Eurographics Conference on Visualization (EuroVis) 2017
B. Kozlíková, T. Schreck, and T. Wischgoll
(Guest Editors)

Short Paper

Natural Language Interfaces for Data Analysis with Visualization: Considering What Has and Could Be Asked

Arjun Srinivasan and John Stasko

School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

Natural language is emerging as a promising interaction paradigm for data analysis with visualization. Designing and implementing Natural Language Interfaces (NLIs) is a challenging task, however. In addition to being able to process and understand natural language expressions, NLIs for data visuailzation must consider other factors including input modalities, providing input affordances, and explaining system results, among others. In this article, we examine existing NLIs for data analysis with visualization, and compare and contrast them based on the tasks they allow people to perform. We discuss open research opportunities and themes for emerging NLIs in the visualization community. We also provide examples from the existing literature in the broader HCI community that may help explore some of the highlighted themes for future work. Our goal is to assist readers to understand the subtleties and challenges in designing NLIs and encourage the community to think further about NLIs for data analysis with visualization.

Categories and Subject Descriptors (according to ACM CCS): H.5.m [Information Interfaces and Presentation (e.g., HCI)]: Miscellaneous—

1. Introduction

Natural language is emerging as a promising interaction paradigm for data analysis with visualization. One reason why natural language interfaces (NLIs) are gaining popularity is because they may help improve the usability of visualization systems. It has been observed that while people (particularly novices) often know the questions they want to ask of their data, they can have difficulties in generating the desired visualizations or answers using existing tools [GTS10]. Further, people may be able to express their questions and intents more freely using natural language [GTS10, AKG*15]. NLIs, however, are challenging to design and implement as these systems not only need to handle inherent issues in parsing natural language, but they also require access to extensive knowledge and sophisticated reasoning to interpret the input expressions.

Over the years, substantial work has been done on developing NLIs for querying databases (NLIDBs) showing how the context of a database can be leveraged to allow people to ask questions of their data (e.g., [PEK03, SKI08, LJ14]). Several ideas presented by NLIDBs likely can be leveraged for visualization. However, designing and implementing NLIs for data analysis with visualization is arguably an even more challenging task. It requires considering several other factors including types of input modalities (e.g., only language-based, language-based + touch, touch + gaze + speech), providing input affordances (e.g., informing people what they can say/ask), explaining system results (e.g., providing feedback on why a chart was shown), among others.

In this article, we examine five existing NLIs for data visualization [CGH*01, SLJL10, GDA*15, SBT*16, AKG*16] with the goal of identifying commonalities and differences between them, and highlighting challenges associated with designing and implementing NLIs for visualization. We develop an analysis framework/structure consisting of multiple categories of tasks users may perform in a natural language driven visualization system. By referring to ongoing research within the visualization community and existing literature in the broader HCI community, we also discuss open research opportunities and themes for emerging NLIs for data visualization to consider. We hope that this article will assist readers in understanding the subtleties and challenges in designing NLIs for data visualization and encourage them to think further about NLIs for data analysis with visualization.

2. Existing NLIs for Data Visualization

NLIs for data visualization have been explored both within the research community and as commercial software (e.g. [IBM, MSP]). We limit our discussions in this article to the five most relevant research prototypes (shown in Table 1). Cox et al. [CGH*01] presented early work in this area and demonstrated that multimodal input provides more expressibility than a single input modality.

		Cox et al. [CGH*01]	Articulate [SLJL10]	DataTone [GDA*15]	Eviza [SBT*16]	Articulate2 [KADE*16]
Visualization-related	Generating Visualizations	□	✎	■		✎
	Visualization specific interactions	□	□	□	■	✎
	Formatting visualizations				*	
Data-related	Low-level analytical operations		□		✎	*
	High-level questions				*	
System control-related		□				□

Table 1: *Characterizing existing NLIs for data analysis with visualization based on tasks they allow people to perform* ***via natural language****.* □ *indicates a system provides minimal support for a task.* ✎ *indicates a system moderately supports a task.* ■ *indicates a system focuses on a task. "*" indicates that a task was identified as a part of an initial/formative study but we are not sure if the task is supported by the current version of the prototype.*

