

WILDVIS: Open Source Visualizer for Million-Scale Chat Logs in the Wild

Yuntian Deng^{1*}, Wenting Zhao², Jack Hessel³,
Xiang Ren⁴, Claire Cardie², Yejin Choi^{5,6*}

¹University of Waterloo ²Cornell University ³Samaya AI
⁴University of Southern California ⁵University of Washington ⁶Nvidia

yuntian@uwaterloo.ca, wzhao@cs.cornell.edu, jmhessel@gmail.com
xiangren@usc.edu, cardie@cs.cornell.edu, yejin@cs.washington.edu

Abstract

The increasing availability of real-world conversation data offers exciting opportunities for researchers to study user-chatbot interactions. However, the sheer volume of this data makes manually examining individual conversations impractical. To overcome this challenge, we introduce WILDVIS, an interactive tool that enables fast, versatile, and large-scale conversation analysis. WILDVIS provides search and visualization capabilities in the text and embedding spaces based on a list of criteria. To manage million-scale datasets, we implemented optimizations including search index construction, embedding precomputation and compression, and caching to ensure responsive user interactions within seconds. We demonstrate WILDVIS’ utility through three case studies: facilitating chatbot misuse research, visualizing and comparing topic distributions across datasets, and characterizing user-specific conversation patterns. WILDVIS is open-source and designed to be extendable, supporting additional datasets and customized search and visualization functionalities.

1 Introduction

While hundreds of millions of users interact with chatbots like ChatGPT (Malik, 2023), the conversation logs remain largely opaque for open research, limiting our understanding of user behavior and system performance. Recently, initiatives such as WildChat (Zhao et al., 2024) and LMSYS-Chat-1M (Zheng et al., 2024) have released millions of real-world user-chatbot interactions, offering rich opportunities to study interaction dynamics. However, the volume and complexity of these datasets pose significant challenges for effective analysis.

To help researchers uncover patterns and anomalies within these vast chat datasets, we introduce

*Work done in large part while at the Allen Institute for Artificial Intelligence.

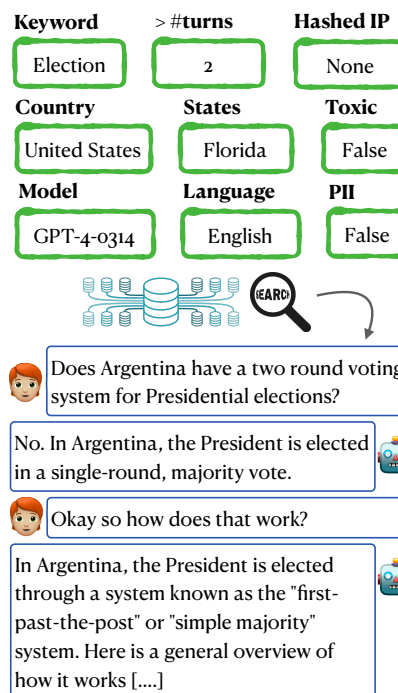


Figure 1: Illustration of an exact, compositional filter-based search in WILDVIS. This example demonstrates the application of multiple criteria, including the keyword “Election,” conversations with more than two turns, and chats from users in Florida.

WILDVIS, an interactive tool for exploring million-scale chat logs. WILDVIS enables researchers to find conversations based on specific criteria, understand topic distributions, and explore semantically similar conversations, all while maintaining efficiency. Figure 1 illustrates an example search using WILDVIS, applying criteria such as the keyword “Election,” conversations with more than two turns, and chats from users in Florida, among others.

WILDVIS features two main components: an exact, compositional filter-based retrieval system, which allows users to refine their search using ten predefined filters such as keywords, geographical location, IP address, and more. The second component is an embedding-based visualization module,

