

User-Adaptive Visualizations: Can Gaze Data Tell Us When a User Needs Them?

Cristina Conati, Giuseppe Carenini, Mona Harati, Dereck Tocker, Nicholas Fitzgerald, Anna Flagg

University of British Columbia, Department of Computer Science, 2366 Main Mall, Vancouver, BC, Canada
{conati, carenini, haraty, derecktocker, nfitz, aflagg}@cs.ubc.ca}

Abstract

The primary goal of our research is to design adaptive information visualization systems that adapt to the specific needs of each individual viewer. Our first step is to explore data sources that could help detect these needs in real-time, including cognitive measures that impact perceptual abilities, interface interactions, eye-tracking, and physiological sensors. In this paper, we focus on current efforts to understand which cognitive measures can be relevant, as well as if/how a viewer's gaze pattern can predict performance on associated visualization task.

Introduction

This Information visualization is a thriving area of research in the study of human/computer communication. However, attempts to measure and formalize visualization effectiveness often lead to inconclusive and conflicting results (Nowell et al., 2002). We believe this is because existing visualizations are designed mostly around the target data set and associated task model, with little consideration for user differences. Both long term user traits like cognitive abilities as well as short term factors like cognitive load and attention have been largely overlooked in the design of information visualizations, despite studies linking individual differences to visualization efficacy for search and navigation tasks (Allen, 2000; Dillon 2000) as well as anecdotal evidence of diverse personal visualization preferences (Baldonado et al., 2000). Thus we plan to explore the possibilities of intelligent, human-centered visualizations that *understand* different users have different visualization needs and abilities, and can *adapt* to these differences.

There is already some evidence of the impact that individual differences can have on visualization effectiveness. For example, Velez et al., (2005) found significant correlations between individual spatial traits and performance on identification of a 3D object from visualizations of its orthogonal projections. Conati and

Maclaren (2008) found that an individual's perceptual speed was a significant predictor of her ease in understanding the same data set with two different visualization types.

Although the benefits of user-adaptive interaction have been shown in a variety of tasks such as operation of menu-based interfaces, web browsing, desktop assistance and human-learning (Jameson 2007), these ideas have rarely been applied to data visualization. This is largely due to the fact that there is limited understanding of which combinations of user traits/goals are relevant for adaptivity. Two notable exceptions are the work by Gotz and Wen (2009), and by Brusilowsky et al. (2006). Gotz and Wen propose a technique to automatically detect a user's changing goals during interaction with a multi-purpose visualization, and adapt the visualization accordingly. In contrast, we focus on adapting the visualizations to other relevant user-dependent factors in addition to goals. Brusilowsky et al. (2006) adapt the *content* of the visualization to the user's state in an educational system, but maintain a fixed visualization technique. In contrast, we are interested in adaptation that involves both selecting *alternative visualizations* for different users, as well as providing *adaptive help* with a given visualization to accommodate changing user needs during interaction (see Figure 1, right).

To achieve this objective, two research questions need to be answered: 1) given a visualization, why do some people perform better than other, and 2) how can a visualization system detect when a user is not performing well. We plan to explore two avenues to answer these questions. One involves further investigating if and how user traits (e.g. spatial/perceptual abilities, personality traits, learning styles) may impact visualization effectiveness. If such measurable features are found and are collectible before interaction, they could be given as input to an adaptive visualization to help it select the visuals for this viewer (See figure 1, bottom left).

