

Touching or Chatting: The Utility of LLMs and Tactile Charts for Learning about Complex Chart Types by BLV Individuals

Tingying He , Maggie McCracken , Daniel Hajas , Sarah Creem-Regehr , Alexander Lex 




Figure 1: Modalities used to support blind and low-vision participants in learning chart types. We compared Tactile+Text+LLM and Text+LLM. (a) Text: A human-authored description of the chart type. In the tactile condition, it guided exploration of the example tactile chart and focused on explaining how the chart works; in the non-tactile condition, tactile-specific content was removed. (b) Tactile: A 3D-printed tactile model of a complex chart—a violin plot (shown) or a clustered heatmap—designed to support chart-type learning. (c) LLM: An LLM-based assistant that answered participants' questions about the chart type, chart, and data; the example shown is from an actual participant session.

Abstract—Visualizations are central to communicating data, yet blind and low-vision (BLV) people often lack support for understanding chart types—knowledge that is essential for interpreting new visualizations and collaborating with sighted peers. Prior work found that BLV individuals viewed example tactile charts as more helpful than text-only approaches and preferred them for learning advanced chart types, particularly for understanding spatial layouts and shapes. Meanwhile, large language models (LLMs) are increasingly used by BLV individuals for chart explanation and question answering (QA), but have been studied primarily for dataset exploration rather than chart-type learning. Existing LLM-based chart QA also shows that users frequently ask about layout and structure, yet struggle with spatial concepts and misdirect questions when mental models are weak. We investigate how LLMs influence chart-type learning and whether tactile learning improves subsequent LLM-supported exploration. We extend our tactile chart learning tools with an LLM chatbot that provides interactive explanations and supports follow-up questions. In an interview study with 12 BLV participants, we compare two learning formats: (1) a tactile chart, a textual explanation, and an LLM chatbot; and (2) a textual explanation and an LLM chatbot. The learning phase was followed by exploration of an unfamiliar dataset using alt text and an LLM. Thematic analysis shows that tactile templates support BLV participants' formation of chart-type mental models, which scaffolds subsequent LLM-mediated data exploration. Text+LLM explanations without tactile support show weaknesses for spatial-reasoning tasks.

Index Terms—Accessibility, tactile representations, LLMs.

1 INTRODUCTION

Visualization is central to scientific communication, decision-making, and public discourse. Accordingly, visualization literacy—the ability to interpret and extract information from charts [37]—is increasingly important for full participation in professional and civic life. People who are blind or have low vision (BLV) need to *read* visualizations [30, 36, 47], *collaborate with co-workers in analyzing* them, and

produce visualizations themselves [60]. This need is prevalent: approximately 7–8 million Americans live with vision impairment or blindness, and globally, an estimated 36 million people are blind, while 217 million have moderate-to-severe vision impairment [1]. Yet BLV users continue to face persistent access barriers when charts are delivered primarily through visual formats [30, 36, 47]. In practice, many BLV users rely on alternative text (alt text) to access charts. However, most alt text is effective only when readers already know the chart types and understand how the encodings work [30]. Shared chart-type understanding is also critical in blind-sighted collaboration, where teams might refer to the chart features during their communication.

Large language models (LLMs) and conversational AI systems are increasingly being explored as interfaces for supporting BLV individuals in accessing data and exploring charts [10, 24, 27, 45]. Recent accessibility research suggests that conversational support can help BLV users understand specific visualizations, but it also shows that chart-type knowledge remains important [24, 33, 45]. For example, prior work has observed that, when using chart question-answering systems, BLV users frequently ask about chart layout, try to understand unfamiliar chart types, and experience greater cognitive load when

- T. He (何汀滢), M. McCracken, S. Creem-Regehr are with the University of Utah, USA. E-mails: tingying.x.he@gmail.com, maggie.mccracken@sarah.creem@psych.utah.edu.
- D. Hajas is with the Global Disability Innovation Hub, UK. E-mail: d.hajas@ucl.ac.uk.
- A. Lex is with Graz University of Technology, Austria and the University of Utah, USA. E-Mail: alex@visdesignlab.net

This is the authors' preprint version of this paper. License: [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) (CC BY 4.0). Please cite: T. He, M. McCracken, D. Hajas, S. Creem-Regehr, A. Lex. *Touching or Chatting: The Utility of LLMs and Tactile Charts for Learning about Complex Chart Types by BLV Individuals*. *IEEE Transactions on Visualization and Computer Graphics*, to appear. 2027.

working with unfamiliar visualizations. These findings suggest that LLM-based chart exploration may be most useful when BLV users have a mental model of the chart type.

Although BLV participants commonly report being familiar with basic charts, such as bar, line, and pie charts, they report much lower familiarity with more advanced or complex chart types, such as area charts or violin plots [17, 18, 53]. A few studies have begun to investigate how to better support BLV individuals' understanding of advanced chart types, mainly focusing on text or tactile charts [26, 32, 41]. Our own prior work found that tactile example charts, particularly when paired with structured exploration instructions, can help BLV people build transferable knowledge of complex chart types [26]. These tactile charts are especially valuable for supporting BLV individuals in understanding the spatial layout and shape of chart elements that are difficult to communicate through text alone.

At the same time, LLMs offer a promising complement to tactile charts, because they can provide interactive, on-demand explanations through a digital interface. While the digital and flexible nature of LLMs makes them an invaluable resource for BLV users to understand specific datasets, it is less clear whether they are beneficial for learning about new chart types in the first place. Prior work on visualization education has explored LLMs for teaching chart types and improving visualization literacy with sighted users [7, 52]. Prior multimodal accessibility work has also established that hybrid techniques combining tactile with speech or sonification reduce workload and improve scalability for complex tasks [57]. Therefore, in this work, we combine tactile charts with exploration instructions and an LLM-based assistant to support BLV individuals in learning advanced chart types. Our goal is to understand the utility of tactile models for chart learning in the context of large language models: whether tactile models still provide unique benefits when LLM explanations are available, whether they are supplanted by LLM-generated textual explanations, or whether their value lies in complementing LLMs by grounding textual explanations in spatial experience. We also examine how chart-type learning shapes participants' subsequent exploration of new datasets in the same chart type with LLM support.

To investigate this question, we conducted an exploratory mixed-methods study with 12 BLV participants. We selected two complex chart types commonly used in scientific contexts, **clustered heatmaps** and **violin plots**. Participants learned two chart types, each under one of two conditions: (1) **tactile model + textual instruction + LLM-based assistant** or (2) **textual instruction + LLM-based assistant**. We used 3D-printed tactile example charts that we developed previously [26], which BLV participants had found to be well designed and preferable for supporting their chart learning, relative to human-authored textual descriptions. We also developed and iteratively refined a study website in collaboration with our blind co-author, where participants could read instructions and alt text and interact with an LLM-based assistant.

In each condition, participants first learned the chart type by reading the instructions, exploring the tactile chart when available, and asking the LLM assistant questions. The tactile charts were intended to support chart-type learning rather than communicate a specific dataset. Participants then explored a new dataset in the same chart type using a manually curated alt-text description and the LLM assistant. We analyzed queries, learning outcomes, and participants' thoughts and preferences on the helpfulness of different chart-type learning methods. We found that learning with tactile charts is perceived as helpful for subsequent data exploration, and participants prefer the combination of the three modalities (tactile chart + instructions + LLM).

The main contribution of this paper is the interview results from 12 BLV participants who learned and explored advanced charts using tactile charts and LLMs. These results show the perceived value of chart-type learning with tactile charts for supporting later exploration of new datasets with LLMs, reveal the benefits and limitations of LLM support in learning about new chart types and datasets, and highlight participants' preferences across learning modalities. We also provide a collection of LLM queries from BLV individuals during the learning and exploration of complex charts, along with insights into the patterns these queries reveal. Our results are useful in an educational context for

BLV individuals, highlighting that to teach about data visualizations, LLMs can supplement, but not replace means of communicating spatial data, such as tactile charts.

2 RELATED WORK

In this section, we review prior work on the accessibility of complex visualizations, chart-type learning and comprehension for BLV individuals, and the use of LLMs to support visualization accessibility.

2.1 Visualization Accessibility of Advanced Chart Types

Visualizations are a powerful medium for exploring data and communicating insights, yet their dependence on vision creates barriers for BLV individuals [8]. To address these barriers, researchers have examined what information BLV readers seek from charts [39], developed a range of sensory substitutions—including speech [39, 41, 59], haptic or tactile representations [19, 23, 56], sonification [11, 21, 28]—that translate data and visual structure of charts into non-visual modalities, and proposed authoring and evaluation guidelines for these modalities. However, most visualization accessibility research has historically focused on basic chart types (e.g., bar, line, pie charts) [34, 54].

Providing equal access to complex visualizations is a significant challenge, and research on making advanced chart forms more accessible is still under-represented [40], but is emerging, recognizing the need to make complex charts accessible for BLV users in an educational and professional context. Recent tools that make visualization specifications navigable via screen readers, such as Olli [4, 61] or Data Navigator [16], have incorporated more complex designs (e.g., multi-series and faceted charts, stacked charts, or maps). Multimodal data representation techniques (e.g., MAIDR [46]) are also expanding to advanced chart types (e.g., heatmaps). Beyond general-purpose systems, some individual advanced chart families have received attention, such as automated alt-text approaches for UpSet plots [41] and genomics visualizations [48], as well as work on tactile representations for network visualizations (e.g., TADA [58]) and accessibility of geospatial visualizations [38]. This work further contributes to making complex visualizations accessible by evaluating chart-type learning methods.

2.2 Chart-Type Understanding for BLV Individuals

Visualization literacy refers to the ability to read, interpret, and extract insights from visualizations [37], as well as create visualizations [22]. Prior work has identified barriers to visualization literacy by examining the mental models that underlie incorrect chart interpretations [44]. To better support BLV individuals in reducing these barriers, Zong et al. argue that researchers need to avoid visual-focused approaches, and instead further develop effective non-visual channels for data exploration [60]. However, many current accessibility approaches still depend heavily on visual representations, which makes chart-type understanding especially important. For example, prior studies show that chart-type understanding is essential for interpreting alt text [31], and that BLV users frequently ask chart-type questions when interacting with chart question-answering systems [33]. In addition, many professional tasks require blind-sighted collaboration. Therefore, chart-type understanding remains a critical skill for BLV individuals' full participation in professional and everyday contexts.

