Prompt Lenses: Improving the Magic of Lenses (for Text Analysis)

Jena Satkunarajan[1], Phillip Wohlfart[1], Samuel Beck[1], Max Franke[2], Steffen Koch[1]

[1]University of Stuttgart, Germany
[2]IsarNet Software Solutions GmbH, Germany

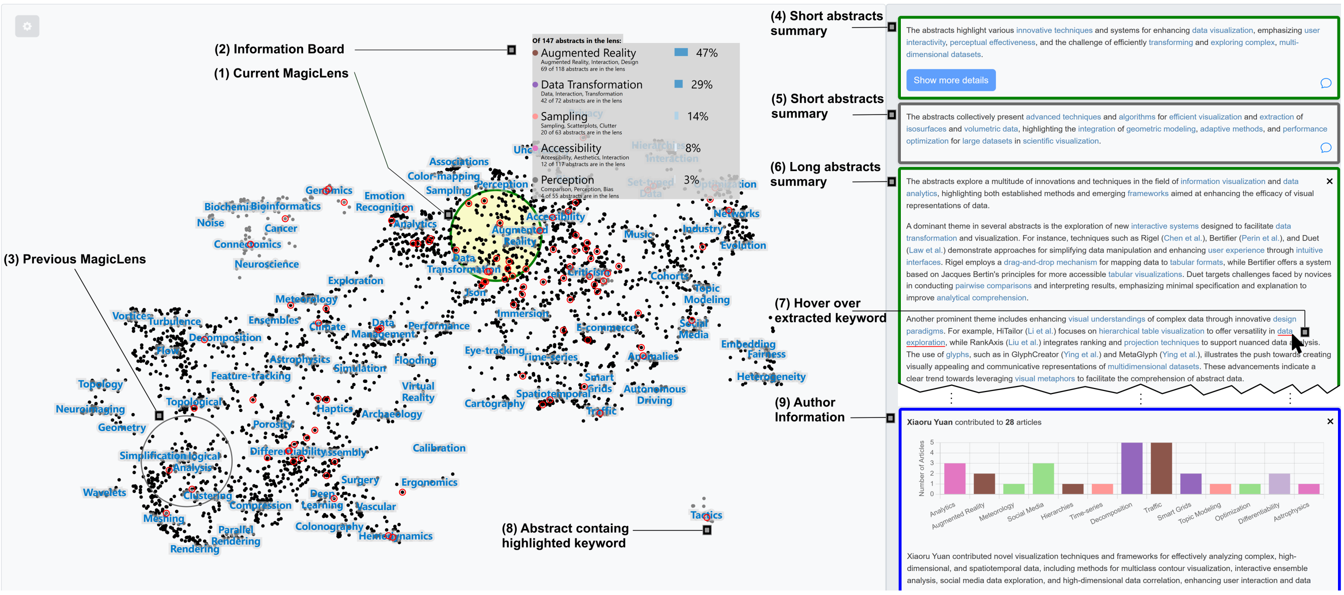

Figure 1: *The Prompt Lens (1) is placed over a dense region with multiple clusters of IEEE VIS abstracts represented in a scatterplot. A tooltip (2) shows the distribution of precomputed cluster labels for hovered abstracts. A separate text box (4) shows a brief LLM-generated summary of all abstracts under the lens. The previous lens position is indicated on the scatter plot (3), and the summary of those abstracts (5) can be compared with the current selection. A detailed summary is shown on demand (6), and important keywords are highlighted and linked back (7) to abstract representations containing the keyword (8). Detailed author information (9) is displayed when clicking on author names.*

Abstract

Incorporating the analytical power of LLMs with the fast-paced interaction of magic lens-based exploration is an intriguing prospect. Unfortunately, the costs of LLM-generated analyses are high, and applying them continuously seems prohibitive at the moment. Accordingly, we suggest an LLM integration into magic lenses that supports the progressive triggering of costly analyses based on users' interest in the data hovered with the lens. We exemplify this approach with a lens technique for exploring dimensionality-reduced embeddings of visualization paper abstracts shown in a scatterplot. Our proposed approach links back analysis results to the explored visualization improving the comprehensibility and the assessment of the shown results.

