Tracking Personal Histories for Knowledge Discovery Tasks

Dennis P. Groth
School of Informatics
Bloomington, Indiana 47408-3912
dgroth@indiana.edu

ABSTRACT

Interactive visualizations provide an ideal setting for exploring the use and exploitation of personal histories. Even though visualizations leverage innate human capabilities for recognizing interesting aspects of data, it is unlikely that two users will follow the exact process for discovery. This results in an inability to effectively recreate the exact conditions of the discovery process, which we call the knowledge rediscovery problem. Because we cannot expect a user to fully document each of their interactions, there is a need for visualization systems to maintain user trace data in a way that enhances a user's ability to communicate what they found to be interesting, as well as how they found it. This project presents a model for representing user interactions that articulates with a corresponding set of annotations, or observations that are made during the exploration. This problem is only made more challenging when pervasive computing and corresponding interactions across devices is factored in.

Keyword

Provenance, history tracking, annotation, collaboration.

INTRODUCTION

Visualization research is frequently presented in terms of a graphic image of a visual representation, along with a verbal description of what the observer should recognize. While this traditional approach in reporting results is necessary and meaningful, it is important to note that, in presenting the result in such a way, the researcher is not reporting a visualization. Rather, the researcher is reporting a presentation graphic. While it is necessary to generate a presentation graphic to report the results, the graphic is substantially inadequate for other researchers or practitioners to apply the results due to the static nature of the information presented in this form.

The final presentation of the visualization output results from a stream of actions performed against the data. A common view of the transformation process is represented by the visualization pipeline, as shown in Figure 1. Under this model, data may be transformed through any number of processes prior to display. The output then may undergo an indeterminate number of user specified view transformations. If we assume a 3D representation, the most basic view transformations are rotation, translation and scaling, with others supported by specific applications.[3]

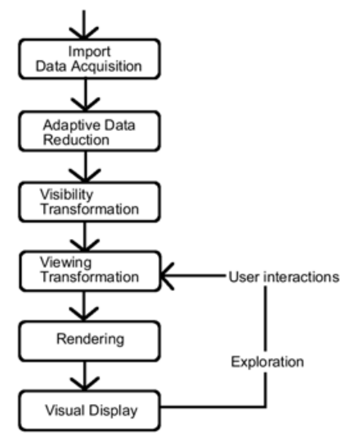

Figure 1: Visualization pipeline.[15]

The main contribution of this research is a conceptual model of user interaction and observation for data visualization. The model is *generic*, in the sense that it concisely captures changes to the state of the visualization made by the user in a way that provides recall of steps the user took to achieve the visual representation. In addition to the conceptual model we describe an instantiation of the model, demonstrating how the model can be adapted to support a variety of modalities of interaction tracking. A prototype implementation of the model is used to demonstrate how the model can be used to enhance user navigation.

Our intent is to consider the interactions with data as knowledge itself. Armed with this knowledge, researchers will be in a position to share not only what was found to be interesting (the discovery), but exactly how it was found (the discovery process). In a broader context, we consider this research to be a fundamental contribution to developing solutions for an emergent research area: *information provenance*.[8] Problems in knowledge and data provenance[2] are gaining interest, with broad applications to the advancement of scientific discovery [13].

Provenance is a term that refers to the lineage of an item. While some people associate the term with artwork, and the lineage of who owned, or possessed the piece, we use it in the context of the information discovery process. The model that we are presenting supports provenance by fully documenting the discovery process. The prototype demonstrates how users can interact with the history of interactions and capture annotations in the same context. Another user may take the interaction data and use it against a different dataset, to see how general the technique may be.

RELATED WORK

The interaction model builds upon work from three areas of research: knowledge discovery in databases (KDD), annotation; and user tracking. We make an assumption that the knowledge discovery process is, indeed, a process, and that the steps that a user takes to discover knowledge are as important as the knowledge itself. Described by Fayyad, et al [5], the KDD process is frequently depicted in terms of a number of iterative steps. There are, of course, obvious similarities between the KDD process and the data visualization pipeline.[15] Our approach is to track all interactions with the data during the KDD process and to provide visualization tools for the KDD process.

