EuroVis Workshop on Visual Analytics (2026)
D. Archambault and A. Vilanova (Editors)

Position Paper: The Opportunity for Visual Analytics in the Age of Generative AI

S. Vink[1,2], K. Byungmoo[1], C. Brumar[1], M. Yang[1], K. Potter[3] and R. Chang[1]

[1]Tufts University, USA
[2]Utrecht University, Netherlands
[3]National Laboratory of the Rockies, USA

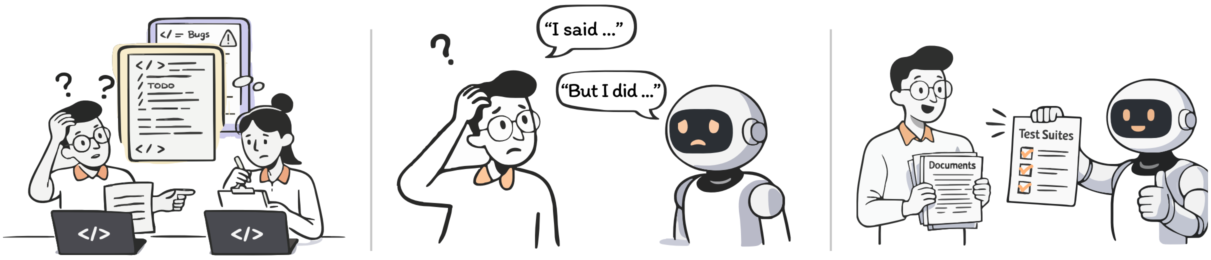

Figure 1: Rapid advances in AI are reshaping how analytical systems are built and used. These changes prompt a broader question for the Visual Analytics community: how should the field evolve as implementation becomes easier and AI becomes an analytical partner?

Abstract

AI is rapidly transforming knowledge work, raising epistemic questions about the future role of many established research areas. For the Visual Analytics (VA) community, these developments raise the same question: how do advances in AI affect the role and direction of the field? In this position paper, we argue that VA is particularly well positioned to respond this. First, we examine the implications of AI for system-oriented VA research. Recent advances in code generation by AI reduce the cost of software implementation, shifting attention toward problem formulation and evaluation. Rather than undermining the role of systems in VA research, this shift highlights that VA systems often serve to instantiate conceptual ideas or to support the evaluation of analytical methods. As implementation becomes easier, greater emphasis is placed on identifying meaningful analytical problems and evaluating whether proposed systems support human tasks. Second, we revisit the broader VA research agenda and show that understanding analytical problems and assessing how computational tools support human reasoning have always been central concerns of the field. Additionally, while the integration of AI introduces new opportunities and challenges, it remains aligned with VA's core objective of supporting humans. We conclude that, rather than disrupting VA, advances in AI reinforce the field's human-centered foundations and align the broader landscape of software development with the challenges VA has long studied.

CCS Concepts

• ***Human-centered computing*** → ***Visualization theory, concepts and paradigms;*** • ***Software and its engineering*** → *Software development methods;*

1. Introduction

Generative AI is rapidly transforming knowledge work across many industries [PKCD23, KYK*25]. It is changing how information is produced, analyzed, and acted upon, raising fundamental questions about the future of many established practices and research areas. For example, advances in code generation by LLMs have prompted concerns within parts of the programming languages community that some traditional research directions may lose relevance [CTJ*21]. For the Visual Analytics (VA) community, a similar question arises: how do these developments affect the role and direction of our field?

In this position paper, we argue that the VA community is in a particularly strong position to respond to these developments. VA was conceived as a discipline aimed at augmenting human analytical capabilities by integrating interactive visual representations with computational methods for exploring, transforming, and communicating data [Tho05]. The VA community has produced a substantial body of research spanning sensemaking, data abstractions, visual representations, interaction designs, and evaluation methodologies. Rather than undermining the foundations of the field, advances in AI highlight the continued importance of many of the core challenges that VA has long addressed. We examine the implications of this through two lenses.

