




# From ‘Here’ to ‘There’: Exploring Proximity Semantics in Multimodal Data Exploration

Dennis Bromley<sup>1</sup> , Diana Wang<sup>1</sup> , and Vidya Setlur<sup>1</sup>   
Tableau Research

## Abstract

Modern data exploration tools often struggle to capture the subtleties of analytical intent, especially when users seek patterns that are difficult to specify using traditional query methods or natural language alone. We introduce a multimodal research probe for querying time-series and geospatial data that integrates free-form sketching, natural language, and visual annotations within a unified interaction space. Users articulate queries by sketching trends or spatial paths and augmenting them with annotations and analytical directives grounded in shared spatial and temporal context. The system employs a hybrid architecture combining geometric sketch matching and visual language models (VLMs) to support queries that interleave pattern matching and semantic constraints. Through a preliminary study with 20 participants, we observed recurring interaction patterns in which participants used spatial, temporal, and visual proximity to relate sketches, annotations, and language. Rather than treating these as isolated inputs, participants relied on their relative placement to disambiguate meaning. We analyze these behaviors as evidence for proximity semantics (PS), a form of deictic disambiguation in which meaning is shaped by the closeness of multimodal elements within a shared interaction space. We present PS as a conceptual lens grounded in observed user behavior, and discuss its implications for the design of future multimodal data exploration systems.

## CCS Concepts

• **Human-centered computing** → **Interaction paradigms; Interaction techniques; Visualization systems and tools;**

## 1. Introduction

Traditional search and query systems depend on users’ ability to articulate analytical intent that aligns with the data schema and query language [BYRNM\*99, MRS08]. This poses challenges when users reason about patterns that are visually intuitive but difficult to express precisely using language-only or query formulations [LUB\*16]. Advances in natural language interfaces (NLIs) enable conversational interaction with data [SBT\*16, GDA\*15, DMN\*17], but remain largely dependent on verbal formulations and offer limited support for visually grounded reasoning. Sketch-based querying shows that freehand drawing can support the expression of visual patterns [SSD01, MA18]; however, these systems typically focus on shape matching and offer limited support for higher-level analytical constraints or multimodal integration. Motivated by the complementary strengths of sketch-based interaction and natural language, our work explores how users combine these modalities within a shared query articulation space. We implemented a research probe that allows users to sketch trends or spatial paths, annotate regions of interest, and write free-form analytical directives on the same canvas (see Figure 1). The supported annotation vocabulary (e.g., circles, cross-outs, arrows, and labels) was iteratively developed informed by prior sketch-based systems and common diagrammatic practices, and kept lightweight to allow reinterpretation. The system supports sketch input to enable a

mathematically precise, user-controllable geometric search, while a visual language model (VLM) interprets handwritten text, symbols, and diagrammatic annotations to generate structured database queries. In a preliminary user study (N=20), we observed recurring interaction patterns in which participants used spatial arrangement and visual proximity to relate sketches, annotations, and text. For example, they placed labels near regions of interest, used arrows to connect text and data, and marked inclusion and exclusion zones directly on the visualization. Rather than presenting a complete interaction paradigm, we treat these observations as empirical grounding for a conceptual contribution. These patterns led us to define *proximity semantics* (PS): a form of deictic disambiguation in which meaning is shaped by the spatial, temporal, and visual closeness of multimodal elements within a shared interaction space. We adopt PS as an analytic lens to characterize these behaviors and inform the design of integrated multimodal analytical interfaces.

## 2. Related Work

Our work spans three themes: (1) text-based and semantic search systems, (2) visual query systems, and (3) sketch-based interaction.

