

See What I Mean: A Technology Probe of AI Assistants for Chart Exploration with Low Vision

Yotam Sechayk

The University of Tokyo
Tokyo, Japan
sechayk-yotam@g.ecc.u-
tokyo.ac.jp

Lucy Jiang

University of Washington
Seattle, Washington, USA
lucjia@uw.edu

Yuhang Zhao

University of
Wisconsin-Madison
Madison, Wisconsin, USA
yuhang.zhao@cs.wisc.edu

Takeo Igarashi

The University of Tokyo
Tokyo, Japan
takeo@acm.org




Figure 1: ChartAEye shares the visible area and indicated point with the assistant context (left). Given a question (center), the generic AI persona (GAI) answers plainly, while the low-vision AI persona (LVAI) adds spatial and semantic guidance (right).

Abstract

Charts remain inaccessible to low vision individuals (LVI), who commonly magnify local details at the cost of the overview. AI assistants can help, but prior work has largely served non-visual access, while LVI prefer using residual vision. We built ChartAEye, a technology probe that lets LVI interact with an AI assistant visually, sharing their magnified view and point of interest as system context. Using the probe, we then asked whether the AI's feedback also needs to be adapted to low vision, comparing a generic AI persona (GAI) with one designed around LVI needs and lived experience (LVAI). In a within-subjects study, 12 LVI used ChartAEye to complement their vision, reduce the effort of viewing chart elements, and practice more independence. Eight of 12 participants preferred the LVAI, whose responses more often carried spatial guidance (71.2% vs. 7.8%) and semantic guidance of visuals (84.8% vs. 29.1%). We close with design recommendations for AI assistants that support LVI in exploring charts with their residual vision.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **User centered design**; *Visualization*; *Accessibility systems and tools*.



This work is licensed under a Creative Commons Attribution 4.0 International License.
ASSETS '26, Vila Nova de Gaia, Portugal
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2521-0/26/10
<https://doi.org/10.1145/3797867.3841222>

Keywords

Low vision, Accessibility, Data visualization, AI assistant, Screen magnification, Technology probe

ACM Reference Format:

Yotam Sechayk, Lucy Jiang, Yuhang Zhao, and Takeo Igarashi. 2026. See What I Mean: A Technology Probe of AI Assistants for Chart Exploration with Low Vision. In *The 28th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '26)*, October 25–28, 2026, Vila Nova de Gaia, Portugal. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3797867.3841222>

1 Introduction and Related Work

Charts communicate complex information and insights across domains such as business, education, healthcare, and journalism [7]. Yet charts remain inaccessible to low vision individuals (LVI), a population estimated at 295 million globally [3]. The most common assistive tool LVI use to access charts is screen magnification [22, 28]. Magnification enlarges local details, but it fragments the chart into partial views and sacrifices the overview needed to identify trends or compare data points [18, 22].

Prior work on chart accessibility has largely supported non-visual access, such as structured screen reader navigation [2, 6, 31] and, more recently, AI assistants that answer questions about charts for screen reader users [8, 11, 13, 24]. However, LVI prefer to use their residual vision for visual access [25, 28]. Work that supports this preference has characterized how LVI read charts and enhanced magnification itself [18, 22, 28], without an AI assistant in the loop. How AI assistants can support chart exploration that builds on residual vision remains largely unexplored.

We built ChartAEye as a technology probe [10], a working prototype used to explore user reactions and surface design needs. The probe lets LVI interact with an AI assistant visually, sharing their magnified view and point of interest as system context. We then asked whether the AI’s feedback also needs adapting to low vision, comparing a generic AI persona (GAI) with one designed around LVI needs and lived experience (LVAI). This question is inspired by ability-based design [29]: even once a system is ability-aware at the input level (the “Context” principle), does it still need to be ability-aware in what it produces (the “Adaptation” principle)?

We conducted a within-subjects study in which $N = 12$ LVI explored charts using ChartAEye with both AI personas. Our qualitative findings show how participants used ChartAEye to complement their vision (e.g., resolving color ambiguity), reduce the effort needed to view chart elements (e.g., retrieving a value with a single click), and practice more independence (e.g., by verifying their own visual observations). In the persona comparison, 8 of 12 participants preferred the LVAI, 2 preferred the GAI, and 2 noticed no difference. In practice, the LVAI included semantic guidance of visuals in 84.8% of responses versus 29.1% for the GAI, and spatial guidance in 71.2% versus 7.8%. Nine of 12 participants acted on and followed this spatial guidance. While questionnaire measures showed at most anecdotal evidence of a difference between the personas, participants’ preferences and behavior suggest that tailoring the AI’s output matters, even when its input is already grounded in the user’s visual context. We close with design recommendations for AI assistants that support LVI in exploring charts with their residual vision.

