 noted, “I didn’t come across any instance where the generated facts were inaccurate or out of context. They were spot on, which is crucial for narrative integrity.” P2 pointed out, “the data facts generated are very, very relevant to what I selected.” Several participants (P7, P9, P13) highlighted that the generated narratives anchored the storytelling process. P7 remarked, “The facts that were generated really guided me through the process, and I felt like they gave me a clear direction to follow in building my narrative.” However, the sheer volume of information was overwhelming for some participants (P2, P6, P12) —“There’s just so much going on... it’s hard to quickly find what’s important without getting lost in all the details. [P6]” To address this issue, participants suggested implementing filtering mechanisms that would allow them to prioritize or refine the data facts based on specific needs. P12 noted, “It would be helpful to have a filter to narrow down the facts based on criteria like time ranges, categories, or trends.” Another suggestion was the use of visual glyphs or other non-textual indicators to help visually scan and identify important trends or outliers without needing to read through all the text-based facts. Additionally, P2 mentioned, “It would be great if the data facts were grouped or tiered, so I could drill down from high-level summaries to more detailed information.” This suggestion underscores the need for a more advanced hierarchical organization of data facts, such as grouping by semantic level [LS22], complexity, or role. This capability would complement DATAWEAVER’s cur-

rent management mechanism that categorizes data facts by types and attributes and sorts them within lists based on significance.

Text-to-vis composition showed promise but also exhibited inconsistent performance due to data limitation. Participants acknowledged the potential of the *text-to-vis* workflow. P3 noted that “the recommended visualizations are useful in terms of helping me expand my narratives.” P5 regarded the chart generation feature as “the most impressive one.” P13 concurred, describing “writing that sentence and having it suggest visualization” as their “favorite feature.” Nevertheless, several participants (P3, P6, P11) encountered challenges in generating meaningful visualizations from text, indicating a limitation in the system’s ability to fully capture the user’s narrative intent. P6 explained, “It’s not always clear what kind of chart will be generated from the text... sometimes I expected a line graph but got something completely different.” In response to that disconnect, participants modified the narrative text to explicitly request a specific chart type or define the visual mapping. Moreover, P11 felt the recommended visualizations “were kind of similar to the charts that already existed.” Such inconsistent performance pointed to a current limitation: DATAWEAVER is restricted to performing basic data operations (e.g., filtering) on user-provided or existing datasets. However, the specific charts some participants expected (e.g., a line chart showing a movie director’s accumulated gross) often exceeded the scope of the provided datasets or required advanced data manipulation (e.g., data rollout).

Customization desired for generated visual and textual content. Besides feedback on the core authoring features, participants highlighted the need for greater customization of the final presentation, encompassing both visual and textual components. For visualizations, they emphasized the importance of flexibility in customizing appearance and formatting, such as color schemes, axis labels, and chart types. P10 stated, “...customization options would be a big thing for me.” Similarly, P11 remarked, “It would be really useful to have more options to tweak the charts that get generated.” Additionally, P9 attempted to convert a generated line chart into a bar chart by typing in the text editor, suggesting the potential for a feature to customize charts using natural language commands. For textual content, participants (P2, P4, P7, P12, P13) praised DATAWEAVER’s ability to quickly adjust generated narratives while maintaining their connection to data facts. Furthermore, P8 suggested incorporating more customizable annotations to better connect visual and textual components, such as overlaying the text on corresponding visuals.

6. Limitations and Future Implications

The user evaluation reveals key limitations, highlighting areas for improvement and broader applicability.

Further Expansion of the Callout Intent Taxonomy. We consider the callout intent taxonomy for visualization a preliminary framework. While the taxonomy covers common chart types and interactions, offering useful guidance for the development of DATAWEAVER, it is not exhaustive. Future work could expand the taxonomy along several dimensions. One apparent avenue is to incorporate a broader repertoire of chart types. We envision two contrasting approaches. For domain-specific visualizations, callout in-

teractions and data facts could be tailored by leveraging conventions and expertise. Alternatively, for dynamically generated visualizations, such as those created with Tableau's ShowMe [Sof24], a more adaptable and sophisticated approach would be necessary. Another direction could involve additional interaction modalities [LSIS21], e.g., sketching [LLY*24]. An advantage of alternative interaction modalities is that they do not interfere with interactions used for other analytical tasks. Beyond designing novel callout interactions, future work can explore multiple callout interactions to further disambiguate callout intent, including *compounded callouts* that combine multiple callouts in a single chart, and *chained callouts*, which link callouts across connected charts.