### 2.1. Text-Based and Semantic Search Systems

Search systems have diversified to cater to a broad spectrum of search intents [CHH\*08, DAC10, FLS\*08]. Recent advances in

LLMs have expanded the expressive power of search and question answering by enabling conversational querying and programmatic query generation [Ope23, Ant25, LLG\*20]. Despite these advances, language-centric approaches remain constrained by verbal formulations and struggle to express analytical intent that is inherently spatial, visual, or relational. Prior work mapping quantitative properties to linguistic descriptors (e.g., “spike,” “bump”) [BS23, BBS24] still relies on predefined vocabularies and offers limited support for visually grounded reference or deictic reasoning. While prior systems support referring expressions (e.g., “this region”), they typically resolve them through explicit identifiers or structured context rather than *spatial or temporal proximity*. As a result, they have limited ability to capture the implicit, spatially grounded references that arise in sketch- and annotation-based interaction.

## 2.2. Visual Query Systems


Visual query systems leverage direct manipulation, graphical metaphors, and visual encodings to lower barriers to data exploration [BCCL91, SKL\*16]. Systems such as Visage [DKR97] and time-box queries [HS04] enable users to specify constraints through spatial selection and interaction, supporting iterative and exploratory workflows. These systems support deictic reference through brushing or bounding boxes, but such references are typically tied to explicit selection primitives rather than flexible, user-defined annotations. Hybrid approaches have integrated graphical interaction with language or attribute selection [KL92], but typically treat modalities as parallel input channels rather than as elements within a shared semantic space. As a result, many visual query systems struggle with open-ended analytical expression involving ambiguous references, relational constraints, or qualitative notions such as “near” or “not this region.” In particular, they offer limited capability for interpreting references that are implied through the relative placement of marks, text, and data rather than through explicit selection or linking mechanisms.

## 2.3. Sketch-based Interaction

Sketch-based interaction has long been studied as a means of facilitating informal, expressive input where rigid syntax is limiting [LM95, SSD01, HD07]. In data-centric contexts, sketching has been used for shape-based querying and pattern matching in time-series and spatial data [MA18, SLW\*21, YXL\*22]. More recent systems extend sketching toward exploratory analysis, sense-making, and communication, including narrative visualization and hypothesis-driven exploration [LKS13, RHRH\*19, LLY\*23]. Recent work by Yen et al. [YZV25] further explores how free-form sketches, annotations, and spatial relationships can be interpreted by AI models to iteratively guide system behavior in a shared workspace. Although the system incorporates multimodal inputs, the focus is on mapping sketch-based annotations to code edits within a localized programming context, rather than supporting broader multimodal queries that combine sketch, language, and deictic reference within a shared interaction space. Emerging work such as SketchGPT [HGS\*25] highlights the potential of combining sketch and language, but the role of spatial proximity and shared context in grounding analytical meaning remains underexplored.

More broadly, prior work on deictic interaction (e.g., pointing, labeling, and gestural reference) has shown that users rely on spatial and temporal context to disambiguate meaning, but these mechanisms are typically treated as explicit signals (e.g., arrows, selections, gestures) rather than as emergent relationships between collocated multimodal elements. Our work builds on these efforts by examining how users combine sketching, annotation, and language when they coexist in a shared interaction space. Through a multimodal design probe and user study, we observed interaction patterns in which users employed spatial, temporal, and visual proximity to relate sketches, annotations, and language. Reflecting on these patterns led us to formalize PS as a lens for reasoning about multimodal analytical query systems.

## 3. Design Probe for Exploring Multimodal Query Articulation



**Figure 1:** Design probe interface showing trend sketch (purple) and annotations (black). 1. Trend sketch; 2. Precision control; 3. Annotation tool; 4. Text input; 5. Eraser; 6. Cancel; 7. Clear; 8. Query.

Our design probe interface is shown in Figure 1. The probe explores how users combine geometric sketch-based search [MA18] with VLM-based interpretation of free-form annotations, with the goal of understanding emergent interaction patterns. The two interpretation pipelines operate separately but can be used in parallel, with their results intersected for final output. Brief implementation details are provided here with additional details in the supplemental materials.

