Toward Usable Interactive Analytics: Coupling Cognition and Computation

Alex Endert
Pacific NW National Lab
Richland, WA USA
alex.endert@pnnl.gov

Chris North
Virginia Tech
Blacksburg, VA USA
north@cs.vt.edu

Remco Chang
Tufts University
Medford, MA USA
remco@cs.tufts.edu

Michelle Zhou
IBM Research
Almaden, CA USA
mzhou@us.ibm.com

ABSTRACT
Interactive analytics provide users a myriad of computational means to aid in extracting meaningful information from large and complex datasets. Much prior work focuses either on advancing the capabilities of machine-centric approaches by the data mining and machine learning communities, or human-driven methods by the visualization and CHI communities. However, these methods do not yet support a true human-machine symbiotic relationship where users and machines work together collaboratively and adapt to each other to advance an interactive analytic process. In this paper we discuss some of the inherent issues, outlining what we believe are the steps toward usable interactive analytics that will ultimately increase the effectiveness for both humans and computers to produce insights.

1. INTRODUCTION
To tackle the onset of big data, visual analytics seeks to marry the human-intuition of visualization with the analytical horsepower of mathematical models. Yet, a critical open question is how humans will interact with, steer, and train these complex mathematical models.

The visual analytics community has worked to provide visual representations of data, as approximated by complex models and analytics [34]. User interaction is critical to the success of such visual data exploration, as it allows users to engage in a process of testing assertions, assumptions, and hypotheses about the information given one's prior knowledge about the world. This cognitive process can be generally referred to as sensemaking. Visual analytics emphasizes sensemaking of large, complex datasets through interactively exploring visualizations generated via a combination of analytic models. Thus, a central focus is understanding how to leverage human cognition in concert with powerful computation through usable visual metaphors.

Initially, the principles of direct manipulation were applied to such models in a simplistic fashion by using control panels to directly manipulate model parameters. Direct manipulation specifies the following three properties for interaction design for information visualization: (1) continuous representation of the object of interest, (2) physical actions or labeled button presses instead of complex syntax, and (3) rapid incremental reversible operations whose impact on the object of interest is immediately visible [31]. Typically, these principles are applied in the form of a control panel, containing visual widgets such as sliders, buttons, or query fields, coupled to 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 (e.g., standard training and labeling tasks).

However, for users and their analytic tasks, these interactions may present significant usability issues by forcing the user out of their cognitive flow or zone [11,22], and may place fundamental limitations on sensemaking activity due to lack of recognition of the depth of interactions which humans apply in their cognitive processes. Exploiting humans merely as data labelers or parameter tuners mis-uses human expertise and skills, forcing humans to adapt to formal algorithmic methods and apriori parameter specifications, when their 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 needed. We suggest *Semantic Interaction* as a potential solution concept, which attempts to bridge these components by binding the user interactions used for visual sensemaking with the training of machine learning techniques [17]. Semantic interaction interfaces produce this coupling by leveraging the visual metaphor as the mapping function, and the visual encoding as the interactive affordance by which users perform their visual data exploration. In this paper we discuss the concept of semantic interaction as a method for systematically learning characteristics about a user and his or her reasoning process, adapting the underlying analytic model, and increasing the usability of incorporating the human in the loop.

2. SEMANTIC INTERACTION
Semantic interaction is an approach to user interaction for visual analytics in which the user's analytical reasoning is inferred and in turn used to steer the underlying models implicitly. The goal of this approach is to enable co-reasoning between the human and the analytic models (coupling cognition and computation) without requiring the user to directly control the models and parameters. This co-reasoning occurs through mutual interaction with a visual medium of communication – the visualization or visual metaphor.

The approach of semantic interaction is to overload the metaphor through which the insights are obtained (i.e., the visualization of information created by computational models) and the interaction metaphor through which hypotheses and assertions are communicated (i.e., interaction occurs within the visual metaphor). Semantic interaction enables users to directly manipulate data within visualizations, from which tacit knowledge about the user is captured, and the underlying analytic models are steered. The analytic models can be incrementally adapted based on the user's incremental sensemaking process and domain expertise explicated via the user interactions with the system. The specifics of the system could include multiple visual metaphors used in concert.