However, the range of questions and visualization types supported were very limited and users were expected to provide fully specified queries. Almost a decade later, the Articulate system [SLJL10] provided a NLI for visualization that mapped user queries to tasks and used these tasks in combination with data attributes to generate required visualizations. More recently, DataTone [GDA*15] allows people to generate visualizations using natural language queries. The system specifically focuses on detecting ambiguity in the input queries and uses a mixed-initiative approach to resolve this ambiguity and help people iteratively construct visualizations. Eviza [SBT*16], on the other hand, presents a visualization as a starting point and allows people to ask questions in the context of the given visualization. By emphasizing the ability for a person to continually revise and update their queries, Eviza seeks to provide a rich dialogue with the visualization. Articulate2 [AKG*16, KADE*16] is a prototype of a conversational interface for visualization that explores the dialogue between a user and a system to generate and interact with visualizations.

We have developed a descriptive framework/structure (Table 1) that encompasses the different tasks a person may seek to perform in a natural language driven data visualization system. We first compiled a list of tasks by considering utterances and query types that existing systems (columns in Table 1) identify and support. We then used an affinity diagramming approach, grouping similar tasks and iteratively refining the groups according to what we believed to be the core end goal of the tasks in each group. We combined groups under broader task categories (rows in Table 1) based on existing visualization task and interaction taxonomies [AES05, YaKS07]. This process resulted in three higher-order categories: *visualization-related tasks*, *data-related tasks*, and *system control-related* tasks. Some sample utterances presented in this section are taken from published articles on existing systems.

Visualization-related tasks are those that focus on generating visualizations or updating an existing visualization. More specifically, *generating visualizations* in Table 1 refers to the task of constructing or requesting new visualizations. *Visualization specific interactions*, on the other hand, refer to tasks that involve performing specific interactions w.r.t. a given visualization. The seven categories of interactions proposed by Yi et al. [YaKS07] (*Select*, *Explore*, *Reconfigure*, *Encode*, *Abstract/Elaborate*, *Filter*, *Connect*) can be considered as examples of interactions people may try to perform. *Formatting visualizations* involves updating a given chart purely from a graphical perspective (e.g., show/hide labels, change color from red to blue). Sample utterances within this category include:

Show me medals for hockey and skating by country. — Zoom in to Italy — Show Africa only — Select the largest bar in View 3. — Sort by average unemployment — Can you show it around the Loop by year broken down by crime type? — Show me a map of crimes in River North and the Loop. — Show y-axis label.

Data-related tasks include two main categories: low-level analytical tasks and operations, and high-level user questions. Examples of *low-level analytical operations* include some of the analytical tasks proposed by Amar et al. [AES05] such as *Retrieve Value*, *Compute Derived Value*, *Find Extremum*, *Determine Range*, *Characterize Distribution*, *Cluster*, and *Correlate*. On the other hand, *high-level questions*, are queries (e.g., "Which stock should I buy today?") that do not express a specific goal or task but request an answer for a more direct question that is typically derived using one or more analytical operations and mathematical functions. Responses for data-related tasks may not need to show or modify a visualization but could be just plain text (e.g., a yes/no response or a number indicating a requested value). However, visualizations can be used to enhance the understanding of the response (e.g., showing a correlation line when the person asks for a correlation value). Sample utterances for data-related tasks include:

What is the range of MPG? — Show top 3 countries — Show me the country with the largest outbreaks — During what time is the crime rate maximum, during the day or the night? — What is the best time to produce decaffeinated coffee? — Is there a seasonal trend for bike usage? — Which stock should I buy today?

Our final category, system control-related tasks, refers to tasks where people leverage natural language to perform operations such as moving or closing a window, inquiring about what questions can be asked, or ask questions to confirm their understanding of the data or visualization they are looking at. Sample utterances for this category include:

Move this window to the top right corner. — Help, what can I ask now? — Can you close the graph? — What other chart can I look at?

Table 1 may give the impression that each of the existing systems addresses just a few of these tasks. However, it is important to realize that each of these tasks is accompanied by challenges of natural language such as ambiguity and underspecified utterances. For example, while DataTone [GDA*15] primarily facilitates generating visualizations, it focuses on detecting and resolving ambiguity associated with input queries. For example, consider the utterance "Show me medals for hockey and skating by country" in the context of an Olympics dataset consisting of medal winning athletes.

Based on the ambiguity space proposed by Gao et al. [GDA*15], this seemingly simple query can lead to over 5000 possible visualizations that the system needs to choose from.