### 3.1. Geometric Trend Search

The geometric trend search interprets a user’s sketch and compares it to trends in the database. User-controllable comparison penalties (e.g., precision, time invariance, sub-shape matching) allow fine-grained control over how sketches are matched (e.g., setting the time penalty to zero enables matches regardless of temporal alignment). Figure 2 illustrates this process. (A) A signal originates from either a user sketch (A.1) or the database (A.2). (B) Signals are simplified using the Douglas-Peucker algorithm [DP73]. (C) Segment properties (e.g., length, angle, position) are extracted. (D) Database signals are similarly processed. (E) Signals are compared

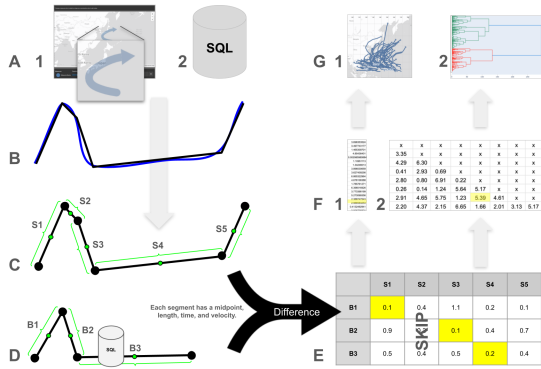


Figure 2: Overview of the probe's geometric processing pipeline.

on a per-segment, per-property basis, with differences weighted by user-defined penalties to compute a final similarity score. If the input is a single sketch, the comparison yields a  $1 \times N$  table (F1), which is sorted to return best matches (G1). If applied across the dataset, it produces an  $N \times N$  matrix (F2) used for analyses such as clustering (G2).

### 3.2. VLM-Interpreted Free-Form Annotations

Users can annotate data using natural language (e.g., “only storms after 1970”), symbols (e.g., arrows, OR), and glyphs such as cross-outs, circles, and boundary lines. While we provide default interpretations, users may repurpose these marks. The annotated image is sent to a VLM (Claude 3.7 Sonnet [Ant25]) along with prompts for task framing, schema grounding, visual context, glyph interpretation, and SQL generation. The VLM produces an annotation-derived SQL query, whose results are intersected with those of the geometric search. We separate the VLM and geometric pipelines because they support different forms of control: annotations require flexible, high-level interpretation suited to language models, while trend matching relies on precise, user-tunable numeric parameters. This separation allows each modality to leverage its strengths while enabling unified query execution through result intersection.

### 4. Preliminary Evaluation and Reflection

To evaluate our design probe, we conducted 45-minute interview sessions with 20 participants, including data scientists, sales consultants, and product managers. Participants were recruited through internal mailing lists and had prior experience with data visualization, though not necessarily with sketch-based querying.

- **Introduction** [~5 min.]: Participants were introduced to the probe. While no formal training was provided, we demonstrated the available tools (e.g., sketching, annotation, and text input) to ensure awareness without prescribing usage strategies or exposing the probe datasets.
- **Task Completion** [~25 min.]: Participants worked with two datasets: **baby names** and **storm paths**, completing five predefined tasks.
- **Line Chart – Baby Names** [~10 min.]: Participants searched for

names matching nuanced temporal patterns (e.g., peaks or resurgence).

- **Map – Storm Paths** [~10 min.]: Participants identified geographic patterns in storm trajectories using shape-based input and region filtering.
- **Open-Ended Task** [~5 min.]: Participants explored the system freely, allowing us to observe emergent behaviors beyond predefined tasks.
- **Feedback Session** [~10 min.]: Semi-structured interviews captured participant impressions. Sessions were recorded and transcribed for analysis.
- **System Usability Scale (SUS)**: Participants completed the SUS survey [Bro96]. We selected SUS as a lightweight, widely used measure to provide a baseline assessment of perceived usability, complementing our qualitative analysis rather than serving as the primary evaluative instrument.

**Analysis method.** We conducted a thematic analysis of transcripts and interaction logs. Two authors independently identified recurring patterns in how participants combined sketches, annotations, and text, then iteratively grouped them into higher-level categories. Disagreements were resolved through discussion. This process revealed a recurring pattern in which participants used spatial, temporal, and visual proximity to relate multimodal elements.

