Exposing LLM-Mediated Data Analysis: Early Visualization Designs

Radu Jianu*, Maeve Hutchinson†, Natalia Andrienko‡, Gennady Andrienko§, Mai Elshehaly¶, and Aidan Slingsby‖

City, St George's, University of London, United Kingdom

Abstract

We extend earlier work on LLM-supported analysis mapping to explore visualization techniques for representing analytic conversations as evolving networks of concepts and artefacts. Building on prior demonstrations that LLMs can extract and maintain semantic structures from analysis transcripts, we investigate how to render these networks in ways that capture temporal sequence, provenance, and iterative refinement. We describe design choices—such as temporally aware layouts, importance propagation, content-rich nodes, and linked conversation timelines—and discuss their potential to support reflection, communication, and provenance tracking in LLM-assisted data analysis. Our contribution is exploratory: we present a set of visualization strategies, highlight their limitations, and situate them as incremental steps toward richer visualizations of human–LLM analytical partnerships.

Index Terms: visual analytics, LLM, networks, modeling analysis, provenance

1 Introduction

Data analysis is inherently iterative—analysts refine questions, test hypotheses, interpret visualizations, and revise assumptions. When this process unfolds through dialogue with LLMs, much of the reasoning, outcomes and analysis artifacts becomes hidden within long, sequential transcripts. Recent research highlights this as both a challenge and an opportunity. For instance, InsightLens demonstrates the benefit of structuring analyst–LLM conversations around emergent insights, providing clarity and easing navigation through complex dialogic content [11]. Likewise, we recently proposed a vision for improving visual data analysis with LLMs, emphasizing the need to represent not only final outcomes but the entire analytic trajectory—capturing reasoning, artefacts, and provenance as a structured, visual object [2].

As a first step toward this vision, we have shown that *analysis maps*—semantic networks that capture the structure of an analysis, with nodes representing questions, datasets, tasks, and findings, and edges expressing their interrelations—can be automatically constructed by LLMs from analyst–LLM dialogues. We tested this approach on analyses conducted entirely through interaction with LLMs, and contributed a methodology for experimenting with different decomposition and prompting strategies to guide the construction of such maps. Importantly, we found that LLMs can reflect on analytic dialogues either post-hoc, by processing captured transcripts, or interactively during live analysis, and in both cases produce meaningful analysis maps [5].

Here, we shift from feasibility to visualization. Our goal is to design visual encodings that make analysis maps easier to interpret, richer with meaning, and more reflective of how analytic conversations unfold. To that end, we experiment with temporally aligned layouts, importance scaling, content-rich nodes that embed analytic artefacts *in situ*, and seamless conversational linkage. These elements together support overviewing analytic trajectories, spotlight provenance, and connect reasoning to its conversational and artefactual grounding—thus bringing us closer to the vision of dynamic, interactive analytic boundary objects.

The contributions of this paper are exploratory but foundational. We present a prototype that illustrates how such visualizations might look in practice, while also helping us reason about challenges, opportunities, and directions for future work.

2 Related Work

A long-standing theme in visual analytics is the externalization of analytic reasoning. Cook and Thomas framed this as a central challenge: how can systems make the invisible process of analysis visible, communicable, and reproducible [10]. Andrienko and Andrienko conceptualized analysis as a process of model building. Federico et al. and Andrienko et al. showed that externalization and structuring of reasoning support interpretability, collaboration, and communication of results [3, 1]. Shrinivasan et al. argued for visualization as a medium for external cognition, where representing reasoning explicitly supports memory and sensemaking [8]. Practical systems such as Jigsaw [9] also demonstrate that encoding analytic reasoning in structured visual forms can enhance both individual and collaborative exploration.

Capturing *provenance* has likewise been recognized as crucial in visual analytics. Provenance helps analysts recall, reproduce, and communicate their reasoning, while also supporting collaboration and knowledge transfer. Ragan et al. provide a comprehensive survey of provenance in visual analytics, covering representations from low-level interaction logs to higher-level conceptual graphs of questions and hypotheses [7]. Other systems have explored interactive histories and knowledge-transfer mechanisms, further illustrating how provenance structures can externalize analytic process and make it reusable [4, 6, 12].