While BLV participants report familiarity with basic charts (e.g., bar, line, or pie charts), familiarity is much lower for more advanced or complex chart types (e.g., stacked bars, area charts, donut charts, violin plots) [17, 18, 53]. To bridge this gap, a few works study how to support BLV individuals in learning unfamiliar chart types. He and Yu [25] adopt multisensory learning strategies to support visualization literacy education for BLV children learning bar, line, and scatter plots. Kim et al. [32] propose an automatically generated explanation system spanning 50 chart specifications and compare pedagogical strategies (i.e., linking to familiar chart types vs. not, declarative vs. procedural framing, abstract vs. concrete descriptions) for explaining charts to BLV people. Smits et al. [49] investigate staged explanation to help BLV individuals learn unfamiliar genomics visualizations. Mei et al. [42] introduce Benthic, a hypergraph-based approach that represents charts and diagrams in perceptually congruent structures for screen reader

users. While not framed as instruction, it supports users in developing a clearer structural understanding of how a chart is represented. Finally, in our previous work, we studied the use of example tactile charts as learning materials for complex visualizations and reported that tactile examples paired with exploration instructions are a preferred and effective learning method compared to text-only explanations [26]. We found tactile models especially helpful in supporting BLV individuals to understand chart layout and the shape of chart elements; therefore, we continue to adopt this method in this study.

Researchers have explored LLMs to improve visualization literacy and help novice users understand unfamiliar chart types in the general population [7, 52]. However, to our knowledge, prior work has not specifically examined how LLMs can support BLV individuals in learning unfamiliar chart types. Our work addresses this gap by integrating LLM-based assistance into a BLV chart-learning workflow.

2.3 LLMs for Visualization Accessibility

Generative AI, particularly LLMs such as ChatGPT, has quickly become a common interface for information access, and BLV individuals are often eager to adopt these tools for everyday tasks. Recent HCI and accessibility studies have examined how BLV users use GenAI tools [2, 51], including preferences and challenges. Within visualization, one line of work uses LLMs to produce textual descriptions of charts. Earlier approaches to automated chart description largely relied on templates and rules [13, 14, 43], and a few recent works have begun using LLMs, which offer greater expressive flexibility [55]. Examples include DATATALES by Sultanum and Srinivasan [50], which supports authoring textual narratives grounded in a given visualization, and AutoVizuA11y by Duarte et al. [15], which uses an LLM to automatically generate descriptions of web-based visualizations and provide statistical insights derived from the underlying dataset. At the same time, McNutt et al. [41] found that LLM-generated textual descriptions are, in some cases, inferior to rules-based generation, especially when considering hallucinations.

Another line of work uses LLMs to support conversational chart question answering, enabling BLV users to query visualizations and iteratively explore data through dialogue. Kim et al. [33] provide foundational evidence for this direction through a Wizard-of-Oz study that analyzed BLV participants' queries about four visualizations. Seo et al. [45] integrate an LLM into a multimodal data representation framework and examine how BLV individuals interact with multimodal LLM assistance for chart exploration. Gorniak et al. [24] present VizAbility, an LLM-based conversational system for exploring common visualization types, including bar charts, line charts, scatterplots, and choropleth maps. Across these studies, researchers observed that BLV users frequently seek chart-type understanding and experience greater cognitive load when interpreting LLM responses about unfamiliar charts. Participants also frequently ask questions about the chart's spatial layout. These findings suggest that users may benefit from becoming familiar with chart types in advance and that this knowledge is transferable between specific datasets.

3 STUDY STIMULI

To prepare for our evaluation of the utility of LLMs and tactile charts for teaching chart types to BLV individuals, we developed the study stimuli in close collaboration with our blind co-author. We used two stimuli: a study website with an LLM-based chatbot, and a set of tactile charts. We describe these components in detail below.

3.1 Tactile Charts

We selected two chart types for our evaluation: a clustered heatmap and a violin plot, both advanced chart types commonly used in scientific reports. We designed the tactile chart models for these two chart types in our previous work [26]. These tactile models are example charts designed specifically for chart-type learning. We developed them in close collaboration with two BLV researchers and followed established best practices for tactile graphics. We also designed accompanying exploration instructions to guide learners' use of the models. These

materials can all be accessed on our website vdl.sci.utah.edu/tactile-charts/. In our previous study, participants responded positively to the design of these materials, and tactile charts paired with exploration instructions were described as a preferred and helpful way for BLV individuals to learn these complex chart types. An example of the violin-plot tactile chart is shown in Figure 1.

For this study, we thus followed the same logistical procedure as before. We 3D-printed a new set of tactile charts and attached the labels from the sighted version, as well as a QR code linking to the corresponding exploration instructions, to the back of each model.

3.2 Study Website and the LLM-based Chatbot

We implemented our online study website for the two study phases: *chart learning* with a simple dataset and *testing* with a more complex dataset that is not available as a tactile chart. The website is implemented with the reVISit framework [9]. The website is available at vdl.sci.utah.edu/tactile-and-LLM-revisit/; the code is open source and available on github.com/visdesignlab/tactile-and-LLM-revisit/. To support data collection for the evaluation study, the website also records audio, user interactions, user prompts, and chatbot responses.

For the chart-type learning phase, the website presents pre-authored instructions tailored either for use with the tactile chart or for text-only learning, depending on the condition (Fig. 1(a)). For the testing phase, the website presents pre-authored alt text for a new dataset using the chart type the participant has just learned. At each step of the study (each page of the website), participants can open a modal window to chat with an assistant, which is an LLM-based chatbot tailored to the current study state (Fig. 1(c)). We use the OpenAI Responses API and the GPT-5.2 model. Below, we describe the implementation details of this chatbot.

Access to Chart-Specific Materials The chatbot can access three types of chart-specific resources: (1) the chart image, (2) the underlying dataset, and (3) the learning materials. Chart images are stored as pre-uploaded PNG files in the OpenAI file system, which the model can consume as an image and leverage GPT-5.2's vision capability to interpret visual details when needed. The underlying dataset is loaded from a CSV file. Learning materials are provided as Markdown.

During the chart-type learning phase, the chatbot may request all three resources (image, dataset, and instructions). During the new-dataset exploration phase, the chatbot can request the chart image and the dataset. The specific image, dataset, and (when applicable) instructions returned are determined by the current study condition (i.e., chart type, learning modality, and study phase). To keep responses fast and to avoid unnecessary data transfer, we provide these resources only on demand via Responses API tool calls.

System Prompt The chatbot's behavior was controlled by a system prompt selected based on: (1) study phase (chart-type learning vs. new-dataset exploration), (2) learning modality during the learning phase (with vs. without the tactile chart), and (3) the current chart type. The selected template is then parameterized with chart-type-specific context so the assistant's guidance matches what the participant is currently learning or exploring.

The system prompt begins by defining the assistant's role and goal in the study. During the learning phase, it explicitly states which materials the participant has access to, so the assistant's guidance remains aligned with the assigned condition. Across all conditions, the prompt specifies which chart-specific artifacts the chatbot can access and enforces a gating policy: the assistant should answer from general knowledge by default and request chart-specific artifacts only when the participant explicitly asks for them or when the question clearly requires exact values or visual details; otherwise, it asks a brief clarifying question. The prompt also constrains response style (e.g., clarity and concision) to keep the interaction consistent and easy to follow.

Maintaining Conversation State We maintain continuity across turns using the Responses API's conversation-state mechanism. It keeps context from earlier turns, and thus allows the model to continue a multi-turn conversation without resending the full history.

Table 1: Participant demographics. IP = degree in progress; $\bar{X}$ is the mean; Y is the mode. Note that P6 and P8 were assigned a participant ID but did not end up participating in the study and are hence absent from the table.

pID	Age (Gender)	First Interview Chart (modality)	Data Intensive Career	Braille Years Exp.	Vision Loss Level	Onset of Vision Loss	Highest Level of Education	AI Tools Used
P1	29 (M)	Heatmap (tactile)	No	25	No residual vision	Blind since birth	Bachelor's	Gemini
P2	48 (M)	Violin (text)	Yes	7	No residual vision	Blind since birth	Master's	BeMyAI, ChatGPT, Gemini
P3	32 (M)	Violin (tactile)	Yes	15	Peripheral vision loss	Blind since birth	Associate's (IP)	BeMyAI
P4	40 (M)	Heatmap (text)	Yes	34	No residual vision	Blind since birth	Not disclosed	ChatGPT, Gemini
P5	34 (F)	Heatmap (tactile)	No	26	No residual vision	Blind since birth	Bachelor's (IP)	ChatGPT, Gemini
P7	23 (M)	Violin (tactile)	No	19	No residual vision	Blind since birth	Bachelor's	ChatGPT, Claude, Gemini
P9	44 (F)	Heatmap (tactile)	No	42	Light perception	Lost vision suddenly	Master's	BeMyAI, Gemini
P10	40 (F)	Violin (tactile)	No	35	No residual vision	Blind since birth	Bachelor's	Aira, ChatGPT, Copilot
P11	34 (F)	Violin (text)	Yes	26	No residual vision	Lost vision suddenly	Master's	ChatGPT, Gemini
P12	40 (M)	Heatmap (text)	Yes	30	No residual vision	Lost vision gradually	Master's	Gemini
P13	29 (M)	Violin (text)	Yes	25	Light perception	Blind since birth	Bachelor's	ChatGPT, Claude
P14	63 (F)	Heatmap (text)	Yes	55	Severe low vision	Blind since birth	Master's	ChatGPT, Gemini
$\bar{X} =$ 38 ± 11			7/12 Yes	$\bar{X} =$ 28 ± 13	$Y =$ No residual vision (8)	$Y =$ Blind since birth (9)	$Y =$ Master's (5)	$Y =$ Gemini (9)

Accessibility for Screen-Reader Users To make sure our study website is easy to use with screen readers, we tested and iterated with our blind co-author. Our final prototype has the following dedicated features: First, the interface uses visually hidden live regions to announce interaction state changes (e.g., “AI is thinking”) and to read out newly available assistant responses as they arrive. Second, to support efficient navigation of the transcript, we mark user messages with Heading 5 and assistant messages with Heading 6. These consistent structural landmarks enable quick navigation by headings without disrupting the page’s primary heading hierarchy, because H5/H6 are typically unused in interfaces. Third, each message is preceded by a screen-reader-only role label (i.e., “You said” and “Assistant said”) to clearly identify the speaker. Finally, we label all interactive controls with aria-labels and manage focus and scroll behavior so keyboard and screen-reader users can keep their place as new content appears.

4 EVALUATION STUDY

With our study website and tactile charts, we conducted interviews with 12 BLV participants to evaluate the use of tactile charts and an LLM for learning unfamiliar chart types and their impact on subsequent exploration of new datasets. The study was reviewed by the University of Utah IRB and approved as exempt from full-board review (№ IRB_00180924).

