



City Research Online

City St George's, University of London

Citation: Elskehaly, M., Jianu, R., Slingsby, A., Andrienko, G. & Andrienko, N. (2025). Designing for Collaboration: Visualization to Enable Human-LLM Analytical Partnership. IEEE Computer Graphics and Applications, 45(5), pp. 107-116. doi: 10.1109/MCG.2025.3583451

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/35390/>

Link to published version: <https://doi.org/10.1109/MCG.2025.3583451>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Designing for Collaboration: Visualization to Enable Human-LLM Analytical Partnership

Mai Elshehaly, *giCentre, City St George's, University of London, UK*

Radu Jianu, *giCentre, City St George's, University of London, UK*

Aidan Slingsby, *giCentre, City St George's, University of London, UK*

Gennady Andrienko, *giCentre, City St George's, University of London, UK*

Natalia Andrienko, *giCentre, City St George's, University of London, UK*

Abstract—

Visualization artifacts have long served as anchors for collaboration and knowledge transfer in data analysis. While effective for human-human collaboration, little is known about their role in capturing and externalizing knowledge when working with large language models (LLMs). Despite the growing role of LLMs in analytics, their linear text-based workflows limit the ability to structure artifacts into useful and traceable representations of the analytical process. We argue that dynamic visual representations of evolving analysis — organizing artifacts and provenance into semantic structures such as idea development and shifts in inquiry — are critical for effective human-LLM workflows. We demonstrate the current opportunities and limitations of using LLMs to track, structure, and visualize analytic processes, and propose a research agenda to leverage rapid advances in LLM capabilities. Our goal is to present a compelling argument for maximizing the role of visualization as a catalyst for more structured, transparent, and insightful human-LLM analytical interactions.

The need for capturing, externalizing, and visualizing accumulated knowledge during data analytics has long been advocated in the Visual Analytics (VA) literature. The seminal report by Thomas and Cook [8] emphasized the need for “*knowledge representations to capture, store, and reuse the knowledge generated throughout the entire analytic process*” [8, p. 42]. Effectively perceivable and explicitly structured knowledge representations, e.g., in the form of *knowledge graphs*, are instrumental for supporting collaboration, communicating analysis results, maintaining provenance, and guiding analytical processes [10]. However, eliciting pre-existing knowledge and capturing new insights generated during analysis remains a persistent challenge [1, sec. 6.6], not least because it cannot be automatically captured easily [9].

The rapid development of Large Language Models


(LLMs) is having impacts across many disciplines, and visual analytics is no exception. There has been significant work investigating the use of LLMs as intelligent user interfaces to support visual analytics [5], including processing data, running analytical techniques, producing visualisations and even extracting insights. However, we see more profound opportunities for LLMs’ reasoning capabilities to **track the activities conducted by human analysts and capture the knowledge generated from these activities**. Namely, LLMs can be leveraged to capture the **provenance** of analytic artifacts, the **state of knowledge** accrued, and the **process** of knowledge evolution. This is becoming more feasible as LLMs advance their ability to track, structure, and externalize knowledge dynamically.

We envision an intelligent assistant that captures and visualizes the process of knowledge co-construction between human and LLM during an analysis workflow. We call for more research into how LLMs can help achieve this goal and capture analysts’ inter-

actions, while augmenting their capabilities to externalize their own thought processes, analytical intent, and progress of their investigation. We see opportunities for LLMs to enable the dynamic generation of these knowledge constructs and reduce the overhead incurred by traditional approaches to analytic provenance. We also see great opportunities for visualization research to shape the agenda for human-LLM interfaces that support their analytical partnership.