CCS Concepts

• ***Human-centered computing*** → ***Visual analytics;***

1. Introduction

Large Language Models (LLMs) add a powerful method to improve interactive analysis, helping users understand and explore visualized data [ZZZ*25]. Interfaces for interacting with LLMs are often based on human natural language dialog, allowing for "chatting" with the model but making the interaction slow. Interactive

DOI: 10.2312/evs.20251086

visualization systems, however, benefit or even rely on the "asymmetry in data rates" [War19] regarding the amount of information that is conveyed to users as compared to their interactive input: With relatively few interactions, often realized as direct manipulations [Shn83], users gain much visual information efficiently. This work aims at improving the data rates by incorporating LLM-based analyses with magic lenses [TGK*17] to support fast exploratory analysis with the powerful capabilities of LLMs. We achieve this by carefully developing prompts that are executed based on the data shown in a visualization to which a lens is applied. This leads to a "define prompt once, apply often" type of approach that combines powerful models with fast-paced interaction. To keep the interaction smooth, we reduce computational costs by adapting the LLM-calls to the user's degree of interest. We employ a progressive [FFS24], multi-stage approach where precomputed information is shown first before less expensive prompts are executed while more expensive ones can be triggered on demand. We improve the traceability of AI-generated results by linking these back to the data hovered by the lens(es). This facilitates fast back-tracking, linking of the information to improve understanding and to assess the generated output. Subsequently, we exemplify our idea by integrating LLMs into magic lenses for exploring abstracts of IEEE VIS papers [IHK*17]. The proposed lens-based solution is employed on a scatterplot that depicts dimensionality-reduced [MHM20] embeddings [RG19] of the abstracts as was suggested in previous works, e.g., by Raval et al. [RWVW23]. By employing a flexible prompt template, we support several analysis goals for this usage scenario, including summarization and comparison of abstracts under one or different lenses or contrastive comparisons. Our contributions comprise: (**1**) the combination of LLM responses with a lens-based technique supporting tasks such as document summary and comparison, (**2**) linking of textual LLM output to the visualization for better exploration and plausibility checking and (**3**) enabling fast LLM-based analysis depending on the user's degree of interest by considering the LLM's computational cost.

2. Background and Related Work

Applying magic lenses [Fur86] in combination with visualization has been extensively researched in the last decades [TGK*17]. In particular, recent works suggested their usage for exploratory data analysis (EDA) in the context of multi-dimensional data visualization. For example, lenses can be used to convey details for regions of interest that show the data in simplified or aggregated form. During the exploration, users indicate their interest through a continuous drag (and drop) interaction of the lens, which typically adds information about the visual data items hovered by the lens. The data under the lens undergoes a task-specific transformation, and the results are often presented locally either in the form of focus+context or cue-based approaches [CKB09]. For the latter, variants exist showing details next to the lens [HJH*16]. Our approach realizes a hybrid variant where brief visual and textual cues are provided next to the lens while more detailed, textual information is shown in a separate view (see Figure 1). Integrating machine learning (ML) approaches with interactive visualization has been a cornerstone of many visual analytics techniques developed in recent years [YCY*21]. Vis4ML and ML4Vis [SDBEA*23] describe the combination of visualization and ML techniques, where our approach clearly addresses the latter case. With the capabilities of foundation models, their incorporation into EDA approaches is a logical consequence [ZZZ*25]. We propose integrating LLMs into magic lenses [TGK*17] to profit from both fast, interactive exploration and more powerful analysis options. Incorporating processing functions into magic lens approaches is limited by performance requirements to facilitate immediate user feedback on the "magical" analysis performed by the lens. Consequently, approaches that incorporate ML results focus on less expensive techniques [SMSL17], preprocessed results, or progressive solutions [KKP*17].