Table 1 also helps identify areas for future exploration and research. For instance, Eviza's initial study found that some participants wanted to format a given visualization using natural language. Several systems in the HCI community have explored NLIs for graphics manipulation [Hau89, PL91, Woo07]. As a first step in this direction, future systems could explore the use of language for formatting visualizations by allowing people to alter and adjust basic graphical encodings such as shapes, color, size, etc.

3. Opportunities and Challenges

While existing systems have taken important steps toward exploring NLIs for data analysis with visualization, these systems have only scratched the surface of what is possible. Drawing upon existing work in the visualization and broader HCI community, in this section we discuss five themes for future research in this area. Table 2 shows how the existing systems relate to the first four themes discussed in this section.

3.1. Exploring Multimodal Interaction in Post-WIMP Settings

The symbiotic relationship between natural language and direct manipulation-based input has been discussed for decades [CDM*89, OC00]. Given the dominance of direct manipulation-based techniques in visualization, it is no surprise that the majority of the five existing systems support some level of multimodal interaction. However, the existing systems have primarily focused on exploring these interactions in WIMP-based settings. An exception would be the data collection phase (a wizard-of-oz study) for Articulate2 that used a large wall-sized tiled display. However, the current version of the developed prototype still requires people to interact with a laptop and propagates responses to a larger screen.

The broader HCI community has more extensively explored multimodal interfaces facilitating natural language in post-WIMP settings [VD97]. Possibly the first, and one of the best known multimodal systems was presented in Bolt's article "Put-that-there" [Bol80] in 1980. The system allowed people to control shapes on a large display using a combination of gestural and voice based input. Hauptmann [Hau89] studied multimodal interaction in the context of graphics manipulation and found that people strongly prefer to use both gestures and speech for graphics manipulation. Subsequent research projects in the HCI community have examined the use of multimodal interfaces facilitating language-based interaction in systems for writing, painting, video browsing, map manipulation, among many others (e.g., [GNS*92, KW92, VW96, CJM*97]).

Exploring multimodal interaction on devices such as tablets, large interactive displays, and projected displays that support input modalities like pen, touch, and gestures, among others is an open research area for NLIs in data visualization. Further, evaluations of such multimodal visualization systems can help confirm/refute earlier findings (e.g. [Ovi97]) and myths of multimodal interaction [Ovi99] within the context of visualization systems.

3.2. Exploring Proactive System Behavior

Existing natural language driven visualization systems are typically "reactive" in the sense that they wait for a person to provide an input (as an utterance/query) and then generate a response. An alternative approach is to explore "proactive" systems that keep track of a person's interactions, are more aware of changes made in the system interface, and leverage interactions and changes to actively engage with users during sessions.

In its simplest form, proactive system behavior in the context of NLIs for visualization could be systems understanding contextual queries, where context is defined by the active system state and components. For example, consider queries asked in the context of a visualization. In a proactive setting, a system would map user queries to the appropriate visualization without the person needing to explicitly refer to the visualization in the query. A slightly advanced variation of proactive behavior could be the system non-intrusively making subtle interface changes and providing suggestions based on user interactions and questions over time. An example of this can include adding suggested visualizations and attributes to examine based on a given query and its response. At its extreme, proactive behavior can be considered an attribute of an intelligent virtual assistant that not only can answer questions but can guide the user through the analysis process. Past work within the HCI community has shown that if suggestions from proactive system are fairly competent, are timed correctly, and presented in non-intrusive ways, they can be beneficial and people even prefer to have such proactive features as part of a system [XCS03, XSC04, HSKK09].

Among existing systems, arguably only Eviza and DataTone exhibit some form proactive behavior by supporting contextual queries (both systems) and a query auto-complete feature (Eviza). An increasing number of systems within the visualization community are exploring visualization recommendations and suggesting potential next steps based on user interactions and preferences (e.g. [GW09, BGV16]). Devising methods of generating similar suggestions, and designing interfaces to enable proactive behavior and assistive features for data exploration is another open opportunity for future NLIs for visualization to explore.