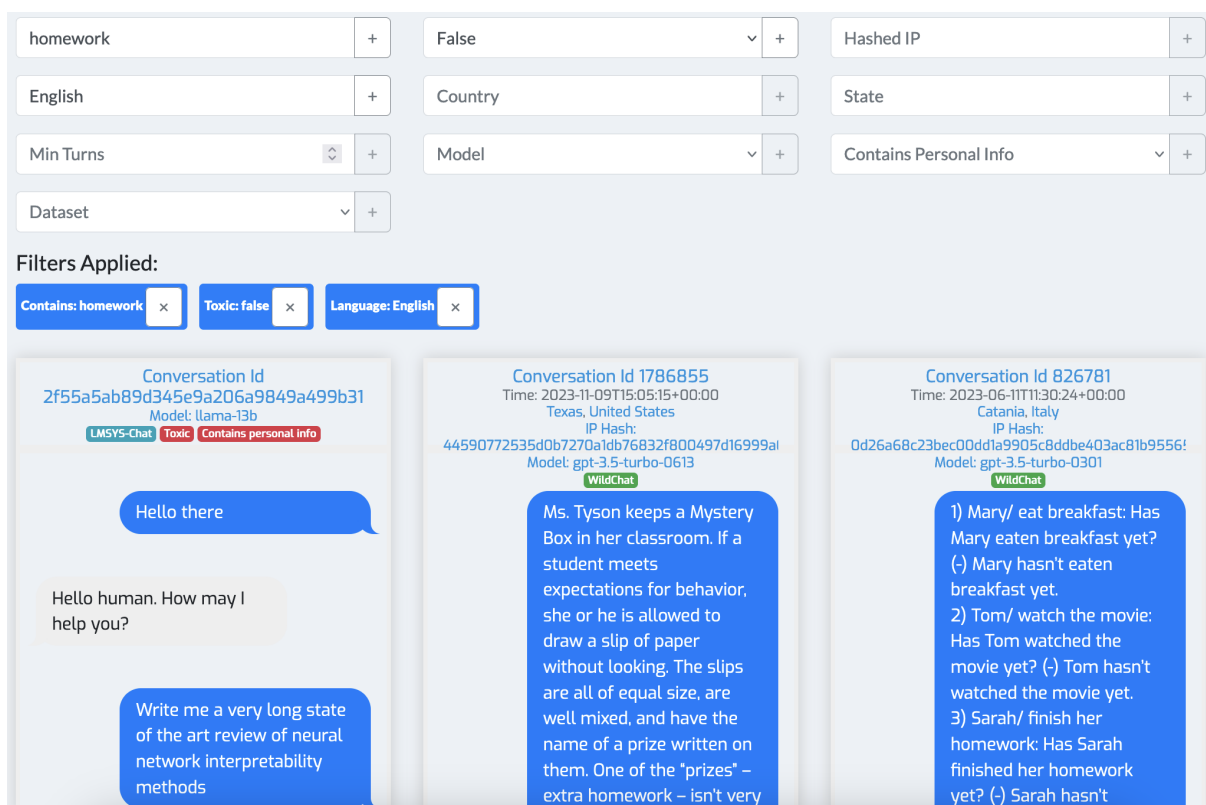


Figure 2: WILDVIS Filter-Based Search Page.¹ This screenshot shows the application of multiple filters, including conversation content (“homework”), non-toxicity, and language (English), to narrow down the search results. The interface displays relevant conversations that match the specified criteria. Users can click on each conversation ID to navigate to the conversation details page. Additionally, metadata in the displayed results, such as the hashed IP address, is clickable, allowing users to filter based on that specific metadata.

which represents conversations as dots on a 2D plane, with similar conversations positioned closer together. Both components are designed to scale to millions of conversations. A preliminary version of the tool, which supported filter-based retrieval for one million WildChat conversations, was accessed over 18,000 times by 962 unique IPs in July and August 2024 alone. The latest release, described in this paper, extends support to both components for WildChat and LMSYS-Chat-1M.

In this paper, we present the design and implementation of WILDVIS, discussing the strategies employed to scale to million-scale datasets while maintaining latency within seconds. We also showcase several use cases: facilitating chatbot misuse research (Brigham et al., 2024; Mireshghalah et al., 2024), visualizing and comparing topic distributions between WildChat and LMSYS-Chat-1M, and characterizing user-specific conversation patterns. For example, WILDVIS reveals distinct topic clusters such as Midjourney prompt generation in WildChat and chemistry-related conversations in LMSYS-Chat-1M. Additionally, we ob-

serve that WildChat exhibits a generally more creative writing style compared to LMSYS-Chat-1M. As an open-source project, WILDVIS is available at github.com/da03/WildVisualizer under an MIT license, and a working demo can be accessed at wildvisualizer.com.

2 User Interface

WILDVIS consists of two primary pages—a filter-based search page and an embedding visualization page—along with a conversation details page. These pages are designed to provide users with both high-level overviews and detailed insights into individual conversations.

2.1 Filter-Based Search Page

The filter-based search page (Figure 2) enables users to filter the dataset based on a list of criteria. Users can input keywords to retrieve relevant conversations or narrow down results using specific

¹This example is available at <https://wildvisualizer.com/?contains=homework&toxic=false&language=English>.

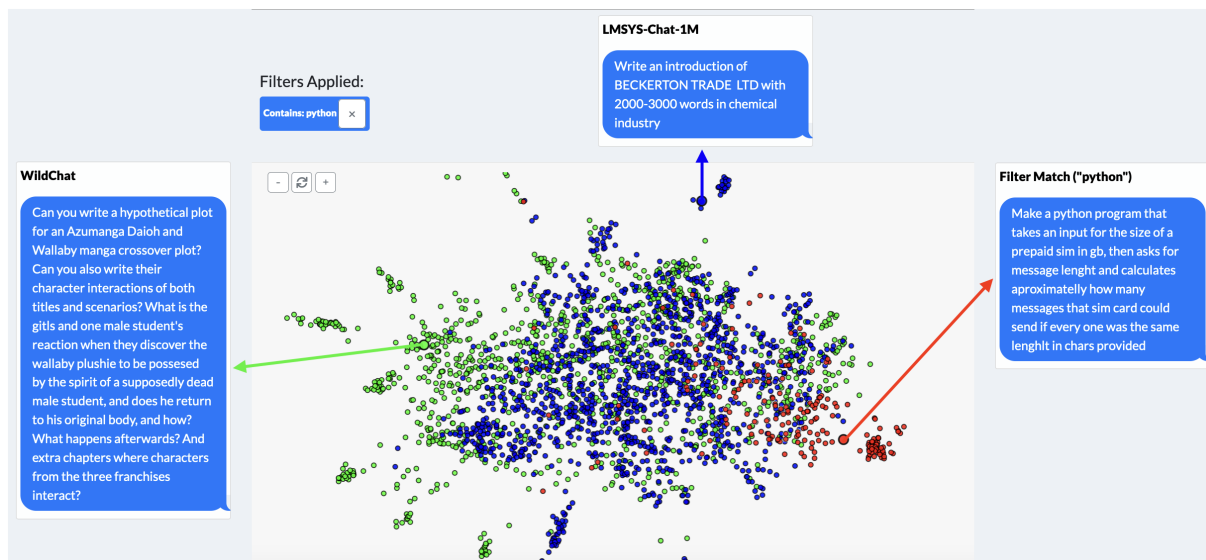


Figure 3: WILDVIS Embedding Visualization page.² Each dot represents a conversation, with green dots from WildChat, blue dots from LMSYS-Chat-1M, and red dots highlighting conversations that match the applied filters (containing “python” in this example). Users can interact with the visualization by hovering over dots to preview a conversation and clicking on a dot to navigate to the full conversation. This figure has been enhanced to show a representative example from each category: “WildChat,” “LMSYS-Chat-1M,” and “Filter Match.”

