The Future of Visualization Dashboards in the Age of Generative AI Agents

VAISHALI DHANOA, Aarhus University, Denmark
DUOSI DAI, Aarhus University, Denmark
NIKLAS ELMQVIST, Aarhus University, Denmark

Generative AI is rapidly transforming data visualization practices, raising fundamental questions about the future of visualization dashboards and the evolving roles of the people who create and consume them. We conducted semi-structured interviews with four domain experts spanning dashboard consulting, AI-driven analytics engineering, and data analytics. Through qualitative analysis, we explored the following topics: (1) the future of dashboards—examining current challenges and how interfaces will evolve, and (2) the future of stakeholder roles—exploring how the responsibilities of dashboard authors and users are shifting. Our findings suggest that dashboards will persist but evolve toward conversational, question-driven interfaces, while the authoring role transitions from technical execution to intent articulation and output validation. We discuss design implications for supporting mixed-initiative interaction, trust calibration, and preserving shared analytical ground in AI-augmented visualization environments.

CCS Concepts: • **Human-centered computing → Empirical studies in visualization**.

Additional Key Words and Phrases: Visualization dashboards, AI, LLMs, data practitioners, agents.

ACM Reference Format:
Vaishali Dhanoa, Duosi Dai, and Niklas Elmqvist. 2026. The Future of Visualization Dashboards in the Age of Generative AI Agents. In *Proceedings of AutomationXP26 Workshop of the 2026 CHI Conference on Human Factors in Computing Systems, April 14, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages.

1 Introduction

Visualization dashboards have long served as a primary medium for communicating data-driven insights across organizations [1, 9]. However, the rapid emergence of generative AI is challenging foundational assumptions about how dashboards are authored, consumed, and maintained. Users can now pose analytical questions in natural language, receive auto-generated visualizations, and interact with data conversationally [10]. This raises several pressing questions: Will dashboards persist or be replaced by conversational AI agents? How will the roles of dashboard authors and users evolve? Will visualization even be needed in a time where AI models are improving by leaps and bounds every year?

Despite growing interest in AI-augmented visualization, there is limited empirical understanding of how practitioners perceive these shifts. In this paper, we begin to address this gap through semi-structured interviews with four domain experts: a dashboard consultant, a co-founding engineer at an AI-analytics company, a Power BI consultant, and a data analytics entrepreneur. Our research questions are

- **RQ1:** What is the future of dashboards in the era of generative AI?
- **RQ2:** What is the future of data practitioners and stakeholders in the era of AI?

Authors' Contact Information: Vaishali Dhanoa, dhanoa@cs.au.dk, Aarhus University, Aarhus, Denmark; Duosi Dai, duosi@cs.au.dk, Aarhus University, Aarhus, Denmark; Niklas Elmqvist, elm@cs.au.dk, Aarhus University, Aarhus, Denmark.

Through qualitative analysis, we address these research questions and find that dashboards will persist in evolved forms and that the roles of practitioners will evolve accordingly.

2 Related Work

Dashboards have been extensively studied in visualization research. Sarikaya et al. [9] characterized dashboards along multiple design dimensions, while Bach et al. [1] proposed a comprehensive design pattern framework. Despite their ubiquity, dashboards face persistent challenges including maintenance burden, unclear stakeholder requirements, and the difficulty of creating adaptive structures that evolve with changing analytical needs [5]. Our work extends this literature by examining how generative AI is reshaping these long-standing challenges.

Natural language interfaces (NLIs) for visualization have gained significant attention [10, 12]. Systems such as NL4DV [8] and Chat2Vis [6] demonstrate how LLMs can translate user queries into visual representations. However, concerns persist around hallucination, loss of transparency, and the risk of users accepting AI outputs uncritically [7]. Our study surfaces practitioner perspectives on these concerns, grounding them in real-world deployment experience.

Research on data work has highlighted the collaborative and sociotechnical nature of visualization practice [4, 11]. As AI capabilities expand, the roles of data analysts, visualization designers, and end users are being challenged. Recent work has started exploring how AI shifts authoring from direct creation toward curation and orchestration [13]. Our findings complement this work by providing empirical evidence from practitioners on the evolving division of labor between humans and AI in the dashboard ecosystem.