Lowering Data Barriers for Text-first Authors. The study participants underlined a significant limitation: supporting authors with minimal data expertise in creating data-driven articles. The 'text-driven' workflow, while reducing the expertise and effort required to create charts, still requires authors' familiarity with chart interaction and interpretation [FS23]. P6 noted that non-data-experts often "*scout over the facts*" during authoring and suggested providing these facts along with their "*correlation or causation*" to the narratives. To further simplify this process, we propose extracting data facts pertinent to a focused narrative. For instance, when an author writes a narrative such as "*Women have made significant strides in achieving equal representation in the Olympics*", automatically extracting relevant data facts, such as "*In 1900, women made up 2.2% of Olympic athletes*" and "*The Paris Olympics achieved 50% female participation*" would assist 'text-first' authors in composing data-driven stories. To mitigate potential confirmation bias during this process, we advocate for a more balanced mechanism that incorporates both 'supporting' and 'refuting' data evidence.

Supporting More Sophisticated Chart Generation. We envision two directions to further enhance DATAWEAVER's capability to generate charts for storytelling. The first involves developing a more advanced data extraction mechanism that can dynamically retrieve and integrate data [KP07] from multiple sources or databases. The second focuses on developing a more robust model capable of proficiently handling complex data manipulations to meet the requirements of the visualization engine. Beyond visualizations that aid data exploration, future research could also explore generating annotated charts [LLJ*20] to accompany textual data facts, providing 'visual data facts' to enrich the data story.

Opportunities and Risks in Expanding Narrative Generation. Another key limitation of DATAWEAVER is its primary focus on assembling together data facts while lacking support for higher-level storylines. This constraint stems from our deliberate prompt design, which restricts text composition to 'accurate facts' rather than allowing LLMs to interpret data, thereby leaving interpretation to human authors. However, our user study revealed that while data facts and descriptions serve as foundational building blocks, adding context and high-level takeaways could enrich storytelling by incorporating deeper semantic meaning and interpretation. Future iterations could relax these restrictions, enabling more layered narratives beyond data facts [LS22, BS24]. For instance, LLMs could be employed to provide context for specific data facts or generate interpretations based on a collection of data points. This approach allows authors to tailor data storytelling for diverse audi-

ences and perspectives, such as low-level data facts for a technical audience and high-level insights for C-level executives. However, relying on LLMs for more creative tasks increases the risk of misinterpretation and bias. Future work should systematically examine the trade-offs between leveraging LLMs for higher-order tasks and mitigating their associated risks.

Support for Advanced Analysis and Presentation. While DATAWEAVER concentrates on the dual composition of visualization and text, enhancing its framework with advanced data analysis and presentation capabilities is crucial for more insightful and compelling storytelling. We envision this integrated framework as a strong foundation for human-AI collaboration in data storytelling, offering flexible extensibility on both ends. For instance, our flow-based architecture could enable more advanced visual data exploration by connecting multiple visualizations and supporting interactions like brushing and linking, making DATAWEAVER compatible with data analysis frameworks like VisFlow [YS17]. With enhanced data processing capabilities, callout interactions could be employed to support *vis-to-vis* composition, i.e., interacting with one chart and recommending others for deeper exploration. Additionally, the auto-established vis-callout-text connections could be adapted to enable more presentation formats (e.g., video [SZZW24]) and potentially support more expressive chart reconfigurations, such as animated unit visualizations [CECX23].

Reflection on Integrating LLMs in Data-driven Storytelling. LLMs have the potential to bridge gaps in the multifaceted data storytelling process. Yet, balancing human agency and creativity with automation, efficiency, and accuracy remains both challenging and crucial. We posit that decomposing the workflow into smaller tasks and strategically delegating them is essential to leveraging the complementary strengths effectively. Implementing such a tool requires understanding the pain points and preferences of different user groups (e.g., journalists vs. data analysts), along with a more comprehensive empirical analysis of LLM performance and the risks associated with various task assignments.