criteria. In total, ten predefined filters are available, including:

- Hashed IP Address: Filter conversations by hashed IP addresses to analyze interactions from the same user.³
- Geographical Data: Filter by inferred state and country to gain insights into regional variations in conversational patterns.
- Language: Restrict results to conversations in specific languages.
- Toxicity: Include or exclude conversations flagged as toxic.
- Redaction Status: Include or exclude conversations with redacted personally identifiable information (PII).
- Minimum Number of Turns: Focus on conversations with a specified minimum number of turns.
- Model Type: Select conversations by the underlying language model used, such as GPT-3.5 or GPT-4.

²This example is available at <https://wildvisualizer.com/embeddings/english?contains=python>.

³IP addresses are hashed to protect user privacy while still allowing the analysis of interactions associated with the same user.

The search results are displayed in a paginated table format, ensuring easy navigation through large datasets. Active filters are prominently displayed above the results and can be removed by clicking the “×” icon next to each filter.

Each result entry displays key metadata, including the conversation ID, timestamp, geographic location, hashed IP address, and model type. Users can interact with these results in multiple ways. Clicking on a conversation ID leads to a detailed view of that conversation. Additionally, all metadata fields, such as the hashed IP address, are clickable, enabling users to quickly search based on specific attributes. For example, clicking on a hashed IP address brings up a list of all conversations associated with that IP, facilitating user-specific analyses.

2.2 Embedding Visualization Page

In addition to traditional search capabilities, WILDVIS offers an embedding visualization page (Figure 3), which allows users to explore conversations based on their semantic similarity. Conversations are represented as dots on a 2D plane, with similar conversations placed closer together.

Basic Visualization Each conversation appears as a dot, with different datasets distinguished by color. Hovering over a dot reveals a preview of the conversation, and clicking on it navigates to

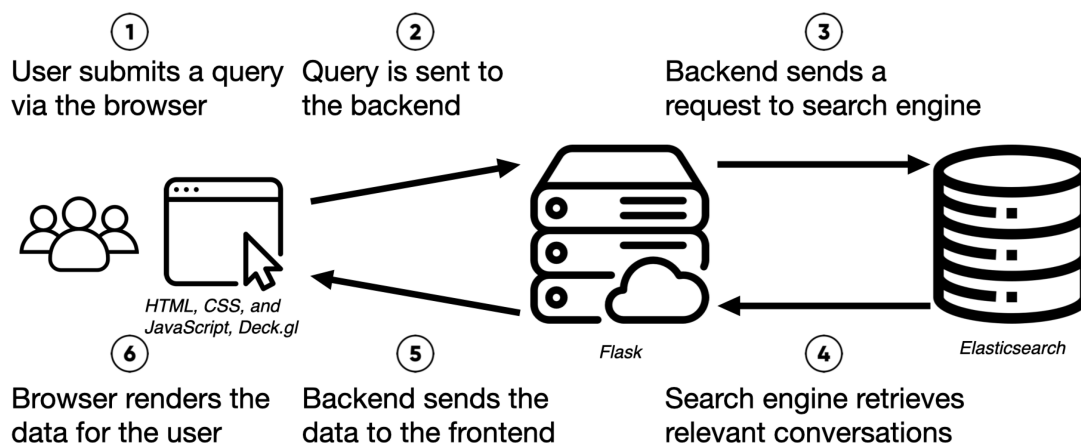


Figure 4: System Architecture: Overview of the data flow from user query submission to result rendering in the browser. The software tools used in the frontend, backend, and search engine are italicized.

the conversation details page.⁴ Users can zoom in, zoom out, and drag the view to explore different regions of the visualization. This spatial arrangement enables users to explore clusters of related conversations and identify structures within the data.

Filter-Based Highlighting Similar to the filter-based search page, users can apply filters to highlight specific conversations on the 2D map, with matching conversations marked in red. This feature helps users locate conversations of interest, such as identifying topics associated with a particular user.

Conversation Embedding To represent each conversation as a point in 2D space, we embed the first user turn of each conversation using OpenAI’s text-embedding-3-small model.⁵ We then trained a parametric UMAP model (Sainburg et al., 2021; McInnes et al., 2020) to project these embeddings into 2D space.⁶ Since initial experiments showed that training a single UMAP model on all embeddings resulted in some clusters driven by language differences (see Figure 7 in Appendix B), in order to create more semantically meaningful clusters, we also trained a separate parametric UMAP model for each language. Users can easily switch between different languages and their corresponding UMAP projections.

The combination of embedding visualization,

⁴On mobile devices, tapping a dot displays a preview with options to view the full conversation or close the preview. See Figure 6 in Appendix A for a screenshot.

⁵We opted to embed only the first user turn, as preliminary experiments showed that embedding the entire conversation led to less intuitive clustering.

