Semantic Interaction for Visual Analytics

Toward Coupling Cognition and Computation

Alex Endert
Pacific Northwest National Laboratory

The world is becoming increasingly instrumented with sensors, monitoring, and other methods for generating data describing social, physical, and natural phenomena. So, data exist that could be analyzed to uncover, or discover, the phenomena from which they were created. However, as the analytic models leveraged to analyze these data continue to increase in complexity and computational capability, how can visualizations and user interaction methodologies adapt and evolve to continue to foster discovery and sensemaking?

User interaction is critical to such visual data exploration's success because it lets users test assertions, assumptions, and hypotheses about the information, given their prior knowledge about the world. This cognitive process can be generally called sensemaking. Visual analytics (VA) emphasizes sensemaking of large, complex datasets through interactively exploring visualizations generated through a combination of analytic models. (For more on this, see the related sidebar.) So, a central focus is understanding how to leverage human cognition in concert with powerful computation through usable visual metaphors.

My PhD dissertation coined the term *semantic interaction* in the context of a user interaction methodology for model steering in VA systems.[1] It made three primary contributions. First, it explained the interactions users commonly employ when analyzing text information spatially without computational layout models, and the meaning they externalize into the manually crafted spatial constructs.[2,3] Second, it enabled bidirectionality of spatializations by inverting popular dimension reduction models.[4–6] Finally, it evaluated semantic interaction's impact on sensemaking through the synchronization of the analytic-model parameters, the visualization, and the user's insights in the text analysis domain.[7]

Semantic Interaction

Semantic interaction aims to enable co-reasoning between the user and the analytic models (coupling cognition and computation) without requiring the user to directly control them. To do this, it utilizes the visual metaphor in two ways:

- the metaphor through which the insights are obtained (that is, the visualization of information created by computational models) and
- the interaction metaphor through which hypotheses and assertions are communicated (that is, interaction occurs within the visual metaphor).

Users directly manipulate data in visualizations; semantic interaction then captures tacit knowledge of the user and steers the underlying analytic models. These models can be adapted incrementally on the basis of the user's sensemaking process and domain expertise explicated through the user's interaction. (For semantic interaction design guidelines, see the related sidebar.)

That is, the visualization's visual constructs expose the underlying analytic models' parameters. On the basis of common visual metaphors (such as the geographic, spatial metaphor in which proximity approximates similarity), we can infer tacit knowledge of the user's reasoning by inverting these analytic models. So, users are shielded from the underlying complexities and can interact with their data through a bidirectional visual medium. The interactions users perform in the visualizations to augment the visual encodings within the metaphor enable the inference of their analytic reasoning, which is systematically applied to the underlying models.

The Semantic Interaction Pipeline

The information visualization pipeline in Figure 1 shows how data characteristics are extracted and

 Published by the IEEE Computer Society

assigned visual attributes or encodings, ultimately creating a visualization.[8] Visualizations following this pipeline exhibit two primary components of the visual interface: the visualization showing the information and a GUI. The GUI's graphical controls (sliders, knobs, and so on) let users directly manipulate the parameters they control.

For example, direct manipulation user interfaces let users directly augment the values of data parameters and see the corresponding change in the visualization.[9] (One example is using a slider to set the range of home prices and observing the filtered results in a map showing homes for sale.) This model has been a successful user interaction framework for information visualizations. Figure 2a shows an example of such an interface.

VA systems have adopted this approach. However, a distinct difference is the added complexity of the models (and their parameters) being controlled. For example, instead of filtering the data by selecting ranges for home prices, users employ graphical controls over model parameters such as weighting the mixture of eigenvectors of a principal component analysis (PCA) dimension reduction model to produce 2D views of high-dimensional data. To users without expertise in such models, this poses fundamental usability challenges. Figure 2b shows an example of this type of direct manipulation interface.

The semantic interaction pipeline (see Figure 3) directly binds model-steering techniques to the interactive affordances created by the visualization. For example, a distance function used to determine the relative similarity between two data points (often visually depicted using distance in a spatial layout) can be the interactive affordance to let users to explore that relationship. So, the user interacts directly with the visual metaphor, creating a bidirectional medium between the user and the analytic models. This interaction method is similar to "by example" interaction because users can directly show their intention using the visualization's structure. This adds to visualization's role in the reasoning process, in that it's not only a method for gaining insight but also one for directly interacting with the information and the system.

