How to Investigate Interaction with Information Visualisation: an Overview of Methodologies

Margit Pohl and Florian Scholz

Vienna University of Technology, Institute for Design and Assessment of Technology

Abstract. Advanced information visualisation systems offer many different forms of interaction. Nevertheless, we do not know how useful these interactions are. Researchers have suggested to develop a science of interaction. In this paper we discuss which research methods might be appropriate to study interaction with information visualisation systems. We suggest that thinking aloud, log files and eye tracking are promising candidates. These methods enable researchers to study interaction in more detail than other methods. All these methods have strengths and weaknesses. A combination of two or three of these methods might help to overcome the weaknesses.

Keywords: thinking aloud, eyetracking, log files, interaction patterns, triangulation

1 Introduction

Advanced IT systems enable users to interact with them in multiple ways. Users of interactive information visualisation systems, for example, can filter data, show data in more or less detail or represent the data in different visual forms (e.g. as scatterplots or as a graph etc). Such interactions have to be designed appropriately to be useful.

The investigation of various forms of interaction has become more important in HCI in recent years[1]. Mirel[2], for example, argues that human problem-solving and open-ended inquiry consist of different high-level activities (e.g. wayfinding, sensemaking, ...). It is necessary to identify interaction patterns, that is "recurring sets of actions and strategies that have a successful record in resolving particular types of problems"[2, p. 35]. Such interaction patterns can form the building blocks of the strategies the users adopt to solve problems or find relevant information. Such investigations are necessary for HCI in general, but they are especially relevant in information visualisation and visual analytics. In these areas, systems are developed which are supposed to support human reasoning processes and open-ended exploration specifically. Therefore, Pike et al[3] argue that a science of interaction is necessary. They assume that interaction and cognition are closely coupled and that InfoVis should be designed as dialogic systems where both users and computers pose questions and answers. To design such dialogic systems it is necessary to investigate interaction processes in detail.

In a previous publication[4], we presented the results of a study of such interaction processes. We identified several usage patterns. A basic activity sequence consists of a combination of exploration and looking for more detail. Scrolling in combination with selecting an item to get more detail would be an example of such an activity sequence. This activity sequence occurs very often and is usually combined with other activities, like e.g. filtering data or reconfiguring the organisation of information on the screen. Processes of exploration and looking for more detail occur fairly often compared to other activities (e.g. using another form of visualisation). Based on this analysis it is possible to compute transition matrices (see e.g. Ratwani et al [5]). Such transition matrices show the probabilities of transition from one activity to another activity. In our work, we used log files as a basis to compute transition matrices, but it is possible to use other methodologies of data collection for this kind of analysis (e.g. eye tracking).

The following paper discusses possible methods of analysis which are especially appropriate for this kind of investigation. In this context, researchers need methods able to represent the various activities which the users engage in while they work with information visualisations. Therefore, methods like interviews or questionnaires are not really appropriate because they cannot give a detailed overview of sequences of activities. In contrast to that, there are other methods like thinking aloud, eye tracking or log files which provide a fairly comprehensive overview of these activities. These methods can be analysed mathematically by similar methods like transition matrices and Markov models. In this paper, we will first describe these methods (eye tracking, log files, thinking aloud). We will discuss their advantages and disadvantages for the analysis of sequences of interactions with information visualisations. We will briefly discuss whether it is beneficial to combine two or more of these methods. Finally, we will present a few examples of the application of these methods in information visualisation and visual analytics. The discussion of the application examples is not exhaustive and can only give a brief idea of how these methods are applied in this area.

2 Eye Tracking

Eye tracking has been used quite extensively in human-computer interaction and usability research (see e.g. [6]). Goldberg and Wichansky[7] summarize usability recommendations based on eye tracking research. These recommendations concern screen elements such as icons, menues, navigation etc. There is also some relevant research on cognitive load. Eye tracking provides information about attention given to different areas of the screen of a computer and the sequence of the users' gaze directions. In this context, it is necessary to distinguish between foveal and peripheral vision. The fovea is a small part of the retina more or less directly opposite the cornea and lens of the eye. Foveal vision is highly acute and represents an area in the centre of the visual field. Gaze direction reflects foveal vision. Peripheral vision, on the other hand, is less detailed and responsible for the perception of the rest of the visual field. Eye tracking only takes foveal vision into account. Researchers, therefore, only get information about the objects which are in the focus of users' attention.

A fundamental assumption of eye tracking research is the so-called eye-mind hypothesis[7][8] which posits that the gaze direction is an indication of what the user is currently thinking about. There is empirical research indicating that this is not always the case. Duchowski[9] argues that peripheral vision also plays an important role in perception. Subjects often remember objects only seen in peripheral vision[10]. This shows that attention is sometimes directed towards objects on the periphery of the visual fields. People apparently sometimes think about things they do not directly look at. Another difficulty in this context is, that it is sometimes challenging to infer what users really are thinking from gaze directions and scanpaths. If a user looks at an object on the screen for an extended period of time, this might indicate that the object is interesting or, on the other hand, that its functionality is not clear.

Eye tracking has several advantages[8]. Eye tracking provides a large amount of objective data about users' attention processes. Such processes sometimes are very fast and unconscious and, therefore, difficult to investigate. It is fairly unobtrusive (in contrast to, e.g., thinking aloud). Eye trackers also usually come with software which provides researchers with interesting visualisations (e.g. heat maps).