3.3. Mixed-initiative Interaction

Allen, in his 2001 article on conversational HCI [AGH99] states "mixed-initiative interaction models human-machine interaction after human collaborative problem solving. Instead of being viewed as a series of commands, the interaction involves defining and discussing tasks, exploring ways to perform the task, and collaborating to get it done". Both DataTone and Eviza have explored the use of mixed-initiative interaction via *ambiguity widgets*. Ambiguity widgets are essentially dynamically appearing user interface widgets such as dropdowns and sliders that help people modify or correct responses generated by the system. An open opportunity for these systems to explore is leveraging mixed-initiative interaction to train themselves and improve over time. Systems can leverage user interactions (including modifications and corrections made in responses) to learn more about the users themselves. By learning from, and about users, these systems can present more personalized responses and avoid repeating mistakes in responses over time.

	Cox et al.	Articulate	DataTone	Eviza	Articulate2
Post-WIMP + Multimodal					*
Proactive vs Reactive	Reactive	Reactive	Minimally proactive	Moderately proactive	Reactive*
Instruction & Feedback	□		□	✎	*
Mixed-initiative Interaction			■	■	

Table 2: *Summarizing existing systems w.r.t. the discussed research themes. Refer to Table 1 for legend.*

Earlier research in HCI has discussed how mixed-initiative interaction can be especially effective at extending the capabilities of dialogue-based interfaces [AGH99, ABD*01]. Particularly for dialogue-based NLIs like Articulate2, exploring mixed-initiative interaction is an exciting open area for research. Earlier work [ABD*01] also has shown that adding behavioral agent-like modules as part of dialogue-based interfaces to facilitate mixed-initiative interaction can help make systems more intelligent and can even help improve users' task performance.

3.4. Instruction and Feedback

It is well-known that a key challenge with NLIs is providing affordances and helping people understand the queries they can pose and interactions they can perform [SM97, YLM95]. While having a visual interface to complement language-based interaction helps, it is important that systems leverage this advantage in the best possible ways. Another key challenge, especially with NLIs for visualization, is explaining (or providing feedback) why a certain interface element (e.g., slider) appeared or why a specific visualization was drawn in response to a query. Some of Eviza's user study participants also voiced concerns about feedback and felt the system was at times behaving like a "black-box" [SBT*16].

Among the existing systems, Eviza and the system by Cox et al. [CGH*01] try to address the challenge of helping people understand what queries they can pose using an auto-complete feature and a "help" command respectively. An issue with both these approaches is that they give people a false sense of limited parsing capabilities. This was also highlighted in Eviza's evaluation where some participants felt like they had to select from the suggested list of sentence completions. Further, auto-complete is not supported in the case of speech-based input.

Both DataTone and Eviza try to address the challenge of feedback by either positioning ambiguity widgets relative to the utterance (in DataTone), or appending descriptive text alongside ambiguity widgets to provide more context (in Eviza). Going forward, an open challenge for NLIs for visualization is to explore more ways to leverage their visual interface to provide affordances and feedback, and mitigate the "black-box" effect.

3.5. Studying the nature of utterances

Existing NLIs for data visualization mostly focus on identifying tasks that people may perform and developing parsers and classifiers to identify these tasks in utterances. An opportunity for further research lies in understanding more about the types of utterances people articulate. For instance, a very basic system may only allow people to communicate command-like utterances while a more sophisticated system may allow people to speak more naturally and present full sentences or ask questions. In addition to identifying tasks (the "what"), formative studies and systems evaluations should also focus on identifying and supporting various utterance types (the "how"). Additionally, understanding usage patterns of utterances can also help inform architectural and algorithmic design decisions for the natural language parsing modules. For instance, by considering follow-up questions and observing how they were asked, Eviza [SBT*16] was able to leverage a finite state machine based approach for its pragmatics module.

From a computational linguistics perspective, understanding more about the types of utterances produced may help invent new language parsing techniques. For instance, existing NLIs for visualization have typically either used a grammar-based or keyword-based parsing approach. Considering how multiple approaches could be combined, or exploring how direct manipulation interaction used to complement natural language can be leveraged to improve a parser's performance are all open areas for exploration.

The notion of how people ask questions becomes even more crucial to understand when considering multimodal interaction. People may prefer one modality over another for different tasks and use multiple modes of input in specific sequences to accomplish the tasks [OCW*00, Ovi03]. Understanding these preferences and more particularly how people specify inputs can help make more informed and smarter architectural design decisions while implementing multimodal input processing modules and frameworks.