To demonstrate our approach, we picked a text document exploration scenario using paper abstracts of vispubdata.org [IHK*17]. Both visual text document exploration [KK15] and the visual analysis of scientific documents [FHKM17] are addressed by many works in the domain. Recent approaches leverage state-of-the-art generative models and their increased reasoning abilities [HTK23] to enrich the exploration and analysis of textual data. Many approaches leverage dimensionality reduction to depict documents' vectors and embeddings in scatterplots [PNML08, CRF*21] and document landscapes [WTP*95, FK14, VMZL22]. In recent years, non-linear projections such as UMAP [MHM20] turned out to be suitable solutions. We apply our magic lens approach to a scatterplot based on a UMAP projection of text embeddings.

Heimerl et al. [HJH*16] introduce a magic lens over a 2D document landscape representing IEEE VIS papers that provides dynamic key term-based and visual cues to explain the data under the lens. Wang et al. [WHC23] provide a map-based visualization of dimensionality-reduced embeddings and adaptive visual summaries that consider the user's current zoom level to determine the level of summary. We employ a comparable approach to create a map-like document landscape and integrate progressive lens interaction that considers the user's degree of interest. As a key difference, our magic lens enables dynamic local and contrastive analysis, where data points under the lens can be compared against the rest of the data. Raval et al. [RWVW23] probe LLMs for explanations for text embeddings, where users lasso-select clusters of interest to compare LLM results with phrases extracted with a traditional model. Nomic Atlas [Nom22] is a commercial tool that combines dimensionality reduction and layout optimization in an interactive landscape for the AI-assisted exploration of unstructured data. We also employ LLM probing and document clustering, but focus on fast, progressive analysis with lens-based interaction. With flexible prompt templates, we support different text-based tasks and provide interactive, internal links for plausibility checks.

3. Approach

We apply our magic lens approach to a spatial mapping of IEEE VIS paper abstracts, where we compute high-dimensional text embeddings using Sentence-BERT [RG19] before we reduce the embeddings to 2D-representations with UMAP [MHM20]. This results in similarity-based projections of the abstracts, which we visualize as a scatterplot. Subsequently, we apply spectral clustering [NJW01] to determine clusters based on the projection of abstracts. Using the abstracts as input, we annotate clusters with LLM-generated labels to provide static user guidance [CGM*17].

Additional steps include the characterization of authors and generating (additional) cluster labels, again using respective prompts (see supplemental material). The summaries are presented to users when inspecting individual abstracts to reduce reading effort. Authors are also described textually with a brief summarization of their contributions to different topics. We use OpenAI's *gpt-4o-mini* [Ope24] model for all LLM-based requests since we did not see a big difference in applying more powerful models such as *gpt-4o*. Still, the approach can be configured to use more powerful models for tasks if required.

3.1. Progressive Data Analysis

Our magic lens incorporates LLM requests to analyze the data under the lens. We designed a multi-stage pipeline that applies different prompts based on users' degree of interest. Costlier prompts are applied as more complex tasks are performed to foster fast exploration and deeper analysis on demand. For arbitrary requests, the computational cost and the resulting response times can render smooth lens-based exploration impossible and significantly delay analyses, even with powerful LLMs that can run on local machines. CO_2 emissions and increased financial costs are other aspects we would like to avoid. Consequently, we apply a progressive analysis approach that employs more expensive approaches only when a user indicates a higher interest in a region hovered by the lens. The interactive lens can be placed and moved in the scatterplot, similar to a drag-and-drop interaction. The displacement of the lens triggers a three-step analysis of the abstracts under the lens. Our three-step pipeline is designed to align with increasing user engagement, progressing from a broad overview for casual interest, to a focused overview for intriguing data points, and finally, detailed analysis for strong user interest.

Analysis Step 1 First, the topical distribution of the abstracts under the lens is displayed in a tooltip next to the lens. The distribution appears as a sorted list of the most prominent clusters by the number of associated abstracts enclosed within the lens. While the lens is dragged across the scatterplot, the tooltip is continuously updated according to the new data under the lens. Since the information is based on the pre-computed clusters, determining cluster association is fast, and immediate feedback can be provided as users explore the data.