4.1 Study Design

Our study design and methodology followed our prior work [26, 41], which used a similar target group. We reused our established protocol of interviewing and providing tactile models [26], but restate our process here for convenience. We employed a mixed-design method with two factors: (1) chart type: violin plot or clustered heatmap, and (2) learning modality: a tactile chart with exploration instructions and an LLM-based assistant (Tactile+Text+LLM), or a textual explanation and an LLM-based assistant (Text+LLM). Each participant learned both chart types, with each chart type using a different learning modality. To minimize order effects, we counterbalanced across the four possible order combinations (2 chart-type orders × 2 learning-modality orders). We adopted this approach because we considered the expected interview length and recruitment challenges, while still allowing participants to experience both chart types and both learning modalities, thereby enabling comparison across modalities. We conducted semi-structured interviews led by two researchers [3]. The average length of interviews was 135 ± 25 min (SD). We compensated each participant with a \$100 Amazon gift card.

4.2 Study Material Preparation

We adopted the datasets, tactile charts, instructions, and alt text from our previous work [26]. Participants encountered two datasets for each chart type: a simple dataset for learning and a more complex dataset

for assessing learning outcomes. The datasets used in the violin plot condition were body mass of penguins (simple) [29] and the Human Development Index across five continents (complex). The datasets used in the clustered heatmap condition were movie genres and actors (simple) and agreements on sociocultural topics in European countries (complex) [20].

For the simple datasets, we created the tactile charts (see Appx. E or vdl.sci.utah.edu/tactile-charts/), wrote the corresponding instructions, and authored alt text (see Appx. B). For the complex datasets, we created the visualizations (see Appx. F) and authored alt text (see Appx. C). Following established guidelines [39, 41], the alt text included a short high-level summary followed by more detailed information, such as statistical details and trend analysis. In this study, because we wanted participants to have open questions for exploration with the LLM assistant, we retained only the high-level summary.

All textual materials and the LLM-based chatbot are available on the study website: vdl.sci.utah.edu/tactile-and-LLM-revisit/. We also prepared the interview questions, including questions about participants’ backgrounds, their understanding of the charts and datasets, and their comparisons of the learning modalities (see Appx. A).

4.3 Participants

We recruited participants by posting in relevant social media groups. To be eligible, participants had to be at least 18 years old, speak English, be legally blind (best-corrected visual acuity of 20/200 or worse, or a visual field ≤20° in the better-seeing eye), read Braille, and use a screen reader on their computer. We further required participants to be working or enrolled in school at least part-time, because our target context was professional data analysis. Participants also needed to have a U.S. mailing address to receive the tactile chart.

We first recruited one eligible participant who had taken part in our previous study and still had the tactile chart in hand for a pilot study. The pilot allowed us to verify the procedure, estimate session duration, and confirm the accessibility and usability of the study website.

We recruited 12 new eligible participants for the formal study. One participant (P6) dropped out on short notice due to a health issue. One participant (P8) attempted to participate but could not due to technical issues with Zoom, which we were unable to resolve within the allotted time. To maintain counterbalancing across conditions, we recruited two additional eligible participants. Ultimately, 12 participants completed the formal interview study, and their results were included in the following analysis. Both the pilot participant and P8, who attempted to participate, received the same compensation as other participants. Table 1 summarizes participants’ demographics, and full participant background information is provided in the supplementary materials.

Table 2: Rates of answer quality for chart understanding, new dataset understanding, and factual observations about the new dataset under the two learning modality conditions (N = 12).

	Training	Correct	Partially Correct	Wrong
Chart type understanding	Tactile+Text+LLM	30.00%	36.67%	33.33%
	Text+LLM	31.67%	45.00%	23.33%
New dataset understanding	Tactile+Text+LLM	62.50%	25.00%	12.50%
	Text+LLM	62.50%	25.00%	12.50%
		Good	OK	Bad
New dataset factual obs.	Tactile+Text+LLM	50.00%	27.78%	22.22%
	Text+LLM	44.44%	25.00%	30.56%

4.4 Procedure

In this section, we present the study procedure for each participant. Depending on the condition and phase, participants completed different activities: they read textual materials, including instructions and alt text; used the LLM-based chatbot on the study website; touched the tactile charts; and answered the study questions. Fig. 2 shows the study procedure. We also describe all study materials in Sec. 4.2 and provide the full set of study questions in Appx. A.

Pre-Interview We screened interested participants for eligibility via chat or email. For eligible participants, we shared the study procedure and consent process and scheduled a 2-hour online interview. Participants then reviewed and signed the consent form in REDCap. Based on each participant’s assigned condition, we sent the corresponding tactile model via mail [12] and provided participants with links to both the tactile-model instructions and the textual explanations for the two chart types ahead of time. We included these links both as QR codes in the package and in a follow-up email. Participants could explore the tactile model and instructions before the session, but we did not require this preparation.

Background Questions (13 ± 2 min) After participants agreed to audio and screen recording, we asked background questions about demographics, vision loss, prior experience with screen readers, Braille, tactile graphics, AI assistants, and familiarity with clustered heatmaps and violin plots. We then sent each participant a link to the study website via Zoom chat or email and asked them to share their screen.

In the rest of the session, we guided participants through the website page by page. Participants could take as much time as needed to read the current page and chat with the assistant. To keep the procedure consistent across participants, we did not allow participants to return to previous pages. We also instructed participants on the chatbot’s functionality to ensure they understood how to use it during the study.

Chart Type Instructions (18 ± 6 min) Participants first read an instruction page about the chart type. Depending on the assigned modality condition, they explored the tactile chart or learned from the text-only materials. Participants could open the chatbot and chat with the assistant at any time. We then asked five questions to assess chart-type understanding: three were shared across both chart types (e.g., “What types of data are best suited for visualization using a clustered heatmap / violin plot?”), and two were chart-specific (e.g., “If a violin’s shape has three peaks, what does that imply about the underlying data?”).

Simple Alt Text (14 ± 3 min) Next, participants read the alt text of the chart used in the instruction page. This step served as practice for using the LLM chatbot to explore a dataset in the chart type they had just learned. We encouraged participants to chat with the AI assistant as much as possible to learn more about the dataset, and we provided three example queries as inspiration. These example queries were adapted from Seo et al.’s study [45] and covered information in Levels 2 to 4 from Lundgard et al.’s semantic levels [39], as Level 1 information is already mostly covered in the alt text. We asked all participants to read the example queries, but we did not require them to ask those questions.

Complex Alt Text (23 ± 7 min) Next, participants read the alt text for a new, more complex dataset in the same chart type and used the LLM assistant to explore it. We asked them to imagine encountering this chart in everyday life and trying to understand it as fully as possible. We encouraged them to use the LLM for in-depth exploration. Afterward, they answered five questions assessing their understanding (e.g., “Imagine you have a friend who has never encountered this chart before. How would you describe this chart to them?”). Participants could take as much time as they needed for the exploration, but they were not allowed to read the alt text or use the LLM while answering these questions. We then asked participants to rate their understanding of the dataset, describe any difficulties they encountered, reflect on whether the assistant helped, and share any additional feedback about their experience with this chart type.

Comparative Questions (12 ± 5 min) Participants completed the procedure for both chart types. We then asked them to compare the two training formats (Tactile+Text+LLM vs. Text+LLM), state their preference, and explain their reasoning. We also asked participants to rate and justify: (1) the helpfulness of the tactile charts for chart-type learning, (2) whether tactile learning affected their subsequent new-dataset exploration, (3) the helpfulness of the LLM during learning and exploration, and (4) the perceived usefulness of each training format for BLV education. Finally, we invited any additional feedback.

4.5 Quantitative Results

In this section, we report our quantitative results, including our assessments of participants’ responses to questions about the chart types and datasets, as well as their subjective ratings of the learning modalities.

Understanding and Factual Questions Two researchers independently coded all participant responses and resolved disagreements through discussion. For questions asked after the instructions that address chart type understanding, we coded responses as correct, partially correct, or incorrect. For questions asked after new dataset exploration, responses about the overall understanding of the dataset were coded using the same categories. For open-ended questions about factual observations in the new datasets, such as numerical characteristics, trends, and interesting or surprising findings, responses were coded as good, OK, or bad. As shown in Table 2, Tactile+Text+LLM training did not improve performance on understanding questions, consistent with our prior results comparing pre-authored alt text and tactile charts [26]. For factual observations about the new datasets, responses in the Tactile+Text+LLM condition were slightly higher in quality (50.00% good vs. 44.44% in the Text+LLM condition).

Subjective Ratings Fig. 3 summarizes participants’ subjective ratings of the two learning modalities. Overall, perceived understanding of the new complex datasets was similar across conditions (Fig. 3(a)). For the violin plot, both conditions received the same mean rating ($M = 4.00$). For the clustered heatmap, participants reported higher perceived understanding after Tactile+Text+LLM training ($M = 4.33$) than after Text+LLM training ($M = 3.83$). Participants strongly preferred tactile-supported learning (Fig. 3(b)): 11 participants preferred Tactile+Text+LLM; one participant indicated that their preference depended on chart type complexity, and said tactile charts are more useful when learning more complex chart types; no participant preferred just Text+LLM. Participants also rated Tactile+Text+LLM as more useful for BLV education ($M = 4.83$) than Text+LLM ($M = 4.25$) (Fig. 3(c)). Consistent with these preferences, participants rated tactile charts as highly helpful for learning chart types ($M = 4.83$) and for subsequent new-dataset exploration with the LLM assistant ($M = 4.42$) (Fig. 3(d)). Participants viewed the LLM assistant positively, but rated it as less helpful for chart-type learning (Fig. 3(e), $M = 4.00$) than tactile charts (Fig. 3(d), $M = 4.83$). Participants rated the LLM assistant more helpful in exploring new datasets (Fig. 3(e), $M = 4.75$) than in chart type learning (Fig. 3(e), $M = 4.00$).

5 THEMATIC ANALYSIS

We conducted a thematic analysis of the corrected transcripts [5, 6, 35]. The first author iteratively coded the data and developed themes through

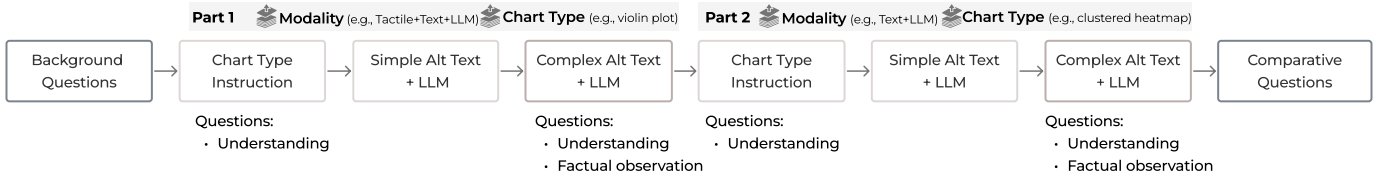


Figure 2: Study procedure. Participants first answered background questions, then completed two parts with different modality/chart type combinations. In each part, they learned the chart type with or without a tactile chart, practiced exploring a simple dataset with alt text and the LLM assistant, and then explored a complex dataset. Participants answered questions after chart-type learning and complex-dataset exploration. The study ended with comparative questions.