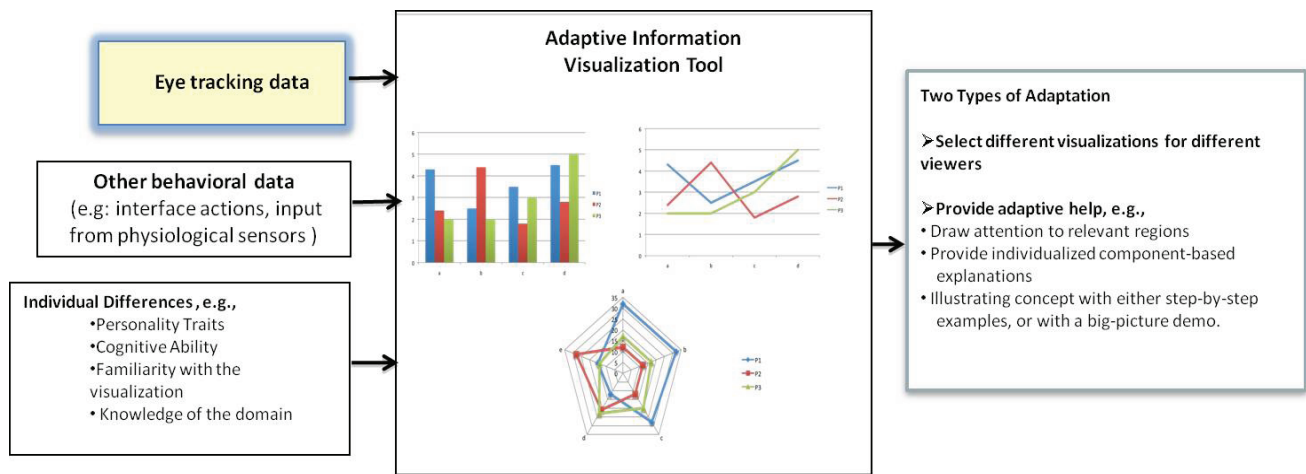


Figure 1: Adaptive visualization framework - data sources that can predict visualization effectiveness for a specific user (left) are processed in real-time by an adaptive visualization and may trigger one or more of the adaptive interventions to the right

Our second approach is to study whether user proficiency with a given visualization can be inferred from her interaction behaviors (see figure 1, top left). We believe that an important window into these behaviors can be provided by eye-tracking information. Thus, we are currently conducting a study to collect initial data on if/how a viewer's gaze patterns can predict performance on associated visualization tasks. In the rest of the paper, we discuss the general objectives and design of the study.

Evaluating Gaze Patterns and Cognitive Measures as Input to Adaptive Visualizations

Several researchers have explored eye-tracking as a source of information for real-time assessment of human/machine interaction performance. Amershi and Conati (2009) used an unsupervised machine learning technique to separate effective and ineffective user behaviors during interaction with a teaching tool for math. The behaviors captured both interface actions as well as attention patterns monitored via eye-tracking. We plan to conduct similar studies to try to reproduce these results in the context of visualizations.

Iqbal and Bailey (2004) found that a given task has a unique signature of eye movement; Goldberg and Helfman (2010) provide initial indications that it is possible to identify specific gaze patterns associated with a variety of tasks performed on different visualizations. Thus, in the current study, we are looking at two alternative visualizations, bar charts and radar graphs, to perform a series of basic tasks based on a set of low-level analysis tasks that Amar et al. (2005) identified as largely capturing people's activities while employing information visualization tools for understanding data.

Figure 3 and 4 show examples of bar charts and radar graphs, respectively, that we use in the study. Each of these visualizations is used to compare the performance of a student against the course average on a series of courses. Sample tasks users are engaged in include, for instance,

comparing two courses (*Is the student stronger in Marine Biology or in Painting*), finding extreme values (*In which course does the student deviate most from the class average?*), computing derived values (*In how many courses is student above the class average?*).

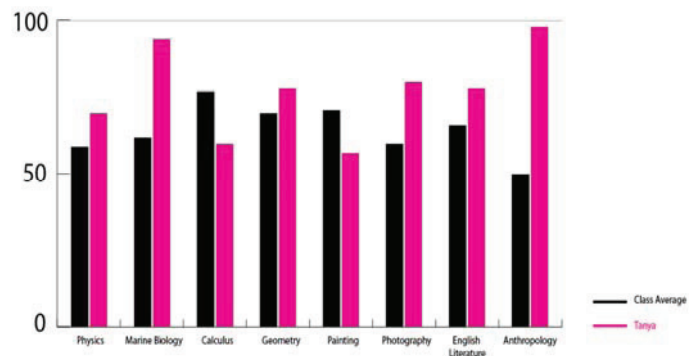


Figure 2: Sample bar chart in the user study




Figure 3: Sample radar graph in the user study

The experiment is run with software that records the user's answers and completion time for each task, as well as gaze

data captured via a Tobii T120 eye-tracker. In addition, we are collecting *retrospective verbal protocols* of the user performance (e.g., Ellig et al., 2011): after performing each task in natural conditions, the user is asked to look at the task again and verbalize how he/she reached her answer. We will code these verbal protocols to see whether we can identify both instances of confusion, as well as specific strategies that users may be adopting to answer the questions. The goal is to use this coded verbal protocols, along with task correctness and completion time, to label the gaze data generated by the user while performing the task; then use the dataset to build a classifier that can identify attention patterns indicative of suboptimal visualization processing and thus calling for adaptive interventions. While we still don't have any concrete result on whether such a classifier can be built, the following findings from previous work provide some indications of the type of gaze features that may predict user visualization performance.