Against this backdrop, our *vision paper* argued that the rise of LLM-mediated analysis creates both a need and an opportunity to revisit these challenges [2]. Dialogues with LLMs naturally capture rich streams of analyst intent, reasoning, and artefacts, but they remain buried in linear transcripts. We proposed that visual analytic research should advance in three directions: (i) provenance tracking that links ideas and artefacts back to their conversational origins, (ii) knowledge state representation that externalizes evolving analytic structures, and (iii) process representation that makes the dynamics of reasoning visible. This framing shares motivation with *InsightLens*, which organizes analyst–LLM conversations around emergent insights to provide clarity and navigability [11].

In recent work we took first steps toward this vision [5]. We showed that LLMs can decompose analytic dialogues into semantic networks of questions, tasks, datasets, and findings, constructed either retrospectively (from archived transcripts) or interactively (during live analysis). This study established feasibility and contributed a methodology for experimenting with decomposition and prompting strategies, but did not yet address how such networks should be visualized.

*e-mail: radu.jianu@citystgeorges.ac.uk
†e-mail: maeve.hutchinson@citystgeorges.ac.uk
‡e-mail: natalia.andrienko@citystgeorges.ac.uk
§e-mail: gennady.andrienko@citystgeorges.ac.uk
¶e-mail: mai.elshehaly@citystgeorges.ac.uk
‖e-mail: a.slingsby@citystgeorges.ac.uk