The bidirectionality afforded by semantic interaction comes through binding the parameter controls traditionally afforded by the GUI directly within the visual metaphor. Through this binding, the system can infer the user's analytic reasoning from the user's interaction with the visualization regarding the underlying mathematical model's parameters. Specifically, a spatial layout is one visual metaphor in which my

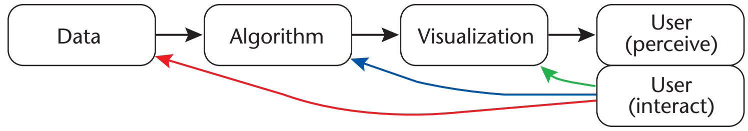

Figure 1. The information visualization pipeline.[8] Users can directly interact with the data (for example, filtering or correcting values), algorithm (for example, adjusting weights of relationships or changing parameter values), or visualization (for example, selecting a different encoding or modifying zoom levels).

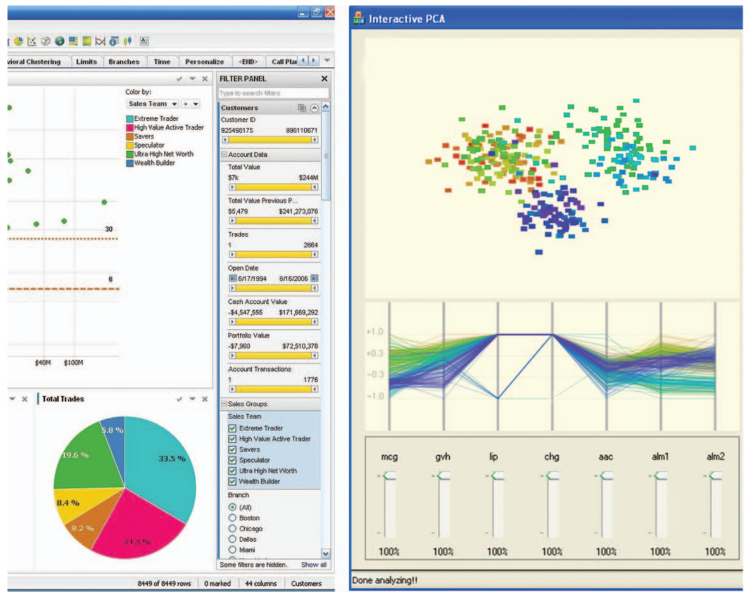

(a) **(b)**

Figure 2. Examples of two types of direct manipulation interfaces. (a) Spotfire employs direct manipulation for dynamic querying (ranges for data values, such as the portfolio value or number of trades) for information visualization. (b) iPCA applies direct manipulation to visual analytics (VA)—for example, directly controlling each dimension's relative contribution for principal component analysis.[10]

colleagues and I have conducted much semantic interaction research.[4,6,7]

Semantic Interaction with Spatializations

A spatial visual metaphor (a spatialization) demonstrates the bidirectionality afforded by semantic interaction. A spatial metaphor lends itself to common dimension reduction models to reduce the dimensionality of complex data to two dimensions. For example, relationships and similarities between high-dimensional data objects can be shown in two dimensions by leveraging such dimension reduction models as PCA, multidimensional scaling, and force-directed layouts. Generally, these models try to approximate the distance between data objects in their true, high-dimensional representation using fewer dimensions.

Researchers have applied semantic interaction methods to this visual metaphor. For example,

Visual Analytics for Sensemaking

Sensemaking is the process of someone acquiring an understanding of the world based on that person's conceptual model of events, actions, and information. Researchers have developed visual-analytics (VA) systems that support aspects of this process. This support can be characterized by the systems' user interactions, especially as they pertain to the visual metaphor and underlying models. Sensemaking has two primary subprocesses: foraging and synthesis.[1]

Foraging

During foraging, users filter and gather collections of interesting or relevant information. Scientists categorize VA tools that support foraging by their ability to pass data through complex analytic and statistical models and visualize the dataset's computed structure for the user to gain insight (see Figure A). So, users interact with these tools primarily by directly manipulating the models' parameters.

For example, interfaces that apply the information visualization interaction methodology of direct manipulation[2] present users with a set of graphical controls (sliders, knobs, and so on) to control and modify the model parameters' values. In VA tools, understanding these parameters (and the result of changing their values) can be difficult and is often outside the area of expertise for an expert in the specific data domain (for example, genomics and international politics). In these cases, users must translate their domain expertise and semantics about the information to determine which parameters to adjust (and by how much)—a fundamental usability concern.