3 Method

We recruited four domain experts with diverse roles in the dashboard ecosystem: P1 (Female, Germany), a dashboard consultant who advises organizations on dashboard design and implementation; P2 (Male, US), a co-founding engineer at an AI-driven analytics company with expertise in building intelligent data tools; P3 (Male, Denmark), a data analytics consultant working primarily within the Power BI ecosystem; and P4 (Male, Brazil), a data analytics consultant and entrepreneur with experience in BI tool adoption and data exploration.

We conducted semi-structured interviews covering seven topic areas: current roles and responsibilities, pain points with dashboards, experiences with LLM integration, observed changes in dashboard usage, promising and concerning AI features, the evolving role of human designers, and predictions for the future of data interfaces. Each interview lasted approximately 45–60 minutes and was recorded and transcribed. The interview protocol is available in our supplementary materials.[1]

Two researchers independently coded the transcripts, generating initial codes that were iteratively refined through discussion. Codes were organized to address our research questions: (1) the *future of dashboards*, encompassing current challenges and interface evolution, and (2) the *future of stakeholder roles*, encompassing the shifting responsibilities of authors and users.

4 Results

We present our findings organized around the two research questions.

[1]https://osf.io/fcdy2

4.1 RQ 1: The Future of Dashboards

4.1.1 Current Practices and Challenges. Participants consistently identified the **maintenance burden** and **inefficiency of current dashboard structures** as key pain points (P1, P2, P3). End users often fail to provide clear requirements and frequently change their requests, forcing costly rebuilds:

> **P3**
>
> *"That was exactly what the stakeholder wanted, of course, two weeks after they want two dimensions more and then it breaks all your visualization and just waste your time of rebuilding that."*

The **authoring process** itself was described as time-consuming and cognitively demanding (P2, P4). P4 struggled with visual design decisions, while P2 emphasized that scientific visualization requires iterative refinement and deep domain understanding. P4 also noted that **limited data processing capabilities** within BI tools force analysts to work across multiple tools—P3 similarly writes code in Visual Studio Code before importing processed data into Power BI.

4.1.2 The Future of the Dashboard Interface. All four participants agreed that dashboards will **persist**, though in evolved forms. P4 drew an analogy to spreadsheets:

> **P4**
>
> *"If Excel sheets existed today, so dashboards will also exist, right? When you look in the broad picture, you have some few people doing many things with AI, but most people are still on sheets, so dashboards are in between—they will exist for a long time."*

What *will* change is the interaction paradigm. Participants foresaw a shift from static, click-based interfaces toward **conversational and question-oriented interaction**. P2 anticipated that basic configuration UIs would become hidden from general users, available only to power users. P3 envisioned that "visualizations will be built by an AI and it'll be possible to talk to the visualization—maybe as a person, as a conductor." P2 emphasized that defining intent and understanding output will still be done by humans.

However, participants raised important **concerns**. P2 warned about dashboard proliferation and the erosion of shared analytical ground:

> **P2**
>
> *"If it becomes super convenient to generate dashboards as quickly as possible... you generate so many dashboards that nobody can [keep track]. Or everybody creates their own personalized dashboard... even though it's the same data, they could look different."*

P2 also raised concerns about AI *sycophancy* reinforcing user biases, drawing parallels to p-value hacking [3]. On the positive side, AI was seen as **lowering barriers to access**—P4 noted that "with very few sentences and very simple prompts, you can make a huge exploration in your data sets." Features expected to persist include core visualizations, public-facing dashboards, and reusable templates. Features likely to diminish include complex configuration interfaces, rigid layouts, and dashboards used solely for intermediate exploratory analysis, as qualified users "cut the dashboards from the process" (P4).

4.2 RQ 2: The Future of Stakeholder Roles

4.2.1 Authors and Users: Current Roles and Tensions. For **authors**, participants described a transition from hands-on design to **orchestration**:

P2

"I would expect it starts to shift away from the technical implementation details, because AI can handle those, more towards handling the intent—the purpose and goal, like what do you actually want to visualize and why."