Figure 3: Participants prefer to learn chart types when all modalities (tactile models, LLM assistants, and text descriptions) are available. (a), (c), (d), and (e): Distributions of participants’ responses to different questions on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree). (b): Participants’ preferred learning modality.

repeated engagement with the transcripts. Emerging codes and themes were discussed and refined with the research team. Below, we report the resulting themes with representative participant quotes. For readability, we removed filler words such as “like” and “you know” from quotes without changing the meaning of the sentences. We also incorporated the analysis of LLM queries collected during the chart-type learning session into our thematic analysis, and report our analysis of queries generated while participants explored new datasets in Appx. D.

5.1 Chart-Type Understanding as a Prerequisite for Meaningful Exploration

BLV participants described chart type understanding as a foundation for LLM-supported exploration.

Without an initial mental model of how an unfamiliar chart type is structured, participants found it difficult to interpret new datasets, ask questions of the LLM assistant and benefit from its responses. This finding highlights the importance of teaching participants unfamiliar complex chart types.

Unfamiliar chart types are difficult to approach directly. Participants described exploring new datasets in unfamiliar chart types as cognitively demanding, and this challenge appeared across participants with different vision histories. For example, P12, who gradually lost their vision, described dataset exploration without prior tactile learning as “The picture in your head may be a little bit difficult if this was a new concept”, and “if you were just kind of fresh and raw and green with it, then it, it could be very overwhelming very quickly for someone that’s not familiar with how [the] data set and things are laid out.” Similarly, P5, who has been blind since birth, noted that some concepts can be

especially difficult to grasp, saying that “blind people, especially people who haven’t seen at all before, struggle with the concept of shading, and densely populated [areas], or what those terms mean in a chart”. These comments suggest that before BLV individuals can effectively interpret data shown in complex visualizations, they first need access to the chart’s conceptual logic.

Knowing what to ask is essential for benefiting from AI. Especially in the context of using LLMs, participants also emphasized that LLM assistance was more useful for chart exploration when they already had some sense of what kinds of questions the chart could support. Understanding the chart type was often the foundation for asking good questions about the chart. P1 stated, “A better understanding when you ask good questions”, and emphasized, “so basically, know what you’re gonna get out of it before you sit down and use the AI. But if you don’t know what questions [to ask], you’re gonna be lost.” P14 also emphasized the importance of asking the correct questions, noting, “I had to really think about how I wanted to ask the question so that the AI had an understanding of what I was looking for.” Similarly, P4 noted that the LLM assistant would be less useful for someone who “doesn’t have more acquaintance with the chart”.

5.2 Tactile Learning Builds Mental Models of Chart Types That Support Subsequent LLM Use

Our prior work [26] showed that tactile charts can help BLV individuals learn unfamiliar chart types. This study confirms and extends that finding: participants reported that prior tactile learning helped them explore a new dataset using the same chart type with an LLM assistant.

BLV participants described the tactile learning experience as helpful for asking more targeted questions and better interpreting the LLM-based assistant’s responses.

Tactile learning gave participants a chart template they could carry into later exploration. Participants repeatedly described the tactile chart as providing a general sense of how the chart type works, which they could reuse when encountering a new dataset in the same format. For example, P1 explained that the tactile chart “basically helps your memory develop more—what is this chart about, what are we doing with it? And having the tactile thing, it kind of drills it into your memory as well, when it comes to cognitive processes.” P12 made a similar point, explaining that “having a tactile [chart] ahead of time, feeling it, I had a lot more [...] understanding of what it was going to look like, as opposed to the heat map [the condition in which the participant didn’t have a tactile chart]. Having that tactile image gave me kind of a leg up on the violin versus the heat map.” We can also see this pattern from participants with lower tactile literacy. For example, P9, who rated her familiarity with tactile charts or graphics as 2 out of 5, said that after learning with the tactile chart, she could approach the new dataset with the expectation that “If it was laid out, it would be in the same way.”

For some participants, this prior tactile experience was especially important for making sense of chart features that were otherwise difficult to picture. P12 described the violin plot as “a foreign concept”, and said of tactile learning, “it did help in that next example to know, okay, this is what a violin model looks like”. Similarly, P14 said that learning with the physical violin plot helped: “so I had the understanding of the

distribution and how it would appear". These comments suggest that tactile learning gave participants a more stable structural expectation before they entered LLM-supported data exploration.

The tactile chart template supported more targeted question asking during LLM-based data exploration. As discussed above, knowing what to ask is important for making effective use of the LLM. Here, we focus on participants' feedback on how tactile learning helped them ask more useful questions when later exploring a new dataset with the assistant. This point is important because conversational exploration unfolds sequentially: what participants can get from the data depends in part on what they choose to ask. P4 explained this relationship especially clearly. After describing how the physical model helped him understand the chart's structure and purpose, P4 said, "*Once you look into the charts physically, then depending on the explanation given from the AI, you can develop more questions. Because you have the understanding of the texture, you can already come up with some questions.*" He also described an example scenario to show how tactile learning supported question formulation by giving him an internal reference for what features to probe further: After learning from the gentoo penguin's violin plot that a bimodal shape could reflect two subgroups, P4 felt able to ask targeted questions when encountering a new violin plot about which subgroups the peaks of the violin represented.

Other participants made similar connections between prior tactile learning and later question formulation with specific reasons. P7 noted that the tactile violin plot helped him understand the central purpose of the chart, explaining that "*It imprints in your mind [...] the concept [of a] distribution; that the main thing is the fat distribution, and the median.*" When exploring the datasets with violin plots, he asked about these, e.g., "*give me a broad overview and trends of the spread in all continents*", "*take only the lowest half of each continent, then tell me the median*". P2 highlighted tactile learning as supporting anticipation during iterative exploration. He explained that "*having it there as a reference, gave you a better idea of what to expect, or to anticipate*". This reference also helped P2 feel more comfortable asking questions during subsequent exploration, as he noted: "*I think I felt more comfortable asking questions, in the second [heatmap] example than I did in the second violin plot example, so I would say absolutely it did [...] help.*"

P9's experience further illustrates that this benefit could also appear for a participant with lower tactile-chart familiarity. P9 said that tactile learning "*helped me to be able to understand what to ask, and what kind of things—statistics—could be shown.*" P9 did ask more specific questions related to statistics in clustered heatmaps (learned with tactile), e.g., extreme values, averages, and ranges in the heatmap, while in the violin plot, she only asked broad questions (e.g., asked for overall patterns, context) or non-targeted questions for information that is beyond the dataset.

Tactile learning helped participants make sense of the assistant's responses. Beyond question asking, participants described the earlier tactile chart as a reference for understanding what the assistant said about the new dataset. P9 summarized this effect directly: "*I kind of knew what kind of things to be listening for.*" P2 made a closely related point, explaining that "*having it there as a reference, gave you a better idea of what to expect, or to anticipate with, with the next [dataset]*", and concluded that "*I think that did assist with assimilating the information more coherently.*" P14 likewise described the first tactile example as helping her interpret the second dataset because "*the first one helped me to explore the second [dataset], in that I had the concepts to work with. I had the previous experience of the concepts that I could then integrate and match up.*"

5.3 Using LLMs for Chart Type Learning

In this section, we discuss participants' opinions on the benefits and limitations of the LLM assistant during chart-type learning.

BLV participants valued LLM-based assistants for supporting interactive explanation, but these assistants could not replace tactile charts for spatial understanding.

5.3.1 Benefits of LLM Assistants

LLMs supported a progressive, learner-driven process of clarification and elaboration. Participants described the value of the LLM assistant in terms of its conversational nature. Rather than delivering a fixed explanation, it allowed participants to begin with a partial understanding and then expand it through follow-up questions. As P5 said, "*I think you'll get a general idea, and then using the AI kind of, like, expands everything.*" P2, as an educator, tried to teach his students to use AI in their daily toolkit, because he believes "*it has the potential to be very explanatory and to get granular when it needs to.*"

Participants particularly valued the LLM because it could clarify ideas when the provided materials were insufficient for learners' needs. For example, P1, who learned the violin plot in the Text+LLM condition, asked, "*Can you give me a better explanation of the chart?*" The chatbot responded, "*You can think of it as a mix of a boxplot (summary) and a smoothed histogram (shape).*" This analogy was not included in our pre-authored instructions. P1 found it helpful and commented, "*Okay, I like the fact that when explaining with the AI, I didn't know the shape before, but now I was able to figure, oh, a histogram, okay, I've seen this before, so it kind of helps me more of explain, visualize in my head, okay, I know what it looks like now.*" Similarly, P7 valued the LLM's ability to provide personalized explanations, stating, "*The AI is very useful because you could basically ask it anything, and it can rephrase and reword things that the instruction may get funny, and you can ask it for additional context*". He then further explained, "*You can add anything, you can come at it from any angle. I think AI is more responsive and more, you know, it's like a back and forth, so you can take your own angle from it, which is good.*"

In this sense, the LLM could complement the pre-authored materials (tactile charts and instructions) by helping participants address their personal confusion, which can be highly diverse and difficult for the pre-authored materials to cover. P1's interaction with an LLM in the tactile clustered heatmap condition illustrates this point. After asking why color was used to represent actors and requesting a clearer explanation, P1 continued by suggesting that textures might work better and explained that this would help "*spot high and low values quicker*". After these turns of conversation with the assistant, P1 positively reflected on the interaction, saying, "*I actually like it like this. They explain more of the chart with the AI when it comes to doing things like this, and it helps you more, okay, cool, it describes it better that way when you actually were able to ask it more questions like this.*" This learner-driven clarification may be especially useful for explaining complex chart components, such as dendrograms in clustered heatmaps. During learning, four participants asked dendrogram-related questions, in both the tactile (P5, P13) and non-tactile (P7, P14) conditions. Related queries included "*Can you explain how the lines on the right side of the chart work?*" and "*does the dendrogram show relationships that are connected?*"

LLMs could guide learning by suggesting next steps. Participants also noted that the LLM often supported learning by prompting the next step in the interaction. Some follow-ups were directly prompted by the assistant, for example when it offered next steps and the participant chose one (e.g., the assistant ended their response with "*If you tell me what you want to focus on—median, spread, or comparing groups—I can explain that part with a quick step-by-step reading strategy.*" and the participant accepted the offer). Others were clarification questions, such as "*Do you mean the average for each species, or the overall average across all penguins in the CSV?*" Participants also initiated their own follow-ups by asking deeper questions based on earlier answers, such as continuing to probe how to interpret the dendrogram (P13's example above). This LLM behavior was observed in both learning and dataset exploration phases, but we considered it as part of the learning process and discussed it here because it helped participants deepen their understanding of the chart.