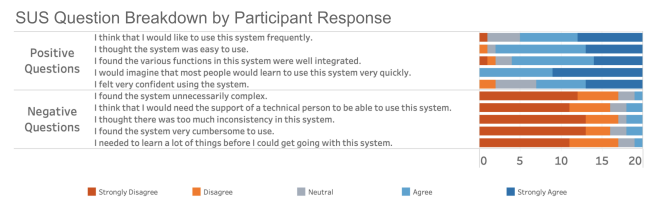


Figure 3: Breakdown of participant responses to SUS questions.

Participants found the probe expressive for exploring complex data trends, with an average SUS score of 81 ( $\geq 68$  indicates good usability [BKM09]). They appreciated the intuitive sketching interface and integration of annotations, while suggesting improvements in query interpretation, feedback, and onboarding. 17 of 20 participants used spatial proximity to convey meaning, often placing labels or glyphs adjacent to regions and referring to them with terms such as “this” or “here.” In addition, participants frequently combined proximity with explicit marks: arrows were commonly used to connect text and data regions, while circles and cross-outs were used to indicate inclusion and exclusion. However, these marks were rarely used in isolation; rather, their interpretation depended on spatial placement relative to candidate referents. For example, P11 labeled and referenced regions textually, forming multimodal queries through spatial grouping. However, three participants relied primarily on a single modality, producing less precise or more ambiguous queries.

Through analysis, we observed interaction patterns in which participants relied on spatial, temporal, and visual proximity to relate sketches, annotations, and language. These relationships were often established implicitly through placement rather than explicit connectors. Proximity functioned as a mechanism for disambiguating references, particularly when multiple candidate referents were

present. These observations motivated us to examine proximity as a mechanism for multimodal query articulation. In the following section, we introduce *proximity semantics* as a lens for describing this form of deictic disambiguation.

## 5. Proximity Semantics


*Proximity Semantics (PS)* is a reference-resolution framework in which relationships are established through chains of proximity bindings across measurable spaces (e.g. Euclidean, semantic, and temporal), yielding resolution sets of zero, one, or multiple candidates; it does not cover, for example, identification by property (e.g. “the largest”) or shared knowledge (e.g. “that thing we saw”). PS thus provides a single formalism for the deictic relationships that underlie common authoring conventions such as labels, arrows, and footnotes. PS has two primitives: *proximity bindings* and *connectors*. Consider Figure 4.A; a proximity binding is a referential relationship between two elements (e.g. a hammer and its label “Hammer”) whose measured distance in a common space falls within an interpretive threshold i.e., the two objects appear ‘close’ together. A *connector* such as an arrow can collapse distance within a space; this *logically* places two items next to each other, reducing complex referential relationships to simple proximity relationships, regardless of the length of the connector. For example, the label “Wrench” is spatially closest to the hammer icon, but the label’s arrow collapses its indicated distance, placing the label *logically* closest to the wrench icon and establishing the proximity binding.

A PS binding chain can cross multiple measurable spaces. In Figure 4.B, the exact semantic match (blue dotted line) between *Query 1* and Figure 4.A’s wrench label (both “Wrench”) renders them proximal in *semantic* space. Together with the *logical* spatial proximity of Figure 4.A’s wrench and label, the PS proximity chain from *Query 1* to wrench icon is established. *Query 2* follows a similar chain despite the misspelling because the semantic distance between “PPliers” and “Pliers” falls within our interpretive threshold, establishing semantic proximity and thus the full PS chain. Finally, *Query 3* is not semantically proximal to any of Figure 4.A’s labels and the PS chain is not established.

Given this, we now consider PS-interpretations of common deictic scenarios. For example, the arrow deixis in Figure 5.A adheres to Kaplan’s [Kap89] formalism of *demonstratives*, *demonstrations*, and *demonstrata*. Within the PS conceptualization however, this becomes a simple proximity chain where the arrow connectors collapse the distance between the words and glyphs, logically placing ‘this’ next to the circle around Japan and ‘here’ next to the exclusion cross-out in the Pacific, establishing disambiguated proximity-based deixis despite word:glyph distances in canvas space.