2 The ChartAEye Technology Probe

We designed ChartAEye to support LVI in exploring charts by pairing their residual vision with an AI assistant that is aware of what is currently displayed and where they point (Figure 1). We grounded the design in prior work on LVI chart access and interaction needs [1, 22, 23, 25, 28] and in the lived experience of an author who identifies as LVI.

ChartAEye is a web application with two persistent regions: the chart area and the chat panel. Magnification (zoom and pan) is integrated directly into the chart area. The prototype shares the portion of the chart visible in the user’s viewport with the assistant. Users indicate a point of interest by double-clicking on the chart [16], an interaction we call *pointing*, which renders a circular indicator that persists across panning and zooming and can be dismissed or replaced. Users can type or speak queries, listen to responses via text-to-speech (reducing eye strain [23]), and adjust the indicator’s size and color, speech parameters, and UI size.

With each click or question, the prototype sends the assistant two images alongside the prompt [12]: (1) the full chart, annotated with a bounding box marking the area currently visible to the user and, when active, a circular marker at their indicated point of interest, and (2) an unannotated crop of the chart matching the visible area. We deliberately exclude the chart’s underlying data table, which is unavailable in many real-world contexts (e.g., news articles or slide decks), so the assistant operates solely on the chart’s visual content. AI interaction uses Google Gemini 3.1 Pro Preview [26], which performed best on chart understanding in our internal testing.

We instructed the assistant to support LVI along two dimensions drawn from prior work [22, 23, 28]: *spatial guidance*, describing where elements are within and beyond the current view, with navigation instructions adjusted to the user’s zoom and position, and *semantic guidance* of visuals, describing colors, patterns, and visual relationships, including what surrounds the current view or indicated point. These instructions define the low-vision AI persona (LVAI), while the generic AI persona (GAI) uses an otherwise identical prompt with both guidance blocks removed, receiving the same inputs, interface, and interaction methods. This design isolates the AI’s output as the only difference between the personas. Full prompts are in the supplementary materials.

3 Methodology

Participants. We recruited 12 LVI over 18 who use residual vision for reading (five men, six women, one non-binary, mean age 39.75, $SD = 14.3$) via low-vision communities and snowball sampling. We screened participants with a questionnaire on assistive technology use. Participants’ diagnosed conditions, visual acuity, and visual fields varied widely, and all but one used screen magnification daily. The authors’ institution’s Institutional Review Board approved the study, and participants were compensated \$45 USD for the roughly two-hour sessions ($M = 115.8$, $SD = 23.4$ minutes).

Procedure. We conducted all sessions remotely via Zoom [32]. Our study had two conditions: GAI and LVAI, both using the same interface. Each session started with an introduction and a training chart to familiarize participants with the interface. Participants then explored two real-world charts (bar [20], choropleth [17]), one per condition, counterbalanced, without knowing which persona was active in each. Per condition, participants had two minutes to answer a starter question about the chart (e.g., whether surpluses were more common before or after 1980), then up to ten minutes to freely explore the chart and describe what they learned to the researcher. After each condition, participants completed NASA-TLX [9] and custom Likert questions inspired by prior work [22]. Finally, participants completed the System Usability Scale (SUS) [5] for the persona-independent interface and a semi-structured interview about the prototype, the two personas, and design suggestions. The detailed procedure, stimuli, and questionnaire items are in the supplementary materials.

Analysis. We video recorded and transcribed all sessions [19], then analyzed them with inductive coding and reflexive thematic analysis [4], foregrounding the perspective of an author who identifies as LVI. Two authors coded two transcripts and met to develop a shared codebook, then split the remaining transcripts equally. They additionally coded participants’ chat logs and video-recorded actions (e.g., zooming, clicking, or following navigation instructions). One author then constructed themes over the generated codes, which the authors discussed and refined together [15]. Separately, an LLM-assisted pipeline classified every assistant message, under both personas, for spatial guidance and semantic guidance of visuals, the dimensions targeted by the LVAI’s prompt additions (§2). An author then verified all labels in a review blinded to persona, correcting 35 of 532 (6.6%), as detailed in the supplementary materials. Lastly, a paired Bayesian analysis (Bayes factors [21])

of NASA-TLX and the custom questionnaire found at most anecdotal evidence of differences between the personas, reported for completeness in [section 4](#).