VISION: Visualized LLM-supported Workflow

- 1) **Initialization:** The analyst poses a question, uploads relevant datasets, and provides references to background information materials, initializing the knowledge model and base representation.
- 2) **Dynamic Evolution:** As the analysis progresses, every action (including prompts, data transformations, visualizations, hypothesis refinements, observations, comments, etc.) is captured by an LLM-powered Analysis and Knowledge Modelling Module. The system continuously updates the current state of analysis and knowledge, tags each artifact with its provenance, and tracks the entire analytical workflow.
- 3) **Visualization and Interaction:** A visualization layer presents the evolving analysis process (e.g., as a timeline, flowchart, or data narrative), the current state of knowledge (e.g., as an interactive knowledge graph), and on-demand provenance links for deeper exploration. Analysts and LLMs use and modify these multi-layered representations to: (i) structure and deepen their understanding of the knowledge being created, (ii) extend and refine the analysis, and (iii) share context and enhance communication.
- 4) **Collaboration and Tailored Representations:** The system facilitates dynamic partnership between analysts and the LLM by enabling adaptive, personalised knowledge representations tailored to different expertise levels and analytical roles. The LLM can generate alternative views of the knowledge model, adjusting the level of detail, terminology, and visualization formats to align with the backgrounds and specializations of different experts. The system also supports collaborative review and annotation, allowing analysts and the LLM to jointly tackle inaccuracies, refine interpretations, suggest new inquiries, and iteratively improve the analysis. Comprehensive reports capturing both human and LLM contributions ensure transparency, reproducibility, and effective knowledge sharing.



This paper promotes the argument that VA research must capitalise on these opportunities. We distinguish between the role of LLMs in tracking the analytical process, maintaining evolving knowledge, and recording provenance and the role of visualisation in representing this information to human analysts and supporting interaction through visual interfaces. We advocate for visualisation research that builds on the

strengths of LLMs in summarization and organization while focusing on effective visual representations of artifacts **provenance**, knowledge **states**, and analysis **process** in LLM-assisted analytics.

To support our argument and test the feasibility of our ideas, we present a proof-of-concept exploration workflow that demonstrates LLM's capabilities in (a) understanding and manipulating human-produced knowledge visualization; and (b) producing its own knowledge visualization. Our experiment is not intended to present an exhaustive characterisation of the fast-evolving LLM capabilities or a typology of knowledge structures it can capture and visualize. Instead, our aim is to demonstrate, with a concrete example, the potential and limitations in the current state of visual knowledge exchange between human and LLM, and offer a perspective on research directions to enable refining, evolving, reusing, and transforming visual analytic artifacts with LLMs.

Opportunities for Visualization in Human-LLM Partnership

For centuries, people have relied on knowledge externalisation artifacts—such as diagrams, whiteboards, post-it notes, and tokens—to represent, explain, and share mental models, facilitate referencing, and establish a shared context for collaboration. As LLMs become more sophisticated, they are shifting from mere tools to active partners in analysis. *Our premise is that just as external representations scaffold human collaboration and knowledge construction (as established in [1]), they can similarly catalyse human-LLM co-construction of knowledge*, providing a shared framework for structured reasoning and interactive, iterative refinement.

Figure 1 presents an extension of the VA knowledge construction framework in [1] to capture this role. In this updated framework, the knowledge model becomes distributed between human cognition and the LLM. Visualization is not only aimed to support the analyst's understanding of data, but also of the entire process of LLM-assisted knowledge construction. VA research has long promoted the benefits of such knowledge externalisation. Specifically, visualisation systems have traditionally sought to fulfil three main knowledge externalisation requirements:

- R1 **Provenance Tracking** – documenting the lineage and semantic rationale behind analytical artifacts.
- R2 **State Representation** – reflecting the current knowledge, including established findings and remaining gaps.




FIGURE 1. We build on an established framework that regards VA as a model building activity [1]. LLM-related components of knowledge construction are highlighted in blue and enclosed with dashed lines. This diagrammatic representation reveals the anticipated role of LLMs in communicating with the analyst and capturing their knowledge, process, and provenance of artifacts.

R3 Process Representation – capturing the overall sequence of analytical transformations.

The advent and increasing use of LLMs in analysis adds a new dimension to such efforts. First, because Natural Language (NL) is their primary input, analytical processes naturally become explicit. Via prompts, *analysts express their thought processes as they specify goals, reflect on outcomes, and define iterations*. This makes thought externalisation an inherent part of analysis rather than an added burden, and creates the premise for enriching captured provenance with deeper semantic context. Second, LLMs excel at summarising, aggregating, and restructuring knowledge, *reducing the effort needed to construct, moderate, and maintain externalised knowledge structures*.

This section highlights opportunities and limitations in state-of-the-art to support the three requirements above and desiderata for visualisation research to better support the human-LLM data analytic partnership.