First, we consider their impact on system-oriented VA research. A large portion of VA research is embodied in analytical systems. Generative AI introduces a structural shift in how such systems can be developed [Sta25]. LLM-based coding agents such as GitHub Copilot, Codex, and Claude Code can generate substantial portions of software systems from natural-language descriptions [KBL25]. Tasks that previously required extensive software engineering and weeks of effort can now be executed through iterative prompting and automated code synthesis in hours [PKCD23]. As a result, the cost of implementation decreases substantially. Within the software development lifecycle, this shift redistributes effort toward activities that precede and follow implementation, particularly *problem formulation* and *evaluation*. Rather than necessarily diminishing the role of systems in VA research, this development makes more explicit what these systems have long served as: artifacts for instantiating conceptual ideas and for evaluating analytical methods.

Second, we consider the implications of AI for the broader research agenda of VA. The scientific core of VA lies in understanding analytical problems, structuring complex tasks, designing representations that support reasoning, and evaluating how computational tools influence human judgment [Tho05, KAF*08]. Viewed through these lenses, AI does not fundamentally alter the objectives of the field. Instead, it changes the technological conditions under which those objectives are pursued.

In doing so, the integration of AI represents a logical continuation of VA's historical trajectory, while introducing new affordances that reshape the field's research focus. The reduced cost of implementation is only one dimension of this transformation. More profoundly, AI systems extend beyond traditional data modeling and analytic automation by acting as generative, dialogical, and adaptive agents within the analytical process. For example, AI systems increasingly function not merely as computational backends but as active components within analytical workflows. Such shifts demand renewed ways of thinking about agency, responsibility, and the distribution of analytical authority. Conversely, the architectural characteristics of AI introduce new epistemic challenges. Because generative models can produce plausible yet incorrect or misleading outputs due to hallucination and other systemic limitations, evaluation can no longer focus solely on usability, correctness of implementation, or task performance. Instead, it must grapple with uncertainty, calibration, trust, and the communicative framing of AI-generated insights. These conditions suggest the need for a re-examination of the research agenda in evaluation that systematically examines how analysts interpret, make sense of, and mitigate AI errors within interactive workflows. By foregrounding these emerging questions, we contend that the VA community must refine its standards and methodologies as the technological conditions under which its research agenda is instantiated continue to evolve.

Taken together, these developments suggest that advances in AI do not undermine the intellectual foundations of VA. Instead, they amplify the importance of the field's longstanding focus on understanding analytical problems, designing computational tools that support humans in analytical reasoning, and evaluating their impact on decision making. In this sense, the rise of AI aligns the broader landscape of software development with the core concerns that VA has addressed since its inception.

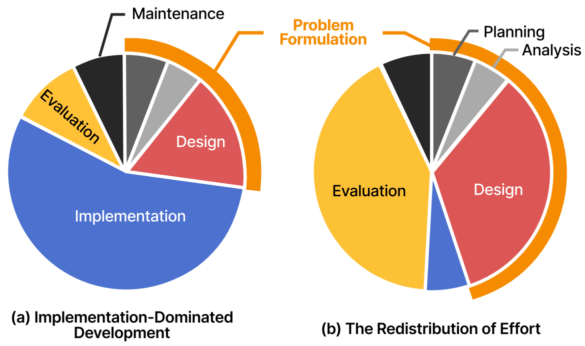

Figure 2: In the software development lifecycle, advances in code generation compress implementation effort and shift emphasis toward problem formulation and evaluation.

2. The Redistribution of Effort in VA System Development

The growing capability of generative AI in the form of LLM-based coding agents has begun to transform how software systems are created. For VA, where research contributions are frequently embodied in interactive systems [WDC*22, Sed16], these developments have direct implications for how systems are designed, implemented, and evaluated. Understanding these implications requires thinking about the structure of the VA system development process itself.

We do this by examining the software development lifecycle model [Man01, Pre05, Som11] that has traditionally structured how we think about systems development. It conceptualizes software creation as an iterative and incremental process. In its common formulation, iteration proceeds through six recurring phases: planning, analysis, design, implementation, evaluation, and maintenance. These phases are not strictly sequential but continuously inform one another across development cycles. For the purposes of this position paper, we treat the planning, analysis, and design phases collectively as *problem formulation*, alongside *implementation* and *evaluation* as the other two higher-level categories. We leave maintenance aside, as it is typically not a research objective in academic VA work.

