



DashGuide: Authoring Interactive Dashboard Tours for Guiding Dashboard Users

N. Hoque¹  and N. Sultanum² 

¹University of Iowa, Department of Computer Science, USA
²Tableau Research, USA

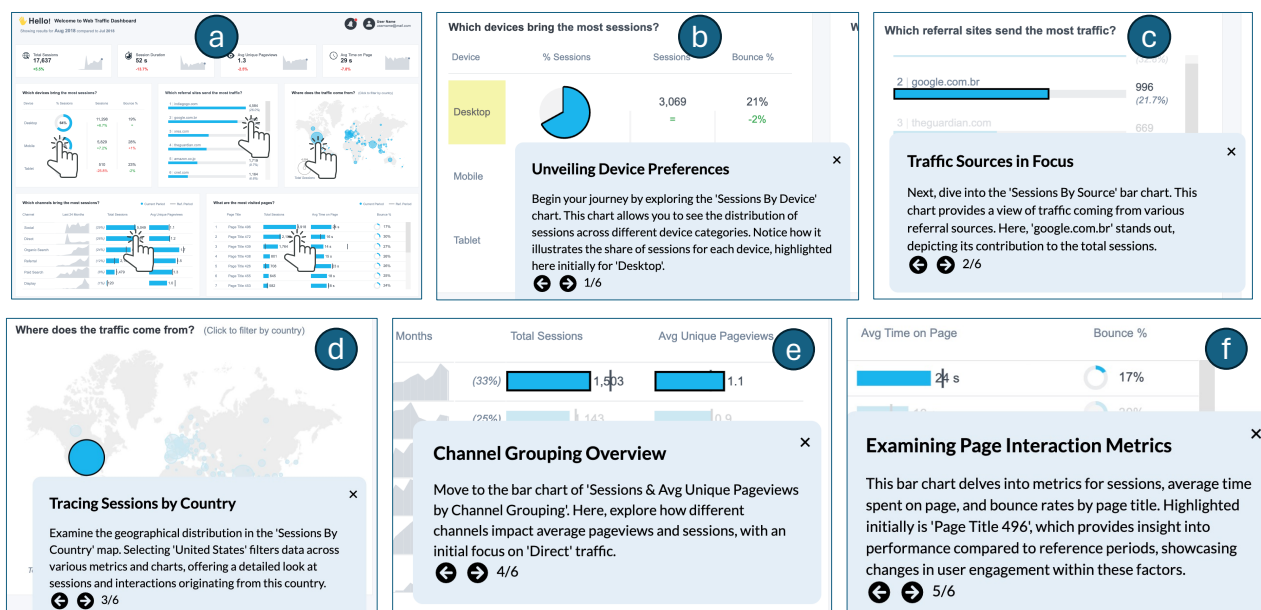


Figure 1: Overview of DASHGUIDE. a) The dashboard in this example visualizes different metrics from web traffic data (Sourced from Tableau Public [link]). To create a guided tour for this dashboard, the dashboard author first selects “Dashboard Semantics and Encoding” as the communication goal for the tour (not shown) and then performs several consecutive click selections on the dashboard (shown with the hand cursors). DASHGUIDE automatically generates content for the tour by prompting a large language model (LLM), which uses the communication goal and interactions provided by the author. The tour is then delivered as a step-by-step tour (b-f) where each step plays back the original interaction while placing the title and description for the step as an overlay near the location of interaction. The author can refine the content manually, provide further instructions to the LLM, or add or remove steps from the tour.

Abstract

Dashboard guidance helps dashboard users better navigate interactive features, understand the underlying data, and assess insights they can potentially extract from dashboards. However, authoring dashboard guidance is a time consuming task, and embedding guidance into dashboards for effective delivery is difficult to realize. In this work, we contribute DASHGUIDE, a framework and system to support the creation of interactive dashboard guidance with minimal authoring input. Given a dashboard and a communication goal, DASHGUIDE captures a sequence of author-performed interactions to generate guidance materials delivered as playable step-by-step overlays, a.k.a., dashboard tours. Authors can further edit and refine individual tour steps while receiving generative assistance. We also contribute findings from a formative assessment with 9 dashboard creators, which helped inform the design of DASHGUIDE; and findings from an evaluation of DASHGUIDE with 12 dashboard creators, suggesting it provides an improved authoring experience that balances efficiency, expressiveness, and creative freedom.

CCS Concepts

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

1. Introduction

Visualization dashboards were conceived as tools for independent, self-service data exploration [TBFC23]. In practice, however, dashboard users often require assistance in understanding components and functionality of a dashboard, from operating data filters to understanding the underlying data, interpreting charts, and navigating the dashboard to answer specific analytical questions [CGM*17, CGM19, KL16]. To this end, dashboard authors routinely provide *guidance* content to help these users, particularly those new to a dashboard [DWH*22, CLL*24, SCW*22].

Common forms of dashboard guidance include embedded text instructions and standard UI elements such as pop-ups and tooltips [CGM*17, SCW*22, SS25]. While such guidance is relatively easy to author and often supported by authoring tools like Tableau, they tend to lack user engagement and may be skipped, missed or ignored by users [SS25, SWG*21]. More recently, researchers have proposed interactive *dashboard tours* to improve user engagement and show diverse analytical capabilities of a dashboard [DHF*24, LSS*23]. These tours show usage scenarios, interactions, and insights in a step-by-step format, with each step incorporating a combination of content elements such as text descriptions and data highlights. Although recent work provides the tooling to instrument such tours [DHF*24], authors still need to manually draft the content itself (e.g., text descriptions and images), which demands significant time and effort and offers limited scalability to supporting multiple tours with diverse usage scenarios.

In view of these challenges, our work contributes: (a) a *framework and system*, DASHGUIDE, to support dashboard creators in authoring dashboard tours and content more easily; (b) *findings from a formative assessment* that guided design and development of DASHGUIDE; and (c) *findings from a qualitative user evaluation of DASHGUIDE* that underscore the value of our approach.

Our user-centered design journey began with formative interviews with 9 dashboard practitioners. Interview findings point to time and effort being the biggest obstacles to crafting engaging and comprehensive guidance. While dashboard creators favor dynamic forms of guidance such as dashboard tours, technical limitations result in guidance being delivered in mostly static form (e.g., technical documentation external to the dashboard). Interviews also informed a range of communication goals for dashboard guidance, from demonstrating coordinated views to presenting key insights and explaining chart encodings.

Based on this feedback, we designed DASHGUIDE, our solution for easier authoring of interactive dashboard tours. To expedite the authoring process while accommodating author oversight and communication goals, DASHGUIDE proposes an *interaction-first* approach. In this approach, a creator can quickly express their intent by choosing a high-level communication goal and then interacting with components of interest in the dashboard. DASHGUIDE captures the interaction, compiles a sequence of key demonstration segments, and links interacted dashboard components and corresponding interactive actions with the respective segments. A Large Language Model (LLM) is then used to generate relevant content for each key segment, steered by the author's informed communication goal. The dashboard tour is then rendered as a step-by-step guide where each step replays the corresponding interactions for

the step, along with an explanatory title and description placed near the interacted components as overlays (Figure 1). After an initial tour is created, DASHGUIDE provides several edit and refinement features, including inserting new steps and regenerating content.

We evaluated DASHGUIDE in a user study with 12 dashboard creators where we asked participants to author dashboard tours for their own (or publicly available) dashboards. Qualitative findings suggest that DASHGUIDE provided them with meaningful authoring support, and that the proposed interaction-first framework is a potentially promising approach for authoring dashboard tours. Our investigation also revealed open challenges around guidance authoring that future embodiments of our framework may help solve.

2. Related Work

Our work builds upon prior works on tutorial systems in HCI, dashboard guidance, and data tours. We discuss these areas below.

