Towards Adapting Information Graphics to Individual Users to Support Recognizing Intended Messages

Richard Burns[1*], Stephanie Elzer Schwartz[2], and Sandra Carberry[3]

[1] Dept of Computer Science, West Chester University, West Chester, PA 19383 USA
[2] Dept of Computer Science, Millersville University, Millersville, PA 17551 USA
[3] Dept of Computer Science, University of Delaware, Newark, DE 19716 USA

Abstract. Our previous body of work has implemented a system for inferring the intended messages of information graphics (bar charts, line graphs) that appear in popular media such as newspapers and magazines. This paper explores how a model of user characteristics could be utilized to tailor information graphics to individual users, thus enabling users with varying degrees of skill and ability to attain the desired information while preserving the graph designer's intended message.

1 Introduction

Information graphics differ from the visual data representations typically studied in the area of information visualization in that information graphics appearing in popular media such as newspapers and magazines are typically designed to convey one or more *messages*. For example, a bar chart might be designed to show that consumers in the US have the maximum amount of credit debt of all the of countries shown. We have posited that graphic designers utilize a variety of communicative signals, such as salience, annotations, and perceptual task effort, to enable the typical user to infer the intended message of the graph [5]. Our research has shown that we can apply this stereotypical user model to automatically infer the intended message of information graphics [6, 14, 2]. This inferred message supports multiple applications, including providing access to information graphics for visually-impaired users.

In this paper, we consider another potential application of our work: facilitating and assisting viewer comprehension of intended messages. There is a burgeoning field of study regarding how individual characteristics influence the processing and comprehension of information visualizations [12, 4, 8]; Steichen and Conati [1] have begun to examine how a model of a user's cognitive abilities might be employed to adapt visualizations to individuals. We hypothesize that information graphics differ from many of the data visualizations that have been studied to date because of the presence and importance of intended messages, and that our message inference system could be extended to include a model of cognitive abilities for facilitating viewer recognition of intended messages.

* rburns@wcupa.edu

2 Using Adaptive Visualizations for Message Recognition

Our long-term goal is novel because we wish to automatically adapt and modify an information graphic visualization based on the most prominent intended messages in a graphic and a user model that takes into account personal traits and preferences. We envision two categories of visual adaptations:

1. Changing the type of graphic (e.g. "linear" to"radial") to facilitate an individual user's strengths and support a user's preferred graph type.
2. Adapting the graphic within the current graphic type (e.g. keeping it a bar chart) and modifying the design through: (a) proximity (modifying the distance between graph entities to affect the degree of visual clutter that is present), (b) salience (e.g. highlighting/coloring a bar or a series of bars to increase salience; annotating elements with their data values), (c) the reorder of elements in a graphic, and (d) proportion shifts (e.g. changing the scale of the dependent axis to alter the difference in height between entities).

In both categories of adapations, we wish to ensure that the original intended message of the graphic designer is still conveyed in the modified graphic. Thus, we believe that our implemented intention recognition systems [6, 14, 2] and perceptual task effort models [7, 3]—that estimate the relative effort for a stereotypical user to recognize a message given a graph—will be vital, because we could run adapted graphics through the systems on-the-fly to ensure that the intended messages remain the same.

3 Towards Modeling Individual Users

We believe that in order to appropriately adapt information graphics to individual users, we need to model users at a finer granularity. Our hypothesis is that this can be performed by applying the work of cognitive research that analyzed the relationship of eye-tracking results, mouse movements, and mouse clicks for individual users performing visualization tasks. The following are components we believe should be incorporated, as well as some of the ways that these components could be applied for adapting information graphics to individuals:

1. **Preference:** Baldonado [10] has suggested that there is anecdotal evidence for diverse personal visualization preferences between individual users. These preferences should influence the selection of the type of information graphic or graphics that are utilized.
2. **Experience:** Shah [11] found that experience was a significant factor for graph comprehension, and that less experienced individuals spend more time in information retrieval and comparison substages. Lewandowsky and Spence [9] found that experience improved the accuracy of performed graph tasks. These results indicate that experience could influence the amount of "cuing"—through salience and annotations—that is necessary for a user.

3. **Personality:** Ziemkiewicz [15] proposed that there is a relationship between LOC (locus of control: one's tendency to view themselves as in control of external events) and compatibility with a graph's layout style (spatial arrangement of graphed entities). She hypothesizes that users with a more external LOC are more willing to adapt their thinking to unfamiliar visual designs. Our analysis of a corpora of information graphics has shown that particular types of information graphics are more commonly utilized for specific messages. Individuals with an external LOC may be more tolerant of less common pairings of information graphics and messages, while individuals with an internal LOC may require more "cuing".
4. **Perceptual Speed and Visual Working Memory:** Velez et al. [13] explored both an individual's perceptual speed and visual working memory capacity in spatial visualization tasks. Our hypothesis is that perceptual speed and visual working memory will influence: (a) the need to annotate graph elements with values rather than having a user interpolate the values, and (b) the desired proximity of graph elements that need to be processed together. We have previously observed that individuals can recognize trends in *groups* of a grouped bar chart faster than in *series*, suggesting that users with slower perceptual speed could benefit from redesigned graphics that convey trends via visual *groups* to maximize the proximity of entities that should be compared.
5. **Verbal Working Memory:** Toker et al. [12] hypothesized that an individual's verbal working memory capacity may also be a significant factor in graph comprehension because of the present textual components in many visualizations, such as legends, labels, wordings in axes, and overall graph titles. For individuals with more limited verbal working memory, it may be desirable to eliminate any extraneous text in the information graphic, either adding it in stages or only including the additional elements if the user requests them.

We believe that an individual user model comprised of these components could be utilized to effectively adapt an information graphic to the user while maintaining the graphic designer's intended message. The components of this data model could be constructed by techniques similar to the research we cite in this paper, including eye tracking, off-line assessments, and recordings of mouse movements and mouse clicks.

4 Conclusion and Future Work

In this paper, we have outlined our ideas for extending our message recognition system for information graphics in order to adapt information graphics to individual users. We would like our system to redesign information graphics, such that (1) the graphic designer's intentions for the original graphic are still very apparent in the redesigned graphic, and (2) the design of the redesigned graphic facilitates an individual's recognition of the graph's intended messages more so than in the design of the original graphic.

References

1. B. Steichen, G. Carenini, C. Conati: Adaptive information visualization: Using gaze data to infer visualization types, tasks and user characteristics. In: Proceedings of IUI 2013. ACM (2013)
2. Burns, R., Carberry, S., Elzer, S., Chester, D.: Automatically recognizing intended messages in grouped bar charts. In: Cox, P., Plimmer, B., Rodgers, P. (eds.) Diagrammatic Representation and Inference, pp. 8–22. No. 7352 in Lecture Notes in Computer Science, Springer Berlin Heidelberg (Jan 2012)
3. Burns, R., Carberry, S., Elzer Schwartz, S.: Modeling a graph viewer's effort in recognizing messages conveyed by grouped bar charts. In: Proceedings of the 21st Conference of User Modeling, Adaptation and Personalization (2013)
4. Conati, C., Carenini, G., Harati, M., Tocker, D., Fitzgerald, N., Flagg, A.: User-adaptive visualizations: Can gaze data tell us when a user needs them? In: Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence. pp. 8–11 (Aug 2011)
5. Elzer, S., Carberry, S., Demir, S.: Communicative signals as the key to automated understanding of simple bar charts. In: Proceedings of the 4th international conference on Diagrammatic Representation and Inference. p. 2539. Diagrams'06, Springer-Verlag, Berlin, Heidelberg (2006)
6. Elzer, S., Carberry, S., Zukerman, I.: The automated understanding of simple bar charts. Artif. Intell. 175(2), 526555 (Feb 2011)
7. Elzer, S., Green, N., Carberry, S., Hoffman, J.: A model of perceptual task effort for bar charts and its role in recognizing intention. User Modeling and User-Adapted Interaction 16(1), 130 (Mar 2006)
8. Goldberg, J., Helfman, J.: Eye tracking for visualization evaluation: Reading values on linear versus radial graphs. Information Visualization 10(3), 182–195 (Jul 2011)
9. Lewandowsky, S., Spence, I.: The perception of statistical graphs. Sociological Methods & Research 18(2-3), 200–242 (Nov 1989)
10. Michelle Q. Wang Baldonado, Allison Woodruff, Allan Kuchinsky: Guidelines for using multiple views in information visualization. In: Proceedings of the Working Conference on Advanced Visual Interfaces. pp. 110–119. ACM (2000)
11. Shah, P., Freedman, E.G.: Bar and line graph comprehension: An interaction of top-down and bottom-up processes. Topics in Cognitive Science 3(3), 560578 (2011)
12. Toker, D., Conati, C., Carenini, G., Haraty, M.: Towards adaptive information visualization: on the influence of user characteristics. User Modeling, Adaptation, and Personalization p. 274285 (2012)
13. Velez, M., Silver, D., Tremaine, M.: Understanding visualization through spatial ability differences. In: IEEE Visualization, 2005. VIS 05. pp. 511–518 (Oct 2005)
14. Wu, P., Carberry, S., Elzer, S., Chester, D.: Recognizing the intended message of line graphs. In: Proceedings of the 6th international conference on Diagrammatic representation and inference. p. 220234. Diagrams'10, Springer-Verlag, Berlin, Heidelberg (2010)
15. Ziemkiewicz, C., Crouser, R., Yauilla, A., Su, S., Ribarsky, W., Chang, R.: How locus of control influences compatibility with visualization style. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST). pp. 81–90 (2011)

User Task Adaptation in Multimedia Presentations

Giuseppe Carenini, Cristina Conati, Enamul Hoque, Ben Steichen

Department of Computer Science, University of British Columbia, Vancouver, Canada
{carenini, conati, enamul, steichen}@cs.ubc.ca

1 Multimedia Presentations Combining Visualizations and Text

It is quite common that documents ranging from newspaper articles to scientific papers convey complex information by combining visualizations with textual material. Presenting information in different modalities not only makes the presentation more engaging, but could also better suit users with different cognitive skills (visual vs. verbal). In these multimedia presentations graphics and text play complementary roles. While graphics can convey large amounts of data compactly and support discovery of trends and relationships, text is much more effective at pointing out and explaining key points about the data, in particular by focusing on specific temporal, causal and evaluative aspects [1]. For illustration, Figure 1 shows an example of a multimedia presentation from The Economist magazine. Notice, for instance, how the sentence "*The end of subsidies to car buyers will lead to a slump in Japan, just as its carmakers' output recovers from the 2011 tsunami.*" provides a causal explanation for the noticeably extreme data about current (year 2012) and forecasted (year 2013) car sales in Japan. Generally speaking, the textual part of a multimedia presentation can be seen as suggesting to the reader a set of visual tasks that can be performed by inspecting the visualization. For example, when reading the two sentences *"India and China will have further strong rises—though not at the double-digit rates seen until 2010. Brazil and Britain will suffer reverses."* the reader is prompted to verify in the visualization (the deviation chart) that all the bars for India and China are on the right side of the chart (i.e., sales are increasing) and less than 10%, while the bars for Brazil and Britain are on the right for 2012, but on the left side (i.e., sales are decreasing) for the 2013 forecast.

2 Task Adaptation in Multimedia Presentations

As we have illustrated, the textual component of a complex multimedia presentation typically specifies several visual tasks for the reader. Although the reader should be able to perform all these tasks in the visualization, the visualization components cannot be designed to favour the performance of any particular one of those tasks. Arguably, designing the visualization to support a specific task will likely hinder the performance of the other ones. The main idea of this short paper is that if a system could track what part of the text the reader is currently reading, and from that text it could infer the corresponding visualization task, such a system could **dynamically adapt the visualization** to support the reader in more effectively performing the inferred current task. For instance, if the user read the sentence *"India and China will*

The global car industry: Wheels of mixed fortune *[Source: The Economist- Dec 22nd 2012]* America will enjoy a fourth consecutive year of growth in car sales in 2013, predicts IHS, a research firm. India and China will have further strong rises—though not at the double-digit rates seen until 2010. Brazil and Britain will suffer reverses. The end of subsidies to car buyers will lead to a slump in Japan, just as its carmakers' output recovers from the 2011 tsunami. In the European Union, car sales will fall for the sixth year in a row: they are now back at early 1990s levels. Although some European car factories face closure, elsewhere assembly lines are being built at a rapid clip. So once again worldwide car production, at 82.8m, will exceed sales, at 81.9m. As the metal stacks up on dealers' forecourts, motorists can look forward to some great deals on wheels in 2013.

Figure 1: Sample Multimedia Presentation Combining a Chart with Text

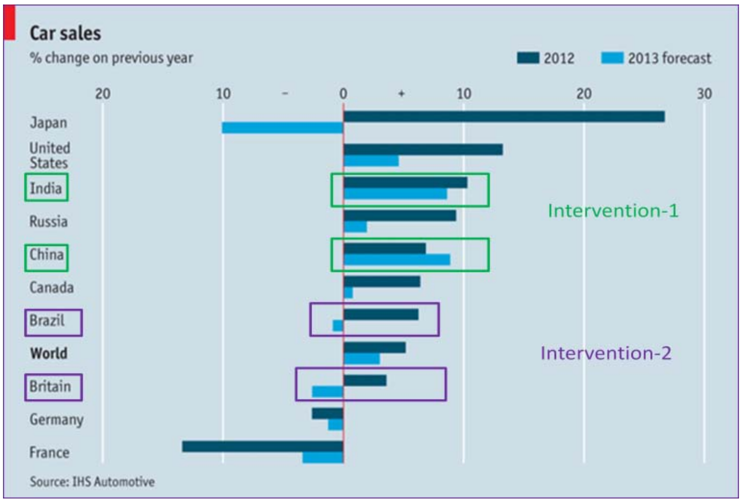

Figure 2: Sample Interventions Adapted to the User Task

have further strong rises—though not at the double-digit rates seen until 2010.", and then turned to the visualization, the system could highlight the information relevant to the corresponding task in the visualization, as shown in Figure 2 by intervention-1. Similarly, if the user turned to the visualization after reading the sentence "*Brazil and Britain will suffer reverses.*", the system could generate intervention-2. Let us now discuss more in detail the modules of a system that can perform task adaptation in multimedia presentations and how each module could be implemented
1) Identify what part of the text the user is currently reading. This module could rely on eye-tracking technology. Although eye tracking devices are still too costly for mass use, more affordable solutions (e.g., based on standard webcams) will be available in the near future [2,3]. **2) Verify if the text the user is currently reading is relevant to the visualization.** Preliminary ideas on how this could be done are presented in [4]. **3) Infer visual tasks for text the user is currently reading.** This is a challenging Natural Language Processing (NLP) problem. Some progress on a related problem was made by Elzer et al. [5] (Sec. 5.2) and [6], while working on using the caption of a visualization as one of the sources to infer the message the visualization was intended to convey. They process the caption to identify and combine verbs (e.g., *lag*), nouns (e.g., *growth),* nouns referring to labels in the visualization, and adjectives (e.g., *soaring*) that are typically used in captions to highlight key points about the displayed data . To be applicable in the system we envision, this work will need to be expanded to deal with more complex and more sophisticated NLP techniques.**4) Provide adaptive interventions.** This module will need to solve at least the following two sub-problems, which also may make use of eye-tracking. First, it needs to decide when an intervention should be triggered and when it should be faded away. A simple approach could be to trigger an intervention every time the user, after reading a chunk of text corresponding to one or more visual tasks, switches her gaze from the text to the visualization. The intervention could then be removed, only if the user has looked at it and possibly returned to the text. The second key problem is selecting the most appropriate intervention for a given task on the specific visualization [7].

3 Further Issues and Questions for Discussion at the Workshop

So far we have described interventions on the visualization to support visual task(s) specified by the text that the user is currently reading. We can also envision interventions on the text, triggered by the user inspection of the visualization. For instance, if eye-tracking data show that the user is inspecting the bottom part of the chart in Figure 1, the sentence about the EU could be highlighted. Such interventions generate a number of research questions: How could the performance of specific visual tasks be detected? (See [8] for preliminary results). Would these graphic to text interventions be useful? What are the implications of allowing both types of interventions?
With the rapid progress in Intelligent User Interfaces, it will become more and more common for multimedia presentations to be generated automatically by computer systems [9, 10]. What are the implications of this for adaptive interventions? Is it the case that for these presentations it may be easier to perform user task adaptation? Next steps in our research include: developing a prototype, run user studies, and also explore adaptation in the context of text-to-speech interaction, which could be more feasible in the short term, as it does not require eye-tracking in the loop.

References

[1] E. R. Tufte. Visual Explanations: Images and Quantities, Evidence and Narrative. 1997.

[2] Sesma, L., Villanueva, A., & Cabeza, R. Evaluation of pupil center-eye corner vector for gaze estimation using a web cam. In Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12), (2012), 217-220.

[3] Samsung Galaxy S4 eye-tracking smartphone unveiled, http://www.bbc.co.uk/news/technology-21791023

[4] Peng Wu and Sandra Carberry. Toward Extractive Summarization of Multimodal Documents. Proceedings of the Canadian AI Workshop on Text Summarization, pages 53-64, 2011

[5]Elzer, S., Carberry, S., & Zukerman, I. The automated understanding of simple bar charts. Artificial Intelligence Journal 175(2), (2011), 526-555.

[6] S. Elzer, S. Carberry, D. Chester, S. Demir, N. Green, I. Zukerman, K. Trnka, Exploring and exploiting the limited utility of captions in recognizing intention in information graphics, in: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics, 2005, pp. 223–230.

[7] Mittal, V.O. Visual Prompts and Graphical Design: A Framework for Exploring the Design Space of 2-D Charts and Graphs. Proc. AAAI/IAAI, 57-63, 1997.

[8] Steichen B., Carenini G., Conati C, User-Adaptive Information Visualization - Using Eye Gaze Data to Infer Visualization Tasks and User Cognitive Abilities. Proceedings of the International Conference on Intelligent User Interfaces, Santa Monica, CA, USA March 19-22, 2013

[9] N. Green., G. Carenini., S. Kerpedjiev, J. Mattis, J. Moore, S. Roth, AutoBrief: an Experimental System for the Automatic Generation of Briefings in Integrated Text and Information Graphics. International Journal of Human-Computer Studies Vol. 61, Issue 1, pag, 32-70, July 2004

[10] Carenini G., Ng R., Pauls A. Interactive Multimedia Summaries of Evaluative Text, *Proceedings of the 10th International Conference on Intelligent User Interfaces* Sydney, Australia, Gen29-Feb1, 2006.

A Study of Emotion-triggered Adaptation Methods for Interactive Visualization

Daniel Cernea[1,2], Achim Ebert[1], and Andreas Kerren[2]

[1]University of Kaiserslautern, Computer Graphics and HCI Group,
P.O. Box 3049, D-67653 Kaiserslautern, Germany
{cernea,ebert}@cs.uni-kl.de
[2]Linnaeus University, Department of Computer Science, ISOVIS Group,
Vejdes Plats 7, SE-35195 Växjö, Sweden
andreas.kerren@lnu.se

Abstract. As the size and complexity of datasets increases, both visualization systems and their users are put under more pressure to offer quick and thorough insights about patterns hidden in this ocean of data. While novel visualization techniques are being developed to better cope with the various data contexts, users find themselves increasingly often under mental bottlenecks that can induce a variety of emotions. In this paper, we execute a study to investigate the effectiveness of various emotion-triggered adaptation methods for visualization systems. The emotions considered are boredom and frustration, and are measured by means of brain-computer interface technology. Our findings suggest that less intrusive adaptive methods perform better at supporting users in overcoming emotional states with low valence or arousal, while more intrusive ones tend to be misinterpreted or perceived as irritating.

Keywords: Adaptive affective visualization, emotion-based adaptation, adaptive methods, human factors.