4 Findings

Participants produced 266 interaction events (127 click-based, 139 typed), and 11 of 12 used pointing (§2). The interface received a mean SUS score of 89.6 ($SD = 9.28$). Most participants (P1, P3–P12) preferred accessing charts visually, whereas P2 generally preferred data tables, and the visual interaction felt natural.

4.1 Exploring Charts with an AI That Shares Visual Context

Complement vision. The assistant gave participants access to information their vision could not reliably provide. P9 used clicks to resolve color ambiguity on the map chart: “once I figured out... that the light purple color was not gray, that helped a lot.” P5 similarly relied on it for color values and axis labels, while P10 and P11 recovered information magnification made them miss, such as labels outside the visible portion (P10) or axis values that require panning and tracking (P11). However, the same visual limits that made this complementary information valuable also made it hard to verify. Errors occurred despite the strong underlying model: the assistant misidentified Bhutan (P7), confused Panama with Costa Rica (P8), and described a clicked country as light blue when it was light orange (P12). P8 captured the asymmetry: “I don’t know if it would have gotten it right or wrong, but I know I wouldn’t have been able to verify it.” Participants adapted their behavior in response. P5 limited queries to larger countries where visual cross-checking was feasible, P10 and P2 trusted broad interpretations more than specific values, and P7 returned to verify previous queries after the AI misidentified Bhutan, “I was quite intrigued when it got Bhutan incorrect. So then I was checking to see whether it got anything else wrong.”

Reduce effort. The assistant also made information participants could already access easier to reach. P7 described how panning, zooming, and cross-referencing labels compound the time chart reading takes with low vision, steps the assistant collapsed into a single click. P3 summarized the consolidation: “usually I have to go scan a chart in, and then manipulate it, and then start looking up the information I’m questioning. This is more of a one-stop shop.” P12 described how this shifted their attention from the logistics of seeing toward interpretation, letting them “explore it and have fun with it.” For P9, unifying magnification, click context, and AI querying in one interface removed the need for the screen reader they otherwise use situationally, “depending on if I’m having a bad vision day.” Latency undermined this benefit. 45% of responses took at least 10 seconds ($M = 11.3$ s, $SD = 7.4$). Some participants wanted click responses to be immediate and convey low-level information (e.g., axis values). P1 expressed: “if I have to wait... like 15 seconds to know the answer, that will make me less willing to click.” When waits grew long, P1 zoomed in and read the value with their own vision instead. P10 named response speed the top usability issue.

Agency and independence. Participants described the assistant as expanding what they could do without taking over. P5 enacted a deliberate division of labor, asking the assistant only about what

was visually inaccessible while continuing independently for what they were confident reading themselves. P7 stated the same preference explicitly: “I don’t think I’d like something to be able to do it all for me. But the more information that it can give you, the easier it is.” P10 formed interpretations with their own vision first, then used the assistant to check them: “it looks like there’s a general trend to more deficits. Let me make sure that interpretation is correct.” The assistant also removed the social cost of asking for help. P9 explained that with human assistance, “there’s only so many questions you can ask someone before they ask, are you sure you know what you’re doing?” Reported agency did not differ between the two personas.

4.2 Adapting AI Feedback to Low Vision

8 of 12 participants preferred the LVAI (P2, P4, P6, P7, P9–P12), P5 and P8 preferred the GAI, and P1 and P3 did not notice meaningful differences. Participants who preferred the LVAI pointed to responses tailored to their visual needs. P2 found the LVAI “much more specific about tying [the visuals] back to numbers or magnitudes” and “better suited to [their] personal needs.” In the coded chat logs (141 GAI, 125 LVAI messages), the LVAI included spatial guidance in 71.2% of its responses and semantic guidance of visuals in 84.8%, compared to 7.8% and 29.1% for the GAI. This guidance shaped exploration in practice. When P12 clicked near Japan, the LVAI reported it was “just outside my view, above the top edge to the right”, letting them pan directly to it. We found that 9 of 12 participants followed the navigation instructions under the LVAI.