VA tools leverage such models as entity extraction, topic modeling, link analysis, dimensionality reduction, clustering, and labeling. These models use various distance metrics to measure similarity between data objects. You can use these models to spatialize data. For example, you can represent unstructured text as a *bag of words*, high-dimensional data in which each dimension is a unique keyword or phrase in the text. Visualizations such as IN-SPIRE's Galaxy View[3] organize points representing text documents such that nearby points represent similar documents. This helps users recognize relationships between documents and between clusters of documents.

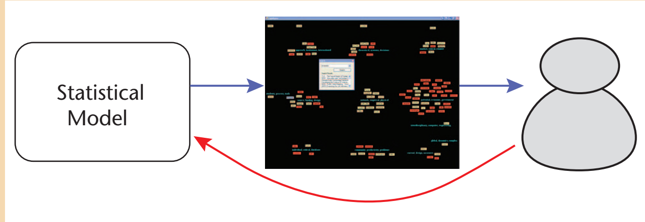

Figure A. Interaction with foraging tools. Users interact directly with the statistical model (red), then gain insight through observing the change in the visualization (blue).

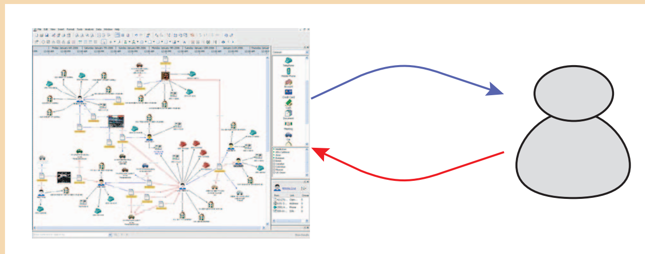

Figure B. Interaction with synthesis tools. Users manually create a spatial layout of the information to maintain and organize their insights about the data.

Synthesis

On the basis of the information acquired from foraging, users advance through the synthesis stages. In these stages, they construct and test hypotheses about how the foraged information might relate to their understanding of the world. Synthesis tools let users organize and maintain their hypotheses and insight regarding the data (see Figure B). These tools often employ a flexible, informal spatial medium or canvas.

For example, by organizing spatial layouts, users can externalize their insights about a dataset on the basis of the information's position.[4] Users frequently organize such layouts by complex schemas and mixed metaphors, often organized topically according to the semantics relevant to their analysis needs. Analysts use tools that support manually constructing spatializations to visually synthesize hypotheses.[5] That is, they create spatial structures (often mixing clusters, timelines, connections, geography, order

inverting PCA, multidimensional scaling, and generative topographic mapping can enable semantic interaction in bidirectional spatializations.[4,11] The ability to understand each model's parameters that can be exposed through the visual encoding (in this case, the relative distance between data points) enabled this affordance. Further research has explored the tradeoffs between the various ways to map the user feedback of changing the relative distance between data objects to the underlying dimension reduction models.[5,12]

Impact: Current and Future

Semantic interaction to increase the usability of complex VA systems has evolved along with VA's growth and maturity as a research discipline. Interactivity has become increasingly important, and users' attempts to communicate their hypotheses and assertions about the data to foster sensemaking have continued to employ (if not depend on) analytic models. Semantic interaction has helped foster this communication between the user and the model, having an impact beyond that at the time of my dissertation.

of discovery, process waypoints, and so on) that carry meaning to them regarding their sensemaking process.

Such informal relationships in the spatial layout are beneficial because they don't require users to overformalize relationships too early in the process. This gradual increase in relationship formality is called *incremental formalism*.[6] This approach directly presents the user interaction to users both in the visual metaphor and on the data. So, the users can leverage their domain expertise to make sense of the information.