### 5.1. PS Query Space Definition and Query Tool Requirements

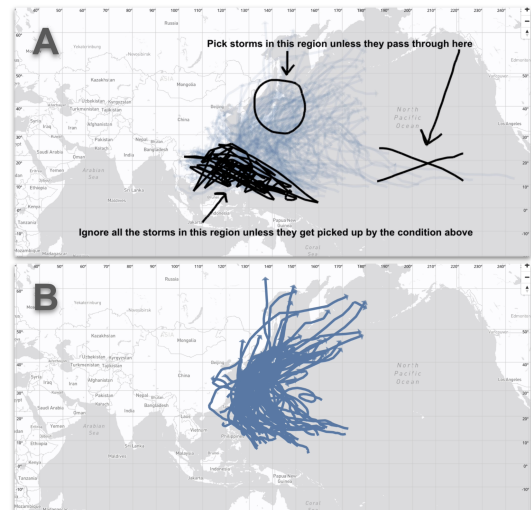
We define a PS Query Space as the set of queries whose correct interpretation depends on proximity-based deictic disambiguation. These are cases in which multiple plausible referents exist and cannot be resolved without establishing a PS chain through one or more measurable spaces. A PS query tool must therefore be capable of using single- or multi-space proximity bindings and connectors to establish a PS chain and resolve semantic ambiguity.



**Figure 4:** A. PS Examples. Simple spatial proximity (labeled hammer) and logical spatial proximity (wrench and pliers references). B. PS chains linking queries to objects across semantic and canvas spaces; red ‘x’ indicates no semantic match.

### 5.2. PS Query Tool Example

We revisit the design probe with these PS definitions and requirements in mind. Figure 5 shows a multimodal query in which geospatial data is filtered using natural language and annotations. The semantics of the query cannot be determined from any single modality. Instead, meaning emerges by resolving the PS chain that links textual references, glyphs, and data regions within the shared visual space.



**Figure 5:** A. A geospatial PS query. Multiple deictic references require proximity-based disambiguation for proper semantic interpretation. B. Results of query interpretation.

## 6. Conclusion

Motivated by interaction patterns observed in our design probe, we introduce the notion of PS and provide examples of both general multimodal interaction and natural language-visual querying contexts. We derive design recommendations for systems that support proximity-based query interpretations, demonstrate a research probe capable of executing these queries, and argue that PS is a unifying framework for deictic reference and a common design vocabulary for multimodal interfaces. Although PS represents a bounded interaction space, our findings suggest that spatial, temporal, and syntactic forms of proximity are meaningful mechanisms through which users express analytical intent.