Analysis Step 2 Dropping the lens over a group of data points triggers an LLM analysis of the data. The LLM is prompted to summarize all affected abstracts with a focus on identifying key facts and phrases to briefly represent the distinct contents of the abstracts. Depending on the number of abstracts, such a request can take a few seconds to process and display the results. We perform the summary task after users settle their exploration on a particular group of data they are interested in, to keep interaction smooth and avoid redundant processing costs. We limit short summaries to a maximum of 300 words to reduce the cost of the prompts.

Analysis Step 3 If the short summary spurs the user's interest, they can explicitly request the third, most costly analysis. This triggers another LLM request to generate a more detailed summary of the abstract contents. The extended summary comprises comprehensive explanations of the central research topics found in the selected abstracts, where key phrases and influential works are highlighted. For the works, we instruct the LLM to return a reference as *[first author] et al.* to backlink the reference to the data point (see 3.3). We allow extended summaries to be up to 600 words long, which requires costlier prompts that generate more tokens.

3.2. Prompting

Our prompting strategy allows us to adapt and reuse the prompt by considering the task and the response length. The response time of the model depends on the overall length of the request, comprising the provided input and the length of the desired answer. The input is a string and consists of the task description and the selected abstracts. We utilized the refinement pattern introduced by White et al. [WFH*23] to design a structured prompt template and elaborate on the iterative refinement in the supplemental material. The template first enumerates and concatenates each abstract, helping the LLM reference specific abstracts in its answer (see 3.3).

We designed the task description to enable changing the specific task and the desired answer length. The LLM is then prompted with the task description and the abstracts. Figure 2 depicts the prompt template on the left with no specific `task` and `wordLimit`. Three example tasks with a fixed `wordLimit` of 300 words are shown on the right, where we instruct the model to summarize the abstracts, identify the most important insights, or analyze controversies. We counteract potential LLM hallucination by linking the results to the visualized data and the originally included abstracts supporting users in assessing the results' validity (see Section 3.3 below). Careful prompt development by including clear rules for the LLM to follow and instructions on how to proceed with negative cases reduces this risk further. This can be helpful for tasks such as *identifying controversies*, where the model can make up controversial statements otherwise. The use of natural language enables us to adapt the prompts for new tasks as well. We elaborate on the different tasks and the prompt refinement in the appendix.

Currently, the lens size and the abstracts' length are set in a way that supports EDA tasks adequately without exceeding the context token limit of *gpt-4o-mini*. This would not work for larger data sets or lens sizes. In such cases, sampling strategies or novel streaming-enabled models could be applied. Data privacy is also not considered for our usage scenario. Using privacy-preserving models or running models locally would be required for analyzing sensitive data.

3.3. Prototype

We realized our LLM-powered interactive lens in a web-based prototype using Angular, TypeScript, and the d3.js library. Figure 1 depicts the main view, which presents the IEEE VIS abstracts as 2D projections in a scatterplot. To help users track their interactive analysis, all data points inspected with the lens are shown in black, while the remaining points are depicted in gray. This realizes a mechanism similar to the "fog of war" idea for magic lenses [HTJ*20] but focuses only on hovered visual items and not on the lens position. To compare the analysis of two different

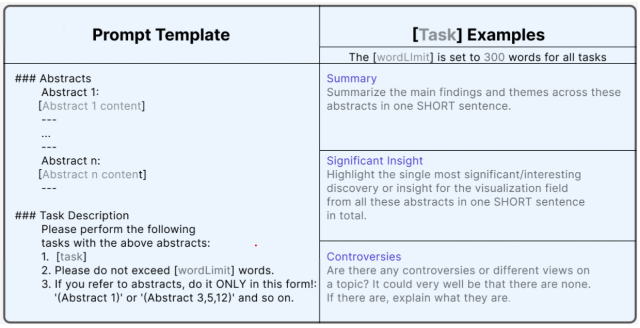

Figure 2: *Left: the template structures the prompt such that the context (abstracts) and the task description are separated. Right: three distinct task examples with the same fixed word limit.*