⁶We chose parametric UMAP over t-SNE (van der Maaten and Hinton, 2008) to enable online dimensionality reduction, which will be discussed in Section 3.2.

filtering, highlighting, and interactive previews enables users to navigate vast amounts of conversation data, uncovering insights and connections that might otherwise remain hidden. For example, users can easily identify outliers and clusters.

2.3 Conversation Details Page

The conversation details page (Figure 9 in Appendix C) provides a detailed view of individual conversations. This page displays all the turns between the user and the chatbot, along with associated metadata. Similar to the filter-based search page, all metadata fields are clickable, allowing users to apply filters based on their values. However, if users arrive at this page by clicking a dot on the embedding visualization page, the filtering will be applied within the embedding visualization context. A toggle switch on the conversation details page allows users to control which page (filter-based search or embedding visualization) clicking on metadata fields will direct them to.

3 System Implementation

WILDVIS is designed to efficiently process large-scale conversational datasets.

3.1 System Architecture

WILDVIS operates on a client-server architecture, where the server handles data processing, search, and conversation embedding, while the client provides an interface for data exploration. The high-level system architecture is illustrated in Figure 4.

Users interact with the frontend web interface, which communicates their queries to the backend server. The backend server is built using

Flask⁷, which processes these queries and constructs search requests for an Elasticsearch⁸ engine. Elasticsearch, known for its scalable search capabilities, retrieves the relevant conversations, which are then sent back to the frontend for rendering. The frontend is developed using HTML, CSS, and JavaScript⁹, with Deck.gl¹⁰ used for rendering large-scale, interactive embedding visualizations.

3.2 Scalability and Optimization

To manage the large volume of data and ensure smooth user interaction, WILDVIS uses several optimization strategies.

Search For search functionalities, an index is built for each dataset with all metadata using Elasticsearch, allowing the backend to efficiently retrieve relevant conversations. To reduce the load during queries with a large number of matches, we employ two strategies: pagination, which retrieves results one page at a time with up to 30 conversations per page, and limiting the number of retrieved matches to 10,000 conversations per search.

Embedding Visualization - Frontend Rendering a large number of conversation embeddings is computationally intensive for a browser, especially on mobile devices, and may lead to visual clutter with overlapping dots. To mitigate these issues, we use Deck.gl to render large numbers of points efficiently. Additionally, we restrict the visualization to a subset of 1,500 conversations per dataset, ensuring smooth rendering and clear visualization.

Embedding Visualization - Backend On the backend, computing embeddings for a large number of conversations can introduce significant delays. To address this, we precompute the 2D coordinates for the subset of conversations selected for visualization. These precomputed results are then compressed using gzip and stored in a file, which is sent to the user upon their first visit to the embedding visualization page. The compressed file is approximately 1 MB in size and only needs to be downloaded once.

Although we only display a subset of conversations, users may still need to search the entire dataset. To support this, we integrate the embedding visualization with the Elasticsearch engine.

⁷<https://flask.palletsprojects.com/>

⁸<https://www.elastic.co/elasticsearch>

⁹The frontend is built on top of MiniConf (Rush and Strobel, 2020).

¹⁰<https://deck.gl/>

When a user submits a query, we first search within the displayed subset of conversations (with an index built for this subset). If sufficient matches are found within the subset (with a default threshold of 100, adjustable up to 1,000), we simply highlight them and do not extend the search further. However, if there are not enough matches, we extend the search to the entire dataset using Elasticsearch, retrieve the relevant conversations (up to the threshold number), and embed and project them into 2D coordinates before sending them to the frontend for visualization. To speed up this process, we cache all computed coordinates in an SQLite database. Due to the need to dynamically compute coordinates for conversations not found in the cache, we chose parametric UMAP over t-SNE, as t-SNE does not learn a projection function, whereas parametric UMAP allows for quick projection of new conversations into lower-dimensional space.

3.3 Performance Evaluation

To evaluate the efficiency of our system, we generated ten random keyword-based search queries and measured the execution time for each using our tool. On the filter-based search page, each query took an average of 0.47 seconds ($\pm 0.06s$). In comparison, a naive for-loop-based approach using the Hugging Face Datasets library took 1148.89 seconds ($\pm 25.28s$). For embedding visualization, the same measurement method was used, and each query took an average of 0.43 seconds ($\pm 0.01s$).

4 Use Cases

This section presents several use cases that demonstrate the potential of WILDVIS. It is important to note that WILDVIS is designed primarily for exploratory data analysis rather than for final quantitative analysis.

Data WILDVIS currently supports two datasets: WildChat (Zhao et al., 2024) and LMSYS-Chat-1M (Zheng et al., 2024). These datasets are integrated into the system by building Elasticsearch indices and precomputing the 2D coordinates of a randomly selected subset of conversations for embedding visualization.

4.1 Facilitating Chatbot Misuse Research

One application of WILDVIS is in facilitating studies on chatbot misuse. We show here that WILDVIS is able to both reproduce existing studies on chatbot misuse and to discover new misuse cases.

