Editor: Theresa-Marie Rhyne

Semantic Interaction: Coupling Cognition and Computation through Usable Interactive Analytics

Alex Endert
Georgia Tech

Remco Chang
Tufts University

Chris North
Virginia Tech

Michelle Zhou
Juji

In June 2014, the Pacific Northwest National Laboratory (PNNL) hosted a workshop on "Semantic Interaction: Coupling Cognition and Computation through Usable Interactive Analytics." The focus of this workshop was user interaction, specifically as it pertains to visual analytics systems. It brought together researchers from visualization, data mining, and human-computer interaction to discuss the current and future challenges for visual analytics.

User interaction has traditionally been a method designed to place a "human in the loop" by augmenting system parameters to explore data and gain insights. However, this workshop focused on the hypothesis that visual analytics systems can leverage user interaction not only as a way for users to directly control the system but also as a way to systematically study and analyze user interactions to learn about the users and their cognitive processes. That is, the goal is to treat user interaction as valuable data from which data and user models can be systematically inferred and used to guide analytic models of the system. This user interaction paradigm is known as *semantic interaction* (see the sidebar for a more detailed definition). As a result, visual analytics systems (including visualization and computation) can provide the external affordances needed to aid in sensemaking.

The continued success of visual analytics is predicated on the ability of users to interactively explore information. Humans think about their data through interactive visual exploration, including testing hypotheses, exploring anomalies, and other cognitive processes of building understanding from data. The claim that these insights are generated as a result of the interaction leads us to posit that user interaction must play a more central role in visual analytics systems, serving as the method for coupling cognition and computation. This report summarizes the discussions generated during the PNNL workshop in the form of claims and design guidelines for the next generation of visual analytics systems.

Visual Analytics and Sensemaking

Visual analytics provides carefully designed combinations of computational and visual affordances to aid cognitive processes of visually exploring and analyzing data. User interaction is fundamental to this process because it enables users to act upon their subject matter expertise by posing questions and hypotheses about the data to the system. Tra-

 Published by the IEEE Computer Society

Semantic Interaction

Semantic interaction is an approach for user interaction tailored to visual analytics systems.[1] By definition, semantic interaction translates interactions performed on high-level visual artifacts of the data into low-level analytic models adjustments. The goal of semantic interaction is to couple cognitive and computational processes through visualization, mapping low-level analytic model adjustments with high-level user interactions. For example, interactions such as highlighting text, organizing documents spatially, note-taking, and searching perform those direct actions, but the system also infers the analytical reasoning embedded in the interactions to reweight the feature space of the document corpus, extract new features, and ultimately calculate a spatialization of documents based on the user and the machine's feedback regarding grouping.[2] As a result, the system's domain experts are abstracted from the low-level parameterization of the underlying analytic model of the system. However, through the inference of the user interactions performed at the high level (that is, native to the domain and within the user's expertise), the power of the computational models can be leveraged without their complexity. For visual data analysis, such an interaction methodology is particularly important because it enables users to act upon their subject matter expertise by posing questions and hypotheses about the data to the system without requiring them to translate their cognitive artifacts into computational actions.

A fundamental component of semantic interaction is the systematic inference performed on the user interactions in order to model specific aspects of the user (interest, task, bias, and so on). For example, studying users analyzing a text dataset using ForceSPIRE, researchers have shown that the adaptation of the model over time serves as a good approximation for the features (such as text terms and phrases) of the dataset that the users reported as interesting.[3] Similarly, data models can also be produced of other high-dimensional datasets via semantic interaction approaches.[4] Also, users can also explore the temporal aspect of data directly within a scatterplot by moving points along their paths, reflecting the change in values temporally.[5]

References

1. A. Endert, "Semantic Interaction for Visual Analytics: Toward Coupling Cognition and Computation," *IEEE Computer Graphics and Applications*, vol. 34, no. 4, 2014, pp. 8–15.
2. A. Endert, P. Fiaux, and C. North, "Semantic Interaction for Visual Text Analytics," *Proc. SIGCHI Conf. Human Factors in Computing* (CHI), 2012, pp. 473–482.
3. 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.
4. E.T. Brown et al., "Dis-Function: Learning Distance Functions Interactively," *Proc. IEEE Conf. Visual Analytics Science and Technology* (VAST), 2012, pp. 83–92.
5. B. Kondo and C. Collins, "DimpVis: Exploring Time-varying Information Visualizations by Direct Manipulation," *IEEE Trans. Visualization and Computer Graphics*, vol. 20, no. 12, 2014, pp. 2003–2012.

