



Supporting Human-Centric Data Exploration Through Semantics and Natural Language Interaction

Vidya Setlur
vsetlur@tableau.com
Tableau Research
Palo Alto, California, USA

Abstract

Data science plays an increasingly central role in decision-making across domains, yet the effectiveness of these decisions hinges not only on sophisticated algorithms but also on how well systems support human interpretation, exploration, and communication of data. This tutorial explores the intersection of semantics, natural language processing (NLP), and human-computer interaction in creating human-centric data exploration tools that promote accessibility, trust, and transparency. This includes techniques for generating perceptually meaningful visual encodings, using NLP for query interpretation and ambiguity resolution, and designing conversational interfaces that align with users' intent. The tutorial will highlight research from a broad set of contributors in the SIGMOD/VLDB, HCI, and visualization communities, spanning natural language interfaces for databases, multimodal interaction systems, semantic search for data repositories, and the use of AI and large language models to augment visual and textual analysis.

As part of this 1.5-hour session, we will present case studies and systems that demonstrate how human-centric design can be integrated into the data analysis pipeline, including tools that support mixed-initiative interaction, adaptive defaults, and subjective query interpretation. We will also discuss open challenges and research opportunities, such as semantic inferencing for unstructured data, retrieval-augmented generation (RAG), and ethical considerations around fairness, explainability, and user agency. Drawing from principles in perception, linguistics, AI, and HCI, this tutorial aims to equip attendees with both conceptual frameworks and practical techniques to build more inclusive, interpretable, and intelligent data systems. It is intended for a broad audience in the SIGMOD/VLDB community interested in designing the next generation of data exploration tools that are aligned with human needs.

CCS Concepts

• **Human-centered computing** → **Visualization; Interactive systems and tools.**

Keywords

Human-centered data exploration, natural language interfaces, semantic enrichment, conversational analytics.

ACM Reference Format:

Vidya Setlur. 2025. Supporting Human-Centric Data Exploration Through Semantics and Natural Language Interaction. In *Companion of the 2025 International Conference on Management of Data (SIGMOD-Companion '25)*, June 22–27, 2025, Berlin, Germany. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3722212.3725628>

1 Themes and Topics

As natural language techniques span various aspects of the visual analysis workflow, the field of information visualization needs to consider the challenges and research opportunities that NLP presents as it continues to grow in relevance. This tutorial will explore how semantics, natural language interfaces, and cooperative interaction principles can support human-centric data exploration. We will address the following three themes of research:

Theme 1: Leveraging Semantics for Visual Communication.

Semantics serve as a critical foundation in making data visualizations both perceptually effective and contextually aligned with user expectations. As part of this tutorial, we will explore relevant research on how semantic reasoning can guide the creation of visualizations that communicate insights more intuitively and support user decision-making.

Semantics influence key design elements such as colors and icons, ensuring their alignment with natural cognitive associations. For example, leveraging linguistic co-occurrence data enables visualization systems to assign perceptually distinct and semantically appropriate colors and icons to categorical data [30, 31]. From a perceptual standpoint, semantic alignment has also been shown to improve user comprehension and accessibility. Research from the visualization community has explored how linguistic annotations and textual context influence visual understanding, with growing interest in integrating structured language descriptors into visual summarization tools [26, 33]. These principles are now converging with work in NLP and knowledge representation to create end-to-end pipelines where semantic cues guide both visual generation and interpretation.

While these ideas are well established in the visualization community, semantics-aware techniques are gaining significant traction in the data management research community as well, particularly in aiding schema understanding, dataset discovery, and automated data explanation. Early foundational work on semantic data models [2, 17, 34] emphasizes the importance of capturing inclusion, aggregation, and attribution relationships between entities to better reflect real-world data structures. These models introduce conceptual taxonomies and ontologies that can enrich data interpretation and visualization defaults (e.g., inheritance in class hierarchies or contextual constraints in object-oriented databases). Automated



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGMOD-Companion '25, June 22–27, 2025, Berlin, Germany*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1564-8/2025/06
<https://doi.org/10.1145/3722212.3725628>