2.1. Guidance and Tutorial Systems in HCI

Guidance in HCI broadly encompasses problems and aids associated with learning to use software and hardware [GF10]. The traditional approach to providing guidance is text documentation [Car90]. However, documentation can be difficult to follow [Ret91], making them less accessible to users. A long line of HCI research has proposed solutions to make guidance more engaging and useful. The general consensus is that users prefer guidance that can show usage scenarios instead of outlining them in a static format [GF10, BCLO05, CFPE21, GAL*09, LGF12]. One popular approach is to adopt a step-by-step guided tour to show usage scenarios. Several works have proposed such tours for mobile applications [ZLCL21, CFPE21], desktop and web applications [CAR*12], and immersive environments (AR/VR) [KNDH19, HQW*21]. Several open-sourced libraries (e.g., Intro.js [Usa24], TourGuide.js [Tou24], and WebTour.js [Web24]) are also available for creating guided tours. This line of work inspired us to use guided tours that can sequentially show usage scenarios for a dashboard (i.e., dashboard tours).

While guided tours are popular, authoring them can be a challenging task. One prominent approach is to construct the steps for the tour by recording user demonstrations. For example, EverTutor [WCC*14] records touch events on a mobile phone to identify the steps for a tour. MixT [CAR*12] features a similar recording process but focuses more on combining images and videos for tutorial steps. Commercial tools such as Scribe [Scr21] also generate steps from user demonstrations. This line of work inspired our *interaction-first* approach where an author's demonstration, along with author-provided metadata, is used to determine the steps and narrative of the tour. We then utilize an LLM to generate content for the steps, which the author can further refine. To our knowledge, our work is the first to introduce such an interaction-driven and LLM-based method to author dashboard tours.

2.2. Characterizing Dashboard Guidance

Ceneda et al. [CGM*17] proposed the most comprehensive model for characterizing dashboard guidance. According to this model,

users may need guidance because they either do not know the desired results or the path to reach desired results. The model discusses five possible domains where users may need guidance: data, tasks, VA methods (i.e., charts, widgets, and interactions), users, and infrastructure. Finally, the model discussed three levels of guidance: *orienting*, *directing*, and *prescribing*. Orienting involves helping users build and preserve their mental map by exposing patterns, data subsets, insights, and provenance [GGL*14, LSS*12, KNS04]. Directing provides recommendations or alternative choices to users. Recommender systems such as VizAssist [BGV16] and Voyager [WMA*16] fall in this bucket. Finally, Prescribing, the focus of this paper, show usage scenarios and analytic processes of a dashboard through demonstration, often in a step-by-step tour format (i.e., dashboard tour) [YLRC10, DHF*24, SRG*23]. This type of guidance is often employed to onboard users in a dashboard and reduce learning curves [DWH*22, DHF*24, SMS*24].

DASHGUIDE falls in the *prescribing* category, as our dashboard tours primarily enable dashboard creators to prescribe usage scenarios to end users. Our work extends previous works focusing on prescribing in two ways. First, despite previous efforts to characterize prescribing, we still lack an understanding of how dashboard creators author such guidance in practice and what kind of support they require in this task. This work fills that gap through a formative study with dashboard creators. Secondly, this work proposes a system to author dashboard tours, an effective way to prescribe guidance to end users [SMS*24, DHF*24]. The system includes several components to expedite the authoring process: an interaction-based approach to capture the author's communication goal, an LLM-supported content generation strategy, and in-situ playback of the dashboard tours as overlays in the dashboard.

It should be noted that we do not limit the purpose of our dashboard tours to onboarding alone, although prescribing-focused guidance is typically designed for onboarding users [CGM*17]. This is because we wanted to understand the design space and potential use cases of dashboard tours as guidance beyond onboarding materials. We also note that we have not conducted studies with end-users in this paper due to our focus on supporting dashboard *authors* prescribe guidance. This is a complex and nuanced task [CGM*17], which made us anticipate that we would need extensive engagement with authors (i.e., dashboard creators) to understand their needs holistically and evaluate our solution. Thus, we decided to prioritize engagements with authors while leaving engagements with end-users for future work.

2.3. Data Tours, Storytelling, and Guidance

Data tours consist of interactive step-by-step data-driven content that walk users through a narrative or story [LSS*23]. From Asimov's "The Grand Tour" [Asi85], several use cases of data tours have emerged. A prominent use of data tours is data storytelling [RHDC18, ZE23, SH10]: news outlets such as The New York Times often use step-by-step format for data-driven stories [XJL*23], plus works such as VizFlow [SCBL21], graph comics [BKH*16], and NetworkNarratives [LSS*23].

Dashboard tours in DASHGUIDE follow the step-by-step format of a data tour. We opted to not use the term "data tour" as it is

widely used in the data storytelling space and may create confusion with our focus on guidance. While not as frequently as for storytelling, prior works have used guided tours as guidance. For example, IBM Cognos Analytics [IBM24] uses a tour to show different functionalities of the platform. The most closely related work to ours is D-Tour [DHF*24]. In this work, the authors proposed a system that allows authors to arrange automatically extracted dashboard components into multiple tours with possible shared steps between the tours. While D-Tour offers a flexible way to author multiple tours, authors have to manually construct the narrative for the tour and create content for each step. In contrast, DASHGUIDE expedites authoring by providing content generation assistance that aligns with authors' communication goals. Our mechanism also simplifies creating the steps and is more suitable for highlighting nuanced insights and scenarios, as authors can set focus to any data points and components just by interacting with the dashboard.

Another closely related work is LEVA [ZZZ*24], an LLM-based system that can generate onboarding materials for a dashboard. But contrary to our system, LEVA is fully automated and dashboard authors have few resources to steer content generation in a way that aligns with their communication goals. Our system prioritizes author agency and the content they want to show in guidance by providing an efficient way to capture those goals, while also providing author-driven tools for content refinement.

3. Formative Study

Our literature review revealed a knowledge gap between our theoretical understanding of prescribing guidance and common practices among dashboard creators in regards to authoring such guidance. To address this gap, we conducted semi-structured interviews with dashboard creators to understand how they currently author guidance focused on prescribing, their challenges, and potential solutions for authoring such guidance.

3.1. Participants and Methodology

We recruited nine participants (P1-P9) by advertising in a specialized community forum of data analysts and professional dashboard creators. Participants were data professionals with occupations ranging from business intelligence (BI) managers (3/9), developers (5/9), and a data consultant (1/9) — more details in supplemental materials. Participants had at least five years of experience producing dashboards, and were familiar with dashboard guidance authoring in various formats, such as technical documentation, text overlays, videos, and demos.

We conducted 45-min semi-structured interviews via videoconferencing, divided into three parts. In the first part, we asked participants about their professional background and target users. The second part focused on participant experiences with creating guidance, specially on prescribing guidance to end users. In the third part, participants brainstormed with the research team to outline requirements and design choices for a guidance authoring system. The 1st author conducted the interviews, while the 2nd author took notes and joined discussions. We conducted thematic analysis [BC06] of interview transcripts and study notes. Responses were aggregated across participants and recurrent themes were identified

under each question. We present our findings below, while establishing parallels with other relevant findings in past literature.

3.2. Findings

Participants shared their experiences on creating dashboards for a broad range of audiences and domains. Some supported dashboarding needs of small-to-medium specialized teams in their respective organizations (P3, P7-9); some supported much larger and diverse audiences with wildly varying levels of visual literacy skills (P1, P4, P5); and some worked with a broad range of clients in a consulting capacity (P2, P6). While their experiences are extremely diverse, they shared many of the same challenges and tasks. We compiled their feedback into the following topics.

3.2.1. Authoring guidance is a nuanced task

Participants stated that guidance is an integral part of their dashboard authoring practices. The primary reason for providing guidance is to attend to the diverse skills and needs of the target user bases or clients (P1-4, P6, P8-9). Prior research has also reported the challenge with visual literacy skills as the primary reason for providing guidance [TLM16, CLL*24, CCE*24]. Participants mentioned that while their dashboards typically have specific analytical purposes, the target user bases may not possess the skills or necessary context to realize those analytical purposes. On that, P2, frames the act of providing guidance as a *translation* exercise: “We assume people have the same understanding of language or words, which is not true.” This means that *good* guidance goes beyond a mechanical description of what dashboard components are and what they do, but is rather a nuanced *authoring* exercise: one that requires being attuned to a user’s context, picking the right words, and delivering them in an accessible way. This means that, instead of pursuing a fully automated guidance generation solutions, we should instead support authors realize their authoring potential.