7. Conclusion

Creating cohesive, interactive data-driven stories is complex, requiring the integration of visualizations and narrative components. DATAWEAVER addresses these challenges by introducing a framework that supports and integrates both visualization-to-text and text-to-visualization compositions. The tool enables users to ground their narratives in data facts derived from call-out interactions, ensuring stronger alignment between the visualizations and the resulting narrative. By offering both data-first and text-first workflows, the tool provides flexibility for different authoring preferences and supports various presentation formats, such as scrollytelling. A user evaluation suggests that DATAWEAVER streamlines the data storytelling process while offering opportunities for further refinement. The study also highlights a key tension in AI-enabled data narratives: balancing the efficiency of automation with the need for customization. Leveraging AI and computational algorithms to enhance efficiency and reduce workload, while preserving the user's ability to control the narrative structure, tone, and depth of insights could foster more meaningful data-driven stories that faithfully reflect both the data and the author's intent.

References

- [AES05] AMAR R., EAGAN J., STASKO J.: Low-level Components of Analytic Activity in Information Visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005*. (2005), pp. 111–117. doi:10.1109/INFVIS.2005.1532136. 4
- [ARL*16] AMINI F., RICHE N. H., LEE B., MONROY-HERNANDEZ A., IRANI P.: Authoring Data-driven Videos with Dataclips. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 501–510. doi:10.1109/TVCG.2016.2598647. 2
- [BKM09] BANGOR A., KORTUM P., MILLER J.: Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale. *J. Usability Studies* 4, 3 (May 2009), 114–123. 8
- [BRCP17] BACH B., RICHE N. H., CARPENDALE S., PFISTER H.: The Emerging Genre of Data Comics. *IEEE Computer Graphics and Applications* 37, 3 (2017), 6–13. doi:10.1109/MCG.2017.33. 2
- [Bro96] BROOKE J.: SUS: A Quick and Dirty Usability Scale. *Usability Evaluation in Industry* (1996). 8
- [BS24] BROMLEY D., SETLUR V.: DASH: A Bimodal Data Exploration Tool for Interactive Text and Visualizations. In *2024 IEEE Visualization and Visual Analytics (VIS)* (2024). doi:10.48550/arXiv.2408.01011. 10
- [CECX23] CAO Y., E J. L., CHEN Z., XIA H.: DataParticles: Block-Based and Language-Oriented Authoring of Animated Unit Visualizations. CHI '23. doi:10.1145/3544548.3581472. 2, 10
- [CLA*20] CHEN S., LI J., ANDRIENKO G., ANDRIENKO N., WANG Y., NGUYEN P. H., TURKAY C.: Supporting Story Synthesis: Bridging the Gap between Visual Analytics and Storytelling. *IEEE Transactions on Visualization and Computer Graphics* 26, 7 (7 2020), 2499–2516. doi:10.1109/TVCG.2018.2889054. 2
- [CRS*14] CARR H., RHEINGANS P., SCHUMANN H., SATYANARAYAN A., HEER J.: Eurographics Conference on Visualization (EuroVis) 2014 Authoring Narrative Visualizations with Ellipsis. doi:10.1111/cgf.12392. 2
- [CTL*18] CHEVALIER F., TORY M., LEE B., VAN WIJK J., SANTUCCI G., DÖRK M., HULLMAN J.: From Analysis to Communication: Supporting the Lifecycle of a Story. In *Data-Driven Storytelling*, Riche N. H., Hurter C., Diakopoulos N., Carpendale S., (Eds.). A K Peters/CRC Press, 2018, pp. 95–115. 2
- [CVTH21] CONLEN M., VO M., TAN A., HEER J.: Idyll Studio: A Structured Editor for Authoring Interactive & Data-Driven Articles. In *ACM Symposium on User Interface Software and Technology* (10 2021), ACM, pp. 1–12. doi:10.1145/3472749.3474731. 2
- [CX22] CHEN Z., XIA H.: CrossData: Leveraging Text-Data Connections for Authoring Data Documents. CHI '22. doi:10.1145/3491102.3517485. 1, 2, 3
- [DMN*17] DHAMDHARE K., MCCURLEY K. S., NAHMIA R., SUNDARARAJAN M., YAN Q.: Analyza: Exploring Data with Conversation. In *International Conference on Intelligent User Interfaces* (2017), IUI, pp. 493–504. doi:10.1145/3025171.3025227. 2