Data Explanation [13] is a system that uses external semantic knowledge bases (e.g., DBpedia and Wikipedia) to automatically generate textual explanations and contextual summaries of relational datasets. Their research shows how semantic enrichment can assist users in interpreting data meaningfully, offering mixed-initiative interactions that align with the goals of accessible and explainable visualization systems. In the context of large-scale data lakes, Fan et al [10] present a contrastive learning framework for semantics-aware dataset discovery. By generating contextualized embeddings for columns using pre-trained language models, their research captures nuanced relationships across heterogeneous datasets, facilitating tasks such as table union search and downstream ML applications. These column-level semantic representations can be foundational in guiding automatic visualization recommendations based on similar structures or themes across data repositories.

This growing convergence between visualization, NLP, and data management research highlights a shared goal: using semantics not only to improve perceptual design but also to structure, annotate, and surface data in ways that reflect human understanding and intent. Specifically, this part of the tutorial session will discuss how semantics can:

- Automate the generation of perceptually meaningful and context-aware data encodings through linguistic signals, ontologies, and knowledge bases.
- Enhance accessibility, usability, and interpretability by aligning visual and textual representations with user expectations and domain semantics.
- Support schema and type inference, semantic annotation, and entity linking to improve dataset understanding and searchability in large-scale data environments.
- Enable semantics-aware defaults and recommendations in data exploration tools by leveraging pre-trained models, metadata, and learned embeddings.
- Facilitate adaptive and personalized data presentations across devices and user contexts while preserving core analytical intent and clarity.

Theme 2: Supporting Analytical Conversation Systems

Facilitating conversational interactions with data is a key direction in building intelligent and accessible data systems. This theme will explore research approaches that integrate natural language interfaces (NLIs), pragmatic reasoning, and system-level query interpretation to support cooperative, mixed-initiative interactions. We will highlight techniques that allow users to issue underspecified or ambiguous queries in natural language and receive meaningful, context-aware responses—bridging the gap between human intent and system execution. These methods are particularly relevant to research challenges in semantic parsing, query generation, provenance tracking, and interactive data exploration at scale.

NLIs to support analytical conversations have explored techniques for interpreting user intent and resolving underspecified or ambiguous queries using contextual inferencing techniques, supported by ambiguity widgets that help users iteratively refine or clarify their inputs [9, 14, 28, 32]. These interactions are grounded in Gricean cooperative principles [15], fostering a mixed-initiative dialogue between the system and the user.

There is also research exploring the translation of natural language into SQL queries[25], addressing linguistic ambiguity through user interaction and feedback. Precision interfaces [39] also explore generating customizable UIs from query logs to enable guided, interactive query refinement—blurring the line between structured querying and conversational interaction. Meanwhile, recent systems bring NLI capabilities to spreadsheet environments [12], demonstrating how task-driven natural language interactions can be tightly integrated with familiar productivity tools.

These systems collectively highlight the technical challenges and opportunities of building robust NLIs for data: interpreting user intent, managing ambiguity, generating executable queries, and providing cooperative feedback. They also demonstrate the diversity of application contexts, ranging from databases to dashboards and spreadsheets, emphasizing the importance of both linguistic understanding and system-level design. To ground these ideas in concrete systems and practical innovations, we will review a range of NLI-based tools that exemplify different approaches to conversational data interaction.

This portion of the tutorial will address the following topics:

- Techniques for semantic parsing and contextual inferencing that effectively resolve underspecification and ambiguity in natural language queries.
- System best practices for implementing cooperative conversational principles, covering strategies for incorporating user clarification and feedback, iterative query refinement, and provenance tracking.
- Guidelines for designing interfaces that integrate natural language with other modalities (e.g., visual, touch, and voice inputs) to create adaptive, user-centric systems.
- Frameworks for assessing the effectiveness of conversational mechanisms in analytical tools. We will explore evaluation metrics for query translation accuracy, response latency, and user engagement, as well as discuss standard benchmarking practices.

Theme 3: Recommendation and Search Systems

Navigating the increasing volume, complexity, and heterogeneity of data in modern data ecosystems, such as data lakes, warehouses, and federated systems requires advanced tools that go beyond keyword search or manual exploration. This theme will focus on recommendation and semantic search systems that integrate techniques from natural language processing, information retrieval, and data management to support intuitive, context-aware exploration of structured and semi-structured data.