3.2.2. Interactive and embedded guidance is the gold standard

Text annotations (4/9) and on-off overlays (5/9) were some of the most common prescribing guidance formats mentioned by practitioners. These forms of embedded guidance are desirable from a usability and convenience standpoint, as “you don’t want [users] to leave the dashboard” (P5). A limitation of these formats is that they are almost always static, and therefore, limited in terms of what aspects of the analytical process they can highlight: “Text overlays are my go-to guidance format. I typically have a button in the dashboard that opens an overlay layer with explanatory text on top of the dashboard. However, the overlays typically provide very high-level guidance and I cannot show any analytical process with that” (P2). Another limitation is that embedded guidance can be very high maintenance, often requiring third-party software (e.g., Figma) to author and manual labor to keep them updated: “The biggest challenge is updating the overlays. Every time I make any changes to the dashboard, I have to update my Figma design and then transform it to my dashboard” (P6). Other forms of guidance include on demo videos (P3-4, P7-8), in-person walkthroughs (P2-3, P7, P9), and phone calls (P1, P4-5).

Interestingly, none of the participants spontaneously mentioned

dashboard tours to be a common practice for prescribing guidance, despite prior research acknowledging their effectiveness [DHF*24, LSS*23, CAR*12]. When asked, all participants stated not knowing how to author tours, but were intrigued by the idea: “Oh! [tours] would be really cool. I kind of take a similar approach when I create videos, but it takes a lot of time to create just one video. If I can create multiple tours quickly, that would be super helpful” (P4). P8 said: “Tours could be really helpful if they can highlight data points and show interactions. But I don’t know how I can do that.” This confirms that dashboard tours are an appealing prescribing format but remain difficult to realize in practice.

3.2.3. Guidance communication goals can vary

Based on participant feedback, we identified four different kinds of communication goals for guidance focused on prescribing. The first kind is *dashboard semantics and encoding*, which explains the different elements of a dashboard, e.g. chart encodings (P2, P4, P6) or purpose of a dropdown menu (P2, P7). The second one is *interaction guidance*, focusing on explaining possible ways to interact with charts and their impacts on other dashboard components (P1-9). The third one is *data facts*, which shows insights and interesting data facts (P1-5, P8-9). Finally, *data backgrounds and context* explains background information and context about the dashboard (P1-3, P8). These findings align with prior research, which identified similar guidance intents [CGM*17, SS25]. For instance, *dashboard semantics* and *interaction guide* overlap with the VA methods and infrastructure categories proposed by Ceneda et al. [CGM*17], both of which focus on explaining charts, widgets, and dashboard infrastructure. Similarly, *data facts* overlap with three categories proposed by Ceneda et al.: *data, task, and users*.

3.3. Design Goals

The interviews suggest authoring effective guidance to prescribe to users is a time-consuming task. Participants were positive towards dashboard tours to prescribe guidance but had limited tools available to author them. We also identified four different communication goals for prescribing guidance. Based on the interviews, we decided to focus on mediating effectiveness and efficiency to facilitate the authoring process for dashboard tours. We coalesce the findings into a list of design goals (DG), as follows.

DG1: Expedite authoring while supporting expressivity and ownership. Dashboard tours should be easy to author, while authors should still retain ownership of the content. Authors should be given quick starting points and reasonable defaults as well as the ability to easily iterate over style and tone. D-Tour [DHF*24] is designed on a similar principle, but offers only template-based content generation. Our tool could use more advanced technology, such as LLMs, to expedite authoring and broaden expressivity in guidance. LEVA [ZZZ*24] showed that LLMs could be useful in generating onboarding materials. However, it does not offer any features to integrate author oversight. It also does not focus on dashboard tours. Our tool should address this gap by using a human-in-the-loop approach to integrate LLMs for dashboard tours.

DG2: Support communication goals. Past work [CGM*17] and formative findings point to a range of relatively well-defined

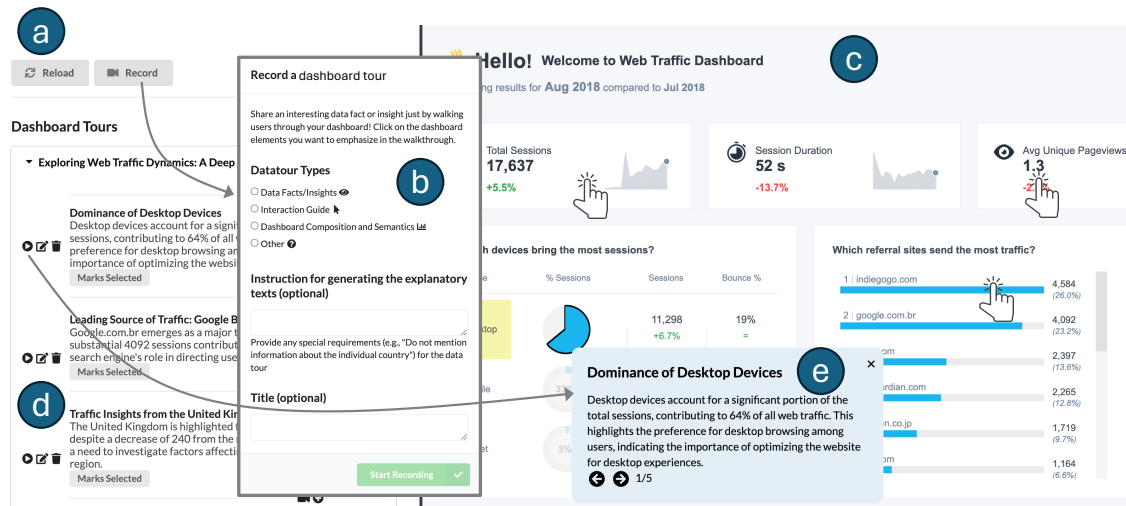


Figure 2: Visual Interface of DASHGUIDE. This example uses the same dashboard as the teaser image but the communication goal of the dashboard tour is **data facts** instead of **dashboard semantics and encoding**. a) Buttons for recording interactions and reloading a dashboard. b) On clicking the Record button (■), a modal opens where authors can select a communication goal, provide instructions to the LLM (optional), and a title (optional). c) After choosing the communication goal, authors can interact with the dashboard for creating the tour (shown with the hand cursors). d) Based on the interactions, a dashboard tour is created using an LLM. Each step in the tour contains three icon buttons for playing back the original interaction, editing the step, and deleting the step. e) Playback of a step in the dashboard. The title and description appear as an overlay in the dashboard.

communication goals (Sec. 3.2.3), which could be useful as guard rails for content generation while maximizing the utility and relevance of generated content. This range of goals should be acknowledged and supported, while still empowering authors to redirect goals as needed. While prior theoretical research has reported on the different goals [CGM*17, SS25], no prior tools explicitly support these communication goals [DHF*24, ZZZ*24].

DG3: Embed guidance into the dashboard. A critical decision for a dashboard tour is the format and content of an individual step. We found that dashboard tours should offer explanatory, data-relevant text in each step as an *embedded* dashboard element (Sec. 3.2.2). Prior research also found embedded elements, often as overlays, to be effective at providing contextual guidance [DHF*24, SWG*21, LSS*23].

4. DashGuide

DASHGUIDE is a web-based and LLM-assisted proof-of-concept system to support authoring tours for Tableau dashboards (Figure 2). We chose Tableau due to its prevalence for dashboard authoring, the large community of publicly available dashboards in Tableau Public [Van24], and its API that allows for Tableau dashboards to be embedded into web-based applications [Hel23].

4.1. Extracting Dashboard Metadata

To expedite authoring, DASHGUIDE automatically extracts several metadata about a dashboard using the Tableau Public API (DG1). This process results in dashboard metadata about charts, images, UI widgets (e.g., dropdowns, filters, etc.), and text content. For instance, for a chart, we extract its position, encoding, and marks.