- [FGB*24] FU Y., GUO S., BURSZTYN V. S., HOFFSWELL J., ROSSI R., STASKO J.: “The Data Says Otherwise” – Towards Automated Fact-checking and Communication of Data Claims. In *ACM Symposium on User Interface Software and Technology* (2024). doi:10.1145/3654777.3676359. 1, 7
- [FS22] FU Y., STASKO J.: Supporting Data-Driven Basketball Journalism through Interactive Visualization. CHI '22, ACM. doi:10.1145/3491102.3502078. 2
- [FS23] FU Y., STASKO J.: More Than Data Stories: Broadening the Role of Visualization in Contemporary Journalism. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–20. doi:10.1109/TVCG.2023.3287585. 1, 3, 10
- [GDA*15] GAO T., DONTCHEVA M., ADAR E., LIU Z., KARAHALIOS K. G.: DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *ACM Symposium on User Interface Software Technology (UIST)* (2015), ACM, pp. 489–500. doi:10.1145/2807442.2807478. 2
- [GGC*21] GADHAVE K., GÖRTLER J., CUTLER Z., NOBRE C., DEUSSEN O., MEYER M., PHILLIPS J. M., LEX A.: Predicting Intent Behind Selections in Scatterplot Visualizations. *Information Visualization* 20, 4 (2021), 207–228. doi:10.1177/14738716211038604. 4
- [HAW08] HEER J., AGRAWALA M., WILLET W.: Generalized Selection via Interactive Query Relaxation. In *SIGCHI Conference on Human Factors in Computing Systems* (2008), CHI '08, ACM, p. 959–968. doi:10.1145/1357054.1357203. 4
- [HD11] HULLMAN J., DIAKOPOULOS N.: Visualization Rhetoric: Framing Effects in Narrative Visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2231–2240. doi:10.1109/TVCG.2011.255. 1
- [IMD24] IMDB: <https://developer.imdb.com/>, 2024. 8
- [KP07] KEZUNOVIC M., POPOVIC T.: Substation Data Integration for Automated Data Analysis Systems. In *2007 IEEE Power Engineering Society General Meeting* (2007), pp. 1–6. doi:10.1109/PES.2007.386177. 10
- [LKS13] LEE B., KAZI R. H., SMITH G.: SketchStory: Telling More Engaging Stories with Data through Freeform Sketching. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2416–2425. doi:10.1109/TVCG.2013.191. 2
- [LLJ*20] LAI C., LIN Z., JIANG R., HAN Y., LIU C., YUAN X.: Automatic Annotation Synchronizing with Textual Description for Visualization. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), CHI '20, ACM, p. 1–13. doi:10.1145/3313831.3376443. 10
- [LLY*24] LIN Y., LI H., YANG L., WU A., QU H.: InkSight: Leveraging Sketch Interaction for Documenting Chart Findings in Computational Notebooks. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 944–954. doi:10.1109/TVCG.2023.3327170. 10
- [LRIC15] LEE B., RICHE N. H., ISENBERG P., CARPENDALE S.: More Than Telling a Story: Transforming Data into Visually Shared Stories. *IEEE Comput. Graph. Appl.* 35, 5 (Sep 2015), 84–90. doi:10.1109/MCG.2015.99. 1, 3
- [LS22] LUNDGARD A., SATYANARAYAN A.: Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (Jan 2022), 1073–1083. doi:10.1109/TVCG.2021.3114770. 9, 10
- [LSIS21] LEE B., SRINIVASAN A., ISENBERG P., STASKO J.: Post-WIMP Interaction for Information Visualization. *Foundations and Trends® in Human-Computer Interaction* 14, 1 (2021), 1–95. doi:10.1561/11000000081. 10
- [LSW*21] LAN X., SHI Y., WU Y., JIAO X., CAO N.: Kineticharts: Augmenting Affective Expressiveness of Charts in Data Stories with Animation Design. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 933–943. doi:10.1109/TVCG.2021.3114775. 2
- [LTW*18] LIU Z., THOMPSON J., WILSON A., DONTCHEVA M., DELOREY J., GRIGG S., KERR B., STASKO J.: Data Illustrator: Augmenting Vector Design Tools with Lazy Data Binding for Expressive Visualization Authoring. In *Conference on Human Factors in Computing Systems - Proceedings* (4 2018), vol. 2018-April, ACM. doi:10.1145/3173574.3173697. 2
- [LWQ24] LI H., WANG Y., QU H.: Where Are We So Far? Understanding Data Storytelling Tools from the Perspective of Human-AI Collaboration. In *CHI Conference on Human Factors in Computing Systems* (5 2024), ACM, pp. 1–19. doi:10.1145/3613904.3642726. 2