P4 echoed this, noting that practitioners now spend more time designing and evaluating results rather than writing code. P3 characterized authors as evolving into "conductors" coordinating AI-supported processes, though noting that practical adoption lags behind public discourse: "It's the influencers who talk about this [AI agents]. It's not the users, it's not the report builders. They're too busy with just making the reports work."

For **users**, the picture is more complex. Participants described a **polarization of trust**—users either over-trust or under-trust AI capabilities:

P2 – Polarized User Trust

"There is this split between people saying, 'AI is all bogus, it's hallucinating all the time, it's so stupid,' and then there are the other folks who somehow blindly [think] it's an agent, so it must be eternally intelligent and answer all the questions."

P2 further warned that natural language interfaces may lead to **reduced analytical reflection**: "You're maybe starting to think less... it's kind of like skipping the middle part and just asking the question directly, which is cool on the one hand; on the other hand, it's tricky." P4 noted that most end users lack the maturity to interact effectively with AI tools: "For the normal users, I don't think they have maturity to talk to an LLM and ask questions."

4.2.2 The Evolution of Stakeholder Roles. Participants unanimously emphasized that certain human competencies will remain **irreplaceable**. **Problem framing** emerged as the most critical skill: P1 emphasized "asking the right questions" and having the "gut feeling" for when to ask them; P2 argued that designers must "spend more time on why do we really want to visualize something"; and P4 stressed that "the ability to understand the problem and to transcribe that in technical requirements—I don't see AI doing that."

Critical evaluation of AI outputs was identified as an emerging core competency. P2 highlighted that AI systems risk producing outputs even when the underlying data is insufficient, noting this is where hallucinations become dangerous:

P2 – AI Hallucination Risk

"It will build [a visualization] even if the answer should be 'you simply don't have the data.' "

P4 raised concerns about **cognitive atrophy**, warning that over-reliance on AI tools may erode foundational analytical skills:

P4 – Cognitive Atrophy

"People will stop learning what needs to be done... there'll be a huge gap of, we have the tools and you don't know how to use them, just because you didn't learn the basics."

Finally, participants noted that the **gap between authors and users is narrowing**. The creation barrier is lowering (P2, P3), enabling non-specialists to generate visualizations. Yet human expertise remains essential—P4 stressed that "data is so particular to each problem, to each company, each system, that [AI] can never get it right."

5 Discussion and Conclusion

Our analysis reveals that dashboards are not disappearing, they are being augmented with conversational AI layers, shifting from static artifacts toward dynamic, question-driven interfaces.

From Builder to Conductor. The authoring role is transforming from technical execution to intent articulation and output validation. Authors are becoming "conductors" (P3) who orchestrate AI-supported workflows. This demands new competencies: framing analytical problems, evaluating AI outputs critically, and translating domain context into actionable specifications.

Trust Calibration as a Design Challenge. The polarization of user trust—between blind faith and wholesale rejection—represents a critical design challenge. Systems must support *trust calibration*: helping users develop evidence-based expectations of AI capabilities. Transparency mechanisms that make AI reasoning visible and auditable are essential.

Preserving Shared Analytical Ground. The ease of generating personalized dashboards risks fragmenting organizational truth (P2). Future systems must balance personalization with mechanisms that preserve common reference points for collaborative decision-making.

Design Implications. The future dashboard should incorporate advanced data processing capabilities that leverage agents to handle the entire pipeline from data processing to presentation within a single tool, enabling authors to build reliable dashboards in one place. This would also allow authors to select different agents for different tasks and orchestrate the data visualization process [2]. In other words, the future of the dashboard is an *agentic dashboard.*

The agentic dashboard should balance control and autonomy. Too much autonomy can leave authors feeling powerless over design decisions, while too much control from the author can severely restrict agentic capabilities. The agentic dashboard should also enforce visualization standards and rules to reduce hallucinations and maintain proper authoring standards. This would enable even end users to build visualizations without producing incorrect analysis results.

6 Future Work

Our findings are based on four expert interviews that provided rich qualitative insights. Future work should expand the participant pool to include end users, managers, and data engineers across diverse organizational contexts. Longitudinal studies tracking how authoring workflows, user behaviors, and trust dynamics evolve as AI-augmented dashboards are deployed would provide further empirical grounding.