For a text object, we extract the raw text data and its position in the dashboard. We store this metadata in a JSON object and feed it as context to the LLM. More details about the metadata and JSON object are available in supplemental materials.

4.2. Recording Dashboard Tour

To outline the steps for a dashboard tour, an author needs to follow two steps: 1) selecting tour metadata; and 2) interacting with dashboard components. This allows the author to craft the skeleton of the tour within minutes and with only a few user actions (DG1).

4.2.1. Tour Metadata

Figure 2b shows the interface for providing metadata about the tour. The primary purpose of this step is to collect the author's communication goal for the tour (DG2). The collected metadata are:

- **Communication Goal.** DASHGUIDE offers authoring support for three different types of communication goals: *dashboard semantics*, *interaction guide*, and *data facts* (DG2). The goal drives the type of content for the data tour. For instance, for a tour with *dashboard semantics* as the goal, we provide the following definition to the LLM: “Focus on explaining chart encoding and markers and the purpose of different filters, dropdowns, and other UI widgets in this dashboard tour.” For *interaction guide* the prompt is: “Focus on explaining interaction and the effects of interactions on other coordinated views for this dashboard tour.”
- **Instruction (optional).** Authors can specify any custom instructions, requirements, or preamble for the tour, such as “Write the tour in third-person point of view.”



Figure 3: Interfaces for editing a tour. a) The top row in the interface shows the title of the tour as well as four editing features for the tour. The play button (▶) hides all other components in the tool, except the dashboard, and then plays all steps one by one. The edit button (✎) opens the setting modal from Figure 2b. The other two buttons allow saving (💾) and deleting (🗑) the tour. Each step of the tour contains three icon buttons. b) The first button (⏮) plays back the interaction associated with the step. c) The second button (✎) opens a modal for editing the setting for an individual step. This interface allows users to change the guidance intent for a step, provide further instructions for regenerating the step, and manually edit the title and description. d) Users can insert a new step without any interaction by clicking the plus icon (+) and then outlining the instruction (e.g., “add a transition step here”) for the step. Authors can also click on the record buttons (📹) in between the steps to add new steps with interaction.

- **Title (optional).** Authors can provide a title for the tour. If provided, DASHGUIDE will use it as a context for the tour. If not, the tool will prompt the LLM to generate a title for the tour.

We refined the prompts for the metadata iteratively, informing several design criteria. For example, we noticed GPT-4 was not producing correct descriptions for the *data backgrounds and context*, one of the categories identified from the formative study (Sec. 3.2.3). This was because a large portion of the information on data typically is hidden in a Tableau Public dashboard for security reasons. For the sake of tapping from the large collection of real-world dashboards Tableau Public offers, and to avoid the incorrect LLM generations, we decided to forgo this category in this proof-of-concept implementation of DASHGUIDE. We also experimented with mixing multiple communication goals for a tour. However, the generated content did not reflect the multiple goals correctly in most cases. Thus, we decided to constraint a tour to one goal although authors can change the goal of an individual step.

4.2.2. Capturing Interaction

After providing the metadata, an author can interact with the dashboard to create steps for the tour. DASHGUIDE tracks common dashboard interactions [YaKSJ07, Shn96], including **data selection, brushing, and filters** applied to other views as a result of these interactions. We do not track hovering actions: this is partly due to a limitation of the Tableau embeddings API, but also because hovering interactions could create unnecessarily (and unintention-

ally) long interaction sequences with a large number of steps that authors would need to retroactively delete. We also track interactions with different UI widgets. These widgets include **buttons, dropdowns, checkboxes, and radio buttons**. We record changes to multiple coordinated views because of the interactions with the UI widgets. We store the interactions in a JSON object. We provide details about the JSON object in the supplemental materials.

4.3. Generating Tour

As mentioned in Sec. 4.1, we first package the metadata as a JSON and pass it to GPT-4 as context within a first prompt. We then prompt the LLM with the author-provided metadata and interactions as a second prompt to generate the initial dashboard tour. We iteratively refined the prompts by experimenting with various Tableau dashboards, all listed in the supplemental materials alongside our refinement strategy. The final prompt is as follows:

“The following is a JSON object outlining the structure of a dashboard tour - <a collection of interactions>. The key ‘steps’ includes a list of author-performed interactions/events. The ‘coordinatedViewChange’ key in each event represents changes to the coordinated views of the dashboard as a result of the event. The author also provided some metadata about the tour including a <communication goal>. Generate a title and short description for each event/step in natural language. Craft a story out of the steps and write the title and description accordingly. Return in the following JSON format {‘title’: <tour title>, steps: {<collection of event object with two keys, ‘title’ and ‘description’>}}.”

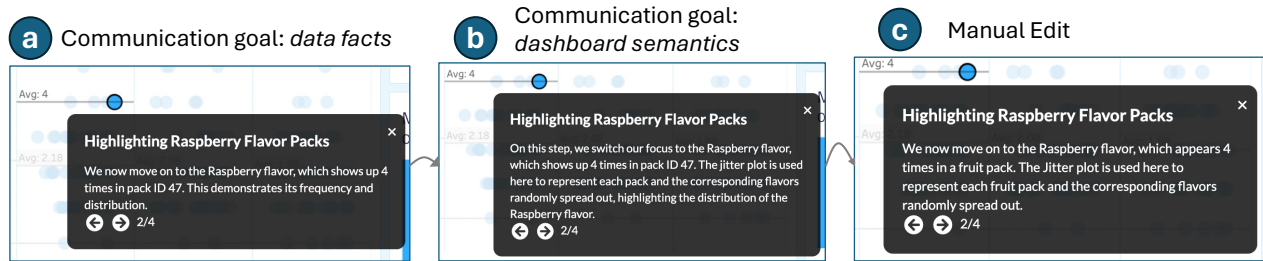


Figure 4: Editing a Dashboard Tour in DASHGUIDE. The dashboard in this example analyzes flavors (e.g., strawberry, peach, etc.) in fruit snacks. Consider that the author created a dashboard tour with the goal of communicating data facts. a) A step of the tour features a jitter plot. Thinking the users may not know about jitter plots, the author regenerates the step with another communication goal, dashboard semantics. b) The step is updated with an explanation on jitter plots. c) The author manually tweaks the description for a clearer explanation.

4.4. Editing Dashboard Tour

Once the initial tour is created, it is crucial that authors are able to verify the content and revise it as they wish (DG1). To support that, DASHGUIDE offers several editing features. The features include:

- **Editing Tour Metadata** (✎). Authors can use the interface discussed in Sec. 4.2.1 to revise the global settings of the tour.
 - **Communication Goal.** Authors can change the communication goal (e.g., from *data facts* to *interaction guide*) to change the overall content of all steps within the tour.
 - **Instruction and Title (optional).** Authors can provide custom instructions (e.g., *make the descriptions shorter*) to the LLM. Authors can also manually edit the title of the tour.
- **Save Tour** (💾). Authors can save the tour in our database.
- **Delete Tour** (🗑️). Authors can permanently delete the tour.
- **Edit Content and Setting for the Steps.** Authors have control over individual steps.
 - **Edit Title and Description.** Authors can use the interface in Figure 3c to manually edit the title and description for the step.
 - **Change Communication Goal and Instruction.** Using the same interface, authors can change the communication goal or add new instruction for the LLM and regenerate the step.
 - **Add New Steps** (➕). In between the steps, there are recording buttons that the author can use to record new interactions and add them as new steps to the guide.
 - **Add a Step Without an Associated Interaction** (⊕). This feature helps authors insert standalone steps that provide context, summary, or transitions to the narrative.

The dashboard tour, including all the editing features, is encapsulated in a UI card (Figure 3). We show an example scenario with different editing features in Figure 4.

4.5. Delivery/Playback

Following DG3, each step for a tour is delivered as an overlay text near the place of interaction (Figure 4). Authors can also freely move the overlay to readjust its position.

Authors can verify the delivery in two ways. First, they can use the play button with a step to play the particular step in the authoring view (Figure 3b). Authors can use the forward and backward icons on the overlay to move back and forth in a tour. The other option is using the play button (▶) at the top of the tour (Figure 3a).

