

# Human-Data Interaction, Exploration, and Visualization in the AI Era: Challenges and Opportunities

Jean-Daniel Fekete<sup>1</sup> Yifan Hu<sup>2</sup> Dominik Moritz<sup>3</sup> Arnab Nandi<sup>4</sup>  
Senjuti Basu Roy<sup>5</sup> Eugene Wu<sup>6</sup> Nikos Bikakis<sup>7</sup> George Papastefanatos<sup>8</sup>  
Panos K. Chrysanthis<sup>9</sup> Guoliang Li<sup>10</sup> Lingyun Yu<sup>11</sup>

<sup>1</sup>Univ. Paris-Saclay & Inria, France <sup>2</sup>Northeastern Univ., USA <sup>3</sup>Carnegie Mellon Univ. & Apple, USA

<sup>4</sup>Ohio State Univ., USA <sup>5</sup>New Jersey Inst. of Technology, USA <sup>6</sup>Columbia Univ., USA

<sup>7</sup>Hell. Med. Univ. & Archimedes/Athena RC, Greece <sup>8</sup>Athena RC, Greece

<sup>9</sup>Univ. of Pittsburgh, USA <sup>10</sup>Tsinghua Univ., China <sup>11</sup>Xi'an Jiaotong-Liverpool Univ., China

## Abstract

The rapid advancement of AI is transforming human-centered systems, with profound implications for human-AI interaction, human-data interaction, and visual analytics. In the AI era, data analysis increasingly involves large-scale, heterogeneous, and multimodal data that is predominantly unstructured, as well as foundation models such as LLMs and VLMs, which introduce additional uncertainty into analytical processes. These shifts expose persistent challenges for human-data interactive systems, including perceptually misaligned latency, scalability constraints, limitations of existing interaction and exploration paradigms, and growing uncertainty regarding the reliability and interpretability of AI-generated insights. Responding to these challenges requires moving beyond conventional efficiency and scalability metrics, redefining the roles of humans and machines in analytical workflows, and incorporating cognitive, perceptual, and design principles into every level of the human-data interaction stack. This paper investigates the challenges introduced by recent advances in AI and examines how these developments are reshaping the ways users engage with data, while outlining limitations and open research directions for building human-centered AI systems for interactive data analysis in the AI era.

## Keywords

Human-AI Collaboration, Visual Analytics, Multimodal Analytics, Generative Visualization, AI-guided Exploration, Foundation Models, Trustworthy AI

## 1 Introduction

The growing influence of AI is fundamentally changing how humans interact with, explore, and reason about data, as well as how human-centered systems are designed and implemented. Advances in AI-driven techniques have expanded the scope of interactive data analysis beyond structured tables to include *large-scale, heterogeneous, and multimodal data* such as text, images, audio, video, and visually rich documents. As a result, data analysis increasingly involves information that is unstructured, weakly indexed, and costly to inspect, which challenges traditional *query-then-explore workflows* that rely on well-defined schemas, stable metadata, and

## ★ Key Insights

- 1. **The AI era introduces new, tightly coupled requirements for human-data interaction systems.** Systems must deliver perceptually aligned responsiveness and scalability, natively handle multimodal and unstructured data, support guidance and emerging interaction and exploration paradigms, leverage visualization capabilities, and remain robust to AI uncertainty through progressive feedback and trustworthy analysis.
- 2. **Latency and scalability are fundamental barriers.** Challenges such as multimodal data, next-generation interfaces, embeddings, complex interactions, and high-velocity streams, require rethinking core data management techniques to meet perceptual latency constraints and support the development of scalable and reliable human-centered systems. Also, DL models must also scale to large training datasets.
- 3. **End-to-end system design is essential.** Interfaces must evolve toward multimodal, action-oriented, and adaptive environments that integrate modalities such as NL, gestures, narrative-rich generative visualizations, and AR/VR. Supporting these capabilities, as well as low-latency performance and scalability, requires coordinated design across system architectures, data operations, AI components, interfaces, and visualization.
- 4. **AI models are powerful but require human-in-the-loop oversight.** AI-driven techniques can enable guidance-driven exploration, semantic understanding, data-oriented and exploratory tasks, and improved system usability. However, their brittleness, error-proneness, unpredictability, and biases underscore that human involvement remains indispensable.
- 5. **Visualization is shifting from static depiction to generative and adaptive guidance.** Visualization is no longer static but an active, generative system component. Progressive, narrative, and aesthetics-aware visualizations guide attention, enhance interpretability, and mediate trust in AI-driven analytics.
- 6. **Sustained cross-disciplinary collaboration is crucial.** The diversity of emerging challenges underscores the need for close collaboration among the systems, databases, AI, information visualization, HCI, computer graphics, and cognitive science communities.

predictable access patterns. These shifts place new demands on human-centered systems.

In practice, users must *iteratively explore data, reformulate queries, and verify intermediate results under uncertainty*; often without a clear understanding of what the data contains or which questions can be meaningfully asked. At the same time, *interactive analysis imposes strict perceptual latency constraints*. Delays of even a few seconds can disrupt analytical reasoning, bias exploration, and reduce the effectiveness of human-AI collaboration. Taken together, these challenges reveal fundamental limitations in existing system architectures, data management techniques, interaction paradigms, and visualization approaches.

Addressing the previously discussed challenges requires a fundamental rethinking of interactive data systems as *end-to-end, human-centered pipelines*. Rather than treating data management, AI models, interfaces, and visualization as loosely coupled components, future systems must be *co-designed* to support iterative, guidance-driven, and mixed-initiative exploration.

*Latency and scalability* emerge as cross-cutting constraints. They shape not only system performance, but also user behavior, cognitive load, and the feasibility of exploratory workflows. Supporting interaction at the *speed of human thought* demands new abstractions, query processing methods, data indexing schemes, and approximation strategies that are explicitly aware of interface structure, human perception, and interaction context.

The growing reliance on foundation models (e.g., LLMs and VLMs) and other deep learning models (e.g., GANs and GNNs) also introduces new concerns, including non-determinism, error propagation, bias, and the cost of verification. These issues reinforce the necessity of *keeping humans meaningfully involved in the analytical loop*, with explicit support for oversight, provenance, and trust calibration.

Interfaces themselves are also evolving. *Next-generation interfaces* increasingly move toward multimodal, action-oriented environments that integrate natural language interaction, gestures, immersive technologies such as augmented and virtual reality (AR/VR), and continuous feedback. These interfaces enable more direct and intuitive engagement with complex data, but they also impose new requirements on underlying systems, including *real-time responsiveness, progressive result refinement, and adaptive guidance*. In such settings, interaction becomes tightly coupled with perception and action, further amplifying the need for co-design across systems, interfaces, and visualization.

In this evolving landscape, *visualization* play an increasingly central role. Visualization is no longer a passive output of analysis, but an *active system component* that supports exploration, query reformulation, guidance, verification, and feedback. Progressive, adaptive, and narrative-driven visual representations can help users navigate uncertainty, focus attention, and assess the reliability of AI-assisted results.

In summary, the AI era ecosystem requires that *human-data interaction systems deliver real-time, cognition-aware performance and scalability, effective multimodal data handling, and advanced visual analytics, while simultaneously supporting user guidance, progressive sensemaking, and trustworthy analysis under AI uncertainty*.

Drawing on perspectives from databases, AI, information visualization, human-computer interaction, and computer graphics, this

paper presents the key challenges and outlines research directions for building scalable, reliable, and human-centered interactive data systems in the AI era.

We also point readers to the following review papers for further details on the discussed topics: (1) Human-AI collaboration [1–4]; (2) Human-centered data management [5–17]; (3) Visual analytics [18–21]; and (4) AI and visual analytics [22–32].

Each section presents the perspective and insights of an author, as follows: Eugene Wu (Sec. 2), Dominik Moritz (Sec. 3), Jean-Daniel Fekete (Sec. 4), Arnab Nandi (Sec. 5), Senjuti Basu Roy (Sec. 6), and Yifan Hu (Sec. 7). The first author contributed equally to the manuscript content, while the last five authors coordinated the preparation, harmonization, and enrichment of the manuscript.

## 2 Interface-System Co-Design

Data management as a discipline is fundamentally about making it easier to see, analyze, and interact with data by designing data models, query abstractions, and execution engines—always with an application class in mind. Data visualization pursues a closely aligned goal, but centers its focus on user experience, perceptually-aware design, and interaction techniques. Systems for human-data interaction sits at the intersection of both disciplines and must account for both usability, resource, and system constraints, which often conflict with each other. However, doing so is difficult at arm’s length. Progress requires not only modeling interfaces and user interactions in ways that make system design tractable, but also leveraging system and query capabilities to expand the interfaces that are possible.

This section summarizes the evolution of how data systems have used increasingly richer models of the interface to inform the system designs. We further argue that system design need not be reactive to *known* application needs, but can be inspiration for altogether new interface and interaction designs. These arguments are discussed in greater depth in Wu et al. [5].

