Noname manuscript No.
(will be inserted by the editor)

Analysing Interactivity in Information Visualisation

Margit Pohl · Sylvia Wiltner · Silvia Miksch · Wolfgang Aigner · Alexander Rind

Received: date / Accepted: date

Abstract Modern information visualisation systems do not only support interactivity but also increasingly complex problem solving. In this study we compare two interactive information visualisation systems: VisuExplore and Gravi++. By analysing logfiles we were able to identify sets of activities and interaction patterns users followed while working with these systems. These patterns are an indication of strategies users adopt to find solutions. Identifying such patterns may help in improving the design of future information visualisation systems.

Keywords evaluation · information visualisation · problem solving · software logging

1 Introduction

Interaction design is one of the most challenging areas in Human-Computer Interaction (HCI). Interactivity is supposed to have many favourable effects for computer users, particularly for complex problem solving processes. In such situations, users often adopt recurring sets of activities to reach a solution. It is necessary to identify such sets of activities to be able to support them specifically by an appropriate design of the user interface. This line of study goes beyond traditional usability research and is more concerned with the issue of utility or usefulness of a system.

Complex problem solving is increasingly supported by modern IT systems. One example is information visualisation (InfoVis). In InfoVis, interactivity can support human reasoning processes. Getting insights from information visualisations is often a continuous process of exploration. This process is still not very well understood. The research presented in this paper tries to clarify what kind of activities users of information visualisations engage in and whether any patterns can be identified in this process. Such patterns are an indication of the strategies users adopt to reach their goal. The analysis of such patterns can help to identify useful features of the systems under consideration and to formulate design guidelines for such systems. In our research we compared two InfoVis systems and found interaction patterns which were adopted by users of both systems.

In this paper we first describe the two InfoVis systems we compared (Section 3) and then the study (Section 4). Then we present the results (Section 5) and draw our conclusions and outline necessary future research (Section 6).

2 Related Work

The design of interaction has always been an important area of HCI (see e.g., Preece et al. [21], Winograd [28]). Preece et al., e.g., emphasized the importance of conceptualising interaction appropriately. In recent years, the increased complexity of IT systems has made a more detailed approach necessary. Mirel [14] points out that complex problem solving has to be supported by systems enabling users to engage in open ended inquiry.

M. Pohl · S. Wiltner
Institute for Design and Assessment of Technology, Vienna University of Technology
Tel.: +43-1-58801-187-53
Fax: +43-1-58801-187-93
E-mail: margit@igw.tuwien.ac.at

S. Miksch · W. Aigner · A. Rind
Institute of Software Technology and Interactive Systems, Vienna University of Technology

This process has to be analysed closely to adapt the interface to the exploration processes of the users. In this context, she distinguishes between low-level (e.g., select and save) and high-level activities (e.g., wayfinding, sensemaking). Early HCI concentrated on low-level activities, but she posits that a meaningful analysis should rather emphasize high-level activities. The goal is to identify patterns of inquiry, that is "recurring sets of actions and strategies that have a successful record in resolving particular types of problems." [14, p. 35]. Mirel also points out that complex problem solving is typically supported by interactive InfoVis systems.

HCI aspects in general have been discussed quite broadly in InfoVis in recent years (for an overview see e.g., Kerren et al. [10]). The goal is to develop effective systems adapted to the needs and the cognitive abilities of the users. The strengths and limitations of human perception and reasoning have to be taken into account (Dix et al. [2]).

The importance of interaction for exploration activities in InfoVis is reflected in several recent publications. Pike et al. [17] describe the outlines of a theory of interaction. They argue 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. The theory of distributed cognition which has frequently been proposed to describe interaction between user and computer in HCI (see e.g., Hollan [7]) can clarify issues arising in this context (Liu et al. [12]). Liu and Stasko [13] point out that it is necessary to investigate the relationship between external and internal representations to get a more comprehensive overview of how users interact with information visualisations. So far, the analysis of internal representations (e.g., mental models) is the domain of cognitive psychology, and the analysis of external representations (visualisation tools) the domain of the InfoVis discipline. How these two forms of representations are related to each other is still not very well understood.