There are also disadvantages. Jacob & Karn [6] summarize disadvantages of eye tracking as follows:

- techical problems
- labor-intensive data extraction
- difficulties in data interpretation

In addition, it should be pointed out that there are usually considerable individual differences concerning the users' scanpaths[11]. This makes generalisation of results of eye tracking studies difficult. A possible solution for this problem is to use small and well defined tasks with clear solutions to get comparable results. Such tasks are usually not very realistic. Results for explorative, open-ended tasks usually differ considerably.

Eye tracking is especially appropriate for investigating interaction with information visualisation tools because it reflects attention to visual stimuli. It can provide detailed information about the users' scanning strategies[12]. Eye tracking can show sequences of the users' activities. Based on this information conclusions about the users' cognitive strategies may be drawn. It should be pointed out, however, that this is not a straightforward process. It has been mentioned above that that the interpretation of eye tracking data is not easy. This problem might be overcome by combining eye tracking with other methodologies (e.g. thinking aloud).

There are some specific challenges concerning the usage of eye tracking in information visualisation. Defining areas of interest (AOIs) or regions of interest (ROI) is especially important for visualisations (we will use the terms interchangeably). AOIs should be related to the research question[10]. Unluckily, it is not possible to provide rules for the definition of AOIs. Another problem is the fact that exploratory tasks are especially important in information visualisation and visual analytics. As mentioned above, it is difficult to investigate such tasks with eye tracking methods. In addition, it is still challenging to analyse dynamic data with eye tracking methodologies[8]. Such data are highly relevant for information visualisation and visual analytics. To conclude, eye tracking seems to have great potential for the analysis of interactions with information visualisations and visual analytics, but it is not yet entirely clear how to use this potential.

3 Log Files

Log files are a very well-known methodology of research in Human-Computer Interaction[13]. They were originally used for the analysis of so-called WIMP interfaces (window, icon, menu, pointer), but nowadays their main application area is the assessment of the usability of websites[14]. Log files enable usability researchers to study website access data of large numbers of web users to find out whether these websites are interesting and attract attention. It should be pointed out, however, that log files resulting from web usage usually only capture navigational behaviour (movement from one webpage to another) or time spent on a site. It is, therefore, difficult to derive valid inferences about the users' cognitive processes from such data. Log files which can serve as data source for cognitive research are often developed by the researchers themselves. The addition of log files to software to analyse the users' behaviour is called instrumenting[15]. The goal of most usability studies using log files, in contrast to that, is to find usability errors and improve the interface. To reach this goal, log files produced automatically by servers or by off-the-shelf software is usually appropriate.

Ivory and Hearst [16] give an overview of various complex systems for the analysis of log files. They distinguish between automated capturing of usage data and the automated analysis of these data. Log file analysis in a narrow sense is the analysis of log file data based on metrics or a mathematical model. The system AMME developed by Rauterberg et al[17] e.g. uses Markov models and Petri nets to investigate the users' problem-solving processes.

Information visualisation systems often offer interaction possibilities going beyond navigation. Users are supported in zooming and panning, filtering the data, choosing different ways of representing data on the screen or other interaction activities (for an overview of categorisations of interactions see e.g. Gotz & Zhou[18], Yi et al[19]). Logging all these activities can enable researchers to get some insights about the users' cognitive processes. It is, for example, interesting to know whether users rather engage in exploration activities without changing the appearance of the visualisation or, in contrast to that, change parameters of the visualisation (through filtering or choosing a different form of visualisation). In addition, users sometimes concentrate on details of the visualisation and sometimes on overall aspects and prediction of the behaviour of the whole system[20]. In addition, log files also provide information about the sequence of activities. This enables researchers to investigate whether there are patterns in the users' behaviour. In this context, it is essential to decide which data to capture and and how to aggregate these data[15]. It is possible to collect data on a keystroke level, but in many cases such data is not very informative. Higher-level data might be more interesting. To decide which data to collect, a clear hypothesis about the cognitive processes involved in the interaction with the information visualisation system is often necessary. These data often have to be aggregated (e.g., when the same goal can be reached by several methods — menus or keyboard shortcuts) and categorised to be useful for an analysis of cognitive processes. Such an analysis process of log files is seldom described or discussed in the literature.

Using log files has several advantages. Compared to other methodologies in Human-Computer Interaction, as e.g. thinking aloud or observation of user actions, log file analysis is less time-consuming, although the amount of work involved in the analysis process should not be underestimated. Log file analysis reflects the actual behaviour of the users, not their attitudes toward a certain piece of software (as in questionnaires or interviews). At times, there can be significant differences in the results between these two approaches. Not every software which is attractive to the users is also efficient. In addition, log file analysis is not intrusive. Users usually do not notice that their actions are being logged. The method, therefore, does not change their behaviour (in contrast to thinking aloud which influences the users' interaction with the system). This is also one of the disadvantages of this methodology, especially in the context of web logs. It is difficult to protect the privacy of a large number of users of web sites in such an analysis. Another serious disadvantage of log files is that it is often difficult to interpret the users' actions without any knowledge of the context in which this interaction happened. When a user repeats an interaction sequence again and again, this might be an indication of a usability problem or an indication of the user's attempt to gain a more thorough mental model of the information represented on the screen. Log files, are, therefore, often used in conjunction with other methods (thinking aloud, observation,...).