**Interfaces as SQL-Workloads.** The most direct approach—what we term the SQL-workload approach—is to treat visualization as a SQL-workload generator. This is sensible because most visualization systems indeed translate user interactions into SQL queries, execute those queries in a DBMS, and then update the interface with the query results. For instance, IDEBench [17] generates SQL queries inspired by cross-filtering visualizations. Under this model, improving query execution throughput or latency directly improves the responsiveness of the interfaces, and variations of classic database techniques such as approximation [33–46], query prefetching [47–51], shared scans [52], precomputation [53–55], and indexing [55–58] have shown success. We refer to [5, 8, 11] for survey.

**Interfaces as System Constraints.** *However, this SQL-workload approach does not focus on the interface’s unique set of constraints.* These constraints motivate new and challenging classes of data management problems. First, as pointed out by Jiang et al. [15], *simple workload statistics like P95 latency or query throughput do not translate into fluid user experiences.* A single click may trigger a dozen parallel queries and the user expectation can vary from all queries responding simultaneously, only the queries “above the fold”, or only the query for the “central chart”. In addition, *when*

components of the interface update in response to interactions matters [12, 59–61].

As Chen et al. [62] argue, there is a *Design Dependence* between how interfaces and interactions are specified and the system designs required to support them. In practice, this dependence cuts both ways: systems are often aggressively optimized for a narrow set of visualization and interaction patterns, which in turn biases designers toward interfaces that are easy to build rather than those that are desirable or expressive. As researchers develop increasingly specialized optimizations—each encodes implicit assumptions about the user needs, interface designs, and interaction semantics—the complexity of selecting, configuring, and composing these optimizations grows combinatorially. Much like physical database design, this burden should not be placed on human designers. This can result in a form of “streetlight effect,” where interface designs converge around what existing systems easily support, rather than exploring what is possible or effective. This motivates the need for interface-aware system optimizers that can reason jointly about interface structure and system design [62].

Second, SQL-workloads are only representative of visualizations created *today*, which themselves are limited by what *today’s* systems can support. For instance, the vast majority of visualization systems, and thus the SQL-workloads, assume that the input is a single table or view. Yet hiding the underlying join dependencies introduces subtle semantic errors [63, 64]. Techniques based on factorized query execution and semiring aggregation have been successfully adapted to scale interactive visual analysis over large join graphs [65], and introduce a variation of view maintenance under incremental changes to the query, rather than the input data.

**Interfaces as First-class Abstractions.** As a full-stack community, we have the option of introducing new system designs, and changing the system abstraction to access new performance frontiers and empowering interfaces with new capabilities. This can vary from introducing a new API so designers can specify new models of user perception to inform approximation algorithms [34], replacing traditional request-response prefetching with a communication model inspired by video streaming [50], and identifying the equivalence between interaction and provenance [66] to simplify interface development [67] and advance provenance systems research [68, 69] at the same time.

Existing data management research naturally translates into new visualization features, but requires a careful co-design with interface and interactivity requirements. For instance, the data management community has explored algorithmic techniques to explain analysis results [70–75], but careful design, interactive latencies, and simple abstractions are needed for wide-spread adoption in *any* visualization [76]. As another example, data cube and comparison theory from two decades ago [77] can motivate new interaction designs to support visual comparison of any combination of charts, marks, and/or values in a data interface [78, 79].

**Rethinking Visualization Interfaces Altogether.** It is also valuable to rethink how visualizations are designed, optimized, and implemented. Today, this requires a unicorn designer who is an expert in the domain, visualization design and implementation, and system optimization, or to coordinate hand-offs of a team of designers and developers—known to be fraught with challenges [20]. What if we

## 📖 Main Takeaway [Section 2]

- Co-design of data systems and interfaces is critical: performance, system abstractions, and interaction design jointly determine user experience, and visualization cannot be optimized independently of data management.
- Traditional workload-centric metrics (e.g., query latency) are insufficient; interface-aware performance models that account for perceptual constraints, interaction structure, and partial results are necessary to support fluid, interactive analytics.
- New system abstractions and AI integration enable innovation. Introducing new APIs, communication models, and provenance-aware designs, combined with automating interface design and extending visualization theory beyond single-table inputs, can make interfaces more powerful, interactive, and capable of fully leveraging AI.

could synthesize the interface design (or thousands of designs) simply from samples of the desired analysis queries [80, 81], annotate the interactions with their latency requirements, and automatically instantiate an end-to-end client-server architecture that, when pointed to a cloud DBMS, guarantees the desired interactivity [62]?

Data visualization is a young and evolving field, and data management research can also inform new visualization abstractions and theories. For instance, what is an intermediate representation of the interface and analysis that is easy to write, amenable to automatic optimization, and can be used to rapidly synthesize new interface designs? Recent work by Chen et al. [82] proposes a Data Interface Grammar, but this remains an open challenge. Another open question is the line between data processing and visualization—how much of the visualization and user analysis process can be modeled as “just a query?” [83].

Finally, as hinted at above, visualization theory that modern visualization grammars and libraries are predicated on, and fundamentally assumes, that the input is a single table. Yet relational theory makes it clear that a single-table does not adequately model all data. Is there a way to extend “single-table” visualization theory to a theory of *Database Visualization* [84, 85]? This potentially could expand the scope of visualization systems much in the same way as the introduction of the relational model [86].

**What About AI?** We would be remiss to not address how these challenges relate to AI and LLMs. Time and again, we have seen evidence that LLMs can mimic syntactic structure and generate natural-looking outputs based on strong priors in its training data, but have difficulty reasoning about logical constraints and semantic structure. As such, AI heavily relies on well-designed and well-curated tools that are appropriate for a given task. This implies that new and clean abstractions are *more important than ever!*

🔗 **Connections to Other Sections.** The importance of tight co-design between backends and interfaces, along with the revision of data management techniques raised in this section, is also highlighted in the article as mandatory to enable cold-start exploration (Sec. 4), order-of-millisecond response (Sec. 3), and real-time feedback in multimodal and AR/VR querying environments (Sec. 5).

### 3 Queries at the Speed of Human Thought

Machine learning (ML) development has shifted from being primarily model-centric to primarily data-centric [87]. In the early days of ML, engineers often trained models on hundreds or thousands of data points and meticulously tweaked the modeling function to avoid overfitting. With the popularization of Deep Learning, however, data became increasingly important, and many models were derived from pre-trained models. Today, foundation models have similar transformer architectures, and big quality differences stem from the data on which they are trained or fine-tuned. The data management community is well equipped to tackle many of the quality and quantity challenges [88]. However, the full AIML lifecycle of requirement elicitation [89], data preparation, monitoring, tuning, and evaluation requires oversight by people who have to slice and dice millions or billions of data records. Since models synthesize patterns, a single record rarely defines the behavior of a model, and therefore ML engineers need to understand patterns and trends in relevant subsets [90]. Relevant subsets are hard to predict and can rarely be precomputed. ML engineers need to grasp often subtle patterns in large datasets full of information. As Herbert A. Simon puts it, “What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently” [91]. *This need to efficiently allocate attention is where data visualization becomes crucial.*

Data visualization leverages our powerful perceptual systems to let us see patterns and trends in data. It is a critical part of any analysis and a precursor to any rigorous statistical analysis [92]. Often, simply looking at the data from various angles uncovers new opportunities for improvement. Data visualization in the AIML lifecycle let us overview and design datasets for training, fine-tuning, and evaluation. It enables the serendipitous discovery of patterns and data issues. At the core of the necessary interactive analysis are interfaces that are fast and fast to use. Well-designed mixed-initiative systems [93] for analysis make people faster [94, 95]. We should continue to invest in good tooling that eliminates barriers to effective analysis. Yet, many data management systems are not designed for interactive (order-of-millisecond) latencies but rather for seconds or minutes. *These delays in interactive systems lead to fewer observations made [96] and could steer analysts towards convenient and fast data along with all the resulting biases.* Fast databases, approximation [33–46], prefetching [47–51], and indexing [55–58] can help developers build fast interactive systems. However, they add complexity for tool builders that may prevent quick adaptation of interfaces to new needs. Modern data architectures abstract from low-level optimizations like Mosaic [97] and are being rapidly adopted. Yet, many opportunities remain for deeper integration with databases, which are often optimized for order-of-second responses rather than real-time order-of-millisecond responses or may not support many of the encoding transformations used in effective data visualizations like cartographic projections. A lot of exciting work remains to explore new way to *effectively present billion-record datasets at the speed of human thought to facilitate effective analysis and communication.*

**Connections to Other Sections.** The demand for interaction at the speed of human cognition discussed here is also highlighted in several parts of the paper, mainly in Sections 2, 4, and 5. Furthermore, the mixed-initiative perspective, as well as similar issues highlighted in this section, are related to the need for succinct visual guidance in multimodal exploration discussed in Section 5, as well as to Section 7, where aesthetics and narrative are suggested as mechanisms for directing user attention.

#### Main Takeaway [Section 3]

- Nowadays, machine learning development depends heavily on data rather than model tweaks. In the AIML lifecycle, data visualization is an effective means for exploring and designing datasets used in training, fine-tuning, and evaluation.
- Data management systems are often optimized for order-of-second responses. On the other hand, interactive analysis requires response times aligned with human cognition (order-of-millisecond); delays distort exploration, bias observations, and reduce analytical effectiveness.
- Data management techniques such as approximation, indexing, and prefetching can offer interactivity but increase complexity for developers. New abstractions and architectures are needed to simplify interfaces adaptation to new requirements.