P2 described this guided interaction process: "*I was basically letting it lead me. Generally, in my interaction, what I've discovered is that it helps to have at least a general idea of what you're looking for, but then you can, based on what it tells you, become more granular in your analysis. You know, if you just treat it as organized brainstorming,*"

essentially.” P12 also described the LLM’s prompts as invitations to dig deeper: “There was a couple [of instances] where it prompted me: Hey, do you want to know more? And I was like, yeah, tell me more, yes. [...] It’s like] digging down into that rabbit hole of Wikipedia, where you can click on one link, and then it’s gonna take you to a whole new thing.” P14 further emphasized that the assistant’s response could redirect attention toward issues they had not initially considered: “So, for example, I could ask a specific question, and then his response would guide me to maybe an area I hadn’t thought of before, or, would provide a suggestion which would extend [...] my understanding, because it would go and describe [it] in detail.” P14 noted that “having AI almost be like a person you’re having a conversation, and rather than having that person look at you kind of confused and being visually impaired, not understanding that that’s what that meant, the AI is giving me the information: Hey what you asked isn’t clear, so I need clarification, and is this what you’re thinking about? And he gives examples, and then I could turn around and say, no, that’s not exactly what I mean. So, it’s like an ongoing conversation to broaden the understanding and to enrich the information.”

5.3.2 Issues with LLM Assistants

Despite the benefits participants identified, they also described several limitations of the LLM assistant for chart-type learning.

LLM explanations make it difficult to understand spatial structure. Like textual explanations, the LLM assistant could not fully convey spatial form and shape. Several participants in the non-tactile condition therefore expressed a need for tactile support. For example, P13, who learned the violin plot in Text+LLM, said while exploring the new dataset: “I struggled to get an actual picture of how many curves, I suppose, the violin has. The alt text gave an average, and the AI wasn’t really able to say, oh, there are 10 points at which the violin, you know, expands on Asia or Africa. So I think for that, I would need to physically see it.”

Effective use of LLMs required some AI literacy and prompting skill. Participants also emphasized that the usefulness of the LLM assistant depended partly on knowing how to interact with it. P2 described this as a learning curve, explaining, “there are the natural limitations that occur when there is learning how to ask the questions in the right way. Sometimes you get the old equation: garbage in, garbage out. If I don’t ask the question the right way, I’m not going to get an answer that is coherent. There’s a small learning curve.” At the same time, P2 is positive about LLMs, saying, “I think that that learning curve, every year, is going to become exponentially smaller as people become more accustomed to working with AI, looking over their shoulder.” P2 noted that the training in this study was sufficient: “I think by the time you get to the second data set, [...] most people will have become sufficiently comfortable that they’re better able to come out of the gate asking the right questions, or at least questions that are better designed to elicit a more coherent response from the chatbot.” P4 shared this opinion, “[...] The helpfulness of the AI really differs. It depends on the skill of the individual on how to use it, and on how they are familiar [...]. The more familiar you are with the AI and with your tools, you can ask more questions, you can get more help. If you are less familiar, it might not even be helpful.” P4 also believed that users’ educational backgrounds could affect their ability to formulate effective prompts and, consequently, obtain useful responses.

Participants worried about trust and information overload. Participants identified two broader challenges in using LLMs for chart learning: the risk of inaccurate responses and the difficulty of processing too much numeric information. These concerns appeared mainly during dataset exploration, but also surfaced in chart-type learning. For example, P7 emphasized that the assistant is most useful when its answers are grounded in the dataset and can be trusted, saying, “if it’s trained with a specific data set like this. I think it is especially useful, because it doesn’t give you random answers, it is already trained, and as long as we can verify that it’s not hallucinating, I think it is extremely useful when it gets to the point.” Several participants also noted that responses with too many numbers were hard to follow. For example, P5 said she would prefer the LLM to “figure out how to consolidate

some of the information so it’s not saying the numbers so rapidly, and, one after the other.”

LLM assistants need to provide more adaptive scaffolding that helps learners recognize and overcome gaps in understanding. Participants noted that LLM assistants need to be more sensitive to learners’ prior knowledge. P4 suggested that, ideally, AI systems should detect knowledge gaps and adjust the depth and pace of explanations. P14 suggested, “The chatbot shouldn’t just have assumptions that, when you go into this, you already know A, and so, it can jump to B. It might need to [...] check for knowledge of A before jumping ahead to B. [...] If I don’t have that [...] the chat bot [should] be able to modify its response to reflect that.” P4 later clarified that he meant both experience with the chart types and with using AI assistants.

Beyond modeling learners’ knowledge, we also observed the need for LLMs to help learners recognize what they do not yet understand. Because learning with an LLM is largely learner-driven, novices may struggle to identify their own knowledge gaps and propose appropriate questions. Therefore, LLM assistants should provide stronger scaffolding. In the chart learning phase, we observed that four participants (P3, P4, P11, P12) did not ask any questions while learning either chart, and five participants (P1, P5, P10, P13, P14) asked questions for only one of the two charts. Most of these participants reported that the provided instructions were sufficiently clear. However, their later responses to the comprehension questions were not always correct, which suggests that their understanding of the chart types could still be improved. P10’s experience illustrates this challenge clearly. Although she found the violin plot difficult to understand, she explained that she did not ask the assistant questions “because I don’t know what to ask it”. She further noted that, even with a human instructor, it would still have been difficult to “figured out how to even know what to ask”. For such learners, the LLM assistant may need to take a more proactive role by offering possible questions related to this chart type, or asking low-stakes diagnostic questions to check learners’ comprehension.

LLM assistants also need to better detect the underlying reason for a learner’s confusion from their questions, and provide explanations that directly target it. P13’s experience with dendrograms was particularly illustrative. He first asked, “Do all heatmaps have dendrograms?”, then followed with more specific questions about interpreting the structure: “It’s definitely clustered. Do the distances between dendrograms also provide information, or are they separated linearly?”, “I see several areas where the dendrograms don’t join with a row. Is that normal? There are other places where two lines are obviously connected,” and “Let’s look at Dwayne. His dendrogram extends far out and appears to eventually make contact with almost everyone else—save for Julia Roberts. Why?” However, even after reading the AI’s responses, P13 still did not understand the hierarchical structure. He explained, “I’m not exactly clear on why it doesn’t look like some of these connected, because my understanding would be you should have a connection from every row to every other row.” The interviewers then explained that the dendrogram represents stepwise grouping: the two most similar actors first form small groups (Julia Roberts and Tom Hanks, Jennifer Lawrence and Leonardo DiCaprio), which then merge into larger groups. P13 immediately reformulated this correctly: “And then Dwayne connects to the group of all four, because he’s so far out there,” and concluded, “That makes sense. Now I think I see what I was missing.” In this case, the LLM did not fully resolve P13’s confusion, likely because it provided too many technical details about hierarchical clustering and multiple possibilities without addressing his core misunderstanding: that dendrograms show groups merging step by step, rather than each row directly connecting to others.

5.4 Multimodal Chart-Type Learning

In our prior work [26], participants preferred a combination of tactile charts with structured exploration instructions for chart-type learning over text-only learning. In this study, we asked whether an LLM assistant, with its interactivity and flexibility, can compensate for the shortcomings of static text, leading participants to prefer it over tactile charts. We found that participants did not consider the assistant as a replacement for tactile charts. Instead, they described the best learning

experience as a combination of tactile charts, exploration instructions, and the LLM assistant.

BLV participants preferred a multimodal learning experience for chart type learning, because tactile charts, text, and LLMs provided complementary forms of support.

Participants preferred multimodal learning because each modality contributed a different kind of understanding. P2 highlighted the importance of reinforcement across modalities, saying that “*using as many different senses as possible to aid in comprehension is always, I think, a bonus.*” P14 also noted, “*Just the AI is just an auditory/visual modality. I don’t really get to incorporate multiple senses.*” She described the combined setting as: “*Three different ways of looking at the information, and each one being its own means of providing a different or adding dimension to it.*” P7 emphasized the value of multimodal learning because of the needs of different people, saying, “*I think people learn in different ways. I think for a lot of blind people, tactile learning is probably one of the best, of course, and AI is interactive, it’s good!*” P4 further shared his strategies for using these modalities together: “*Usually, for me, I’d love to have the physical chart, and then go to the AI whenever I need more explanation, and I believe that more work can be done on the physical chart as well.*” P7 also envisioned a more interactive design to combine these modalities: “*AI is connected to the refreshable graphics board, and that graphics board can change based on AI, and AI can draw the line for you from place to place if you ask. If you ask about the trend of, [or] about the peak of something. It can have a moving dot there to show your finger where it is, and maybe even show a number there to describe to you numbers.*”

When tactile support is unavailable, LLMs can partially compensate for its absence and provide greater support than text alone. Although participants preferred multimodal learning, they also recognized the practical value of LLM assistance. P7 framed this explicitly as a real-world constraint, noting that “*In real life, it’s going to be a while until we can do the [3D] print on a whim. So, I think a lot of blind people, including myself, are going to be quite happy with text description and an AI chatbot, if we can make it this accurate.*” He further stated that the combination of text and LLM “*is better than text only.*” P1’s experience in the Text+LLM condition shows how the assistant could partly compensate for missing tactile support: as we discussed in Sec. 5.3.1, the LLM’s explanation that the violin plot is similar to a histogram helped him understand the chart.

6 DISCUSSION

Our findings suggest that BLV participants preferred a multimodal learning experience: across the interviews and subjective ratings, participants valued Tactile+Text+LLM, because each modality contributed differently to the learning process. Tactile charts were especially important for conveying spatial information, which is difficult to communicate through other modalities. This spatial understanding appeared to provide an important foundation for subsequent data exploration. The textual introduction helps participants understand the tactile model and sets up the dataset. LLMs, in turn, were valued for their flexibility. Compared with static descriptions such as alt text or instructions, LLMs can respond to a learner’s immediate confusion, elaborate on partially understood concepts, support exploration through dialogue, rely on “world knowledge” for broader questions, and answer queries based on patterns in the data that may not be encoded in a visualization. This flexibility seems especially useful once learners have formed an initial mental model and want to refine or extend their understanding.