LLM-based coding agents redistribute effort within the software development lifecycle. As implementation becomes faster for certain tasks, the relative effort devoted to it within the development process may decrease, although this does not necessarily reduce the overall complexity of developing VA systems. Greater relative emphasis falls on problem formulation and evaluation (see Figure 2), which shifts the primary attention both upstream toward specification and downstream toward evaluation. In particular, specifying complex interactive behavior and ensuring correctness may introduce new forms of effort that were previously embedded in implementation. In the upstream stage, the emphasis shifts to the precise specification of system requirements, as LLM-based coding agents require instructions that are sufficiently explicit to reduce ambiguity in generated outputs. Natural-language instructions must clearly describe analytical tasks, interaction logic, and data constraints in sufficient detail for coding agents to produce meaningful outputs. In the downstream stage, the emphasis shifts toward evaluation. As systems can be generated rapidly, the limiting factor becomes the ability to determine whether a generated system meaningfully and correctly supports the analytical task at hand.

3. The Central Role of Problem Formulation and Evaluation

This redistribution of effort highlights that the intellectual core of VA system research does not lie in the technical act of implementation. Rather, system-oriented work has long focused on identifying meaningful analytical problems and assessing whether proposed solutions effectively support them [Sed16, WDC*22]. In design studies, for example, systems serve primarily as concrete instantiations of conceptual ideas grounded in domain needs, enabling researchers to demonstrate how analytical tasks, abstractions, and interactions support real-world problem solving [MD19]. In other cases, systems function as experimental artifacts used to evaluate algorithms, representations, or interaction techniques.

When systems can be generated from high-level descriptions by LLMs, the specification of requirements and evaluation increasingly take concrete form in how the generation process is structured. In practice, the intellectual work shifts toward either specifying what the system should do or defining how its correctness and usefulness should be evaluated. While these approaches offer a useful way to structure generation, they also rely on the ability to precisely specify requirements and evaluation criteria, which can be challenging in complex VA settings. We observe two approaches that reflect these two activities: specification-driven and test-driven generation.

1. **Specification-Driven Generation.** In this paradigm, insights from problem formulation are translated into a structured requirements document that serves as the instruction set for the coding agent. The document specifies analytical tasks, data abstractions, interaction logic, constraints, and expected system behavior. The coding agent then generates an implementation that follows this specification. This approach aligns naturally with system-oriented work grounded in design studies, where the central challenge is translating domain needs into appropriate analytical abstractions and interactions.
2. **Test-Driven Generation.** In this paradigm, the researcher defines a comprehensive suite of tests that encode correctness, performance, and robustness requirements. This approach closely aligns with the test-driven development paradigm in software engineering [Bec03]. The generative model then produces an implementation that must satisfy these tests. It is particularly suitable for algorithmic contributions, where the central challenge is defining clear evaluation criteria that determine whether a method behaves correctly and improves measurable outcomes.

In practice, these two paradigms are often combined, as VA systems typically contain components that require both. Consider a VA system with an interactive interface that has a compiler for a domain-specific language. The overall architecture, abstractions, and interaction design can be defined through a structured product requirements document, while the compiler is governed by a suite of correctness and performance tests (e.g., unit and integration tests).

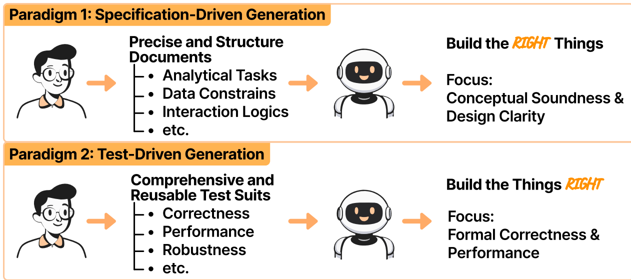

Figure 3: Two complementary paradigms for AI-assisted system generation. Systems can be produced either from explicit specifications of intended behavior (specification-driven generation) or from test suites that define correctness and performance criteria (test-driven generation).

Framed through Munzner's nested model for visualization design and validation [Mun09], the specification-driven paradigm primarily operates at the domain problem and data/operation abstraction levels, where analytical tasks and interaction concepts are defined. The test-driven paradigm operates at the algorithm level, where implementations are validated through executable tests such as unit and performance tests, without attempting to directly test the outer conceptual layers. The interface is thus generated from a conceptual specification of user needs, whereas the compiler is generated under strict testable requirements.