4 Thinking Aloud

Thinking Aloud is a methodology which was developed by Ericsson and Simon[21]. The original goal was the investigation of cognitive processes, especially in the context of problem solving activities. It is difficult to analyse such processes exclusively on the basis of observation of visible behaviour because cognitive activities usually happen inside the head, and only the results of such activities can be perceived, not the process itself. Ericsson and Simon looked for a methodology to get more information about the processes which happen during problem solving and the strategies problem solvers adopt to reach their goals. They realised that asking problem solvers retrospectively about their strategies might be misleading because subjects often do not remember what happens in such situations and/or retrospectively interpret their activities according to common assumptions about cognitive activities. They argue that the only way to get some glimpse of actual cognitive activities is to let subjects think aloud, that is, utter any thought which comes to their mind. These utterances represent thought processes because they are not modified by the subjects' assumptions and interpretations. Ericsson and Simon also point out that coding of thinking aloud protocols is an important aspect of the research process. The categories used to interpret the material usually depend on the underlying research question. Ericsson and Simon also describe several coding schemas tailored to clarify problem solving behaviour.

This model is based on several assumptions. The theoretical context of thinking aloud is the information processing model of cognition (see e.g. [22]). Related to this is the assumption that thinking is a serial process which takes place in the working memory and the assumption that thinking aloud provides a complete overview of the cognitive processes. These assumptions are controversial (for a description of these discussions see e.g. [23]). It has been argued that thinking might comprise several different parallel processes. If that is true, thinking aloud would not provide a complete representation of cognitive activities but just one specific strand of the thought processes. It is, for example, difficult to use thinking aloud in the context of highly practised activities like doing arithmetic calculations because such activities are not completely conscious. It is still an open question whether the assumptions of thinking as a serial process and the accuracy and completeness of reports resulting from thinking aloud are correct.

Boren and Ramey [24] give a comprehensive overview of problems which might arise when thinking aloud is applied in usability research. They point out that cognitive psychology is quite different to usability research, and some of the problems encountered when thinking aloud is used in usability research is due to that fact. In usability research, researchers often encounter system crashes or bugs in using the system which is being investigated. The consequence of such problems is that thinking aloud is interrupted, and it is not clear whether an investigation where many such crashes and bugs happen really gives an accurate impression of the users' thought processes. In addition, difficulties in using novel and unstable prototypes often necessitates that researchers talk to subjects to explain relevant issues of using the system. Such behaviour is not acceptable in the context of the original methodology. Boren and Ramey [24] suggest another theoretical framework — speech communication — to allow researchers to conduct more consistent and well defined investigations in usability. The theoretical framework they propose accommodates the current practice in usability research much better than the original theoretical positions formulated by Ericsson and Simon[21]. In addition, we would like to point out that there is another problem with thinking aloud which was already mentioned by Ericsson and Simon[21]. They point out that verbalisation is difficult when problems are presented in a physical form (e.g. the problem of the towers of Hanoi has to be solved by manipulating the disks physically). Subjects concentrate on the manipulation of the objects and are less able to verbalise their thought processes. Using a computer program to solve problems might be a similar situation. Therefore, this problem is especially relevant for usability investigations. In our research, we encountered the problem that subjects often uttered statements like "And now I am moving the cursor from one sector of the visualisation to the next sector...". Such utterances are probably less useful for the analysis of thought processes than the utterances reported by Ericsson and Simon[21] where subjects had to solve problems in their minds.

Nevertheless, thinking aloud has significant advantages compared to methods like eye tracking and log file analysis. It gives insights into the thought processes, goals and motivation of the users of information visualisations. It also provides context to activities of the users and helps researchers to understand what strategies users adopt. Thinking aloud also provides a good impression of sequences of actions, although the granularity of these actions is usually more coarse than the one of log files and eye tracking.

Thinking aloud also has some disadvantages. As mentioned above, the application of thinking aloud in usability research, and more generally to investigate interactions with computer programs poses some problems. In addition, thinking aloud is disruptive and unnatural. Ericsson and Simon[21] argue that thinking aloud only makes the problem solving process longer. Otherwise, the procedure of problem solving is unchanged. There is some reason to assume that this is sometimes not the case[13]. We would also argue that thinking aloud is not a natural behaviour, although people adapt to it fairly quickly.

Despite all these difficulties, we think that results gained from the application of thinking aloud can yield valuable insights about the nature of the interaction processes of users of information visualisations.

5 Mixing Methodologies

The methodologies described above all have strengths and weaknesses. It has often been suggested that a combination of these methodologies might yield more valid results. Lazar et al[15], e.g., argue that log files are difficult to interpret and seldom provide contextual information about users and their cognitive processes. They suggest to combine log file analysis with video recordings or direct observation to get more information about the problem at hand. Gilhooly and Green[25], on the other hand, suggest the combination of thinking aloud and log files. Log files enable the researchers to get more detailed and fine grained information about the users' activities. Sometimes, such data can clarify what users meant with their utterances. Thinking aloud can also be combined with eye tracking[8] to be able to interpret the results of eye tracking studies. Subjects are often motivated to talk while they work with the system. A problem occuring in this context is that their eye movements are affected by thinking aloud, and eye movement recordings do not provide a realistic representation of what people are attending to any more. Webb and Renshaw suggest that retrospective methods should be applied in this case (e.g. discussing gaze plots with the users after the experiment). Holmqvist et al[10] argue that the accuracy of eye tracking data especially suffers in the case of a remote eye tracking system because such systems cannot compensate the effects of fast head movements which usually accompany verbal behaviour. In the case of head-mounted systems, the problem is less serious.