The value of LLM support is also reflected in comparison with our previous work [26]. Although the two studies involved different participants, both targeted the same population and used the same chart types and new datasets. Participants in the present study reported higher perceived understanding of the complex datasets, with mean ratings ranging from 3.83 to 4.33 under the Tactile+Text+LLM and Text+LLM conditions, compared with 3.50 to 3.83 under the Tactile+Text and Text-Only conditions in the prior work. Similarly, perceived usefulness in education was higher for Text+LLM ($M = 4.25$) than for Text-Only ($M = 4.00$).

In addition, LLMs have the practical advantage of being immediately available as a digital support tool and, obviously, have enormous flexibility across a wide range of contexts. However, they still cannot convey what a chart feels or looks like. Therefore, we suggest using multimodal systems to support BLV chart learning, especially when formally educating BLV individuals about data and data visualization, such as in college classes. These modalities should be closely connected. In particular, LLM support for chart learning should be grounded in both the tactile chart and the underlying data, so learners can smoothly move between physical exploration and conversational explanation.

Although participants’ answers to factual observation questions, their subjective ratings, and their interview comments all pointed in a more positive direction, our quantitative accuracy measures did not show an improvement for Tactile+Text+LLM on chart understanding questions. This mismatch between perceived value and measured performance could have several potential reasons. First, the effect of tactile chart type learning on subsequent data exploration may be difficult to detect with our small sample size. Second, the chart type learning phase was relatively short, lasting 18 ± 6 minutes. It is possible that short-term chart type learning does not strongly transfer to measurable performance improvements during data exploration. Future work could examine this possibility through longer-term studies, for example in collaboration with BLV educational organizations. Third, the questions we asked may not have been the most suitable measures for capturing learning outcomes or participants’ mental models of the chart type. This limitation points to an important methodological opportunity for future work: to develop evaluation methods that better capture how BLV individuals benefit from visualizations and non-visual representations in supporting data analysis.

Our findings also suggest that the LLM assistants in this study may not always be sensitive enough to learners’ prior knowledge. For example, when explaining complex components such as dendrograms, the LLM became lost in technical details, whereas a human interviewer could intervene more effectively because they were better able to recognize the source of the difficulty for participants in understanding the chart. In this study, we intentionally limited customization mostly to the system-prompt level in order to approximate how LLMs might be used in practice. We did not provide explicit instruction-customization options, as prior work showed that participants with similar backgrounds rarely customized the assistant themselves [45]. However, our findings suggest that future LLM systems for BLV chart learning could benefit from more adaptive behavior, such as detecting knowledge gaps during conversation, accurately identifying where participants are getting stuck, and adjusting explanations accordingly.

7 CONCLUSION AND FUTURE WORK

Our thematic analysis suggests that tactile charts and LLMs complement each other in supporting BLV individuals’ learning of unfamiliar complex chart types: tactile charts provide essential spatial understanding, while LLMs support flexible clarification and follow-up exploration. This combination points to a promising multimodal approach for improving chart learning and thereby enabling better participation in data analysis by BLV individuals.

Our study focused on short-term learning in a controlled setting rather than long-term learning in real educational contexts. However, participants recognized the potential value of these learning modalities in real educational settings. As one participant noted about Tactile+Text+LLM, “*I wish I had this when I was in school, that would have really helped out.*” This response suggests promising directions for future research on applying these learning formats in real educational settings, including schools for BLV students and other real-world learning environments.

Participants also expressed interest in making tactile access more adaptive. This points to promising opportunities for future work on refreshable tactile displays combined with LLM assistants, which could provide a more responsive and interactive learning experience than static tactile materials, although cost and access remain important practical considerations.

ACKNOWLEDGMENTS

The authors thank Aaron Delahunta and Sofia Djunic for correcting interview transcripts.

REFERENCES

- [1] P. Ackland, S. Resnikoff, and R. Bourne. World blindness and visual impairment: despite many successes, the problem is growing. *Community Eye Health*, 30(100):71–73, 2017. 1
- [2] R. Adnin and M. Das. “I look at it as the king of knowledge”: How blind people use and understand generative AI tools. In *Proc. ASSETS*, art. no. 64, 14 pp. ACM, 2024. doi: 10/qzfk 3
- [3] D. Akbaba and M. Meyer. “Two heads are better than one”: Pair-interviews for visualization. In *Short Paper Proc. VIS*, pp. 206–210. IEEE, 2023. doi: 10/gtcvp3 4
- [4] M. Blanco, J. Zong, and A. Satyanarayan. Olli: An extensible visualization library for screen reader accessibility. In *Posters Proc. VIS*, 2022. 2
- [5] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qual Res Psychol*, 3(2):77–101, 2006. doi: 10/fswdcx 5
- [6] V. Braun and V. Clarke. Toward good practice in thematic analysis: Avoiding common problems and be(com)ing a knowing researcher. *Int J Transgend Health*, 24(1):1–6, 2023. doi: 10/gq658d 5
- [7] K. Choe, C. Lee, S. Lee, J. Song, A. Cho, N. W. Kim et al. Enhancing data literacy on-demand: LLMs as guides for novices in chart interpretation. *IEEE Trans Vis Comput Graph*, 31(9):4712–4727, 2025. doi: 10/n98f 2, 3
- [8] J. Choi, S. Jung, D. G. Park, J. Choo, and N. Elmqvist. Visualizing for the non-visual: Enabling the visually impaired to use visualization. *Comput Graph Forum*, 38(3):249–260, 2019. doi: 10/gh525h 2
- [9] Z. Cutler, J. Wilburn, H. Shrestha, Y. Ding, B. Bollen, K. A. Nadib et al. ReVISit 2: A full experiment life cycle user study framework. *IEEE Trans Vis Comput Graph*, 32(1):13–23, 2026. doi: 10/hbkxwp 3
- [10] A. K. Das, M. Tarun, and K. Mueller. Making charts speak: LLM-based conversational chart question answering for blind and low-vision users. In *Proc. CHI Extended Abstracts*, art. no. 446, 5 pp. ACM, New York, NY, USA, 2026. doi: 10/rb9j 1
- [11] G. Daunys and V. Lauruska. Sonification system of maps for blind. In *Adv Hum Comput Interact*, chap. 16. IntechOpen, Rijeka, 2008. doi: 10/g9bdwv 2
- [12] L. de Greef, D. Moritz, and C. Bennett. Interdependent variables: Remotely designing tactile graphics for an accessible workflow. In *Proc. ASSETS*, art. no. 36, 6 pp. ACM, New York, NY, USA, 2021. doi: 10/psz6 5
- [13] S. Demir, S. Carberry, and K. F. McCoy. Generating textual summaries of bar charts. In *Proc. INLG*, pp. 7–15. ACL, Salt Fork, OH, USA, 2008. doi: 10/cd7k7q 3
- [14] S. Demir, D. Oliver, E. Schwartz, S. Elzer, S. Carberry, and K. F. McCoy. Interactive SIGHT into information graphics. In *Proc. W4A*, art. no. 16, 2010. doi: 10/b3wpsb 3
- [15] D. Duarte, R. Costa, P. Bizarro, and C. Duarte. AutoVizuA1ly: A tool to automate screen reader accessibility in charts. *Comput Graph Forum*, 43(3):e15099, 2024. doi: 10/qvt4 3
- [16] F. Elavsky, L. Nadolskis, and D. Moritz. Data navigator: An accessibility-centered data navigation toolkit. *IEEE Trans Vis Comput Graph*, 30(1):803–813, 2024. doi: 10/g9bdw3 2
- [17] C. Engel and G. Weber. Improve the accessibility of tactile charts. In *Proc. INTERACT*, pp. 187–195. Springer, 2017. doi: 10/g8q2w8 2
- [18] C. Engel and G. Weber. A user study to evaluate tactile charts with blind and visually impaired people. In *Proc. ICCHP*, pp. 177–184. Springer, 2018. doi: 10/g8q2xc 2
- [19] C. Engel and G. Weber. User study: A detailed view on the effectiveness and design of tactile charts. In *Proc. INTERACT*, pp. 63–82. Springer, 2019. doi: 10/g8q2xb 2
- [20] EVS. European values study longitudinal data file 1981–2008. GESIS Data Archive, Cologne. ZA4804 Data file Version 3.1.0, 2020. doi: 10/pt4j 4
- [21] K. Franklin and J. Roberts. Pie chart sonification. In *Proc. IV*, pp. 4–9. IEEE, 2003. doi: 10/bq9rz 2
- [22] L. W. Ge, Y. Cui, and M. Kay. AVEC: An assessment of visual encoding ability in visualization construction. In *Proc. CHI*, pp. 1–16. ACM, New York, NY, USA, 2025. doi: 10/qz dq 2
- [23] C. Goncu and K. Marriott. GraVVITAS: Generic multi-touch presentation of accessible graphics. In *Proc. INTERACT*, pp. 30–48. Springer, 2011. doi: 10/bbns6v 2
- [24] J. Gorniak, Y. Kim, D. Wei, and N. W. Kim. VizAbility: Enhancing chart accessibility with LLM-based conversational interaction. In *Proc. UIST*, art. no. 89, 19 pp. ACM, New York, NY, USA, 2024. doi: 10/g8nwhp 1, 3
- [25] S. He and L. Yu. Charting beyond sight with DataStory: Sensory substitution and storytelling in visual literacy education for visually impaired children. In *Proc. CHI Extended Abstracts*, art. no. 73, 8 pp. ACM, New York, 2024. doi: 10/gt4w3w 2
- [26] T. He, M. McCracken, D. Hajas, S. Creem-Regehr, and A. Lex. Using tactile charts to support comprehension and learning of complex visualizations for blind and low-vision individuals. *IEEE Trans Vis Comput Graph*, 32(1):199–209, 2026. doi: 10/hbdjdj 2, 3, 4, 5, 6, 8, 9
- [27] C. Hohenwalde and A. Hazim. Enhancing graphic accessibility for blind and visually impaired people: A systematic comparison of vision language models (VLM). In *Proc. INFORMATIK*, p. 613. Gesellschaft für Informatik e.V., Bonn, 2025. doi: 10/rb9k 1
- [28] M. N. Hoque, M. Ehtesham-UI-Haque, N. Elmqvist, and S. M. Billah. Accessible data representation with natural sound. In *Proc. CHI*, art. no. 826, 19 pp. ACM, 2023. doi: 10/g9bdwz 2
- [29] A. M. Horst, A. Presmanes Hill, and K. B. Gorman. *palmerpenguins: Palmer Archipelago (Antarctica) Penguin Data*, 2020. R package version 0.1.0. doi: 10/qzfm 4
- [30] S. C. S. Joyner, A. Riegelhuth, K. Garrity, Y.-S. Kim, and N. W. Kim. Visualization accessibility in the wild: Challenges faced by visualization designers. In *Proc. CHI*, art. no. 83, 19 pp. ACM, 2022. doi: 10/g8q2xh 1
- [31] C. Jung, S. Mehta, A. Kulkarni, Y. Zhao, and Y.-S. Kim. Communicating visualizations without visuals: Investigation of visualization alternative text for people with visual impairments. *IEEE Trans Vis Comput Graph*, 28(1):1095–1105, 2022. doi: 10/g59b9m 2
- [32] G. Kim, J. Kim, and Y.-S. Kim. “Explain what a treemap is”: Exploratory investigation of strategies for explaining unfamiliar chart to blind and low vision users. In *Proc. CHI*, art. no. 805, 13 pp. ACM, 2023. doi: 10/g8q2w6 2
- [33] J. Kim, A. Srinivasan, N. W. Kim, and Y.-S. Kim. Exploring chart question answering for blind and low vision users. In *Proc. CHI*, art. no. 828, 15 pp. ACM, 2023. doi: 10/g8q2xf 1, 2, 3, 15
- [34] N. W. Kim, S. C. Joyner, A. Riegelhuth, and Y.-S. Kim. Accessible visualization: Design space, opportunities, and challenges. *Comput Graph Forum*, 40(3):173–188, 2021. doi: 10/g7h7gm 2
- [35] J. Lazar, J. Feng, and H. Hochheiser. *Research Methods in Human-Computer Interaction*. Wiley, 2014. 5
- [36] B. Lee, E. K. Choe, P. Isenberg, K. Marriott, and J. Stasko. Reaching broader audiences with data visualization. *IEEE Comput Graph Appl*, 40(2):82–90, 2020. doi: 10/gpbhbm 1
- [37] S. Lee, S.-H. Kim, and B. C. Kwon. VLAT: Development of a visualization literacy assessment test. *IEEE Trans Vis Comput Graph*, 23(1):551–560, 2017. doi: 10/f92d38 1, 2
- [38] C. Li, R. Y. Pang, A. Sharif, A. Chheda-Kothary, J. Heer, and J. E. Froehlich. Alteoviz: Facilitating accessible geovisualization. In *Short Paper Proc. VIS*, pp. 61–65, 2024. doi: 10/qwz5 2
- [39] A. Lundgard and A. Satyanarayan. Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content. *IEEE Trans Vis Comput Graph*, 28(1):1073–1083, 2022. doi: 10/gm5pdj 2, 4, 5
- [40] K. Marriott, B. Lee, M. Butler, E. Cutrell, K. Ellis, C. Goncu et al. Inclusive data visualization for people with disabilities: A call to action. *Interactions*, 28(3):47–51, 2021. doi: 10/gngfxc 2
- [41] A. McNutt, M. K. McCracken, I. J. Eliza, D. Hajas, J. Wagoner, N. Lanza et al. Accessible text descriptions for UpSet plots. *Comput Graph Forum*, 44(3):e70102, 2025. doi: 10/pt4n 2, 3, 4
- [42] C. Mei, J. Pollock, D. Hajas, J. Zong, and A. Satyanarayan. Benthic: Perceptually congruent structures for accessible charts and diagrams. In *Proc. ASSETS*, art. no. 27, 17 pp. ACM, New York, 2025. doi: 10/qv vj 2
- [43] P. Moraes, G. Sina, K. McCoy, and S. Carberry. Evaluating the accessibility of line graphs through textual summaries for visually impaired users. In *Proc. ASSETS*, pp. 83–90, 2014. doi: 10/qzfs 3
- [44] C. Nobre, K. Zhu, E. Mörtz, H. Pfister, and J. Beyer. Reading between the pixels: Investigating the barriers to visualization literacy. In *Proc. CHI*, art. no. 197, 17 pp. ACM, New York, NY, USA, 2024. doi: 10/n98g 2
- [45] J. Seo, S. S. Kamath, A. Zeidieh, S. Venkatesh, and S. McCurry. MAIDR meets AI: Exploring multimodal LLM-based data visualization interpretation by and with blind and low-vision users. In *Proc. ASSETS*, art. no. 57, 31 pp. ACM, 2024. doi: 10/g8qpcf 1, 3, 5, 9
- [46] J. Seo, Y. Xia, B. Lee, S. Mccurry, and Y. J. Yam. MAIDR: Making statistical visualizations accessible with multimodal data representation.