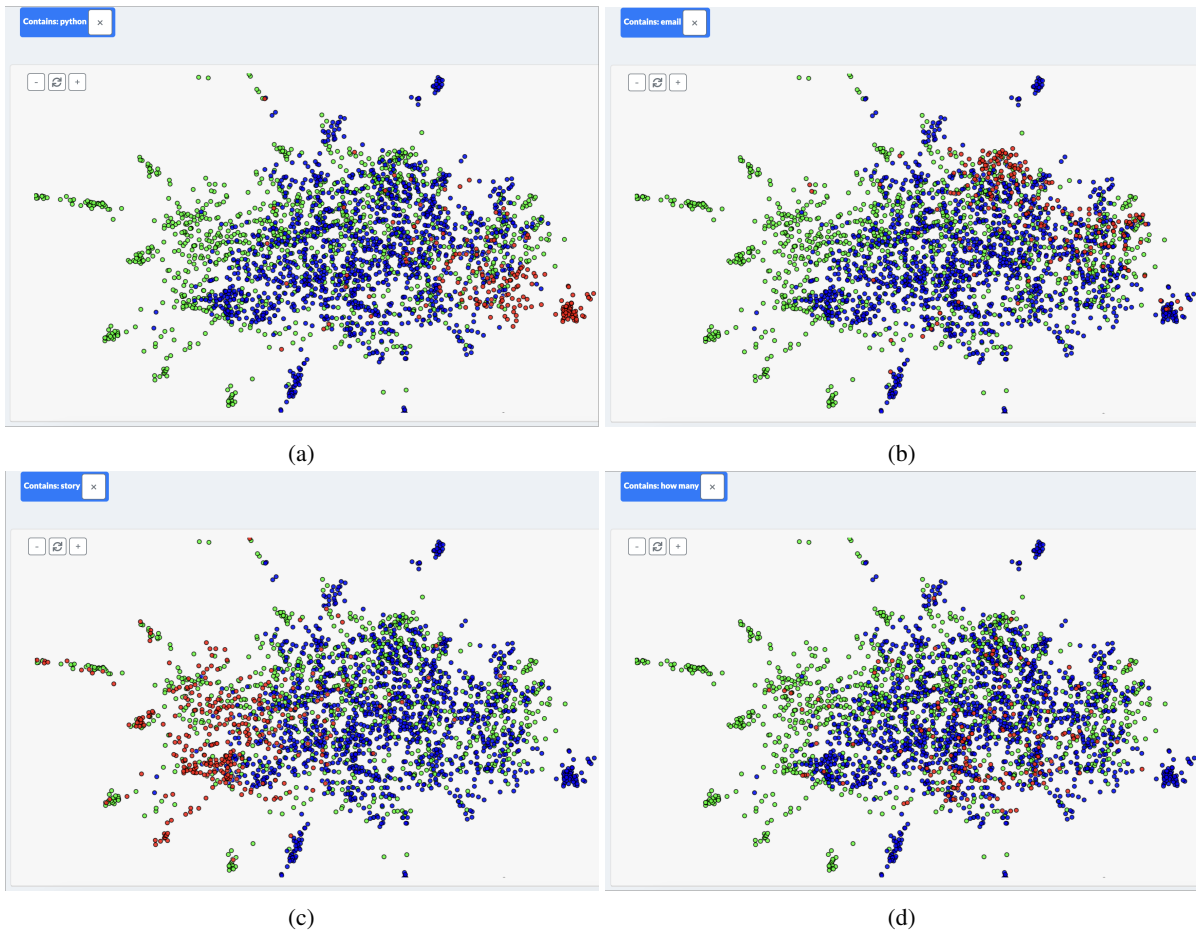


Figure 5: Major topic clusters.¹¹ (a) Coding (identified by searching for “python”). (b) Writing assistance (identified by searching for “email”). (c) Story generation (identified by searching for “story”). (d) Math question answering (identified by searching for “how many”).

Reproducing a Study on Journalist Misuse In this use case, we replicate the findings of [Brigham et al. \(2024\)](#), which identified instances of journalists misusing the chatbot behind WildChat to paraphrase existing articles for their work. To locate a specific instance mentioned in the study, we use the following quote from the original research:

write a new article out of the information in this article, do not make it obvious you are taking information from them but in very sensitive information give them credit.

To find this conversation, we enter the phrase *you are taking information from them* in the “Contains”

¹¹These examples can be found at <https://wildvisualizer.com/embeddings/english?contains=python>, <https://wildvisualizer.com/embeddings/english?contains=email>, <https://wildvisualizer.com/embeddings/english?contains=story>, and <https://wildvisualizer.com/embeddings/english?contains=how%20many>.

field on the search page and execute the search.¹² The search returns a single result, matching the case mentioned in the original paper. By clicking on the hashed IP address, we can view all conversations from this user, identifying all 15 conversations analyzed in the original study ([Brigham et al., 2024](#)).

Reproducing a Study on User Self-Disclosure In another example, we replicate findings from a study on user self-disclosure behaviors by [Mireshghallah et al. \(2024\)](#). We search for a key phrase from that paper: *I have invited my father*.¹³ Again, the search returns a single result, allowing us to find the conversation discussed in the study.

Discovering Additional Misuse Cases WILD-VIS also facilitates the discovery of additional misuse cases. For instance, by searching for conver-

¹²This case can be found at <https://wildvisualizer.com/?contains=you%20are%20taking%20information%20from%20them>.

¹³This case can be found at <https://wildvisualizer.com/?contains=I%20have%20invited%20my%20father>.

sations that contain both personally identifiable information (PII) and the term “Visa Officer”¹⁴, we identified multiple entries from the same IP address. Further filtering based on this IP address revealed that the user appears to be affiliated with an immigration service firm and has disclosed sensitive client information.¹⁵

4.2 Visualizing and Comparing Topics

A powerful feature of the embedding visualization page in WILDVIS is its ability to visualize the overall distribution of topics, with conversations of similar topics positioned close to each other. In our previous discussion on embedding conversations, we illustrated language-specific clusters (Figure 7 in Appendix B). As another example, for English data, this visualization reveals that the embedding space can be roughly divided into four regions: coding (by searching for “python”), writing assistance (by searching for “email”), story generation (by searching for “story”), and math question answering (by searching for “how many”), as illustrated in Figure 5. This observation aligns with the findings in Merrill and Lerman (2024).