Such combinations of methodologies enable researchers to get more detailed and accurate data. Such an approach is called triangulation (see e.g. [15]). This also increases the reliability and validity of the results. In their book on mixing methods, Creswell and Plano Clark[26] give an overview of how different research methodologies can be combined. Usually, quantitative and qualitative approaches are mixed. There are several reasons why it could be necessary to combine methodologies. Sometimes, one data source is not sufficient (concurrent approach), or initial results have to be explained by another approach (sequential approach). It should be pointed out, however, that mixing methods is always time-consuming.

One example for mixing methods is the study conducted by Jakobsen and Hornbaek[27]. They combined grounded theory, thinking aloud, activity logging, probes and interviews. They argue that the methods in combination "provide stronger evidence of particpants' adoption and use" of a specific software. Activity logging, e.g., does not provide any information about the subjects' intents and the context of their work. Thinking aloud, therefore, complements the data from activity logging.

6 Applications in Information Visualization

To illustrate how these methods could be used we will now give some examples of how they were already applied in the field of information visualisation and Visual Analytics.

6.1 Eyetracking

Goldberg and Helfman[12], in comparing three different graph types using eye tracking, defined Regions-of-Interest (Areas-of-Interest) for the sub-elements of the graphs concerned. They measured the time until their participants' fixations first met the ROIs in which the information necessary for the given task was encoded. They also looked at sequences of fixations, observing, for example, that "The second viewed bar graph was generally scanned left to right (except the first AOI), but the first viewed bar graph was not regularly scanned in a particular direction".[12, p77]

A similar approach is used by Siirtola et al[28] where ROIs are defined for elements of a parallel coordinate visualisation, divided in ROIs deemed necessary and relevant to a given task and those who are not. From the eyetracking data they derive the number of fixations that happened before a fixation in certain ROIs, as well as the total number of fixations and the total time spent in it, comparing these for ROI- and task-type, and relating their results to the interactions offered by the visualisation.[28]

Interestingly, the delay between the start of a task and a fixation on a specific ROI is measured differently in these two studies: Seconds and the number of fixations, respectively. These choices seem to reflect nicely the differing goals and setups of these studies. While Goldberg and Helfman were interested in an direct comparison of performance[12], Siirtola et al conducted their study to assess not absolute performance as such, but qualities of use especially for novice users[28]. For the later objective a proxy for the number of items that were looked at before the first possibility arose of finding the required information might arguably be more useful than a mean of the seconds elapsed.

Another interesting way to use eyetracking is found in Convertino et al[29] who investigate multiple-view visualisations in different configurations. They use the data to derive a measure of how often their participants moved their focus between views, and relate these values to other measurements taken.[29]

Huang et al[30] on the other hand, consciously avoided using eyetracking data for quantitative analysis, after observing that "currently available measures" are "difficult to relate [...] to specific graph elements such as nodes, links or paths" and decided to employ "Eye movement videos", that is, a video of of the screen content overlaid with a marker for the gaze direction during the experiment[30, p3:3]. By analyzing these they developed new theories concerning human behavior reading these graphs and then tested said theory by designing and implementing experiments employing classical time-and-error measurements.[30]

Pointing out, that "[..] eye movement data only tell us how, or what strategies are used in performing tasks. However, they do not tell us why."[30, p3:3] Huang et al also considered Thinking-Aloud as a possible method of investigation. Noting that eye movements happen on a timescale too small for it to be a good fit for Thinking-Aloud and also taking into account possible interference on the gaze behavior of the participants they did not implement Thinking-Aloud as such, but followed an interesting approach of having their participants watch and explain their eye-tracking videos after the core tasks.[30]

6.2 Thinking Aloud

As with eye-tracking there are different kind data and results one can get with Thinking Aloud. In evaluating information visualisation and Visual Analytics systems often something like "insight" is coded, though definitions (for a discussion see, for example, Chang et al [31]) as well as categorization approaches vary, depending among other things on the goals and questions of the study concerned.

Bautista and Carenini[32], for example, use this approach as part of a qualitative and quantitative triangulation in an effort to demonstrate usefulness and improve a Visual Analytics tool for preferential choice. They code their participants' insights for "value" and category, with the categorization scheme closely aligned to their specific research questions, and using both total number and value-weighted number of insights as one factor to compare the two versions of their tool tested.[32].

Coding verbal data for insights generally seems a popular approach in evaluating visual analytics systems, and Saraiya et al[33] conducted a study comparing this kind of evaluation to the more classical method of measuring time and error in the visual analytics domain. They compared three visual analytics techniques of the same dataset and the two methods between subjects. Concerning the insights they use the number of insights (total and in categories, developed in cooperation with domain specialists) and the time spent in the open ended exploration scenario in an analysis of variance. They also note that even though they did not prompt their participants for it they additionally got feedback concerning the usability and visualisations in the insight condition, which might be another interesting factor to consider in the context of methodology choice.[33]

Thinking aloud data is usually collected in sessions spanning hours at most, but it can also be a part of a longer term field study, as with Jakobsen and Hornbaek[27], who studied real-life workplace adoption of a source code visualisation tool spanning multiple weeks employing thinking aloud and logging methods, among others. Here, the thinking aloud sessions with the participants happened at the beginning and the end of the periods the individual subjects were participating, in their respective workplaces. While the aforementioned studies focused on "Insights", here the authors derived categories of visualisation use using open coding. They reason that with programming, "any insights from visualisations are likely to be secondary in relation to higher-level programming objectives"[27, p1580] and targeted the participants "intent and experience of use", and ecologically validity in the face of a real, "cognitively complex task" by being conservative in nudging participants to verbalise. Comparing and contrasting the thinking aloud results with the other methods deployed, they mention that thinking aloud "showed use of the interface across a range of tasks, including some surprising ad-hoc uses"[27, p1586].