ditionally, the principles of direct manipulation were simply applied to such models by using control panels to directly manipulate model parameters.[1] Typically, these principles are applied in the form of a control panel, containing visual widgets such as sliders, buttons, or query fields that are coupled with the parameters of a visual representation in the main view. For the purpose of interactive machine learning, these interfaces provide feedback in an expressive and formal way (such as parameters to the algorithms or for standard training and labeling tasks).

However, for users and their analytic tasks, these interactions may present significant usability issues because they force them out of their cognitive flow or zone and may place fundamental limitations on sensemaking activities. Reasoning about data is an inherently cognitive activity, where the mental artifacts that we leverage to reason can manifest themselves at different semantic and symbolic levels of detail. Thus, a gap exists between the cognitive constructs and manipulations humans employ to think and reason about information and the interactive affordances that user interfaces offer them. Exploiting humans merely as data labelers or parameter tuners misuses human expertise and skills, forcing us to adapt to formal algorithmic methods and a priori parameter specifications when our strengths are in incremental informal reasoning. More importantly, it misses a major opportunity for the potential benefits of coupling cognition and computation.

We contend that a new methodology to couple the cognitive and computational components of such systems is necessary. Therefore, we need a science of interaction to enable visual analytics systems to couple cognitive and computational processes. We suggest semantic interaction as a potential user interaction paradigm for visual analytics.

Interaction Methodologies for Visual Analytics

As a result of the discussions at the PNNL workshop, we contend that there are two categories of

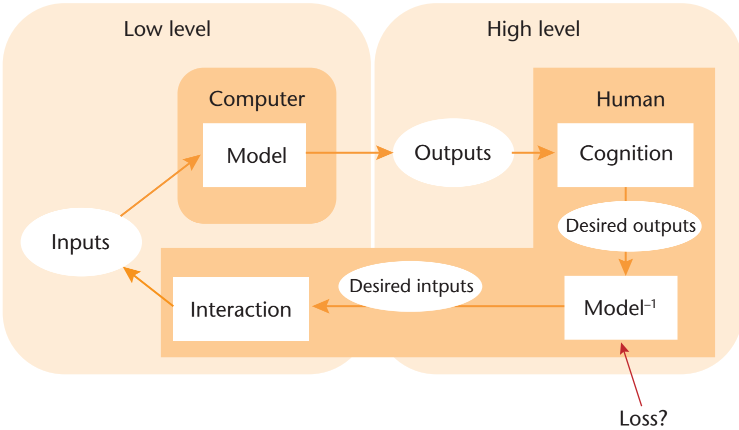

Figure 1. Model of a direct manipulation interface. Users are responsible for translating their desired output and thus must understand how to augment the model. This is a potentially lossy process.

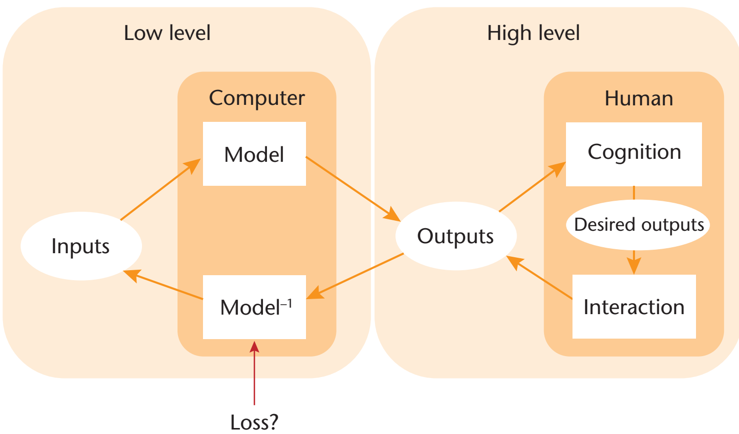

Figure 2. Model of a semantic interaction interface. The interaction is designed to allow users to maintain a higher level of cognitive state. The translation happens on the computer side.

user interaction design with which the user can communicate with the system: direct manipulation and semantic interaction. The differences in the principles between these two interaction methodologies can be found in earlier work;[2] for the purposes of this article, we focus on the implications that the choice of design has on the user's analytic discourse.