This feature will hide all authoring controls on the interface and only show the dashboard and tour. This feature helps authors get a sense of how the tour will appear to an end user.

5. Evaluation

Following the design and development of DASHGUIDE, we conducted a user evaluation to validate its proposed approach for authoring dashboard tours. Participants used the tool to author tours for two dashboards (including one of their own, for those with Tableau Public profiles) and to share their impressions on the experience. We outline our study methodology and present our findings.

5.1. Recruitment and Methodology

We recruited 12 dashboard creators with Tableau experience—P3 and P6 from the formative study (Sec. 3.1), plus ten new participants P10-P19—via special interest groups. Participant backgrounds span data analysts and engineers (8), data consultants (2) and managers in relevant data/Bi roles (2), all with at least five years of professional dashboard authoring experience. Additional participant details can be found in supplemental materials.

Sessions were 1 hour long, conducted via video-conferencing with an experimenter and a note taker. Participants accessed DASHGUIDE on their own browsers while sharing their screen and thinking aloud. They performed two authoring tasks, each 8 minutes at most. Tasks were preceded by a demonstration of the tool, and were each followed by a brief discussion to reflect on participant's impressions and experiences. For the 1st task, authors were given a dashboard out of a list of two dashboards (counterbalanced across participants) to create a dashboard tour focusing on a given guidance goal (between *data facts*, *interaction guide*, and *dashboard semantics*, also counterbalanced across participants). This setup ensured a diversity of content generation exposure, but also a degree of consistency across experiences. After the 1st task, participants were asked for their impressions on the authoring experience, how they envisioned this tool fitting into their current practices, and alternative ways that guidance may be provided to their users. For the 2nd task, 10 participants created a dashboard tour for a Tableau Public dashboard that they authored, and 2 used a Tableau Public dashboard they were familiar with. This helped provide more ecologically valid context for authoring. Follow-up questions focused

on perceived time or quality gains with a tool like DASHGUIDE in their practice, and suggestions for improvement. More details on the study protocol can be found in supplemental materials.

Sessions were screen-recorded, audio-recorded, and automatically transcribed. Extensive notes were also taken, and supported the initial stages of thematic analysis along with the transcripts. Usage logs were collected to inform basic usage patterns, including the first and last stages of their authored tours. Qualitative findings drive most of the analysis, contextualized by illustrative quotes, relevant log stats, and frequency counts in the format ($_ / 12$) denoting the number of participants who supported the theme. Frequency counts reflect *spontaneous* comments from participants and are thus a *lower bound*, and not a full signal, of theme relevance.

5.2. Findings

Overall, participants reactions to DASHGUIDE were very positive. They appreciated the ease of use (5/12), and how quickly they can craft tours with it (5/12). When asked about foreseeable gains if using DASHGUIDE in their practice, most (11/12) concurred on at least one aspect of efficiency and quality, e.g., “*it would add value for the same time invested*” (P15), and “*guidance is the last mile of dashboarding, the quicker I can make it the better*” (P13). A majority of participants also commented on the higher-than-expected quality of content generation (7/12), e.g., “*I was pleasantly surprised when it did what I wanted it to do*” (P14), and “*I like this much more than I thought I would*” (P13).

Participants remarked on the importance of interactive and relevant hands-off guidance for dashboard users, as “*there’s that group of people for whom data is extremely overwhelming. And giving this to someone (...) this feeling of ‘someone has guided me around this insight’, I feel like it would be awesome*” (P11). Also, “*having this as a dashboard guide is super helpful, because it’s a way to cut back on [users] needing me to go through things over and over again*” (P6). Beyond being “*really helpful in training new users*” (P16), one participant argued DASHGUIDE’s usefulness may extend to experienced users as well: “*maybe you haven’t used the dashboard in a long time, you come back to it and you want to re-orient yourself to it*” (P11).

Out of all participants, only one — who had significant experience in recording instructional videos for guidance — saw limited value in DASHGUIDE’s assistance: “*it has to be easier to use than any default video recording software (...) I was expecting the tool to do something similar [to a video recording], but it’s more of a screenshotting tool.*” (P17). On that, even other participants who were enthusiastic about DASHGUIDE also noted several points of friction, including usability and workflow issues. More notably, they identified contextual challenges beyond the current scope of DASHGUIDE that point to compelling research opportunities for future work, which we cover later in Sec. 6.1.

Here, we discuss strengths and opportunities for improvement for DASHGUIDE. We organize our findings under four themes: observations on authoring workflows (Sec. 5.2.1), impressions on AI-generated content (Sec. 5.2.2), reactions to DASHGUIDE’s approach to communication goal (Sec. 5.2.3), and comments on formatting and delivery of guidance to end users (Sec. 5.2.4).

5.2.1. On Authoring Workflows and Interaction Capture

Findings suggest that even as a proof-of-concept, DASHGUIDE’s interaction-centered workflow not only worked well as a tour authoring approach but also supported a diversity of authoring workflows. Within the allotted 8 minutes for each task, participants were able to quickly internalize and appropriate the authoring features. After the initial recording, on average, participants appended 1.50 (CI: 0.00, 3.09) new interactive steps and 2.20 (CI: 0.58, 3.82) new non-interactive steps (see supplemental materials). On average, participants re-generated 2.00 (CI: 0.52, 3.48) steps with the help of LLM and manually edited 3.33 (CI: -6.71, 13.37) steps. Finally, participants regenerated the full tour 2.43 (CI: 0.30, 4.55) times on average. A chart is available in supplemental materials.

Some participants found DASHGUIDE’s interaction-centered approach to be intuitive, both from a guidance perspective, e.g., “*I want to be able to show them exactly where to go*” (P6); and from a usage perspective, e.g., “*it makes sense for me to do the action that I think the user will take*” (P10). Many also stated that authoring flexibility is important to them and praised features of the tool that support this freedom, e.g., inserting steps in the middle (P6), re-prompting (P11, P14), and direct editing (P13). This diversity underscores the wide range of workflows DASHGUIDE supports.

Participants also highlighted limitations with DASHGUIDE’s proposed workflow. First, it requires an initial sequence of interactive actions before content is generated, which imposes a higher degree of deliberation and planning from the author and is thus less amenable to narrative and flow exploration. P14 also anticipated challenges arising with more spontaneous interaction workflows: “*I would find that having to delete stray marks, because I click around too much, would annoy me after a while. Especially if I was going through multiple multi-step tours.*” (P14). Suggested improvements include pre-filtering interactions of interest, means for easy editing of captured step sequences before generation, and the ability to group steps into sections or segments.

5.2.2. On Content Generation

Impressions on content generation were generally positive, with 7/12 spontaneously stating they appreciated the auto-generated assistance, including self-proclaimed “skeptics of AI”: “*I’m typically quite skeptical of everything that has AI baked into it, but this works*” (P15). One benefit they found was expedited authoring: “*‘Data facts’ and ‘dashboard semantics’ are the most labor intensive. I could see that saving a lot of time, a lot of headache, a lot of maintenance work, and a lot of repeated education*” (P18). Another benefit was helping overcome the “*blank screen*” (P14) paralysis: “*These are things that I would probably spend a lot of time writing up and thinking about and I don’t feel like I’d have to make a lot of changes.*” (P16). These findings are echoed by think aloud utterances, where 10/12 participants shared positive reactions toward content generation at least once across the two tasks. To many, it largely fulfilled expectations from the get go, as about 72% of all authored steps were kept as is: e.g., “*the insights that came out, right out the gate, I either didn’t have any edits, or the edits that I did have were very easy*” (P18). This suggests our proposed AI-assisted workflow is helpful and that content generation does bring some efficiency. Some also liked the responsiveness to custom in-

structions, e.g., “it was about 50% but changing that prompt got it to about 90-95% of where I’d want it” (P6). Some also appreciated the external knowledge brought in, e.g., the “extra context for the analysis that I do not provide on this dashboard” (P12).

