Scholarly Exploration via Conversations with Scholars-Papers Embedding

Ryan Yen* University of Waterloo
Yelizaveta Brus† University of Waterloo
Leyi Yan‡ University of Waterloo
Jimmy Lin§ University of Waterloo
Jian Zhao¶ University of Waterloo

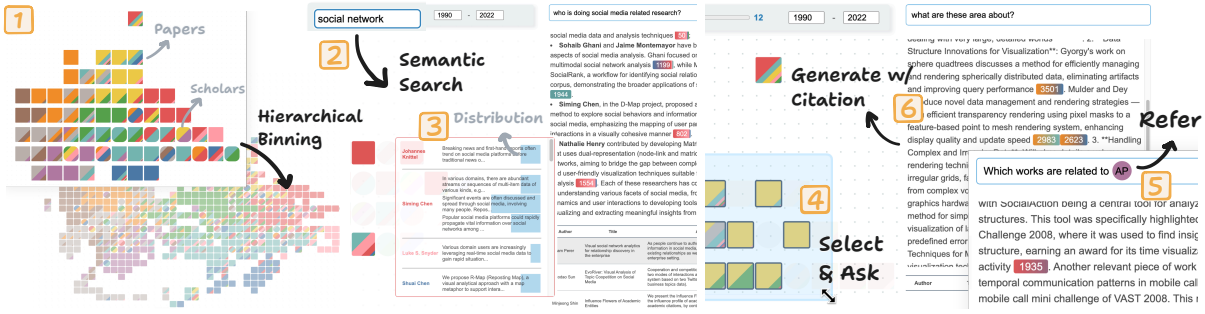

Figure 1: Key interactions with Scholet in an interactive system: (1) Scholars' works and academic papers are visualized in a 2D space, using hierarchical binning to reduce overlap; (2) Semantic searches can be performed to find relevant scholars and their works; (3) Users can explore the distribution of scholars and papers related to the search query; (4) Users can select specific Scholets of scholars or papers and ask detailed questions; (5) Users can refer to specific works or scholars in their queries; (6) The system generates responses grounded in the retrieved context, providing citations that link to relevant scholars and papers.

Abstract

We propose a system that supports contextually aware, controllable, and interactive exploration of academic publications and scholars. By enabling bidirectional interaction between question-answering components and Scholets, the 2D projections of scholarly works' embeddings, our system enables users to textually and visually interact with large amounts of publications. We report the system design and demonstrate its utility through an exploratory study with graduate researchers.

Index Terms: academic, retrieval augmented generation

1 Introduction

Exploring academic publications is an essential step toward understanding a specific domain or academic community. The rapid growth in publications from various sub-domains necessitates the development of visual analytics systems that enable researchers to efficiently navigate and retrieve relevant information. Previosu examples include 2D projections of documents [1], visualization of research trends [3], and bibliometric networks [5, 4, 2] that convey relationships between scholars.

These tools help researchers explore large volumes of publications with visual designs but their search features often fall short of providing dynamic and customized exploration of academic content. To fill in this gap, recent advancements in retrieval augmented generation (RAG) enable users to 'talk' to academic publications by asking complex, context-driven questions that go beyond simple searches. For example, users can inquire about the relationships between different groups of scholars, request summaries of multiple papers, or ask for overviews of specific research trends. Despite these advancements, current approaches often lack sufficient user control over the retrieval and generation process and do not integrate visualizations with the question-answering process, limiting their effectiveness in facilitating customized academic exploration.

We introduce an interactive system designed to assist users in exploring academic publications and scholars through contextually aware conversations integrated with visualizations (Figure 1). Our system facilitates bidirectional interaction between the question-answering components and 2D projections of scholarly embeddings, aggregated into bins that we term Scholets. The system allows users to select and refer to these Scholets along with their query. The generated results can also cite these Scholets from the visualizations, helping users explore the semantic relationships between cited works or scholars within the domain. This integration ensures that users can visually and contextually navigate the academic landscape while interacting with the data based on their own needs. In the following sections, we report on the system design and demonstrate the usefulness of our system with a user study.

2 Visualizing and Interacting with Scholets

Our system enables a conversational-based exploration of academic publications and bi-directional interaction with two main views.

The *Exploration View* visualizes the latent space of scholars' works on a 2D canvas, reflecting the semantic relationships based on their abstract (Figure 1-1). To reduce overlap and information load, we group the data into Scholet as bins and reorder these bins to snap into the grid. We then calculate the mean of each scholar's works to plot scholars within the view using rounded shapes. Users can use the pan and zoom to adjust the view. Imagine Celine, an explainable AI scholar, seeking collaboration with individuals in the visualization field with similar experiences. She can import data about visualization communities papers into the system and use the semantic search feature to find Scholet that are related to cross-disciplinary scholars in AI and visualization (Figure 1-2). Celine can easily locate scholars and their works within the highlighted area and hover over the Scholet to get more details of the works and the distribution of works related to the query (Figure 1-3). Additionally, Celine can adjust the year range to filter for the most recent works and modify the bin size to determine how many data points are included in a single Scholet. A larger bin size encompasses more data within a single Scholet, while a smaller bin size

*e-mail: r4yen@uwaterloo.ca
†e-mail: ybrus@uwaterloo.ca
‡e-mail: l52yan@uwaterloo.ca
§e-mail: jimmylin@uwaterloo.ca
¶e-mail: jianzhao@uwaterloo.ca

User Query Example	Decomposition
Authors & Publications: List publications on [topic A] and [topic B].	Filter papers related to [topic A] and [topic B].
Collaborations: Identify frequent co-authors with [scholar A] (2020-2024).	Filter papers with [scholar A] (2020-2024). Extract and count co-authors, excluding [scholar A].
Popularity: Show the evolution of research in [field/discipline] (last decade).	Filter and group papers by year and keywords in [field/discipline] (last decade).