References

1. P. Pirolli and S.K Card, "Information Foraging Models of Browsers for Very Large Document Spaces," *Proc. 1998 Int'l Working Conf. Advanced Visual Interfaces* (AVI 98), 1998, pp. 83–93.
2. B. Shneiderman, "Direct Manipulation: A Step beyond Programming Languages," *Computer*, vol. 16, no. 8, 1983, pp. 57–69.
3. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Visual Text Analytics," *Proc. 2012 SIGCHI Conf. Human Factors in Computing Systems* (CHI 12), 2012, pp. 473–482.
4. C. Andrews, A. Endert, and C. North, "Space to Think: Large, High-Resolution Displays for Sensemaking," *Proc. 2010 SIGCHI Conf. Human Factors in Computing Systems* (CHI 10), 2010, pp. 55–64.
5. A. Endert et al., "The Semantics of Clustering: Analysis of User-Generated Spatializations of Text Documents," *Proc. Int'l Working Conf. Advanced Visual Interfaces* (AVI 12), 2012.
6. F.M. Shipman III and C.C. Marshall, "Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems," *Computer Supported Cooperative Work*, vol. 8, no. 4, 1999, pp. 333–352.

Making Insights in Big Data Accessible

The ForceSPIRE system demonstrates how a spatialization of text documents can be the primary interface for user interaction (see Figure 4).[6] ForceSPIRE uses relative distance to indicate documents' similarity. It computes the distances through force-directed layout. The single spatial layout is the primary view, through which most interaction occurs. We chose the user interactions specifically to correspond with those found during studies observing users performing text analysis using a spatial metaphor.[2,3] The studies found that users reposition documents, highlight phrases, take notes, and perform text searches while actively reading. ForceSPIRE couples each of these interactions with model updates.[6]

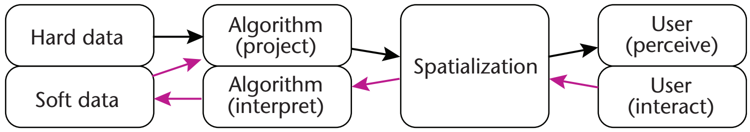

Figure 3. The semantic interaction pipeline. Users interact directly with the visualization, from which inferences are made to update the model or algorithm. Semantic interaction uses the stored "soft data" in conjunction with the "hard data" (raw data) to incorporate the user's expertise into the VA system.

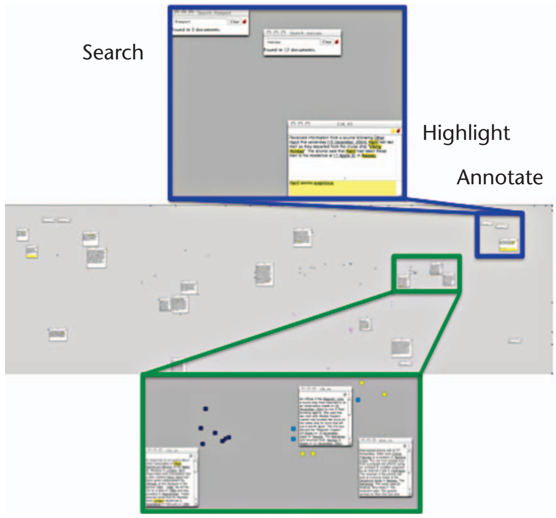

Figure 4. With ForceSPIRE, users can search, highlight, annotate, and reposition documents spatially. Documents can appear as minimized rectangles (see the yellow, blue, and teal rectangles in the enlarged region at the bottom) and as full-detail windows (resizable by the user). ForceSPIRE makes model inferences on each user interaction, creating machine and human co-reasoning.

My colleagues and I directly extended the findings from this research into work in analyzing large volumes of text. We used multiple tiers and styles of analytic and mathematical models to process and retrieve data, extract features, and so on. Each of these stages in the data-processing pipeline presents opportunities to steer the model on the basis of the inference of the user interaction.[13] For example, a challenge in large data volumes is retrieving only the most relevant subset of the data to maintain locally and visualize. Thus, how can semantic interaction steer information retrieval techniques to locally maintain and visualize only the most relevant information with respect to the user's analytic process? Many such techniques can benefit

Semantic Interaction Design Guidelines

Here are guidelines for semantic interaction for spatializations:[1]

- A visual "near = similar" metaphor supports analysts' spatial cognition and is generated by statistical models and similarity metrics.[2]
- Use semantic interactions within the visual metaphor, based on common interactions occurring in spatial analytic processes[3] such as searching, highlighting, annotating, and repositioning documents.
- Interpret and map the semantic interactions to the model's underlying parameters, by updating weights and adding information.
- Shield users from the complexity of the underlying mathematical models and parameters.
- Models should learn incrementally by taking into account interaction during the entire analytic process, supporting analysts' process of incremental formalism.[4]
- Provide visual feedback of the updated model and learned parameters within the visual metaphor.
- Reuse learned model parameters in streaming data or future data.