Several different frameworks have been proposed to describe the activities of users when they interact with InfoVis systems (see e.g., Wehrend and Lewis [27], Zhou and Feiner [30], Pillat [16]). The goal of such systems is, among others, to assess the quality of InfoVis systems and to describe the activities of the users of such systems. Gotz and Zhou [5] also distinguish between low-level and high-level activities. They propose an intermediate level of granularity bridging the gap between the semantic level and low-level visual interaction events. Yi et al. [29] developed a categorization reflecting the user's side of view and uses intermediate level categories which can be used to categorize users' activities. Kang et al. [9] demonstrate in their study that, based on a categorisation of users' activities, different usage patterns can be identified which, in their case, reflect the properties of the various tools which were used by subjects and indicate that these tools afford different usage strategies. These usage patterns can be compared to the patterns of inquiry Mirel [14] mentions.

Many of the studies above were conducted in the context of the discipline of InfoVis, but, as Liu et al. [12] point out it is necessary to relate this research to investigations coming from cognitive psychology to clarify interaction processes also from the point of view of the humans involved in this process. In this context, work from the area of graph comprehension and visuospatial reasoning might yield an interesting input (see e.g., Tversky [26], Patel [15], Shah et al. [24]).

Based on this research we concluded that the investigation of interaction patterns is a promising research area. Pike et al. [17] especially point out that temporal relationships between user activities (e.g. activities occuring in close temporal proximity) might indicate insights. The analysis of user interactions has to be based on a system of comprehensive categories which allow generalisation of results across different visualisations. Mirel [14] as well as Gotz and Zhou [5] argue that such an analysis should be conducted at a higher level because such activities are semantically rich and can be interpreted meaningfully. The frameworks to describe user activities cited above are all fairly similar. We decided to use the Yi et al. [29] categorization. This categorization is based on a thorough study of previous research in that area and, in addition, reflects users' intentions. It is, therefore, especially appropriate for research on users' strategies. It should be pointed out that, due to the similarity of these frameworks, the results gained with different kinds of these frameworks will be quite similar.

3 Description of the Two Systems

So far we evaluated two InfoVis systems (VisuExplore and Gravi++) that visualised time-oriented data. VisuExplore (Fig. 1) combines well-known visualisations techniques (line chart, bar chart, etc.) and interaction techniques. The purpose of the system is to visualise patients' data collected during long term therapy, e.g., for diabetes. In contrast to Gravi++, VisuExplore shows one patient's data that was collected over the course of the therapy at a glance. All diagrams are positioned on a common time scale thus aiding in comparing data. The default view shows the diagrams of all single variables (e.g., blood sugar, cholesterol) one below the other so that it is possible to compare all the values over time, but variables can also be combined in one diagram. If there are too many diagrams on the screen, scrolling is necessary to see additional variables. Diagrams are customisable and can be adapted by the user (close, new diagram, move diagram, etc.). It is also possible to zoom into the time axis. Tooltips, a measure tool to measure a time interval, and a table panel to show the exact values in a separate window are also available.

Gravi++ (Fig. 2) was designed to support therapists by visualising the development of anorexic patients during therapy. Gravi++ allows to visualise all patients and questionnaires (or if desired a selection of both) at once. To see the patients' development over time animation is used. Users can choose specific time steps one by one or auto-play through them. Two types of icons (patient icons and questionnaire icons) are used. Underlying the visualisation is a spring metaphor; the questionnaire icons attract the patient icons depending on the score patients' answers reached on the questionnaires. Various visualisation methods are available (Starglyph, Attraction Fields, Traces). Add or remove person or questionnaire icons, highlight and tooltip are also available [6]. Furthermore, users can arrange or re-arrange questionnaire icons and adjust the strength with which patient icons are attracted by questionnaire icons. The visual feedback is immediate for all interactions.

The main differences between the two systems are that VisuExplore visualises one single patient's data over the whole time while Gravi++ visualises the data of many patients at a certain point in time (to see other points in time one has to interact with the system). Because VisuExplore uses diagrams it is easy to see approximate values without using any interaction, while Gravi++'s spring metaphor might lead to misinterpretations. On the other hand, Gravi++ supports clustering of several patients' data.

For a more detailed description of VisuExplore and Gravi++ see [19], [23] and [6].

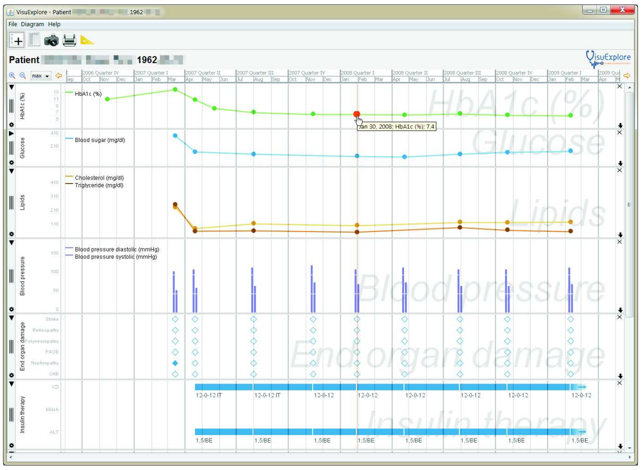

Fig. 1 A screenshot of VisuExplore. The various types of diagrams share a common time axis.