1 Introduction

Massive datasets are being generated on many different occasions, e.g., in genome or climate research, huge software systems, or economy. At the same time, our dependency on a good understanding of data and effective analyses increases because of the need to discover, for example, new drugs, correlations in social networks, or bugs in complex software systems. As such, analyzing and gaining insight into these large multivariate datasets through visualization is one of the major challenges of our days [15]. However, the complexity of the data and the task can exercise pressure and influence the user in a negative way, inducing various emotional states: some of discomfort or frustration, and others of boredom or the feeling of being overwhelmed. At the same time, most computer systems are currently designed to execute the same operations "regardless of whether you are sitting forward eagerly in your seat, or have begun to emit loud snoring sounds" [18].

This has never been more true than in the case of visualizations that have the potential to closely interact with users and adapt to their needs, e.g. mental models, cognitive processes and affective states, in order to improve performance and increase user satisfaction [14, 21]. Still, visualization systems are mostly oblivious of user emotions and needs, thus limiting the adaptability of the representation and the overall impact of visualization effectiveness. In this paper we will focus on investigating the efficiency and intrusiveness of a set of emotion-triggered adaptation techniques for visualizations. This is achieved through a study in which participants are interacting with a visualization tool while wearing brain-computer interface (BCI) devices capable of interpreting electrical signals generated by their brains as user emotions.

In the following sections, we shortly highlight related work in the field of affective visualization, adaptive approaches and BCI-based emotion detection. Next, we describe the design of our study and WebComets [7], the visualization tool used by the participants. We finally highlight our findings and open questions, as well as expose our conclusions.

2 Emotion and Visualization Tailoring

In order to employ user emotions as a driving feature behind visualization tailoring, one must first consider how these emotional states could be detected. While there are various approaches involving speech, facial expressions and eye movement [4], one of the real-time approaches that has gained popularity in the recent years is represented by EEG measurements executed with lightweight neuroheadsets [5, 6, 12, 17]. Moving towards aspects of emotions in Information Visualization, such wireless non-intrusive BCIs have been used in evaluating the cognitive workload induced to the user by different visualization techniques [2]. At the same time, lightweight passive BCIs have been used to detect emotional states that can be indicative of moments of insight during the exploration of visualization systems [5]. Related to our study task, the work [9] highlights possible correlations of webpage complexity with emotional valence.

On the other side, adaptive systems have been investigated in the context of user emotion mostly in affective computing and emotion-based adaptive interfaces [3, 13, 16, 20]. Narrowing the view down to the field of visualization, user-adaptive representations have considered multiple human characteristics, e.g., shaping visualizations based on user behavior [11], visualization context [10] and user models [1]. However, emotion—a powerful inner force that influences humans in all their activities—remains widely unexplored as a user attribute employed for adaptive visualization.

3 Study on Emotion-Triggered Adaptive Visualization

The following study focuses on estimating the effectiveness of various emotion-triggered adaptation techniques applied on an information visualization system.

In order to allow a visualization to adapt to the user emotions, we need to support the real-time detection of these emotions. In our study, this is achieved through the Emotiv EPOC headset, a wireless, non-intrusive EEG device that is capable of capturing EEG signals from the brain and interpreting them as a set of user emotions through a software framework. The accuracy of the emotion detection achieved with this BCI device and its framework has been previously investigated showing promising results [5, 6] and therefore is not the focus of our current study. Although capable of detecting a wider range of emotions through the included software framework (detailed in [5]), our attention went towards a subset of available emotions, namely *boredom* and *frustration*. These two emotions have been selected as they can influence the user experience negatively, but also due to their negative positioning on the arousal (boredom) and valence (frustration) axes of Russell's circumplex model of affect [19]. Furthermore, non-basic emotions like frustration and boredom have been recently reconfirmed as highly relevant in affective computing [8].

Our study involved 5 participants (2 female, 3 male) with an average age of 23 years. The majority of these participants had some experience with visualizations and have participated in previous BCI-based studies. In order to explore the potential of emotion-triggered adaptation of visualization, we decided to use a visualization system that enabled—through a wider range of supported functionality and interaction—the implementation of multiple adaptation approaches: WebComets [7] is an interactive visualization tool for multi-session, multi-user parallel browsing histories. The information displayed in the visualization corresponds to the websites that one or multiple users have accessed during their online navigation. The accessed websites are represented as circles (called nodes), and positioned on horizontal lines encoding the individual tabs that were employed to load the page, see Fig. 1 (top). Horizontal lines can be interconnected vertically, suggesting that one browser tab has branched off from another one when a user clicked on a hyperlink, thus generating a deep tree-like browsing structure. The temporal dimension flows from left to right.

For the purpose of our study, the WebComets visualization has been adapted to react to the emotional readings detected by the EPOC headset. More precisely, when the headset signals the detection of boredom or frustration, the visualization tries to support the user and make some changes to the system in order to influence the user's state. Note that in order to avoid triggering the visualization adaptations too often, the emotion triggers are only activated if the boredom and frustration readings are 75% above the individual baseline for a period of over 3 seconds. The baselines for both emotions have been individually established prior to the study, while each participant executed mundane tasks that did not influence the emotional readings from the BCI and during which the participants themselves reported a balanced emotional state.

The participants received a set of tasks of identifying particular node sequences (i.e., browsing patterns). Furthermore, they were given an introduction of the visualization, by covering theoretical background, functionality and sample tasks until the point were every user felt confident using the system. They

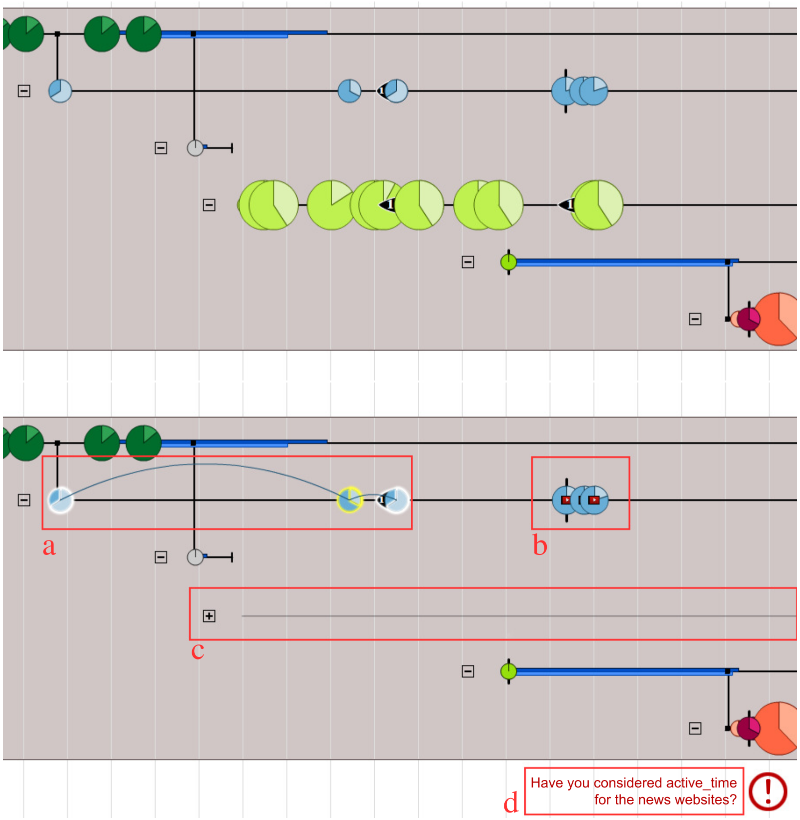

Fig. 1. A snapshot of the WebComets visualization (top). The set of considered adaptation techniques (bottom): (a) highlighting visualization elements (e.g., nodes), (b) adding or removing details through the control of displayed attributes (e.g., website icons), (c) adding or removing width information through the control of branching, and (d) showing custom hint messages.

were also informed of the fact that the visualization will make automatic changes based on their subjective states, and that these changes will involve either showing/highlighting relevant information or hiding irrelevant one. The detection of boredom and frustration were considered in two distinct sessions for all participants. In terms of adaptation, the following techniques have been considered:

– adding or removing details (through the control of attributes for multidimensional data, as each website is a multivariate node with over ten dimensions),
– adding or removing width information (through the control of branching),
– highlighting visualization elements (e.g., nodes or connections), and
– showing custom hint messages.

Fig. 1 highlights each of the six approaches (adding and removing were considered separately in both cases). In order to support a smooth transition and the

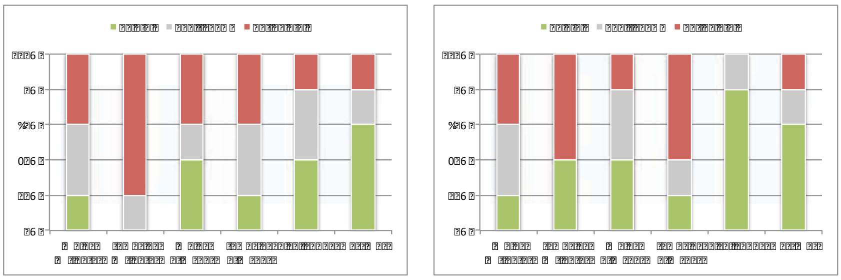

(a) user boredom (b) user frustration

Fig. 2. User rating of the automated adaptation triggered by the visualization system.

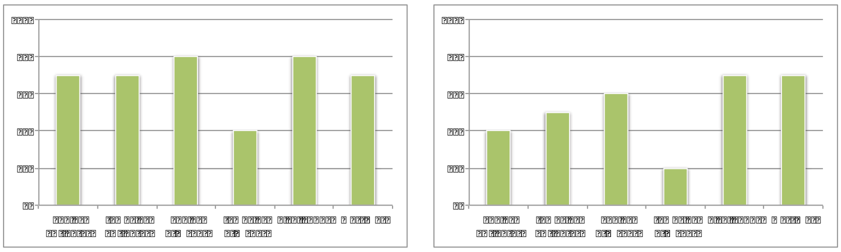

(a) user boredom (b) user frustration

Fig. 3. BCI measurements after each system adaptation capturing the ratio of significant drops in boredom and frustration measured in a time period of one minute after an adaptation was executed.

users' contextual awareness, changes were animated or faded in over a period of three seconds. Additionally, whenever a change that has been induced by the adaptive mechanism was executed, an exclamation sign would be displayed in the corner of the system (Fig. 1 (bottom)) to improve user awareness.

The hint messages were generated based on the different keywords present in the task statement. The various adaptation techniques were used in a random order, such that each user would see each adaptation twice. Also, tasks as well as the datasets had varying complexity, ranging from simple tasks to difficult ones, and from small datasets to complex ones. The increasing difficulty ensured that users would experience both boredom and frustration in a relatively short amount of time, as neither session exceeded 30 minutes for any of the participants.

4 Results

After having experienced all six adaptation approaches, the users were inquired about their usefulness and intrusiveness. Specifically, the participants were asked about each technique if they found that it was useful for influencing their emotional state positively. Their answers are highlighted in Fig. 2. It seems that in the case of user boredom, the adaptation techniques perceived the most helpful were element and connection highlighting, hint messages, and adding width information. Removing attributes seems to be least helpful when trying to overcome boredom. In the case of user frustration, element and connection highlighting as well as hints have again been perceived as most helpful, while hiding attributes and width information less so.

While participant feedback is vital, we also wanted to inspect the effects that the various adaptation methods had on the user emotional states. Therefore, we considered the emotional states for each user for a period of one minute after an adaptation technique has been triggered, in order to see if a significant drop in boredom, respectively frustration, could be detected. We defined a significant drop as emotional readings that were less than a 25% increase from the individual baseline for longer than half of the one minute period. The findings are captured in Fig. 3. Note that the drops in boredom or frustration may have multiple reasons. However, our results capture interesting differences. The drop in boredom is mostly consistent for all adaptation methods, except for the hiding of width information, after which a larger number of participants still experienced the same emotional state. This same pattern is visible for frustration, where removing width information did not seem to positively influence the frustration levels of many users. At the same time, hints and highlighting seem to be coupled to higher rates of frustration reduction.

When correlating the user rating with our measurements, we notice that participants tended to perceive the adaptation techniques more negatively or without any noticeable effects compared to our BCI readings. One possible reason for this could be the fact that four participants had difficulties accepting that the system would sometimes control which information would be added and which considered irrelevant and hidden. Thus, it seems that in terms of both effectiveness and intrusiveness, hint messages and element highlighting are the most reliable adaptive methods triggered to counteract boredom and frustration. Further research is required in order to investigate how visualizations could be adapted dynamically based on user emotions as well as how various emotions influence the performance of visualization users.

5 Conclusion

The current study focuses on the effectiveness of a set of adaptation techniques applied on an interactive visualization to influence user boredom and frustration. Our results show that approaches involving less intrusive methods like highlighting and help messages offer better results in guiding the user towards a more positive and aroused emotional state, and implicitly, towards a more focused interaction with the visualization tool. Further, we plan to extend this study to a larger set of emotions involving multiple participants as well as inspect the importance of emotional self-awareness in visualization.

References

1. Ahn, J.W., Brusilovsky, P.: Adaptive Visualization of Search Results: Bringing User Models to Visual Analytics, Information Visualization, 8(3):167–179, Palgrave Macmillan (2009)
2. Anderson, E. W., Potter, K. C., Matzen, L. E., Shepherd, J. F., Preston, G. A., Silva, C. T.: A user study of visualization effectiveness using EEG and cognitive load. In Proceedings of the 13th Eurographics / IEEE – VGTC Conference on Visualization (EuroVis '11), pp. 791–800, Aire-la-Ville, Switzerland (2011)
3. Axelrod, L., Hone K.: Uncharted passions: User displays of positive affect with an adaptive affective system. In Proceedings of the 1st International Conference on Affective Computing and Intelligent Interaction (ACII '05), pp. 890–897, Springer Berlin Heidelberg (2005)
4. Castellano, G., Kessous, L., Caridakis, G.: Emotion recognition through multiple modalities: Face, body gesture, speech. In Affect and Emotion in Human-Computer Interaction, volume 4868 of Lecture Notes in Computer Science, pp. 92–103, Springer Berlin Heidelberg (2008)
5. Cernea, D., Olech, P.-S., Ebert, A., Kerren, A.: EEG-based measurement of subjective parameters in evaluations. In Proceedings of the 14th International Conference on Human-Computer Interaction (HCII '11), poster paper, volume 174 of CCIS, pp. 279–283. Springer (2011)
6. Cernea, D., Olech, P.-S., Ebert, A., Kerren, A.: Measuring subjectivity supporting evaluations with the Emotiv EPOC neuroheadset. In Künstliche Intelligenz (KI), Special Issue on Human-Computer Interaction, 26(2):177–182. Springer (2012)
7. Cernea, D., Truderung, I., Kerren, A., Ebert, A.: WebComets: A tab-oriented approach for browser history visualization. In Proceedings of the International Conference on Information Visualization Theory and Applications (IVAPP '13), pp. 439–450, SciTePress (2013)
8. D'Mello, S., Calvo, R.A.: Beyond the basic emotions: What should affective computing compute? In CHI 2013 Extended Abstracts (2013)
9. Goldberg, J.H.: Relating Perceived Web Page Complexity to Emotional Valence and Eye Movement Metrics. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting 56, pp. 501–505 (2012)
10. Golemati, M., Halatsis, C., Vassilakis, C., Katifori, A., Lepouras, G.: A Context-Based Adaptive Visualization Environment. In Proceedings of the Conference on Information Visualization (IV '06), pp. 62–67, IEEE Computer Society Press (2006)
11. Gotz, D., Wen, Z.: Behavior-driven visualization recommendation. In Proceedings of the 13th International Conference on Intelligent User Interfaces (IUI '09), pp. 315–324, ACM (2009)
12. Horlings, R., Datcu, D., Rothkrantz, L. J. M.: Emotion recognition using brain activity. In Proceedings of the 9th International Conference on Computer Systems and Technologies and Workshop for PhD Students in Computing (CompSysTech '08), pp. 6:II.1–6:1, New York, NY, USA (2008)
13. Hudlicka, E., Mcneese, M.: Assessment of User Affective and Belief States for Interface Adaptation: Application to an Air Force Pilot Task. User Modelling and User-Adapted Interaction, 12(1):1–47, Kluwer Academic Publishers, Netherlands (2002)
14. Schneider-Hufschmidt, M., Khme, T., Malinowski, U., Eds.: Adaptive user interfaces: Principles and practice. Human Factors in Information Technology, North-Holland, Amsterdam (1993)

15. Kerren, A., Schreiber, F.: Toward the Role of Interaction in Visual Analytics. In Proceedings of the 2012 Winter Simulation Conference (WSC '12), pages 420:1-420:13, IEEE Computer Society Press. (2012)
16. Mähr, W.: eMotion: An outline of an emotion aware self-adaptive user interface. In Proceedings of the 2nd Scandinavian Student Interaction Design Research Conference (SIDeR '06), Gothenburg, Sweden (2006)
17. Mikhail, M., El-Ayat, K., El Kaliouby, R., Coan, J., Allen, J.J.B.: Emotion detection using noisy EEG data. In Proceedings of the 1st Augmented Human International Conference (AH '10), pp. 7:1–7:7, New York, NY, USA (2010)
18. Picard, R. W.: Affective Computing. MIT Press, Cambridge, MA, USA (1997)
19. Russell, J. A.: A circumplex model of affect. Journal of Personality and Social Psychology 39, pp. 1161–1178 (1980)
20. Shahid, S., Krahmer, E., Swerts, M., Melder, W., Neerincx, M.: You make me happy: Using an adaptive affective interface to investigate the effect of social presence on positive emotion induction. In Affective Computing and Intelligent Interaction and Workshops (ACII '09), pp. 1–6, IEEE Computer Society (2009)
21. Toker, D., Conati, C., Carenini, G., Haraty, M.: Towards adaptive information visualization: on the influence of user characteristics. In Proceedings of the 20th International Conference on User Modeling, Adaptation, and Personalization (UMAP '12), Masthoff, J., Mobasher, B., Desmarais, M.C., Nkambou, R. (Eds.), pp. 274–285, Springer Berlin Heidelberg (2012)

When to Adapt: Detecting User's Confusion During Visualization Processing

Cristina Conati, Enamul Hoque, Dereck Toker, Ben Steichen

Department of Computer Science, University of British Columbia
{conati, enamul, dtoker, steichen}@cs.ubc.ca

Abstract. In this paper, we discuss an approach to collect data on instances of user confusion during visualization processing. The long-term goal is to use this data to train classifiers that can detect instances of user confusion in real time, as triggers for adaptive interventions aimed at alleviating the confusion.

1 Introduction

The benefits of user-adaptive interaction have been shown in a variety of tasks and applications such as operation of menu based interfaces, web search, desktop assistance, and human learning [19]. There are three key decisions that need to be made when designing a user-adaptive system: (1) *what to adapt to*, namely understanding which individual user features should be considered for adaptation, including stable, long-term user traits (e.g., cognitive abilities, expertise, personality), as well as transitory, short-term states (e.g., current task, cognitive load, attention); (2) *when to adapt*, namely understanding when it is appropriate and/or necessary to provide adaptive support to the user, by identifying those situations in which the benefits of providing adaptive interventions outweigh their cost; (3) *how to adapt,* namely understanding how adaptation should be provided.

In this paper, we discuss issues related to the *when to adapt* decision in the context of designing user-adaptive visualizations. While there has been extensive work in investigating how to detect when a user needs help in fields such as Intelligent Tutoring Systems [1] or Adaptive Games [23], this is not the case in visualization. To our knowledge, the work by Gotz & Wen [22] is so far the only one that actively monitors real-time user behavior in order to infer such needs for intervention. In their work, interface action data (e.g., mouse clicks) are constantly tracked in order to detect suboptimal usage patterns. Once these repetitive patterns (determined empirically a priori) are detected, the system then triggers adaptive help.