R1: Provenance Tracking

Understanding the process by which an analytic insight is generated is often referred to as “insight provenance”. More broadly, “analytic provenance” captures different semantic layers of this process, from low-level keyboard strokes and mouse events to high-level sensemaking tasks and analytical intent [9]. The purpose of visualisation has been to reduce cognitive overhead for analysts by automating the capture and representation of these steps to support recall, replication, action recovery, collaborative communication, presentation, and meta-analysis [6]. Specific to human-




FIGURE 2. The PromptTHIS system visualizes an image variant graph to support provenance in artistic creation [4].

AI collaboration, Guo *et al.* expanded the definition of provenance to include “*text and images involved in the prompt history*” in an AI-supported creative process [4]. The purpose of provenance in this context is to enable the organisation, review, and comparison of pairs of human-generated and AI-generated variants of knowledge or artistic artifacts, and the decisions made throughout their production (e.g., Figure 2).

This semantic layer of provenance is particularly central to our vision, as it provides deeper insight into the reasoning, intent, and rationale behind analytical choices. Instead of merely logging interactions and outputs, provenance should capture the motivations and




FIGURE 3. A visual language for visualisation of knowledge state: nodes, links, tags, comments, and references [10].

justifications behind transformations or refinements, offering a structured representation of decision-making. Extracting rationales seamlessly from analyst-LLM dialogues presents an opportunity to integrate richer semantic context into provenance without overburdening analysts. Additionally, LLMs can facilitate contextual linking between interactions, prompts, and generated insights, enhancing the interpretability and traceability of the analytical workflow.

Thus our **Research Desideratum [D1]** for provenance is to develop a framework for interactive multi-level provenance visualisation that captures and traces the full analytic lineage—including prompts, interactions, and data transformations—while enriching provenance with semantic context, explicitly representing the intents, rationales, and interpretations driving analytical decisions.

R2: State Representation

The need for capturing, externalising, and visualising the **state** of accumulated knowledge during an analytic process was repeatedly advocated in the VA literature [2, 3]. Such state can be expressed in the form of rules, models, and computational constructs that inform data selection, analyses, and interpretation (e.g., Figure 3). The importance of knowledge state externalization is further discussed in [1], which reviews state-of-the-art approaches for representing the analyst’s mental model. A key challenge consistently highlighted is that *“the level of facilitation of the knowledge externalisation job is insufficient; it still requires much effort of the analyst beside the analysis itself.”* [1]. While the role of visualization in knowledge internalization (e.g., sensemaking) is well established, its role in supporting knowledge externalization remains less clear [3].

Similar to provenance tracking, LLMs can assist in capturing and maintaining the evolving state of knowledge, thereby reducing the manual effort required from analysts. Their ability to synthesize, summarize, and

organize knowledge structures facilitates a more seamless co-construction process, where humans and LLMs collaboratively refine and update shared knowledge.

It is important to emphasise the personal nature of knowledge construction. We hypothesise that design aesthetics that account for elements of **affect** could enhance analysts’ engagement with externalizing their knowledge. Recognizing the analyst’s experience and cognitive involvement in this process may inspire new approaches to facilitating knowledge externalization.

Beyond individual reasoning, externalized analysis states also function as a shared medium of communication between human analysts and LLMs. By providing a structured, persistent context, these representations help sustain analytical conversations, ensuring that both the analyst and the LLM can refer to, interpret, and build upon prior knowledge. To fulfil this role, visualization of externalized analysis states should be designed with LLM interpretability in mind, and incorporate interactive mechanisms that facilitate fluid exchanges between analysts and LLMs.

Thus, our **Research Desideratum [D2]** for analysis state visualisation is to develop visually compelling, adaptive representations of analysis states that serve as both a shared medium for human-LLM communication and a scaffold for reasoning. These representations should capture the cognitive and emotional dimensions of the analyst’s knowledge acquisition process to motivate active externalisation and sharing of evolving thoughts.

R3: Process Representation

Current visual analytics systems typically represent the analysis process as a sequence of *visual states* corresponding to particular configurations of a visual display. Transitions between states are defined by the interactive operations that modify the display. This representation, often referred to as a “navigation view”, allows analysts to retrace their steps, revisit previous states, and branch into new lines of exploration [7]. While being useful, this view primarily operates at a low level, focusing on changes in the visualisation itself rather than the underlying evolution of knowledge.

Our vision shifts the focus from a record of interface states to a conceptual model of knowledge evolution, where consecutive knowledge states are explicitly linked by reasoning processes and derivation of new analytical artifacts. By structuring the analytical process in terms of evolving insights, hypotheses, and discoveries rather than just visual transformations, we aim to support a more semantic, high-level understanding of how knowledge develops throughout an analysis.