References

1. A. Endert, "Semantic Interaction for Visual Analytics: Inferring Analytical Reasoning for Model Steering," PhD dissertation, Dept. Computer Science, Virginia Tech, 2012.
2. A. Skupin, "A Cartographic Approach to Visualizing Conference Abstracts," *IEEE Computer Graphics and Applications*, vol. 22, no. 1, 2002, pp. 50–58.
3. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Visual Text Analytics," *Proc. 2012 SIGCHI Conf. Human Factors in Computing Systems* (CHI 12), 2012, pp. 473–482.
4. F.M. Shipman III and C.C. Marshall, "Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems," *Computer Supported Cooperative Work*, vol. 8, no. 4, 1999, pp. 333–352.

from the information inferred about the user to more accurately query within, and across, databases containing relevant information. That is, how can semantic interaction scale the inferred reasoning of the user into the larger data volumes through the malleability of information retrieval techniques? Furthermore, this might require additional visual representations (or aggregations) of information.

Semantic interaction has impacted projects at Pacific Northwest National Laboratory that stem from user needs to understand these large volumes of text data. Semantic interaction's capability to capture the analytic reasoning associated with a user interaction and amplify that reasoning into the analytic model lets users extend their reach and coverage into the larger data scales. These users' domain expertise generally does not include knowledge in statistics or the data sciences. So, placing their user interaction directly onto the visual data representations enables them to reason on the data using the visualization and to communicate their hypotheses and assertions directly in the visualization. Anecdotal feedback from these users has been positive, with a user evaluation in progress. Similarly, research at Virginia Tech is investigating how semantic interaction can help steer information retrieval techniques to address big-data challenges. This research is fundamentally advancing our understanding of semantic interaction and evolving ForceSPIRE as a testbed for prototyping and evaluating specific pairings of user interaction and computation.

Semantic interaction techniques have also affected big-data challenges that emphasize a variety of data (for example, multimedia). Phenomena that are captured, collected, and encoded digitally often span multiple media types. So, promoting sensemaking through VA technologies often requires users to reason across multiple media types. One challenge with such heterogeneous datasets is to correlate, or fuse, the data types' feature spaces that represent a cognitively cohesive concept or topic. Through inferring the higher-level analytic reasoning from user interaction tailored toward each of these data types, the opportunity exists to successfully decode phenomena whose discovery and understanding require multiple data types.

From Streaming Data to Streaming Insights

The continuous sensing and collecting of information poses streaming-data challenges and opportunities. A specific challenge is how to understand evolving and changing phenomena in real time. In terms of steering and adapting the underlying models using semantic interaction, challenges exist regarding the temporal nature of the data and the reasoning process. As users generate hypotheses and reason about the data, how can the models interpret the temporal nature of those hypotheses and assertions? How can VA systems working with streaming data understand the temporal importance of what information to retain and what to delete as a user progresses through sensemaking?

Researchers are applying semantic interaction to streaming-data challenges (following the last design guideline in the sidebar "Semantic Interaction Design Guidelines"). Instead of using semantic interaction to understand the features users are interested in over time, the goal here might be to understand the features or data that users don't show interest in. So, semantic interaction enables streaming models to determine what information to "forget." For example, dimension reduction models can understand what dimensions carry little if any weight, given the user's context. Similarly, reasoning models can learn what rules or assertions are no longer valid, weight them according to a belief propagation network implicit from the user, or create new ones from the user's domain expertise.

Evaluating Visual Analysis

VA technology has evolved from visualizing information to visualizing an analytic model's approximation of data (the model's output). Such a model approximates, or fits, the information given a specific parameterization of that model. The ability to steer, select, and refine such models is critical because they result in the generated visualizations. This poses the challenge of measuring the similarity between a user's conceptual model of a topic or domain and the analytic model's approximation of the information.

My colleagues and I have evaluated how semantic interaction affects the analytic process. Semantic interaction is intended primarily for sensemaking and discovery tasks. So, the goal is to foster the creation of insight. Prior research has investigated the challenges of evaluating visualizations intended for open-ended discovery.[14,15] Thus, to evaluate semantic interaction (specifically, in the context of text analysis using ForceSPIRE), we can observe the analytic process and the analytic product. For example, my colleagues and I evaluated semantic interaction's ability to couple cognition and computation through visualization by analyzing the evolution of three components throughout a user study: the model parameter weights, the visualization, and the user's insights.[7] Our research showed that semantic interaction could incrementally steer the underlying model, and in turn the visualization, to coincide with the user's analysis and insights.