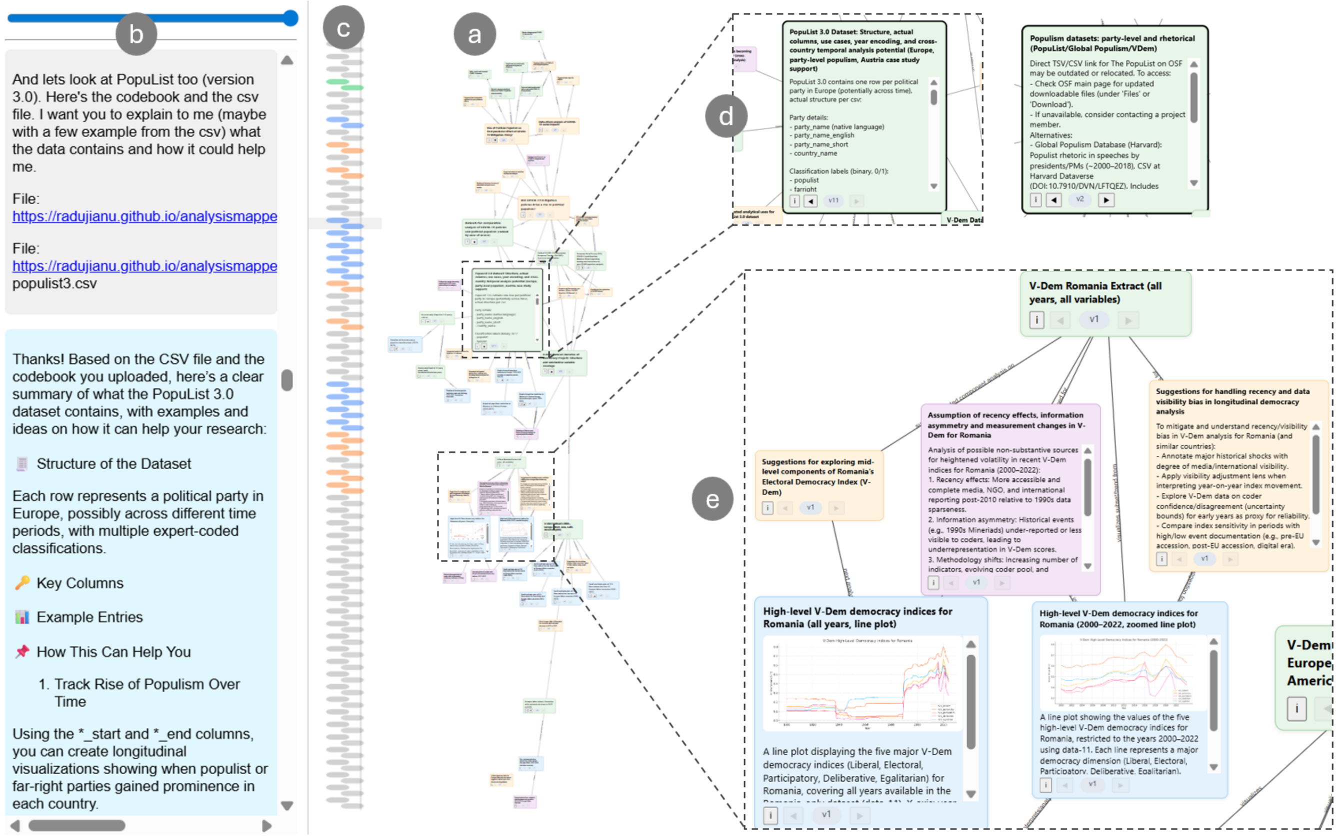

Figure 1: Prototype visualization of *analysis maps*. The system integrates three coordinated views: (b) the chat history showing the analyst–LLM dialogue, (a) the evolving analysis network, aligned vertically to reflect temporal progression and revealing clusters that correspond to distinct analytic foci, and (c) a conversation timeline, also aligned to temporal progression, that captures the full duration of the session and displays interlocked request–response pairs. Selecting a node (d) reveals its temporal footprint on the timeline: green marks creation, blue marks refinements, and orange marks references. Frequent blue marks indicate iterative updates; the node view shows the latest version, while earlier versions can be accessed via navigation arrows at the bottom of the node. Clicking on an exchange in the timeline scrolls the chat history, while scrolling the chat highlights the corresponding region on the timeline. Panel (e) illustrates a local analysis thread on Romanian democracy indicators, with expanded nodes for two visualizations, an observation, and a suggestion.

Here we build on that foundation. Rather than focusing on whether networks can be generated, we ask how they can be visualized to expose trajectory, provenance, and linkage to artefacts and conversation. In doing so, we begin addressing the research challenges posed in our vision paper—particularly the calls for provenance tracking and process representation—while also complementing approaches such as InsightLens by experimenting with richer, multi-faceted representations of analytic reasoning.

3 Methods

3.1 Mapping Data-Analysis Dialogues

We used test data from our previous work [5]. There, two exploratory analyses were carried out entirely through dialogue with an LLM. One analysis examined the relationship between COVID-19 and populism, lasting about four hours and producing 96 exchanges, while the other focused on stop-and-search practices in London, lasting around three hours with 57 exchanges. Analyses (exchanges and artefacts produced - tables, plots) were captured in full to create a corpus for offline experimentation. Analyses conducted fully through LLM interaction are not unusual, with recent research such as on InsightLens [11] adopting a similar setup.

We then used LLMs to distill these transcripts into structured, semantic *analysis maps*. To this end we experimented with different generation (and by extension prompting) strategies. These included (i) *segmentation*—how much of the dialogue is interpreted in a single step; (ii) *concept refinement*—how iterative refinements to the same analytic concepts are captured in the network; and (iii) the use of lightweight *corrections*—small adjustments supplied by the analyst to fix suboptimal mappings. The resulting maps can be explored online at `https://observablehq.com/@rdjianu/mindmapping-analyses`.

For this study, we used the networks generated with the most effective strategies: **paired segmentation**, where each analyst request and corresponding LLM response are processed together, and **in-place refinement**, where subsequent refinements to a concept are integrated into the same node rather than producing new ones. We also applied a small number of **corrections** (e.g., merging duplicate nodes or renaming them when labels were unclear).

Analysis maps were captured as *dynamic networks*. Each analyst–LLM exchange resulted in an incremental update consisting of (i) new nodes and links (for newly introduced concepts), (ii) edits to existing nodes (to capture refinements to concepts already discussed), and (iii) references to previously mentioned nodes (concepts). Analysis resources or artefacts such as tables and plots were attached as links to external resources in an online repository. Our visualization system was designed to integrate these updates progressively, so that the evolving network reflected the unfolding of the analysis over time.

3.2 Visual Design

The extracted networks provided a semantic backbone, but our goal was to design a visualization that would make them useful to analysts. Building on the desiderata articulated in our earlier vision paper [2], we focused on three aims: *(i)* providing a clear overview of the trajectory of the analysis, *(ii)* capturing provenance of ideas and artefacts as they were introduced and refined, and *(iii)* tightly linking the evolving network to the analytic conversation itself. The following design elements reflect these goals and are illustrated in Figure 1.