In general, data and analytic models transform low-level data characteristics into high-level concepts or topics. In visual analytics, these are typically communicated to the user via visualization. As a result, humans using the system can reason about the information at a higher level. Furthermore, the act of sensemaking is an interactive process, where thinking is supported by both static visuals and the interactive affordances of the tool. As such, users have the choice of interacting with—and thus thinking about—their data with either direct manipulation or semantic interaction (or both, depending on the system design).

Figure 1 models this analytic process when direct manipulation interfaces are used. Direct manipulation interfaces allow users to engage in analytic discourse with their data through graphical affordances that control various attributes of the analytic model and visualization (such as sliders to filter values and menus to change visual encodings).[1] Direct manipulation specifies the following three properties for interaction design for information visualization:

- continuous representation of the object of interest,
- physical actions or labeled button presses instead of complex syntax, and
- rapid incremental reversible operations whose impact on the object of interest is immediately visible.[1]

By design, this requires users to translate their cognitive artifacts and goals into low-level interactions in the system. This translation could be lossy depending on the usability issues that could arise out of a mismatch of domains during this translation—that is, when the users perform the translation from their goals to their plan of action (or interaction, in this case).

In comparison, Figure 2 illustrates a semantic interaction interface. Semantic interaction uses interaction as data about the users, their processes, and their subject matter expertise. This data source is then systematically analyzed to learn data and user models, from which visual analytics systems can adapt accordingly. Semantic interaction interfaces produce this coupling by leveraging the visual metaphor as the mapping function and the visual encoding as the interaction interface by which users perform their visual data exploration.

The fundamental difference between direct manipulation and semantic interaction interfaces is that, in the latter, the interactions are designed on the higher-level artifacts of the visualization and the system is responsible for performing the translation into the adaptations to the underlying model(s) it uses. These could include user, data, and other models used for adaptive computation. This translation is likely to incur some signal loss as well. However, we hypothesize that, for some applications leveraging complex analytic models, the loss is reduced over the direct manipulation alternative. Maintaining the user interaction at the higher level can also lead to the inference of the user's higher-level cognitive processes for the purpose of visual analytics systems.

Design Principles for Visual Analytics Systems

An outcome of the PNNL workshop was an understanding of the design guidelines for visual analytics systems that seek to use semantic interaction:

1. Systems must be able to capture and understand users' actions.
2. Systems must be able to make inferences based on users' interactions.
3. Systems must be able to react and take initiative based on the inferences at three levels: interface, computation, and cognitive.

This section summarizes the discussions that resulted in these guidelines. They include claims from the workshop as well as supporting discussion and select literature to support the claims.

1. Capturing User Actions

> *Claim*: We can capture and store a stream of user interaction data, including both virtual interactions and physical actions.

User interactions are a quantitative realization of a user's cognitive processes during data exploration. As such, the interaction logs captured during analysis can be used as data. This interaction data about the user can be captured from two sources: virtual interactions and physical actions.

Logs of virtual interactions are produced from the interactions that users perform within a user interface. Prior work has studied user interaction in this context. For example, Wenwen Dou and his colleagues showed that interaction logs can be analyzed to understand specific low-level analytical processes of the users who performed them.[3] Most importantly, these results indicate that a detectable connection exists between the low-level user interaction and the high-level analytic processes of users when it comes to visual data exploration. Fundamentally, the understanding of how user interaction fits into the exploratory process frames the science of interaction.[4]

The physical environment in which users perform their analysis can also be instrumented with various sensors, which can populate a dataset of physical actions or attributes that humans exhibit while analyzing data. For example, visualizations on large displays let users physically navigate the visualization (that is, panning and zooming is replaced with walking).[5] These physical movements can be captured by various motion-tracking devices outfitted in the room. Furthermore, the analysis of these physical navigation actions and strategies can reveal the effectiveness of the visualization design decisions. These, as well as other physiological measures (such as EEG, functional MRI [fMRI], and functional near-infrared spectroscopy [fNIRS]), can increase the amount of information about a user that can be modeled and ultimately recast into interactions with mixed-initiative analytics systems.[6]

2. Inferring Analytical Reasoning

> *Claim*: The analytic reasoning is embedded in the user interaction data and can be systematically analyzed to recreate facets of the user's reasoning processes.