Tailored feedback did not fit all participants equally. Five participants (P1, P5, P8, P10, P11) found the LVAI more verbose than the GAI, matching the custom questionnaire ratings. P10 found the detail useful, feeling it answered questions they were about to ask, while P1 preferred brevity: “sometimes I want an easy answer, value, name.” Combined NASA-TLX task load averaged $M = 37.9$ ($SD = 20.5$) for the GAI and $M = 30.5$ ($SD = 14.6$) for the LVAI. On the custom questionnaire (GAI; LVAI), participants rated the LVAI more favorably on five of seven dimensions, most notably effectiveness ($M = 5.33$; 6.17) and personal fit ($M = 4.42$; 5.08). Ratings were near-similar on agency, trust, and access to visuals, and favored the GAI on control ($M = 5.67$; 5.50) and verbosity ($M = 5.17$; 4.75). A Bayesian analysis found no measure reached more than anecdotal evidence for a difference, with full ratings in the supplementary materials.

4.3 Expectations and Desired Extensions

The probe prompted participants to articulate concrete extensions, which converged on four low-vision needs. First, participants wanted the assistant’s output anchored visually on the chart: color-coded pins as visual landmarks (P9), highlights on the elements being described (P7, P8), and on-chart annotations as a more accessible alternative to long textual responses (P8). Second, participants wanted the assistant to proactively surface what their vision missed, as P11 requested: “describe some of the things on this map that appear very small... what in this chart is potentially visibly difficult to notice.” P12 similarly imagined the assistant flagging missed content with meaningful highlight colors, easier to catch at a glance than text. Third, participants wanted structured, scalable output:

multi-point selection with values listed to offload memory (P11), scalable response detail (P5), and response scope that follows the zoom level (P10). Fourth, participants asked to post-process the chart's visuals, inverting or changing colors as they already do in other contexts (P3, P9).

5 Discussion and Design Recommendations

ChartAEye probed two layers of accommodation: sensing the user's visual context, and adapting the AI's feedback to low vision. Our findings suggest both layers matter. The three experience themes emerged under both personas, so sharing visual context supported participants regardless of how the AI phrased its feedback. The tailored persona, in turn, changed what participants received, how they explored, and what most preferred. Response content is therefore an accessibility design surface in its own right, the "Adaptation" principle of ability-based design [29]. However, this preference did not appear in the questionnaire ratings. Single-session exposure, instrument sensitivity, and preferences driven by salient moments rather than aggregate load can all mute such differences. The 8 of 12 split is therefore a directional signal, not a measured effect. From the findings and the desired extensions, we derive four design recommendations (DRs) for AI assistants that support LVI in accessing visual content.

DR1: Anchor AI output on the chart itself. Participants' desired extensions converged on making the assistant's responses visible: pins, highlights, and annotations placed on the chart, in colors that carry meaning at a glance. Future systems should complement responses with persistent visual anchors in the user's viewport, and proactively highlight content the user is likely to have missed (e.g., small or low-contrast elements).

DR2: Design for verifiability, not only accuracy. Errors occurred despite the strong underlying model, consistent with benchmarks showing that VLMs degrade on fine-grained perceptual chart tasks [27, 30]. Participants often could not check the assistant's answers with their own vision, the same access gap the assistant was compensating for. Without the underlying data, every reported value is a visual estimate. Reducing the verification burden is therefore a design problem, not only a model-quality problem. Systems should highlight what the assistant is referring to on the chart, ground answers in visual evidence, and scaffold participants' testing strategies, such as probing with questions whose answers they already knew.

DR3: Scale responses to the interaction. Clicks called for fast, minimal answers while typed questions invited detail, yet responses were uniformly long and 45% took at least 10 seconds. Latency may also have muted the persona comparison, as the LVAI's longer responses take longer to generate and read. Response depth should follow the query: quick lookups for clicks, potentially served by smaller and faster models, and fuller detail for typed questions. Participants also suggested letting users scale response detail and tying response scope to the zoom level, describing a region when zoomed out and a point when zoomed in.

DR4: Augment vision, preserve agency. Participants explicitly rejected full delegation, divided labor between their own vision and the assistant, and verified their own interpretations against it. Systems should keep every exchange user-initiated and treat the

user's current view as authoritative rather than navigating on their behalf.