Visualization recommendation (VizRec) systems play a crucial role in making data exploration more efficient and accessible, especially in large and unfamiliar datasets. These systems aim to bridge the gap between raw data and user insights by automatically surfacing meaningful visual patterns, summaries, or suggestions based on the data's statistical properties or metadata [5, 24, 38, 40]. Other systems recommend visualizations based on statistical heuristics, data patterns (e.g., correlations, outliers, clusters), or user-specified tasks [8, 22, 27, 36, 37]. These systems often precompute summaries or insights, relying on lightweight profiling techniques or materialized views to scale to large datasets. More recent systems [1] go a

step further by incorporating user goals and hypotheses into the recommendation process.

As these visualization recommendation systems increasingly incorporate user intent and data semantics, they naturally intersect with research in search and data discovery, wherein the goal is not just to visualize data, but to intelligently retrieve, interpret, and guide exploration across complex data repositories. Early work in search systems, ranging from structured query interfaces [6, 11, 20] to keyword [3, 19] and natural language-based approaches [4, 7], explore techniques for semantic augmentation in data repositories. This includes using curated metadata (e.g., attribute names, value types), synonym expansion, and external ontologies to enrich search indexing and query interpretation. To support question and answering (Q&A), systems, such as QUERIX[21] and NaLIR[25] demonstrate how natural language inputs can be parsed using dependency parsers (e.g., Stanford CoreNLP [23]) and mapped to structured query representations. These approaches often involve the detection of analytical intents, such as temporal, spatial, or aggregation operations. A unique challenge in data-centric search systems, as opposed to document retrieval, is the sparse textual footprint of many datasets and visualizations. Similar to techniques in multimedia retrieval [18, 35], where image and video content is indexed using external descriptors or metadata, we apply ontological enrichment and visual metadata parsing (e.g., extracting chart types, axes, XML descriptors) to augment searchability and semantic coverage across data collections. Semantic search techniques, originally developed for improving precision in document retrieval, have increasingly influenced how we explore and retrieve structured and semi-structured data from data repositories. These methods aim to interpret user intent and understand the contextual significance of query terms, going beyond literal keyword matching to retrieve more relevant and conceptually aligned results [16, 29]. This portion of the tutorial will explore:

- Techniques for generating real-time, context-aware suggestions during query formulation, leveraging schema profiling, metadata, and learned embeddings. We will highlight approaches that support progressive query construction over structured datasets, enhancing usability in large and heterogeneous data repositories.
- Systems that combine Q&A, exploratory search, and design-oriented search patterns. These approaches support a range of analytical workflows by integrating structured query generation, visual summarization, and pre-authored content retrieval.
- Methods for detecting and operationalizing vague or subjective language in search queries, using techniques from word co-occurrence models and large language models.

Future research topics, discussion, and wrap-up

As the tutorial concludes, we will turn our attention to emerging research directions and open challenges in human-centric data exploration, particularly at the intersection of semantics, data systems, and artificial intelligence (AI). These discussions are intended to spark new ideas and collaborations across the community, where core advances in query processing, data modeling, and system architecture can meaningfully contribute to more intelligent and accessible data experiences.

One major area of interest is the semantic understanding of unstructured and semi-structured data. As data systems increasingly ingest heterogeneous content—from text and spreadsheets to dashboards, logs, and sensor data—semantic enrichment becomes vital for tasks such as schema discovery, metadata generation, entity resolution, and intent modeling. We will explore how Generative AI models, particularly RAG frameworks, can be used to interpret and surface insights from diverse sources, such as documents, reports, and transcripts, enriching query interfaces and recommendation systems with contextual knowledge. At the same time, we face pressing questions about the role of human creativity and agency in increasingly automated analytical workflows. While LLMs can suggest queries, generate visualizations, and summarize results, future tools must move beyond passive automation to actively support user-driven exploration and decision-making. Ultimately, these future directions point toward a new generation of data systems that combine the strengths of declarative querying, AI augmentation, and interactive design. The session will close with an open discussion, inviting participants to reflect on the tutorial's themes and contribute their perspectives on future research opportunities.