Regarding opportunities for improvement, suggestions called for more implicit and explicit context management over generative AI. The most common ask was having more control over context to be passed to the large language model (5/12). Participants wanted specific information to be taken into account, e.g., recent or specialized knowledge, which could be potentially mediated by data dictionaries and shared semantic layers. But when internal labels used in the Tableau workbooks appeared in the content, some participants wanted them *excluded* from the generation context as they were not always user friendly: “what we name our worksheets probably isn’t relevant to users” (P16). One participant (P3) suggested interpreting other visual dashboard elements, e.g., a magnifying glass for a search icon, and factor their meaning into the generation. Other ideas included more proactive interventions over content (3/12), e.g., suggesting initial steps for an *interaction guide* tour based on interactive widgets present on the dashboard (P3), and recommending follow-up steps to append to existing tours (P11).

AI skepticism aside, several participants (4/12) working in data-sensitive domains (e.g. governmental) also shared practical obstacles to the use of large language models in their dashboard authoring practices, from security restrictions to high cost and sustainability concerns. As such, they hoped for an option to “skip the AI” (P13), or potentially have it replaced by classic template-based approaches (P16), arguing DASHGUIDE provides value even without the generative piece: “If you were to launch this annotation experience without ChatGPT integration you would still have huge gains, because you can enter your own insights. The ability to create an annotated walkthrough in the manner you created is a huge win” (P11).

5.2.3. On Communication Goal

Participants chose a variety of communication goals on their free form task (2nd task), including 6 *data facts*, 4 *interaction guides*, and 2 *dashboard semantics*. Feedback points to these types being largely aligned with communication needs of dashboard authors: “I really think the 3 categories are 90% of what an annotation guidance would even need” (P18). Several participants (4/12) also appreciated DASHGUIDE’s potential for storytelling beyond guidance, such as presentations (P14). Overall, we found the *communication goal* approach to be an effective and versatile authoring scaffold, giving authors a range of guidance options while also equipping content generation with suitable guard rails.

As next steps, several participants (5/12) suggested expanding the list of communication goals to support *data backgrounds and context*, the specific communication goal we chose not support in DASHGUIDE’s current implementation (as covered in Sec. 4.2.1) but is still feasible within our framework: “In my day to day, it’s something that we’re always asking: ‘what’s the numerator’, ‘the denominator’, ‘what’s feeding into these 1%’, etc.” (P18). One participant also suggested experimenting with multiple active communication goals per tour: “I imagine there being times I’d want to combine them and, instead of one or the other, create a guide that does all of them together” (P16). Others also suggested having

more control over global tone and wording choices (3/12), e.g., a slider for positive or negative tone (P10).

5.2.4. On Guidance Delivery

DASHGUIDE’s embedded step-by-step format was informed from our formative investigation presented in Sec. 3.2.2. In this evaluation, participants echoed similar impressions; they praised DASHGUIDE’s co-located delivery (3/12), how actions are replayed while the dashboard is still interactable (3/12), and the styling of content pop-ups (4/12). The ability to create multiple data tours for a dashboard was also acknowledged (3/12) and considered useful.

Other priorities participants had for guidance delivery include customizing text and pop-up styling (6/12), but also considering styling in the context of accessibility, for things like font size and contrast (2/12). On that, a text-based delivery may be considered a potential boon for accessibility provided that “screen readers can tap into those blurbs” (P6). Other opportunities included consistent entry points in dashboards for triggering tours (P6); flexible user controls to choose between multiple tours or from which step to start the tour (P10); multi-modal presentation, e.g. with video excerpts and animated *gifs* (P6, P19); and a user-accessible table of contents (P10). Finally, some participants underscored the value of DASHGUIDE as a more general authoring tool and wished to see it applied to formats such as reports (P12) and presentations (P14).

6. Discussion

Study findings suggest DASHGUIDE is both an effective and an efficient alternative for authoring dashboard tours. Participants were generally enthusiastic about the prospects of assisted authoring, and were often positively surprised by the generated content. This suggests that the guard rails we designed, which in turn were informed by feedback from practitioners and prior work, are well aligned with dashboard authors’ communication needs. In the following segments, we take a step back and discuss opportunities for further investigation that emerged from our evaluation. We also reflect on the impact of this work in the space of dashboard guidance and LLM-assisted authoring, and discuss limitations of our work.

6.1. Open Ended Challenges and Opportunities

Study point to DASHGUIDE addressing a relevant need, but also informed areas for improvement. Some of these areas touched on open research questions, which could in turn inform compelling opportunities for future work. We discuss them as follows.

6.1.1. Handling data updates

In our formative study, dashboard updates were considered a significant pain point for guidance maintenance. In line with this, one of the most frequent questions participants asked this time around was whether DASHGUIDE handled data updates (6/12), which naturally leads to the question of how generative AI may reconcile those changes with the guidance content. While generative AI could potentially update this guidance automatically, it should not do so without oversight, and the ideal extent and format for such oversight are not immediately clear. Simple solutions exist, such as automatic alerts on data changes (P13), but additional guard rails to

improve quality of generated content would be worth investigating. Another solution is to combine generation with template-based approaches and text-data links, such as the dynamic text placeholders often used in dashboard tooltips and data labels [SS25]: “*I would expect to be able to use field references in the descriptions.*” (P14)

6.1.2. Reusing guidance content

One participant remarked on the opportunity to reuse previous tours: “*If you’ve got dashboards that are very similar, it would be really nice to reuse a tour.*” (P10). Following up on Sec. 6.1.1, previous content could serve as guard rails to generate new guidance content, constituting a form of reuse. Other reuse scenarios could include repurposing tours for different target audiences (e.g., sales and marketing”), by following the same interaction route but adapting the language and takeaways; collectively repurposing tours from the same source (e.g., an author or team) for structure and tone, to create new content that is aligned and consistent with past content; and having end-users reapplying style and interaction flow from past guidance they liked onto a new dashboard.

6.1.3. Adapting to users

Supporting audiences with different backgrounds, goals, and skills is one of the many challenges guidance authors face. DASHGUIDE enables that via rapid authoring of multiple tours for a dashboard, which is often a desirable setup, but still requires authors to proactively create tours ahead of time and know their users well: “*I don’t want business users from this department looking at [a dashboard] the same way as the accountants, especially if I am trying to use a single dashboard to serve multiple purposes*” (P14). Beyond further automation via content reuse (Sec. 6.1.2), there is also an opportunity to tweak guidance at delivery time. One idea is to pick up contextual signals from the user (such as a trace of dashboard interaction) to identify relevant guidance steps from existing tours and selectively present them. Another idea is to use generative AI to create transitions between steps, adding cohesive flow while still largely respecting pre-vetted guidance content. Finally, one participant argued collective interaction traces from many users could also provide useful signals to guide dashboard tour authoring: “*I’d like to see where all those clicks are distributed*” (P12).

6.1.4. Building on the exploration-explanation continuum

In view of DASHGUIDE’s applicability to storytelling-related guidance, we argue it could also be used to chronicle an analytical exploration workflow performed on the dashboard. This idea is inspired by prior research in provenance of user interactions for sensemaking [XOW*20] and mediating sensemaking and communication [CLA*20], as well as a participant’s comment on early stages of authoring: “*when you’re in an authoring flow, when you are clicking somewhere, you don’t know whether the insight you discovered is gonna be worth sharing. You are not ready to add it to your tour yet, but you think it might be. (..) You want to add bookmarks in the flow of analysis of something that you might want to refer to later*” (P11). On this, DASHGUIDE could be extended with means to retrieve, organize, regroup and contextualize interaction traces that can be flexibly mapped dashboard tour steps later.

6.2. Reflections on the Authoring Workflow in DASHGUIDE

Despite the popularity of prompt-based interaction, they can also be misaligned to how humans perceive tasks and goals [KLK*24]. Subramonyam et al. [SPP*24] discussed the case of “gulf of envisioning”, or the challenges of tweaking a prompt to improve output quality. In DASHGUIDE, authors prompt the LLM through a mix of UI widgets, dashboard interactions, and direct instructions, thus deconstructing the prompting task into smaller visual tasks. We believe this approach helped reduce the “gulf of envisioning” for authors. We hope our work, along with other recent works in HCI [ASV*24, KLCK23], helps motivate future research to consider visually constructing prompts with smaller tasks.