6 Conclusion

We presented ChartAEye, a technology probe that lets LVI explore charts with an AI assistant through their residual vision, sharing their magnified view and point of interest as system context. Our findings are formative: 12 LVI, diverse in visual conditions and access strategies [14], each used the probe for a single session, with two pre-selected charts and one vision-tailored prompt configuration. While future work should examine longer-term use, broader visual media, alternative tailoring configurations, and larger, stratified samples, our findings suggest supporting LVI requires accommodation on both sides of the interaction, in what the AI senses and what it says. We hope our design recommendations guide future AI assistants that support LVI in using their residual vision, on charts and beyond.

References

- [1] Rubén Alcaraz Martínez, Mireia Ribera Turró, and Toni Granollers Saltiveri. 2022. Methodology for Heuristic Evaluation of the Accessibility of Statistical Charts for People with Low Vision and Color Vision Deficiency. *Universal Access in the Information Society* 21, 4 (Nov. 2022), 863–894. doi:10.1007/s10209-021-00816-0
- [2] Matthew Blanco, Jonathan Zong, and Arvind Satyanarayan. 2022. Olli: An Extensible Visualization Library for Screen Reader Accessibility. In *IEEE VIS Posters*. IEEE.
- [3] Rupert Bourne, Jaimie D Steinmetz, Seth Flaxman, et al. 2021. Trends in Prevalence of Blindness and Distance and near Vision Impairment over 30 Years: An Analysis for the Global Burden of Disease Study. *The Lancet Global HEALTH* 9, 2 (Feb. 2021), e130–e143. doi:10.1016/s2214-109x(20)30425-3
- [4] Virginia Braun and Victoria Clarke. 2019. Reflecting on Reflexive Thematic Analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (Aug. 2019), 589–597. doi:10.1080/2159676X.2019.1628806
- [5] John Brooke. 1996. SUS: A 'quick and Dirty' Usability Scale. In *Usability Evaluation in Industry*. CRC Press.
- [6] Frank Elavsky, Lucas Nadolskis, and Dominik Moritz. 2024. Data Navigator: An Accessibility-Centered Data Navigation Toolkit. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (Jan. 2024), 803–813. doi:10.1109/TVCG.2023.3327393
- [7] Steven L. Franconeri, Lace M. Padilla, Priti Shah, Jeffrey M. Zacks, and Jessica Hullman. 2021. The Science of Visual Data Communication: What Works. *Psychological Science in the Public Interest* 22, 3 (Dec. 2021), 110–161. doi:10.1177/15291006211051956
- [8] Joshua Gorniak, Yoon Kim, Donglai Wei, and Nam Wook Kim. 2024. VizAbility: Enhancing Chart Accessibility with LLM-based Conversational Interaction. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. ACM, Pittsburgh PA USA, 1–19. doi:10.1145/3654777.3676414
- [9] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183. https://doi.org/10.1016/S0166-4115(08)62386-9
- [10] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Alison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. Association for Computing Machinery, New York, NY, USA, 17–24. doi:10.1145/642611.642616
- [11] Jiho Kim, Arjun Srinivasan, Nam Wook Kim, and Yea-Seul Kim. 2023. Exploring Chart Question Answering for Blind and Low Vision Users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3544548.3581532
- [12] Chu Li, Rock Yuren Pang, Ather Sharif, Arnavi Chheda-Kothary, Jeffrey Heer, and Jon E. Froehlich. 2024. AltGeoViz: Facilitating Accessible Geovisualization. In *2024 IEEE Visualization and Visual Analytics (VIS)*. 61–65. doi:10.1109/VIS55277.2024.00020
- [13] Chu Li, Rock Yuren Pang, Arnavi Chheda-Kothary, Ather Sharif, Henok As-salif, Jeffrey Heer, and Jon E. Froehlich. 2026. GeoVisA11y: An AI-based Geovisualization Question-Answering System for Screen-Reader Users. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*