## References

- [Ant25] ANTHROPIC: Claude 3.7 Sonnet. <https://www.anthropic.com/claude/sonnet>, 2025. 2, 3
- [BBS24] BENDECK A., BROMLEY D., SETLUR V.: SlopeSeeker: A Search Tool for Exploring a Dataset of Quantifiable Trends. In *IUI* (2024), ACM, p. 817–836. doi:<https://doi.org/10.1145/3640543.3645208>. 2
- [BCCL91] BATINI C., CATARCI T., COSTABILE M. F., LEVIALDI S.: Visual Query Systems: A Taxonomy. In *Visual Database Systems* (1991). doi:<https://doi.org/10.1006/jvlc.1997.0037>. 2
- [BKM09] BANGOR A., KORTUM P., MILLER J.: Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale. *J. Usability Studies* 4, 3 (2009), 114–123. 3
- [Bro96] BROOKE J.: SUS: A Quick and Dirty Usability Scale. *Usability Evaluation in Industry* (1996). 3
- [BS23] BROMLEY D., SETLUR V.: What Is the Difference Between a Mountain and a Molehill? Quantifying Semantic Labeling of Visual Features in Line Charts. *IEEE TVCG* (2023). doi:<https://doi.org/10.1109/VIS54172.2023.00041>. 2
- [BYRNM\*99] BAEZA-YATES R., RIBEIRO-NETO B., MILLS D., BONN O., JUAN S., MEXICO M., TAIPEI C., WESLEY A., LIMITED L.: Modern Information Retrieval. 1
- [CHH\*08] CIMIANO P., HAASE P., HEIZMANN J., MANTEL M., STUDER R.: Towards Portable Natural Language Interfaces to Knowledge Bases - The Case of the ORAKEL System. *Data & Knowledge Engineering* 65, 2 (2008), 325–354. doi:<https://doi.org/10.1016/j.datak.2007.10.007>. 1
- [DAC10] DAMLIJANOVIC D., AGATONOVIC M., CUNNINGHAM H.: Natural Language Interfaces to Ontologies: Combining Syntactic Analysis and Ontology-Based Lookup through the User Interaction. In *Semantic Web Conference* (2010). doi:[https://doi.org/10.1007/978-3-642-13486-9\\_8](https://doi.org/10.1007/978-3-642-13486-9_8). 1
- [DKR97] DERTHICK M., KOLOJEJCHICK J., ROTH S.: An Interactive Visual Query Environment for Exploring Data. doi:<https://doi.org/10.1145/263407.263545>. 2
- [DMN\*17] DHAMDHERE K., MCCURLEY K. S., NAHMIA R., SUNDARARAJAN M., YAN Q.: Analyza: Exploring Data with Conversation. In *IUI* (2017), pp. 493–504. doi:<https://doi.org/10.1145/3025171.3025227>. 1
- [DP73] DOUGLAS D. H., PEUCKER T. K.: Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or its Caricature. *The Canadian Cartographer* 10, 2 (1973), 112–122. doi:<https://doi.org/10.1002/9780470669488.ch2>. 2
- [FLS\*08] FERNANDEZ M., LOPEZ V., SABOU M., UREN V., VALLET D., MOTTA E., CASTELLS P.: Semantic Search Meets the Web. In *IEEE International Conference on Semantic Computing* (2008), pp. 253–260. doi:<https://doi.org/10.1109/ICSC.2008.52>. 1
- [GDA\*15] GAO T., DONTCHEVA M., ADAR E., LIU Z., KARAHALIOS K. G.: DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *UIST* (2015), ACM, pp. 489–500. doi:<https://doi.org/10.1145/2807442.2807478>. 1
- [HD07] HAMMOND T., DAVIS R.: LADDER: A Sketching Language for User Interface Developers. *SIGGRAPH*, ACM, p. 35–es. doi:<https://doi.org/10.1145/1281500.1281546>. 2
- [HGS\*25] HUANG Z., GAO C., SHAN Y., HU H., LI Q., DENG X., MA C., LAI Y.-K., LIU Y.-J., TIAN F., DAI G., WANG H.: SketchGPT: A Sketch-Based Multimodal Interface for Application-Agnostic LLM Interaction. In *UIST* (2025), ACM. doi:<https://doi.org/10.1145/3746059.3747598>. 2
- [HS04] HOCHHEISER H., SHNEIDERMAN B.: Dynamic Query Tools for Time Series Data Sets: Timebox Widgets for Interactive Exploration. *Information Visualization* 3 (2004), 1 – 18. doi:<https://doi.org/10.1057/palgrave.ivs.9500061>. 2