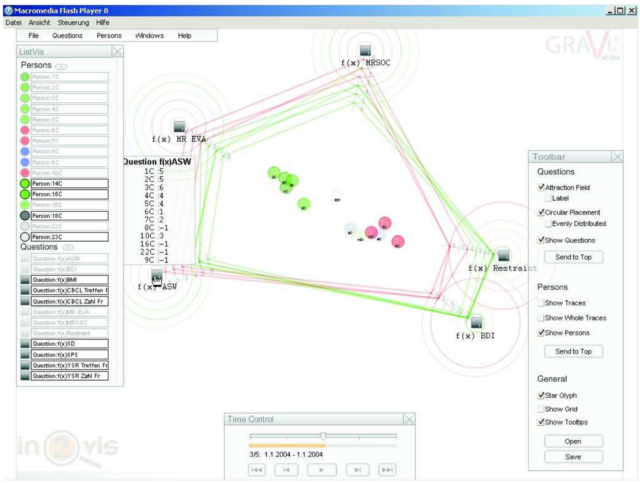

Fig. 2 A screenshot of Gravi++ showing questionnaire icons (squares) and patients icons (circles). Attraction Fields and Starglyph are activated.

4 Description of the Study

The goal of the study described in this paper is to find out whether there are usage patterns or patterns of inquiry which are independent of the InfoVis system used. To reach this goal we compare the logfiles from two studies using two different systems. Both systems represent medical data. In the first case, 32 students interacted with the Gravi++ system for about one hour. In the second case, nine physicians from two hospitals interacted for about 20 minutes with the VisuExplore system, and furthermore 16 students interacted up to one hour with the same system. In both cases, the users had to solve several specific tasks. These tasks were developed by domain experts and were fairly brief and straightforward. Nevertheless, they were rather high-level (find predictors for the success of a therapy; interpret the data of a certain patient). In addition, they included a certain amount of exploration. There was no given answer. Subjects could, to a certain extent, generate their own view of the data, but the insights they produced had to be plausible. It is certainly necessary to compare these results with InfoVis systems from other domains and with other types of tasks to be able to generalize these results. In the case of Gravi++, it was not possible to test physicians. The students making up the sample got a detailed one-hour instruction into the domain and what kind of information they had to look for. Their task was to find information in the way real therapists would do. In contrast to that, the task of the student sample in the VisuExplore study was to find usability problems of the system. We are aware of the fact that it might be problematic to compare data of students and professionals, but we got the impression that the tasks had a more profound influence on activity patterns than the characteristics of the subjects (see also Section 5.3).

For the analysis of users' interactions with InfoVis systems we adopted software logging. One of the advantages of software logging is that it is not intrusive and records activities on a level of detail not accessible to human observers. Software logging has been used in the context of HCI and usability research frequently for quite a long time (see e.g., [11], [20]). Ivory and Hearst [8] provide a detailed overview of data capturing based on logfile analysis. Their emphasis is on usability testing, therefore time spent on task completion and number of errors are important variables in the type of investigation they describe. In explorative interaction with InfoVis systems, other variables are more relevant. Software logging has been used occasionally in evaluation studies of InfoVis systems. Shrinivasan and van Wijk [25] used logfiles to provide an overview of previous analysis processes based on an information visualisation. Dou et al. [3] use software logging to retrieve and understand analysts' reasoning processes. Cowley et al. [1] developed the Glass Box system which also records what analysts do during their work. Gotz and Zhou [5] use logfiles to find out how users of information visualisations get insights. They distinguish between tasks, subtasks (both are domain specific), actions (e.g., query, filter) and events (e.g., mouse-drag, menu-select). They argue that the action level is the most appropriate for the analysis of interaction with information visualisations because it is on the one hand domain independent and on the other hand general enough to represent meaningful user activities.

In the logfiles investigated in our study, we captured activities like "add patient's icon to the visualisation" or "show tooltip" or "draw new diagram" (see also Table 1). A small program was written to count those acitivities. Sometimes, when there were several ways to reach a goal (e.g., via a menu or by drag and drop), the numbers for these activities were added. We found that capturing users' activities through logfiles has certain advantages. On the other hand, it also poses some problems. Some activities of the users are, e.g., not made on purpose. Scrolling or moving the mouse across the screen might be made in a desultory fashion without any intent. Mouseover can, for example, show concrete values even when users do not want to see these values. We took this into account in the interpretation of our data.