The implications of this shift are therefore not that VA system development becomes less relevant. On the contrary, the field is well positioned in this changing landscape. As LLM-based coding agents can reduce parts of the implementation effort, the relative value increasingly lies in problem formulation and evaluation. These activities correspond directly to the longstanding strengths of the VA community. Rather than necessarily disrupting system-oriented VA research, this development can be seen as realigning the broader landscape of software and analytical development with the questions that VA has long addressed.

4. Revisiting the Visual Analytics Research Agenda

Beyond the implications for VA system development, we now look at how advances in generative AI affect the broader VA research agenda. One of the earliest and most influential articulations of this agenda is the *Illuminating the Path* report [Tho05], which defined VA as the study of human reasoning facilitated by interactive visual interfaces. Viewed through this agenda, the central concern of VA research is understanding how humans reason with data and how computational tools can support that process. For example, analysts may use generative AI systems to propose hypotheses, generate transformations, or summarize findings, which introduces new dynamics in how reasoning unfolds. The field investigates how analytical problems are formulated, how tasks are structured, how visual and interactive representations support reasoning under uncertainty, and how these tools influence human judgment. In this sense, the scientific core of VA lies in understanding human cognition and evaluating how computational systems support sensemaking.

Proceedings published by Eurographics - The European Association for Computer Graphics.

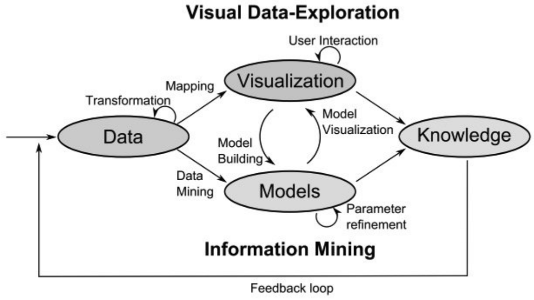

Figure 4: Keim's VA model [KAF*08] to illustrate that replacing "Models" with "AI" does not fundamentally alter the structure of the VA model, but highlights new challenges within it.

This perspective is also reflected in the VA process model proposed by Keim et al. [KAF*08], which describes analysis as an iterative process connecting data transformation, computational models, visualization, and human reasoning (Figure 4). The model explicitly positions the human analyst as the central element in the analytical loop, emphasizing that knowledge emerges through the interaction between computational processes and human interpretation.

These perspectives highlight a defining characteristic of VA: the central role of the human analyst. Computational methods provide powerful capabilities for interpreting data, but the objective of VA has never been to replace human reasoning. Instead, the field focuses on designing computational tools that augment human interpretation, allowing analysts to explore evidence, construct explanations, and make informed judgments in complex analytical settings. This human-centered orientation is particularly important because VA often addresses complex and ill-structured problems. Many domains studied in VA, such as fraud analysis or intelligence assessment, are characterized by evolving requirements, adversarial dynamics, and ambiguity. They do not admit clean or fully specified problem formulations. The objective is not merely to compute a solution, but to support humans in forming, revising, and defending judgments in changing contexts. Such problems inherently require human interpretation, domain expertise, and accountability.

From this perspective, the rise of generative AI does not necessarily constitute a categorical break with the past. The field has long incorporated increasing levels of automation, moving from database querying to data mining, and from statistical modeling to machine learning. AI represents a continuation of this trajectory, while also introducing new challenges for human-machine collaboration rather than fully redefining the field's objectives. For example, in fraud detection or cybersecurity, both defenders and adversaries increasingly have access to similar analytical and AI-driven tools. The degree of automation increases, but the fundamental objective of supporting human analytical reasoning in complex settings remains unchanged. At the same time, analysts must increasingly reason about both the data and the behavior of AI systems, introducing an additional layer of complexity. For this reason, the emergence of AI does not invalidate the VA research agenda. Instead, it changes the technological conditions under which this agenda is pursued.

5. Visual Analytics Research In The Era Of AI

Under these evolving technological conditions, the core components of the VA research agenda remain highly relevant, but may need to be reconsidered in light of AI-augmented analytical workflows. As AI systems increasingly participate in analytical tasks, they influence how reasoning is conducted, how analysts interact with computational tools, how data are processed and represented, and how analytical findings are communicated. In this section, we revisit four key components of the VA agenda [Tho05], and discuss how the inclusion of AI reshapes the challenges and design opportunities.