2 Audience

This tutorial is designed for researchers, practitioners, and graduate students in data management, information retrieval, natural language processing, and human-centered data systems who are interested in the integration of semantics, AI, and interaction in data exploration workflows. It is especially relevant for those working on or interested in natural language interfaces for databases, semantic search, visualization recommendation systems, data discovery, and AI-augmented analytics.

Attendees should have a foundational understanding of data management principles (e.g., query processing, schema design, or indexing), along with a basic familiarity with natural language processing and data visualization. While advanced knowledge of machine learning or visualization is not required, some experience with AI models (e.g., transformers, embeddings, or retrieval-augmented generation) will help participants engage more deeply with the content. The tutorial aims to bridge communities across SIGMOD/VLDB, HCI, and visualization disciplines, providing conceptual frameworks, system-level insights, and emerging research opportunities for building human-centric, intelligent data systems.

3 Presenter

Vidya Setlur is the senior director of Tableau Research. Her area of expertise is in natural language processing and computer graphics. She earned her doctorate in 2005 at Northwestern University and loves pushing the boundaries of product innovation, drawing from both research and engineering. Vidya's interests are in developing novel computer algorithms and user interfaces that enhance visual communication and the understanding of user intent through semantics. Her research draws inspiration from information retrieval, computational linguistics, human perception, and cognitive science to help users effectively interact with devices and information in their environment.