### 4 Scalable Visual Analytics

Visualization is a recognized component of data science, but unlike its sibling domains, such as databases, simulation, and machine learning, its research articles are not as concerned with scalability. Only 2% of the articles published at the IEEE VIS conference use a keyword related to scalability, compared to the other domains where scalability and management of large amounts of data concern most of the published articles. Yet, in recent years, several articles have tried to improve the situation, either by describing novel “scalable” visualization systems [97], better defining what “scalability” means in visualization [9], or introducing novel programming paradigms for scaling visualization, such as progressive visual analytics [98–100] and approximate query processing (AQP) [33–46]. However, *the big picture of scalability in visualization is not clear yet.* Is there a single, consensual measure to use in determining whether we have achieved scalability in visualization, or several? What are the next challenges we need to address to fill the scalability gap between visualization and the other data science domains, and is it realistic to believe we will eventually fill that gap? We identify two key issues to address for achieving scalability in visualization, in addition to handling large amounts of data.

**Prepared vs. Cold-Started Visualization.** Visualization systems designed to explore data are interactive and should provide a controlled latency to be usable [96]. But what is an acceptable latency to *start* data analysis? Most visualization systems advertised for their scalability *do not consider data preparation time a major issue*, assuming that the exploration time will far exceed the preparation

time or can be done by someone other than the data analyst. However, this assumption may no longer be true. With the growing number of public datasets available, finding out which dataset or group of datasets could help answer a question becomes part of the data exploration process [101], and the time from deciding to explore a dataset to visualization cannot be considered as *preparation* but should rather be included in the exploration process. The prepared systems are well-suited to recurring explorations with well-defined queries that can be optimized using indexes. Cold-starting systems are suited to one-time or more open-ended explorations [10, 102, 103], such as journalist investigations or discovering datasets and topics within them. The prepared systems correspond to traditional use cases, whereas the cold-starting systems have become more important with the recent profusion of open datasets on the web. Scalability should now consider *the new use cases made possible by integrating and joining unexplored datasets on-the-fly. Scalable systems should mention their assumptions regarding preparation time when reporting on their scalability.*

Cold-starting visualization comes with several new challenges that are well-suited to the progressive approach but will require more research and technical development to become mainstream since the necessary algorithms and data management systems are not yet ready off the shelf.

**Scalability of Visualization Techniques.** Most visualization techniques are not designed to scale to large amounts of data. When dealing with large-scale data, uncontrollable issues arise, such as large numbers of categories, outliers, and pathological distributions with uneven value densities and possible long tails. Some of these issues have been documented by Wickham [104], but none of the popular visualization libraries provide mechanisms to overcome them. Over the last fifteen years, the visualization community has been relying on the Grammar of Graphics (GoG) [105] to express visualizations; however, the GoG specifications are not designed to be resilient to big data issues. It would require the visualization programmer to foresee potential data issues and design the right GoG specification before examining the data, which is unrealistic. When dealing with large amounts of data, visualizations should be designed to be resilient to data issues: applying appropriate aggregations when necessary, filtering out outliers without hiding them completely, and showing aggregated data in dense areas and non-aggregated data in sparse areas. Building GoG specifications for such operations might be possible, but certainly much more complex than the concise specifications the GoG has been designed for.

Scalable visualization systems should rely on data analysis and hybrid visualizations—such as overlaying outliers onto a density map for scatterplots [106]—to adapt to data issues. Wickham’s Bin-summarise-smooth framework describes operators such as “peel” to automatically remove low-density extrema from the visualizations. The GoG allows data transformations to better adapt channels to data, such as using a logarithmic transformation for the color scale or for positional axes when the data distribution is not normal. However, these transformations are explicit in the GoG at the specification level. Scalable visualizations should help analysts in inferring reasonable default values. Scalable visualizations should provide more automation than the GoG allows, and this

#### 📖 Main Takeaway [Section 4]

- Visualization research needs to place greater emphasis on scalability, starting with a clear definition of what scalability entails.
- Cold-start visualization requires further development to enable real-time exploration of new or large datasets on the fly.
- Existing visualization techniques must become adaptive and data-aware, automatically handling outliers, skewed distributions, dense regions, and large category counts through aggregation and hybrid representations.
- New paradigms are needed for scalable visual analytics, including progressive computation and approximate query processing.

automation should rely on data analysis and hybrid visualization techniques.

Visualization scalability will not come solely from speeding up algorithms and data structures, or predicting user actions; many paradigms must change. *Optimizing the traditional visualization pipeline will not scale to the three-to-five orders of magnitude required to keep up with the other data science domains;* progressive and AQP programming should take over. For specifying visualizations, the GoG is too simple to adapt to the unavoidable issues found in large-scale datasets. *More complex visualization specifications are needed, based on data analysis and adaptation rules.*

🔗 **Connections to Other Sections.** The need for perceptually aligned latency and progressive computation mentioned in this section is also highlighted in Sections 2 & 3. The need for adaptive, summary-driven visual representations aligns with guidance-based multimodal exploration (Sec. 5) and interactive query reformulation in exploratory workflows (Sec. 6). Finally, in a different context, scalability also emerges in aesthetic deep learning models, as discussed in Section 7.

## 5 Querying Multimodal Data

Interacting with data increasingly goes beyond tables and text and spans multimodal corpora such as images, video, audio, and visually rich documents. This shift exposes a fundamental challenge in query formulation: *users cannot craft meaningful queries until they understand what the data contains, yet they cannot understand the data without querying it.* For structured data, this chicken-and-egg problem is manageable with a quick exploratory data analysis (EDA) step or skimming [107] through some fragments of the data.

For multimodal data such as video, this bootstrapping is prohibitively expensive. Skimming a video archive is very different from skimming a table – watching even 1% of a 10,000-hour video corpus would take several days, and file-level metadata (timestamps, codec, resolution) offer little semantic guidance. Humans become their own bottleneck when consuming such data, exacerbating the query formulation problem.

Here, recent multimodal LLMs and vision-language models (VLMs) enable a crucial unlock: they can “peek inside” multimodal content and produce high-quality, zero-shot semantic descriptions. Rather than expecting users to arrive with well-formed queries, we can now help users discover what they can ask, guiding [108]

them through candidate query concepts using familiar interaction patterns such as autocompletion [109, 110]. Guidance can be built into complex querying as well, e.g., ZELDA [111] uses VLMs to extract event semantics from surveillance footage, enabling queries like “show me unusual pedestrian behavior”, and SketchQL [112] allows users to sketch what they’re looking for, and VLMs match it against frames. Query synthesis approaches [113] infer compositional video-event queries from limited user interactions, reducing the need for expertise in declarative query languages. These systems shift the interaction model from *query-then-explore* to *explore-with-guidance*, where the system continuously suggests what is discoverable as users navigate. However, making VLM-driven guidance practical for interactive visual analytics exposes challenges across database systems, HCI, and visualization.

**Trust and Feedback.** These unlocks enabled by multimodal LLMs and VLMs come hand-in-hand with hallucination and non-determinism challenges [114]. Guidance provided to the user is inherently brittle and may erode trust when below a quality threshold. For instance, a VLM might confidently tag a video frame as containing “bicycles” when the image actually shows motorcycles, potentially misleading the user’s data exploration. Unlike text where users can quickly read an answer to verify correctness, validating multimodal suggestions is expensive; users cannot watch every suggested clip. Systems must provide fast verification mechanisms to ameliorate such issues, e.g., thumbnails or “highlights reels” [115] rather than requiring full video playback, or risk users abandoning guidance due to high verification costs that rival manual exploration. Beyond verification, maintaining provenance and feedback mechanisms throughout the stack allows users to point out incorrect suggestions for retraining. Alternatively, the system itself can monitor its suggestions, and self-reflect on possible mistakes as a self-improvement step. These concepts are aligned with the interaction, lineage, and feedback components discussed in prior sections.

**Next-Generation Interfaces Amplify the Need for Guidance.** As interfaces move toward augmented and virtual reality (AR/VR) and “in-the-world” querying [116, 117], guidance becomes even more critical. Users wearing AR headsets in warehouses or medical settings cannot type complex queries; they need systems that understand gestural or spoken intent and immediately suggest what’s discoverable. Consider a warehouse worker asking “show me damaged pallets” while looking at inventory. The system must process live video, recognize relevant objects, and suggest sub-clauses (“damaged on left side?” “within last week?”) in real-time, guiding the user to the right questions to ask. Camera-first interaction paradigms [118] and gestural query specification [119] require query abstractions that accept multimodal inputs (sketches, gestures, exemplar images, natural language) and provide real-time feedback (e.g., via AQP [120]) on what’s queryable, or risk disorientation within such highly interactive interfaces. In settings such as augmented reality, guidance also comes with hard latency and performance requirements since the user is interacting within the real-world and mistakes can be safety-critical.