This feature also allows for the comparison of topic distributions across different datasets. By inspecting regions with different colors, users can identify outliers, regions where one dataset is well-represented while the other is not, and areas where both datasets overlap. By hovering over these regions, patterns in the types of conversations can be observed. For example, we found that WildChat contains more conversations related to creating writing and an outlier cluster of Midjourney prompt generation (see Figure 8a) compared to LMSYS-Chat-1M, while LMSYS-Chat-1M has outlier clusters of conversations about chemistry (see Figure 8b).

4.3 Characterizing User-Specific Patterns

WILDVIS can also be used to visualize the topics of all conversations associated with a specific user on the embedding map. For example, Figure 10 displays all conversations of a single user, revealing two main topic clusters: coding-related and email writing-related.

¹⁴<https://wildvisualizer.com/?contains=Visa%20Officer&redacted=true>

¹⁵https://wildvisualizer.com/?hashed_ip=048b169ad0d18f2436572717f649bdeedac793967fb63ca6632a2f5dca14e4b8

5 Related Work

Hugging Face Dataset Viewer Hugging Face’s Dataset Viewer (Lhoest et al., 2021) provides basic search functionalities for datasets hosted on Hugging Face. However, it is designed for general dataset visualization and is not specifically tailored for conversational datasets. For example, while it offers useful statistics, navigating JSON-formatted conversations in a table format can be cumbersome and lacks the intuitive visualization needed for exploring conversational data.

Paper Visualization Tools The ACM Fellows’ Citation Visualization tool¹⁶ embeds ACM Fellows based on their contribution statements. While its interface shares many similarities with the embedding visualization page of WILDVIS, it focuses on publication data rather than conversational data. Another relevant work is Yen et al. (2024), which visualizes papers in a similar manner, with an added conversational component that allows users to interact with the visualizations by asking questions. However, it is also primarily designed for academic papers rather than large-scale chat datasets.

6 Conclusion

In this paper, we introduced WILDVIS, an interactive web-based tool designed for exploring large-scale conversational datasets. By combining powerful search functionalities with intuitive visualization capabilities, WILDVIS enables researchers to uncover patterns and gain insights from vast collections of user-chatbot interactions. The system’s scalability optimizations ensure efficient handling of million-scale datasets, while maintaining a responsive and user-friendly experience.

WILDVIS fills a gap in existing tools by providing a specialized platform for visualizing and exploring chat datasets, which are inherently challenging to analyze using generic dataset viewers. Our use cases demonstrate the tool’s potential to replicate and extend existing research on chatbot misuse and user self-disclosure, as well as to facilitate topic-based conversation exploration.

Acknowledgments

This work is supported by ONR grant N00014-24-1-2207, NSF grant DMS-2134012, and an NSERC Discovery grant. We also thank Bing Yan, Pengyu Nie, and Jiawei Zhou for their valuable feedback.

¹⁶<https://mojtabaa4.github.io/acm-citations/>

References

- Natalie Grace Brigham, Chongjiu Gao, Tadayoshi Kohno, Franziska Roesner, and Niloofar Miresghallah. 2024. [Breaking news: Case studies of generative ai’s use in journalism](#). *Preprint*, arXiv:2406.13706.
- Jonathan P. Chang, Caleb Chiam, Liye Fu, Andrew Wang, Justine Zhang, and Cristian Danescu-Niculescu-Mizil. 2020. [ConvoKit: A toolkit for the analysis of conversations](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 57–60, 1st virtual meeting. Association for Computational Linguistics.
- Zixin Chen, Jiachen Wang, Meng Xia, Kento Shigyo, Dingdong Liu, Rong Zhang, and Huamin Qu. 2024. [Stugptviz: A visual analytics approach to understand student-chatgpt interactions](#). *Preprint*, arXiv:2407.12423.
- Shachar Don-Yehiya, Leshem Choshen, and Omri Abend. 2024. [The sharelm collection and plugin: Contributing human-model chats for the benefit of the community](#). *Preprint*, arXiv:2408.08291.
- Yanai Elazar, Akshita Bhagia, Ian Helgi Magnusson, Abhilasha Ravichander, Dustin Schwenk, Alane Suhr, Evan Pete Walsh, Dirk Groeneveld, Luca Soldaini, Sameer Singh, Hannaneh Hajishirzi, Noah A. Smith, and Jesse Dodge. 2024. [What’s in my big data?](#) In *The Twelfth International Conference on Learning Representations*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aisha Malik. 2023. [OpenAI’s ChatGPT now has 100 million weekly active users](#). Accessed: 2024-08-04.
- Leland McInnes, John Healy, and James Melville. 2020. [Umap: Uniform manifold approximation and projection for dimension reduction](#). *Preprint*, arXiv:1802.03426.
- Jeremy B. Merrill and Rachel Lerman. 2024. [What do people really ask chatbots? it’s a lot of sex and homework](#). *The Washington Post*. Accessed: 2024-08-27.
- Niloofar Miresghallah, Maria Antoniak, Yash More, Yejin Choi, and Golnoosh Farnadi. 2024. [Trust no bot: Discovering personal disclosures in human-llm conversations in the wild](#). *Preprint*, arXiv:2407.11438.
- Alexander M. Rush and Hendrik Strobelt. 2020. [Mini-conf – a virtual conference framework](#). *Preprint*, arXiv:2007.12238.
- Tim Sainburg, Leland McInnes, and Timothy Q Genter. 2021. [Parametric umap embeddings for representation and semi-supervised learning](#). *Preprint*, arXiv:2009.12981.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Ryan Yen, Yelizaveta Brus, Leyi Yan, Jimmy Lin, and Jian Zhao. 2024. [Scholarly exploration via conversations with scholars-papers embedding](#).
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. [Wildchat: 1m chatGPT interaction logs in the wild](#). In *The Twelfth International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. 2024. [LMSYS-chat-1m: A large-scale real-world LLM conversation dataset](#). In *The Twelfth International Conference on Learning Representations*.