6.3. Limitations

DASHGUIDE has several limitations. The tool is built upon Tableau Embedding API, which allowed us to concretely evaluate the tool on a very diverse ecosystem of real-world dashboards, but with limited support for capturing interactions (e.g., hovering actions). Nevertheless, the *interaction-first* approach proposed in this work is platform-independent and could be implemented for other platforms by following the structure of the JSON files and LLM-prompts. Detailed instruction for adapting DASHGUIDE to other platforms can be found in the supplemental materials.

In its current form, DASHGUIDE lacks a few post-capture editing features, such as merging steps. We also acknowledge that mapping one step per one interaction may not capture complex insights, generated by a sequence of interactions. Similar to other LLM-based applications, there is a chance of the LLM generating incorrect or non-existent information (i.e., hallucinations) [Bea21]. Finally, while our results show promises about how end-users might benefit from the tours, we have not formally evaluated DASHGUIDE with them, something which would help characterize the extent of DASHGUIDE’s benefits and better assess its delivery format.

7. Conclusion

We presented DASHGUIDE, a framework and system that can support dashboard creators in authoring interactive dashboard tours for guiding end users. This paper advances the state-of-the-art for dashboard tours by demonstrating how we can utilize LLMs to flexibly author the content of the tours. We conducted formative interviews to identify design requirements for LLM-assisted authoring of dashboard tours. Based on author feedback, we designed different aspects of the tool, from communication goals to guidance delivery and authoring control over the content of the tours. We conducted a summative user evaluation with 12 dashboard creators, where they authored guided tours for existing dashboards using our tool. DASHGUIDE’s *interaction-first* authoring framework was found to be efficient and intuitive, which we argue may lower barriers for authoring dashboard guidance. We also identified challenges and opportunities for future investigations to guide further research on the topic. As a community, we are still in the early stages of learning how to utilize this technology in ways that uphold human agency, awareness, creativity, and trust. We hope our work can help inspire other LLM-assisted work striving for the same goals.

References

- [Asi85] ASIMOV D.: The grand tour: A tool for viewing multidimensional data. *SIAM Journal on Scientific and Statistical Computing* 6, 1 (1985), 128–143. doi:10.1137/0906011.3
- [ASV*24] ARAWJO I., SWOOPES C., VAITHILINGAM P., WATTENBERG M., GLASSMAN E. L.: Chainforge: A visual toolkit for prompt engineering and LLM hypothesis testing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2024), ACM, pp. 304:1–304:18. doi:10.1145/3613904.3642016.10
- [BC06] BRAUN V., CLARKE V.: Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp063oa.3
- [BCLO05] BERGMAN L. D., CASTELLI V., LAU T. A., OBLINGER D.: Docwizards: a system for authoring follow-me documentation wizards. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2005), ACM, pp. 191–200. doi:10.1145/1095034.1095067.2
- [Bea21] BOMMASANI R., ET AL.: On the opportunities and risks of foundation models. *CoRR abs/2108.07258* (2021). doi:10.48550/arXiv.2108.07258.10
- [BGV16] BOUALI F., GUETTALA A. E., VENTURINI G.: Vizassist: an interactive user assistant for visual data mining. *Vis. Comput.* 32, 11 (2016), 1447–1463. doi:10.1007/s00371-015-1132-9.3
- [BKH*16] BACH B., KERRACHER N., HALL K. W., CARPENDALE S., KENNEDY J., RICHE N. H.: Telling stories about dynamic networks with graph comics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2016), ACM, pp. 3670–3682. doi:10.1145/2858036.2858387.3
- [Car90] CARROLL J. M.: *The Nurnberg funnel: Designing minimalist instruction for practical computer skill*. MIT press, 1990. 2
- [CAR*12] CHI P., AHN S., REN A., DONTCHEVA M., LI W., HARTMANN B.: Mixt: automatic generation of step-by-step mixed media tutorials. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), ACM, pp. 93–102. doi:10.1145/2380116.2380130.2,4
- [CCE*24] CENEDA D., COLLINS C., EL-ASSADY M., MIKSCH S., TOMINSKI C., ARLEO A.: A heuristic approach for dual expert/end-user evaluation of guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 997–1007. doi:10.1109/TVCG.2023.3327152.4
- [CFPE21] CHI P., FREY N., PANOVICH K., ESSA I.: Automatic instructional video creation from a markdown-formatted tutorial. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2021), ACM, pp. 677–690. doi:10.1145/3472749.3474778.2
- [CGM*17] CENEDA D., GSCHWANDTNER T., MAY T., MIKSCH S., SCHULZ H., STREIT M., TOMINSKI C.: Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 111–120. doi:10.1109/TVCG.2016.2598468.2,3,4,5
- [CGM19] CENEDA D., GSCHWANDTNER T., MIKSCH S.: A review of guidance approaches in visual data analysis: A multifocal perspective. *Computer Graphics Forum* 38, 3 (2019), 861–879. doi:10.1111/CGF.13730.2
- [CLA*20] CHEN S., LI J., ANDRIENKO G. L., ANDRIENKO N. V., 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 (2020), 2499–2516. doi:10.1109/TVCG.2018.2889054.10
- [CLL*24] CHOE K., LEE C., LEE S., SONG J., CHO A., KIM N. W., SEO J.: Enhancing data literacy on-demand: Llms as guides for novices in chart interpretation. *IEEE Transactions on Visualization and Computer Graphics* (2024), 1–17. doi:10.1109/TVCG.2024.3413195.2,4
- [DHF*24] DHANOA V., HINTERREITER A., FEDIUK V., ELMQVIST N., GROELLER E., STREIT M.: D-tour: Semi-automatic generation of interactive guided tours for visualization dashboard onboarding. *IEEE Transactions on Visualization and Computer Graphics* 1, 1 (2024), 1–11. doi:10.1109/TVCG.2024.3456347.2,3,4,5
- [DWH*22] DHANOA V., WALCHSHOFER C., HINTERREITER A. P., STITZ H., GRÖLLER M. E., STREIT M.: A process model for dashboard onboarding. *Computer Graphics Forum* 41, 3 (2022), 501–513. doi:10.1111/CGF.14558.2,3
- [GAL*09] GRABLER F., AGRAWALA M., LI W., DONTCHEVA M., IGARASHI T.: Generating photo manipulation tutorials by demonstration. *ACM Trans. Graph.* 28, 3 (2009), 66. doi:10.1145/1531326.1531372.2
- [GF10] GROSSMAN T., FITZMAURICE G. W.: Toolclips: an investigation of contextual video assistance for functionality understanding. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2010), ACM, pp. 1515–1524. doi:10.1145/1753326.1753552.2
- [GGL*14] GRATZL S., GEHLENBORG N., LEX A., PFISTER H., STREIT M.: Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2023–2032. doi:10.1109/TVCG.2014.2346260.3
- [Hel23] HELP T.: Tableau Embedding API V3, 2023. URL: https://help.tableau.com/current/api/embedding_api/en-us/index.html.5
- [HQW*21] HUANG G., QIAN X., WANG T., PATEL F., SREERAM M., CAO Y., RAMANI K., QUINN A. J.: AdapTutAR: An adaptive tutoring system for machine tasks in augmented reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM, pp. 417:1–417:15. doi:10.1145/3411764.3445283.2
- [IBM24] IBM: Ibm Cognos Analytics, 2024. URL: <https://www.ibm.com/at-de/products/cognos-analytics>.3
- [KL16] KWON B. C., LEE B.: A comparative evaluation on online learning approaches using parallel coordinate visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2016), ACM, pp. 993–997. doi:10.1145/2858036.2858101.2
- [KLCK23] KIM T. S., LEE Y., CHANG M., KIM J.: Cells, generators, and lenses: Design framework for object-oriented interaction with large language models. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2023), ACM, pp. 4:1–4:18. doi:10.1145/3586183.3606833.10
- [KLK*24] KIM Y., LEE J., KIM S., PARK J., KIM J.: Understanding users' dissatisfaction with chatgpt responses: Types, resolving tactics, and the effect of knowledge level. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (New York, NY, USA, 2024), ACM, pp. 385–404. doi:10.1145/3640543.3645148.10
- [KNDH19] KUMARAVEL B. T., NGUYEN C., DIVERDI S., HARTMANN B.: TutoriVR: A video-based tutorial system for design applications in virtual reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2019), ACM, p. 284. doi:10.1145/3290605.3300514.2
- [KNS04] KREUSELER M., NOCKE T., SCHUMANN H.: A history mechanism for visual data mining. In *IEEE Symposium on Information Visualization* (2004), Ward M. O., Munzner T., (Eds.), IEEE Computer Society, pp. 49–56. doi:10.1109/INFVIS.2004.2.3
- [LGF12] LI W., GROSSMAN T., FITZMAURICE G. W.: GamiCAD: a gamified tutorial system for first time autocad users. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), ACM, pp. 103–112. doi:10.1145/2380116.2380131.2
- [LSS*12] LEX A., STREIT M., SCHULZ H., PARTL C., SCHMAL-