That is, the parameters of the underlying analytic models are exposed through the visual constructs of the visualization. Based on common visual metaphors (such as the geographic, spatial metaphor where proximity approximates similarity), tacit knowledge of the user's reasoning can be inferred through inverting these analytic models. As a result, users are shielded from the underlying complexities, and able to interact with their data through a bi-directional visual medium. The interactions users perform within the visualizations to augment the visual encodings within the metaphor enable the inference of their analytic reasoning, which are systematically applied to the underlying models. The visual metaphor helps define the mapping between the model parameters and the visualization, and the visual encoding provides the visual interactive affordance by which users can interact. Thus, the process of visual data exploration and models steering occur on the same set and sequence of interactions.

The semantic interaction pipeline (shown in Figure 1) takes an approach of directly binding 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 (visually depicted as distance in a spatial layout), can serve as the interactive affordance to allow users to explore that relationship. Therefore, the user interaction is directly in the visual metaphor, creating a bi-directional medium between the user and the analytic models. This method of user interaction is also similar to the "by example" method of interaction, as users can directly show their intention using the structure of the visualization. This adds to the role of visualization in the reasoning process, in that it is no longer intended to be solely a method for gaining insight, but also one for directly interacting with the information and the system. The bi-directionality afforded by semantic interaction comes via binding the parameter controls traditionally afforded by the GUI directly within the visual metaphor. It is through this binding that an inference can be made about the user's analytic reasoning from the user interaction with the visualization with regards to the parameters of the underlying mathematical model.

For example, a spatial layout is one specific visual metaphor where existing research on semantic interaction has been conducted, described in [14,15,16]. The spatial visual metaphor (i.e., a spatialization) is one where the bi-directionality afforded by semantic interaction has been demonstrated. A spatial metaphor lends itself well 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 dimension reduction models including: principal-component analysis, multi-dimensional scaling, force-directed layouts, etc. In general, these models attempt to approximate the distance between data objects in their true, high-dimensional representation using a smaller number of dimensions (e.g. two dimensions in the case of spatial visualization).

Prior work has applied semantic interaction methods to this visual metaphor. For example, inverting multi-dimensional scaling, principal-component analysis, and generative topographic mapping can enable bi-directional spatializations to afford semantic interaction [4,16]. The ability to understand the parameters of each of the models that can be exposed through the visual encoding (in this case, relative distance between data points) enabled this affordance. Further work 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 [24,27].

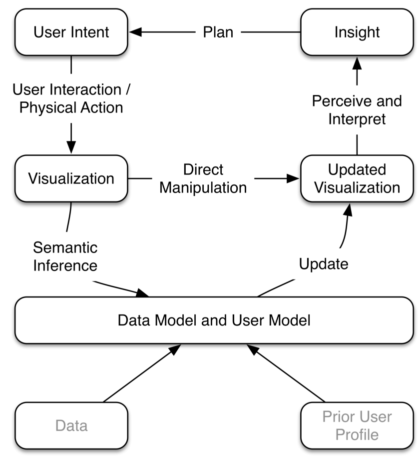

Figure 1 A generalizable model for coupling cognition and computation. Plans generate intents that are externalized by users via interactions and physical actions. Data and user models can be inferred from these actions, and used to update a visualization to continue the analytic process.

3. RESEARCH AGENDA

Based on the promising initial results of current research on semantic interaction for visual analytics, the sections below describe open areas of research to advance the field in usable interactive analytics. These sections describe current work in each topic, as well as illuminate open areas of research that can advance the goal of creating usable interactive analytics via semantic interaction. The areas of research can be depicted in a generalizable model for semantic interaction interfaces, shown in Figure 1.

3.1 Sensing and Capturing User Interaction

Semantic interaction interfaces are grounded in the concept of treating user interaction as data from which models about the user are created. This interaction data about the user can be captured from two categories of sources: virtual interactions, and physical actions.

Virtual interactions refer to those that a user performs within a user interface. These have been previously studied for the purposes of understanding the user. For example, Yi et al. presented an extensive categorization of user interactions available in popular exploratory visualization tools [35]. Further, Dou et al. have shown that through logging user interactions in a visualization of financial data, low-level analytical processes can be reconstructed [9,26]. 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. The advancement of understanding how processes and knowledge from users manifest in user interaction forms the *science of interaction* [29].