Building upon claim 1, the participants of the PNNL workshop contend that there are artifacts of the user's reasoning process that can be reverse-engineered from the user interaction logs. These methods (which are not mutually exclusive) include data models, user models, task models, role-based models, and more.

User interactions are a quantitative realization of a user's cognitive processes during data exploration.

One area that this discussion draws from is user modeling, where the goal is to reconstruct a user's relevant profile by analyzing their interactions with a complex visualization tool. For example, Eli Brown and his colleagues demonstrated that a user's performance during a visual search task, as well as aspects of a user's personality profile, can be inferred and predicted in real time.[7] Similarly, biometrics used for authenticating the user can be inferred from the person's raw mouse movements, suggesting a connection between the movements we produce and our holistic cognitive and motor systems as individuals. Also, studying eye-tracking data can reflect a user's personality traits and cognitive abilities. For example, the System U tool can automatically identify users' full personality profiles by examining as little as 200 of their Twitter postings.[8]

Data modeling can also be leveraged and might consist of a weighting of data features applied in a weighted dimensionality reduction algorithm. Such a model can more accurately represent the approximation of the data with respect to the user's expertise in the domain. Thus, instead of requiring users to directly manipulate the input weighting of features, semantic interaction enables users to manipulate the output visualization of the information (the visualization itself). The weighting of data features and dimensions can be inferred from these interactions within the visual metaphor and encoding generated algorithmically. Within a spatial metaphor, prior work has shown how user interactions such as reorganizing the information directly within the spatial layout can be used to infer the weighting of the features that make up the high-dimensional data, such as observation-level interaction (OLI)[9] and Dis-Function.[10]

Many other models have been explored in the visual analytics community to characterize and categorize user interactions and tasks.

The science of interaction is critical to the continued success of visual analytics.

3. Machine-Initiated Action

> *Claim*: The system can take initiative based on the inferred analytical reasoning.

In defining future directions of visual analytics, the PNNL workshop discussed how to leverage the rich information within a system. That is, how can the system become an analytically valuable collaborator given what it knows about the user? We summarize that there are three types of adaptive system behaviors:

- interface-level actions are visible and "in line" with what users' needs or expects,
- computation-level actions are those that users might not be aware of but that will assist their analysis process, and
- cognitive-level actions steer or guide the users' analytic processes.

Adaptive user interfaces and visualization systems have been an important research topic in the HCI community. Interfaces such as SUPPLE have demonstrated that a system can learn a user's motor disabilities or the limitations of the device (such as smartphone and tablet) and then automatically adapt the size and positioning of UI elements to generate an interface that is optimal to the user and the device.[11] Moving beyond visual-level adaptations, systems can also adapt based on the amount of information presented to the user.

In addition to interface-level adaptation, adaptive analytics algorithms and computation systems can also benefit from having knowledge about the user's analysis processes. For example, user-guided computation that leverages knowledge of the user's analysis process and goals can lead to advances in efficient, online algorithms that compute only the information the user needs. As the user explores the data, these algorithms can incrementally increase (decrease) detail by incorporating more (less) data. Such an analytic engine can maintain a small memory footprint while providing the user with rich information throughout the user's exploration process.

Lastly, we posit that there are adaptive systems that can be aware of the user's cognitive states, abilities, and limitations and appropriately supplement the user's analysis accordingly. This adaptive behavior may often be contradictory to the user's analytic process or strategy, but it may benefit the final analysis outcomes nonetheless. For example, if a system can detect that a user is suffering from confirmation bias during an investigation, it should not continue to seek supporting evidence but should instead challenge the user's beliefs by presenting alternate, possibly conflicting information and hypothesis. For example, if a user is focusing an investigation on a large document corpus on one specific theme or topic (as detected by the system), an adaptive system may recommend alternative hypotheses or evidence to support more than one theme. Furthermore, as users seek evidence to support one hypothesis, the system could collect refuting evidence on behalf of the users and present it at an opportune monument. Incorporating the knowledge of cognitive processes into adaptive visualizations has not been extensively investigated by the visual analytics community, and it remains one of the open challenges outlined by James Thomas and Kristin Cook.[12]