Thus, our **Research Desideratum [D3]** for process is to develop a unified, interactive process visualisation framework that dynamically integrates provenance and state representations into a coherent, multi-layered narrative of the entire analytic process, enabling users to trace lines of reasoning, transformation sequences, compare alternative analytic paths, and iteratively refine their workflow in real time.

Proof-of-concept Exploration

To offer a preliminary illustration of our vision's feasibility, we engaged GPT-4 in a process of visual knowledge exchange between human and LLM. While not a systematic evaluation, this exploration demonstrates the potential of current LLMs to support structured knowledge co-construction. First, we supplied GPT-4 with a detailed specification of the desired LLM behaviour as an analytic partner. A full listing of the prompt, the LLM's response and subsequent prompts can be found at <https://dx.doi.org/10.17605/OSF.IO/SU9MB>. Subsequent prompts instructed the LLM to break down the problem into smaller parts including: dynamic knowledge representation (STEP 1), and progressive refinement towards specific capabilities, such as identifying insights, externalizing knowledge, and visualizing processes (STEP 2).

Step 1: Dynamic Knowledge Representation

We provided GPT-4 with three manually created graphs illustrating stages of a data analytic process: initial exploration, extraction of relevant information, and hypothesis testing. Figure 4 (left) shows an example graph representing the first stage of analysis. The model was able to interpret the provided graphs, identifying key analytic elements such as questions, observations, hypotheses, and relationships between them. It also generated semi-structured JSON-like representations capturing the evolving analysis process. Although the captured structures were partial and lacked full semantic depth, the results indicate that LLMs can assist in externalizing analytic states and provenance in a form amenable to visualization. This presents a good starting point for our research desiderata.

We argue that the visualization of knowledge **state** should also include high-level knowledge constructs, in addition to the captured elemental knowledge artifacts. Such higher constructs should include **pattern artifacts**, derived from collections of information such as temporal and spatial patterns, **higher-order knowledge artifacts**, such as *argument*, *causality*, *estimation*, etc. and complex reasoning constructs, such as *hypothesis*, *rationale*, *conclusion* or *inference*

(deduced from some premises). It is also important to capture the attributes of knowledge such as *confidence* (e.g., agreement/ importance per user), and *origin* (references to external sources), etc.

Next, we prompted the LLM to expand a data-analytic scenario ("a story"). A full listing of the LLM-generated story is included in supplementary material. Here, we focus on elements of interactive visualization that the LLM was able to capture (Figure 4 (right)).

Node Sequencing- The LLM detected a sequential order among graph nodes that represent knowledge states. It assigned numeric labels to those nodes in the generated scenario. Figure 4 (right) overlays the LLM-generated node numbers on the graph.

Edge Labels- The LLM identified new labels for some of the existing graph edges. For example, it defined a *refinement* relationship between the *initial_questions* and *refined_questions*.

New Nodes- The LLM generated new nodes that were not in the initial representation of each **state**. Node 7 in Figure 4 (right) is an example of an LLM-added node, capturing observations found by the analyst from data ("*high workplace mobility correlates with rising cases; stay at home policies reduce cases*").

New Edges- The LLM introduced new types of relationships such as *context*, *requires_analysis*, and *supports_hypothesis* (see Figure 4).

Interaction- Several interactions were described by the LLM including: (i) highlighting nodes (e.g., the original question) to adjust relationships or propose refinements; (ii) filtering nodes, e.g., to focus on the nodes relevant to the user's refined question; (iii) graph annotation, linking findings like *periodic variations* to supporting datasets; and (iv) graphical interfaces. In addition to these graphical interactions, LLM-supported natural language interactions were also described like: (i) drilling down into specific observations and their provenance; (ii) using the LLM to generate alternative hypotheses, such as "*Does vaccination coverage affect patterns?*"; and (iii) exploring counterexamples or anomalies flagged by the LLM.

Key Takeaways- The LLM understood the graphical representation of the knowledge generation process and used our manually-created examples to generate a natural language description of an analytic workflow. It was able to produce some, but not all, necessary nodes and edges for knowledge representation. It was also able to describe anchors for collaboration (e.g., sharing with colleagues who can annotate, suggest refinements, or query for additional insights), and human communication (e.g., present findings to stakeholders).