The physical actions or attributes that humans exhibit while analyzing data may also provide cues from which models can be generated and adapted. For example, research has shown that navigating large information spaces using physical navigation with large displays is significantly advantageous over virtual navigation with small displays [2]. These physical actions, or strategies for interacting, can also be analyzed to identify effectiveness of analytic strategies on such displays [12]. For example, the sensing of office chair rotation relative to a large display can provide an approximation of the user's primary focus of attention [13]. These, as well as other physiological measures, such as EEG, fNIRS, and fMRI, can increase the amount of information about a user that can be modeled, and ultimately re-cast into interactions with mixed-initiative analytics systems [1,28,32].

Open questions within this topic include:

- What additional visual metaphors and user interfaces can be sources of user interaction data to add breadth to the science of interaction?
- How can the directness of the virtual interactions (with respect to the interface and task) be coupled with the passiveness of the physical actions? What are the tradeoffs between the passive sensing of physical actions and the direct sensing of virtual interactions?

3.2 Inferring User Models

As visualization systems become more complex, so do the user's ability to express their reasoning process through these complex interfaces. These reasoning processes reflect a user's cognitive abilities [7] and personality traits [36], and are often influenced by the user's cognitive and mental state (such as emotion and cognitive load) [23,28].

The research goal of User Modeling is to reconstruct the relevant profile of a user by analyzing their interactions with a complex visualization tool. For example, Brown et al. 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 [5]. Similarly, the physical motions of a user's mouse movement have been shown to be effective as biometrics to authentic a user's identity for security purposes [30].

Beyond analyzing a user's virtual interactions (mouse and keyboard interactions), other user-generated data has also been used to infer models of a user. For example, eye-tracking data has been shown to reflect a user's cognitive abilities and personality traits [33]. More broadly, Gou et al. developed a tool called *System U* that can automatically identify a user's full personality profile by examining as little as two hundred of the user's Twitter postings [21]. These user modeling techniques give rise to the possibility of mixed-initiative visual analytics systems in which the computer can understand and support the user's analysis needs in real time [34].

Open research areas include:

- What other forms of models can be inferred, steered, and created (e.g., task models, role-based models, etc.)?
- How can we detect artifacts of cognitive processes that may be less desired (e.g., forms of bias, cognitive depletion, etc.)?

3.3 Inferring Data Models

Semantic interaction interfaces can implicitly map to, train, and steer underlying data models. One method to do this is to enable users to manipulate the output of the model, and then computationally invert the model to learn optimized inputs that would produce the desired outputs.

For example, a data model might consist of a weighting of data features applied in a weighted dimensionality reduction algorithm. Instead of requiring users to directly manipulate the input weighting of features, semantic interaction enables users to manipulate the output visualization of the information, from which the weighting of features can be inferred. Prior work has shown how such user interactions (e.g., re-organizing data within a spatial layout) can map to the weighting of features, in tools such as OLI [4,16,27] and Dis-Function [4]. Term weights for text analytcs can be learned from users interactions with spatial organizations of documents, highlighting, annotations, reading patterns, eye gaze, etc. in ForceSPIRE [14,15] and StarSpire [3]. iCluster demonstates learning of a document clustering model through users' incremental cluster membership choices [10]. Apolo demonstrates learning network belief-propagation models through textual sensemaking interactions [6].

Open questions include:

- What additional data processing models can be steered or created?
- How can we consider models that function on different scales of data (i.e., from overview to detail, but also from detail to meaningful context)?

3.4 Adaptive Visualization

Techniques for User and Data Modeling would inform the visualization and the analytics system's high-level information about the user's analysis goals and needs. Responding to these inputs, the visualization system can adapt the information and representation presented to the user. Similarly, the analytics engine can also modify its behavior to achieve more efficient and accurate analysis results (see Section 3.5).

Adaptive user interfaces and visualization systems have been an important research topic in HCI. Interfaces such as SUPPLE have demonstrated that a system can learn a user's motor disabilities (such as Parkinson's) or the limitations of the device (such as smart phone and tablet), and automatically adapt the size and positioning of UI elements to generate a user interface that is optimal to the user and the device [20].

In adaptive visualization, researchers have examined the relationship between visual metaphors and the user's personality traits [23,33,36]. Moving beyond interface-level adaptations, systems have also adapted based on the amount of information presented to the user. In the context of games and training, these types of adaptations are often referred to as "dynamic difficulty" adjustments [25], but the same techniques have been more broadly applied to real-world scenarios such as assisting operators of unmanned vehicles and robots [1,8,32].

Open questions include:

- How do we ensure that, in mixed-initiative systems, both the system and the user have equal opportunities to provide feedback?