- [LWS*22] LAN X., WU Y., SHI Y., CHEN Q., CAO N.: Negative Emotions, Positive Outcomes? Exploring the Communication of Negativity in Serious Data Stories. In *Conference on Human Factors in Computing Systems - Proceedings* (4 2022), ACM. doi:10.1145/3491102.3517530. 2
- [LYZ*23] LI H., YING L., ZHANG H., WU Y., QU H., WANG Y.: Notable: On-the-fly Assistant for Data Storytelling in Computational Notebooks. In *CHI Conference on Human Factors in Computing Systems* (2023), CHI '23, ACM. doi:10.1145/3544548.3580965. 2
- [LZK*22] LATIF S., ZHOU Z., KIM Y., BECK F., KIM N. W.: Kori: Interactive Synthesis of Text and Charts in Data Documents. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 184–194. doi:10.1109/TVCG.2021.3114802. 2
- [MNCV23] MASSON D., MALACRIA S., CASIEZ G., VOGEL D.: Charagraph: Interactive Generation of Charts for Realtime Annotation of Data-Rich Paragraphs. In *Conference on Human Factors in Computing Systems* (4 2023), ACM. doi:10.1145/3544548.3581091. 2
- [NSS21] NARECHANIA A., SRINIVASAN A., STASKO J.: NL4DV: A Toolkit for Generating Analytic Specifications for Data Visualization from Natural Language Queries. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 369–379. doi:10.1109/TVCG.2020.3030378. 2
- [NYT] New York Times. <https://nytimes.com>. 3
- [PSS23] PANDEY A., SRINIVASAN A., SETLUR V.: MEDLEY: Intent-based Recommendations to Support Dashboard Composition. vol. 29, pp. 1135–1145. doi:10.1109/TVCG.2022.3209421. 2
- [RBL*17] REN D., BREHMER M., LEE B., HÖLLERER T., CHOE E. K.: ChartAccent: Annotation for Data-Driven Storytelling. In *2017 IEEE Pacific Visualization Symposium (PacificVis)* (2017), pp. 230–239. doi:10.1109/PACIFICVIS.2017.8031599. 2
- [RHDC18] RICHE N. H., HURTER C., DIAKOPOULOS N., CARPENDALE S. (Eds.): *Data-Driven Storytelling*. AK Peters Visualization Series. CRC Press, 2018. 1
- [RQSR24] RAHMAN M. D., QUADRI G. J., SZAFIR D. A., ROSEN P.: Exploring Annotation Taxonomy in Grouped Bar Charts: A Qualitative Classroom Study. *Information Visualization* (2024). doi:10.1177/14738716241270247. 4
- [SBL21] SULTANUM N., BYLINSKII Z., LIU Z.: Leveraging Text-Chart Links to Support Authoring of Data-Driven Articles with VizFlow. In *Conference on Human Factors in Computing Systems - Proceedings* (5 2021), ACM. doi:10.1145/3411764.3445354. 2
- [SBT*16] SETLUR V., BATTERSBY S. E., TORY M., GOSSWEILER R., CHANG A. X.: Eviza: A Natural Language Interface for Visual Analysis. In *Annual Symposium on User Interface Software and Technology* (10 2016), ACM, pp. 365–377. doi:10.1145/2984511.2984588. 2
- [SCC*23] SUN M., CAI L., CUI W., WU Y., SHI Y., CAO N.: Erato: Cooperative Data Story Editing via Fact Interpolation. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 983–993. doi:10.1109/TVCG.2022.3209428. 2
- [SDS19] SRINIVASAN A., DRUCKER S. M., ENDERT A., STASKO J.: Augmenting Visualizations with Interactive Data Facts to Facilitate Interpretation and Communication. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 672–681. doi:10.1109/TVCG.2018.2865145. 2, 6
- [SH10] SEGEL E., HEER J.: Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1139–1148. doi:10.1109/TVCG.2010.179. 1
- [SH14] SATYANARAYAN A., HEER J.: Authoring Narrative Visualizations with Ellipsis. *Computer Graphics Forum* 33, 3 (2014), 361–370. doi:10.1111/cgf.12392. 2
- [SH24] SNYDER L. S., HEER J.: DIVI: Dynamically Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (1 2024), 403–413. doi:10.1109/TVCG.2023.3327172. 4