FIGURE 4. Manual representation of the first phase of the analysis (left) and the LLM-envisioned variant (right). Dashed arrows represent new relationships generated by the LLM.

The scenario suggests that LLM reasoning can support multi-level provenance visualization (relevant to D1). This was shown through its ability to describe possible interactions that connect with users' analytic intent. The cognitive dimension of the knowledge acquisition process (relevant to D2) was captured in its description of the analysts' iteration (e.g., refining questions and hypotheses, linking observations to hypotheses within the graph, etc.). Finally, the LLM produced a multi-layered narrative (albeit an incomplete one) to capture a coherent sequence of analytic steps that make up the **process** (relevant to D3). It was able to integrate nodes representing both knowledge and analysis states in its natural language representation.

The examples supplied to the LLM in this step are meant to communicate expectations and requirements for an intelligent assistant that is capable of generating dynamic knowledge representations ready for visualization. However, these examples are limited in number and diversity. Therefore, we do not expect them to showcase the full capability of what LLMs are able to produce. Nevertheless, we anticipate that the LLM would be able to produce similar representations when presented with a new data analytic problem.

Step 2: Knowledge Visualization

Our next set of prompts aimed to assess the LLM's capability to apply the reasoning it learned in Step 1 to generate visual knowledge representations for a new data analytic problem. We supplied the LLM with an email trail in which a real-world data analytic problem was discussed. The emails were from stakeholders at a local authority describing questions they would like to answer from data about air quality across their district.

Similar to Step 1, the LLM was able to successfully extract textual descriptions of the knowledge nodes that existed in the email trail. The extracted nodes in-

cluded research questions, assumptions, data sources, etc. However, when asked to create a JSON structure of this knowledge graph and generate code for its visualization, the results demonstrated the LLM's limited capability in both generating and visualizing the knowledge graph.

Figure 5 shows an example of the produced graph. The LLM was able to produce JavaScript code to visualize the graph as a Sankey diagram. It was not able to suggest more appropriate design alternatives for representing the process. Additionally, there were several limitations in knowledge capture and representation. For example, the LLM was not able to connect data and task nodes with any other node categories. Therefore, these nodes are represented as disconnected blank nodes to the right of the graph. It was not able to encode different types of relationships in the visualization, despite being able to describe *some* types of relationships in natural language. Information relating to data provenance was very limited, despite the email trail having sufficient detail about the data sources and capturing mechanisms.

Research Agenda

Although our vision for LLM-mediated externalizations of analytical processes to enhance human-LLM dialogue is clear, how these externalizations should be (i) modeled and visualized; (ii) evolved over the course of an analysis; and (iii) used to support analysis activities is far from obvious. Here, we discuss these themes and highlight research opportunities to tackle them.

(i) Modeling and representing analyses

Large context window sizes of modern LLMs enables them to track the evolution of knowledge and maintain coherence across multiple complex workflows. In addition, LLMs can integrate external memory with




FIGURE 5. LLM-generated knowledge graph for a new data analytic process.

knowledge retrieval mechanisms to cross-reference prior findings, obtain additional information, and build a structured provenance trail. A question then is:

RQ1: How can we effectively externalize analytical processes in ways that are amenable to visualization and interpretation by humans? First, flexible knowledge representation models are needed to represent non-linear analytic processes in ways interpretable to humans and usable by LLMs. Temporal Knowledge Graphs (TKGs) show promise, but their effectiveness needs further exploration. Second, visualizations must adapt to different goals: analysis, communication, or collaboration. They should highlight understanding, uncertainties, and rationale, while tailoring detail to specific users. For example, analysts track and address uncertainties, while stakeholders need a narrative linking evidence to decisions.

RQ2: How can we effectively represent the multi-modal and non-linear nature of real-world analysis? Analytic building blocks span text, data, code, and visuals. Insights often arise from visual patterns, like trends or anomalies. These patterns should be treated as first-class artifacts, and linked to reasoning and provenance. Such non-linear and multi-modal relationships among artifacts present several visualization challenges. Furthermore, the very definition of analytic building blocks can be challenging. For example, they can follow normative frameworks (e.g., goals, methods, evidence) for structure or emerge organically through interaction, reflecting natural reasoning. A hybrid approach can also offer both rigour and flexibility, while adapting to real-world workflows.