## Main Takeaway [Section 5]

- Multimodal data such as video collections present a query formulation barrier: users cannot query what they don’t know exists, yet they cannot discover what’s queryable without prohibitively expensive manual inspection.
- VLMs unlock this dependency, enabling guidance-driven exploration by proposing zero-shot query suggestions and autocompletion-style querying based on semantic understanding of content. This shifts interaction from *query-then-explore* to *explore-with-guidance*.
- Making guidance practical requires addressing trust while designing for next-generation interfaces (AR/VR, gesture, camera-first) where guidance becomes critical, and representing suggestions as succinct visual cues rather than verbose text to support rapid exploration.

**Representing Guidance.** Making guidance effective for multimodal exploration requires addressing how to shape VLM outputs for human consumption under tight interaction budgets. The challenge is not just inference speed, but usability: textual descriptions of visual content are verbose and cognitively difficult to skim quickly. A VLM might return “a person wearing a red jacket riding a bicycle on a paved path near trees during daytime” for a single frame. While this information is accurate, it is difficult for a human to process when scanning hundreds of suggestions. Effective guidance needs *succinct, readable visual cues*, possibly in the form of icons, thumbnails, or even emoji-like glyphs [109] that let users quickly assess what a suggestion means and whether it matches their intent. This also raises fundamental questions about *guidance granularity*: Should systems propose coarse concepts (e.g., “sports,”) that characterize the overall content of the data, but risk being too repetitive across queries, or fine-grained evidence (*a specific thumbnail showing a bicycle at timestamp 1:23*) that points to concrete queries? What level of specificity best helps users formulate their first query without overwhelming them with detail?

Addressing these challenges requires cross-disciplinary collaboration, bridging concepts in query optimization, representation design, mixed-initiative interaction, and progressive visualization to build systems that reliably guide users through the query formulation barrier in multimodal analytics.

**Connections to Other Sections.** At the systems level, the real-time guidance requirements described here are in line with the performance, co-design principles and new exploration paradigm (e.g., cold-started exploration) outlined in Sections 2, 3 and 4. The guidance-driven interaction model discussed is closely related to mixed-initiative analysis systems described in Section 3, as well as to query reformulation and exploration techniques in Section 6. Additionally, the need for compact, cognitively efficient visual cues resonates with the adaptive visualization strategies discussed in Section 4 and relates to the discussion in Section 7 on aesthetic principles and the generation of narrative-rich visual representations to guide user attention. Finally, the broader need for human oversight discussed in Section 6.

## 6 Interactive Data Exploration

Interactive data exploration plays a pivotal role in discovery-driven applications, where the process of data discovery is ad hoc and highly interactive. In these applications, users write multiple queries with varying criteria, aiming to strike a balance between gathering all relevant information and avoiding an overwhelming volume of data. The data management community has closely investigated the need to help users reformulate their queries during data exploration so that the retrieved sets are *not too large* (many answer problems) [121], *not too small* (empty answer problem) [122]; or are *not off-target and diverse* (ensuring relevance and diversity) [123, 124].

When data is structured, existing work has leveraged metadata or attributes to understand the user’s intent and seek their preferences accordingly. Algorithms were designed to navigate users to serendipitous, relevant, diverse representatives of data or to help user find the entity of interest while minimizing effort, e.g., minimize #clicks users provided. These algorithms aided users to reformulate their original queries so that users could interactively explore related data before landing on a small sample of their final interests.

AI has led to an unprecedented explosion in the volume and complexity of unstructured. With the rise of unstructured data, ranging from text and images to video and audio, the need for interactive ad hoc data exploration has never been greater. Users frequently engage in query reformulation and seek relevant, diverse [125], and sometimes serendipitous information. However, the lack of structured metadata, which traditionally guided exploration, presents a significant challenge. As AI continues to reshape the landscape of data exploration, the following question therefore arises: “*Can we still rely on traditional data exploration tools to navigate and uncover insights from the increasingly complex data?*”.

The classical data exploration techniques can still make a headway in this new landscape, provided there is a seamless integration between the structured and unstructured data worlds. Identifying structures from unstructured data requires a combination of advanced techniques in NLP, machine learning, image and speech recognition, and data integration [126]. LLMs are particularly useful, too, as they excel at understanding context, extracting meaning, and processing data of different modalities including audio and video. *Although LLMs have demonstrated impressive capabilities in processing unstructured data, several key research challenges still prevail.* Non-determinism of LLMs is one such key challenge refers to the inherent unpredictability or variability in the output generated by these models, even when the same input is provided multiple times. This introduces a level of unpredictability, requiring the data exploration algorithms to be reconsidered through a *probabilistic lens* [127]. Other challenges include the high cost of invoking these models [128], their potential biases [129], limited contextual understanding, inability to adapt to highly specialized domains, and significant computational resource demands. To overcome these challenges, it is crucial to complement LLM-driven processes with expert human oversight, careful data curation, and transparency in model training and usage. Overall, this is an exciting time for data researchers to collaborate, bringing together expertise across fields to create more nuanced and adaptable solutions for the evolving challenges of data exploration.

**Connections to Other Sections.** The exploration challenges described here instantiate many of the issues raised earlier, particularly multimodal guidance and query formulation barriers from Section 5. Effective exploration further depends on the low-latency interaction and scalable infrastructures discussed in Sections 2, 3 and 4.

### Main Takeaway [Section 6]

- The rise of large-scale unstructured and multimodal data challenges traditional exploration methods mainly due to the lack of metadata, increasing the need for interactive, ad-hoc exploration and continuous query reformulation.
- LLMs are powerful tools for extracting structure from unstructured data; however, they introduce uncertainty due to non-determinism, biases, cost, and limited domain adaptation and contextual understanding, all of which must be addressed.
- Addressing LLMs drawbacks requires combining LLM-driven processes with expert human oversight, careful data curation, and transparent model training and usage.

## 7 Aesthetic Principles and User Guidance in Generative Data Visualization

Information visualization enables the graphical representation of data and information that helps users find patterns, trends, and relationships within the data, and to communicate complex data sets clearly and effectively. In the field of network visualization, traditionally [130–132], visualization of undirected graphs relies on heuristics, especially those based on modeling graphs as physical systems, with the belief that minimizing the energy of such systems leads to aesthetically pleasing layouts that help illustrate the underlying unstructured data. However, it has not been proven that human aesthetic preference can be well modeled by such physical systems. With the advent of DL, there is a growing interest in taking advantage of the power of neural networks to help visualize network data, e.g., [133, 134]. In particular, it is now possible to use DL to expand the horizon of traditional graph visualization in a number of directions.

It was demonstrated that a Graph Neural Network (GNN) model can be used to optimize arbitrary differentiable objective functions. Once trained, such a model can be applied to arbitrary graphs never seen in the training data and create visualizations that optimize the objective function even better than traditional benchmark algorithms [133, 135–137]. Recently, it was shown that by using the Generative Adversarial Network (GAN), we can even train neural networks to optimize non-smooth objective functions [138]. In fact, this approach does not require access to the objective function at all and only requires a comparative function that can choose between two visualizations the “better” one (based on a hidden objective function unknown to GAN).

This opens avenues for future investigation: *instead of optimizing arbitrary objectives such as edge crossing or stress, we should be able to model the human sense of aesthetics directly.* For example, if reducing edge crossing is most important for humans, Tiezzi et al. [136]

demonstrated that a fully connected network can classify whether two edges cross; on the other hand, Wang et al. [138] have shown that the discriminator in a GAN setup can be trained to choose drawings with a lower number of edge crossings, learning from pairs of bad and good drawing examples. The next step is to train a model on a collection of human-labeled example visualizations to learn human visual preference. We believe such a system is possible, but the challenges include scaling GNN, GAN, or other DL-based models to very large unstructured data. Some progress in scalability has already been made [139–141]. More work is needed to collect large amounts of training data, make network visualization adapt to human preferences, incorporate human-specified constraints, and run even faster than traditional network layout algorithms.

A remaining bottleneck is the availability of large-scale, high-quality human preference data. While learning aesthetics directly from human judgments is conceptually appealing, collecting such labels is expensive and slow, which has so far limited empirical studies to relatively small datasets or to machine-generated proxy labels [138]. This raises a fundamental question: can we approximate human aesthetic judgment at scale without relying exclusively on human annotators? Recent advances in large language models (LLMs), vision-language models, and vision foundation models suggest a promising alternative, as these models have demonstrated strong capabilities in visual understanding and comparative reasoning. If properly aligned with human preferences, they could serve as scalable proxies for human labelers and substantially reduce the cost of data collection.

Beyond aesthetics and human preference alone, an equally compelling possibility is to use AI to learn which visualizations actually help people understand the data and perform visual tasks more effectively. Rather than asking only which layout looks better, AI models could be trained to predict which visualizations enable humans to perform such visual analytic tasks (e.g., finding shortest paths or nearest neighbors) more accurately and efficiently. By incorporating task performance signals (e.g., correctness, response time) alongside preference judgments, AI-based discriminators could learn a richer notion of visualization quality that couples perceptual clarity with cognitive effectiveness, allowing future systems to automatically generate visualizations that adapt to both the user’s goals and their perceptual strengths.