This approach, however, does not easily transfer to situations in which it is hard to define a priori a set of appropriate interaction behaviors to perform given tasks with a visualization, as well as their suboptimal counterparts. This is the case, for instance, for visualizations that support open ended or exploratory tasks, or when one wants to consider interaction data beyond mouse or keyboard events, such as gaze data. Gaze data has been shown to have a great potential for providing information on a user's task, expertise, and other cognitive measures relevant for adaptation [14, 15], but it is more erratic in nature than interface actions, and it is less well understood in terms of what constitutes a priori effective/ineffective interface actions.

In this paper, we explore an alternative approach that involves collecting ground truth labels for specific salient episodes during interaction with a visualization, that may indicate the need of adaptive interventions. The long-term goal is to use these labels to train classifiers on interaction data consisting of both action logs and eye-tracking data, and to leverage these classifiers to detect in real time, for a new user, when adaptive interventions may be needed.

Collecting ground truth labels for building classifiers on relevant user states or processes can be a challenging endeavour. Here we propose one possible approach to collect labels relevant for building user-adaptive visualizations, which we are currently testing in a user study. However, this paper's primary aim is to open the discussion on the issue of *when to adapt* in user-adaptive visualizations, as opposed to provide well-defined solutions.

In the rest of the paper, we first briefly describe ValueCharts, the visualization that used as a test-bed for this research, and the study we are running to test, among other things, the label collection method. Then, we discuss the labelling approach that we have developed and present some preliminary results on its effectiveness.

2 ValueCharts and User Study

A ValueChart is a set of visualizations and interactive techniques intended to support decision-makers in inspecting linear models of preferences and evaluation [4]. Linear models are popular decision-making tools designed to help the decision-maker perform preferential choice under conflicting objectives, i.e., select the best option out of a set of alternatives. However, as models and their domain of application grow in complexity, model analysis can become a very challenging task. ValueCharts are intended to help decision makers deal with this complexity, based on a design driven by a detailed task model for preferential choice [2]. They have been extensively evaluated and shown to be quite effective [3, 12,17].

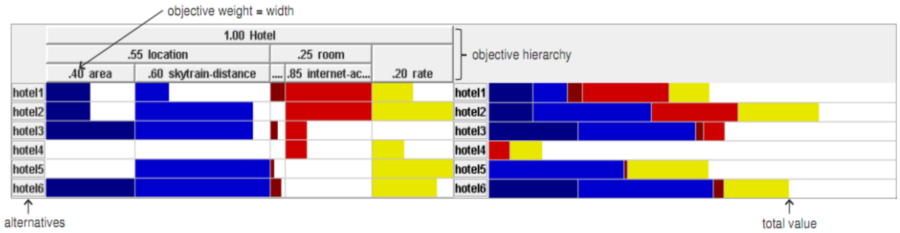

Figure 1: Sample Value Chart

Figure 1 shows an example of Value Chart for the simple preferential choice of selecting an hotel when traveling to a new city, out of six available alternatives. For the sake of simplicity, here we just describe the key features of ValueCharts. The relevant hotel attributes or *objectives* (e.g., area, skytrain distance, internet access, etc.) are arranged hierarchically and are represented in the top-left quadrant of the figure, forming the columns in the ValueChart display. The width of each column indicates the relative *weight* assigned to the corresponding objective (e.g., sky-train distance is much more important than area). The available alternatives (hotels here) are represented as the rows in the display. The cells in each row specify how the corresponding alternative fares with respect to each objective (i.e., the *value* of that objective for that alternative), indicated by the amount of filled color in the cell. So for instance, *hotel1* is far from the sky-train, but it has excellent internet access. In the rightmost quadrant, all values for each alternative are accumulated and presented as horizontal stacked bars, displaying the overall value of each alternative (e.g., in Figure 1, *hotel2* is the best alternative). Several interactive techniques are available in ValueCharts to support the inspection of the preference model. For instance, users can inspect the specific domain value of each objective (e.g., actual distance from the sky-train of *hotel1*); sensitivity analysis of objectives' weight is enabled by allowing the user to change the width of the corresponding column.

In the contest of an on-going project on devising theories and techniques for user adaptive visualizations, we are currently running a user study designed to evaluate the impact of a variety of user traits (e.g., perceptual speed, visual/verbal working memory, visualizations expertise, locus of control, etc.) on the effectiveness of two different versions of ValueCharts. The first version uses an horizontal layout to show the components of the decision making problem (see Figure 1), while the second version displays the same information by using a vertical layout (see Figure 2). We are comparing these two layouts because previous studies with ValueCharts suggest that they may not be equivalent with respect to the user's performance and preference. We test the impact of the aforementioned user traits because they were shown to have an effect during interaction with other visualizations [e.g., 6,16,18].

During the study we conduct, participants use each of the two Value Chart versions in two phases. The first phase (also known as *structured phase*) involves performing a selection of specific tasks in one of four available domains. The tasks are mainly related to retrieving information on the available decision alternatives (e.g., "how far is *hotel1* from the sky-train?", "How many hotels have better internet access than *hotel3*?", "List the 3 highest valued hotels"). The second phase (also known as *open-ended phase*) involves having a participant select a new domain and exploring it until the participant can identify a preferred alternative. Throughout the two phases, we track participants' gaze with a Tobii T120 desktop-mounted eye-tracker, similar to the study described in [16], because that study showed that gaze data can provide useful information on a user's individual differences and on the user's tasks [14, 15].

While one goal of this study is to ascertain whether the tested set of individual differences affect a user's performance with the two ValueChart layouts (i.e., help with the decision of *what to adapt to* in the context of using ValueCharts), we also wanted to leverage the study to provide data toward the question on *when to adapt*. Namely, we wanted to see whether we could find ways to collect information on salient points of the interactions that may benefit from adaptive interventions. The next section describes the approcah that we tested in this study.

3 Collecting labels of user confusion

The aim of providing real-time adaptive interventions is to help a user overcome situations that may generate a sub-optimal experience with an interface. For instance, adaptive interventions in an educational application can be generated when the user makes a mistakes or otherwise shows that she is not learning from the interaction [1]. Intuitively, adaptive interventions to improve a user's experience with a visualization would be suitable when the user is not processing the visualization appropriately, for instance when the user is uncertain or confused about where to look or how to interpret the visualization. Thus, in our ValueChart study we tried to devise a way to capture instances of user confusion during interaction, with the long-term goal of building a classifier user model for confusion detection, trained on these instances and on the related action and gaze patterns.

There have been a variety of methods proposed in the literature to capture confusion, mostly in the context of emotion modeling during the interaction with educational software. *Concurrent verbal protocols* involve having participants verbalize their thinking or feelings during interactions [e.g., 13]. We discarded this approach because of existing research indicating that concurrent protocols may alter a user's gaze patterns in unexpected ways, thus generating gaze data not representative of the user's attention patterns during a more naturalistic interaction [e.g., 9] *Retrospective verbal protocols* involve having participants look at a replay of the target interaction and try to verbalize their thinking or feelings at that time. While this approach has shown good results for collecting labels of emotion valence/arousal [11] or on the occurrence of one specific emotion (not including confusion) triggered on purpose via selected movie clips [10], it showed to be quite unreliable when subjects had to identify their naturally occurring emotions (including confusion), during interaction with an educational system [8]. We actually tried this approach in the study described in [16], quite similar in design to the study described here, and also found it inadequate. During pilot phases of that study, we first tried to ask subjects to generate retrospective protocols after each individual task, but because there are many rather short tasks, subjects quickly grew tired and the process interfered with the primary task. We then resorted to ask subjects to generate retrospective protocols at the very end of their study session, but at that point subjects had a difficult time re-generating their relevant states over the course of the complete interaction.

An alternative to verbal protocols is to *obtain labels from subjects via interface input*, e.g., buttons, pop-up windows, or other affordances that allow participants to select the relevant labels when the related episodes occur during interaction. This method has been successfully used to elicit information on user motivation and emotions during interaction with educational systems [7]. However, it has two main drawbacks. The first is that it can be hard to strike a balance between leaving it to the user to provide the information (e.g., via an ever-present interface button), which may results in not collecting a sufficient number of labels, vs forcing the user to provide as many labels as needed (e.g., via pop-up windows that cannot be dismissed), which may disrupt the interaction. The second drawback is that this approach does not provide as much information on the episodes of interest as verbal protocols do[1].

[1] An additional drawback when the user's gaze is tracked is the that presence of the interface affordance that allows for label provision changes the user's gaze patterns that would happen when the affordance is not present. The changes, however, are predictable and can be dealt with during gaze data processing.

As far as the first problem is concerned, in this study we decided to be conservative in terms of intrusiveness and we rely on the user's willingness to provide the labels via a *confusion button* placed on the side of the currently displayed ValueChart (see Figure 2 for an example with a vertical ValueChart). At the beginning of each study

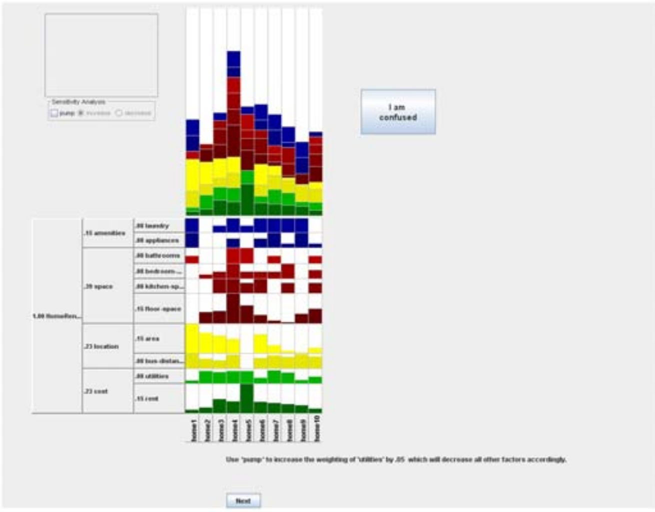

Figure 2: Vertical value chart with confusion button

session, the experimenter reads the participant a script that includes the following instructions to elicit usage of the confusion button:

"…We ask that you press this button any time you feel even slightly confused while performing the task. For instance, if you feel that you want to ask the experimenter a question about something, click the confusion button. If you are confused about the interface, click the confusion button, if you are confused about the wording of a question, click the confusion button. If you find a glitch or typo that confuses you, click the confusion button. These are just a few examples, to show that confusion can occur in many unforeseeable ways, and we want consider any type of confusion as being an OK reason to click the confusion button. Note that the system will not be able to give you any help to resolve your confusion. Pressing the confusion button will simply allow us to record moments in which the tasks or the visualizations make you confused. It is very important for the objectives of the experiment that we collect this information, so please take the time to press the confusion button when appropriate."

Preliminary data from the study participants that we ran so far (eight) indicate that the current set up is quite effective in eliciting presses of the confusion button. Five of the eight participants pressed the confusion button at least once. The total number of clicks by the 8 participants is 22, with an average of 2.75. presses per participant (Std dev =3.32). If this trend continues through all of the 30 participants planned to evaluate the impact of individual differences on ValueChart effectiveness, we will run additional subjects with the sole purpose of collecting enough episodes of user confusion for training a classifier that detects confusion from users' action and gaze data

To address the second problem in obtaining labels from user interface input, (i.e., lack of details on what may be causing a user to be confused), we combine the confusion button approach with a form of *focused retrospective verbal protocol*. Namely, we show to each participant video replays of interaction segments centered around presses of the confusion button, and for each of these segments, we ask the participant to explain why the confusion button was pressed. The elicited participant speech is then audio recorded. These verbal protocol sessions happen after each pair of structured and open-ended tasks performed with one of the two ValueChart versions, thus each participant undergoes at most two of these sessions. Pilots of this approach showed that participants are quite capable of generating explanations of confusion that happened during a structured task, after the subsequent open-ended tasks.

We expect that the information collected via focused retrospective verbal protocols will help us qualify the labels obtained via button presses in terms of *the reasons* for confusion, a fundamental piece of information to identify potential adaptive interventions that can help users resolve their confusion. For instance, one of the study's pilot subjects pressed the confusion button 3 times, always during a structured task with a Vertical ValueChart. For the first confusion button event, the participant said that she was confused because of the alternatives' names being vertical, which caused her difficulties in reading them. If this reason for confusion could be automatically identified, it might be alleviated by enabling a functionality that allows the user to mouse over the names to see them displayed more clearly (e.g., horizontally). The two subsequent confusion button events were both explained by the participant as due to difficulty during two different instances of a structured task that requires comparing alternatives with respect to a higher-level dimension in the objective hierarchy (e.g., location or room quality in our hotel selection example). This task requires to visually aggregate the values of the objectives under the target dimension and then comparing them. It could be facilitated, if confusion is detected, by visual props that help identify the aggregated blocks of values and draw the comparison. Another pilot subject clicked the confusion button twice, and for both occurrences the given reason was that the participant had not be able to tell the difference between the overall values of two alternatives, because they were placed in non-contiguous rows and were too similar to tell which one was greater/lower. A suitable adaptive intervention for this type of confusion might be a visual prop that, as before helps draw a comparison, but focusing on discriminating between small differences.

Discussion and Conclusions

In this paper, we discussed an approach to collect data on instances of user confusion during visualization processing, with the long-term goal to use this data to train classifiers that can detect instances of user confusion in real time. There are several open questions on this approach, that we would like to discuss at the workshop including: (i) how much data will be required to reliably identify user's confusion? (ii) Will it be possible to identify a taxonomy of confusion types, along with a mapping between elements of this taxonomy and types of adaptive interventions adequate to alleviate them? (iii) Which other user states in addition to confusion, or which other interaction episodes could serve as triggers for adaptive interventions? (iv) Which other approaches could be explored to collect data on the relevant user states?

References

1. B. P. Woolf. Building Intelligent Interactive Tutors, Student-Centered Strategies for Revolutionizing E-Learning, Published by Elsevier & Morgan Kaufmann, (2008).
2. Bautista J. and Carenini G., An Integrated Task-Based Framework for the Design Evaluation of Visualizations to Support Preferential Choice," Proc. Advanced Visual Interfaces (AVI 2006), (2006), 217-224.
3. Bautista J., Carenini G., An Empirical Evaluation of Interactive Visualization Techniques for Preferential Choice. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, (AVI08), Naples, Italy, (2008).
4. Carenini G. and Loyd J. ValueCharts: Analyzing Linear Models Expressing Preferences and Evaluations, Proc. The International Working Conference on Advanced Visual Interfaces (AVI 2004), (2004), 150-157.
5. Conati C. and Maclaren H.. Modeling User Affect from Causes and Effects. In *Proceedings of UMAP 2009*, First and Seventeenth International Conference on User Modeling, Adaptation and Personalization, Springer, (2008).
6. Conati, C., and Maclaren, H., Exploring the Role of Individual Differences in Information Visualization. In Proceedings of the working conference on Advanced visual interfaces (AVI '08). ACM, New York, NY, USA, (2008), 199-206.
7. deVincente, A., & Pain, H.. Motivation Self-Report in ITS. AIED'99, 9th International Conference on Artificial Intelligence in Education, Le Mans, France, (1999), 651-659.
8. D'Mello, S.K., S. K., Craig, S.D., Witherspoon, A. W., McDaniel, B. T., and Graesser, A. C. Automatic Detection of Learner's Affect from Conversational Cues. User Modeling and User-Adapted Interaction, 18(1-2), (2008) , 45-80.
9. Kim B., Dong Y., Kim S., and Lee K. Development of integrated analysis system and tool of perception, recognition, and behavior for web usability test: with emphasis on eye-tracking, mouse-tracking, and retrospective think aloud. In *Proceedings of the 2nd international conference on Usability and internationalization (UI-HCII'07)*. Springer-Verlag, Berlin, Heidelberg, (2007), 113-121.
10. Lisetti, C. L., & Nasoz, F.. Using Non-invasive Wearable Computers to Recognize Human Emotions from Physiological Signals. EURASIP Journal on Applied Signal Processing,11, (2004), 1672-1687.
11. Peter, C., & Herbon, A.. Emotion Representation and Physi-ology Assignments in Digital Systems. Interacting with Computers,18(2), (2006), 139-170.
12. Pommeranz A., Broekens J., Wiggers P., Brinkman W.P., and Jonker C.M., "Designing Interfaces for Explicit Preference Elicitation: a User-Centered Investigation of Preference Representation and Elicitation Process," J. User Modeling and User-Adapted Interaction, vol. 22, no. 4-5, (2012), 357-397.
13. Sidney K. D'Mello, Scotty D. Craig, Jeremiah Sullins, Arthur C. Graesser: Predicting Affective States expressed through an Emote-Aloud Procedure from AutoTutor's Mixed-Initiative Dialogue. I. J. Artificial Intelli-gence in Education 16(1), (2006), 3-28.
14. Steichen, B., Carenini G., and Conati C. Adaptive Information Visualization: Using Gaze Data to Infer Visualization Types, Tasks, and User Characteristics. In

Proceedings of IUI 2013, International Conference on Intelligent User Interfaces, ACM, (to appear), (2013).
15. Toker D., Conati C., Steichen B., and Carenini G.. Individual User Characteristics and Information Visualization: Connecting the Dots through Eye Tracking. In *Proceedings of CHI 2013, ACM SIGCHI Conference on Human Factors in Computing Systems*, (to appear), (2013).
16. Toker, D., Conati, C., Carenini, G., and Haraty, M.. Towards Adaptive Information Visualization: On the Influence of User Characteristics. In *Proceedings of UMAP 2012, the 20th International Conference on User Modeling, Adaptation, and Personalization*. Springer LNCS 7379, (2012), 274-285.
17. Wongsuphasawat K., Plaisant C., Taieb-Maimon M., and Shneiderman B., "Querying Event Sequences by Exact Match or Similarity Search: Design and Empirical Evaluation," J. Interacting with Computers, 24(2), (2012), 55-68.
18. Ziemkiewicz, C., et al. How Locus of Control Influences Compatibility with Visualization Style. In *Proc. IEEE VAST*, (2011), 81-90.
19. Jameson, A. "Adaptive Interfaces and Agents" in Human-Computer Interface Handbook, eds J.A. Jacko and A. Sears, 305-330, 2003.
20. Grawemeyer, B. Evaluation of ERST – an external representation selection tutor. In Proceedings of the 4th international conference on Diagrammatic Representation and Inference, 154-167, 2006.
21. Green, T. M. & Fisher, B. Towards the personal equation of interaction: The impact of personality factors on visual analytics interface interaction. In IEEE Visual Analytics Science and Technology, 203-210, 2010.
22. Gotz D., & Wen, Z.. Behavior Driven Visualization Recommendation. ACM Int. Conf. on Intelligent User Interfaces, 315-324, 2009.
23. Peirce, N., Conlan, O., Wade, V. Adaptive Educational Games: Providing Non-invasive Personalised Learning Experiences. In Second IEEE International Conference on Digital Games and Intelligent Toys Based Education, 2008.

Adaptive Visualization of Research Communities

Martijn de Jongh, Patrick M. Dudas, and Peter Brusilovsky

University of Pittsburgh
{mad159,pmd18,peterb}@pitt.edu

Abstract. *Adaptive visualization approaches attempt to tune the content and the topology of information visualization to various user characteristics. While adapting visualization to user cognitive traits, goals, or knowledge has been relatively well explored, some other user characteristics have received no attention. This paper presents a methodology to adapt a traditional cluster-based visualization of communities to user individual model of community organization. This class of user-adapted visualization is not only achievable, but expected due to real world situation where users cannot be segmented into heterogeneous communities since many users have affinity to more than one group. An interactive clustering and visualization approach presented in the paper allows the user communicate their personal mental models of overlapping communities to the clustering algorithm itself and obtain a community visualization image that more realistically fits their prospects.*

Keywords: Community detection, user model, social network, adaptive visualization

1 Introduction

The increased popularity of social networking research attracted attention to the problem of community discovery and visualization. Since the work by Girvan & Newman [1] many different approaches to discover communities (i.e., clusters of similar users) in social networks and other social systems were suggested and explored - see Fortunato [2] for a comprehensive overview. In addition, a number of packages such as Gephi [3], Pajek [4], and Ucinet [5] were developed to visualize the results of community discovery to the end users.