- [Kap89] KAPLAN D.: Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and other Indexicals. 4
- [KL92] KEIM D., LUM V.: Visual Query Specification in a Multimedia Database System. In *Visualization* (1992), pp. 194–201. doi:<https://doi.org/10.1109/VISUAL.1992.235208>. 2
- [LKS13] LEE B., KAZI R. H., SMITH G.: SketchStory: Telling More Engaging Stories with Data through Freeform Sketching. *IEEE TVCG* 19, 12 (2013), 2416–2425. doi:<https://doi.org/10.1109/TVCG.2013.191>. 2
- [LLG\*20] LEWIS M., LIU Y., GOYAL N., GHAZVININEJAD M., MOHAMED A., LEVY O., STOYANOV V., ZETTLEMOYER L.: BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL* (2020), pp. 7871–7880. doi:<https://doi.org/10.18653/v1/2020.acl-main.703>. 2
- [LLY\*23] LIN Y., LI H., YANG L., WU A., QU H.: InkSight: Leveraging Sketch Interaction for Documenting Chart Findings in Computational Notebooks. *IEEE TVCG* 30 (2023), 944–954. doi:<https://doi.org/10.1109/TVCG.2023.3327170>. 2
- [LM95] LANDAY J. A., MYERS B. A.: Interactive Sketching for the Early Stages of User Interface Design. In *CHI* (1995). doi:<https://doi.org/10.1145/223904.223910>. 2
- [LUB\*16] LI X., URICCHIO T., BALLAN L., BERTINI M., SNOEK C. G. M., BIMBO A. D.: Socializing the Semantic Gap: A Comparative Survey on Image Tag Assignment, Refinement, and Retrieval. *ACM Computing Surveys* 49, 1 (2016). doi:<https://doi.org/10.1145/2906152>. 1
- [MA18] MANNINO M., ABOUZIED A.: Qetch: Time Series Querying with Expressive Sketches. In *SIGMOD* (2018), ACM, p. 1741–1744. doi:<https://doi.org/10.1145/3183713.3193547>. 1, 2
- [MRS08] MANNING C. D., RAGHAVAN P., SCHÜTZE H.: *Introduction to Information Retrieval*. Cambridge University Press, Cambridge, UK, 2008. URL: <http://nlp.stanford.edu/IR-book/information-retrieval-book.html>. 1
- [Ope23] OPENAI: *GPT-4 Technical Report*. Tech. rep., OpenAI, 2023. URL: <https://openai.com/research/gpt-4>. 2
- [RHRH\*19] ROMAT H., HENRY RICHE N., HINCKLEY K., LEE B., APPERT C., PIETRIGA E., COLLINS C.: ActiveInk: (Th)Inking with Data. In *CHI* (2019), ACM, p. 1–13. doi:<https://doi.org/10.1145/3290605.3300272>. 2
- [SBT\*16] SETLUR V., BATTERSBY S. E., TORY M., GOSSWEILER R., CHANG A. X.: Eviza: A Natural Language Interface for Visual Analysis. In *UIST* (2016), ACM, pp. 365–377. doi:<https://doi.org/10.1145/2984511.2984588>. 1
- [SKL\*16] SIDDIQUI T., KIM A., LEE J., KARAHALIOS K., PARAMESWARAN A.: Effortless Data Exploration with Zenvisage: An Expressive and Interactive Visual Analytics System. 457–468. doi:<https://doi.org/10.14778/3025111.3025126>. 2
- [SLW\*21] SIDDIQUI T., LUH P., WANG Z., KARAHALIOS K., PARAMESWARAN A. G.: From Sketching to Natural Language: Expressive Visual Querying for Accelerating Insight. *SIGMOD Rec.* 50, 1 (2021), 51–58. doi:<https://doi.org/10.1145/3471485.3471498>. 2
- [SSD01] SEZGIN T. M., STAHOVICH T., DAVIS R.: Sketch-Based Interfaces: Early Processing for Sketch Understanding. In *Perceptive User Interfaces* (2001), ACM, p. 1–8. doi:<https://doi.org/10.1145/971478.971487>. 1, 2
- [YXL\*22] YAN L., XU N., LI G., BHOWMICK S. S., CHOI B., XU J.: SENSOR: Data-driven Construction of Sketch-Based Visual Query Interfaces for Time Series Data. *Proc. VLDB Endow.* 15, 12 (2022), 3650–3653. doi:<https://doi.org/10.14778/3554821.3554866>. 2
- [YZV25] YEN R., ZHAO J., VOGEL D.: Code Shaping: Iterative Code Editing with Free-form AI-Interpreted Sketching. In *CHI* (2025), ACM. doi:<https://doi.org/10.1145/3706598.3713822>. 2