1. *The duration of fixations on each area of interest is an indicator of the complexity of that area*

A study by Crowe and Narayanan (2000) found that—as one might expect—an unusually long fixation on one component of a visualization indicates lack of understanding of that component. Identifying these areas may make for a more focused adaptation, because it allows the system to target the specific area that is perplexing the viewer.

2. *Degree of pupil dilation has been proved to be a valid and reliable measure of cognitive load (e.g., Loewenfeld and Wilhelm, 2002).*

We plan to investigate if pupil dilation as measured via an eye-tracker can be a reliable indication of cognitive load during visualization processing. For example, detecting high cognitive load could prompt the system to take steps to simplify the data presentation or the viewer's task.

3. *Users do not look at all areas of interest*

Analyzing gaze locations might be a good first step to identifying when a viewer is having trouble with a given visualization. Lohmann, et al. (2009) used this approach to compare relative effectiveness of alternative tag cloud visualizations in the context of drawing attention to the areas of greatest interest. Gaze locations, and the locations that have been overlooked, can inform the design of the adaptive help. For instance, after becoming aware the viewer is not looking at an area of crucial importance, the visualization could emphasize this area to attract attention

Utilizing Gaze Data For Two Types of Adaptation

We are interested in adaptation that involves both selecting different visualizations for different viewers, as well as providing adaptive help within a visualization to accommodate changing user needs during interaction. For example, given a set of alternative visualizations, our adaptive system would monitor the interaction and may

switch visualizations if the current display does not appear to be working for the viewer. During the interaction itself, the system would focus more on providing explicit interactive help, such as by drawing attention to certain important areas or by explaining explicitly how to derive a given piece of information from the current visualization.

These proposals for adaptation must be thoroughly tested within the context of information visualization before they can be realistically applied. Thus, in addition to conducting studies to validate the use of eye-tracking data in detecting when a viewer is having difficulties, we are investigating the benefits and feasibility of a variety of adaptive interventions within the context of information visualizations. For instance, we are designing a study in which we compare alternative ways to highlight relevant visualization elements, e.g., segments indicating differentials in a radar graph.

References

- Allen, B. 2000. Individual differences and the conundrums of user-centered design: Two experiments. *Journal of the American Society for Information Science* 51, 508-520.
- Amar, R., Eagan, J., and Stasko, J. 2005. Low-level components of analytic activity in information visualization, 15.
- Amershi, S. and Conati, C. 2007. Unsupervised and supervised machine learning in user modeling for intelligent learning environments. *Proceedings of the 12th international conference on Intelligent user interfaces*, 72–81.
- Baldonado, M.Q.W., Woodruff, A., and Kuchinsky, A. 2000. Guidelines for using multiple views in information visualization. *Proceedings of the working conference on Advanced visual interfaces*, ACM (2000), 110-119.
- Conati, C. and Maclaren, H. 2008. Exploring the role of individual differences in information visualization. *Proceedings of the working conference on Advanced visual interfaces*, 199–206.
- Dillon, A. 2000. Spatial-semantics: how users derive shape from information space. *J. Am. Soc. Inf. Sci.* 51(6), 521-528.
- Elling, S., Lentz, L. and de Jong, M. 2011: Retrospective think-aloud method: using eye movements as an extra cue for participants' verbalizations. CHI 2011: 1161-1170
- Goldberg, J.H. and Helfman, J.I. 2010. Comparing Information Graphics: A Critical Look at Eye Tracking, *BELIV '10*.
- Iqbal, S.T. and Bailey, B.P. 2004. Using eye gaze patterns to identify user tasks. *The Grace Hopper Celebration of Women in Computing*.
- Loewenfeld, I. and Wilhelm, H. 2000. The Pupil: Anatomy, Physiology, and Clinical Applications, Volumes I and II. *Archives of Ophthalmology*, 118864–864.
- Nowell, L., Schulman, R., and Hix, D. 2002. Graphical encoding for information visualization: an empirical study. *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002*, 43–50.
- Velez, M.C., Silver, D., and Tremaine, M. 2005. Understanding visualization through spatial ability differences. *Proceedings of Visualization 2005*.

Crowe, Eric C., and Narayanan, Hari N. 2000. Comparing Interfaces Based on What Users Watch and Do. ETRA '00 Proceedings of the 2000 symposium on Eye tracking research & applications

Lohmann, S., Ziegler, J., and Tetzlaff, L. 2009. Comparison of Tag Cloud Layouts: Task-Related Performance and Visual Exploration.