Despite a relatively large volume of work on the topic, little attention was paid to take into account user mental models and domain knowledge when presenting visual structure of the community. Existing visualization programs tend to represent a simplified community organization formed by a number of distinct, non-overlapping communities that are displayed universally to all users of the system.

The novelty of our approach is the understanding that different users can form different models of community organization. They can recognize different subcommunities within the same large community due to their unique personal knowledge and domain expertise. For example, a researcher working on the application of machine learning to user modeling could be considered as a user modeling researcher by one group of users and as a machine learning researcher by another. Existing visualizations do not recognize these individual preferences and as a result produce community structure that might be acceptable only by a subset of the target users.

This paper presents our adaptive platform to create an interactive community visualization that can take into account user preferences on community organization and provide dynamic adaption to the user model of community structure. To collect user preferences, the system allows the users to identify cliques of researchers that, from their prospect, should belong to the same community. These user-defined cliques are considered by an interactive clustering approach developed by us along with the original data, describing similarity between researchers, to produce a user-adapted cluster visualization.

The presentation of our approach is organized as follows. We start with the interface part of the approach explaining how our system allows the users to specify their preferences. Then we provide the methodology of our approach explaining the interactive clustering algorithm and the visualization approach that we use to present its results. We conclude the paper after a discussion of similar project and future work.

2 Interactive Visualization

In the process of user-adaptive clustering, the users interact directly with a community visualization that shows the community topology and the currently identified set of groups. Figure 1 shows a mapping of authors that have published in the UMAP conference series connected by co-authorship and similarity links[1].The visualization provides special affordances to guide the user through the interaction process [6].

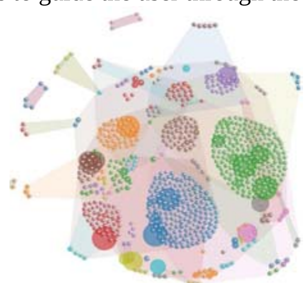

Fig. 1. The community visualization showing UMAP (User Modeling, Adaptation and Personalization) data

[1] This data was extracted from the DBLP [9] bibliography database, which created a 766 vertex and 8038 edge network. The similarity measure used is based on the Jaccard index [10].

The node saliencies include 1) nodes size based on degree centrality and 2) appended pie charts to highlight overlap and to what degree. The edges reflect co-authorship Jaccard similarity by increases the thickness of the line. The use of both convex hulls and inner-cluster distance minimized versus inter-cluster distance to outline and exacerbate group overlap and to differentiate the clusters. Vertices and text boxes provide a pointer cursor to make them known as selectable objects. To provide their individual views on community organization, the user has the option (and is encouraged) to select multiple vertices as a group and declare it as a clique that should belong to the same sub-community. The pictures below explain this process in detail.

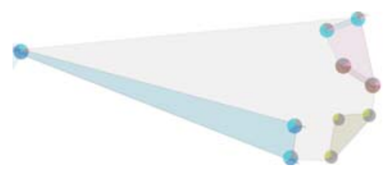

Fig. 2. Pie Charts Showcasing LPA correlation

Figure 2 shows a highlighted section of the graph that will be subject of user-defined groupings. The user selects all the nodes using control-clicks which is comparable to traditional multi-file select [7]. The nodes change their glyphs (circles to rhombus) and distinctly change to their most dominant group identity as highlighted in Figure 3. This provides two-dimensions of clustering aesthetics, both color (machine-derived clusters) and glyph changes (user-defined clusters).

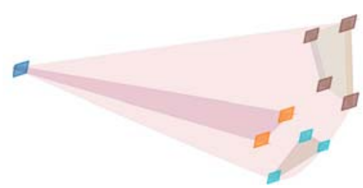

Fig. 3. User-Defined Clusters

After a set amount of seconds the new user-defined clique is automatically passed to our interactive clustering algorithm, presented in details in the next section, and after its completion the cluster assignments of the current dataset are updated in the visualization display. Finally, the user-defined group is differentiated by changing the opacity to be completely opaque, Figure 4. This process continues and the user-defined groups will be differentiated from one another by using new glyphs.

Fig. 4. Final Aesthetic to Highlight Completed Nodes[2]

3 Extending Label Propagation

To cluster the graph nodes, we applied an extended version of the Label Propagation Algorithm (LPA), which was proposed by Raghavan, Albert, and Kumara [8]. LPA iteratively determines the final cluster assignments. It initializes all vertices with a unique label, and then proceeds to update vertex labels by checking the labels of their neighbors. The most frequently occurring label among a vertex's neighbors is chosen as its new label. Ties are broken randomly. During all iteration, all vertices are (possibly) assigned a new label asynchronously. The process continues until it converges on a stable set of label assignments, which usually is within a few iterations.

We have extended the original LPA to allow it to be an active part in our interactive clustering visualization. We added to ability to run the algorithm on weighted graphs, changing the label picking criterion from the most frequent label of the neighbors to a version where the frequencies are modified by the edge weights. Furthermore, the labeling process was extended to allow for discovery of overlapping clusters. The calculated cluster label weights are further weighted by the proportions they appear in the label distribution of the adjacent vertices. A weight w_j for cluster label j is calculated using edge weights e_i of adjacent vertex i and cluster label proportion p_{ij} of adjacent vertex i for label j:

$$w_j = \sum_{i=1}^{n} e_i \frac{p_{ij}}{\sum_{j=1}^{c} p_{ij}}$$

The proportion weights p_{ij} for the new set of cluster labels of a vertex *i* are calculated after selecting the top *n* labels. The calculated weights of these *n* labels are normalized and stored with the new cluster labels of the vertex.

$$p_{ij} = \frac{w_{ij}}{\sum_{j=1}^{n} w_{ij}}$$

A very useful property of LPA is its near-linear time complexity [8]. Its swiftness makes it feasible to do the cluster calculations online. Specifically, we have the algorithm running in the background, waiting for updated information coming from

[2] This is a force-based graph and with new cluster assignments the subsequent topology will change. To keep consistent with explanation of the interactive process, we rotated the aforementioned areas to make it easier to follow.

the visualization component. To facilitate a fluid, adaptive cluster visualization experience, we have improved LPA to make this type of interaction feasible. The original LPA paper proposed an approach to seed a few vertices with cluster labels and leaving the rest unlabeled, allowing for clusters to form around those seeds. We have modified this approach by allowing vertices to be fixed. Once a vertex is fixed, it will no longer update its cluster label. This allows clusters to grow around vertices in a similar way to the original approach with the added benefit that we can choose to fix vertices in an already existing complete set of cluster assignments and rerun the algorithm on this cluster assignment to update it. Second, featured prominently in our visualization, is granting the user the ability to group vertices. Effectively, this allows grouped vertices to behave as a single vertex. When any of the vertices of the group updates its (set of) label(s), all others follow suit. Groups can also be fixed and if one of its members is a fixed vertex, the whole group becomes fixed.

4 Related Work

As mentioned above, mainstream software packages available for network data such as Gephi [3], Pajek [4] or Prefuse [11] are all limited to non-overlapping clustering approaches and do not allow the user the ability to put vertices into new groups without having to change meta-data about the vertex itself. Some research systems, however, explore both interactive clustering and overlapping communities that are distinguishing features of our approach.

Apolo [12] provides an approach in incorporating both user interactions and machine learning in large datasets. They accomplish this by building on a single vertex and as users provide a paper of interest to interface, the network-like visualization builds by providing cited work and visual clues based on number of citations and relevance. Building on this visualization, TourViz [13] divides sub-domains of interest (to the user) into convex hulls to help segment multiple topics of interest. In this context, our advancement takes into account the entire network structure, providing overall topology. Also, we provide not only the utilization of color changes (to distinguish group assignment), but also glyph distinction by user-defined clusters and opacity changes to discern vertices already selected and grouped.

5 Future Work

To continue this work we want to validate first the claim that our approach allows for easy understanding of community identity using convex hulls and adaptive clustering. To do this, we will need to provide a user-study of this mechanism and show that a mixture between user-defined and machine learned clustering can build an optimal and accurate model of the network topology and the user's mental model.

There is also a novel and yet strikingly obvious need to understand what it means to belong to multiple overlapping communities. Much work has been done in social capital that illustrates the strength of weak ties in bridging multiple communities [14]. We are interested in studying these networks more in depth to see if these vertices that fall between multiple communities can be defined more precisely using arguments from social capital. Gilbert [15] supports the claims made in this paper in regards to a spectrum to vertex-to-vertex variability and we believe that social capital and overlapping communes go hand in hand.

6 Conclusion

This paper presents an approach that allows to adapt a traditional cluster-based visualization of communities to user individual model of how the communities are organized. This kind of user-adapted visualization is possible because in a real world situation there are many alternative ways to segment users into heterogeneous groups since many users have affinity to more than one group. An interactive clustering and visualization approach presented in the paper allows the user to communicate their personal mental models of the communities to the clustering the algorithm itself and obtain community visualization picture that fits their expectations. To provide the user a fluid user-interface, both the visualization and the modified LPA was adapted to handle the size of the network and the interactions. Modifications were made to the visualization to showcase the communities in better quality and minimize edge crossing. The LPA was adjusted to allow for fixed vertices within the algorithm, allowing it to obtain an optimal solution in a relatively small amount of time. We believe that as networks are examined more in-depth, that platforms like this that take both visual information and user involvement can balance out both the human mental model of the network and the machine learning techniques used for efficiency.

7 Reference

1. Newman, M.E.J. and M. Girvan, *Finding and evaluating community structure in networks.* Physical review E, 2004. **69**(2): p. 026113.
2. Fortunato, S., *Community detection in graphs.* Physics Reports, 2010. **486**(3): p. 75-174.
3. Bastian, M., S. Heymann, and M. Jacomy, *Gephi: An open source software for exploring and manipulating networks.* 2009.
4. Batagelj, V. and A. Mrvar, *Pajek-program for large network analysis.* Connections, 1998. **21**(2): p. 47-57.
5. Borgatti, S.P., M.G. Everett, and L.C. Freeman, *Ucinet for Windows: Software for social network analysis.* 2002.
   networks visualization. in *Smart Graphics*. 2008: Springer.
6. Dieberger, A., et al., *Social navigation: techniques for building more usable systems.* interactions, 2000. **7**(6): p. 36-45.
7. Ritter, A. and S. Basu. *Learning to generalize for complex selection tasks*. in *Proceedings of the 14th international conference on Intelligent user interfaces*. 2009: ACM.
8. Raghavan, U.N., R. Albert, and S. Kumara, *Near linear time algorithm to detect community structures in large-scale networks.* Physical Review E, 2007. **76**(3): p. 036106.

9. Ley, M. *The DBLP computer science bibliography: Evolution, research issues, perspectives*. in *String Processing and Information Retrieval*. 2002: Springer.
10. Jaccard, P., *Etude comparative de la distribution florale dans une portion des Alpes et du Jura*. 1901: Impr. Corbaz.
11. Heer, J., S.K. Card, and J.A. Landay. *Prefuse: a toolkit for interactive information visualization*. in *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2005: ACM.
12. Chau, D.H., et al. *Apolo: making sense of large network data by combining rich user interaction and machine learning*. 2011: ACM.
13. Chau, D.H., et al. *TourViz: interactive visualization of connection pathways in large graphs*. in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2012: ACM.
14. Burt, R.S., *The social capital of structural holes*. The new economic sociology: Developments in an emerging field, 2002: p. 148-190.
15. Gilbert, E. and K. Karahalios. *Predicting tie strength with social media*. in *Proceedings of the 27th international conference on Human factors in computing systems*. 2009: ACM.

Visual Variables in Adaptive Visualizations

Kawa Nazemi and Jörn Kohlhammer

Fraunhofer Institute for Computer Graphics Research (IGD), Darmstadt, Germany
{kawa.nazemi, joern.kohlhammer}@igd.fraunhofer.de

Abstract. Visualizations provide various variables for the adaptation to the usage context and the users. Today's adaptive visualizations make use of various visual variables to order or filter information or visualizations. However, the capabilities of visual variables in context of human information processing and tasks are not comprehensively exploited. This paper discusses the value of the different visual variables providing beneficial and more accurately adapted information visualizations.

Keywords: Adaptive Information Visualizations, Visual Variables, Visualization Tasks

1 Introduction

Information visualization and visual analytics provide valuable techniques for interacting with huge amounts of data and solving information-related tasks. [9] In the recent past, researchers from different disciplines recognized the need for a more human-centered design process in visualizations ([1], [15]). These human-centric thoughts resulted in research and development of adaptive information visualizations (AIV), which consider various contextual aspects to adapt visualizations ([1], [2], [23], [13], [14]). In particular, the user plays an increasing role, with her behavior, pre-knowledge and aptitudes ([15], [23], [1], [13]). For involving users in the visualization adaptation process, various systems were developed and successfully evaluated [1]. Although, the users' behavior was analyzed with various techniques and different modeling approaches, the adaptation capabilities of visual information representation were not comprehensively exploited. The focus of visual adaptation was rather on only one value of the available visual variables. Similar to recommendation systems, filtering [12], selection ([13], [14]) or ordering [2] was adopted to these systems.

This paper investigates the distinguishability of visual variables based on the early definition of Bertin [3]. Further findings from the area of vision perception and various task classifications are investigated to outline the need for the adaptation of both, *layout* and *presentation* as proposed in [18]. We believe that one main question for adaptive visualizations still remains, namely how to adapt the visual representation [6]. The main contribution of this paper is an enhancement and application of cognitive models to the adaptation of visual representations. Observations already showed that the use of various visual variables results in different level of task-solving accuracy [8].

2 Interaction and Tasks in Visualizations

Interaction with visualizations enables the dialog between user and the visual representation of the underlying data. The interactive manipulation of the data, the visual structure or the visual representation provides the ability to solve various tasks and discover insights. The main goal of interactive visual representations still remains the acquisition of knowledge [16]. The term “task” in context of information visualization is often used ambiguously. A dissociation of interactions and tasks in visualizations is rarely performed, whereas the knowledge about the task to be solved with the visualization is of great importance for its design and therewith for the adaptation. We have investigated various task and interaction classifications ([4], [7], [10], [11], [17], [19], [22], [26], [28], [33], [34]) to a find an abstract view on visual tasks for a mapping to the human information processing. Therefore all the tasks and interactions were categorized into three abstract levels: “search”, “explore”, and “analyze”. Figure 1 illustrates the identified high-level tasks and their assigned interactions and subtasks derived from the existing visual task classifications.

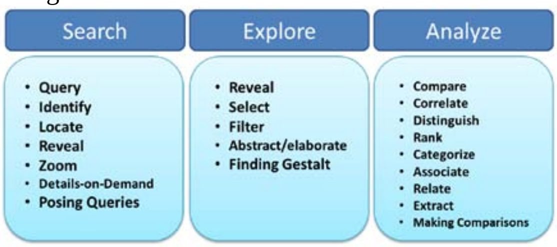

Fig. 1. High-level tasks with assigned subtasks and interactions

3 Visual Variables and Human Perception

The differentiated investigation of visual variables allows a more goal-directed adaptation to users’ needs and tasks. [20] An early definition and differentiation of visual variables was proposed by Bertin [3]. He differentiates between visual variables that use the two dimensions of a plane to encode information through graphical marks and those, which encode information through their relationship above the plane. A graphical mark is defined by basic geometrical elements of points, lines and areas. The position of a mark indicates a meaning between the values of the two dimensions. Marks could be changed through their *size*, *saturation*, *texture*, *color*, *orientation* and *shape*. These features (retinal variables) [27] can further be classified in ordinal, quantitative, selective and associative variables [3].

The second class of Bertin’s visual variables (*imposition*) encodes information through their relationships to each other above the plane. He differentiates this based on how these relationships can be visually illustrated in *diagrams*, *networks*, *maps*, and *symbols*. [3] The main value of Bertin’s classification in the context of this paper is the differentiation between the graphical *layout* and visual (or retinal) variables. The differentiation is of great importance for adapting visualizations, which is also supported by results in cognitive science (e.g., *feature integration theory or guided search model*). However, the established model of Card et al. [5], make use of visual variables (*Visual Structure*) but does not differentiate the layout from presentation. Different and independent studies illustrated a rapid and parallel processing of the retinal variables by the low-level human vision ([24], [25], [29], [30], [31], [32]). The so called "pop-out effect" makes use of the human's parallel vision processing and guides the attention to the related location on the screen [32]. Ware proposes a three-tiered model by considering both the pre-attentive parallel processing and attentive stages of human vision [27]. He subdivides the attentive processing of visual information into a serial stage of *pattern recognition* and a further stage of *sequential goal-directed processing*. [27]. While the pre-attentive stage refers to the retinal variables, the attentive stages (or post-attentive stages) require a serial (or sequential) processing of information, which can be provided by visual information of object relationships over the plane [27]. This aspect of attentive serial processing, in particular by separating the visual retinal variables and layout information was also investigated by Rensink ([20], [21]). In his *coherence theory* and the *triadic architecture* the strict differentiation of *layout* and the low-level retinal variables was proposed in terms of the dynamic generation of a visual representation ([20], [21]). Rensink's *triadic architecture* starts with the low-level vision (pre-attentive) and is generally similar to Ware's model. The most important aspect in this context is the unification of *layout*. Rensink proposes that one important aspect of the scene structure is *layout*, "without regards to visual properties or semantic identity" (p. 36, [20]).

Based on the three introduced models and the results on research of parallel and serial processing we introduce a model for visual adaptation on an abstract level by considering the high-level visual tasks as a foundation for discussing the adaptable variables of visualizations (Figure 2). This is a refined model of the previous work on an adaptive reference model [18] and proposes the use of the different visual variables for adapting visualizations to the users' knowledge and tasks.

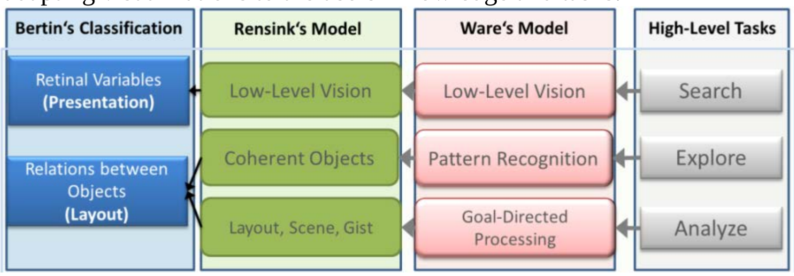

Fig. 2. Model of layout and presentation in adaptive visualizations

4 Conclusions

We introduced various models for differentiating visual variables in context of human information processing and lined out that a separated view on layout (visualization types) and presentation (retinal variables) is important for an accurate and beneficial adaptation of visualizations.