groups of data points, the previous lens position and LLM answers are still displayed after dropping the currently used lens (green) at a new position. The precomputed cluster labels are shown by default, but can be toggled off to focus on visual patterns of the scatterplot only. The tooltip seen in Figure 1(2) assigns a distinct color to each cluster, indicated by a small circle to the left of the label. This helps to associate clusters more easily, for example, in the author bar chart (see Figure 1(9)). Clustering helps with generating aggregated labels for guidance and getting an overview. We deliberately refrain from showing clusters explicitly by color, since abstracts might be mixtures of different topics, and coloring would indicate their association with a single specific cluster.

The LLM-generated summaries are shown to the right in a scrollable panel as seen in Figure 1(4). Key phrases extracted by the LLM and article references (indicated by the first author's name) are depicted as blue links (see Figure 1(7)). Clicking on a phrase produces a text box with further explanations generated by the LLM. If a user clicks on a reference, shown as *[first author] et al.*, the article's authors are listed, and the article's content is shown as a one-sentence, LLM-generated summary. The summary can be extended to display the actual, full abstract of the article, and a link below the abstract leads interested users to the publication for further reading. Figure 1(9) shows that clicking on an author in the author list depicts their main research areas by the number of published articles per area as a bar chart. We also show an LLM-generated summary of the author's main contributions and a list of frequent co-authors based on the author-involved abstracts.

Both important research phrases and author contributions are linked back to data points by marking the corresponding abstracts with a colored circle when users hover over the phrases or author bar chart respectively (see Figure 1(8)). To differentiate the links, we color backlinks of phrases in red and backlinks of authors in blue. The linking of abstract summaries to the abstracts and LLM facts to the data points allows users to keep track of the explored data and helps them assess the correctness of the LLM analysis. There are two ways of linking aspects of the analysis back to the visual representation: let the LLM perform the matching, or perform keyword-based matching. Again, we apply a hybrid approach by matching authors using a string comparison and letting the LLM define the backlinks for extracted important phrases and concepts using LLM output, plus string matching to increase robustness. The rationale of our approach is that author names can be string-matched straightforwardly, excluding typos, ambiguities, or matching problems related to hallucination. For the phrases, we want the LLM's results to allow for non-exact matching of synonymously described concepts and phrases. In this case, we only apply string-matching to the returned pseudo-structured links to abstracts to make the linking more robust if the structure of the identified abstracts is not perfectly well-formed (see template in Figure 2). The analysis tasks to be accomplished with the lens can be switched. By default, it creates a short summary of the hovered abstracts when dropped. Available tasks include stating the most significant insight, naming three salient facts, and identifying controversies, as seen with the prompts depicted in Figure 2. The appendix provides an application example of our prototype and details our suggested workflow for scientific literature exploration using the magic lens.

4. Discussion & Future Work

We believe that fast, interactive visual data exploration has benefits that cannot easily be compensated for with chat-based interfaces. Our suggested integration of LLMs into a magic lens-based technique is one step toward making fast interaction approaches more powerful. We speed up the analysis and circumvent high LLM processing delays by using a progressive solution. Applying LLM analyses more directly in a continuous manner might become possible as soon as powerful foundation models can be run locally on cost-efficient hardware.

We utilize openly available but not open-source LLMs in our approach since these models currently provide the best performance, at the risk of data leaks or unauthorized data usage when using private or sensitive data. As open-source models become more capable, one possible solution can be to train our own model. The LLM analyses of our approach are mainly shown as textual output. This seems adequate for analyzing textual documents and the details-on-demand approach we currently employ. However, it is not ideal for fast interaction since reading text takes time and makes interaction slower. Accordingly, fast, visually aggregated feedback, including textual labels, would be more desirable. The flexibility of our approach is currently limited by predefined tasks, but can be adapted with small changes to the prompt templates. Clearly, it would be great to let users define their own tasks via prompts to utilize them in a "prompt once, apply often" manner as well. Our approach supports fast and focused analysis of textual data in its current form, with the IEEE VIS abstracts as one example. With changes to the prompting and the linking mechanism, we can adapt the approach to a variety of textual datasets and applications, e.g., to progressively analyze large news corpora as demonstrated in the appendix. For future work, we intend to increase the generalizability of our approach by supporting the exploration of other, non-textual data and researching LLM-based visual feedback solutions using magic lenses. By linking analysis artifacts to the explored visualization, analysis errors might become understandable and detectable for users. However, we believe that such linking is even more important for the interactive analysis itself, revealing additional regions of interest.