- In *Proc. CHI*, art. no. 211, 22 pp. ACM, 2024. doi: [10/pp53](#) 2
- [47] A. F. Siu, D. Fan, G. S.-H. Kim, H. V. Rao, X. Vazquez, S. O'Modhain et al. COVID-19 highlights the issues facing blind and visually impaired people in accessing data on the web. In *Proc. W4A*, art. no. 11, 15 pp. ACM, 2021. doi: [10/gpztzv](#) 1
 - [48] T. C. Smits, S. L'Yi, A. P. Mar, and N. Gehlenborg. AltGosling: Automatic generation of text descriptions for accessible genomics data visualization. *Bioinformatics*, 40(12):btae670, 2024. doi: [10/g9bdw7](#) 2
 - [49] T. C. Smits, S. L'Yi, H. N. Nguyen, A. P. Mar, and N. Gehlenborg. Explaining unfamiliar genomics data visualizations to a blind individual through transitions. In *Proc. AccessViz workshop at IEEE VIS*, pp. 24–28, 2024. doi: [10/g9bdws](#) 2
 - [50] N. Sultanum and A. Srinivasan. DATATALES: Investigating the use of large language models for authoring data-driven articles. In *Short Paper Proc. VIS*, pp. 231–235. IEEE, Los Alamitos, CA, USA, 2023. doi: [10/qvt5](#) 3
 - [51] X. Tang, A. Abdolrahmani, D. Gergle, and A. M. Piper. Everyday uncertainty: How blind people use genai tools for information access. In *Proc. CHI*, art. no. 63, 17 pp. ACM, 2025. doi: [10/qzfv](#) 3
 - [52] L. Wang, Z. Wang, S. Xiao, L. Liu, F. Tsung, and W. Zeng. VizTA: Enhancing comprehension of distributional visualization with visual-lexical fused conversational interface. *Comput Graph Forum*, 44(3):e70110, 2025. doi: [10/g94xd4](#) 2, 3
 - [53] R. Wang, C. Jung, and Y. Kim. Seeing through sounds: Mapping auditory dimensions to data and charts for people with visual impairments. *Comput Graph Forum*, 41(3):71–83, 2022. doi: [10/g8q2w5](#) 2
 - [54] B. L. Wimer, L. South, K. Wu, D. A. Szafir, M. A. Borkin, and R. A. Metoyer. Beyond vision impairments: Redefining the scope of accessible data representations. *IEEE Trans Vis Comput Graph*, 30(12):7619–7636, 2024. doi: [10/g8q2xj](#) 2
 - [55] C. Yan, H.-P. Hutter, F. M. Schmitt-Koopmann, and A. Darvishy. Chart accessibility: A review of current alt text generation. *IEEE Access*, 13:94040–94056, 2025. doi: [10/qvt6](#) 3
 - [56] Y. Yang, K. Marriott, M. Butler, C. Goncu, and L. Holloway. Tactile presentation of network data: Text, matrix or diagram? In *Proc. CHI*, pp. 1–12. ACM, 2020. doi: [10/g8q2xm](#) 2
 - [57] W. Yu and S. Brewster. Evaluation of multimodal graphs for blind people. *Univ Access Inf Soc*, 2(2):105–124, 2003. doi: [10/ff3vn2](#) 2
 - [58] Y. Zhao, M. A. Nacent, M. A. Sukhai, and S. Somanath. TADA: Making node-link diagrams accessible to blind and low-vision people. In *Proc. CHI*, art. no. 45, 20 pp. ACM, 2024. doi: [10/pq4n](#) 2
 - [59] J. Zong, C. Lee, A. Lundgard, J. Jang, D. Hajas, and A. Satyanarayan. Rich screen reader experiences for accessible data visualization. *Comput Graph Forum*, 41(3):15–27, 2022. doi: [10/g9bdw2](#) 2
 - [60] J. Zong, I. P. Piner, M. K. Chen, D. Hajas, and A. Satyanarayan. Umwelt: Accessible Structured Editing of Multi-Modal Data Representations. In *Proc. CHI*, art. no. 46, 20 pp. ACM, New York, 2024. doi: [10/pcz5](#) 1, 2
 - [61] J. Zong, I. P. Piner, M. K. Chen, D. Hajas, and A. Satyanarayan. Semantic scaffolding: Augmenting textual structures with domain-specific groupings for accessible data exploration, 2025. arXiv:2506.15883. 2

Touching or Chatting: The Utility of LLMs and Tactile Charts for Learning about Complex Chart Types by BLV Individuals

Appendix

In this appendix, we provide additional information and figures that we could not include in the main paper due to space limitations.

A INTERVIEW QUESTIONS

A.1 Background Questions

Demographics

- If you would like, please let us know your preferred pronouns.
- Year of birth
- Gender
- What is the highest level of education you have completed?
- What is your current status? Are you working or studying, and is it full-time or part-time? If so, what is your job or major?

Vision loss level

- How would you describe your vision-loss level? (Blind since birth / Lost vision suddenly / Lost vision gradually)
- How would you describe your vision level (e.g., completely blind, light perception, central vision loss, etc.)?
- (If the participant is not totally blind) If you know, what is your corrected visual acuity in either Snellen (e.g., 20/200) or LogMAR (e.g., 1.3)?
- If you are comfortable sharing, please indicate your visual pathology diagnosis. This information is optional and not required.

Screen Reader Experience

- Which screen reader do you use with your computer or other devices (e.g., NVDA, JAWS, VoiceOver, etc.)?
- How long have you been using a screen reader?
- When using a screen reader, what is your preferred rate of speech?
- Do you use other accessibility devices or software in combination with a screen reader, such as screen magnification or a Braille display? If yes, please describe all accessibility devices or software you use in combination with a screen reader.

Braille experience

- How long have you been reading Braille?
- How would you rate your skill in reading Braille (1 - I can't read Braille / 5 - I'm very proficient at reading Braille)
- In what contexts did you read Braille (e.g., work, personal use)?

Tactile Chart Experience

- How would you describe your familiarity with tactile charts or graphics? (1 - Not at all familiar / 5 - Extremely familiar)
- Can you describe in which contexts you have interacted with tactile charts or graphics?