Temporally Aligned Layout: We used a force-directed network layout powered by D3, but with an added constraint to emphasize how the analysis unfolded over time. Specifically, each node's vertical position is tied to its temporal footprint in the dialogue: we compute the average of the indices where the node was created, updated, or referred to, and use this as its target y coordinate. The force-directed layouter then pulls the node toward this position while still optimizing for readability. This hybrid approach yields a coherent network structure while visually aligning the analysis map with its chronological development (Figure 1a). To further aid interpretation, nodes are color-coded by type: *green* for datasets, *blue* for charts and visualizations, *orange* for analytic process elements such as research questions or goals, and *purple* for observations, insights, and assumptions.

Conversation Timeline: Alongside the network, we show a timeline of the analytic dialogue. Selecting a node highlights the conversational exchanges where it was created (green), updated (blue), or referred to (orange) (Figure 1c). Brushing and linking synchronize the timeline with the chat window: scrolling the transcript highlights its visible range on the timeline; conversely clicking a timeline point scrolls the transcript to that exchange. This connection reinforces provenance and allows users to retrace analytic reasoning seamlessly.

Importance Propagation: In any extended analysis, some concepts are only touched upon briefly or mentioned in passing—for instance, a suggested course of action that is never followed up, or an inconsequential observation. Others become central, being revisited, refined, and exerting influence over subsequent reasoning. As large analysis maps can easily become visually cluttered, it is important to make influential concepts salient.

To this end, we assign an *importance score* to each node and map it to the node's visual size. A node's score increases whenever it is created, updated, or referred to. In addition, a portion of this score is recursively propagated to earlier connected nodes—those that preceded and influenced it—so that the impact of central ideas is reflected not only in their own prominence but also in the prominence of the concepts that inspired them. This propagation allows viewers to quickly identify which ideas drove the analysis forward.

Content-rich Nodes: Nodes are rendered as HTML containers that can be collapsed to a concise label or expanded to reveal extended description and analytic artifacts (plots, tables, code). This supports an *information-on-demand* approach, where detail is available when needed but does not overwhelm the overall view. Analytic artifacts can be viewed *in situ*, directly within the flow of the analysis, rather than requiring navigation to separate resources. The mechanism further helps users decide which parts of the analysis should be more salient: important nodes—or those currently under discussion—can be expanded and kept open, while less relevant ones can remain collapsed. The layout dynamically rearranges itself to make space for enlarged nodes, ensuring that the expanded content remains legible without obscuring the surrounding network (Figure 1e).

Nodes also contain all of their different versions. By default, the final version is displayed, but users can toggle through earlier states if they wish to see how a concept evolved (Figure 1d). This version history makes it possible to reconstruct the iterative refinement of ideas, helping to trace not only end results but also the steps by which they were reached.

4 Learnings and Discussion

We previously introduced three key aims for our visualization design: (i) providing an overview of the trajectory of the analysis, (ii) capturing provenance of ideas and artifacts as they were introduced and refined, and (iii) tightly linking the evolving network to the analytic conversation itself. Here we reflect on what we learned with respect to these aims, discuss open questions, and outline directions for future work.

Analysis Overview and Artifacts: Temporal alignment seems effective in making the unfolding of the analysis legible. As shown in Figure 1, clusters of related activity emerge naturally and the vertical ordering reveals when different analytic strands were explored. The artifacts displayed *in situ* further enriched this overview, allowing tables and graphs to be interpreted directly within the flow of the analysis. We see promise in these approaches. At the same time however, such analysis maps remain complex, and narrative structures are not trivial to extract. Future work could explore additional structuring devices, such as grouping nodes hierarchically, supporting multiple node levels, or organizing maps around higher-level narratives (similar to data stories). More broadly, it remains an open question whether analytic processes should be represented around multiple types of information (concepts, artefacts, and insights together), or whether maps should instead emphasize one aspect, such as insights (as in InsightLens), artefacts, or narrative structures.