A Embedding Visualization on Mobile Devices

Figure 6 shows the embedding visualization page on mobile devices. Since mobile devices do not support hover interactions, we adapted the interface by using a tap gesture for displaying previews.

B Language-Specific Clusters

When visualizing all conversations together on the embedding visualization page, clusters based on language emerge, such as the Spanish, Arabic, Chinese, and Russian clusters in Figure 7.

C Conversation Details Page

Figure 9 shows a screenshot of the conversation details page, where all metadata fields are displayed alongside the dialogue content. Clicking any metadata field will filter the conversations based on the selected value. Depending on how the user navigated to this page—either from the filter-based search page or the embedding visualization page—the filtering action will redirect the user back to the respective page. A toggle switch at the top allows users to control this behavior.

D Visualizing and Comparing Topic Distributions

The embedding visualization highlights distinct outlier clusters in the dataset. One notable cluster in the WildChat dataset involves Midjourney prompt engineering, where users ask the chatbot to generate detailed prompts for use with Midjourney, as shown in Figure 8a (this was also noted by Merrill and Lerman (2024)). Another distinct outlier cluster comprises chemistry-related questions in LMSYS-Chat-1M, illustrated in Figure 8b.¹⁷

E Characterizing User-Specific Patterns

WILDVIS can be used to visualize the topics of all conversations associated with a specific user on the embedding map. For example, Figure 10 displays all conversations from a single user, revealing two main topic clusters: coding-related and email writing-related.

F Additional Related Work

Chat Visualization Tools Several browser-based tools exist for chat visualization, such as

¹⁷Yao Fu discovered this and shared it with the authors.

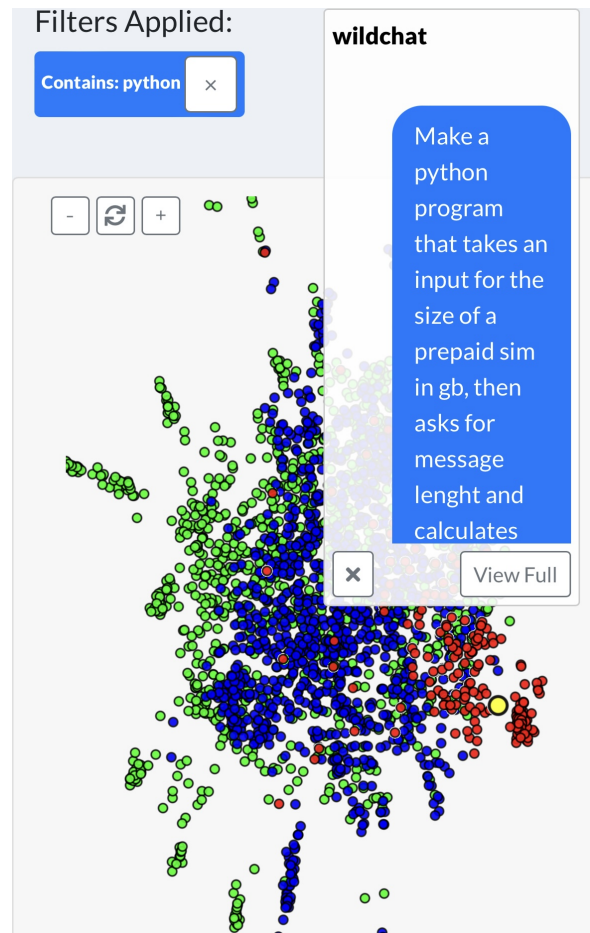


Figure 6: Embedding visualization on mobile devices. Tapping a dot displays a preview with options to view the full conversation or close it. This example can be viewed at <https://wildvisualizer.com/embeddings/english?contains=python> on a mobile device.

ShareGPT¹⁸, which allows users to share their conversations. Similarly, browser extensions like ShareLM (Don-Yehiya et al., 2024) enable users to upload and view their conversations, and ChatGPT History Search¹⁹ offers search functionality for a user’s personal conversations. In addition, Chen et al. (2024) developed StuGPTViz, a visual analytics system for analyzing student-ChatGPT interactions in educational settings.

Large-scale Data Analysis Tools Specialized tools like ConvoKit (Chang et al., 2020) provide a framework for analyzing dialogue data. Another notable tool, WIMBD (Elazar et al., 2024), supports the analysis and comparison of large text corpora, offering functionalities such as searching for documents containing specific queries and counting statistics like n-gram occurrences.