LLM experience

- Have you ever used an AI chatbot assistant or a large language model (LLM), like ChatGPT, Google Gemini, Claude, or something similar?
- Which tool or tools do you use?
- How often do you use these tools?
- What do you usually use them for?
- Have you used an AI assistant to interpret data—like a dataset, table, chart, or visualization?

Data Visualization Experience

- How many hours do you use a computer or smartphone each day?
- Would you consider your career to be data-intensive or numbers-driven (e.g., regularly work with large datasets, perform statistical analyses, or make decisions based on quantitative information)?
- How often do you interact with data visualizations, such as those for work, from news articles, in video games, etc.?
- In which contexts do you encounter data and data visualizations (e.g., work, news, leisure)?
- How do you typically view datasets? (e.g., download to excel, you don't, etc.)

Familiarity With the Two Chart Types

- How familiar are you with **clustered heatmaps** - a visualization of tabular data that encodes values by color? (1 - Not at all familiar / 5 - Extremely familiar)
- Have you seen or touched **clustered heatmaps** through a tactile display or an embossed paper before? (Yes / No / I don't know what a clustered heatmap is)
- How familiar are you with **violin plots** - a visualization of a distribution of values? (1 - Not at all familiar / 5 - Extremely familiar)
- Have you seen or touched **violin plots** through a tactile display or an embossed paper before? (Yes / No / I don't know what a violin plot is)
- After you received our package, have you explored the model, or heard any of the instructions?
- Researcher B, do you have any additional questions?
- To ensure that we are both on the same link during today's discussion, do you mind also sharing your screen?

A.2 Questions After Chart Type Instructions

- Can you describe what a violin plot / clustered heatmap is, in your own words?
- What types of data are best suited for visualization using a violin plot / clustered heatmap?
- Can you provide an example scenario where you would use a violin plot / clustered heatmap?

[Only For Violin Plot]

- If a violin's shape has three peaks, what does it mean for the underlying data?
- If two distributions have the same median and overall range but different spreads around the center (one is tightly clustered, one is more dispersed), how would their violins differ in shape?

[Only For Clustered Heatmap]

- In a clustered heatmap, if two rows (or columns) are very similar, how would you expect them to appear in the heatmap and in the dendrogram?
- If one row (or column) looks very different from all the others, how would that show up in the heatmap and in the dendrogram?

A.3 Questions After Complex Dataset Exploration

- What is the dataset about?
- Imagine you have a friend who has never encountered this chart before. How would you describe this chart to them?
- What are the numerical characteristics being represented by this chart?
- What are the trends represented by this chart?
- What did you find interesting or surprising about the dataset?

A.4 Comparative Questions After Both Charts Were Explored

- You have learned about visualization types in two different ways: one by exploring a tactile model with instructions and the LLM, and another by only hearing a textual explanation and the LLM. Which training format do you prefer?
- How helpful is the tactile model used in the training for understanding the chart type?
- Do you think the tactile model helped you explore and interpret the new dataset (the second dataset for each chart type) with the AI assistant?
- How helpful is the AI assistant for interpreting new datasets (the second dataset for each chart type)?
- How helpful is the AI assistant used in the training for understanding the chart type?
- How useful is the training format “tactile model with exploration instructions and AI assistant” for BLV education in your opinion?
- How useful is the training format “textual explanation and AI assistant” for BLV education in your opinion?
- Do you have any other feedback or comments that we didn’t touch on today?

B SIMPLE ALT TEXT

B.1 Clustered Heatmap

This is a **clustered heatmap** showing how frequently actors appear in different movie genres.

- **Rows:** Actors.
- **Columns:** Movie genres.
- **Cells (color intensity):** The number of appearances for each actor–genre pair. **Darker** cells indicate **more** appearances, and **lighter** cells indicate **fewer** (or none).
- **Clustering:** Actors and genres are **clustered by similarity**. A **hierarchical tree (dendrogram)** next to the rows shows which actors cluster together (shorter branch distances indicate more similar genre-appearance patterns). Another dendrogram above the columns shows which genres cluster together.

Overall pattern:

- Julia Roberts and Tom Hanks show similar patterns, focusing primarily on Drama and Comedy
- Jennifer Lawrence and Leonardo DiCaprio share a strong emphasis on Drama
- Dwayne Johnson is an outlier in this dataset, with a strong specialization in Action and no Romance films

B.2 Violin Plot

This is a violin plot showing the distribution of body mass (in pounds) across three species of penguins: Adelie, Chinstrap, and Gentoo.

- **Y-axis:** Body mass (pounds).
- **X-axis:** Species, from left to right: **Adelie, Chinstrap, Gentoo**.
- **Violins:** Each violin shows the distribution of body mass for that species.

Overall pattern: Adelie and Chinstrap have unimodal distributions with similar median values. Gentoo shows a bimodal distribution and a higher median than the other two species.

C COMPLEX ALT TEXT

C.1 Clustered Heatmap

This is a **clustered heatmap** showing how citizens in **15 European countries** rate the relative importance of **9 sociocultural, institutional, and religious topics** (for example, **confidence in government, belief in God, and attitudes toward homosexuality**) based on a survey.

- **Rows:** Topics.
- **Columns:** Countries.
- **Cells (color intensity):** The agreement level for each country–topic pair. Darker cells indicate higher agreement, and lighter cells indicate lower agreement.
- **Clustering:** Both countries and topics are **clustered by similarity**. A **hierarchical tree (dendrogram)** next to the rows shows which topics cluster together (shorter branch distances indicate more similar response patterns). Another dendrogram above the columns shows which countries cluster together.

Overall pattern: The Nordic countries are the most confident in state institutions and most liberal. Religious values are strongest in Russia, Italy, Poland, Germany, Greece and Portugal. Poland, Italy, and Russia are most conservative on social issues.

C.2 Violin Plot

This is a violin plot showing the distribution of the Human Development Index (HDI) for countries in 2021, grouped by continent: Africa, Americas, Asia, Europe, and Oceania.

- **Y-axis:** HDI, ranging from 0 (low) to 1 (high).

- **X-axis:** Continents.

- **Violins:** Each violin shows how countries' HDI values are distributed within that continent (wider sections indicate more countries at that HDI level). A central marker indicates the median.

Overall pattern: Europe exhibits the highest and most consistent HDI values, while Asia has the widest variability. Africa and Oceania have non-normal distributions, indicating groups of countries with different levels of development. HDI is measured from 0 (low) to 1 (high).

D QUERIES DURING DATA EXPLORATION

We collected 276 queries in total. Of these, 13 were practice queries used by participants to familiarize themselves with the chatbot at the beginning of the study (e.g., “Hello”) and were not related to the charts. The remaining 263 queries were made during chart learning or data exploration and were used in the analysis reported in this section. In some cases, queries were not logged due to technical issues, such as browser refreshes. We recovered these queries from the Zoom video recording and included them in the dataset.

In this section, we analyze the queries collected from the complex alt text phase. We incorporated the analysis of queries collected during the chart-type instruction phase into our thematic analysis, reported in Sec. 5.

We collected 140 queries from the complex alt text phase, where participants explored the new datasets across the two chart types. Of these, 138 were substantive, excluding responses such as “thank you.” Because the tactile model did not encode the new data, the types and frequencies of exploration queries were similar across conditions, so we do not distinguish them here. Kim et al. [33] characterized BLV participants’ chart-QA queries in a Wizard-of-Oz study using prepared answer sheets. Here, we focus on patterns not covered in their analysis, such as interactions with a real LLM, while also comparing our results with theirs on points most relevant to our goals, especially chart-type learning.

Chart-/Dataset-Answerable Questions Of the 138 substantive queries, 99 could be answered using the chart and the dataset we provided to the LLM. Thirty queries needed to be answered based on the dataset, using information that was not directly encoded in the chart. This was especially common in the violin plot, where the dataset contains country-level information, whereas the chart shows only continent-level distributions. Some participants asked country-level questions, such as “*List the highest and lowest countries HDI [human development index]*”, or “*compare contrast northern and southern Europe for me*”. One participant also directly asked the LLM to provide the dataset—“*Please output the table in CSV format*”—because he wanted to analyze it in Excel. For some queries that were answerable from the chart, the LLM relied on the dataset to produce a more precise response. For example, for a query such as “*take the lowest half of Europe and the top half of Africa, contrast and compare, give me any interesting insights*”, the assistant tended to support its answer with exact values such as ranges or medians, drawn from the dataset. This highlights a strength of LLMs that have access to data: unlike a predefined chart, the LLM can answer more detailed questions users may have.

Questions Leveraging LLM’s World Knowledge For queries that were not answerable from the chart or the underlying dataset, several were answerable with the LLM’s world knowledge. These included requests for definitions (e.g., “*How is HDI [human development index] determined?*”), underlying reasons (e.g., “*why do Northern European countries have such a high HDI?*”), and perspective-based interpretations (e.g., “*tell me more about this data set from a christian perspective?*”), again showing the benefits of using an LLM that can incorporate world knowledge beyond what even an interactive visualization or the full dataset could reveal.

Out-of-Scope Questions We observed 11 queries about details beyond the dataset, which Kim et al. [33] characterize as out of scope. Most of these came from P3 and P12. For example, P3 asked for information not contained in the dataset, such as “*how does population data work in this chart*”, or “*how is it determined which country is conservative*”. P12 asked three questions related to Kansas, e.g., “*How does the HDI of Kansas compare to the USA HDI?*”, but he later explained that he simply wanted to test whether the LLM was limited to the provided dataset or could use outside knowledge. In these cases, the LLM typically acknowledged the limitation in its response (e.g., “I can’t tell from the dataset I have: it only includes country, Human Development Index (value), and region—there’s no population column here.”) and sometimes suggested uploading additional data.

Unclear Questions We observed 10 unclear queries. Some involved genuinely ambiguous references, similar to those reported by Kim et al. [33]. For example, in “*when it comes to the chart, give me the data from high to low*”, it is unclear what “data” refers to, so the LLM asked for clarification and suggested possible interpretations. Other unclear queries reflected intuitive expressions that made sense to participants but required further specification for the LLM. For example, in “*which continents have a similar HDI*”, the LLM requested a more explicit definition of “similar”, such as similarity in mean or range.

Follow-Up Questions During dataset exploration, some participants relied more heavily on the LLM’s guidance and followed the LLM through multi-turn exchanges. For example, while exploring the violin plot, P5 replied “Yes” three times to continue along the LLM’s suggested directions.

E TACTILE CHARTS




Figure 4: The 3D-printed tactile chart for the clustered heatmap, front view.




Figure 5: The 3D-printed tactile chart for the violin plot, front view.

F COMPLEX CHARTS




Figure 6: The clustered heatmap with the complex dataset.




Figure 7: The violin plot with the complex dataset.