Science of Analytical Reasoning. The integration of AI into VA expands the participants and mechanisms involved in analytical reasoning, calling for a redefinition of collaborative analytics [HA08] and mixed-initiative systems [CCI*15].

Analytical reasoning is no longer solely human-centered: LLMs increasingly function as active components within analytical workflows. In many systems, they contribute to sensemaking by generating hypotheses, synthesizing evidence, and proposing interpretations [ZWX*24, ZSF*25, GGMM*24, HS20]. They may also suggest data operations, recommend procedures, or decompose complex tasks [ZWX*24, ZSF*25, GGMM*24, HS20]. In these roles, the AI system moves beyond passive assistance to become an active participant in the analytical process. However, their contributions are not always reliable and require careful interpretation and validation by analysts.

While powerful, these capabilities can shift analytical effort toward human–AI coordination. Unlike human collaborators [AHKGF11], AI systems can produce detailed traces of their reasoning processes. However, the volume and complexity of this information often exceed what analysts can interpret, creating new challenges for managing cognitive load and maintaining situational awareness.

Visual Representations and Interaction Technologies. The inclusion of AI as an analytic partner introduces new requirements for visualization and interaction in VA systems. Analysts need to monitor and guide AI components while accounting for hallucination, bias, and uncertainty. As AI systems gain greater agency, such as initiating analytical inquiries, suggesting transformations, or proposing new directions of exploration, analysts need mechanisms that make these processes visible and steerable. The central challenge is enabling AI initiative while maintaining human control.

Addressing this challenge requires visual representations that communicate not only the underlying data but also the reasoning processes of AI. As LLMs contribute intermediate steps, assumptions, and uncertainties, these processes must be exposed through visual and interactive representations. Prior work in explainable AI has explored techniques for visualizing model behavior and uncertainty [HKPC18, CMJ*20]. In AI-augmented VA systems, however, such representations introduce the need to become more central components of the analytic interface rather than auxiliary diagnostic tools, allowing analysts to inspect, question, and refine AI-generated outputs.

Interaction paradigms must also evolve to support this form of collaboration. Prompting alone is often insufficient because natural language input can be ambiguous and sensitive to subtle phrasing differences [Mor24, SSL*22]. Analysts require structured mechanisms that allow them to express intent, constrain model behavior, incorporate domain knowledge, and systematically explore alternative analytical paths. Hybrid interaction paradigms that combine visual controls, formal specifications, and iterative refinement can provide more precise ways to shape AI-assisted analysis while maintaining transparency and control.

Data Representation and Transformation. In addition to analytical reasoning, LLMs also participate in data cleaning, processing, and transformation. As AI systems increasingly assist in constructing data-processing pipelines, VA systems need to support transparency, traceability, and control over how data are transformed and represented. While AI-assisted tools can automate the construction of transformation workflows, they can also obscure how raw data are cleaned, structured, abstracted, and ultimately mapped to visual representations, making transparency and auditability essential for maintaining reliable analytical workflows.

For example, LLM-based coding agents can generate transformation pipelines directly from natural-language specifications. Although this lowers the barrier to building analytical workflows, it can obscure the logic governing how data are processed [WXC*25]. Decisions about filtering, aggregation, or restructuring may no longer be explicitly authored by analysts, making it difficult to understand how intermediate representations were derived. Transformation steps must therefore be logged and inspectable so analysts can audit how results were produced [Ye23, Bar25]. Documenting the sequence of transformations, from raw input to final representation, also supports reproducibility and allows analysts to verify intermediate steps and intervene when outputs diverge from analytical intent [SNTH13, GFI16, MNB*17, ZJW*25].

These concerns become particularly important when considering bias. Transformation decisions can embed assumptions about how data should be structured, filtered, or aggregated [MR22]. Generative models trained on historical data may reproduce dominant conventions that marginalize populations or suppress signals that deviate from expected patterns. Without mechanisms that expose transformation logic and support human intervention, such biases can propagate silently through the analytical workflow [Car24, JKV*22]. VA systems should therefore aim to surface the assumptions underlying data transformations and provide analysts with opportunities to inspect, question, and revise these decisions as part of the analytical process.