References

1. Ahn J.-W.: Adaptive Visualization for Focused Personalized Information Retrieval. PhD thesis, School of Information Sciences, University of Pittsburgh, 2010.
2. Ahn, J.-W., Brusilovsky, P.: Adaptive visualization of search results: Bringing user models to visual analytics. Information Visualization 8, 8 (3) (September 2009), 180–196.
3. Bertin, J.: Semiology of graphics. University of Wisconsin Press, 1983.
4. Buja, A., Cook, D., Swayne, D. F.: Interactive high-dimensional data visualization. Journal of Computational and Graphical Statistics 5 (1996), 78–99.
5. Card, S. K., Mackinlay, J. D. and Shneiderman B. (Eds.). Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 1999.
6. Chen, C.: Information Visualization: Beyond the Horizon. Springer-Verlag London Limited 2006.
7. Chuah, M. C., Roth, S. F.: On the semantics of interactive visualizations. In INFOVIS '96: Proceedings of the 1996 IEEE Symposium on Information Visualization, IEEE Computer Society.
8. Cleveland, W. S., McGill, R.: Graphical Perception: Theory, experimentation and application to the graphical methods. J. Am. Stat. Assoc. 79, 387 1984.
9. Fekete, J.-D., Wijk, J. J., Stasko, J. and North C. (2008): The Value of Information Visualization. In *Information Visualization*, A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North (Eds.). Lecture Notes in Computer Science, Vol. 4950. Springer-Verlag, Berlin, Heidelberg 1-18.
10. Fluit, C., Sabou, M., Van Harmelen, F.: Supporting user tasks through visualisation of light-weight ontologies. In Steffen Staab and Rudi Studer (eds.) Handbook on Ontologies. International Handbooks on Information Systems. Springer, 2004, pp. 415–434.
11. Fluit, C., Sabou M., Van Harmelen, F.: Ontology-based information visualization: Towards semantic web applications. In In Geroimenko, V. and Chen, C. (eds.): Visualizing the Semantic Web. XML-Based Internet and Information Visualization. Second Edition. Springer-Verlag London Limited, 2006.
12. Golemati, M., Halatsis C., Vassilakis, C., Katifori, A., Peloponnese, U. O.: A Context-Based Adaptive Visualization Environment. In Proceedings of the conference on Information Visualization (Washington, DC, USA, 2006), IV '06, IEEE Computer Society, pp. 62–67.
13. Gotz, D., and Wen, Z.: Behavior-driven visualization recommendation. Proc. 13th intl. conf. on Intelligent user interfaces, 2009, pp. 315-324.
14. Gotz, D., When, Z., Lu, J., Kissa, P., Cao, N., Qian, W. H., Liu S. X., Zhou M. X.: HARVEST: an intelligent visual analytic tool for the masses. In Proceedings of the first international workshop on Intelligent visual interfaces for text analysis (New York, NY, USA, 2010), IVITA '10, ACM, pp. 1–4.
15. Keim, D., Andrienko, G., Fekete, J.-D., Görg, C., Kohlhammer J., Melancon G.: Visual analytics: Definition, process, and challenges. In Information Visualization, Kerren A., Stasko J., Fekete J.-D., North C., (Eds.), vol. 4950 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2008, pp. 154–175.
16. Keim , D., Mansmann, F., Schneidewind, J., Thomas, J. and Ziegler, H.: Visual Analytics: Scope and Challenges. In *Visual Data Mining*, Simeon J. Simoff, M. H. Böhlen, and A. Mazeika (Eds.). Lecture Notes In Computer Science, Vol. 4404. Springer-Verlag, Berlin, Heidelberg 76-90.

17. Keller, P. R., Keller, M. M.: Visual Cues: Practical Data Visualization. Los Alamitos, CA. IEEE Computer Society Press, 1994.
18. Nazemi, Kawa; Stab, Christian; Kuijper, Arjan: A Reference Model for Adaptive Visualization Systems. In: Jacko, J. A. (Ed.): Human-Computer Interaction: Part I: Design and Development Approaches. Berlin, Heidelberg, New York : Springer, 2011, pp. 480-489 (LNCS) 6761).
19. Pike, W. A., Stasko, J., Chang, R., O'Connell, T. A.: The science of interaction. Information Visualization 8, 4 (Dec. 2009), 263–274.
20. Rensink, R. A.: The dynamic representation of scenes. Visual Cognition. 7 (2000), 17–42.
21. Rensink, R. A.: Change detection. Annual Review of Psychology 53 (2002), 245–277.
22. Shneiderman, B.: The eyes have it: A task by data type taxonomy for information visualizations. In VL (1996), pp. 336–343.
23. Steichen, B., Carenini, G., Conati, C.: Adaptive Information Visualization - Predicting user characteristics and task context from eye gaze. Poster and Workshop Proceedings of the 20th conference on User Modeling, Adaptation and Personalization (UMAP 2012), CEUR Workshop Proceedings, ISSN 1613-0073, Vol-872.
24. Treisman, A. M., Gelade, G.: A feature-integration theory of attention. Cognitive Psychology 12:1 (1980), 97–136.
25. Treisman, A. M., Gormican, S.: Feature analysis in early vision: Evidence from search asymmetries. Psychological Review 95:1 (1988), 15–48.
26. Ward, M., Grinstein G., Keim, D.: Interactive Data Visualizations Foundations, Techniques, and Applications. A K Peters, Ltd. Natick, Massachusetts, 2010.
27. Ware, C.: Information Visualization Perception for Design. Morgan Kaufmann (Elsevier), 2013.
28. Wehrend, S., Lewis, C.: A problem-oriented classification of visualization techniques. In VIS '90: Proceedings of the 1st conference on Visualization '90 (Los Alamitos, CA, USA, 1990), IEEE Computer Society Press, pp. 139–143.
29. Wolfe J. M.: Guided search: An alternative to the feature integration model for visual search. Journal of Experimental Psychology Human Perception & Performance 15, 3 (1989), 419–433.
30. Wolfe J. M.: Guided search 2.0: A revised model of visual search. Psychonomic Bulletin & Review 1,2 (1994), 202–238.
31. Wolfe, J., Treismann, A. M.: What shall we do with the preattentive processing stage use it or lose it? Poster presented at the Third Annual Meeting of the Vision Sciences Society, May 2003. Sarasota, FL.
32. Wolfe J. M.: Guided search 4.0: Current progress with a model of visual search. W. Gray (Ed.), Integrated Models of Cognitive Systems (2007), 99–119.
33. Yi, J. S., Ah Kang, Y., Stasko, J., Jacko J.: Toward a deeper understanding of the role of interaction in information visualization. Visualization and Computer Graphics, IEEE Transactions on 13, 6 (nov.-dec. 2007), 1224 –1231.
34. Zhou, M. X., Feiner, S. K.: Visual task characterization for automated visual discourse synthesis. In CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems (New York, NY, USA, 1998), ACM Press/Addison-Wesley Publishing Co., pp. 392–399.

Seeing how you're looking - Using Real-Time Eye Gaze Data for User-Adaptive Visualization

Ben Steichen, Oliver Schmid, Cristina Conati, Giuseppe Carenini

Department of Computer Science, University of British Columbia
{steichen,oschmid,conati,carenini}@cs.ubc.ca

1 Introduction

In the field of cognitive and perceptual psychology, the use of eye tracking has long been established as a suitable means for analyzing user attention patterns in information processing tasks [11]. More recently, the fields of human-computer interaction and information visualization have similarly started to use eye-tracking technology to investigate trends and differences in user attention patterns and cognitive processing, e.g. for different interfaces/visualization types [13][7], task types (e.g. reading vs. mathematical reasoning) [8], or user abilities (e.g. cognitive abilities [16], expertise [12]). Researchers have also started to use machine-learning techniques on gaze data to predict, for example, user intents [3], cognitive processes [5][14], or student learning [10][4]. Similarly, in our own work [15], we have shown that we can predict a number of visualization task types and user characteristics using simple machine learning techniques on a broad set of eye gaze measures and statistics.

However, these studies have generally only attempted to gain insights or generate models using *off-line* processes. In terms of actually using eye tracking in *real-time* systems, most early research has focused on designing gaze-directed interaction, i.e. using gaze data as a direct input to control a system [9]. Such interaction techniques are now also appearing in mainstream, commercial applications, such as TVs [2] or mobile phones [1]. Researchers have recently started to build systems that go beyond simple gaze-direction, by performing dynamic adaptations to basic gaze patterns (e.g. providing feedback to 'unattentive' students looking away from the screen [6]). However, such systems typically only use basic eye fixation counts/lengths, rather than the full spectrum of eye gaze features explored in the machine-learning experiments described above. They are hence quite limited in terms of adaptation potential. In the next section, we present a system that bridges this gap by providing a real-time, feature-rich eye-gaze analysis service that can be leveraged by user-adaptive systems.

2 Real-Time Gaze Analysis Service

As shown in Figure 1, our service consists of a web server application that continuously receives raw gaze data from an eye tracker. When a client application (which may or may not be located on the same machine as the web server and/or eye-tracker) places a request for analysis, the server calculates real-time statistics, either starting from a specific start time (set previously by the client application), or for a specific time window (as specified in the server configuration file, e.g. the last 10 seconds).

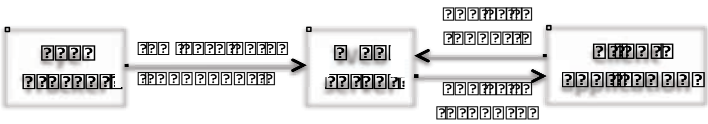

Figure 1. Real-time eye-tracking architecture

The web server application integrates a feature-rich eye-gaze analysis module, which is capable of calculating numerous summative statistics. Example features that can be requested by client applications include a user's fixation-rate (i.e. number of fixations per second), mean absolute saccade angles (i.e. angles of the trajectory between two fixations), proportionate amount of time in an *Area of Interest* (AOI) (the location of AOIs is set in a configuration file), transitions between AOIs, etc. (see [15] for a more detailed description of all features).

There are two main advantages of this architecture over previous systems described in section 1. First of all, the web service has been designed to be application-independent, and may therefore be reused for different application domains and purposes. Requests and responses are handled using a simple REST model, which can be easily integrated into any network-connected application. This also means that client applications do not have to run on the same machine as the eye-tracker (see section 3, use case 1). Secondly, in contrast to previous gaze-based adaptive systems presented in section 1, the service integrates a broad range of eye gaze statistics and is therefore capable of driving adaptive systems that are based on feature-rich classification techniques (an example of this is shown in section 3, use case 3).

3 Use-Cases of User-Adaptive Visualization

There are a number of use-case scenarios where the above architecture can be leveraged for real-time adaptive applications. In particular, since visual scanning and processing are fundamental components of working with any form of visualizations (and in fact the only components for non-interactive visualizations), eye tracking may be particularly useful in this field. Below is a quick summary of three scenarios that we have devised in our own research on adaptive, gaze-based visualizations.

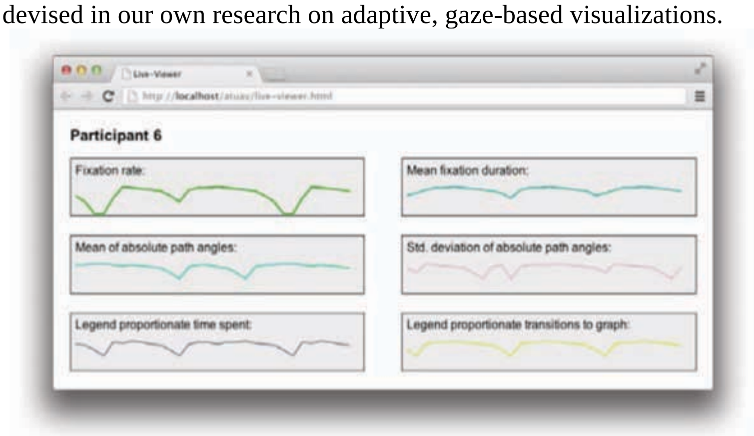

Figure 2. Real-time visualization of eye gaze statistics

Use-Case 1: Real-Time Gaze Visualization

The first use-case is to simply display the different analysis values in a real-time visualization. Figure 2 provides a sample screenshot of a 'spark lines' visualization that we have implemented to showcase the underlying service's real-time processing capabilities. As an end-user 'looks' at a screen displayed on the eye-tracker, this live statistics visualization can be viewed, for example, by a remote experimenter in order to see the real-time changes of the end-user's eye gaze. This setup may be useful for system pilot-testing, real-time gaze analysis, or Wizard-of-Oz-type studies.

Use-Case 2: Gaze-based triggers in Visualization Experiments

A second application scenario is to use simple gaze-based trigger rules to provide basic forms of adaptation. For example, in a recent study we wanted to force users to first read the experimental task question before viewing the corresponding visualization. We used simple fixation-based rules that specified that the 'number of fixations' in the 'text area of interest' had to be above a certain threshold before triggering the visualization display. This use case resembles simple gaze-based adaptation as presented in section 1, however, it is worth noting again that a client application can choose from a much larger array of features to be used for triggering actions.

Use-Case 3: User-Adaptive Visualization

The most sophisticated use-case scenario combines the gaze-based service with rich classification and intervention techniques to predict and adapt to individual user traits and states (e.g., task, individual user characteristics, transient states). In order to implement such a system, the client application (e.g. an adaptive visualization system) requests the full set of eye gaze measures over a specific period of time (e.g. last 10 seconds). Secondly, the measures are run through a previously trained classifier to predict a user's task and/or characteristics (as shown in our previous work [15]). Lastly, the system provides an adaptive intervention that is suitable for the predicted context. Interventions could, for example, consist of reference lines (to facilitate comparisons), bolding (to highlight important parts of the visualization), or de-emphasizing unimportant information (to help users who are overwhelmed by visual clutter).

4 Conclusions and Road Ahead

This paper has presented a quick overview of an architecture that is capable of providing real-time eye gaze statistics to adaptive systems. The service has been fully implemented and has already been used for the development of real-time gaze statistic visualizations (use-case 1), as well as an application that uses simple gaze-based triggers (use-case 2). The next steps include the integration of an end-to-end user-adaptive visualization system, which combines this real-time eye-tracking service with a machine-learning-based classifier and an adaptive visualization intervention framework (use-case 3).

References

1. BBC News. Samsung Galaxy S4 eye-tracking smartphone unveiled. Last accessed April 2 2013. http://www.bbc.co.uk/news/technology-21791023.
2. BBC News. Eye-controlled Gaze TV unveiled by Haier and Tobii. Last accessed April 2 2013. http://www.bbc.co.uk/news/technology-19441860.
3. Bednarik, R., Vrzakova, H., & Hradis, M. What do you want to do next: a novel approach for intent prediction in gaze-based interaction. In Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12).
4. Bondareva, D., Conati, C., Feyzi-Behnagh, R., Harley, J., Azevedo, R., & Bouchet, F. Inferring Learning from Gaze Data during Interaction with an Environment to Support Self-Regulated Learning. AIED 2013 (to appear).
5. Courtemanche, F., Aïmeur, E., Dufresne, A., Najjar, M., & Eboa, F.H.M. Activity recognition using eye-gaze movements and traditional interactions. In Proceedings of Interacting with Computers, (2011), 202-213.
6. D'Mello, S., Olney, A., Williams, C., & Hays, P. Gaze tutor: A gaze-reactive intelligent tutoring system. International Journal Human-Computer Studies, 70(5), 370-389.
7. Goldberg, J.H. & Helfman, J.I. Eye tracking for visualization evaluation: Reading values on linear versus radial graphs. Information visualization 10(3), (2011), 182-195.
8. Iqbal, S.T., & Bailey, B.P. Using eye gaze patterns to identify user tasks. In: The Grace Hopper Celebration of Women in Computing, (2004).
9. Jacob, R. J. K., & Karn, K. S. Eye tracking in Human-Computer Interaction and usability research: Ready to deliver the promises. The mind's eye: Cognitive and applied aspects of eye movement research, 573-605.
10. Kardan, S., & Conati, C. (2012). Exploring Gaze Data for Determining User Learning with an Interactive Simulation. Proceedings of UMAP 2012, the 20th International Conference on User Modeling, Adaptation, and Personalization. Springer LNCS 7379, 126-138.
11. Keith, R. Eye movements and cognitive processes in reading, visual search, and scene perception. Eye Movement Research Mechanisms, Processes, and Applications. North-Holland Press, (1995).
12. Li, R., Pelz, J., Shi, P., Alm, C.O., & Haake. A.R. Learning eye movement patterns for characterization of perceptual expertise. In Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12).
13. Plumlee, M.D., & Ware, C. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. ACM Trans. Comput.-Hum. Interact, 13, 2, (2006), 179-209.
14. Simola, J., Salojarvi, J., & Kojo, I. Using hidden Markov model to uncover processing states from eye movements in information search tasks. Cognitive Systems Research. 9, (2008), 237-251.
15. Steichen, B., Carenini, G., Conati, C. User-Adaptive Information Visualization - Using eye gaze data to infer visualization tasks and user cognitive abilities. Proceedings of the International Conference on Intelligent User Interfaces (IUI 2013).
16. Toker D., Conati C., Steichen B., & Carenini G. Individual User Characteristics and Information Visualization: Connecting the Dots through Eye Tracking. CHI 2013.

Exploring the potential of neurophysiological measures for user-adaptive visualization

Susanne Tak, Anne-Marie Brouwer, Alexander Toet, and Jan van Erp

TNO, Kampweg 5, 3769 DE, Soesterberg, the Netherlands
{susanne.tak,anne-marie.brouwer,lex.toet,jan.vanerp}@tno.nl

Abstract. User-adaptive visualization aims to adapt visualized information to the needs and characteristics of the individual user. Current approaches deploy user personality factors, user behavior and preferences, and visual scanning behavior to achieve this goal. We argue that neurophysiological data provide valuable additional input for user-adaptive visualization systems since they contain a wealth of objective information about user characteristics. The combination of neurophysiological data with other information like eye movement data can significantly improve system reliability by reducing the inherent uncertainty in the interpretation of the user data. Moreover, neurophysiological data can be obtained continuously and unobtrusively without disturbing the interaction of the user with the system.

Keywords: visualization, user adaptation, neurophysiology

1 Introduction

User-adaptive visualization is a novel approach to adapt an information visualization to individual user differences. For example, it can mean adapting the visualization to general (and typically static) user traits such as speed of perceptual processing or user expertise. In this case, adaptation to the user happens only once or a limited number of times. Adaption can also occur more or less continuously, e.g. when the visualization is adapted to the current mental and/or cognitive state of the user (attention, emotion), or to the characteristics of the user-visualization interaction (such as history of user actions).

One of the current research questions in user-adaptive visualization is which individual characteristics can be used as 'input' for adaptation. In this paper we explore the possibilities to adapt visualizations to the continuously changing mental state of the user as can be estimated by (neuro)physiological variables such as heart rate and brain signals. The advantages of using neurophysiological variables are that they provide a continuous, online measure, and do not involve potentially distorted or subjective post-hoc judgment. Also, ongoing miniaturization and development of wireless sensing techniques will allow minimal user interference in the near future.

In this paper, we first provide a short overview of related work in user-adaptive visualization, focusing on the individual characteristics that are typically employed and how these are gathered. Next, we examine the state of the art in neurophysiological measurement. Finally, we provide ideas on how these neurophysiological measures could be used to adapt the visualization, i.e. what characteristics of the visualization can be adapted and how. This includes ideas on how neurophysiological measures can complement other commonly used measures of individual characteristics, such as eye gaze data.

2 Adapting to the user

A variety of factors can be used to adapt a visualization to an individual user. In this section, we briefly review the current state of research in user-adaptive visualization, focusing on *which* individual differences can be used to personalize a given visualization for a particular user. We distinguish differences in personality factors, user actions and preferences, and visual scanning behavior.

Previous work [1, 2] has found that personality factors (locus of control, extraversion and neuroticism) influence performance with different visualization types. Also, several cognitive abilities correlate with various aspects of visualizations. For example, spatial ability correlates with comprehension of (3D) information visualizations [3]. Toker et al. [4] show that perceptual speed (also see [5]), verbal working memory, visual working memory and user expertise have a significant effect on task efficiency, user preference and ease of use of different visualization types. Toker et al. [6] show that perceptual speed and verbal working memory also have a significant effect on eye-gaze behavior when viewing a visualization. The authors suggest that adaptive interventions can be driven by these individual characteristics, for example, by giving more emphasis to certain elements of the visualization (e.g., more emphasis on text for users scoring low on verbal working memory). We note that personality and cognitive abilities are typically assessed with computer-based or paper-and-pencil tasks and questionnaires [3–6], or self-reports [4]. However, self-reports are potentially unreliable, as several memory errors can undermine their accuracy [7]. Also, they are susceptible to social desirability biases. Steichen et al. [8, 9] address this problem by using eye gaze to infer these cognitive abilities. Finally, we note that adapting to these relatively static user traits is less suitable for interactive or continuous adaption of the visualization.