Website: <https://www.tableau.com/research/people/vidya-setlur>

References

- [1] Fatma Bouali, Abdelheq Et-tahir Guettala, and Gilles Venturini. 2016. VizAssist: An Interactive User Assistant for Visual Data Mining. *The Visual Computer: International Journal of Computer Graphics* 32, 11 (2016), 1447–1463. <https://doi.org/10.1007/s00371-015-1132-9>
- [2] Mokrane Bouzeghoub and Elisabeth Métais. 1991. Semantic Modeling of Object Oriented Databases. In *17th International Conference on Very Large Data Bases, September 3–6, 1991, Barcelona, Catalonia, Spain, Proceedings*, Guy M. Lohman, Amílcar Sernadas, and Rafael Camps (Eds.). Morgan Kaufmann, 3–14. <http://www.vldb.org/conf/1991/P003.PDF>
- [3] Gong Cheng, Weiye Ge, and Yuzhong Qu. 2008. Falcons: Searching and Browsing Entities on the Semantic Web. In *Proceedings of the 17th International Conference on World Wide Web (Beijing, China) (WWW '08)*. Association for Computing Machinery, New York, NY, USA, 1101–1102. <https://doi.org/10.1145/1367497.1367676>
- [4] Philipp Cimiano, Peter Haase, Jörg Heizmann, Matthias Mantel, and Rudi Studer. 2008. Towards Portable Natural Language Interfaces to Knowledge Bases - The Case of the ORAKEL System. *Data & Knowledge Engineering* 65, 2 (May 2008), 325–354. <https://doi.org/10.1016/j.datak.2007.10.007>
- [5] Christopher Collins, Natalia Andrienko, Tobias Schreck, Jing Yang, Jaegul Choo, Ulrich Engelke, Amit Jena, and Tim Dwyer. 2018. Guidance in the human-machine analytics process. *Visual Informatics* 2, 3 (2018), 166–180.
- [6] Olivier Corby, Rose Dieng-Kuntz, and Catherine Faron-Zucker. 2004. Querying the Semantic Web with the CORESE search engine. *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI'2004)*, 705–709.
- [7] Danica Damjanovic, Milan Agatonovic, and Hamish Cunningham. 2010. Natural Language Interfaces to Ontologies: Combining Syntactic Analysis and Ontology-Based Lookup through the User Interaction. In *Extended Semantic Web Conference*.
- [8] Çağatay Demiralp, Peter J Haas, Srinivasan Parthasarathy, and Tejaswini Pedapati. 2017. Foresight: Rapid Data Exploration through Guideposts. *arXiv preprint arXiv:1709.10513* (2017).
- [9] Kedar Dhamdhare, Kevin S. McCurley, Ralfi Nahmias, Mukund Sundararajan, and Qiqi Yan. 2017. Analyza: Exploring Data with Conversation. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces (IUI 2017)*. 493–504.
- [10] Grace Fan, Jin Wang, Yuliang Li, Dan Zhang, and Renée J. Miller. 2023. Semantics-Aware Dataset Discovery from Data Lakes with Contextualized Column-Based Representation Learning. *Proc. VLDB Endow.* 16, 7 (March 2023), 1726–1739. <https://doi.org/10.14778/3587136.3587146>
- [11] Tim Finin, Li Ding, Rong Pan, Anupam Joshi, Pranam Kolari, Akshay Java, and Yun Peng. 2005. Swoogle: Searching for Knowledge on the Semantic Web. In *Proceedings of the 20th National Conference on Artificial Intelligence - Volume 4 (Pittsburgh, Pennsylvania) (AAAI'05)*. AAAI Press, 1682–1683.
- [12] Derek Flood, Kevin Mcdaid, and Fergal McCaffery. 2009. NLP-SIR: A Natural Language Approach for Spreadsheet Information Retrieval. *CoRR abs/0908.1193* (01 2009).
- [13] Sainyam Galhotra and Udayan Khurana. 2022. Automated Relational Data Explanation using External Semantic Knowledge. *Proc. VLDB Endow.* 15, 12 (Aug. 2022), 3562–3565. <https://doi.org/10.14778/3554821.3554844>
- [14] Tong Gao, Mira Dontcheva, Eytan Adar, Zhicheng Liu, and Karrie G. Karahalios. 2015. DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software Technology (UIST 2015)*. ACM, New York, NY, USA, 489–500.
- [15] H. Paul Grice. 1975. Logic and Conversation. In *The Logic of Grammar*, Donald Davidson (Ed.). Dickenson Pub. Co., 64–75.
- [16] Ramanathan V. Guha, Rob McCool, and Eric Miller. 2003. Semantic Search. In *The Web Conference*.
- [17] Michael Hammer and Dennis McLeod. 1978. The Semantic Data Model: A Modelling Mechanism for Database Applications. In *Proceedings of the 1978 ACM SIGMOD International Conference on Management of Data (Austin, Texas) (SIGMOD '78)*. Association for Computing Machinery, New York, NY, USA, 26–36. <https://doi.org/10.1145/509252.509264>
- [18] Samira Hammiche, Salima Benbernou, Mohand-Saïd Hacid, and Athena Vakali. 2004. Semantic Retrieval of Multimedia Data. In *Proceedings of the 2nd ACM International Workshop on Multimedia Databases (Washington, DC, USA) (MMDB '04)*. Association for Computing Machinery, New York, NY, USA, 36–44. <https://doi.org/10.1145/1032604.1032612>
- [19] A. Harth, Aidan Hogan, Renaud Delbru, Jürgen Umbrich, Seán O'Riain, and Stefan Decker. 2007. SWSE: Answers Before Links!. In *Semantic Web Challenge*.