4. Conclusion

We examined existing NLIs for data analysis with visualization, and highlighted commonalities and differences between them based on the tasks they let people perform via natural language interaction. We draw connections between ongoing work within the visualization community and existing research in the broader HCI community to highlight open research opportunities and themes for future work to consider. Another key challenge with NLIs for visualization is effective evaluation of prototypes. While we do not explicitly discuss evaluation in this article, development of creative new evaluation techniques (e.g., jeopardy-style evaluation [GDA*15]) for these systems, and measuring metrics such as engagement, enjoyment, and flow in addition to usability and performance [SSK15, SSK16] is another important open area of research for future work to consider.

References

[ABD*01] ALLEN J. F., BYRON D. K., DZIKOVSKA M., FERGUSON G., GALESCU L., STENT A.: Toward conversational human-computer interaction. *AI magazine 22*, 4 (2001), 27. 4

[AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *Proceedings of the 2005 IEEE Symposium on Information Visualization* (2005), pp. 111–117. 2

[AGH99] Allen J., Guinn C. I., Horvtz E.: Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications 14*, 5 (1999), 14–23. 3, 4

[AKG*15] Aurisano J., Kumar A., Gonzales A., Reda K., Leigh J., Di Eugenio B., Johnson A.: "Show me data": Observational study of a conversational interface in visual data exploration. In *IEEE VIS '15 (Poster paper)* (2015). 1

[AKG*16] Aurisano J., Kumar A., Gonzalez A., Leigh J., Di Eugenio B., Johnson A.: Articulate2: Toward a conversational interface for visual data exploration. In *IEEE VIS '16 (Poster paper)* (2016). 1, 2

[BGV16] Bouali F., Guettala A., Venturini G.: Vizassist: an interactive user assistant for visual data mining. *The Visual Computer 32*, 11 (2016), 1447–1463. 3

[Bol80] Bolt R. A.: “put-that-there”: Voice and gesture at the graphics interface. In *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques* (1980), SIGGRAPH '80, pp. 262–270. 3

[CDM*89] Cohen P. R., Dalrymple M., Moran D. B., Pereira F., Sullivan J. W.: Synergistic use of direct manipulation and natural language. In *ACM SIGCHI Bulletin* (1989), vol. 20, pp. 227–233. 3

[CGH*01] Cox K., Grinter R. E., Hibino S. L., Jagadeesan L. J., Mantilla D.: A multi-modal natural language interface to an information visualization environment. *International Journal of Speech Technology 4*, 3-4 (2001), 297–314. 1, 2, 4

[CJM*97] Cohen P. R., Johnston M., McGee D., Oviatt S., Pittman J., Smith I., Chen L., Clow J.: Quickset: Multimodal interaction for distributed applications. In *Proceedings of the 5th ACM International Conference on Multimedia* (1997), ACM, pp. 31–40. 3

[GDA*15] Gao T., Dontcheva M., Adar E., Liu Z., Karahalios K. G.: DataTone: Managing Ambiguity in Natural Language Interfaces for Data Visualization. In *Proceedings of UIST '15* (2015), ACM, pp. 489–500. 1, 2, 3, 4

[GNS*92] Gourdol A. P., Nigay L., Salber D., Coutaz J., et al.: Two case studies of software architecture for multimodal interactive systems: Voicepaint and a voice-enabled graphical notebook. *Engineering for Human-Computer Interaction 92* (1992), 271–84. 3

[GTS10] Grammel L., Tory M., Storey M.-A.: How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 943–952. 1

[GW09] Gotz D., Wen Z.: Behavior-driven visualization recommendation. In *Proceedings of IUI* (2009), ACM, pp. 315–324. 3

[Hau89] Hauptmann A. G.: Speech and gestures for graphic image manipulation. In *ACM SIGCHI Bulletin* (1989), vol. 20, ACM, pp. 241–245. 3

[HSKK09] Hong J., Suh E.-H., Kim J., Kim S.: Context-aware system for proactive personalized service based on context history. *Expert Systems with Applications 36*, 4 (2009), 7448–7457. 3

[IBM] IBM watson analytics. `http://www.ibm.com/analytics/watson-analytics/`. 1

[KADE*16] Kumar A., Aurisano J., Di Eugenio B., Johnson A., Gonzalez A., Leigh J.: Towards a dialogue system that supports rich visualizations of data. In *17th Annual Meeting of the Special Interest Group on Discourse and Dialogue* (2016), p. 304. 2