The science of interaction is critical to the continued success of visual analytics. For visual data analysis, the participants of the PNNL workshop contend that user interaction is as important to study as the visualization of the information itself or the analytic models being developed and optimized. The field of information visualization has made great strides in understanding methods for visualizing data characteristics, attributes, relationships, and so on. The next challenges for visual analytics reside in understanding how these visual representations can be used as external, digital aids for sensemaking and thus how they can afford reasoning and interaction. This workshop illuminated a number of such challenges, which we pose as open problems to both the visualization community and beyond.

The participants of this workshop were primarily from the visualization community, and we recognize that the visualization community cannot realize the entire research agenda of developing a science of interactions alone. As data size and complexity continue to increase and as data and visual analytics systems become more sophisticated, continued engagement and involvement from disciplines including the cognitive sciences, data sciences, and mathematics will become even more important as we strive to address these open challenges and make high-impact breakthroughs.

Acknowledgments

This workshop was sponsored by and held at PNNL, which is managed for the US Department of Energy by Battelle under contract DE-AC05-76RL01830. This work was partially supported by US National Science Foundation grants IIS-1218346 and IIS-1218170. We also thank the attendees for their contributions to these ideas: Alex Endert, Christopher Collins, Polo Chau, Eric Alexander, Wenwen Dou, Dave Gunning, Jaegul Choo, John Stasko, Jordan Crouser, Barry Drake, Chris North, Michelle X Zhou, Remco Chang, Bill Pike, and Mark Greaves.

References

1. B. Shneiderman, "Direct Manipulation: A Step Beyond Programming Languages," *Computer*, vol. 16, no. 8, 1983, pp. 57–69.
2. 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.
3. W. Dou et al., "Recovering Reasoning Processes from User Interactions," *IEEE Computer Graphics and Applications*, vol. 29, 2009, pp. 52–61.
4. W.A. Pike et al., "The Science of Interaction," *Information Visualization*, vol. 8, no. 4, 2009, pp. 263–274.
5. R. Ball, C. North, and D.A. Bowman, "Move to Improve: Promoting Physical Navigation to Increase User Performance with Large Displays," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI), 2007, pp. 191–200.
6. E. Solovey et al., "Brainput: Enhancing Interactive Systems with Streaming Fnirs Brain Input," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2012, pp. 2193–2202.
7. E.T. Brown et al., "Finding Waldo: Learning about Users from their Interactions," *IEEE Trans. Visualization and Computer Graphics*, vol. 20, no. 12, 2014, pp. 1663–1672.
8. L. Gou, M.X. Zhou, and H. Yang, "KnowMe and ShareMe: Understanding Automatically Discovered Personality Traits from Social Media and User Sharing Preferences," *Proc. 32nd Ann. ACM Conf. Human Factors in Computing Systems*, 2014, pp. 955–964.
9. A. Endert et al., "Observation-Level Interaction with Statistical Models for Visual Analytics," *Proc. IEEE Conf. Visual Analytics Science and Technology* (VAST), 2011, pp. 121–130.
10. E.T. Brown et al., "Dis-Function: Learning Distance Functions Interactively," *Proc. IEEE Conf. Visual Analytics Science and Technology* (VAST), 2012, pp. 83–92.
11. K.Z. Gajos, J.O. Wobbrock, and D.S. Weld, "Improving the Performance of Motor-Impaired Users with Automatically-Generated, Ability-Based Interfaces," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2008, pp. 1257–1266.
12. J.J. Thomas and K.A. Cook, eds., *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, National Visualization and Analytics Center, 2005.

***Alex Endert** is an assistant professor in the School of Interactive Computing at Georgia Tech. Contact him at endert@gatech.edu.*

***Remco Chang** is an assistant professor at Tufts University. Contact him at remco@cs.tufts.edu.*

***Chris North** is a professor at Virginia Tech. Contact him at north@vt.edu..*

***Michelle Zhou** was at IBM Research when she co-organized the workshop, and she is now a cofounder and CEO of Juji. Contact her at mzhou@acm.org.*

Contact department editor Theresa-Marie Rhyne at theresamarierhyne@gmail.com.