RQ3: How can we balance granularity and usability when capturing and visualizing complex analysis histories? Capturing analysis in too much detail can lead to complex, hard-to-use models. Key choices include whether to include minor steps, dead ends,

or aggregate insights. Visualizations should balance clarity and depth, using overviews with expandable detail. Focus+context techniques can highlight relevant content while preserving the full analytical flow.

(ii) Co-evolution of Knowledge

In our proof-of-concept exploration, we demonstrated LLM capabilities for extracting analysis structures from static visual and NL descriptions of data analytic problems. However, in real-world analytic scenarios we envision that for the most part analysis structures would evolve continuously and in full symbiosis with text prompting and visual interactions.

Designing effective mechanisms for integrating LLMs into the creation and evolution of diagrammatic representations presents several research questions.

RQ4: What interactions best support the co-construction and assessment of analytical structures with LLMs? Interactions should enable fluid and intuitive changes to the analysis structure, keeping the process smooth and natural. Models should make uncertainty and ambiguity explicit while enhancing interpretability and supporting decision-making. Further research is needed to explore how LLMs and interactive visualization can best augment human reasoning.

RQ5: How can LLMs integrate updates across visual and textual representations while maintaining consistency and interpretability? LLMs should integrate prompt elements into existing structures without losing key details, ensuring that analysts can detect inconsistencies or omissions. Furthermore, LLMs must cohesively combine text-based and visual representations, optimizing multi-modal reasoning to maintain consistency and interpretability in model updates.

(iii) Collaboration and Evaluation

We envision that co-evolved diagrams of the analytical process, created collaboratively by human analysts and LLMs, could facilitate multi-modal interactions and

serve as a shared reference, much like traditional schematics help establish common ground in collaborative problem-solving.

Analysts could use these diagrams to seamlessly integrate visual elements into their prompting strategies, enhancing the fluidity and precision of interactions. However, the benefits would not be one-directional. Instead of relying solely on text-based context, the LLM could reference specific, well-defined diagram elements when making analytical suggestions, validating insights, or tracking changes. Updates, refinements, and recommendations generated by the LLM could thus be explicitly anchored to visual components, making it easier for analysts to track the model's contributions, detect hallucinations, and understand how analytic artifacts integrate into the evolving analysis.

Beyond aiding human understanding, structured schematics could also enhance LLM reasoning. Human analysts benefit from externalized visual structures to organize thoughts and processes [4], and similarly, supplementing prompts with structured visualizations could improve the LLM's ability to retain key elements, reduce ambiguity, and enhance reasoning. By embedding analysis diagrams into the collaborative workflow, we can create a more transparent, interpretable, and cognitively effective interaction paradigm for human-LLM analytical partnerships. A key research question must be addressed before we can reap these benefits:

RQ6: How can interaction design and evaluation methodologies be leveraged to establish the effectiveness of multi-modal interactions in human-LLM partnerships? Interaction design should support intuitive exchanges, with LLMs navigating and referencing structured representations. Effective prompting techniques are key for meaningful interactions with both text and diagrams. A question here is: *How can interaction design increase the effectiveness of the human-LLM collaboration through these diagrammatic representations?* Finally, LLMs may reason better with structured diagrammatic context. Evaluation methods should assess whether this hybrid approach improves analytical workflows and outcomes. Therefore, we must answer the question: *How can the impact of multi-modal interactions on analytical outcomes be evaluated to establish the effectiveness of the human-LLM collaboration?*

CONCLUSION

While visualization artifacts have been shown to improve human-human collaboration, their role in human-

LLM partnership remains underexplored. Our exploration highlights how visualization can serve as an anchor for this partnership, enabling LLMs to both understand and produce structured representations of knowledge and knowledge generation processes. By structuring analytical workflows through visual knowledge representations, we can enhance transparency, traceability, and shared understanding between human analysts and LLMs.

However, realizing this vision requires addressing several challenges. We propose an agenda and desiderata for new frameworks for provenance tracking, state representation, and process visualization that seamlessly integrate with LLM-driven analysis. Further, research is needed to refine LLMs' ability to generate, manipulate, and reason over structured visual artifacts dynamically. By advancing these capabilities, we can move beyond linear, text-based interactions to richer, more intuitive multi-modal collaborations.