6.3 Log Files

There are many things that can be logged in an evaluation of a Visual Analytics system. For example, a study looking at time and error rates would need a record of these somewhere. Here we are more specifically concerned with log files that enable the researcher to inspect the interactions during a task or scenario in more detail. As with Thinking Aloud and eye-tracking, the results gained from these can be both quantitative and qualitative in nature, and might be compared and related to results gained with other methods. Sadly, it is all to easy to imagine that this might create a situation where the space given to the reporting of the analysis of log files has to be balanced with reporting about other approaches and the bigger picture.

Still, even without a detailed description of the analysis processes utilised some examples of log file analysis we found are quite instructive. Bautista and Carenini[32], for example, mention that their log files revealed to them the reason why a certain task showed no significant difference in time-to-complete, even though this task required a functionality that was hidden in one condition and quite prominent in the other. According to them the log files showed that their participants traded the time gained by quicker access to the function for time spent using it.[32]

A study conducted at our institute[34] is an example for a more in-depth report of a statistical analysis of log files created during a study of an visual analytics tool (Gravi++). The log files concerned had a level of detail going down to single tooltip-activations on mouse hover. The study looked at distributions of different interactions in their sample, and also compared subgroups of interactions, within and between. To find out if the task (scenario) influenced participants' interactions, and if so, in what way, they conducted of an analysis of variance, finding a significant effect on the use of certain interactions.[34]

Building on this, a later study[4] explored the issue of emerging patterns of interactions. Beside the aforementioned log file data they also analysed log files generated by another visual analytics application (VisuExplore) for comparison. While the specific interactions afforded by different tools usually also differ, they were able to use a variant of the general categorisation scheme proposed by Yi et al[19] to group the logged events into higher-level activities, to see if the same general patterns could be found in both tools. To that end they qualitatively identified recurring patterns (see figure 1) and also computed transition probabilities between categories.[4]

In the study by Convertino et al[29] already mentioned as an example for eye tracking investigations the authors also statistically analysed their log file data, using the number of different interactions as dependent measure in an ANOVA and finding some significant effects for condition and task type, and also relating them to completion time measurements.[29]

One advantage of using log files is that the data can be potentially collected over a longer period of time and that the method is relatively non-intrusive, therefore viable for medium or long term field studies, as Jakobsen and Hornbaek[27] demonstrated in the study already mentioned above concerning the in-the-field adoption of a fisheye source code visualisation. Visualising and quantitatively analyse the logs, they could observe the real life adoption of the visualisation component over time and in the context of the surrounding workflow.[27]

7 Conclusion

In this paper, we describe three methods for the analysis of interactions with information visualisations. All these methods have advantages and disadvantages. A possibility to overcome the disadvantages would be to combine some of these methods. A careful study of the users' interaction processes with information visualisations could provide more information about cognitive processes accompanying reasoning and problem-solving. It must be mentioned, however, that research in this area is still at the beginning and formulating guidelines based on such insights is still difficult.

We think that in the future it is necessary to discuss methodological problems regarding the investigation of interactions with information visualisations in more detail. Important research has been done to clarify research methods for the evaluation of information visualisations in general[35]. Unluckily, it is sometimes difficult to gain information about methodological issues from publications in the field. Such discussions would be highly relevant to improve our understanding of interaction with information visualisations. The overview given in this paper only addresses three selected methodologies often mentioned in the literature. We think that other methodologies might also be used to analyse interaction processes. In addition, we think that these methodologies should be compared to get a good idea of strenths and weaknesses of these methodologies regarding information visualisation and visual analytics.

8 Acknowledgments

This work is conducted in the context of the CVAST - Centre of Visual Analytics Science and Technology project. It is funded by the Austrian Federal Ministry of Economy, Family and Youth in the exceptional Laura Bassi Centres of Excellence initiative.

References

1. Preece, J., Rogers, Y., Sharp, H.: Interaction Design. John Wiley, New York (2002)
2. Mirel, B.: Interaction Design for Complex Problem Solving. Elsevier/Morgan Kauffman, Amsterdam, Boston, Heidelberg (2004)
3. Pike, W.A., Stasko, J., Chang, R., O'Connell, T.A.: The science of interaction. Information Visualization **8**(4) (December 2009) 263–274
4. Pohl, M., Wiltner, S., Miksch, S., Aigner, W., Rind, A.: Analysing interactivity in information visualisation. KI - Künstliche Intelligenz **26** (2012) 151–159
5. Ratwani, R.M., Trafton, J.G., Boehm-Davis, D.A.: Thinking graphically: Connecting vision and cognition during graph comprehension. Journal of Experimental Psychology: Applied **14**(1) (Mar 2008) 36–49
6. 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. Elsevier, Amsterdam, Boston, London (2003) 573–603
7. Goldberg, J.H., Wichansky, A.: Eye Tracking in Usability Evaluation: A Practitioner's Guide. In: The Mind's Eye – Cognitive and Applied Aspects of Eye Movement Research. Elsevier, Amsterdam, Boston, London (2003) 493–516
8. Webb, N., Renshaw, T.: Eyetracking in HCI. In: Research Methods for Human-Computer Interaction. Cambridge University Press, Cambridge, New York, Melbourne (2008) 35–69
9. Duchowski, A.: Eye Tracking Methodology - Theory and Practice. 2nd edn. Springer, London (2007)
10. Holmqvist, K., Nystrom, M., Andersson, R., Dewhurst, R., Jarodzka, H., van de Weijer, J.: Eye Tracking - A Comprehensive Guide to Methods and Measures. Oxford University Press, Oxford (2011)
11. Nielsen, J., Pernice, K.: Eyetracking Web Usability. New Riders Press, Berkeley, CA (2010)