- (CHI '26). Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3772318.3790334
- [14] Kelly Mack, Emma McDonnell, Dhruv Jain, Lucy Lu Wang, Jon E. Froehlich, and Leah Findlater. 2021. What Do We Mean by “Accessibility Research”? A Literature Survey of Accessibility Papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3411764.3445412
 - [15] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 72 (2019), 23 pages. doi:10.1145/3359174
 - [16] Vishnu Nair, Hanxiu 'Hazel' Zhu, and Brian A. Smith. 2023. ImageAssist: Tools for Enhancing Touchscreen-Based Image Exploration Systems for Blind and Low Vision Users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Chi '23)*. Association for Computing Machinery, New York, NY, USA, Article 76, 17 pages. doi:10.1145/3544548.3581302
 - [17] Our World in Data. [n. d.]. Gender Ratio. <https://ourworldindata.org/gender-ratio> Accessed: February 21, 2026.
 - [18] Yash Prakash, Pathan Aseef Khan, Akshay Kolgar Nayak, Sampath Jayarathna, Hae-Na Lee, and Vikas Ashok. 2025. Towards Enhancing Low Vision Usability of Data Charts on Smartphones. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (Jan. 2025), 853–863. doi:10.1109/TVCG.2024.3456348
 - [19] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust Speech Recognition via Large-Scale Weak Supervision. arXiv:2212.04356 [eess.AS] <https://arxiv.org/abs/2212.04356>
 - [20] Reuters. [n. d.]. USA Trump Tax Graphic. <https://www.reuters.com/graphics/USA-TRUMP/TAX-GRAPHIC/dwvkjoznpm/> Accessed: February 21, 2026.
 - [21] Jeffrey N. Rouder, Paul L. Speckman, Dongchu Sun, Richard D. Morey, and Geoffrey Iverson. 2009. Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review* 16, 2 (2009), 225–237. doi:10.3758/PBR.16.2.225
 - [22] Yotam Sechayk, Hennes Rave, Max Rädler, Mark Colley, Zhongyi Zhou, Ariel Shamir, and Takeo Igarashi. 2026. Improving Low-Vision Chart Accessibility via On-Cursor Visual Context. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)* (Barcelona, Spain) (CHI '26). Association for Computing Machinery, New York, NY, USA, 21 pages. doi:10.1145/3772318.3791165
 - [23] Yotam Sechayk, Ariel Shamir, Amy Pavel, and Takeo Igarashi. 2025. VeasyGuide: Personalized Visual Guidance for Low-Vision Learners on Instructor Actions in Presentation Videos. In *Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '25)*. Association for Computing Machinery, Denver, CO, USA. arXiv:2507.21837 [cs] doi:10.1145/3663547.3746372
 - [24] JooYoung Seo, Sanchita S. Kamath, Aziz Zeidieh, Saairam Venkatesh, and Sean McCurry. 2024. MAIDR Meets AI: Exploring Multimodal LLM-Based Data Visualization Interpretation by and with Blind and Low-Vision Users. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '24)*. Association for Computing Machinery, New York, NY, USA, 1–31. doi:10.1145/3663548.3675660
 - [25] Sarit Felicia Anais Szpiro, Shafeka Hashash, Yuhang Zhao, and Shiri Azenkot. 2016. How People with Low Vision Access Computing Devices: Understanding Challenges and Opportunities. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '16)*. Association for Computing Machinery, New York, NY, USA, 171–180. doi:10.1145/2982142.2982168
 - [26] Gemini Team. 2024. Gemini: A Family of Highly Capable Multimodal Models. arXiv:2312.11805 [cs.CL] <https://arxiv.org/abs/2312.11805>
 - [27] Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. Eyes Wide Shut? Exploring the Visual Shortcomings of Multimodal LLMs. arXiv:2401.06209 [cs] doi:10.48550/arXiv.2401.06209
 - [28] Yanan Wang, Yuhang Zhao, and Yea-Seul Kim. 2024. How Do Low-Vision Individuals Experience Information Visualization?. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Chi '24)*. Association for Computing Machinery, New York, NY, USA, Article 43, 43:1–43:15 pages. doi:10.1145/3613904.3642188
 - [29] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. 2011. Ability-Based Design: Concept, Principles and Examples. *ACM Trans. Access. Comput.* 3, 3 (April 2011), 9:1–9:27. doi:10.1145/1952383.1952384
 - [30] Yifan Wu, Lutao Yan, Leixian Shen, Yunhai Wang, Nan Tang, and Yuyu Luo. 2024. ChartInsights: Evaluating Multimodal Large Language Models for Low-Level Chart Question Answering. arXiv:2405.07001 [cs] doi:10.48550/arXiv.2405.07001
 - [31] Jonathan Zong, Crystal Lee, Alan Lundgard, JiWoong Jang, Daniel Hajas, and Arvind Satyanarayan. 2022. Rich Screen Reader Experiences for Accessible Data Visualization. *Computer Graphics Forum* 41, 3 (2022), 15–27. doi:10.1111/cgf.14519
 - [32] Zoom Video Communications, Inc. 2026. Zoom. <https://zoom.us> Computer software.