- STIEG D., PARK P. J., GEHLENBORG N.: StratomeX: Visual analysis of large-scale heterogeneous genomics data for cancer subtype characterization. *Computer Graphics Forum* 31, 3pt3 (2012), 1175–1184. doi:10.1111/J.1467-8659.2012.03110.X. 3
- [LSS*23] LI W., SCHÖTTLER S., SCOTT-BROWN J., WANG Y., CHEN S., QU H., BACH B.: NetworkNarratives: Data tours for visual network exploration and analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2023), ACM, pp. 172:1–172:15. doi:10.1145/3544548.3581452. 2, 3, 4, 5
- [Ret91] RETTIG M.: Nobody reads documentation. *Commun. ACM* 34, 7 (1991), 19–24. doi:10.1145/105783.105788. 2
- [RHDC18] RICHE N. H., HURTER C., DIAKOPOULOS N., CARPENDALE S.: *Data-driven storytelling*. CRC Press, 2018. 3
- [SCBL21] SULTANUM N., CHEVALIER F., BYLINSKII Z., LIU Z.: Leveraging text-chart links to support authoring of data-driven articles with vizflow. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM, pp. 16:1–16:17. doi:10.1145/3411764.3445354. 3
- [Scr21] SCRIBE: Scribe: Turn any process into a step-by-step guide, instantly., 2021. URL: <https://scribehow.com>. 2
- [SCW*22] STOIBER C., CENEDA D., WAGNER M., SCHEITNER V., GSCHWANDTNER T., STREIT M., MIKSCH S., AIGNER W.: Perspectives of visualization onboarding and guidance in VA. *Visual Informatics* 6, 1 (2022), 68–83. doi:10.1016/J.VISINF.2022.02.005. 2
- [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. 3
- [Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages* (1996), IEEE Computer Society, pp. 336–343. doi:10.1109/VL.1996.545307. 6
- [SMS*24] STOIBER C., MOITZI D., STITZ H., GRASSINGER F., PRAKASH A. S. G., GIRARDI D., STREIT M., AIGNER W.: VisAhoi: Towards a library to generate and integrate visualization onboarding using high-level visualization grammars. *Visual Informatics* 8, 1 (2024), 1–17. doi:10.1016/J.VISINF.2024.06.001. 3
- [SPP*24] SUBRAMONYAM H., PEA R., PONDOC C. L., AGRAWALA M., SEIFERT C. M.: Bridging the gulf of envisioning: Cognitive challenges in prompt based interactions with llms. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2024), ACM, pp. 1039:1–1039:19. doi:10.1145/3613904.3642754. 10
- [SRG*23] STOIBER C., RADKOHL S., GRASSINGER F., MOITZI D., STITZ H., GOLDGRUBER E., GIRARDI D., AIGNER W.: Authoring tool for data journalists integrating self-explanatory visualization onboarding concept for a treemap visualization. In *Proceedings of the Biannual Conference of the Italian SIGCHI Chapter, CHIItaly* (New York, NY, USA, 2023), ACM, pp. 2:1–2:14. doi:10.1145/3605390.3605394. 3
- [SS25] SULTANUM N., SETLUR V.: From instruction to insight: Exploring the functional and semantic roles of text in interactive dashboards. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2025), 382–392. doi:10.1109/TVCG.2024.3456601. 2, 4, 5, 10
- [SWG*21] STOIBER C., WALCHSHOFER C., GRASSINGER F., STITZ H., STREIT M., AIGNER W.: Design and comparative evaluation of visualization onboarding methods. In *The 14th International Symposium on Visual Information Communication and Interaction* (2021), ACM, pp. 11:1–11:5. doi:10.1145/3481549.3481558. 2, 5
- [TBFC23] TORY M., BARTRAM L., FIORE-GARTLAND B., CRISAN A.: Finding their data voice: Practices and challenges of dashboard users. *IEEE Computer Graphics and Applications* 43, 1 (2023), 22–36. doi:10.1109/MCG.2021.3136545. 2
- [TLM16] TANAHASHI Y., LEAF N., MA K.: A study on designing effective introductory materials for information visualization. *Computer Graphics Forum* 35, 7 (2016), 117–126. doi:10.1111/CGF.13009. 4
- [Tou24] TOURGUIDE: Tourguide js: Intuitive and customizable tours for user onboarding, 2024. URL: <https://tourguidejs.com>. 2
- [Usa24] USABLICA: Intro.js: User onboarding and product walkthrough, 2024. URL: <https://introjs.com>. 2
- [Van24] VANDERAA K.: 11 Most-Favorited Data Visualizations on Tableau Public, 2024. URL: <https://www.tableau.com/blog/8-most-favorited-data-visualizations-tableau-public>. 5
- [WCC*14] WANG C., CHU W., CHEN H., HSU C., CHEN M. Y.: Evertutor: automatically creating interactive guided tutorials on smartphones by user demonstration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), ACM, pp. 4027–4036. doi:10.1145/2556288.2557407. 2
- [Web24] WEBTOUR: Webtour js: Step-by-step guide in vanilla javascript, 2024. URL: <https://votch18.github.io/webtour.js/>. 2
- [WMA*16] WONGSUPHASAWAT K., MORITZ D., ANAND A., MACKINLAY J. D., HOWE B., HEER J.: Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 649–658. doi:10.1109/TVCG.2015.2467191. 3
- [XJL*23] XIAO M., JHAVERI I., LUTZ E., KOETTL C., BARNES, E. J.: Tracking the chinese balloon from space, 2023. URL: <https://www.nytimes.com/interactive/2023/03/20/science/chinese-space-balloon-incident.html>. 3
- [XOW*20] XU K., OTTLEY A., WALCHSHOFER C., STREIT M., CHANG R., WENSKOVITZ J. E.: Survey on the analysis of user interactions and visualization provenance. *Computer Graphics Forum* 39, 3 (2020), 757–783. doi:10.1111/CGF.14035. 10
- [YaKSJ07] YI J. S., AH KANG Y., STASKO J. T., JACKO J. A.: 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. 6
- [YLRC10] YU L., LU A., RIBARSKY W., CHEN W.: Automatic animation for time-varying data visualization. *Computer Graphics Forum* 29, 7 (2010), 2271–2280. doi:10.1111/J.1467-8659.2010.01816.X. 3
- [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. 3
- [ZLCL21] ZHONG M., LI G., CHI P., LI Y.: HelpViz: Automatic generation of contextual visual mobile tutorials from text-based instructions. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2021), ACM, pp. 1144–1153. doi:10.1145/3472749.3474812. 2
- [ZZZ*24] ZHAO Y., ZHANG Y., ZHANG Y., ZHAO X., WANG J., SHAO Z., TURKAY C., CHEN S.: LEVA: Using large language models to enhance visual analytics. *IEEE Transactions on Visualization and Computer Graphics* (2024), 1–17. doi:10.1109/TVCG.2024.3368060. 3, 4, 5