- [SKH*23] SHIN M., KIM J., HAN Y., XIE L., WHITELAW M., KWON B. C., KO S., ELMQVIST N.: Roslingifier: Semi-Automated Storytelling for Animated Scatterplots. *IEEE Transactions on Visualization and Computer Graphics* 29, 6 (2023), 2980–2995. doi:10.1109/TVCG.2022.3146329. 2
- [Sof24] SOFTWARE T.: Show Me: Automatic Presentation for Visual Analysis. Online Resource, 2024. Tableau Software. URL: <https://www.tableau.com/showme>. 10
- [SS18] SRINIVASAN A., STASKO J.: Orko: Facilitating Multimodal Interaction for Visual Exploration and Analysis of Networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 511–521. doi:10.1109/TVCG.2017.2745219. 2
- [SS23a] SRINIVASAN A., SETLUR V.: BOLT: A Natural Language Interface for Dashboard Authoring. In *Eurographics Conference on Visualization* (2023). URL: <https://api.semanticscholar.org/CorpusID:267751137>. 2
- [SS23b] SULTANUM N., SRINIVASAN A.: DataTales: Investigating the use of Large Language Models for Authoring Data-Driven Articles. In *2023 IEEE Visualization and Visual Analytics (VIS)* (2023), pp. 231–235. doi:10.1109/VIS54172.2023.00055. 2
- [SXS*21] SHI D., XU X., SUN F., SHI Y., CAO N.: Calliope: Automatic Visual Data Story Generation from a Spreadsheet. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 453–463. doi:10.1109/TVCG.2020.3030403. 2, 4, 6
- [SZZW24] SHEN L., ZHANG Y., ZHANG H., WANG Y.: Data Player: Automatic Generation of Data Videos with Narration-Animation Interplay. *IEEE Transactions on Visualization and Computer Graphics* 30, 01 (jan 2024), 109–119. doi:10.1109/TVCG.2023.3327197. 2, 10
- [tab] Tableau Public. <https://public.tableau.com/>. 3
- [TLS21] THOMPSON J. R., LIU Z., STASKO J.: Data Animator: Authoring Expressive Animated Data Graphics. In *CHI Conference on Human Factors in Computing Systems* (2021), CHI '21, ACM. doi:10.1145/3411764.3445747. 2
- [WSZ*20] WANG Y., SUN Z., ZHANG H., CUI W., XU K., MA X., ZHANG D.: DataShot: Automatic Generation of Fact Sheets from Tabular Data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 895–905. doi:10.1109/TVCG.2019.2934398. 2, 4
- [XHRC*18] XIA H., HENRY RICHE N., CHEVALIER F., DE ARAUJO B., WIGDOR D.: DataInk: Direct and Creative Data-Oriented Drawing. In *CHI Conference on Human Factors in Computing Systems* (2018), CHI '18, ACM, p. 1–13. doi:10.1145/3173574.3173797. 2
- [YKSJ07] YI J. S., KANG Y. A., STASKO J., JACKO J.: Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231. doi:10.1109/TVCG.2007.70515. 4
- [YS17] YU B., SILVA C. T.: VisFlow - Web-based Visualization Framework for Tabular Data with a Subset Flow Model. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 251–260. doi:10.1109/TVCG.2016.2598497. 10
- [YS20] YU B., SILVA C. T.: FlowSense: A Natural Language Interface for Visual Data Exploration within a Dataflow System. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (1 2020), 1–11. doi:10.1109/TVCG.2019.2934668. 2
- [YXL*22] YANG L., XU X., LAN X., LIU Z., GUO S., SHI Y., QU H., CAO N.: A Design Space for Applying the Freytag's Pyramid Structure to Data Stories. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 922–932. doi:10.1109/TVCG.2021.3114774. 2
- [ZE23] ZHAO Z., ELMQVIST N.: The Stories We Tell About Data: Surveying Data-driven Storytelling Using Visualization. *IEEE Computer Graphics and Applications* 43, 4 (2023), 97–110. doi:10.1109/MCG.2023.3269850. 2
- [ZHC24] ZHOU T., HUANG J., CHAN G. Y.-Y.: Epigraphics: Message-Driven Infographics Authoring. ACM (ACM), pp. 1–18. doi:10.1145/3613904.3642172. 2