This raises the question, does temporal synchronization between the model parameter weights, the visual representation, and the user's insights represent a valuable metric for evaluating VA tools for discovery? Such an approach for evaluating visual data exploration performs well in conjunction with methods such as insight-based evaluation,[16] to understand the evolution of a user's insight over time. Additionally, whereas much semantic interaction research has focused on implicitly steering models, the holistic design of VA tools will likely combine explicit, direct manipulation interactions in concert with semantic interactions to provide users with both direct manipulation controls and implicit knowledge amplification when desired.

Toward a Science of Interaction

The need to understand, measure, and quantify the analysis process has created a study, or science, of interaction.[17] (For more on this, see the sidebar "Inferring Reasoning from User Interaction.") The underlying claim is that user interactions embody and externalize aspects of the analysis process. Semantic interaction can help further this scientific understanding of user interaction by systematically quantifying the interaction and binding it to model parameters.

The research I've been describing has looked at how to analyze user interaction directly within the visual metaphor to reveal analytical reasoning. However, other scientific areas have studied user interaction data captured from other metaphors, such as clickstream data for Web browsing, physical and social movement for behavioral analysis of groups of people, and product purchase trends for marketing.

Also, beyond direct model steering, the analysis of the user interaction can include understanding user biases and cognitive stages during a sensemaking task. The "soft data" (see Figure 3) collected from user interaction can be the basis of study. Theories and models for analysis (for example, task models and user models) can be developed in a data-driven way. That is, through exploring additional mappings between user interaction data and cognitive processes, this science of interaction can continue to evolve and solidify as theories continue to form.

Other Visual Metaphors and Analytic Models

Semantic interaction research has focused largely on spatializations that show similarity using the relative distance between data objects. Other visual metaphors and representations can be leveraged in a similar, bidirectional nature. In transitioning semantic interaction design guidelines (see the related sidebar) to such metaphors, a critical component is the model used for generating the visualization.

For example, using a sparkline to show the temporal trend of the abundance of a specific term or hashtag on social media might not directly benefit from semantic interaction. This is because the visualization generation doesn't involve model translation (the count of terms is directly visualized). However, if we use an analytic model to determine the trend as an aggregate of terms or

Inferring Reasoning from User Interaction

Two challenges for information visualization are how to gain a deeper understanding of how users interact with visualizations and, what's more important, how these interactions are integrated into their analytic process.[1] To investigate this science of interaction, researchers have used several methodologies. Ji Soo Yi and his colleagues extensively categorized the user interactions available in popular exploratory visualization tools.[2] However, categorizing interaction in visualization is inherently complicated.[3] Wenwen Dou and her colleagues showed that by logging user interactions with a financial-data visualization, they could reconstruct low-level analytical processes.[4] These results indicate that during visual data exploration, a detectable connection exists between low-level user interaction and the users' analytic processes.

References

1. A. Endert et al., "The Human Is the Loop: New Directions for Visual Analytics," *J. Intelligent Information Systems*, Jan. 2014, pp. 1–25.
2. J.S. Yi et al., "Toward a Deeper Understanding of the Role of Interaction in Information Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, 2007, pp. 1224–1231.
3. C. Plaisant et al., "Evaluating Visual Analytics at the 2007 VAST Symposium Contest," *IEEE Computer Graphics and Applications*, vol. 28, no. 2, 2008, pp. 12–21.
4. W. Dou et al., "Recovering Reasoning Processes from User Interactions," *IEEE Computer Graphics and Applications*, vol. 29, no. 3, 2009, pp. 52–61.

hashtags (or even a weighting of some terms as more important than others), we can steer the technique for determining that trend. For example, users could directly manipulate the sparkline's height at specific places to inject their subjective domain expertise around that time or topic.

Extending semantic interaction to additional visual metaphors and encodings is particularly valuable in VA, which typically leverages one or more analytic models to produce a visualization. Such advancements could improve interactivity for exploratory data analysis in visualization tools such as Tableau or Spotfire.