Another commonly used approach, more suited continuous adaption of the visualization, is to adapt to user actions and (implied or explicitly given) preferences (e.g. [10–12]). This information can be used to build a user model, which can be updated over time based on new information, i.e. dynamic. An example of a dynamic model is one that learns from expressed user dislike of a visualization [13]. However, information on these actions and (dis)likes does not reveal *why* these occurred, as they do not give information on the user's (mental) state.

The use of eye gaze measures can provide information on user state. For example, Steichen et al. [8, 9] use eye gaze to predict the user's task. Also, they use eye gaze to infer cognitive abilities such as perceptual speed, visual working memory and verbal working memory. As noted above, these cognitive abilities have a significant effect on the use of different visualization types. Conati et al. [14] investigate the possibilities of using eye gaze to detect when the user needs an adapted visualization. For example, duration of a fixation can indicate complexity, pupil dilation may indicate cognitive load and eye gaze data can be used to determine areas that the user has not looked at. Also, eye gaze can reveal informative patterns such as repeated scanning of the same area in a visualization [15]. In general, for analyzing data generated by an eye-tracker, several metrics are potentially informative: number of fixations (a large number of fixations generally implies a less efficient search [16], or a large number of fixations on a particular area of the screen can be indicative of high interest in that area), fixation duration (long fixation duration often means the user has difficulty extracting information [16]), number of saccades (more saccades indicate a larger amount of visual search), scanpath metrics (such as length, duration and convex hull [17]), or saccade direction changes (a direction change larger than 90 degrees could imply that the user's goals have changed or the user interface is not the way the user expected [18]). Eye gaze measures can be extracted online and continuously without interrupting the user. However, there are several problems associated with eye gaze measures, such as how to define a fixation, how to account for errors in gaze location and how to handle scanning interruptions [15].

Summarizing, current approaches to user adaptive visualizations are based on differences in personality factors, user actions and preferences, and visual scanning behavior. However, since measures of personality factors are inherently unreliable and biased, and since eye movements need additional information for a correct interpretation, additional sources of information are needed to correctly estimate a user's state of mind. In the next section we argue that this information can partly be obtained from neurophysiological measures.

3 Neurophysiological measures

In this section we will show which neurophysiological measures can provide objective information about the mental and cognitive characteristics of users.

A multitude of neurophysiological variables can be measured and analyzed more or less continuously and non-invasively in an office environment. The advantages of using neurophysiological variables are that they are a continuous, online measure, and do not involve potentially distorted or subjective post-hoc judgment. Previous work has already suggested that user-state monitoring is the next potential breakthrough in the use brain-computer interfaces [19]. Examples include electrical brain activity as measured at the scalp (electroencephalography or EEG), oxygen-bound hemoglobin in the brain (near-infrared spectroscopy or NIRS), a combination of EEG and NIRS [20], cardiovascular measures (e.g. heart rate and blood pressure), respiratory measures (e.g. respiration rate) and electrodermal measures (electrical conductance of the skin which varies with sweat excretion).

Non-invasive portable equipment that records brain signals (EEG and NIRS) can only do so from the surface of the brain (the cortex). The signals themselves, electrical activity and the amount of oxygen-bound hemoglobin, are weakened and smeared out by the tissue lying between the sensor and the origin of the signal, and the signals themselves are rather indirect measures of what the brain actually is doing. Still, they can provide us with potentially useful information. Much of the cognitive and sensory functions are localized in the cortex. Roughly speaking, the frontal area is involved in cognitive processing, the area at the back of the brain in vision, auditory processing is at the sides of the brain and touch is at the upper part. This localization knowledge can be combined with general indicators of brain activity or inhibition. For EEG, the power of frequencies around 12Hz (alpha band) indicate idling [21] or inhibition [22] of the recorded part of the brain. For NIRS, a relatively large amount of oxygenated blood generally corresponds to brain activity. For example, experiments on workload indeed indicate that frontal alpha activity [20, 23, 24] and oxygenated hemoglobin increases with task difficulty. For a more detailed overview of the usefulness of brain-based indices for effort, vigilance, workload and engagement see [25].

Another common type of measurement extracted from EEG besides power in frequency bands are event-related potentials (ERPs). These are the peaks and valleys as observed in averaged EEG traces that are locked to an external event, such as the onset of the presentation of an image. The P300 is a peak occurring 300 ms (or somewhat later) after the presentation of a stimulus that attracts special attention. This special attention can be because it either pops out with respect to previous and subsequent stimuli (e.g. [26]), because of inherent meaning (e.g. one's own name in between others names [27]) or because the stimulus is one that the individual is asked to consciously attended to [28]. Error related potentials [29–32] are ERPs associated with the onset of making a mistake that is realized to be a mistake or with the onset of an unexpected outcome of an action.

The correlates of cognitive and attentional processes in EEG as described above cannot be observed in raw data with the naked eye. While classical EEG studies make use of averaging over many trials to visualize and determine effects, applied neuroscientific studies (for which it is crucial to extract information from one individual over a relatively short time period) commonly use classification techniques (e.g. linear discriminant analyses, support vector machines). Van Gerven et al. [33] give an overview of using classification techniques in brain-computer interfaces. In short, EEG is recorded while an individual experiences the states of interest at known time intervals (e.g. low and high workload, or image of interest present or not). With these labeled EEG data a classification model is trained. The trained model can then classify new unseen data into the trained categories. Usually, models are personalized, though work is ongoing create models that generalize over participants [34].

Whereas cognitive and perceptual processing occur mostly at the cortex, emotional processing occurs mostly in the center areas of the brain. Therefore, to track emotions, other (physiological) measures are needed. For example, an increase in skin conductance (i.e., sweating) and a decrease of high-frequency heart rate variability are associated with stress or arousal [35–38]. Kreibig [39] gives an extensive review of physiological correlates of different basic types of emotion such as anger and disgust. However, it is important to note that physiology does not correspond one-to-one to different emotions, but stimuli and context matter. Alternatively, emotional valence and arousal could be dervied from facial expressions (facial expression analysis software is available).

4 Neurophysiological data - novel opportunities for visual adaptation

In the previous sections, we have shown that neurophysiological data can be employed to learn more about user characteristics. Naturally, a follow-up question is how this can be linked to adaptation of the visualization. In other words, when we know more about the user, what aspects of the visualization can be adapted such that the user is supported better. In general, aspects of the visualization that can be adapted to the user include:

- which data is shown,
- how the data is shown (visualization type, e.g. [5]),
- how the visualization is parameterized (how data attributes map to visual attributes, e.g., size and color)
- how the data is laid out spatially (e.g., [10])
- what portion of the data is shown (e.g., zoom and filter settings),
- which details of the data are shown,
- which elements of the visualization are given more visual emphasis (suggested by [6, 14]), and
- help provided to the user.

As noted, neurophysiological measures can be used to estimate workload and stress/arousal. This information about the user could be used to adapt the visualization. For example, if stress levels and workload are high, the level of detail of the visualization, or the amount of data shown could be reduced. Also, more help could be provided when workload is high.

While we are not aware of studies on the use of (neuro)physiological measures in adaptive visualization, several studies have examined the use of information as revealed by eye movements. Eye movements can also reflect mental state and, as mentioned previously, share the advantages of (neuro)physiological measures. Neurophysiological measures can be used to complement eye tracking data. In general, we note that a *combination* of information sources (neurophysiological, eye tracking or also facial expressions or body posture [40]) will likely prove most effective. For example, while a fixation on a certain area of the visualization can mean the user is interested in that area, it can also mean that the user does not understand something in that area or that (s)he is simply staring. With neurophysiological measures, these different states can be differentiated better. Conversely, eye tracking can reveal information that neurophysiological measures might not so easily capture. For example, point of gaze is a measure that is most easily captured using an eye tracker. Another example of how combining neurophysiological measures with eye-tracking could lead to more insight into current user behavior is differentiation between searching and exploring behavior. Many gaze shifts in combination with hardly any long fixations can imply that a person is searching for an item [18]. If this is found to be paired with a high workload as indicated by EEG signals, it may imply that the user cannot find what he/she is looking for. If paired with a low workload the user may just be (casually) exploring the visualization.

As for ERP measures, we note that in the vast majority of ERP studies, experimental participants do not move their eyes at the time that the stimuli of interest are presented, so that perception onset is fixed and EEG artifacts caused by eye movements are limited. This conflicts with every day (visualization) behavior where individuals look around freely. However, when EEG is examined following eye fixation onset rather than image onset, similar ERPs can be distinguished [41]. This offers the possibility to deduct from EEG and eye movements where the user is looking when he/she realized that there was mistake or an unexpected outcome. This information can then be used to detect anomaly or error, to present a pop-up help query, or to automatically highlight (or tag) the area where this happened so the user can easily return to it later.

Because the neurophysiological measures discussed in this paper are continuous the effect of such adaptations can be measured on-the-fly, and different adaptations can be experimented with. Combined with an adaptive (or learning) user model this can lead to highly personalized visualizations.

5 Conclusion

User-adaptive visualization aims to adapt visualized information to the needs and characteristics of the individual user. Current approaches typically deploy user personality factors, user behavior and preferences, and visual scanning behavior to achieve this goal.

We have shown that neurophysiological measures can provide information about the mental state of an individual. Challenges include noise and interpretation difficulties caused by body movements and speech. These can cause 'actual' disturbances to the signal as well as confounds (e.g. in a high workload situation individuals move more than when workload is low, causing higher heart rate due to movement rather than mental workload). However, since visualizations are often used in a relatively quiet environment, with relatively stationary users, there is minimal risk of noise and confounds in the neurophysiological data. We therefore believe that user-adaptive visualizations are a relatively good case for exploring the potential neurophysiological measures.

Finally, we note that robust, user friendly and high quality measurement equipment needs to be developed, especially to reliably and easily record EEG. Impressive progress is being made [42, 43] and first user friendly EEG measurement equipment is on the market. Also, we note that some methods that are more time- and resource-demanding, such as fMRI, whilst unsuitable for adapting visualizations online could be used for the evaluation of visualization.

References

1. Green, T., Fisher, B.: Towards the personal equation of interaction: The impact of personality factors on visual analytics interface interaction. In: Proc. of VAST '10, pp. 203–210. IEEE (2010)
2. Ziemkiewicz, C., Crouser, R., Yauilla, A., Su, S., Ribarsky, W., Chang, R.: How locus of control influences compatibility with visualization style. In: Proc. of VAST '11, pp. 81–90. IEEE Computer Society (2011)
3. Velez, M., Silver, D., Tremaine, M.: Understanding visualization through spatial ability differences. In: Proc. of IEEE Visualization (VIS) 2005, pp. 511–518. IEEE Computer Society (2005)
4. Toker, D., Conati, C., Carenini, G., Haraty, M.: Towards adaptive information visualization: on the influence of user characteristics. In: Proc. of UMAP '12, pp. 274–285. Springer-Verlag, Berlin, Heidelberg (2012)
5. Conati, C., Maclaren, H.: Exploring the role of individual differences in information visualization. In: Proc. of AVI '08, pp. 199–206. ACM, New York, NY, USA (2008)
6. Toker, D., Conati, C., Steichen, B., Carenini, G.: Individual user characteristics and information visualization: Connecting the dots through eye tracking. In: Proc. of CHI'13 (2013)
7. Tourangeau, R.: Remembering what happened: Memory errors and survey reports. In: Stone, A.A., Turkkan, J.S., Bachrach, C.A., Jobe, J.B., Kurtman, H.S., Cain, V.S. (eds.) The Science of Self-Report: Implications for Research and Practice, pp. 29–47. Lawrence Erlbaum (2001)
8. Steichen, B., Carenini, G., Conati, C.: Adaptive information visualization - predicting user characteristics and task context from eye gaze. In: Poster and Workshop Proc. of UMAP '12. Springer-Verlag, Berlin, Heidelberg (2012)
9. Steichen, B., Carenini, G., Conati, C.: User-adaptive information visualization - using eye gaze data to infer visualization tasks and user cognitive abilities. In: Proc. of IUI '13, pp. 317–328. ACM, New York, NY, USA (2013)
10. Ahn, J.W., Brusilovsky, P.: Adaptive visualization of search results: bringing user models to visual analytics. Information Visualization 8(3), 167–179 (2009)
11. Ahn, J.W., Brusilovsky, P.: Can concept-based user modeling improve adaptive visualization? In: Proc. of UMAP '10, pp. 4–15. Springer-Verlag, Berlin, Heidelberg (2010)
12. Ahn, J.W., Brusilovsky, P.: What you see is what you search: adaptive visual search framework for the web. In: Proc. of WWW '10, pp. 1049–1050. ACM, New York, NY, USA (2010)
13. Golemati, M., Halatsis, C., Vassilakis, C., Katifori, A., Peloponnese, U.o.: A context-based adaptive visualization environment. In: Proc. of IV '06, pp. 62–67. IEEE Computer Society, Washington, DC, USA (2006)
14. Conati, C., Carenini, G., Harati, M., Tocker, D., Fitzgerald, N., Flagg, A.: User-Adaptive Visualizations: Can Gaze Data Tell Us When a User Needs Them? In: Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence, pp. 8–11 (2011)
15. Goldberg, J.H., Helfman, J.I.: Comparing information graphics: a critical look at eye tracking. In: Proc. of BELIV '10, pp. 71–78. ACM, New York, NY, USA (2010)
16. Jacob, R.J.K., Karn, K.S.: Eye tracking in human-computer interaction and usability research: Ready to deliver the promises. In: The Mind's Eye: Cognitive and Applied Aspects of Eye Movement Research, pp. 573–603. Elsevier Science (2003)

17. Goldberg, J., Kotval, X.: Computer interface evaluation using eye movements: methods and constructs. International Journal of Industrial Ergonomics 24(6), 631–645 (1999)
18. Poole, A., Ball, L.: Eye tracking in human-computer interaction and usability research. In: Encyclopedia of human computer interaction. Idea group
19. van Erp, J., Lotte, F., Tangermann, M.: Brain-computer interfaces: Beyond medical applications. Computer 45(4), 26–34 (2012)
20. Coffey, E.B., Brouwer, A.M., van Erp, J.B.: Measuring workload using a combination of electroencephalography and near infrared spectroscopy. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 56(1), 1822–1826 (2012)
21. Pfurtscheller, G., Stancák, A.J., Neuper, C.: Event-related synchronization (ERS) in the alpha band–an electrophysiological correlate of cortical idling: a review. International Journal of Psychophysiology 24(1-2), 39–46 (1996)
22. Klimesch, W., Sauseng, P., Hanslmayr, S.: EEG alpha oscillations: The inhibition-timing hypothesis. Brain Research Reviews 53(1), 63–88 (2007)
23. Brouwer, A.M., Hogervorst, M., van Erp, J., Heffelaar, T., Zimmerman, P., Oostenveld, R.: Estimating workload using EEG spectral power and ERPs in the n-back task. Journal of Neural Engineering 9(4), 045008 (2012)
24. Sterman, M.B., Mann, C.A., Kaiser, D.A., Suyenobu, B.Y.: Multiband topographic EEG analysis of a simulated visuomotor aviation task. International Journal of Psychophysiology 16(1), 49–56 (1994)
25. van Erp, J., Veltman, H.J., Grootjen, M.: Brain-based indices for user system symbiosis. In: Brain-Computer Interfaces, pp. 201–219. Springer, London, UK (2010)
26. Polich, J., Brock, T., Geisler, M.W.: P300 from auditory and somatosensory stimuli: probability and inter-stimulus interval. International Journal of Psychophysiology 11(2), 219–223 (1991)
27. Rosenfeld, J.P., Biroschak, J.R., Furedy, J.J.: P300-based detection of concealed autobiographical versus incidentally acquired information in target and non-target paradigms. International Journal of Psychophysiology 60(3), 251–259 (2006)
28. Brouwer, A.M., van Erp, J.: A tactile P300 brain-computer interface. Frontiers in Neuroscience 4(19), 1–11 (2010)
29. Falkenstein, M., Hoormann, J., Christ, S., Hohnsbein, J.: ERP components on reaction errors and their functional significance: a tutorial. Biological Psychology 51(23), 87–107 (2000)
30. Gehring, W.J., Goss, B., Coles, M.G.H., Meyer, D.E., Donchin, E.: A neural system for error detection and compensation. Psychological Science 4(6), 385–390 (1993)
31. Ferrez, P.W., del R Millan, J.: Error-related EEG potentials generated during simulated brain-computer interaction. IEEE transactions on biomedical engineering 55(3), 923–929 (2008)
32. Lehne, M., Ihme, K., Brouwer, A.M., van Erp, J., Zander, T.: Error-related EEG patterns during tactile human-machine interaction. In: Proc. of ACII '09, pp. 1–9 (2009)
33. van Gerven, M., Farquhar, J., Schaefer, R., Vlek, R., Geuze, J., Nijholt, A., Ramsey, N., Haselager, P., Vuurpijl, L., Gielen, S., Desain, P.: The braincomputer interface cycle. Journal of Neural Engineering 6(4), 041001 (2009)
34. Reuderink, B.: Robust brain-computer interfaces. Dissertation, University of Twente (2011)
35. Lang, P., Greenwald, M., Bradley, M., Hamm, A.: Looking at pictures: Affective, facial, visceral, and behavioral reactions. Psychophysiology 30(3), 261–273 (1993)

36. Carter, J.R., Durocher, J.J., Kern, R.P.: Neural and cardiovascular responses to emotional stress in humans. American Journal of Physiology - Regulatory, Integrative and Comparative Physiology 295(6), R1898–R1903
37. Aasman, J., Mulder, G., Mulder, L.J.M.: Operator effort and the measurement of heart-rate variability. Human Factors: The Journal of the Human Factors and Ergonomics Society 29(2), 161–170 (1987)
38. Delaney, J., Brodie, D.: Effects of short-term psychological stress on the time and frequency domains of heart-rate variability. Perceptual and Motor Skills 91, 515–524 (2000)
39. Kreibig, S.D.: Autonomic nervous system activity in emotion: a review. Biological Psychology 84(3), 394–421 (2010)
40. D'Mello, S., Graesser, A.: Mining bodily patterns of affective experience during learning. In: Proceedings of the 3rd International Conference on Educational Data Mining, pp. 31–40 (2010)
41. Dimigen, O., Sommer, W., Hohlfeld, A., Jacobs, A., Kliegl, R.: Coregistration of eye movements and EEG in natural reading: analyses and review. Journal of Experimental Psychology: General 140(4), 552–572 (2011)
42. Patel, S., Park, H., Bonato, P., Chan, L., Rodgers, M.: A review of wearable sensors and systems with application in rehabilitation. Journal of NeuroEngineering and Rehabilitation 9, 1–17 (2012)
43. Zander, T.O., Lehne, M., Ihme, K., Jatzev, S., Correia, J., Kothe, C., Picht, B., Nijboer, F.: A dry EEG-system for scientific research and braincomputer interfaces. Frontiers in neuroscience 5(53), 1–10 (2011)

Capturing and Reusing Empirical Visualization Knowledge

Martin Voigt[1] and Martin Franke[2] and Klaus Meißner[1]

[1] Senior Chair of Multimedia Technology
[2] Software Engineering of Ubiquitous Systems,
Technische Universität Dresden
Dresden, Germany
martin.voigt,martin.franke,klaus.meissner@tu-dresden.de

Abstract. The context-aware discovery and ranking of visualization components is a crucial part of an adaptive, automated information visualization system. Since existing approaches allow for using expert knowledge formalized a priori, insights gained during the visualization processes by the users, e. g., suitable data-visualization combinations, are mostly neglected. In this paper, we propose a concept to capture and formalize these insights. Furthermore, we enhance a context-aware ranking approach using this knowledge by applying the well-known collaborative filtering. Thus, we are able to employ ratings also if data-visualization combinations are new for the current user.