- How do we ensure the system responds in a way that amplifies the cognitive processes, and aids them, instead of deteriorating the performance of the person?

3.5 Adaptive Computation

In addition to user-level adaptation, analytics algorithms and systems can also benefit from having knowledge about the user's cognitive style and analysis processes.

As datasets get larger, it becomes increasingly difficult for visualizations and analytic systems to provide both interactivity and complete data analysis simultaneously. The old design adage of "Overview First, Details on Demand" limits the size of the data that a visual analytics system can support at an interactive rate. The "Big Data" challenge requires new computational techniques and paradigm shifts. In the case of visual analytics, one potentially rich and fruitful approach is to integrate User and Data Modeling into novel adaptive computational techniques.

"Approximate computing" can generate an overview of a large dataset in real-time. Approximate computing will, by definition, be less accurate than traditional statistical or machine learning techniques, but will deliver sufficient information for the user to perceive high-level patterns within the data in a fraction of the time. Some plausible factors for approximate computing include the consideration of human perception properties such as just noticeable difference (JND), or cognitive limitations based on attention, working memory capacity, or cognitive load [18].

In addition, "user-guided computation" that leverages knowledge of the user's analysis process and goals can lead to advancements in efficient, online algorithms that compute only the information needed by the user. As the user explores the data, these algorithms can incrementally increase (or decrease) in detail by incorporating more (or 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 [19].

Open questions include:

- How do we perform model selection over a set of models has been created, selecting the one (or combination) that is most appropriate given the context of the analysis?
- What are other forms of models or computation that lend themselves to the semantic interaction methods outlined in this position statement?

4. CONCLUSION

Achieving effective coupling of cognition and computation for interactive analytics will require significant research attention towards usability and interaction issues. Clearly, we must go well beyond existing simple human-in-the-loop methods. We have outlined a research agenda that we believe will be critical to enabling insight in the big data era.

ACKNOWLEDGEMENTS

PNNL is managed for the US Department of Energy by Battelle under Contract DE-AC05-76RL01830.

REFERENCES

1. Afergan, D., Peck, E.M., Solovey, E.T., et al. Dynamic Difficulty Using Brain Metrics of Workload. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2014), 3797–3806.
2. Ball, R., North, C., and A. Bowman, D. Move to improve: promoting physical navigation to increase user performance with large displays. *ACM CHI*, ACM (2007).
3. Bradel, L. and North, C. StarSpire: Multi-scale Semantic Interaction. *IEEE VAST*, (2014).
4. Brown, E.T., Liu, J., Brodley, C.E., and Chang, R. Dis-function: Learning Distance Functions Interactively. *IEEE VAST*, (2012).
5. Brown, E.T., Ottley, A., Zhao, J., et al. Finding Waldo: Learning about Users from their Interactions. *Trans. Vis. Comput. Graph*, (2014).
6. Chau, D.H., Kittur, A., Hong, J.I., and Faloutsos, C. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2011), 167–176.
7. Chen, C. Individual differences in a spatial-semantic virtual environment. *Journal of the American Society for Information Science 51*, 6 (2000), 529–542.
8. Donmez, B., Nehme, C., and Cummings, M.L. Modeling Workload Impact in Multiple Unmanned Vehicle Supervisory Control. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans 40*, 6 (2010), 1180–1190.
9. Dou, W., Jeong, D.H., Stukes, F., Ribarsky, W., Lipford, H.R., and Chang, R. Recovering Reasoning Processes from User Interactions. *IEEE Computer Graphics and Applications 29*, (2009), 52–61.
10. Drucker, S.M., Fisher, D., and Basu, S. Helping users sort faster with adaptive machine learning recommendations. Springer-Verlag (2011), 187–203.
11. Elmqvist, N., Moere, A.V., Jetter, H.-C., Cernea, D., Reiterer, H., and Jankun-Kelly, T. Fluid interaction for information visualization. *Information Visualization 10*, (2011), 327–340.
12. Endert, A., Andrews, C., and North, C. Visual Encodings that Support Physical Navigation on Large Displays. *Graphics Interface*, (2011).
13. Endert, A., Fiaux, P., Chung, H., Stewart, M., Andrews, C., and North, C. ChairMouse: leveraging natural chair rotation for cursor navigation on large, high-resolution displays. ACM (2011), 571–580.
14. Endert, A., Fiaux, P., and North, C. Semantic Interaction for Visual Text Analytics. *ACM Human Factors in Computing (CHI)*, (2012).
15. Endert, A., Fiaux, P., and North, C. Semantic Interaction for Sensemaking: Inferring Analytical Reasoning for Model Steering. *Visualization and Computer Graphics, IEEE Transactions on 18*, (2012), 2879–2888.
16. Endert, A., Han, C., Maiti, D., House, L., Leman, S.C., and North, C. Observation-level Interaction with Statistical Models for Visual Analytics. *IEEE VAST*, (2011), 121–130.
17. Endert, A. Semantic Interaction for Visual Analytics: Inferring Analytical Reasoning for Model Steering. 2012. http://scholar.lib.vt.edu/theses/available/etd-07112012-123927/.
18. Fisher, D., Drucker, S.M., and König, A.C. Exploratory Visualization Involving Incremental, Approximate Database Queries and Uncertainty. *IEEE Computer Graphics and Applications 32*, 4 (2012), 55–62.
19. Fisher, D., Popov, I., Drucker, S., and schraefel, m. c. Trust Me, I'M Partially Right: Incremental Visualization Lets Analysts Explore Large Datasets Faster. *Proceedings of the*