Production, Presentation, and Dissemination. The integration of AI into VA expands how analytical findings are produced and communicated, introducing new opportunities for automation while raising challenges for transparency, accountability, and trust.

AI systems can rapidly summarize analytical results, generate narrative explanations, and tailor reports to specific audiences [ZZZ*24, CLA*18, TKL*24], allowing analysts to translate exploratory findings into communicable outputs with minimal effort and share insights more broadly. However, the automation of analytical communication also introduces risks. AI-generated summaries may present findings with unwarranted confidence, omit important nuances, or frame results in ways that alter interpretation [HD11]. Because LLMs prioritize coherent narratives rather than epistemic accuracy, they may overstate conclusions or obscure uncertainty. Personalized reporting further complicates accountability. AI systems can adapt explanations and visual emphasis for different audiences [BGK*22], potentially leading stakeholders to receive differently framed interpretations of the same analysis. Ensuring that such communication remains faithful to the underlying evidence becomes an important design concern.

These challenges highlight the importance of maintaining clear analytical provenance [XOW*20]. Analysts and decision makers must be able to trace how conclusions were derived, including data transformations, reasoning steps, and AI contributions [Cri24]. Distinguishing between human and AI-generated content helps preserve accountability and supports auditability. As a result, trust, interpretability, and communicative fidelity become central design concerns [MPO25]. While dissemination may become increasingly dynamic, it should remain transparent and auditable to support defensible decision making.

6. Conclusion: It Has Always Been About Humans

The rapid progress of AI raises understandable questions about the future role of many research communities whose work has historically centered on knowledge work. For the VA community, these developments do not undermine the field's scientific foundations. Instead, they can reinforce the central importance of the problems that VA has long addressed. Throughout this paper we examined the implications of AI from two perspectives.

First, we considered the impact on system-oriented VA research. Advances in code generation by AI can reduce parts of the cost of software implementation and shift attention toward tasks that have traditionally defined the core of VA: identifying meaningful analytical problems and evaluating how computational tools support human reasoning. Second, we revisited the broader VA research agenda and argued that, while the technologies used to construct analytical systems continue to evolve, the central objective of supporting human reasoning remains unchanged.

The rise of AI introduces new opportunities and responsibilities. At the same time, these developments also introduce new complexities, as analysts must increasingly reason about both data and AI behavior. As analytical systems increasingly incorporate generative and adaptive capabilities, questions of trust, uncertainty, accountability, and human-AI collaboration become central. These questions lie squarely within the domain of VA, which has long investigated how humans interact with computational tools in analytical settings.

For this reason, the VA community is not merely well positioned to adapt to these developments. It is well positioned to help shape them. The field's longstanding expertise in problem formulation and evaluation provides a foundation for designing and assessing the next generation of AI-augmented analytical systems. In this sense, the growing capabilities of AI do not diminish the relevance of VA. Instead, they highlight how the challenges of designing and evaluating analytical systems remain central to the field.

Proceedings published by Eurographics - The European Association for Computer Graphics.

References

[AHKGF11] ARIAS-HERNANDEZ R., KAASTRA L. T., GREEN T. M., FISHER B.: Pair analytics: Capturing reasoning processes in collaborative visual analytics. In *2011 44th Hawaii international conference on system sciences* (2011), IEEE, pp. 1–10. 4

[Bar25] BARRAK A.: Traceability and accountability in role-specialized multi-agent llm pipelines. In *2025 40th IEEE/ACM International Conference on Automated Software Engineering Workshops (ASEW)* (2025), pp. 315–322. 5

[Bec03] BECK K.: *Test-driven development: by example*. Addison-Wesley Professional, 2003. 3

[BGK*22] BULLOCK B., GUO S., KOH E., ROSSI R., DU F., HOFFSWELL J.: Let's get personal: Exploring the design of personalized visualizations. In *2022 IEEE Visualization and Visual Analytics (VIS)* (2022), IEEE, pp. 85–89. 5

[Car24] CARNAT I.: Human, all too human: accounting for automation bias in generative large language models. *International Data Privacy Law 14*, 4 (2024), 299–314. 5

[CCI*15] COOK K., CRAMER N., ISRAEL D., WOLVERTON M., BRUCE J., BURTNER R., ENDERT A.: Mixed-initiative visual analytics using task-driven recommendations. In *2015 IEEE conference on visual analytics science and technology (VAST)* (2015), IEEE, pp. 9–16. 4