A critical aspect of the process is the implied interaction with a user. Obviously, the user is involved with problem selection, as well as the interpretation of the results. Often, the user may review the results and develop a more refined problem statement, which initiates further exploration. In the context of this research, annotation is the adding of information to existing data by a user. For visualizations, numerous approaches have been described. Marshall and Brush [12] discuss the issue of shared annotations.

User tracking involves the recording of actions taken by users in the course of completing a task. Scaife and Rogers [14] critically examine the linkages between external representations (e.g. visualizations) and a user's corresponding internal representation of the information. They describe the concept of a cognitive trace, which may include explicit marks, or highlighting, of information. A need to record the corresponding parameter settings for the software is also identified. Fitzgerald, et al [6] define a framework for describing event-tracking for multimodal user interfaces. Such a framework can be used to develop a more comprehensive model of user interaction. Franklin, et al [7] describe a tracking mechanism for an electronic classroom environment, in which the users actions are tracked for playback purposes.

Tracking of user interactions within a visualization environment has been studied by Lee and Grinstein [10] and, more recently, Jankun-Kelly, et al [9] and Lowe, et al [11]. These previous efforts are particularly relevant to our work at the conceptual level. Where we differ from them is in the capturing of meta-information, such as annotations, along with the interactions.

CONCEPTUAL MODEL AND PROTOTYPE

We base the model [8] on directed graphs, with nodes signifying measurable states of the system, and edges denoting transitions between the states. The states of the system are generically captured in the model, leaving it up to the implementation to define the specific contents of the state and transition information. For example, as described in [9], the transitions might contain discrete interactions, such as zoom, rotate, or translate. Pictorially, the graph can be depicted as shown in the example in Figure 2.

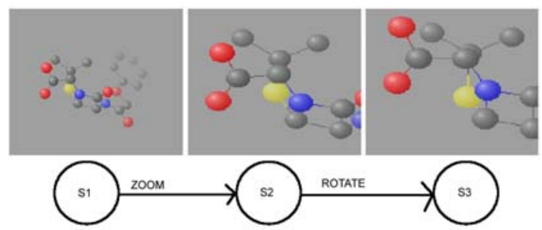

Figure 2: A simple interaction graph and user views.

The model allows for the articulation of annotation data with the interactions that are being applied to the visualization. For example, we will assume that the user follows the general process of: 1) observing the display; 2) making an annotation; and 3) applying an interaction. For this example, the annotation data is represented by text. However, there is no restriction on the mode of input used to perform the annotation.

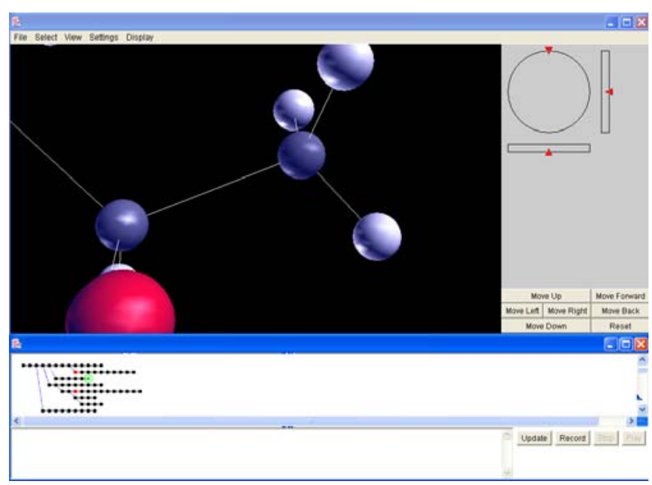

Figure 3: Prototype implementation and the history panel.

Our prototype is implemented according to the architecture shown in Figure 4. We modified a system for visualizing multivariate data to have the CheckState, GetState and SetState methods to generate the graph nodes. The state information that we tracked within the system was a viewpoint model - camera position and direction within a 3D environment.