SIGCHI Conference on Human Factors in Computing Systems, ACM (2012), 1673–1682.
20. Gajos, K.Z., Wobbrock, J.O., and Weld, D.S. Improving the Performance of Motor-impaired Users with Automatically-generated, Ability-based Interfaces. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2008), 1257–1266.
21. Gou, L., Zhou, M.X., and Yang, H. KnowMe and ShareMe: Understanding Automatically Discovered Personality Traits from Social Media and User Sharing Preferences. *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, ACM (2014), 955–964.
22. Green, T.M., Ribarsky, W., and Fisher, B. Building and applying a human cognition model for visual analytics. *Information Visualization 8*, (2009), 1–13.
23. Harrison, L., Skau, D., Franconeri, S., Lu, A., and Chang, R. Influencing Visual Judgment Through Affective Priming. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 2949–2958.
24. Hu, X., Bradel, L., Maiti, D., House, L., North, C., and Leman, S. Semantics of Directly Manipulating Spatializations. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2052–2059.
25. Hunicke, R. and Chapman, V. AI for dynamic difficulty adjustment in games. .
26. Jeong, D.H., Dou, W., Lipford, H.R., Stukes, F., Chang, R., and Ribarsky, W. Evaluating the relationship between user interaction and financial visual analysis. *IEEE Symposium on Visual Analytics Science and Technology, 2008. VAST '08*, (2008), 83–90.
27. Leman, S.C., House, L., Maiti, D., Endert, A., and North, C. A Bi-directional Visualization Pipeline that Enables Visual to Parametric Interaction (V2PI). *PLOS One*, (2011).
28. Peck, E.M.M., Yuksel, B.F., Ottley, A., Jacob, R.J.K., and Chang, R. Using fNIRS Brain Sensing to Evaluate Information Visualization Interfaces. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 473–482.
29. Pike, W.A., Stasko, J., Chang, R., and O'Connell, T.A. The science of interaction. *Information Visualization 8*, (2009), 263–274.
30. Pusara, M. and Brodley, C.E. User Re-authentication via Mouse Movements. *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, ACM (2004), 1–8.
31. Shneiderman, B. Direct Manipulation: A Step Beyond Programming Languages. *Computer 16*, 8 (1983), 57–69.
32. Solovey, E., Schermerhorn, P., Scheutz, M., Sassaroli, A., Fantini, S., and Jacob, R. Brainput: Enhancing Interactive Systems with Streaming Fnirs Brain Input. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 2193–2202.
33. Steichen, B., Carenini, G., and Conati, C. User-adaptive Information Visualization: Using Eye Gaze Data to Infer Visualization Tasks and User Cognitive Abilities. *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, ACM (2013), 317–328.
34. Thomas, J.J. and Cook, K.A. Illuminating the path. 2005. http://nvac.pnl.gov/agenda.stm#book.
35. Yi, J.S., Kang, Y. ah, Stasko, J., and Jacko, J. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics 13*, (2007), 1224–1231.
36. Ziemkiewicz, C., Ottley, A., Crouser, R.J., et al. How Visualization Layout Relates to Locus of Control and Other Personality Factors. *IEEE Transactions on Visualization and Computer Graphics 19*, 7 (2013), 1109–1121.