Ultimately, our viewpoint highlights the potential of visualization to elevate human-LLM analytics into a more structured, interactive, and interpretable process. We call for deeper interdisciplinary research to develop visualization-driven methodologies that maximize LLMs' analytical contributions while ensuring human agency, oversight, and interpretability remain central. This agenda paves the way for the next generation of intelligent analytical assistants, where visual representations act as cognitive scaffolds for complex reasoning, decision-making, and collaborative knowledge discovery.

ACKNOWLEDGMENT

This work was supported by Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as part of the *Lamarr Institute for Machine Learning and Artificial Intelligence* (Lamarr22B), and by EU in project *CrexData* (grant agreement no. 101092749).

References

- [1] N. Andrienko et al. "Viewing Visual Analytics as Model Building". In: *Computer Graphics Forum* 37.6 (2018), pp. 275–299. DOI: <https://doi.org/10.1111/cgf.13324>.
- [2] Chaomei Chen. "Top 10 unsolved information visualization problems". In: *IEEE computer graphics and applications* 25.4 (2005), pp. 12–16. DOI: [10.1109/MCG.2005.91](https://doi.org/10.1109/MCG.2005.91).

- [3] Paolo Federico et al. "The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics". In: *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2017, pp. 92–103. DOI: [10.1109/VAST.2017.8585498](https://doi.org/10.1109/VAST.2017.8585498).
- [4] Yuhan Guo et al. *PromptThis: Visualizing the Process and Influence of Prompt Editing during Text-to-Image Creation*. 2024. arXiv: [2403.09615](https://arxiv.org/abs/2403.09615).
- [5] Maeve Hutchinson et al. "LLM-Assisted Visual Analytics: Opportunities and Challenges". In: (2024). arXiv: [2409.02691](https://arxiv.org/abs/2409.02691).
- [6] Eric D Ragan et al. "Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes". In: *IEEE transactions on visualization and computer graphics* 22.1 (2015), pp. 31–40. DOI: [10.1109/TVCG.2015.2467551](https://doi.org/10.1109/TVCG.2015.2467551).
- [7] Yedendra Babu Shrinivasan and Jarke J. van Wijk. "Supporting the analytical reasoning process in information visualization". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. Florence, Italy: Association for Computing Machinery, 2008, pp. 1237–1246. DOI: [10.1145/1357054.1357247](https://doi.org/10.1145/1357054.1357247).
- [8] J.J. Thomas and K.A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE, 2005.
- [9] Kai Xu et al. "Analytic provenance for sense-making: A research agenda". In: *IEEE computer graphics and applications* 35.3 (2015), pp. 56–64. DOI: [10.1109/MCG.2015.50](https://doi.org/10.1109/MCG.2015.50).
- [10] J. Zhao et al. "Supporting Handoff in Asynchronous Collaborative Sensemaking Using Knowledge-Transfer Graphs". In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (2018), pp. 340–350. DOI: [10.1109/TVCG.2017.2745279](https://doi.org/10.1109/TVCG.2017.2745279).

Mai Elshehaly is a Senior Lecturer in Data Visualization at giCentre, Department of Computer Science, City St. George's, University of London. She holds a Ph.D. and M.Sc. in Computer Science from the Center for Human-Computer Interaction at Virginia Tech. She is a member of IEEE. Contact her at Mai.Elshehaly@city.ac.uk

Radu Jianu is a Lecturer in Data Visualization at giCentre, Department of Computer Science, City St. George's, University of London. Contact him at Radu.Jianu@city.ac.uk.

Aidan Slingsby is a Senior Lecturer in Visual and Analytic Computing at giCentre, Department for Computer Science, City St. George's, University of London, London, U.K. Contact him at a.slingsby@city.ac.uk.

Gennady Andrienko is a lead scientist responsible for visual analytics research with Fraunhofer Institute IAIS in Sankt Augustin, Germany, a Co-PI of the Lamarr Institute for Machine Learning and Artificial Intelligence, and a professor with City St. George's, University of London, London, U.K. Contact him at gennady.andrienko@iais.fraunhofer.de.

Natalia Andrienko is a lead scientist with Fraunhofer Institute IAIS in Sankt Augustin, Germany, the Area Chair and Co-PI of the Lamarr Institute for Machine Learning and Artificial Intelligence, and a professor with City St. George's, University of London, London, U.K. Contact her at natalia.andrienko@iais.fraunhofer.de.