- [20] Jeff Hefflin, James A. Hendler, and Sean Luke. 2003. SHOE: A Blueprint for the Semantic Web. In *Spinning the Semantic Web*.
- [21] Esther Kaufmann, Abraham Bernstein, and Renato Zumstein. 2006. Querix: A Natural Language Interface to Query Ontologies Based on Clarification Dialogs. (01 2006).
- [22] Alicia Key, Bill Howe, Daniel Perry, and Cecilia Aragon. 2012. VizDeck. *Proceedings of the 2012 International Conference on Management of Data - SIGMOD '12* (2012), 681. <https://doi.org/10.1145/2213836.2213931>
- [23] Dan Klein and Christopher D. Manning. 2003. Accurate Unlexicalized Parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Sapporo, Japan, 423–430. <https://doi.org/10.3115/1075096.1075150>
- [24] Doris Jung-Lin Lee, Vidya Setlur, Melanie Tory, Karrie Karahalios, and Aditya Parameswaran. 2021. Deconstructing Categorization in Visualization Recommendation: A Taxonomy and Comparative Study. *arXiv:2102.07070 [cs.HC]*
- [25] Fei Li and H. Jagadish. 2014. Constructing an Interactive Natural Language Interface for Relational Databases. *Proceedings of the VLDB Endowment* 8 (09 2014), 73–84. <https://doi.org/10.14778/2735461.2735468>
- [26] Alan Lundgard and Arvind Satyanarayan. 2022. Accessible Visualization via Natural Language Descriptions: A Four-Level Model of Semantic Content. *IEEE Trans. Vis. Comput. Graph.* 28, 1 (2022), 1073–1083. <https://doi.org/10.1109/TVCG.2021.3114770>
- [27] Jock Mackinlay, Pat Hanrahan, and Chris Stolte. 2007. Show Me: Automatic Presentation for Visual Analysis. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1137–1144.
- [28] Vidya Setlur, Sarah E. Battersby, Melanie Tory, Rich Gossweiler, and Angel X. Chang. 2016. Eviza: A Natural Language Interface for Visual Analysis. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16)*. 365–377. <https://doi.org/10.1145/2984511.2984588>
- [29] Vidya Setlur, Andriy Kanyuka, and Arjun Srinivasan. 2023. Olio: A Semantic Search Interface for Data Repositories. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (San Francisco, CA, USA) (UIST '23)*. Article 95, 16 pages. <https://doi.org/10.1145/3586183.3606806>
- [30] Vidya Setlur and Jock D. Mackinlay. 2014. Automatic Generation of Semantic Icon Encodings for Visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Toronto, Ontario, Canada) (CHI '14)*. New York, NY, USA, 541–550. <https://doi.org/10.1145/2556288.2557408>
- [31] Vidya Setlur and Maureen C. Stone. 2016. A Linguistic Approach to Categorical Color Assignment for Data Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 698–707. <https://doi.org/10.1109/TVCG.2015.2467471>
- [32] Arjun Srinivasan and John Stasko. 2018. Orko: Facilitating Multimodal Interaction for Visual Exploration and Analysis of Networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 511–521.
- [33] Chase Stokes, Cindy Xiong Bearfield, and Marti A. Hearst. 2024. The Role of Text in Visualizations: How Annotations Shape Perceptions of Bias and Influence Predictions. *IEEE Transactions on Visualization and Computer Graphics* 30, 10 (Oct. 2024), 6787–6800. <https://doi.org/10.1109/TVCG.2023.3338451>
- [34] Veda C. Storey. 1993. Understanding Semantic Relationships. *The VLDB Journal* 2, 4 (Oct. 1993), 455–488. <https://doi.org/10.1007/BF01263048>
- [35] Gonzalo Vaca-Castano and Mubarak Shah. 2015. Semantic Image Search From Multiple Query Images. In *Proceedings of the 23rd ACM International Conference on Multimedia (Brisbane, Australia) (MM '15)*. Association for Computing Machinery, New York, NY, USA, 887–890. <https://doi.org/10.1145/2733373.2806356>
- [36] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2016. Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 649–658. <https://doi.org/10.1109/TVCG.2015.2467191>
- [37] Kanit Wongsuphasawat, Zening Qu, Dominik Moritz, Riley Chang, Felix Ouk, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2017. Voyager 2 : Augmenting Visual Analysis with Partial View Specifications. (2017). <https://doi.org/10.1145/3025453.3025768>
- [38] Aoyu Wu, Yun Wang, Xinhuan Shu, Dominik Moritz, Weiwei Cui, Haidong Zhang, Dongmei Zhang, and Huamin Qu. 2021. Survey on Artificial Intelligence Approaches for Visualization Data. *arXiv preprint arXiv:2102.01330* (2021).
- [39] Haoci Zhang, Viraj Raj, Thibault Sellam, and Eugene Wu. 2018. Precision Interfaces for Different Modalities. In *Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD '18)*. Association for Computing Machinery, New York, NY, USA, 1777–1780. <https://doi.org/10.1145/3183713.3193570>
- [40] Sujia Zhu, Guodao Sun, Qi Jiang, Meng Zha, and Ronghua Liang. 2020. A Survey on Automatic Infographics and Visualization Recommendations. *Visual Informatics* 4, 3 (2020), 24–40.