12. Goldberg, J.H., Helfman, J.I.: Comparing information graphics: a critical look at eye tracking. In: Proceedings of the 3rd BELIV'10 Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization. BELIV '10, New York, NY, USA, ACM (2010) 71–78
13. Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S., Carey, T.: Human-Computer Interaction. Addison Wesley, Wokingham, England, Reading, Mass., Menlo Park, California (1994)
14. Kuniavsky, M.: Observing the User Experience - A Practitioner's Guide to User Research. Morgan Kaufmann, San Francisco, San Diego, New York (2003)
15. Lazar, J., Feng, J.H., Hochheiser, H.: Research Methods In Human-Computer Interaction. Wiley, Chichester (2010)
16. Ivory, M.Y., Hearst, M.A.: The state of the art in automating usability evaluation of user interfaces. ACM Computing Surveys **33**(4) (December 2001) 470–516
17. Rauterberg, M.: A petri net based analysing and modelling tool kit for logfiles in human-computer interaction. In: Proceedings of Cognitive Systems Engineering in Process Control – CSEPC '96 conference. 268–275
18. Gotz, D., Zhou, M.X.: Characterizing users' visual analytic activity for insight provenance. Information Visualization **8**(1) (2009) 42–55
19. 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
20. Friel, S.N., Curcio, F.R., Bright, G.W.: Making sense of graphs: Critical factors influencing comprehension and instructional implications. Journal for Research in Mathematics Education **32**(2) (march 2001) 124–158
21. Ericsson, K.A., Simon, H.A.: Protocol Analysis – Verbal Reports as Data – Revised edition. MIT Press, Cambridge, Mass., London, England (1993)
22. Anderson, J.: Kognitive Psychologie. 3rd edn. Spektrum Akademischer Verlag, Heidelberg (2001)
23. Gilhooly, K., Green, C.: Protocol analysis: theoretical background. In: Handbook of Qualitative Research Methods for Psychology and the Social Sciences. BPS Blackwell, Malden, Mass., Oxford, England, Victoria, Australia (1996) 43–54
24. Boren, T., Ramey, J.: Thinking aloud: reconciling theory and practice. Professional Communication, IEEE Transactions on **43**(3) (sep 2000) 261–278
25. Gilhooly, K., Green, C.: Protocol analysis: practical implications? In: Handbook of Qualitative Research Methods for Psychology and the Social Sciences. BPS Blackwell, Malden, Mass., Oxford, England, Victoria, Australia (1996) 55–74
26. Creswell, J., Plano Clark, V.: Designing and Conducting Mixed Methods Research. Sage, Los Angeles, London, New Delhi (2011)
27. Jakobsen, M.R., Hornbaek, K.: Fisheyes in the field: using method triangulation to study the adoption and use of a source code visualization. In: Proceedings of the 27th international conference on Human factors in computing systems. CHI '09, New York, NY, USA, ACM (2009) 1579–1588
28. Siirtola, H., Laivo, T., Heimonen, T., Raiha, K.J.: Visual perception of parallel coordinate visualizations. In: Information Visualisation, 2009 13th International Conference. (july 2009) 3 –9
29. Convertino, G., Chen, J., Yost, B., Ryu, Y.S., North, C.: Exploring context switching and cognition in dual-view coordinated visualizations. In: Coordinated and Multiple Views in Exploratory Visualization, 2003. Proceedings. International Conference on. (july 2003) 55 – 62

30. Huang, W., Eades, P., Hong, S.H.: Beyond time and error: a cognitive approach to the evaluation of graph drawings. In: Proceedings of the 2008 conference on BEyond time and errors: novel evaLuation methods for Information Visualization. BELIV '08, New York, NY, USA, ACM (2008) 3:1–3:8
31. Chang, R., Ziemkiewicz, C., Green, T., Ribarsky, W.: Defining insight for visual analytics. Computer Graphics and Applications, IEEE **29**(2) (march-april 2009) 14 –17
32. Bautista, J., Carenini, G.: An empirical evaluation of interactive visualizations for preferential choice. In: Proceedings of the working conference on Advanced visual interfaces. AVI '08, New York, NY, USA, ACM (2008) 207–214
33. Saraiya, P., North, C., Duca, K.: Comparing benchmark task and insight evaluation methods on timeseries graph visualizations. In: Proceedings of the 3rd BELIV'10 Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization. BELIV '10, New York, NY, USA, ACM (2010) 55–62
34. Pohl, M., Wiltner, S., Miksch, S.: Exploring information visualization: describing different interaction patterns. In: Proceedings of the 3rd BELIV'10 Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization. BELIV '10, New York, NY, USA, ACM (2010) 16–23
35. Carpendale, S.: Evaluating information visualization. In: Information Visualization – Human-Centered Issues and Perspectives. Springer (2008)