Our results point toward mixed-initiative systems blending conversational AI with direct manipulation as a critical research direction, particularly around trust calibration and transparency. The concern that AI reliance may erode foundational analytical skills also warrants focused investigation comparing reasoning quality between AI-assisted and unassisted workflows. Finally, the tension between personalization and shared analytical ground has organizational implications that require further research.

Acknowledgments

This work was partially supported by Villum Investigator grant VL-54492 by Villum Fonden. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

References

[1] Benjamin Bach, Euan Freeman, Alfie Abdul-Rahman, Cagatay Turkay, Saiful Khan, Yulei Fan, and Min Chen. 2022. Dashboard Design Patterns. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 1–11. doi:10.1109/TVCG.2022.3209448

[2] Vaishali Dhanoa, Anton Wolter, Gabriela Molina León, Hans-Jörg Schulz, and Niklas Elmqvist. 2025. Agentic Visualization: Extracting Agent-Based Design Patterns From Visualization Systems. *IEEE Computer Graphics and Applications* 45, 6 (Nov. 2025), 89–100. doi:10.1109/MCG.2025.3607741

[3] Graham Elliott, Nikolay Kudrin, and Kaspar Wüthrich. 2022. Detecting p -Hacking. *Econometrica* 90, 2 (2022), 887–906. doi:10.3982/ECTA18583

[4] Sean Kandel, Andreas Paepcke, Joseph M. Hellerstein, and Jeffrey Heer. 2012. Enterprise Data Analysis and Visualization: An Interview Study. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2917–2926. doi:10.1109/TVCG.2012.219

[5] Eser Kandogan, Aruna Balakrishnan, Eben M. Haber, and Jeffrey S. Pierce. 2014. From Data to Insight: Work Practices of Analysts in the Enterprise. *IEEE Computer Graphics and Applications* 34, 5 (Sept. 2014), 42–50. doi:10.1109/MCG.2014.62

[6] Paula Maddigan and Teo Susnjak. 2023. Chat2VIS: Generating Data Visualizations via Natural Language Using ChatGPT, Codex and GPT-3 Large Language Models. *IEEE Access* 11 (2023), 45181–45193. doi:10.1109/ACCESS.2023.3274199

[7] Andrew M Mcnutt, Chenglong Wang, Robert A Deline, and Steven M. Drucker. 2023. On the Design of AI-powered Code Assistants for Notebooks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–16. doi:10.1145/3544548.3580940

[8] Arpit Narechania, Arjun Srinivasan, and John Stasko. 2021. NL4DV: A Toolkit for Generating Analytic Specifications for Data Visualization from Natural Language Queries. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 369–379. doi:10.1109/TVCG.2020.3030378

[9] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2019. What Do We Talk About When We Talk About Dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 682–692. doi:10.1109/TVCG.2018.2864903

[10] Melanie Tory and Vidya Setlur. 2019. Do What I Mean, Not What I Say! Design Considerations for Supporting Intent and Context in Analytical Conversation. In *2019 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, Vancouver, BC, Canada, 93–103. doi:10.1109/VAST47406.2019.8986918

[11] Jagoda Walny, Christian Frisson, Mieka West, Doris Kosminsky, Soren Knudsen, Sheelagh Carpendale, and Wesley Willett. 2020. Data Changes Everything: Challenges and Opportunities in Data Visualization Design Handoff. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 12–22. doi:10.1109/TVCG.2019.2934538

[12] Yun Wang, Zhitao Hou, Leixian Shen, Tongshuang Wu, Jiaqi Wang, He Huang, Haidong Zhang, and Dongmei Zhang. 2022. Towards Natural Language-Based Visualization Authoring. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 1–11. doi:10.1109/TVCG.2022.3209357

[13] Aoyu Wu, Yun Wang, Xinhuan Shu, Dominik Moritz, Weiwei Cui, Haidong Zhang, Dongmei Zhang, and Huamin Qu. 2022. AI4VIS: Survey on Artificial Intelligence Approaches for Data Visualization. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (Dec. 2022), 5049–5070. doi:10.1109/TVCG.2021.3099002