References

[CGM*17] Ceneda D., Gschwandtner T., May T., Miksch S., Schulz H.-J., Streit M., Tominski C.: Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 111–120. doi:10.1109/TVCG.2016.2598468. 2

[CKB09] Cockburn A., Karlson A., Bederson B. B.: A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys 41*, 1 (2009), 1–31. doi:10.1145/1456650.1456652. 2

[CRF*21] Caillou P., Renault J., Fekete J.-D., Letournel A.-C., Sebag M.: Cartolabe: A web-based scalable visualization of large document collections. *IEEE Computer Graphics and Applications 41*, 2 (2021), 76–88. doi:10.1109/MCG.2020.3033401. 2

[FFS24] Fekete J.-D., Fisher D., Sedlmair M.: *Progressive Data Analysis*. Eurographics, 2024. doi:10.2312/pda.20242707. 2

[FHKM17] Federico P., Heimerl F., Koch S., Miksch S.: A survey on visual approaches for analyzing scientific literature and patents. *IEEE Transactions on Visualization and Computer Graphics 23*, 9 (2017), 2179–2198. doi:10.1109/TVCG.2016.2610422. 2

[FK14] Fried D., Kobourov S. G.: Maps of computer science. In *IEEE Pacific Visualization Symposium* (2014), IEEE, pp. 113–120. doi:10.1109/PacificVis.2014.47. 2

[Fur86] Furnas G. W.: Generalized fisheye views. *ACM SIGCHI Bulletin 17*, 4 (1986), 16–23. doi:10.1145/22339.22342. 2

[HJH*16] Heimerl F., John M., Han Q., Koch S., Ertl T.: DocuCompass: Effective exploration of document landscapes. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2016), pp. 11–20. doi:10.1109/VAST.2016.7883507. 2

[HTJ*20] Han Q., Thom D., John M., Koch S., Heimerl F., Ertl T.: Visual quality guidance for document exploration with focus+context techniques. *IEEE Transactions on Visualization and Computer Graphics 26*, 8 (2020), 2715–2731. doi:10.1109/TVCG.2019.2895073. 3

[HTK23] Hämäläinen P., Tavast M., Kunnari A.: Evaluating large language models in generating synthetic hci research data: a case study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023), CHI '23. doi:10.1145/3544548.3580688. 2

[IHK*17] Isenberg P., Heimerl F., Koch S., Isenberg T., Xu P., Stolper C. D., Sedlmair M., Chen J., Möller T., Stasko J.: Vispubdata.org: A metadata collection about IEEE Visualization (VIS) publications. *IEEE Transactions on Visualization and Computer Graphics 23*, 9 (2017), 2199–2206. doi:10.1109/TVCG.2016.2615308. 2

[KK15] Kucher K., Kerren A.: Text visualization techniques: Taxonomy, visual survey, and community insights. In *IEEE Pacific Visualization Symposium (PacificVis)* (2015), IEEE, pp. 117–121. doi:10.1109/PACIFICVIS.2015.7156366. 2

[KKP*17] Kim M., Kang K., Park D., Choo J., Elmqvist N.: TopicLens: Efficient multi-level visual topic exploration of large-scale document collections. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2017), 151–160. doi:10.1109/TVCG.2016.2598445. 2

[MHM20] McInnes L., Healy J., Melville J.: UMAP: Uniform manifold approximation and projection for dimension reduction, 2020. arXiv:1802.03426, doi:10.48550/arXiv.1802.03426. 2