[KW92] Kuijpers E., Wilson M.: A multi-modal interface for man machine interaction with knowledge based systems-mmi 2. *Knowledge Based Systems 2001* (1992), 2. 3

[LJ14] Li F., Jagadish H.: Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment 8*, 1 (2014), 73–84. 1

[MSP] Microsoft Power BI. `https://powerbi.microsoft.com/en-us/`. 1

[OC00] Oviatt S., Cohen P.: Perceptual user interfaces: multimodal interfaces that process what comes naturally. *Communications of the ACM 43*, 3 (2000), 45–53. 3

[OCW*00] Oviatt S., Cohen P., Wu L., Duncan L., Suhm B., Bers J., Holzman T., Winograd T., Landay J., Larson J., et al.: Designing the user interface for multimodal speech and pen-based gesture applications: state-of-the-art systems and future research directions. *Human-Computer Interaction 15*, 4 (2000), 263–322. 4

[Ovi97] Oviatt S.: Multimodal interactive maps: Designing for human performance. *Human-Computer Interaction 12*, 1 (1997), 93–129. 3

[Ovi99] Oviatt S.: Ten myths of multimodal interaction. *Communications of the ACM 42*, 11 (1999), 74–81. 3

[Ovi03] Oviatt S.: Multimodal interfaces. *The human-computer interaction handbook: Fundamentals, evolving technologies and emerging applications 14* (2003), 286–304. 4

[PEK03] Popescu A.-M., Etzioni O., Kautz H.: Towards a theory of natural language interfaces to databases. In *Proceedings of IUI* (2003), ACM, pp. 149–157. 1

[PL91] Pausch R., Leatherby J. H.: An empirical study: Adding voice input to a graphical editor. In *Journal of the American Voice Input/Output Society* (1991), Citeseer. 3

[SBT*16] Setlur V., Battersby S. E., Tory M., Gossweiler R., Chang A. X.: Eviza: A natural language interface for visual analysis. In *Proceedings of UIST '16* (2016), ACM, pp. 365–377. 1, 2, 4

[SKI08] Simitsis A., Koutrika G., Ioannidis Y.: Précis: from unstructured keywords as queries to structured databases as answers. *The VLDB Journal - The International Journal on Very Large Data Bases 17*, 1 (2008), 117–149. 1

[SLJL10] Sun Y., Leigh J., Johnson A., Lee S.: Articulate: A semi-automated model for translating natural language queries into meaningful visualizations. In *International Symposium on Smart Graphics* (2010), Springer, pp. 184–195. 1, 2

[SM97] Shneiderman B., Maes P.: Direct manipulation vs. interface agents. *interactions 4*, 6 (1997), 42–61. 4

[SSK15] Saket B., Scheidegger C., Kobourov S.: Towards understanding enjoyment and flow in information visualization. In *Computer Graphics Forum (EuroVis 15)* (2015). 4

[SSK16] Saket B., Scheidegger C., Kobourov S.: Comparing node-link and node-link-group visualizations from an enjoyment perspective. In *Computer Graphics Forum* (2016), IEEE. 4

[VD97] Van Dam A.: Post-wimp user interfaces. *Communications of the ACM 40*, 2 (1997), 63–67. 3

[VW96] Vo M. T., Wood C.: Building an application framework for speech and pen input integration in multimodal learning interfaces. In *Proceedings of ICASSP-96* (1996), vol. 6, IEEE, pp. 3545–3548. 3

[Woo07] Woolfe G.: Making color adjustment accessible to non-experts through the use of language. In *Color and Imaging Conference* (2007), no. 1, pp. 3–7. 3

[XCS03] Xiao J., Catrambone R., Stasko J.: Be quiet? evaluating proactive and reactive user interface assistants. In *Proceedings of INTERACT* (2003), vol. 3, pp. 383–390. 3

[XSC04] Xiao J., Stasko J., Catrambone R.: An empirical study of the effect of agent competence on user performance and perception. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1* (2004), IEEE Computer Society, pp. 178–185. 3

[YaKS07] Yi J. S., ah Kang Y., Stasko J.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1224–1231. 2

[YLM95] Yankelovich N., Levow G.-A., Marx M.: Designing speechacts: Issues in speech user interfaces. In *Proceedings of ACM SIGCHI* (1995), ACM Press/Addison-Wesley Publishing Co., pp. 369–376. 4