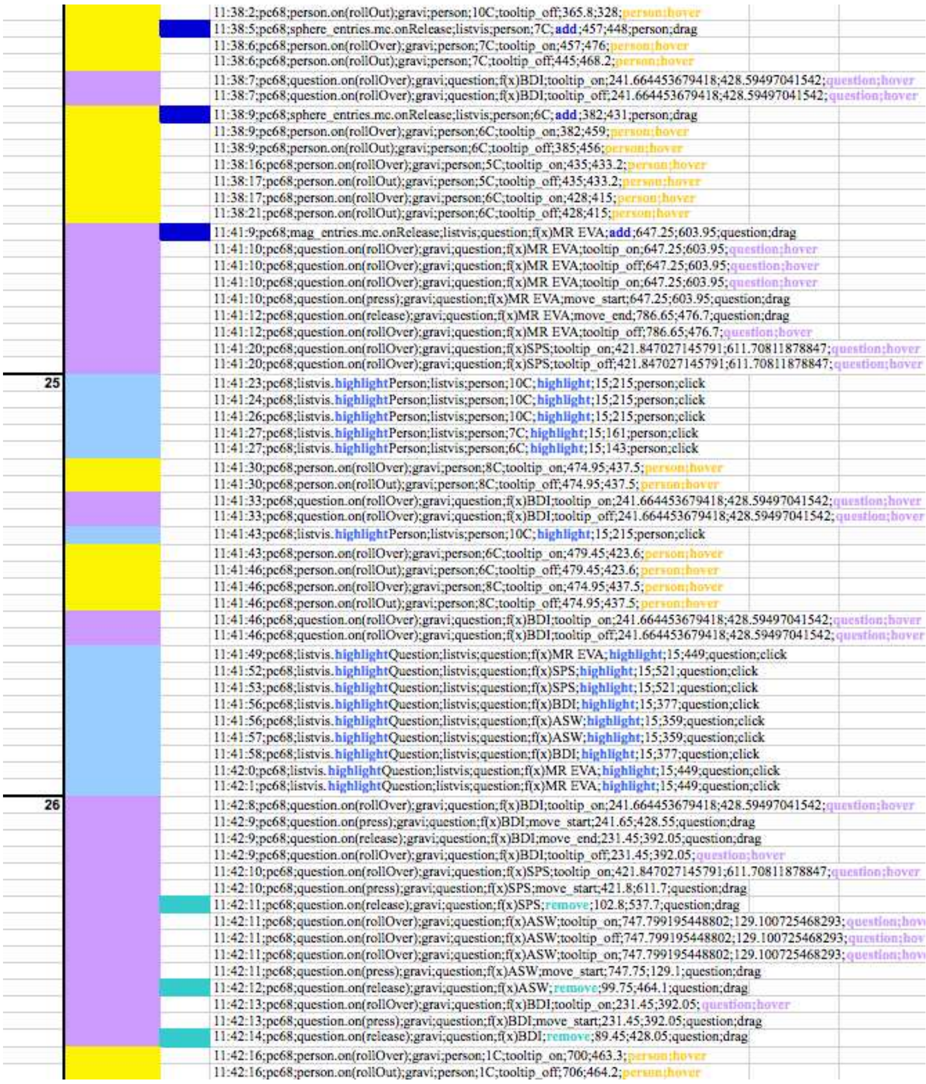

11:38:2;pc68;person.on(rollOut);gravi;person;10C;tooltip_off;365.8;328;person;hover
11:38:5;pc68;sphere_entries.mc.onRelease;listvis;person;7C;add;457;448;person;drag
11:38:6;pc68;person.on(rollOver);gravi;person;7C;tooltip_on;457;476;person;hover
11:38:6;pc68;person.on(rollOut);gravi;person;7C;tooltip_off;445;468.2;person;hover
11:38:7;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_on;241.664453679418;428.59497041542;question;hover
11:38:7;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_off;241.664453679418;428.59497041542;question;hover
11:38:9;pc68;sphere_entries.mc.onRelease;listvis;person;6C;add;382;431;person;drag
11:38:9;pc68;person.on(rollOver);gravi;person;6C;tooltip_on;382;459;person;hover
11:38:9;pc68;person.on(rollOut);gravi;person;6C;tooltip_off;385;456;person;hover
11:38:16;pc68;person.on(rollOver);gravi;person;5C;tooltip_on;435;433.2;person;hover
11:38:17;pc68;person.on(rollOut);gravi;person;5C;tooltip_off;435;433.2;person;hover
11:38:17;pc68;person.on(rollOver);gravi;person;6C;tooltip_on;428;415;person;hover
11:38:21;pc68;person.on(rollOut);gravi;person;6C;tooltip_off;428;415;person;hover
11:41:9;pc68;mag_entries.mc.onRelease;listvis;question;f(x)MR EVA;add;647.25;603.95;question;drag
11:41:10;pc68;question.on(rollOver);gravi;question;f(x)MR EVA;tooltip_on;647.25;603.95;question;hover
11:41:10;pc68;question.on(rollOver);gravi;question;f(x)MR EVA;tooltip_off;647.25;603.95;question;hover
11:41:10;pc68;question.on(rollOver);gravi;question;f(x)MR EVA;tooltip_on;647.25;603.95;question;hover
11:41:10;pc68;question.on(press);gravi;question;f(x)MR EVA;move_start;647.25;603.95;question;drag
11:41:12;pc68;question.on(release);gravi;question;f(x)MR EVA;move_end;786.65;476.7;question;drag
11:41:12;pc68;question.on(rollOver);gravi;question;f(x)MR EVA;tooltip_off;786.65;476.7;question;hover
11:41:20;pc68;question.on(rollOver);gravi;question;f(x)SPS;tooltip_on;421.847027145791;611.70811878847;question;hover
11:41:20;pc68;question.on(rollOver);gravi;question;f(x)SPS;tooltip_off;421.847027145791;611.70811878847;question;hover
25 11:41:23;pc68;listvis.highlightPerson;listvis;person;10C;highlight;15;215;person;click
11:41:24;pc68;listvis.highlightPerson;listvis;person;10C;highlight;15;215;person;click
11:41:26;pc68;listvis.highlightPerson;listvis;person;10C;highlight;15;215;person;click
11:41:27;pc68;listvis.highlightPerson;listvis;person;7C;highlight;15;161;person;click
11:41:27;pc68;listvis.highlightPerson;listvis;person;6C;highlight;15;143;person;click
11:41:30;pc68;person.on(rollOver);gravi;person;8C;tooltip_on;474.95;437.5;person;hover
11:41:30;pc68;person.on(rollOut);gravi;person;8C;tooltip_off;474.95;437.5;person;hover
11:41:33;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_on;241.664453679418;428.59497041542;question;hover
11:41:33;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_off;241.664453679418;428.59497041542;question;hover
11:41:43;pc68;listvis.highlightPerson;listvis;person;10C;highlight;15;215;person;click
11:41:43;pc68;person.on(rollOver);gravi;person;6C;tooltip_on;479.45;423.6;person;hover