Another future direction involves *not just producing the visualization itself but also crafting a narrative, or even combining it with animations/videos* that elucidate the patterns depicted within the visualization, thus directing the user’s attention to the most crucial aspects of generated visualization.

Accomplishing this task will *necessitate leveraging multimodal LLMs and VLMs trained extensively on a diverse corpus consisting of source data (e.g., CSV files, unstructured data), their visualizations, and their corresponding captions* extracted from past visualization and other scientific literature. If we can solve the above challenges, we can then apply such a model to any data to be analyzed, and automatically generate sample visualizations, animations, and narratives that are not only aesthetically pleasing to look at, but also informative, with a well-told story about the visualization.

**Connections to Other Sections.** Generative and learning-based visualization techniques build directly on the human-centered design principles emphasized in Sections 2 and 5. Their scalability limitations mirror the broader visualization scalability concerns discussed in Section 4. Finally, the focus on narrative, animation, and attention-guiding aesthetics complements the guidance-driven interaction and exploratory workflows discussed in Sections 5 and 6.

### Main Takeaway [Section 7]

- Deep learning techniques can significantly advance graph visualization by optimizing layouts and learning human aesthetic preferences beyond traditional metrics such as edge crossings.
- Training aesthetic deep learning models introduces major scalability challenges, when large-scale training datasets are involved. Additionally, collecting sufficient human preference data remains a significant challenge.
- Visualization systems should generate not only layouts but also narratives and animations that highlight crucial data patterns and guide user attention. Achieving this requires multimodal models trained on datasets that include both visualizations and explanatory context, such as captions.

## 8 Discussion & Lessons Learned

A central takeaway of this paper is that AI-driven ecosystems represent a fundamental inflection point in the design of interactive human-centered data systems. Moving forward, interactive systems can no longer be built by independently optimizing data management, AI, and interaction; instead, they require rethinking efficiency and scalability, redefining human-machine collaboration, and embedding cognitive and design principles into every layer of the human-data interaction stack.

In this emerging paradigm, performance, human reasoning, and machine intelligence must be treated as highly interconnected concepts that evolve and adapt together. As AI increasingly mediates how data is explored, interpreted, and validated, systems must support continuous human involvement through interaction, feedback, and verification. This shift places new demands on interfaces, visualization, and system architectures, and challenges long-standing assumptions about how interactive human-centered systems are built and evaluated. Among these components, visualization plays a critical role by evolving from a passive presentation mechanism into an active, adaptive element that supports exploration, sensemaking, and trust in AI-assisted analysis.

The following paragraphs explore the lessons learned in the context of *systemic integration, performance and scalability, data exploration and interaction paradigms, next-generation interfaces, AI oversight, generative visualization, and cross-disciplinary collaboration*.

**New Fundamental, Tightly Coupled Requirements for Human-Data Interaction Systems.** Throughout the paper, several fundamental requirements are identified for the design of

effective human–data interaction systems. In particular, modern systems must simultaneously align real-time, cognition-aware performance with scalability under interactive workloads, provide robust support for large-scale, heterogeneous, multimodal, and predominantly unstructured data, and enable emerging exploration scenarios and interaction paradigms that are guidance-driven and mixed-initiative. In addition, systems must incorporate sophisticated visual analysis that support sensemaking, verification, and explanation, while offering continuous user guidance, progressive feedback, and mechanisms for producing transparent and trustworthy results in the presence of AI uncertainty.

**Holistic, End-to-End System Design is Essential.** The paper makes it clear that isolated improvements, whether in query engines, interfaces, visualization, or AI models, are not enough. The demands of the emerging AI landscape require co-design across storage structures, indexing methods, interface models, query processing techniques, perceptual workflows, and AI-assisted reasoning. *These point toward holistic, end-to-end system architectures, where interfaces shape system requirements just as much as systems shape interface capabilities.*

✂️ **Synthesis Across Sections.** This dependency is highlighted throughout the paper, from interface-driven workloads in Section 2 to millisecond-level interaction constraints (Sec. 3), cold-start visualization (Sec. 4), and real-time multimodal guidance (Sec. 5).

**Latency and Scalability are Cross-cutting Constraints on Human–Data Interactive Systems.** Achieving perceptually aligned response times and scalability remains a fundamental technical and design challenge. Multimodal data, massive and heterogeneous datasets, high-velocity streams, next-generation interface requirements, and AI models introduce critical data-oriented challenges that must be addressed to enable the development of efficient and scalable human-centered interactive systems. In this context, *the involved communities face multiple challenges, requiring a re-consideration of core problems and the handling of numerous new ones.*

✂️ **Synthesis Across Sections.** The importance of latency and scalability is highlighted across multiple sections. Sections 3 and 4 show that latency and scalability function as cognitive constraints as well as technical ones, with effects that propagate to multimodal analytics (Sec. 5), new-generation interfaces (Sec. 5), interactive exploration over unstructured data (Sec. 6), and training models in large-scale datasets (Sec. 7).

**Emerging Interaction and Exploration Paradigms Drive Systems Development.** Multimodal and unstructured data fundamentally reshape how users interact with data, exposing the limitations of traditional query–then–explore workflows. *When data is opaque, large-scale, or costly to inspect, users struggle to articulate precise queries, prompting a shift toward guidance-driven and mixed-initiative interaction.* Furthermore, the importance of reducing data-to-analysis time reinforces the development of cold-start exploration systems (a.k.a. in-situ exploration), enabling on-the-fly exploration over unseen datasets. Enabling such interaction and

exploration paradigms places new demands on data management systems, which must support multimodal data as first-class citizens, provide fast semantic access, and deliver progressive, low-latency feedback. At the same time, *guidance introduces challenges around trust, verification effort, and cognitive load, underscoring the need for carefully designed interfaces and systems that preserve meaningful human oversight throughout the exploration process.*

✂️ **Synthesis Across Sections.** The necessity of guidance-driven, mixed-initiative interaction when data cannot be easily inspected is discussed in Sections 5 and 6, but these interaction paradigms introduce new challenges in data management, trust, verification effort, and cognitive load. Finally, Section 4 highlights the importance of enabling on-the-fly exploration over unseen datasets.

**Next-Generation Interfaces are Crucial for Unlocking New Capabilities.** Interfaces must move beyond traditional text and graphical paradigms toward multimodal, action-oriented environments that integrate natural language, rich visual capabilities, gestures, and augmented or virtual reality. *These new modalities support more intuitive and immersive human–data engagement, but they also introduce new challenges related to latency, cognitive load, interaction paradigms, real-time guidance, adaptive feedback, and interface design.* The aforementioned challenges further highlight the strong dependencies between backend systems and frontend design, as well as the necessity for new exploration paradigms discussed earlier.

✂️ **Synthesis Across Sections.** The shift to next-generation interfaces is grounded in the co-design principles of Section 2 and the perceptually aligned interaction requirements of Section 3, and are further developed in the contexts of multimodal and unstructured data exploration, as well as AR/VR querying in Sections 5 and 6.

**AI Reinforces the Necessity of Human-in-the-Loop Oversight and Transparency.** Foundation models, such as LLMs and VLMs offer powerful capabilities for guidance-driven exploration, data and semantic understanding, zero-shot multimodal analysis. They also assist with data-oriented tasks, such as data integration and discovery, while improving interface usability. However, *their brittleness, prompt sensitivity, and risk of hallucination underscore the need for human involvement and supervision to ensure results that are trustworthy, meaningful, and actionable.* This positions AI models not as autonomous problem solvers but as components within interactive systems that enhance usability and support explanation, correction, and accountability. These challenges motivate explainable and transparent AI mechanisms that support verification, interpretation, and trust.

✂️ **Synthesis Across Sections.** The need for human involvement is discussed across Sections 5 and 6, motivating interactive feedback, provenance, and verification mechanisms that keep humans meaningfully in control.

**From Static to Generative and Adaptive Visualization.** *Visualization has evolved from a passive, static output into an active,*

*generative component of analytical systems.* By incorporating progressive, narrative, and aesthetics-aware approaches, visualizations guide user attention, enhance interpretability, and foster trust in AI-driven, multimodal analytics. Deep learning-based aesthetic modeling, narrative construction, and automated visualization generation create new opportunities for expressiveness. At the same time, these methods rely heavily on training data quality, availability of human preferences, and scalable layout models. *The broader lesson is that visualization must be explanatory, interpretable, human-guided, and integrated into interactive analytic workflows rather than optimized solely for visual appeal.*

✂ **Synthesis Across Sections.** Visualization evolution is discussed in Section 7, framing it as a generative and adaptive component, and is reflected in guidance-driven multimodal analytics (Sec. 5) and interactive exploration over unstructured data (Sec. 6), where visual representations support explanation, sensemaking, and trust.

**The Necessity of Sustained Cross-Disciplinary Collaboration.** The challenges discussed throughout the paper span system architecture and performance, data management, AI modeling, visualization, interaction design, and human cognition, and cannot be addressed in isolation. Progress therefore critically depends on long-term coordination among the systems, databases, AI, information visualization, HCI, computer graphics, and cognitive science communities, rather than on isolated advances within individual fields.