[NJW01] Ng A., Jordan M., Weiss Y.: On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems* (2001), vol. 14, MIT Press. 2

[Nom22] Nomic AI: Nomic atlas. https://github.com/nomic-ai/nomic, 2022. 2

[Ope24] OpenAI: Gpt-4o mini: advancing cost-efficient intelligence. https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/, 2024. 3

[PNML08] Paulovich F. V., Nonato L. G., Minghim R., Levkowitz H.: Least square projection: A fast high-precision multidimensional projection technique and its application to document mapping. *IEEE Transactions on Visualization and Computer Graphics 14*, 3 (2008), 564–575. doi:10.1109/TVCG.2007.70443. 2

[RG19] Reimers N., Gurevych I.: Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing* (2019), Association for Computational Linguistics, pp. 3982–3992. 2

[RWVW23] Raval S., Wang C., Viégas F., Wattenberg M.: Explain-and-Test: An interactive machine learning framework for exploring text embeddings. In *IEEE Visualization and Visual Analytics (VIS)* (2023), pp. 216–220. doi:10.1109/VIS54172.2023.00052. 2

[SDBEA*23] Schetinger V., Di Bartolomeo S., El-Assady M., McNutt A., Miller M., Passos J. P. A., Adams J. L.: Doom or deliciousness: Challenges and opportunities for visualization in the age of generative models. *Computer Graphics Forum 42*, 3 (2023), 423–435. doi:10.1111/cgf.14841. 2

[Shn83] Shneiderman B.: Direct manipulation: A step beyond programming languages. *Computer 16*, 08 (1983), 57–69. doi:10.1109/MC.1983.1654471. 2

[SMSL17] Shao L., Mahajan A., Schreck T., Lehmann D. J.: Interactive regression lens for exploring scatter plots. *Computer Graphics Forum 36*, 3 (2017), 157–166. doi:10.1111/cgf.13176. 2

[TGK*17] Tominski C., Gladisch S., Kister U., Dachselt R., Schumann H.: Interactive lenses for visualization: An extended survey. *Computer Graphics Forum 36*, 6 (2017), 173–200. doi:10.1111/cgf.12871. 2

[VMZL22] Voigt H., Meuschke M., Zarriess S., Lawonn K.: KeywordScape: Visual document exploration using contextualized keyword embeddings. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (2022), Association for Computational Linguistics, pp. 137–147. doi:10.18653/v1/2022.emnlp-demos.14. 2

[War19] Ware C.: *Information Visualization: Perception for Design*. Morgan Kaufmann, 2019. doi:10.1016/C2016-0-02395-1. 2

[WFH*23] White J., Fu Q., Hays S., Sandborn M., Olea C., Gilbert H., Elnashar A., Spencer-Smith J., Schmidt D. C.: A prompt pattern catalog to enhance prompt engineering with chatgpt, 2023. arXiv:2302.11382, doi:10.48550/arXiv.2302.11382. 3

[WHC23] Wang Z. J., Hohman F., Chau D. H.: Wizmap: Scalable interactive visualization for exploring large machine learning embeddings. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)* (2023), Association for Computational Linguistics, pp. 516–523. doi:10.18653/v1/2023.acl-demo.50. 2

[WTP*95] Wise J. A., Thomas J. J., Pennock K., Lantrip D., Pottier M., Schur A., Crow V.: Visualizing the non-visual: Spatial analysis and interaction with information from text documents. In *Proceedings of Visualization 1995 Conference* (1995), IEEE, pp. 51–58. doi:10.1109/INFVIS.1995.528686. 2

[YCY*21] Yuan J., Chen C., Yang W., Liu M., Xia J., Liu S.: A survey of visual analytics techniques for machine learning. *Computational Visual Media 7* (2021), 3–36. doi:10.1007/s41095-020-0191-7. 2

[ZZZ*25] 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 (2025), 1830–1847. doi:10.1109/TVCG.2024.3368060. 1, 2