Semantic interaction principles can also apply to models that might not be directly invertible. For example, topic detection models might not have a clean, mathematical inversion. However, augmenting their parameterization is still critical to foster sensemaking and discovery. So, we can couple their parameterization with the visual affordances created in the visualization to enable semantic interaction. For example, we can show topics in a word cloud, in which users can create and adjust the hierarchical topic detection and aggregation methods used by the specific model. We can then infer the reasoning of such visual augmentations and use it to parameterize the model in accordance with the user's domain expertise. Additional such models might include those used for anomaly detection, standard deviation and error, entity extraction, storytelling, and network structure detection.

The work on semantic interaction has presented the visual-analytics community a set of contributions that can seed idea spaces for further work. This work poses challenges to multidisciplinary research projects and institutions to understand couplings between not only cognitive and computational processes of systems being built but also disciplines including human–computer interaction, information visualization, data mining, and statistics. In reflecting on this work in the context of the visual-analytics community's current needs and directions, opportunities exist to continue to promote the usability and effectiveness of systems that enable users to gain insights in impactful domains.

Acknowledgments

I thank all the coauthors, peers, and colleagues who helped make this research possible. Special thanks go to my advisor, Chris North. Without his guidance, advice, and friendship over the years, this research wouldn't have been possible. The US National Science Foundation's Foundations of Data Analysis and Visual Analytics (FODAVA) program largely funded the research described in this article.

References

1. A. Endert, "Semantic Interaction for Visual Analytics: Inferring Analytical Reasoning for Model Steering," PhD dissertation, Dept. Computer Science, Virginia Tech, 2012.
2. C. Andrews, A. Endert, and C. North, "Space to Think: Large, High-Resolution Displays for Sensemaking," *Proc. 2010 SIGCHI Conf. Human Factors in Computing Systems* (CHI 10), 2010, pp. 55–64.
3. A. Endert et al., "The Semantics of Clustering: Analysis of User-Generated Spatializations of Text Documents," *Proc. 2012 Int'l Working Conf. Advanced Visual Interfaces* (AVI 12), 2012.
4. A. Endert et al., "Observation-Level Interaction with Statistical Models for Visual Analytics," *Proc. 2011 IEEE Symp. Visual Analytics Science and Technology* (VAST 11), 2011, pp. 121–130.
5. S.C. Leman et al., "Visual to Parametric Interaction

(V2PI)," *PLOS One*, vol. 8, no. 3, 20 Mar. 2013; www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0050474.
6. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Visual Text Analytics," *Proc. 2012 SIGCHI Conf. Human Factors in Computing Systems* (CHI 12), 2012, pp. 473–482.
7. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Sensemaking: Inferring Analytical Reasoning for Model Steering," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, 2012, pp. 2879–2888.
8. J. Heer, *Prefuse User's Manual*, 2006; http://prefuse.org/doc/manual.
9. B. Shneiderman, "Direct Manipulation: A Step beyond Programming Languages," *Computer*, vol. 16, no. 8, 1983, pp. 57–69.
10. D.H. Jeong et al., "iPCA: An Interactive System for PCA-Based Visual Analytics," *Computer Graphics Forum*, vol. 28, no. 3, 2009, pp. 767–774.
11. E.T. Brown et al., "Dis-Function: Learning Distance Functions Interactively," *Proc. 2012 IEEE Symp. Visual Analytics Science and Technology* (VAST 12), 2012, pp. 83–92.
12. X. Hu et al., "Semantics of Directly Manipulating Spatializations," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 12, 2013, pp. 2052–2059.
13. A. Endert, L. Bradel, and C. North, "Beyond Control Panels: Direct Manipulation for Visual Analytics," *IEEE Computer Graphics and Applications*, vol. 33, no. 4, 2013, pp. 6–13.
14. C. Plaisant et al., "Evaluating Visual Analytics at the 2007 VAST Symposium Contest," *IEEE Computer Graphics and Applications*, vol. 28, no. 2, 2008, pp. 12–21.
15. J. Scholtz et al., "Evaluation of Visual Analytics Environments: The Road to the Visual Analytics Science and Technology Challenge Evaluation Methodology," *Information Visualization*, June 2013.
16. N. Chris, "Toward Measuring Visualization Insight," *IEEE Computer Graphics and Applications*, vol. 26, no. 3, 2006, pp. 6–9.
17. W.A. Pike et al., "The Science of Interaction," *Information Visualization*, vol. 8, no. 4, 2009, pp. 263–274.

Alex Endert *is a visualization researcher at Pacific Northwest National Laboratory. Contact him at alex.endert@pnnl.gov.*

Contact department editor Jim Foley at foley@cc.gatech.edu.