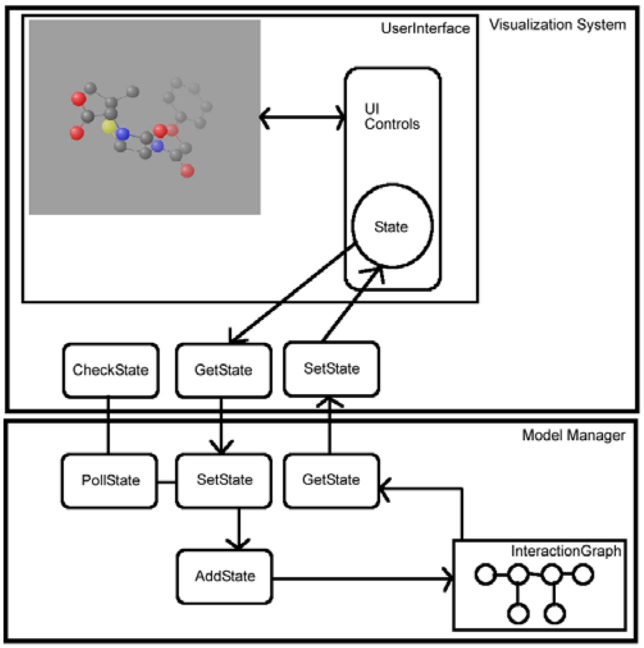

Figure 4: Architecture for tracking interaction history.

One of the most interesting aspects of our prototype is the model manager interface, which exposes the interaction graph to the user. The resulting application allows the user to interact with the visualization system as well as the interaction graph. Figure 3 shows a screen capture from the model manager. The user interaction was a simple sequence of zooming operations to display an overview of the entire dataset.

The prototype model manager supports annotation directly through either typed comments, or recorded voice. This capability saves the visualization system the effort of performing annotation capabilities. We have developed a tablet PC based interface to support direct scribbling of annotations.[4] The interaction graph displays visual cues to indicate current position within the interaction history as well as the location of annotations.

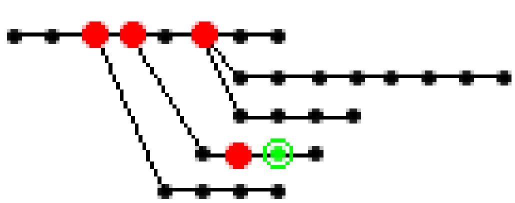

Figure 5: An example interaction graph.

It is worth pointing out that there is no restriction in the model for branching, or non-linear behavior represented in the interaction graph. However, in order for the model to support non-linear behavior tracking the model manager must keep track of where the user is relative to the interaction graph. The prototype supports this feature by creating branches in the graph if the current state is not a leaf node. An example graph is shown in Figure 5. Note the use of color to provide visual cues to the user. Larger, red nodes signify the location of annotation data, a single green node with a larger circle around the node is the current location in the graph.

Users can edit their histories by cutting segments of the graph out with the mouse. Histories can be combined by inserting graphs as sub-graphs of an active history. Also, we have developed intelligent pruning techniques to collapse the graph into only annotated nodes.

OPPORTUNITIES AND CHALLENGES

The interaction model demonstrates one way of implementing personal history tracking. There are a number of directions that this project can take. These opportunities are detailed here.

Question: Can interactions be modeled generically?

A generic model of interaction history allows the separation of interactions from the objects that are the target of interaction. The interactions can then be replayed against different datasets, or in different contexts.

Question: To what extent does collaborative, or shared, explorations influence discovery?

Collaborations may occur asynchronously, when users send their histories to each other, or synchronously, when users are simultaneously interacting within a shared space.

Question: Can recommender systems be developed to leverage historical interactions?

The opportunity exists to point users to views that other people have taken of the data. However, we need to determine what makes a particular view "interesting" enough to warrant suggesting.

Question: Should recommender systems be developed to leverage historical interactions?