[CLA*18] CHEN S., LI J., ANDRIENKO G., ANDRIENKO N., WANG Y., NGUYEN P. H., TURKAY C.: Supporting story synthesis: Bridging the gap between visual analytics and storytelling. *IEEE transactions on visualization and computer graphics 26*, 7 (2018), 2499–2516. 5

[CMJ*20] CHATZIMPARMPAS A., MARTINS R. M., JUSUFI I., KUCHER K., ROSSI F., KERREN A.: The state of the art in enhancing trust in machine learning models with the use of visualizations. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 713–756. 4

[Cri24] CRISAN A.: We don't know how to assess llm contributions in vis/hci. *2024 IEEE Evaluation and Beyond-Methodological Approaches for Visualization (BELIV)* (2024), 115–118. 5

[CTJ*21] CHEN M., TWOREK J., JUN H., YUAN Q., PINTO H. P. D. O., KAPLAN J., EDWARDS H., BURDA Y., JOSEPH N., BROCKMAN G., ET AL.: Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021). 1

[GFI16] GOODMAN S. N., FANELLI D., IOANNIDIS J. P. A.: What does research reproducibility mean? *Science Translational Medicine 8*, 341 (2016), 341ps12–341ps12. 5

[GGMM*24] GU K., GRUNDE-MCLAUGHLIN M., MCNUTT A., HEER J., ALTHOFF T.: How do data analysts respond to ai assistance? a wizard-of-oz study. In *Proceedings of the 2024 CHI conference on human factors in computing systems* (2024), pp. 1–22. 4

[HA08] HEER J., AGRAWALA M.: Design considerations for collaborative visual analytics. *Information visualization 7*, 1 (2008), 49–62. 4

[HD11] HULLMAN J., DIAKOPOULOS N.: Visualization rhetoric: Framing effects in narrative visualization. *IEEE transactions on visualization and computer graphics 17*, 12 (2011), 2231–2240. 5

[HKPC18] HOHMAN F., KAHNG M., PIENTA R., CHAU D. H.: Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE transactions on visualization and computer graphics 25*, 8 (2018), 2674–2693. 4

[HS20] HAN W., SCHULZ H.-J.: Beyond trust building—calibrating trust in visual analytics. In *2020 IEEE workshop on trust and expertise in visual analytics (TREX)* (2020), IEEE, pp. 9–15. 4

[JKV*22] JAUNET T., KERVADEC C., VUILLEMOT R., ANTIPOV G., BACCOUCHE M., WOLF C.: Visqa: X-raying vision and language reasoning in transformers. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (2022), 976–986. 5

[KAF*08] KEIM D., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: Visual analytics: Definition, process, and challenges. In *Information visualization: Human-centered issues and perspectives*. Springer, 2008, pp. 154–175. 2, 4

[KBL25] KRINGS K., BOHN N. S., LUDWIG T.: (r) evolution of programming: Vibe coding as a post-coding paradigm. *arXiv preprint arXiv:2510.12364* (2025). 2

[KYK*25] KENNEDY B., YAM E., KIKUCHI E., PULA I., FUENTES J.: *How Americans View AI and Its Impact on People and Society*. Tech. rep., Pew Research Center, September 2025. URL: https://www.pewresearch.org/wp-content/uploads/sites/20/2025/09/PS_2025.9.15_AI-and-its-impact_report.pdf. 1

[Man01] MANIFESTO A.: Manifesto for agile software development, 2001. 2

[MD19] MEYER M., DYKES J.: Criteria for rigor in visualization design study. *IEEE transactions on visualization and computer graphics 26*, 1 (2019), 87–97. 3

[MNB*17] MUNAFÒ M. R., NOSEK B. A., BISHOP D. V., BUTTON K. S., CHAMBERS C. D., PERCIE DU SERT N., SIMONSOHN U., WAGENMAKERS E.-J., WARE J. J., IOANNIDIS J. P.: A manifesto for reproducible science. *Nature human behaviour 1*, 1 (2017), 0021. 5

[Mor24] MORRIS M. R.: Prompting considered harmful. *Communications of the ACM 67*, 12 (2024), 28–30. 5