Keywords: visualization, recommendation, adaptation, knowledge, collaborative filtering

1 Introduction

Due to the inexorable growth of data in all areas of life, humans face more and more the problem to understand the datasets they are confronted with. Thus, information visualization (InfoVis) tools are required to assist end users who are not familiar in creating effective graphical representations. Towards such an application tailored to visualization novices, we already proposed a visualization ontology (VISO) [8] to capture expert knowledge of the interdisciplinary domain, e. g., about data, graphic vocabulary, or human activity. This formal model is the foundation of a semantics-based InfoVis process [12] which guides the end user from identifying interesting parts of a dataset, over the context-aware selection of visualization components to the final configuration and perception of the visual representation. Its core is a discovery algorithm [11] to match the selected data but also to rank the suitable components according to factual visualization, domain, and context knowledge. Since the latter also includes users explicit preferences stored in a user model, we understand that it is hard for novices to explain which visualization techniques they like. Furthermore, we figured out that the knowledge created within the visualization process, i. e., which data-visualization combinations are appropriate, or inappropriate, in the current context, is lost so far.

Hence, our goal is not to neglect this kind of knowledge anymore and to use it for a better adaptation of the ranking of suitable visualization components. Therefore, we had to answer mainly two questions. Firstly, how could we extract the knowledge overburdening the end user? Secondly, how could we store, formalize, consolidate, and reuse the assembled information within our semantics-based visualization process? Thus, we present our solutions to tackle the mentioned problems by adapting concepts from the area of knowledge-assisted visualization as well as the domain of recommender systems. The concrete contributions are twofold: First, we propose a concept to *externalize* users visualization-specific insights by using implicit and explicit ratings, which are stored in a semantic knowledge-base. Second, we adapt our ranking algorithm to employ this empirical knowledge using collaborative filtering which allows for the indirect *collaboration* of end users within the visualization system.

In the following, we discuss the related work (Sect. 2), present our concept (Sect. 3), and finally conclude with some findings (Sect. 4).

2 Related Work

The work related to ours comes from the broad field of automated visualization systems. Here, the first sophisticated concepts and tools arose over 20 years ago. Unfortunately, *adaptation* of these systems, for instance to user needs and the device, has attracted less attention. With this regard, Golemati et al. [3] proposed a concept for a visualization environment for document collections taking into account user's, system's, and document's context. Similar to our work, they allow for an explicit rating of visualizations in the current context. This empirical knowledge is used for adaption in upcoming visualization selections of this user.

In contrast, Gotz et al. [4] and Nazemi et al. [5] present adaptive InfoVis systems building primarily on tracking users interactions and thus on their implicit ratings. Based on the interaction patterns they infer users task, match if the current presentation fits well or not, and if necessary adapt the graphic representation. These interaction-driven approaches leverage implicit state information, but they consider neither task information or preferences that are explicitly expressed by the user. Furthermore, in our concept the user's knowledge is reused in a collaborative manner as a global shared knowledge.

Since the adaption builds on many interrelated parameters, the knowledge-management within automated InfoVis tools is crucial. Thus, we rely on concepts from the domain of *knowledge-assisted visualization* as well. Its goal is to overcome the burden of learning complex visualization techniques by formalizing and sharing domain and visualization knowledge [1]. Wang et al. [13] propose a knowledge conversion process in visual analytics system. It comprises the *internalization* of knowledge of the user based on the graphical representation, the *externalization* of users insights based on interactions, or the direct *collaboration* of users during the visualization to foster the internalization. Bearing this theoretical framework in mind, we present an adaptive visualization system building on expert and empirical knowledge modeled using standardized ontology languages to foster their reuse and sharing.

As mentioned above, collaboration is a helpful concept to support users within the visualization process. Hence, intelligent InfoVis systems building on formalized expert and empirical knowledge may help. While the first could be captured and modeled a priori, the latter needs to be tracked, formalized, and reused on runtime. With this in mind, *collaborative filtering* is a well-known approach [9] which investigates similarity of ratings for items, in our case visualizations, given by users. Hence, no content analysis or tagging by experts is required. To the best of our knowledge, collaborative filtering has not been applied within automated visualization systems so far. In our opinion, the main reason is that they may identify suitable graphical representations based on users ratings but do not allow for an automated mapping of data to visual objects. However, if this matching is already provided by the visualization system, collaborative filtering introduces a new facet for rating and recommending visualization techniques: users conclusion if a mapping is suitable in the current context. Thus, it allows for tracking and storing user knowledge, enables the sharing of insights and the adaption of the visualization process due to this empirical knowledge.

3 Concept

Building on a context-aware recommendation algorithm from [11], which employs VISO [8] as ontological knowledge-base, in the following, we present concepts to externalize and reuse empirical knowledge evolving in the visualization process.

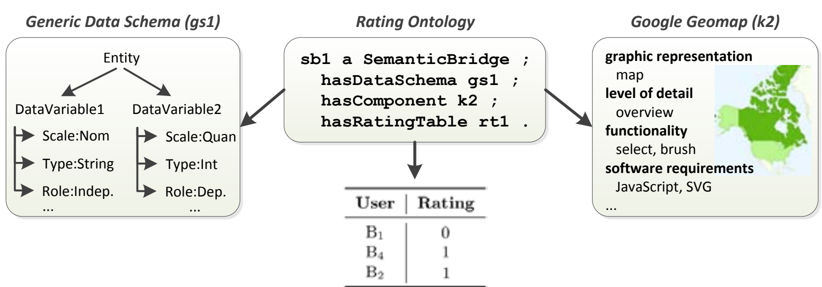

Fig. 1. Overview of the function of the rating ontology.

3.1 Externalization of Empirical Knowledge

The *externalization* process describes the storage of user's internal knowledge within the system. This process can be distinguished in acquiring *implicit* and *explicit* knowledge. The tracking and interpretation of interactions as knowledge is called *implicit rating* and is used for gathering knowledge, before the user marks down the end of the adaption loop by an *explicit rating*. In our concept, such a rating is always saved for a combination of a generic data schema and concrete visualization component using the *rating ontology* which semantically links both resources, see Fig. 1. The data schema is generated during the matching [11] and is an abstraction of the data a user has chosen for visualization including meta information like the scale of measurement or the data type for each data variable. The description of a visualization component holds all semantical information, e. g., its kind of graphic representation or interaction abilities. Finally, the ontology comprises the ratings of the users for a data-component combination. Based on this semantic interlinking, we are able to query for instance "only good rated components which visualize data on a map and allow for linking and brushing".

As mentioned above, we distinguish between explicit and implicit rating. To harvest implicit ratings, we rely on three actions from [6].

- **Repeated Use:** A visualization component is *used* more than three times by the same user. The usage of the component is recognized by doing a specific count of interactions or within a defined time interval. Since the repeated use is a sign that the user favors a component, it is temporary assigned with $r_{user} = 0.75$
- **Glimpse:** If the chosen visualization component is discarded without reaching a defined time interval or a count of interactions for recognizing the *Repeated Use*, it is downgraded and temporary assigned with $r_{user} = 0.25$
- **Related Rate:** The visualization component was explicit rated by the user, but in a different data combination. Since the user knows its characteristics, it is possible that he likes it for other data selections, which are distinct from the generic data schema, as well. In case of a good rating, the temporary is $r_{user} = 0.75$, otherwise $r_{user} = 0.25$.

Beside this implicit knowledge capturing, we gather *explicit* ratings. Thus, the user can explicitly decide whether the visualization is applicable for its purpose or not. Since we like to stimulate the user to rate, we employ a simple scale of *applicable* (1) or *not applicable* (0). We fulfill thereby the requirement to give the user an adequate possibility to rate, without an excessive demand.

3.2 Indirect Collaboration

The *collaboration* process describes a direct cooperation of two or more users [13], such as a chat or co-browsing. We broaden this scope by including indirect sharing of knowledge between users as collaboration. For our purpose, we need to reuse the collected knowledge within our ranking algorithm presented below. It comprises originally three kinds of ratings [11] according the factual visualization knowledge (r_{v_i}), the quality of the domain assignments (r_{d_j}), and the context knowledge (r_{c_k}). Now, it is extended by the user rating (r_{user}) to employ the user-generated visualization knowledge by using collaborative filtering. Thus, if the user has not rated the visualization component in the specific combination with a dataset, the algorithm tries to foresee a possible rating.

The combination of these factors is done using an arithmetic mean. The overall rating R which is calculated in terms of

$$R = \frac{1}{3|4}\left(\frac{1}{x}\sum_{i=1}^{x} r_{v_i} + \frac{1}{y}\sum_{j=1}^{y} r_{d_j} + \frac{1}{z}\sum_{k=1}^{z} r_{c_k}\ (+\ r_{user})\right)$$

has, therefore, a range between 0 and 1. We weight all three, respectively four, rating types equivalently for two reasons. First, the assignment of a (quantitative) rating is often subjective. Second, a complex user study is needed to evaluate the impact of each knowledge base in users visualization selection process what will be future work. The meaning factor $1/n$ is assigned with 1/3, if neither a user rating can be found, nor can be calculated. In all other cases it is assigned with 1/4 to keep the equivalently rating of all factors.

	K$_1$	**K$_2$**	**K$_3$**	**K$_4$**	**K$_5$**
B$_1$	1	0	-	0	0
B$_2$	0	1	0	1	1
B$_3$	0	r_{rs}	1	1	1
B$_4$	1	1	1	1	-

Fig. 2. Example set for CFRS prediction

To identify an appropriate collaborative filtering algorithm, we analyze them by taking the accuracy, efficiency, stability, justification, and serendipity into account [2]. Hence, we decided to employ the item-based approach for our use case. It calculates the similarity between the ratings of visualization components, given by different users. With this information, the algorithm can predict the rating for the current visualization. An example is given in Fig. 2. It has to calculate the prediction r_{rs} for user B_3 and visualization component K_2. In this setting, the prediction is assessed based on the rating distances between K_2 and all other visualization components (K_1 - K_5). The algorithm chooses K_4 and K_5 as *nearest neighbours* and forecasts a rating of $r_{rs} = 1$. The formal calculation, which could be used without adaption, is considered in [10].

3.3 Architecture

To realize the concepts discussed above, we specified an architecture shown in Fig. 3 . It comprises three layers: ontological knowledge bases, loosely-coupled web services, and a component-based user interfaces. The first layer consists of three interconnected knowledge-bases. The *VISO* ①, contains all visualization specific knowledge, e. g., to describe the interaction and visualization capabilities of components. The API and meta information apart from the InfoVis domain of a component is described by the *Mashup Component Description Ontology* (MCDO) ② [7] which links to VISO concepts. As mentioned in Sect. 3.1, we designed an ontological knowledge base to store user's *ratings* ③ for the mapping of selected data to the chosen component. Therefore, it refers to concepts of the

VISO and MCDO. We build on different web service which heavily make use of the mentioned knowledge bases. The *Component Repository* 4 allows for the semantic-driven management of visualization components based on the MCDO. Further, the recommendation algorithm for appropriate components [11] is integrated within this service. Hence, it gathers amongst others the assessments stored or calculated within the *Rating Repository* 5 according to our proposed concepts. The visualization workflow of a user is accomplished by a composite web application called *VizBoard* 6 [12]. For example, it allows for searching for graphical representations, to integrate suitable ones and, finally, represented the data. At this last stage, it enables to explicitly and implicitly acquire users knowledge, like explained in Sect. 3.1, and save it using the Rating Repository.

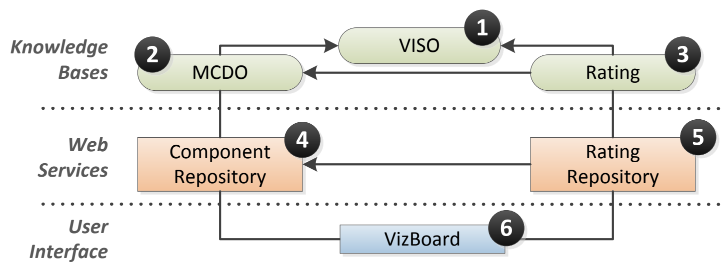

Fig. 3. Overview of the software architecture.

4 Conclusion and Further Work

The reuse of empirical knowledge evolving in a visualization process, especially proper mappings of data to visualization components within a context, was neglected and got lost so far. Hence, we proposed a concept for its externalization to capture, formalize and integrate the insights in an existing expert knowledge base. Furthermore, it becomes an essential part within our context-based ranking approach. Due to the application of collaborative filtering, we are able to employ explicit and implicit ratings also if data-visualization combinations are new for the current user. To the best of our knowledge, our approach is the first which employs formalized, inferred expert knowledge but also empirical, evolving knowledge from users to identify the most suitable visualization components.

Currently, we are also planning to conduct an exhaustive user study to identify and model the interdependencies between the knowledge bases employed within the ranking. Furthermore, we are working on a concept to use the a priori and empirical knowledge to assist the user in interpreting the visualized data what will underpin the usefulness of knowledge-assisted visualization.

Acknowledgements

This work is partly funded by the German Federal Ministry of Education and Research (01IA09001C); and by the European Social Fund and the Federal State of Saxony in Germany within project VICCI (ESF-100098171)

References

1. Chen, M., Hagen, H.: Guest editors' introduction: Knowledge-assisted visualization. Computer Graphics and Applications, IEEE 30(1), 15 –16 (jan-feb 2010)
2. Desrosiers, C., Karypis, G.: A comprehensive survey of neighborhood-based recommendation methods. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) Recommender Systems Handbook, pp. 107–144. Springer US (2011)
3. Golemati, M., Halatsis, C., Vassilakis, C., Katifori, A., Lepouras, G.: A context-based adaptive visualization environment. In: Information Visualization, 2006. IV 2006. Tenth International Conference on. pp. 62–67 (July 2006)
4. Gotz, D., Wen, Z.: Behavior-driven visualization recommendation. In: IUI '09: Proceedings of the 13th international conference on Intelligent user interfaces. pp. 315–324. ACM, New York, NY, USA (2009)
5. Nazemi, K., Stab, C., Kuijper, A.: A reference model for adaptive visualization systems. In: Proc. of the 14th Int. Conf. on Human-computer interaction: design and development approaches - Volume Part I. pp. 480–489. HCII'11, Springer (2011)
6. Nichols, D.M.: Implicit rating and filtering. In: In Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering. pp. 31–36 (1998)
7. Pietschmann, S., Radeck, C., Meißner, K.: Semantics-Based Discovery, Selection and Mediation for Presentation-Oriented Mashups. In: Proc. of the 5th Int. Workshop on Web APIs and Service Mashups. ACM (2011)
8. Polowinski, J., Voigt, M.: VISO: A shared, formal knowledge base as a foundation for semi-automatic infovis systems. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 2013 (CHI'13) (2013)
9. Ricci, F., Rokach, L., Shapira, B.: Introduction to recommender systems handbook. In: Recommender Systems Handbook, pp. 1–35. Springer (2011)
10. Sarwar, B., Karypis, G., Konstan, J., Reidl, J.: Item-based collaborative filtering recommendation algorithms. In: Proceedings of the 10th international conference on World Wide Web. pp. 285–295. WWW '01, ACM, New York, NY, USA (2001)
11. Voigt, M., Pietschmann, S., Grammel, L., Meißner, K.: Context-aware recommendation of visualization components. In: Proc. of the 4th Int. Conf. on Information, Process, and Knowledge Management (eKNOW 2012) (2012)
12. Voigt, M., Pietschmann, S., Meißner, K.: A semantics-based, end-user-centered information visualization process for semantic web data. In: Hussein, T., Paulheim, H., Lukosch, S., Ziegler, J., Calvary, G. (eds.) Semantic Models for Adaptive Interactive Systems, pp. 81–106. Springer Berlin / Heidelberg (2013)
13. Wang, X., Jeong, D.H., Dou, W., Lee, S.W., Ribarsky, W., Chang, R.: Defining and applying knowledge conversion processes to a visual analytics system. Computers & Graphics 33(5), 616 – 623 (2009)

Adaptive Security Event Visualization for Continuous Monitoring

Anatoly Yelizarov, Dennis Gamayunov

Lomonosov Moscow State University
Department of Computational Mathematics and Cybernetics
Information Systems Security Lab
{tolya, gamajun}@lvk.cs.msu.su

Abstract. The field of information security routinely produces the need for a security information and event management system operator who would be capable of durable and extensive (e.g., workday-long) monitoring of the system in his control with well-timed decision making in emergencies. The obvious concern is that such continuous exertion is bound to lead to the operator's increased fatigue, reduced attention span, and flawed decision making. This paper proposes methods of the visualization system's adaptation to these changes for improving the operator's efficiency in terms of speed and accuracy.

Keywords: Adaptive User Interfaces, Event Visualization, Information Security

1 Introduction

The problem of increasing the effectiveness of the "human-machine" system is tackled from two different aspects: by better preparing the human for work with a given machine (through training and building up experience) and by better *adapting the machine* to the specific objectives of the human. In the latter case, special attention is paid to the parts of the system where human directly interacts with machine or, in other words, the parts responsible for information reception and transmission. These parts go beyond the visualization system, screen, mouse, and keyboard, to include also the human himself as the one who perceives and interprets the information, and then interacts with the machine to convey his decisions. Differently put, there must be close attention to properties of the human as a "data channel".

One relevant example of such a "human-machine" system is exactly the "human operator and information security visualization system". Considering the ever increasing number of cyber-attacks, there is a growing need, on both corporate and government security levels, for the human operator of the attack detection system. Frequently, such an operator is tasked with *continuous* (through the entire working day) monitoring of the networks and services under his control. Besides, such an operator must make timely decisions in emergencies by performing typical actions that protect the monitored resources (e.g., blocking external hosts and subnets, running and shutting down services on the protected hosts and simple tweaking of the system's operation logic). Hereinafter such human operator will be referred to as "*the operator*".

The obvious concern is that workday-long exertion would result in increased operator fatigue, reduced attention span, and flawed decision making. Although we have not yet received sufficient supporting test data, further discussion proceeds from the assumption that *psychophysical condition of the operator varies over time* as the fatigue builds up, blunting his responses and worsening the perception quality of the same cognitive load level; though, these characteristics may change in the opposite direction as well (e.g., after a break or in the state of heightened alert).

As the operator is the only one responsible for making well-timed and correct decisions, it can be presumed that careful *consideration of the current psychophysical condition and individual cognitive features* of different operators would benefit the speed and accuracy of their decisions. However, the visualization systems now available (e.g., the monitoring modules in the security information and event management solutions from major vendors) provide versatile visualization tools designed for the "average" user. Whereas such systems do usually provide means for interface customization, they typically ignore current user individual abilities.

This paper discusses and proposes the methods to adapt the visualization system's functionality and presentation capabilities to the psychophysical condition of the operator; in particular, it focuses on ways in which the visualization system could improve the decision making process by taking into account changes of human-computer interaction (HCI) indicators according to operator's fatigue, reduced attention, etc. A further discussion narrows this focus to the informational security visualization systems and theirs operators, leaving the door open to generalization to any event-based visualization system and its continuously working user.

2 Background and Related Work

2.1 Information Security Visualization Systems

The hardware and software systems which support information security are commonly referred to as *security information and event management* (SIEM) solutions. A typical SIEM architecture consists of a security events collection module (that receives information from various sources, e.g., intrusion detection systems (IDS), intrusion prevention systems, firewalls, operating systems, databases, various applications, etc.), correlation and analysis module, database of security events, and monitoring module. The latter is generally divided into reporting module (to generate reports with information required for security administrators), and visualization module (responsible for displaying security events and status of network devices in real-time, visualization of past attacks, providing an interface for working with security event database, and providing opportunity for efficient management and resolution of incidents).