## References

- [1] B. Shneiderman, *Human-Centered AI*, Oxford University Press, 2022.
- [2] M. Raees, I. Meijerink, I. Lykourantzou, V. Khan, K. Papangelis, From Explainable to Interactive AI: a Literature Review on Current Trends in Human-AI Interaction, *Int. J. Hum. Comput. Stud.* (2024).
- [3] H. Li, Y. Wang, H. Qu, Where Are We So Far? Understanding Data Storytelling Tools from the Perspective of Human-AI Collaboration, in: *ACM CHI*, 2024.
- [4] S. Amershi, D. S. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. T. Iqbal, P. N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, E. Horvitz, *Guidelines for Human-AI Interaction*, in: *ACM CHI*, 2019.
- [5] E. Wu, Y. Chen, H. Mohammed, Z. Huang, A Decade of Systems for Human Data Interaction (2025). *arXiv*: 2511.15585.
- [6] S. Amer-Yahia, L. Battle, Y. Hu, D. Moritz, A. Parameswaran, N. Bikakis, P. K. Chrysanthos, G. Li, G. Papastefanatos, L. Yu, Data Exploration and Visual Analytics Challenges in AI Era, *ACM SIGMOD Blog* (Jan. 2025). URL <https://wp.sigmod.org/?p=3820>
- [7] N. Bikakis, P. K. Chrysanthos, G. Li, G. Papastefanatos, L. Yu, Visual Analytics Challenges and Trends in the Age of AI: The BigVis Community Perspective, *ACM SIGMOD Record* (2025).
- [8] L. Battle, C. Scheidegger, A Structured Review of Data Management Technology for Interactive Visualization and Analysis, *IEEE TVCG* 27 (2) (2021).
- [9] G. Richer, A. Pister, M. Abdelaal, J. Fekete, M. Sedlmair, D. Weiskopf, Scalability in Visualization, *IEEE TVCG* 30 (7) (2024).
- [10] A. Ulmer, M. Angelini, J. Fekete, J. Kohlhammer, T. May, A Survey on Progressive Visualization, *IEEE TVCG* 30 (9) (2024).
- [11] X. Qin, Y. Luo, N. Tang, G. Li, Making Data Visualization More Efficient and Effective: a Survey, *VLDBJ* 29 (1) (2020).
- [12] E. Wu, F. Psallidas, Z. Miao, H. Zhang, L. Rettig, Combining Design and Performance in a Data Visualization Management System, in: *CIDR*, 2017.
- [13] N. Bikakis, *Big Data Visualization Tools*, in: *Encyclopedia of Big Data Technologies*, 2nd Ed., Springer, 2022.
- [14] G. Andrienko, N. Andrienko, S. Drucker, J.-D. Fekete, D. Fisher, S. Idreos, T. Kraska, G. Li, K.-L. Ma, J. D. Mackinlay, A. Oulasvirta, T. Schreck, H. Schmann, M. Stonebraker, D. Auber, N. Bikakis, P. K. Chrysanthos, G. Papastefanatos, M. Sharaf, *Big Data Visualization and Analytics: Future Research Challenges and Emerging Applications*, in: *Workshop on Big Data Visual Exploration and Analytics (BigVis 2020)*, 2020.
- [15] L. Jiang, P. Rahman, A. Nandi, Evaluating Interactive Data Systems: Workloads, Metrics, and Guidelines, in: *ACM SIGMOD*, 2018.
- [16] L. Battle, P. Eichmann, M. Angelini, T. Catarci, G. Santucci, Y. Zheng, C. Binnig, J. Fekete, D. Moritz, Database Benchmarking for Supporting Real-Time Interactive Querying of Large Data, in: *ACM SIGMOD*, 2020.
- [17] P. Eichmann, E. Zraggen, C. Binnig, T. Kraska, IBEDench: A Benchmark for Interactive Data Exploration, in: *ACM SIGMOD*, 2020.
- [18] A. Wu, D. Deng, M. Chen, S. Liu, D. A. Keim, R. Maciejewski, S. Miksch, H. Strobel, F. B. Viégas, M. Wattenberg, *Grand Challenges in Visual Analytics Applications*, *CG&A* (2023).
- [19] Y. Lian, J. Hao, W. Zeng, Q. Luo, A Survey of Visual Insight Mining: Connecting Data and Insights via Visualization, *Visual Informatics* (2025).
- [20] J. Walny, C. Frisson, M. West, D. Kosminsky, S. Knudsen, S. Carpendale, W. Willett, Data Changes Everything: Challenges and Opportunities in Data Visualization Design Handoff, *IEEE TVCG* 26 (1) (2020).
- [21] E. Adamakis, G. Margetis, S. Ntoa, C. Stephanidis, A Methodological Approach Towards Human-Centered Visual Analytics, *Visual Informatics* (2025).
- [22] R. C. Basole, T. Major, Generative AI for Visualization: Opportunities and Challenges, *IEEE TVCG* 44 (2) (2024).
- [23] J. Wang, S. Liu, W. Zhang, Visual Analytics for Machine Learning: A Data Perspective Survey, *IEEE TVCG* 30 (12) (2024).
- [24] A. Wu, Y. Wang, X. Shu, D. Moritz, W. Cui, H. Zhang, D. Zhang, H. Qu, AI4VIS: Survey on Artificial Intelligence Approaches for Data Visualization, *IEEE TVCG* 28 (12) (2022).
- [25] Q. Wang, Z. Chen, Y. Wang, H. Qu, A Survey on ML4VIS: Applying Machine Learning Advances to Data Visualization, *IEEE TVCG* 28 (12) (2022).
- [26] N. V. Andrienko, G. L. Andrienko, L. Adilova, S. Wrobel, T. Rhyne, *Visual Analytics for Human-Centered Machine Learning*, *CG&A* (2022).
- [27] Y. Ye, J. Hao, Y. Hou, Z. Wang, S. Xiao, Y. Luo, W. Zeng, Generative AI for Visualization: State of the Art and Future Directions, *Visual Informatics* 8 (1) (2024).
- [28] W. Yang, M. Liu, Z. Wang, S. Liu, Foundation Models Meet Visualizations: Challenges and Opportunities, *Comput. Vis. Media* 10 (3) (2024).
- [29] N. Elmqvist, C. N. Klokmoose, Automating the Path: An R&D Agenda for Human-Centered AI and Visualization, *IEEE TVCG* 45 (3) (2025).
- [30] J. Yuan, C. Chen, W. Yang, M. Liu, J. Xia, S. Liu, A Survey of Visual Analytics Techniques for Machine Learning, *Comput. Vis. Media* 7 (1) (2021).
- [31] F. Hohman, M. Kahng, R. S. Pienta, D. H. Chau, Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers, *IEEE TVCG* 25 (8) (2019).
- [32] N. Chen, Y. Zhang, J. Xu, K. Ren, Y. Yang, VisEval: A Benchmark for Data Visualization in the Era of Large Language Models, *IEEE TVCG* 31 (1) (2025).
- [33] V. Raman, J. M. Hellerstein, Potter's Wheel: An Interactive Data Cleaning System, in: *VLDB*, 2001.
- [34] D. Alabi, E. Wu, PFunk-H: Approximate Query Processing Using Perceptual Models, in: *HILDA*, 2016.
- [35] F. Li, B. Wu, K. Yi, Z. Zhao, W. H. Hand, Online Aggregation via Random Walks, *ACM TODS* 44 (1) (2019).
- [36] B. Ding, S. Huang, S. Chaudhuri, K. Chakrabarti, C. Wang, Sample + Seek: Approximating Aggregates with Distribution Precision Guarantee, in: *ACM SIGMOD*, 2016.
- [37] T. Kraska, Northstar: An Interactive Data Science System, *PVLDB* 11 (12) (2018).
- [38] J. M. Hellerstein, P. J. Haas, H. J. Wang, Online Aggregation, in: *ACM SIGMOD*, 1997.
- [39] D. Moritz, D. Fisher, B. Ding, C. Wang, Trust, but Verify: Optimistic Visualizations of Approximate Queries for Exploring Big Data, in: *ACM CHI*, 2017.
- [40] A. Kim, E. Blais, A. G. Parameswaran, P. Indyk, S. Madden, R. Rubinfeld, Rapid Sampling for Visualizations with Ordering Guarantees, *PVLDB* (2015).
- [41] S. Rahman, M. Aliakbarpour, H. Kong, E. Blais, K. Karahalios, A. G. Parameswaran, R. Rubinfeld, I've Seen "Enough": Incrementally Improving Visualizations to Support Rapid Decision Making, *PVLDB* (2017).
- [42] Y. Park, M. J. Cafarella, B. Mozafari, Visualization-aware Sampling for Very Large Databases, in: *IEEE ICDE*, 2016.
- [43] S. Maroulis, N. Bikakis, V. Stamatopoulos, G. Papastefanatos, Adaptive Indexing for Approximate Query Processing in Interactive Data Analysis, *arXiv:2505-19872* (2025).
- [44] Y. Wang, Y. Wang, X. Chen, Y. Zhao, F. Zhang, E. Wu, C. Fu, X. Yu, OM3: An Ordered Multi-level Min-Max Representation for Interactive Progressive