¹⁸<https://sharegpt.com>

¹⁹<https://chatgpthistorysearch.com/en>

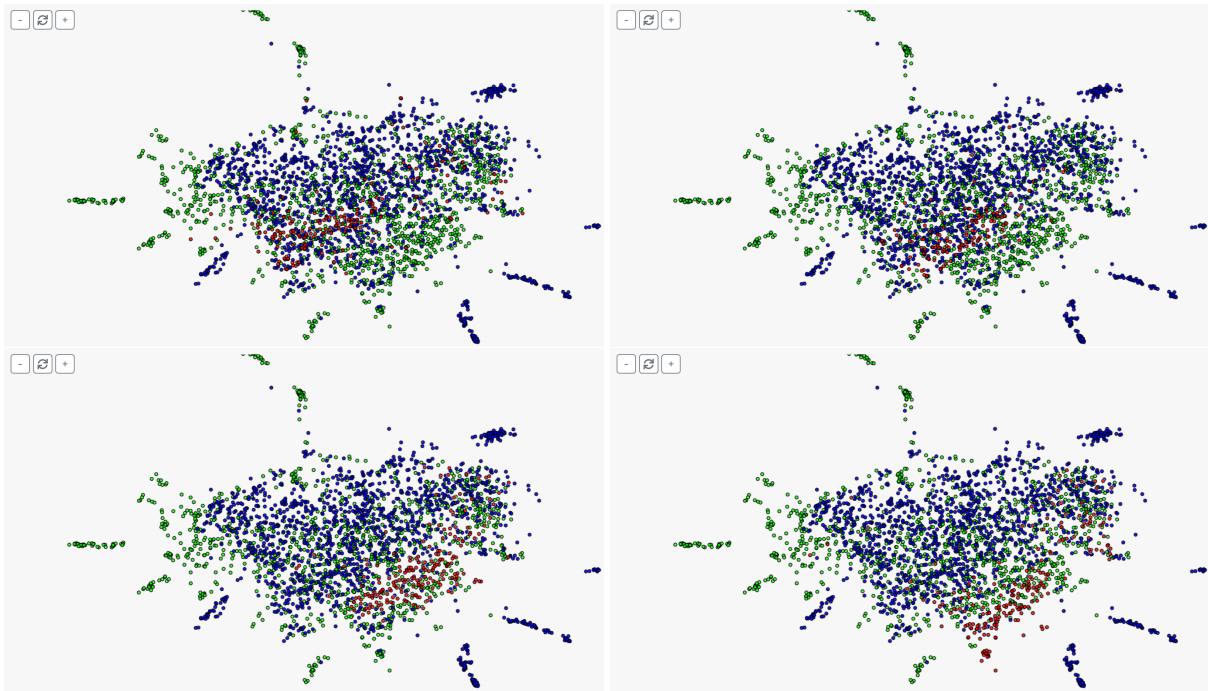


Figure 7: Language-specific clusters (left to right, top to bottom): Spanish, Arabic, Chinese, and Russian. These examples can be found at <https://wildvisualizer.com/embeddings?language=Spanish>, <https://wildvisualizer.com/embeddings?language=Arabic>, <https://wildvisualizer.com/embeddings?language=Chinese>, and <https://wildvisualizer.com/embeddings?language=Russian>.



Figure 8: (a): An outlier cluster related to Midjourney prompt engineering in WildChat. (b): An outlier cluster related to chemistry in LMSYS-Chat-1M. These can be found at <https://wildvisualizer.com/embeddings/english?contains=Midjourney> and <https://wildvisualizer.com/embeddings/english?contains=chemical>.

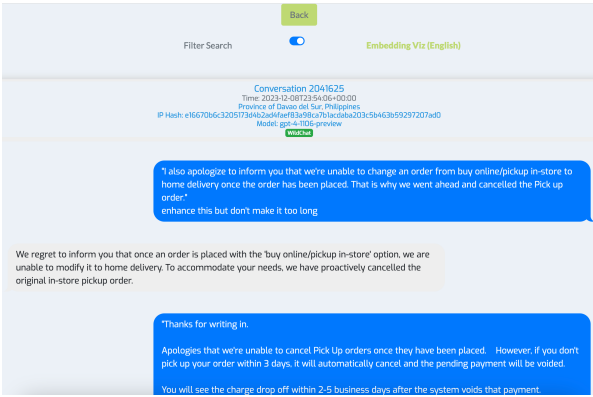


Figure 9: Conversation details page. Clicking any meta-data field filters based on its value. <https://wildvisualizer.com/conversation/wildchat/2041625?from=embedding&lang=english>.

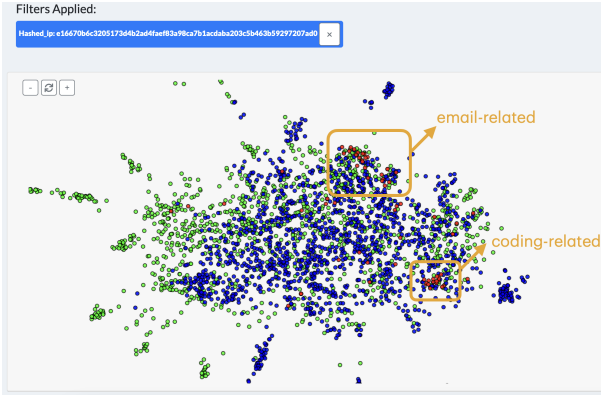


Figure 10: Visualization of conversations from a single user. https://wildvisualizer.com/embeddings/english?hashed_ip=e16670b6c3205173d4b2ad4faef83a98ca7b1acdaba203c5b463b59297207ad0.