11:41:46;pc68;person.on(rollOut);gravi;person;6C;tooltip_off;479.45;423.6;person;hover
11:41:46;pc68;person.on(rollOver);gravi;person;8C;tooltip_on;474.95;437.5;person;hover
11:41:46;pc68;person.on(rollOut);gravi;person;8C;tooltip_off;474.95;437.5;person;hover
11:41:46;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_on;241.664453679418;428.59497041542;question;hover
11:41:46;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_off;241.664453679418;428.59497041542;question;hover
11:41:49;pc68;listvis.highlightQuestion;listvis;question;f(x)MR EVA;highlight;15;449;question;click
11:41:52;pc68;listvis.highlightQuestion;listvis;question;f(x)SPS;highlight;15;521;question;click
11:41:53;pc68;listvis.highlightQuestion;listvis;question;f(x)SPS;highlight;15;521;question;click
11:41:56;pc68;listvis.highlightQuestion;listvis;question;f(x)BDI;highlight;15;377;question;click
11:41:56;pc68;listvis.highlightQuestion;listvis;question;f(x)ASW;highlight;15;359;question;click
11:41:57;pc68;listvis.highlightQuestion;listvis;question;f(x)ASW;highlight;15;359;question;click
11:41:58;pc68;listvis.highlightQuestion;listvis;question;f(x)BDI;highlight;15;377;question;click
11:42:0;pc68;listvis.highlightQuestion;listvis;question;f(x)MR EVA;highlight;15;449;question;click
11:42:1;pc68;listvis.highlightQuestion;listvis;question;f(x)MR EVA;highlight;15;449;question;click
26 11:42:8;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_on;241.664453679418;428.59497041542;question;hover
11:42:9;pc68;question.on(press);gravi;question;f(x)BDI;move_start;241.65;428.55;question;drag
11:42:9;pc68;question.on(release);gravi;question;f(x)BDI;move_end;231.45;392.05;question;drag
11:42:9;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_off;231.45;392.05;question;hover
11:42:10;pc68;question.on(rollOver);gravi;question;f(x)SPS;tooltip_on;421.847027145791;611.70811878847;question;hover
11:42:10;pc68;question.on(press);gravi;question;f(x)SPS;move_start;421.8;611.7;question;drag
11:42:11;pc68;question.on(release);gravi;question;f(x)SPS;remove;102.8;537.7;question;drag
11:42:11;pc68;question.on(rollOver);gravi;question;f(x)ASW;tooltip_on;747.799195448802;129.100725468293;question;hov
11:42:11;pc68;question.on(rollOver);gravi;question;f(x)ASW;tooltip_off;747.799195448802;129.100725468293;question;hov
11:42:11;pc68;question.on(rollOver);gravi;question;f(x)ASW;tooltip_on;747.799195448802;129.100725468293;question;hov
11:42:11;pc68;question.on(press);gravi;question;f(x)ASW;move_start;747.75;129.1;question;drag
11:42:12;pc68;question.on(release);gravi;question;f(x)ASW;remove;99.75;464.1;question;drag
11:42:13;pc68;question.on(rollOver);gravi;question;f(x)BDI;tooltip_on;231.45;392.05;question;hover
11:42:13;pc68;question.on(press);gravi;question;f(x)BDI;move_start;231.45;392.05;question;drag
11:42:14;pc68;question.on(release);gravi;question;f(x)BDI;remove;89.45;428.05;question;drag
11:42:16;pc68;person.on(rollOver);gravi;person;1C;tooltip_on;700;463.3;person;hover
11:42:16;pc68;person.on(rollOut);gravi;person;1C;tooltip_off;706;464.2;person;hover

Fig. 1. Figure 1 shows a section of a logfile. Various activities are colour coded (yellow: explore, blue: select, green/dark blue: reconfigure)