- Visualization of Time Series, *ACM SIGMOD* 1 (2) (2023).
- [45] U. Jugel, Z. Jerzak, G. Hackenbroich, V. Markl, M4: A Visualization-Oriented Time Series Data Aggregation, *PVLDB* 7 (10) (2014).
- [46] U. Jugel, Z. Jerzak, G. Hackenbroich, V. Markl, Faster visual analytics through pixel-perfect aggregation, *PVLDB* 7 (13) (2014).
- [47] P. R. Doshi, E. A. Rundensteiner, M. O. Ward, Prefetching for Visual Data Exploration, in: *DASFAA*, 2003.
- [48] B. Shneiderman, Extreme Visualization: Squeezing a Billion Records into a Million Pixels, in: *ACM SIGMOD*, 2008.
- [49] L. Battle, R. Chang, M. Stonebraker, Dynamic Prefetching of Data Tiles for Interactive Visualization, in: *ACM SIGMOD*, 2016.
- [50] H. Mohammed, Z. Wei, R. Netravali, E. Wu, Continuous Prefetch for Interactive Data Applications, *PVLDB* 13 (11) (2020).
- [51] F. Tauheed, T. Heinis, F. Schürmann, H. Markram, A. Ailamaki, SCOUT: Prefetching for Latent Feature Following Queries, *PVLDB* 5 (11) (2012).
- [52] M. A. Khan, A. Nandi, Flux Capacitors for JavaScript Deloreans: Approximate Caching for Physics-Based Data Interaction, in: *ACM IUI*, 2019.
- [53] C. A. de Lara Pahins, S. A. Stephens, C. Scheidegger, J. L. D. Comba, Hashedcubes: Simple, Low Memory, Real-Time Visual Exploration of Big Data, *IEEE TVCG* 23 (1) (2017).
- [54] S. Agarwal, B. Mozafari, A. Panda, H. Milner, S. Madden, I. Stoica, Blinkdb: Queries with Bounded Errors and Bounded Response Times on Very Large Data, in: *EuroSys*, 2013.
- [55] Z. Liu, B. Jiang, J. Heer, *imMens*: Real-Time Visual Querying of Big Data, *Comput. Graph. Forum* 32 (3) (2013).
- [56] W. Tao, X. Liu, Y. Wang, L. Battle, C. Demiralp, R. Chang, M. Stonebraker, Kyrix: Interactive Pan/Zoom Visualizations at Scale, *Comput. Graph. Forum* 38 (3) (2019).
- [57] D. Moritz, B. Howe, J. Heer, Falcon: Balancing Interactive Latency and Resolution Sensitivity for Scalable Linked Visualizations, in: *ACM CHI*, 2019.
- [58] S. Maroulis, N. Bikakis, G. Papastefanatos, P. Vassiliadis, Y. Vassiliou, Resource-Aware Adaptive Indexing for In-situ Visual Exploration and Analytics, *VLDBJ* (2022).
- [59] Y. Wu, R. Chang, J. M. Hellerstein, A. Satyanarayan, E. Wu, DIEL: Interactive Visualization Beyond the Here and Now, *IEEE TVCG* (2019).
- [60] Y. Wu, R. Chang, J. M. Hellerstein, E. Wu, Facilitating Exploration with Interaction Snapshots under High Latency, *IEEE VIS* (2018).
- [61] A. D. Fekete, I. Gupta, A. G. Parameswaran, F. Chen, C. D. Leon, T. Wattanawaroon, J. Yun, S. Seshadri, Z. Yi, C. Lien, R. Sun, R. Durrani, J. E. Gonzalez, A. D. Joseph, Transactional Panorama: A Conceptual Framework for User Perception in Analytical Visual Interfaces (2023). [arXiv: 2302.05476](https://arxiv.org/abs/2302.05476).
- [62] Y. Chen, X. Li, J. Tao, L. Ramjit, R. Netravali, S. Mitra, A. Parameswaran, J. Ghaderi, D. Rubenstein, E. Wu, Physical Visualization Design: Decoupling Interface and System Design, *ACM SIGMOD* (2025).
- [63] Z. Huang, P. K. Damalapati, E. Wu, Aggregation Consistency Errors in Semantic Layers and How to Avoid Them, in: *HILDA*, 2023.
- [64] J. Hyde, J. Fremlin, Measures in SQL, *ACM SIGMOD* (2024).
- [65] Z. Huang, E. Wu, Lightweight Materialization for Fast Dashboards Over Joins, *ACM SIGMOD* 1 (2023).
- [66] F. Psallidas, E. Wu, Provenance for Interactive Visualizations, in: *HILDA*, 2018.
- [67] F. Psallidas, E. Wu, Demonstration of Smoke: A Deep Breath of Data-Intensive Lineage Applications, in: *ACM SIGMOD*, 2018.
- [68] F. Psallidas, E. Wu, Smoke: Fine-Grained Lineage at Interactive Speed, *PVLDB* 11 (2018).
- [69] H. Mohammed, C. Summers, S. Kaushik, E. Wu, SmokedDuck Demonstration: SQLStepper, *ACM SIGMOD* (2023).
- [70] S. Gathani, P. Lim, L. Battle, Debugging Database Queries: A Survey of Tools, Techniques, and Users, in: *ACM CHI*, 2020.
- [71] S. Roy, L. Orr, D. Suci, Explaining Query Answers with Explanation-Ready Databases, *PVLDB* 9 (4) (2015).
- [72] Z. Huang, E. Wu, Reptile: Aggregation-Level Explanations for Hierarchical Data, in: *ACM SIGMOD*, 2022.
- [73] W. Wu, L. Flokas, E. Wu, J. Wang, Complaint-Driven Training Data Debugging for Query 2.0, in: *ACM SIGMOD*, 2020.
- [74] E. Wu, S. Madden, Scorpion: Explaining Away Outliers in Aggregate Queries, *PVLDB* (2013).
- [75] F. Abuzaid, P. Kraft, S. Suri, E. Gan, E. Xu, A. Shenoy, A. Ananthanarayan, J. Sheu, E. Meijer, X. Wu, et al., Diff: a Relational Interface for Large-Scale Data Explanation, *PVLDB* 12 (4) (2018).
- [76] H. Mohammed, A. Yao, L. Flokas, C. Summers, H. Zhong, G. Chan, S. Mitra, E. Wu, FaDE: More Than a Million What-ifs per Second, *PVLDB* (2025).
- [77] S. Sarawagi, G. Sathe, i3: Intelligent, Interactive Investigation of Olap Data Cubes, *SIGMOD Record* (2000).
- [78] E. Wu, View Composition Algebra for Ad Hoc Comparisons, *IEEE TVCG* (2022).
- [79] E. Wu, Extending the View Composition Algebra to Hierarchical Data (2022). [arXiv: 2205.01283](https://arxiv.org/abs/2205.01283).
- [80] Y. Chen, E. Wu, PI2: End-to-End Interactive Visualization Interface Generation from Queries, in: *ACM SIGMOD*, 2021.
- [81] H. Zhang, T. Sellam, E. Wu, Mining Precision Interfaces From Query Logs (2017). [arXiv: 1904.02344](https://arxiv.org/abs/1904.02344).
- [82] Y. Chen, J. Tao, E. Wu, DIG: The Data Interface Grammar, in: *HILDA*, 2023.
- [83] E. Wu, R. Chang, Design-Specific Transformations in Visualization, *IEEE VIS BELIV Workshop* (2024).
- [84] E. Wu, X. Y. Tuang, A. Li, V. Bainwala, A Formalism and Library for Database Visualization (2025). [arXiv: 2504.08979](https://arxiv.org/abs/2504.08979).
- [85] E. Wu, Database Theory in Action: Database Visualization, in: *ICDT*, 2025.
- [86] E. F. Codd, A Relational Model of Data for Large Shared Data Banks, *Commun. ACM* 13 (1970).
- [87] D. Zha, Z. P. Bhat, K. Lai, F. Yang, Z. Jiang, S. Zhong, X. Hu, Data-Centric Artificial Intelligence: A Survey, *CoRR abs/2303.10158* (2023).
- [88] S. E. Whang, Y. Roh, H. Song, J. Lee, Data Collection and Quality Challenges in Deep Learning: a Data-Centric AI Perspective, *VLDBJ* 32 (4) (2023).
- [89] S. Robertson, Z. J. Wang, D. Moritz, M. B. Kery, F. Hohman, Angler: Helping Machine Translation Practitioners Prioritize Model Improvements, in: *ACM CHI*, 2023.
- [90] Á. A. Cabrera, E. Fu, D. Bertucci, K. Holstein, A. Talwalkar, J. I. Hong, A. Perer, Zeno: An Interactive Framework for Behavioral Evaluation of Machine Learning, in: *ACM CHI*, 2023.
- [91] H. A. Simon, Designing Organizations for an Information Rich World, in: *Computers, Communications, and the Public Interest*, Johns Hopkins Press, 1971.
- [92] J. W. Tukey, M. B. Wilk, Data Analysis and Statistics: an Expository Overview, in: *American Federation of Information Processing Societies (AFIPS)*, 1966.
- [93] J. E. Allen, C. I. G. and Eric Horvitz and, Mixed-Initiative Interaction, *IEEE Intelligent Systems and their Applications* 14 (5) (1999).
- [94] K. Wongsuphasawat, D. Moritz, A. Anand, J. D. Mackinlay, B. Howe, J. Heer, Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations, *IEEE TVCG* 22 (1) (2016).
- [95] W. Epperson, V. Gorantla, D. Moritz, A. Perer, Dead or Alive: Continuous Data Profiling for Interactive Data Science, *IEEE TVCG* 30 (1) (2024).
- [96] Z. Liu, J. Heer, The Effects of Interactive Latency on Exploratory Visual Analysis, *IEEE TVCG* 20 (12) (2014).
- [97] J. Heer, D. Moritz, Mosaic: An Architecture for Scalable & Interoperable Data Views, *IEEE TVCG* 30 (1) (2024).
- [98] M. N. Hoque, S. Shin, N. Elmqvist, Visualization for Human-Centered AI Tools (2024). [arXiv: 2404.02147](https://arxiv.org/abs/2404.02147).
- [99] C. D. Stolper, A. Perer, D. Gotz, Progressive Visual Analytics: User-Driven Visual Exploration of in-Progress Analytics, *IEEE TVCG* 20 (12) (2014).
- [100] J.-D. Fekete, D. Fisher, M. Sedlmair, Progressive Data Analysis: Roadmap and Research Agenda, *Eurographics*, 2024.
- [101] S. Castelo, R. Rampin, A. S. R. Santos, A. Bessa, F. Chirigati, J. Freire, Auctus: A Dataset Search Engine for Data Discovery and Augmentation, *PVLDB* 14 (12) (2021).
- [102] N. Bikakis, S. Maroulis, G. Papastefanatos, P. Vassiliadis, In-situ Visual Exploration over Big Raw Data, *Information Systems Journal* 95 (2021).
- [103] I. Alagiannis, R. Borovica, M. Branco, S. Idreos, A. Ailamaki, Nodb: Efficient Query Execution on Raw Data Files, in: *ACM SIGMOD*, 2012.
- [104] H. Wickham, Bin-Summarise-Smooth: a Framework for Visualising Large Data, <https://vita.had.co.nz/papers/bigvis.pdf> (2013).
- [105] L. Wilkinson, *The Grammar of Graphics*, Springer, 2005.
- [106] A. Mayorga, M. Gleicher, Splatplots: Overcoming Overdraw in Scatter Plots, *IEEE TVCG* 19 (9) (2013).
- [107] M. Singh, A. Nandi, H. Jagadish, Skimmer: rapid scrolling of relational query results, in: *ACM SIGMOD*, 2012.
- [108] A. Nandi, H. Jagadish, Guided interaction: Rethinking the query-result paradigm, *PVLDB* 4 (12) (2011).
- [109] H. Yoo, A. Nandi, Emojis in autocompletion: Enhancing video search with visual cues, in: *HILDA*, 2025.