Provenance and Evolution: Synchronizing nodes with the dialogue reinforced provenance, letting us trace how ideas were introduced, refined, and reused. However, iterative updates to nodes, while accessible, are not trivial to track as it is difficult to see what had changed between toggle-able versions. A more explicit "diff" representation, showing what information was added or removed at each step, could make concept evolution clearer.

Interaction: For now, our prototype supported only passive exploration. A natural next step is to enable interactive use in live analysis sessions, where users can point to nodes, refer to them directly in conversation, or request that the LLM restructure or summarize their content. Our earlier work already showed that LLMs can manipulate extracted networks. The richer node representations here suggest opportunities for deeper interaction, where co-construction applies to how information is structured within nodes.

Limitations: As with our earlier study, these findings are preliminary. They are based on only two analyses conducted by the authors, with no systematic evaluation beyond our own impressions. Heuristics such as temporal alignment and importance propagation were assessed informally. These constraints limit the generality of our conclusions, but the prototype nonetheless illustrates potential directions for representing analytic processes.

5 Conclusion

We extended our previous work on generating analysis maps from analyst–LLM dialogue with a focus on visualization. The enhancements introduced here include temporally aligned layouts, importance-based node scaling, content-rich nodes that display artifacts *in situ*, and integration with a conversation timeline. Together, these elements provide clearer overviews of the analytic trajectory, capture provenance, and connect reasoning to its conversational grounding. Although our findings remain preliminary, they illustrate how dynamic analysis maps can represent the unfolding of analytic processes and point toward systems that support their interactive use in practice.

REFERENCES

[1] N. Andrienko, T. Lammarsch, G. Andrienko, G. Fuchs, D. Keim, S. Miksch, and A. Rind. Viewing visual analytics as model building. In *Comput. Graph. Forum*, vol. 37, pp. 275–299, 2018. 1

[2] M. Elshehaly, R. Jianu, A. Singsby, G. Andrienko, and N. Andrienko. Designing for collaboration: Visualization to enable human-llm analytical partnership. *IEEE Comput. Graph. Appl. (to appear)*, 2025. 1, 3

[3] P. Federico, M. Wagner, A. Rind, A. Amor-Amoros, S. Miksch, and W. Aigner. The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics. In *Proc. IEEE Conf. Visual Analytics Sci. Technol. (VAST)*, pp. 92–103, 2017. 1

[4] D. Gotz and M. X. Zhou. Characterizing users' visual analytic activity for insight provenance. In *2008 IEEE symposium on visual analytics science and technology*, pp. 123–130. IEEE, 2008. 1

[5] R. Jianu, M. Hutchinson, N. Andrienko, G. Andrienko, M. Elshehaly, and A. Slingby. Mind-mapping data analysis with llms: From vision to first steps. In *Proc. Computer Grapihcs and Visual Computing (CGVC)*, p. to appear, 2025. 1, 2

[6] C. North, R. Chang, A. Endert, W. Dou, R. May, B. Pike, and G. Fink. Analytic provenance: process+ interaction+ insight. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pp. 33–36. 2011. 1

[7] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen. Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. *IEEE Trans. Vis. Comput. Graph.*, 22(1):31–40, 2015. 1

[8] Y. B. Shrinivasan and J. J. Van Wijk. Supporting the analytical reasoning process in information visualization. In *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, pp. 1237–1246, 2008. 1

[9] J. Stasko, C. Gorg, Z. Liu, and K. Singhal. Jigsaw: Supporting investigative analysis through interactive visualization. In *Proc. IEEE Symp. Vis. Analytics Sci. Technol. (VAST)*, pp. 131–138, 2007. 1

[10] J. J. Thomas and K. A. Cook, eds. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Comput. Soc., 2005. 1

[11] L. Weng, X. Wang, J. Lu, Y. Feng, Y. Liu, H. Feng, D. Huang, and W. Chen. Insightlens: Augmenting llm-powered data analysis with interactive insight management and navigation. *IEEE Trans. Vis. Comput. Graph.*, 31(6):3719–3731, 2025. 1, 2

[12] J. Zhao, M. Glueck, P. Isenberg, F. Chevalier, and A. Khan. Supporting handoff in asynchronous collaborative sensemaking using knowledge-transfer graphs. *IEEE Trans. Vis. Comput. Graph.*, 24(1):340–350, 2017. 1