What personal or privacy concerns are there? Will users only follow the paths of others? If so, exploiting personal histories may actually stifle discovery instead of enabling it. A different view of the problem suggests that coverage of the interaction space can be maximized by suggesting views that have not occurred yet.

Question: Can histories be compared?

One way of comparing histories is to break the interactions into discrete events and disregard temporal relationships. This approach is akin to leaving a fingerprint on an object – we can tell that something has been touched. The temporal aspect, or the context of when objects are interacted with, is much more interesting. We will need to develop a similarity metric to quantify history comparisons.

CONCLUSION

We have presented a model for tracking of interactions for knowledge discovery tasks. An implementation of the model for visualization tasks shows how personal histories are captured, edited and shared. The prototype further explores how users interactions with their personal histories introduces a new style of interaction. Clearly, there are many interesting and challenging problems to be addressed in this research space.

REFERENCES

1. Bargeron D. and Moscovich, T. Reflowing Digital Ink Annotations. *Proceedings of the Conference on Human Factors in Computing Systems (*2003), 385-393.
2. Buneman, P., Khanna, S., and Tan, W. Why and Where: A Characterization of Data Provenance. *Proceedings of the International Conference on Database Theory* (2001), 316-330.
3. Card, S. and MacKinlay, J. The Structure of the Information Visualization Design Space. *Proceedings of the IEEE Symposium on Information Visualization* (1997), 92-99.
4. Ellis, S. and Groth, D. A Collaborative Annotation system for Data Visualization. *Proceedings of the ACM Conference on Advanced Visual Interfaces.* (2004), 411-414.
5. Fayyad, U., Piatetsky-Shapiro, G. and Smyth, P. From Data Mining to Knowledge Discovery: An Overview. *Advances in Knowledge Discovery and Data Mining*. (1996), 1-34.
6. Fitzgerald, W., Firby, R. and Hanneman, M. Multimodal Event Parsing for Intelligent User Interfaces. *Proceedings of the International Conference on Intelligent User Interfaces.* (2003), 53-60.
7. Franklin, D., Budzik, J. and Hammond, K. *Plan-Based Interfaces: Keeping Track of User Tasks and Acting to Cooperate. Proceedings of the International Conference on Intelligent User Interfaces.* (2002), 79-86.
8. Groth, D. Information Provenance and the Knowledge Rediscovery Problem. *Proceedings of the International Conference on Information Visualization.* (2004), 345-351.
9. Jankun-Kelly, T.J., Ma, K., and Gertz, M. A Model for the Visualization Exploration Process. *Proceedings of the IEEE Conference on Visualization.* (2002), 323-330.
10. Lee, J. and Grinstein, G. An Architecture for Retaining and Analyzing Visual Explorations of Databases. *Proceedings of the IEEE Conference on Visualization.* (1995), 101-109.
11. Lowe, N., Tory, M. Potts, S., Datta, A. and Moller, T. A Parallel Coordinates Style Interface for Exploratory Volume Visualization. *IEEE Transactions on Visualization and Computer Graphics.* (2005), 71-80.
12. Marshall, C. and Brush, A. From Personal to Shared Annotations. *Proceedings of the Conference on Human Factors in Computing Systems.* (2002), 812-813.
13. Myers, J., Chappell, A., Elder, M., Geist, A. and Schwidder, J. Reintegrating the Research Record. *IEEE Computing in Science and Engineering.* (2003), 44-50.
14. Scaife, M. and Rogers, Y. External Cognition: How Do Graphical Representations Work. *International Journal of Human-Computer Studies.* (1996), 185-213.
15. Wright, H. Brodlie, K. and Brown M. The Dataflow Visualization Pipeline as a Problem-Solving Environment. *Virtual Environments and Scientific Visualization.* (1996), 267-276.

AUTHOR BIOGRAPHY

Dennis Groth has a B.S. in Computer Science from Loyola University of Chicago and a Ph.D. in Computer Science from Indiana University. His research focuses on the intersection of databases, visualization and interaction. He is an Assistant Professor at the Indiana University School of Informatics in Bloomington, Indiana.