- [110] H. Yoo, A. Nandi, Guided Querying over Videos Using Autocompletion Suggestions, in: HILDA, 2024.
- [111] F. Romero, C. Winston, J. Hauswald, M. Zaharia, C. Kozyrakis, Zelda: Video analytics using vision-language models (2023). [arXiv:2305.03785](#).
- [112] R. Wu, P. Chunduri, D. J. Shah, A. J. Aravind, A. Payani, X. Chu, J. Arulraj, K. Rong, Sketchql demonstration: Zero-shot video moment querying with sketches (2024). [arXiv:2405.18334](#).
- [113] E. Zhang, M. Daum, D. He, B. Haynes, R. Krishna, M. Balazinska, Equi-vocal: Synthesizing queries for compositional video events from limited user interactions, *PVLDB* 16 (11) (2023).
- [114] Y. Fu, R. Xie, X. Sun, Z. Kang, X. Li, Mitigating hallucination in multimodal large language model via hallucination-targeted direct preference optimization, in: Findings of the Association for Computational Linguistics: ACL 2025, 2025.
- [115] D. Winecki, A. Nandi, V2V: Efficiently Synthesizing Video Results for Video Queries, in: IEEE ICDE, 2024.
- [116] C. Burley, A. Nandi, Arquery: Hallucinating analytics over real-world data using augmented reality, in: CIDR, 2019.
- [117] M. Khan, A. Nandi, Dreamstore: A Data Platform for Enabling Shared Augmented Reality, in: IEEE Virtual Reality and 3D User Interfaces (vr), 2021.
- [118] K. Wolf, D. Winecki, A. Nandi, Camera-First Form Filling: Reducing the Friction in Climate Hazard Reporting, in: HILDA, 2023.
- [119] A. Nandi, L. Jiang, M. Mandel, Gestural Query Specification, *PVLDB* 7 (4) (2013).
- [120] S. Jo, I. Trummer, Thalamusdb: Approximate query processing on multi-modal data, *ACM SIGMOD* 2 (3) (2024).
- [121] S. Basu Roy, H. Wang, G. Das, U. Nambiar, M. Mohania, Minimum-effort Driven Dynamic Faceted Search in Structured Databases, in: ACM CIKM, 2008.
- [122] D. Mottin, A. Marascu, S. B. Roy, G. Das, T. Palpanas, Y. Velegrakis, A Probabilistic Optimization Framework for the Empty-Answer Problem, *PVLDB* 6 (14) (2013).
- [123] M. M. Islam, M. Asadi, S. Amer-Yahia, S. B. Roy, A Generic Framework for Efficient Computation of top-k Diverse Results, *VLDBJ* 32 (4) (2023).
- [124] S. Nikookar, M. Esfandiari, R. M. Borromeo, P. Sakharkar, S. Amer-Yahia, S. Basu Roy, Diversifying Recommendations on Sequences of Sets, *VLDBJ* 32 (2) (2023).
- [125] M. R. Vieira, H. L. Razente, M. C. Barioni, M. Hadjieleftheriou, D. Srivastava, C. Traina, V. J. Tsotras, On Query Result Diversification, in: IEEE ICDE, 2011.
- [126] V. Shah, S. Li, A. Kumar, L. Saul, Speakql: Towards Speech-Driven Multimodal Querying of Structured Data, in: ACM SIGMOD, 2020.
- [127] D. Suciu, D. Olteanu, C. Ré, C. Koch, Probabilistic Databases, Springer, 2022.
- [128] S. N. Nia, S. Ghosh, S. B. Roy, S. Amer-Yahia, Personalized Top-k Set Queries Over Predicted Scores (2025). [arXiv:2502.12998](#).
- [129] M. M. Islam, M. Asadi, S. Basu Roy, Equitable top-k Results for Long Tail Data, *ACM SIGMOD* 1 (4) (2023).
- [130] P. Eades, A Heuristic for Graph Drawing, *Congressus Numerantium* 42 (1984).
- [131] T. Kamada, S. Kawai, An Algorithm for Drawing General Undirected Graphs, *Inf. Process. Lett.* 31 (1) (1989).
- [132] E. R. Gansner, Y. Koren, S. C. North, Graph Drawing by Stress Majorization, in: Graph Drawing, International Symposium (GD), Vol. 3383, 2004.
- [133] Y. Wang, Z. Jin, Q. Wang, W. Cui, T. Ma, H. Qu, DeepDrawing: A Deep Learning Approach to Graph Drawing, *IEEE TVCG* 26 (1) (2020).
- [134] O. Kwon, K. Ma, A Deep Generative Model for Graph Layout, *IEEE TVCG* 26 (1) (2020).
- [135] X. Wang, K. Yen, Y. Hu, H. Shen, DeepGD: A Deep Learning Framework for Graph Drawing Using GNN, *IEEE TVCG* 41 (5) (2021).
- [136] M. Tiezzi, G. Ciravegna, M. Gori, Graph Neural Networks for Graph Drawing, *IEEE Trans. Neural Networks Learn. Syst.* 35 (4) (2024).
- [137] L. Giovannangeli, F. Lalanne, D. Auber, R. Giot, R. Bourqui, Toward Efficient Deep Learning for Graph Drawing (DL4GD), *IEEE TVCG* 30 (2) (2024).
- [138] X. Wang, K. Yen, Y. Hu, H. Shen, SmartGD: A GAN-Based Graph Drawing Framework for Diverse Aesthetic Goals, *IEEE TVCG* 30 (8) (2024).
- [139] K. Yan, T. Zhao, M. Yang, Graphuly: graph U-Nets-Based Multi-Level Graph Layout, *IEICE Trans. Inf. Syst.* 105-D (12) (2022).
- [140] F. Grötschla, J. Mathys, R. Veres, R. Wattenhofer, CoRe-gd: A Hierarchical Framework for Scalable Graph Visualization with gnns, in: Intl. Conf. on Learning Representations (ICLR), 2024.
- [141] Z. Huang, S. Zhang, C. Xi, T. Liu, M. Zhou, Scaling Up Graph Neural Networks Via Graph Coarsening, in: ACM KDD, 2021.