Of the SIEM architecture just described, this paper focuses solely on the visualization module (hereinafter such module is referred to as "*the visualization system*"). Visualization for information security is a relatively young domain that studies, designs new, and adapts common visualization techniques for information security data. These data are characterized by a large volume, lots of parameters and the need for real-time display. Comprehensive coverage of that topic, with emphasis on the visualization systems, is presented in surveys [16] and [10].

2.2 Cognition in Human-Computer Interaction

Cognitive psychology and cognitive science are the major research areas addressing the mental processes such as attention, perception, and problem solving, which influence the operator's decision making efficiency. *Human-computer interaction* (HCI) studies and designs interaction between people and machine. The application of cognitive sciences to problems of HCI is called cognitive ergonomics [12]. HCI consists of factors associated both with the computer (such as equipment, performance, and software) and the human (such as training, experience, and individual differences) [15]. The term “individual differences” denotes the user's personal cognitive abilities rather than demographic description, though it may and does affect his cognitive abilities. Recent researches have confirmed that some of these individual *cognitive abilities*, such as working memory capacity [11],[20],[18], perceptual speed [20],[19], spatial ability [20], and locus of control [24], correlate with the user's performance as to speed and accuracy. The nature of human *visual perception* specific to visualization is covered in depth in [22]. Limitations of human cognition, and their impact on information security visualization are discussed in [3] by reviewing vulnerabilities of operator's visual perception and how they can be exploited in cyber-attacks.

2.3 User-Adaptive Visualization

Successful design of HCI starts from understanding the user, his tasks, and context of his actions [9]. An *adaptive interface* is generally defined as an interface that automatically varies its layout and elements adapting to the user's needs, task and context. Early adaptive interfaces could adapt only to their tasks or the data to be displayed, ignoring any information about the current user. With progress of research and technology, new user tracking techniques were developed — click-streams and eye gaze processing, physical and biomedical sensors, and user models, to name a few. Recent studies provide encouraging evidences that user's cognitive ability could be reliably detected in real time using eye gaze information [19] or [17].

One example of the adaptive interface is the *attentive user interface* that is a user interface aimed to support the user’s attention capacities [21]. Such an interface arises where the visualization system is blended with technologies allowing to track and infer priorities of user attention [14]. Current attentive interfaces are able to maximize the expected utility from the information user receive, thus increasing the efficacy of the HCI [8]. But the fundamental challenge

[7] of any new such attentive technique is reasoning about its interruption costs (e.g., periphery animation cause distraction from attention- and motion-intensive tasks [4]).

Another approach includes *adaptation to the user's experience*. For example, [19] presumes that in some cases non-experts may benefit from adaptive intervention (based on different effect on the label/legend access for different visualizations); or [1] that adapts the content of visualization in an educational system according to the user's domain knowledge. A *behavior-based adaptation* approach is presented in [6], where the system relies on the use of click-stream analysis to detect semantically meaningful patterns, and recommends a specific user support visualization for his current task.

3 Security Event Visualization

Further discussion is centered around the *security event visualization system we are developing* as a part of our research. Visualization, for our purposes here, is taken to mean a way of presenting information in the form of optical images only. Our other premise is that the mouse and keyboard are the only equipment we use to receive the HCI indicators from (more advanced mentioned tracking technologies such as eye gaze, physical and biomedical sensors are beyond this research). For simplicity, it can be presumed that visualization system receives messages from an IDS containing all *information to be displayed*. Any such message should include the following information on the detected event: the IP addresses of the target and source hosts; the type (e.g., scan, remote root attempt, denial-of-service (DoS), etc.) and severity level of the event (low, medium, high, or information messages), as well as its time of occurrence and links to related events. This latter piece of information is received from an IDS correlation module, and is essential for displaying complex [23] or multi-step attacks. All information received from IDS is stored in a separate database that operates with two main entities: hosts and attacks (every attack is unique and may be described recursively: it is either a single event or an array of attacks).

3.1 Use Case Scenario

We will demonstrate our approaches using an exemplary scenario. Let us assume three simultaneous complex attacks detected within ≈10 seconds in the monitored network of ≈1000 hosts. The first of these is a distributed denial-of-service (DDoS) attack against certain network segment, the second is a DDoS attack against a certain host, and the third is a multi-step attack consisting of distributed scanning, root access to monitored host, and a DoS attack against another monitored host. The messages hitting the visualization system during this period fall into the following types: scan (low severity), remote root attempt (medium severity), successful remote root (high severity), DoS (medium severity). In the situation just described, the visualization system receives thousands of messages from the IDS, and the *operator has to come up with the following decisions* as quickly as possible: block the autonomous system and/or ranges of the IP addresses that appear to be involved in DDoS attacks and distributed scanning;reconfigure the captured host and transfer its services to another non-compromised host; and restore functionality of the hosts hit by the DDoS attacks if required.

We decided that the *highest priority data* for operator's decision making include the addresses of both the target host and the attacker, severity of the event, and links to related events. The time of occurrence and type of event were relegated to secondary parameters. The *time of event* for continuously working operator is not that important as he has to resolve first the most dangerous problem rather than the most recent one. As to the *type of event*, it is less important for decision making than the severity level, because the operator does not need any in-depth analysis of the situation; all he has to do is resolve the attack of the highest severity. On the other hand, it is the type of event that defines which decisions the operator can make in a specific situation — the factor taken into account by the visualization system when it prompts the operator to act (as shown in Figure 2).

3.2 Visualization

The traditional approach to visualization system design induces operator to switch among the modes, views and tabs in order to perceive all relevant nuances to making a decision (not to mention that the operator usually cannot interact directly with the visualization to implement his decision). We have assumed that the time taken to find a solution depends, among other factors, on the number of "*search steps*" (or, in other words, on the number of switching context of perception). We therefore stepped aside from the conventional approach and made our visualization system display all its data within one OpenGL-widget. More specifically, we have built our system using Qt library components, with the visualization taking place *within a single QGLWidget*.

It can be considered [22] that the most separable dimensions for perception are color, elements of form (size and orientation), position in the 2D-space and simple motion. We presumed that the most natural, and therefore easiest-to-perceive, visualization of the monitored *network is a graph* with topology corresponding to the reality. Thus, the edges of the graph are consistent with data transmission channels and the vertices correspond to the hosts. That is, of the four most separable dimensions, the position in the 2D-space is set aside for host allocation.

We color-coded the *severity level* of event as it has only four values. Each vertex of the monitored network graph is colored in line with the most severe related event and highlighted with a color shade of the second most severe related event.

External hosts are displayed outside of the space occupied by the network graph. The entire external host space can be represented as a stripe with the X-axis corresponding to the two most significant bytes of an external IP-address (first two in the big-endian notation), and the Y-axis consistent with the two least significant bytes. The more external hosts from close subnets are involved in a certain attack, the greater is their display size (we believe that this feature helps to *identify botnets and distinguish false positive events on legitimate external users*). The stripe may be divided into up to 10 areas, corresponding to different attacks those can be binded to digit keys on a keyboard (see the Interaction section for details). The corners of that space are utilized to display the external hosts which are involved in several attacks.

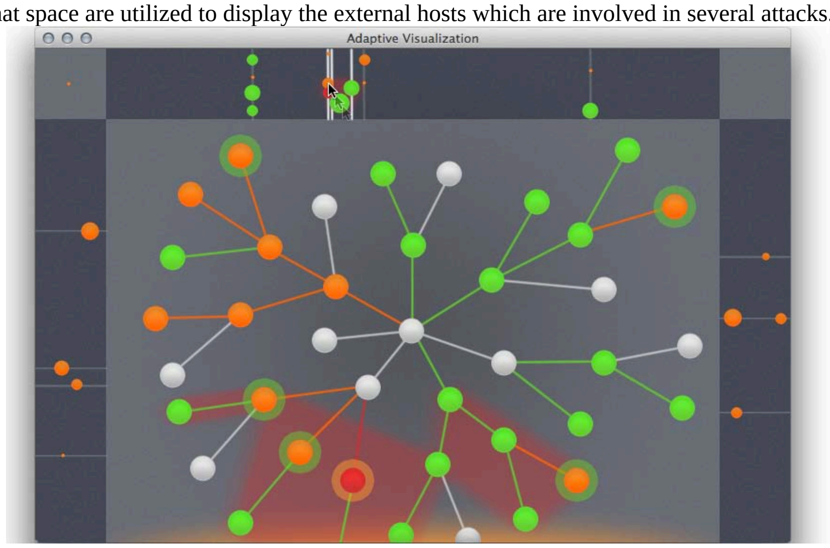

Fig. 1. Visualization system displaying the use case scenario

We decided not to explicitly present such entity as an *event* (in contrast to common event log systems representing each event as a row in their tables), firstly, because of their huge number, and, secondly, because the perception of an event as such has no effect on the sequence of actions by the operator (again, the operator should react not to the event itself, but to a breach in security likely to be resulting from the complex attack consisting of several events).

So too we decided not to explicitly display the *connections among related events* within the same complex attack in order to avoid occlusion due to the large amount of links. Instead, we highlight all hosts involved in a certain attack on operator's interactions with the system.

All events are displayed in the visualization system in *real-time*. New events displayed either through animation of the target host when it falls within the field of vision, or by highlighting the screen border beyond which the target host is hidden (as shown in Figure 1). The latter technique engages the operator's spatial reasoning abilities and prevents his amazement at seeing a few new compromised hosts as he navigates through the network map. The animation fades with time, letting the operator perceive the order of recent events.

3.3 Interaction

Owing to the human *pre-attentive color processing* and the fact that 2D-location and color form the most separable dimension pair [22], the operator may easily perceive the hosts that require his attention. All he has to do is just select, e.g., the red hosts located in a certain area in the network map of his current interest. Generally speaking, the *highest priority breach may not be the severest attack* but the attack that targets the specific network segment (e.g., where most valuable hosts are located). Once he has determined the highest priority problem facing him, the operator moves the mouse pointer over the host of interest. As soon as this host is hovered, a *pie menu* [2] appears around it displaying the severest events related to that host (this technique stems from Fitts's law [13]).

After hovering the menu item of concern, every host related to the attack holding the chosen event becomes highlighted. The highlighting changes with hovering another element of the pie menu. Oftentimes the operator needs to compare host groups related to different attacks. Our visualization system enables operator to save up to ten selections, simply by holding the desired digit key on the keyboard while the hosts of interests are being highlighted. It will be possible to recall the saved group later by pressing the appropriate key. If the operator clicks on the selected event type, the system prompts him to make a decision based on the type of event.

Fig. 2. Interaction with the pie menu

4 Adaptation for Continuous Monitoring

In our visualization system we have implemented the following four adaptation methods: adaptation to current task, adaptation to operator's fatigue level, adaptation to operator's cognitive overload, and adaptation to operator's experience level.

4.1 Task-Centered Adaptation

Generally speaking, any visualization system enables the operator to take only those decisions which have been predefined. It can be presumed that the operator implements his decisions by choosing which *predefined command sequence* to run with what parameters, e.g., which host (as parameter) to block (as predefined command sequence).

According to the object of action, all possible decisions may be divided into two categories — those targeting internal or external hosts. Of these, the *operations with the internal hosts* are performed via direct interaction with the network graph vertices. For example, the operator can transfer functionality of the compromised host to another host in a couple of clicks, with the system highlighting the hosts which are acceptable for the transfer (as shown in Figure 2).

In their turn, the operations with external hosts can be performed in either indirect or direct way. The *indirect operations with external hosts* are performed via the described pie menu around the internal host with an event related to the external hosts of interest (e.g., the operator can block all external hosts those have scanned certain internal host in one click). The *direct operations with external hosts* are performed via direct manipulations with the space of an external hosts. When the operator decides to block a certain group of external hosts, he just hovers one of them and the system automatically highlights proximate external hosts related to the hovered one; if the operator still needs to extend the group, he can simply hold the mouse button and move the pointer towards other external hosts, thus expanding the scope of highlighting (as shown in Figure 1).

4.2 Adaptation to Fatigue Level

The operator's psychophysical condition is taken in this paper to mean the visualization system's notion about the operator's psychophysical state based on following collected HCI data: the *number of actions per various time intervals*, *response speed* (measured by how fast the operator responds to a newly appeared menu or recently displayed event), and *accuracy* in hitting interface control elements (according to Fitts's model [5]).

The system monitors *keyboard keystrokes* and mouse actions. The latter include *mouse button presses* and *mouse movements*. E.g., the interface element is considered to be hit accurately if the mouse pointer was stopped after entering the space occupied by the element and before leaving that space. Mouse movement tracking relies on the standard Qt mechanisms and is possible by enabling the `QWidget::mouseTracking` property.

As fatigue builds up, the visualization system notices changes in the operator's interaction characteristics and adjusts the following visualization parameters: *increases size* of the control elements, *lowers saturation* of the color palette, *diminishes animation* intensity, and automatically *leads the mouse pointer* to the most significant displayed object in the nearby area.

4.3 Adaptation to Cognitive Overload

Since in practice a cognitive overload has a direct bearing on the level of fatigue, we utilize the same HCI data as mentioned above. Besides that, the visualization system is able to estimate a current *cognitive load* by the number of displayed hosts, events and control elements. Sometimes the operator may need an extra time to make a decision, not because of fatigue, but because of the vast amount of displayed information he has to absorb.

The visualization system measures the time of the operator's continuous engagement with it and estimates a likely fatigue level, recognizing the situations of cognitive overload. This done, the visualization system decreases the cognitive load by adjusting the following visualization parameters: by *zooming* the network graph (so there is fewer displayed elements), by *increasing transparency* level of insignificant hosts (e.g., those uninvolved in the most severe attacks, or the ones tagged as low-priority in the network topology setup), by *decreasing the pie menu elements*, and by *aggregating information about a subnetwork on some host* on a higher network hierarchy level and then hiding that subnetwork.

4.4 Adaptation to Experience Level

The operator's experience level is taken in this paper to be a metric of his skills in handling the system's functionality and interface. The visualization system measures that level by analyzing such collected HCI data as the *keyboard shortcuts usage rate* and *time interval between appearance of an assistance tooltip and the operator's response action*. According to these data, the system varies both the cognitive load and its level of assistance.

The visualization system has a predefined base of average operator *activity patterns* (e.g., once there is a new event, the operator hovers the related hosts in one minute if he has no other decisions to make). Whenever the current operator demonstrates a deviation from that pattern, his experience level is considered to be increased (or decreased). Thus, when the system estimates the level to be low, it provides *assistance* in the form of tooltip that points to the element of interest according to the current situation.

5 Conclusions

In this paper we have presented a security event visualization system able to recognize changes in the indicators of its interaction with the operator and adapt its functionality and presentation capabilities to those changes to bring about an improved decision making performance as to speed and accuracy. We have presented methods of the visualization system's adaptation to the current task, as well as the operator's fatigue level, cognitive overload, and experience level. We believe that these adaptation methods can go beyond security visualization to be generalized to any event visualization system and its continuously working operator. As this paper is a part of ongoing research that is on its early stage, we are not ready yet to present any evaluation numbers, but this is our next step.

References

1. Brusilovsky, P., Ahn, J. W., Dumitriu, T., & Yudelson, M. (2006, July). Adaptive Knowledge-Based Visualization for Accessing Educational Examples. In Information Visualization, 2006. IV 2006. Tenth International Conference on (pp. 142–150). IEEE.
2. Callahan, J., Hopkins, D., Weiser, M., & Shneiderman, B. (1988, May). An Empirical Comparison of Pie vs. Linear Menus. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 95–100). ACM.
3. Conti, G., Ahamad, M., & Stasko, J. (2005, July). Attacking Information Visualization System Usability Overloading and Deceiving the Human. In Proceedings of the 2005 Symposium on Usable Privacy and Security (pp. 89–100). ACM.
4. Dutta, S., McCrickard, D. S., Deshmukh, S., & Jouenne, V. (2002, June). Evaluating Benefits and Distractions of Animated Secondary Displays for Attention-Centric Primary Tasks. In 2002 International Conference on Imaging Science, Systems, and Technology (CISST'02), Las Vegas, NV.
5. Fitts P.M. (1954, June). The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. In Journal of Experimental Psychology, 47(6), 381–391.
6. Gotz, D., & Wen, Z. (2009, February). Behavior-Driven Visualization Recommendation. In Proceedings of the 14th International Conference on Intelligent User Interfaces (pp. 315–324). ACM.
7. Horvitz, E., Jacobs, A., & Hovel, D. (1999, July). Attention-Sensitive Alerting. In Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (pp. 305–313). Morgan Kaufmann Publishers Inc..
8. Huberman, B. A., & Wu, F. (2008). The Economics of Attention: Maximizing User Value in Information-Rich Environments. Advances in Complex Systems, 11(04), 487-496.
9. Karat, J., Karat, C. M., & Ukelson, J. (2000). Affordances, Motivation, and the Design of User Interfaces. Communications of the ACM, 43(8), 49-51.
10. Kohlhammer J., et al. (2011). Visual Analytic Representation of Large Datasets for Enhancing Network Security: D1.1 Analysis of Current Practices.
11. Lohse, G. L. (1997). The Role of Working Memory on Graphical Information Processing. Behaviour & Information Technology, 16(6), 297-308.
12. Long, J. (1989). Cognitive Ergonomics and Human Computer Interaction (Vol. 1). Cambridge University Press.
13. MacKenzie, I. S. (1992). Fitts' Law as a Research and Design Tool in Human-Computer Interaction. Human-Computer Interaction, 7(1), 91-139.
14. McCrickard, D. S., & Chewar, C. M. (2006). Designing Attention-Centric Notification Systems: Five HCI Challenges. Cognitive Systems: Human Cognitive Models in Systems Design, 67-89.
15. Merchant, S. (2002). Customizing the Human-Computer Interface to Compensate for Individual Cognitive Attitude: An Exploratory Study. Informing Science, 1043-1049.
16. Shiravi, H., Shiravi, A., & Ghorbani, A. A. (2012). A Survey of Visualization Systems for Network Security. Visualization and Computer Graphics, IEEE Transactions on, 18(8), 1313-1329.
17. Steichen, B., Carenini, G., & Conati, C. (2013). User-Adaptive Information Visualization — Using Eye Gaze Data to Infer Visualization Tasks and User Cognitive Abilities. In Int. Conf. on Intelligent User Interfaces.
18. Toker, D., Conati, C., Carenini, G., & Haraty, M. (2012). Towards Adaptive Information Visualization: On the Influence of User Characteristics. In User Modeling, Adaptation, and Personalization (pp. 274–285). Springer Berlin Heidelberg.
19. Toker, D., Conati, C., Steichen, B., & Carenini, G. (2013). Individual User Characteristics and Information Visualization: Connecting the Dots through Eye Tracking. In Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems,(CHI 2013).
20. Velez, M. C., Silver, D., & Tremaine, M. (2005, October). Understanding Visualization Through Spatial Ability Differences. In Visualization, 2005. VIS 05. IEEE (pp. 511–518). IEEE.
21. Vertegaal, R., Shell, J. S., Chen, D., & Mamuji, A. (2006). Designing for Augmented Attention: Towards a Framework for Attentive User Interfaces. Computers in Human Behavior, 22(4), 771-789.
22. Ware, C. (2012). Information Visualization: Perception for Design. Morgan Kaufmann Pub.
23. Yelizarov, A., & Gamayunov, D. (2009, October). Visualization of Complex Attacks and State of Attacked Network. In Visualization for Cyber Security, 2009. VizSec 2009. 6th International Workshop on (pp. 1–9). IEEE.
24. Ziemkiewicz, C., Crouser, R. J., Yauilla, A. R., Su, S. L., Ribarsky, W., & Chang, R. (2011, October). How Locus of Control Influences Compatibility With Visualization Style. In Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on (pp. 81–90). IEEE.