Table 1: Sampled user queries and their decomposition.

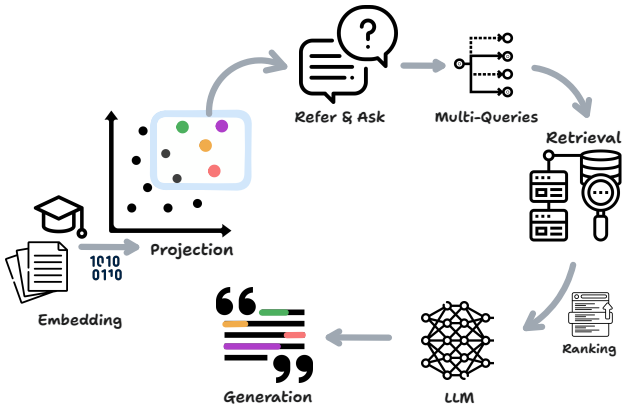

Figure 2: Our retrieval-augmented generation pipeline.

allows for exploration at a finer granularity.

The *Conversation View* assists users in 'talking' to publications or scholars through customized queries they input [6]. Responses generated in this view include clickable citations referring to corresponding Scholets from the *Exploration View*, offering contextual answers customized to users' queries (Figure 1-6). Users can also select or refer to Scholets in the input box to target conversations on specific subsets of scholars and publications (Figure 1-4/5). For instance, Celine can ask about *"the explainable AI research related to the* 1945*"*, receive detailed answers citing the 1945, and quickly navigate to the relevant Scholets to explore further.

System Architecture. We employ an adapted RAG process to integrate context retrieval with text generation, involving several steps (Figure 2): generating query prompts from the user's input, constructing database commands to filter results (e.g., transforming a query about related research into finding authors with topics related to "Adam"), retrieving similar contexts using vector similarity search, and ranking documents with Reciprocal Rank Fusion. The top-k results serve as a context for the generation model, grounding responses in the relevant information. The system, implemented in TypeScript with Svelte, uses Python for backend processing, MongoDB for data storage, and GPT-4.5 for response generation. The system employs binning and lazy loading to support a smooth interaction with a large dataset on the web. User prompts are decomposed into keywords, mapped to database columns, and ordered by semantic relevance to retrieve contextual information, as exemplified in database queries like identifying scholars with similar research interests (Table 1). For instance, the user query, '*Who have worked with similar research interest as Professor A?*" would be processed as follows: first, filter papers authored by Professor A; second, generate keywords representing Professor A's primary research areas; third, filter papers based on these keywords; and finally, retrieve the authors of the relevant papers.

3 Exploratory User Study

We demonstrated the usefulness of our system through a exploratory study involving four graduate students (P1-4) who used our system to explore scholars and papers within their institute. They were asked to use the system freely to explore scholars and their work within the institution for 15 minutes and to conduct a 15-minute semi-formal interview. The dataset was compiled from the 50 most recent publications of each faculty member in an AI institution, resulting in a collection of 190 unique authors across eight departments and six faculties. Abstracts were obtained via the OpenAlex API, Semantic Scholar API, and web scraping.

Participants approached the dataset with diverse objectives beyond merely understanding the data. For example, P2 used the system to identify related research pertinent to his own work, while P3 aimed to gain a holistic view of the research she intended to investigate in the future. Overall, participants appreciated the generated results' accuracy, particularly in providing factual citations that could be referenced back to the visualizations. P4 noted that the "Select and Ask" feature facilitated more in-depth discussions on specific sub-domains by generating relevant answers with more control. All participants reported that they felt more control in the generation process when they saw that the citations were consistent with their choices. While P1 enjoyed exploring different domains using the system, he mentioned that viewing details through tooltips one by one was inefficient. He expressed a preference for having more detailed information directly available within the Scholet, suggesting enhancements for future iterations of the system.

4 Discussion and Future Work

While our system assists users in dynamic and context-aware exploration, other existing visualizations focused on exploring scholar relationships, citation networks, or research trends might be more suitable for tailored scenarios. Our system also has several limitations that present opportunities for future work. One aspect is that while we provide data scraping tools, it is still necessary to manually enter a data entry describing the scholarly information wanted to be included. Additionally, enhancing support for more flexible data importing, such as customized column configurations, would make the system more adaptable to various datasets. To facilitate community engagement and further development, the source code can be accessed at `https://github.com/ryanyen2/Scholet` and we will update it with the above-mentioned improvements.

References

[1] M. Berger, K. McDonough, and L. M. Seversky. cite2vec: Citation-driven document exploration via word embeddings. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):691–700, 2017. doi: 10.1109/TVCG.2016.2598667 1

[2] N. Elmqvist and P. Tsigas. Citewiz: a tool for the visualization of scientific citation networks. *Information Visualization*, 6(3):215–232, dec 2007. doi: 10.1145/1375939.1375943 1

[3] F. Heimerl, Q. Han, S. Koch, and T. Ertl. Citerivers: Visual analytics of citation patterns. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):190–199, 2016. doi: 10.1109/TVCG.2015.2467621 1

[4] L. Shi, H. Tong, J. Tang, and C. Lin. Vegas: Visual influence graph summarization on citation networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(12):3417–3431, 2015. doi: 10.1109/TKDE.2015.2453957 1

[5] N. J. Van Eck and L. Waltman. Citnetexplorer: A new software tool for analyzing and visualizing citation networks. *Journal of informetrics*, 8(4):802–823, 2014. 1

[6] R. Yen and J. Zhao. Memolet: Reifying the reuse of user-ai conversational memories. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, October 2024. doi: 10.1145/3654777.3676388 2