[MPO25] MCKINLEY O. G., PANDEY S., OTTLEY A.: Trustworthy by design: The viewer's perspective on trust in data visualization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (2025), pp. 1–17. 5

[MR22] MOSCOVICH A., ROSSET S.: On the cross-validation bias due to unsupervised preprocessing. *Journal of the Royal Statistical Society Series B: Statistical Methodology 84*, 4 (06 2022), 1474–1502. 5

[Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics 15*, 6 (2009), 921–928. 3

[PKCD23] PENG S., KALLIAMVAKOU E., CIHON P., DEMIRER M.: The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* (2023). 1, 2

[Pre05] PRESSMAN R. S.: *Software engineering: a practitioner's approach*. Palgrave macmillan, 2005. 2

[Sed16] SEDLMAIR M.: Design study contributions come in different guises: Seven guiding scenarios. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), pp. 152–161. 2, 3

[SNTH13] SANDVE G. K., NEKRUTENKO A., TAYLOR J., HOVIG E.: Ten simple rules for reproducible computational research. *PLoS computational biology 9*, 10 (2013), e1003285. 5

[Som11] SOMMERVILLE I.: Software engineering 9th edition. *ISBN-10 137035152* (2011), 18. 2

[SSL*22] SHEN L., SHEN E., LUO Y., YANG X., HU X., ZHANG X., TAI Z., WANG J.: Towards natural language interfaces for data visualization: A survey. *IEEE transactions on visualization and computer graphics 29*, 6 (2022), 3121–3144. 5

[Sta25] STACK OVERFLOW: Stack overflow developer survey 2025. https://survey.stackoverflow.co/2025/, 2025. Accessed: January 2026. 2

[Tho05] THOMAS J. J.: *Illuminating the path:[the research and development agenda for visual analytics]*. IEEE Computer Society, 2005. 1, 2, 3, 4

[TKL*24] TANG X., KROKOS E., LIU C., DAVIDSON K., WHITLEY K., RAMAKRISHNAN N., NORTH C.: Steering llm summarization with visual workspaces for sensemaking. *arXiv preprint arXiv:2409.17289* (2024). 5

[WDC*22] WU A., DENG D., CHENG F., WU Y., LIU S., QU H.: In defence of visual analytics systems: Replies to critics. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (2022), 1026–1036. 2, 3

[WXC*25] WANG F., XI X., CUI Z., DAI H., WANG X.: Embedding traceability in large language model code generation: Towards trustworthy ai-augmented software engineering. In *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering* (New York, NY, USA, 2025), FSE Companion '25, Association for Computing Machinery, p. 1760–1763. 5

[XOW*20] XU K., OTTLEY A., WALCHSHOFER C., STREIT M., CHANG R., WENSKOVITCH J.: Survey on the analysis of user interactions and visualization provenance. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 757–783. 5

[Ye23] YE J.: Traceability link recovery for relations in natural language software architecture documentation and software architecture models, 2023. 5

[ZJW*25] ZENG Q., JIN C., WANG X., ZHENG Y., LI Q.: AIRepr: An analyst-inspector framework for evaluating reproducibility of LLMs in data science. In *Findings of the Association for Computational Linguistics: EMNLP 2025* (Suzhou, China, Nov. 2025), Christodoulopoulos C., Chakraborty T., Rose C., Peng V., (Eds.), Association for Computational Linguistics, pp. 10170–10201. 5

[ZSF*25] ZHAO Y., SHU X., FAN L., GAO L., ZHANG Y., CHEN S.: Proactiveva: Proactive visual analytics with llm-based ui agent. *IEEE Transactions on Visualization and Computer Graphics* (2025). 4

[ZWX*24] ZHAO Y., WANG J., XIANG L., ZHANG X., GUO Z., TURKAY C., ZHANG Y., CHEN S.: Lightva: Lightweight visual analytics with llm agent-based task planning and execution. *IEEE Transactions on Visualization and Computer Graphics 31*, 9 (2024), 6162–6177. 4

[ZZZ*24] ZHAO Y., ZHANG Y., ZHANG Y., ZHAO X., WANG J., SHAO Z., TURKAY C., CHEN S.: Leva: Using large language models to enhance visual analytics. *IEEE transactions on visualization and computer graphics 31*, 3 (2024), 1830–1847. 5

Proceedings published by Eurographics - The European Association for Computer Graphics.