As mentioned above our categories are based on the ones proposed by Yi et al. [29] for InfoVis systems. We think that this set of categories is most appropriate for the type of investigation we are doing, because it reflects activities of the users, not the system, and uses, in addition, intermediate level categories. These categories were used to make the results of different information visualisations more comparable. Yi et al. suggest seven categories of interaction and describe them as follows [29, p. 1226]:

Select: mark something as interesting
Explore: show me something else
Reconfigure: show me a different arrangement
Encode: show me a different representation
Abstract/Elaborate: show me more or less detail
Filter: show me something conditionally
Connect: show me related items

The categorisation of activities according to these categories can be difficult sometimes. As we also dealt with time-oriented data we had to expand them to include time. The addition of the time element, e.g., can be achieved in two ways. Either, the original categories are used and tagged with a flag indicating that this category is time-oriented (e.g., time scroll bar move) or an additional category is introduced. We decided to use the original categories with a time flag because we noticed that there were practically no time related activities which could not be captured through the original categories. The categorisation of logfile activities into interaction categories is shown in Table 1. As the goal of this study is to find general interaction patterns across different visualisations, we decided to retain the original seven categories. The flags can be used for a specific analysis of interaction with time-oriented data.

5 Results

In this section we describe the results from the logfile analysis and the interaction patterns subjects adopted while working with the two InfoVis systems. On the left side Table 1 lists the interactions users were able to perform in VisuExplore and Gravi++ [18]. On the right side are the corresponding categories of interaction by Yi et al. [29]. Our practical experience with two different visualisations indicates that this system of categorisation can be applied usefully in the analysis of users' activities. The usage of such systems can help to get more general insights into users' activities across several different information visualisation systems. Categorisation implies that different activities are combined in a category. Nevertheless, these categories express some

Table 1 Interactions and corresponding categories of interaction [29]

Interaction	Categories of Interaction
VisuExplore	
vertical scroll bar moved	explore(scroll)
pan	explore
time scroll bar moved	explore
diagram resized	encode
new diagram	encode
select diagram/ data point	select
timemeasure tool result	time.select
diagram moved	reconfigure
closing diagram, deleting all diagrams	reconfigure
diagram collaps/expand	reconfigure
tooltip shown	abstract/elaborate
opening table panel	abstract/elaborate
zoom	abstract/elaborate(zoom)
Gravi++	
starglyph show	encode
attraction fields show	encode
traces show	encode
move (drag)	reconfigure
add	filter
remove	filter
time (time control)	time.explore
highlight	select
tooltip (hover)	abstract/elaborate

Table 2 Interaction patterns

No. Strategy	Description
VisuExplore	(physicians)
1. scroll-tooltip	scrolling, tooltips
2. scroll-tooltip-move	scrolling, tooltips, moving diagram
3. scroll-tooltip-table	scrolling, tooltips, looking at table
4. scroll-tooltip-new diagram	scrolling, tooltips, creating new diagram
Gravi++	
1. hover-drag	tooltips, moving icons
2. add-remove-hover	adding/removing icons, tooltips
3. time-hover-drag	time control, tooltips, moving icons
4. time-hover	time control, tooltips
5. highlight-hover	highlight, tooltips

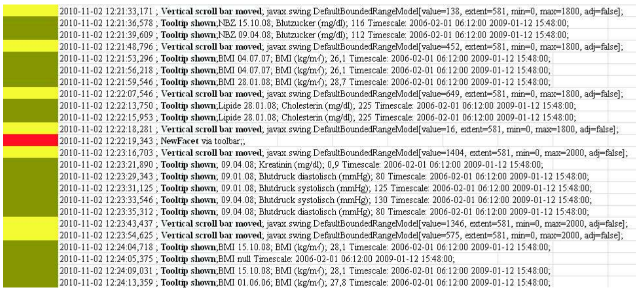
2010-11-02 12:21:33,171 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=138, extent=581, min=0, max=1800, adj=false];
2010-11-02 12:21:36,578 ; **Tooltip shown**;NBZ 15.10.08; Blutzucker (mg/dl); 116 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:21:39,609 ; **Tooltip shown**;NBZ 09.04.08; Blutzucker (mg/dl); 112 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:21:48,796 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=452, extent=581, min=0, max=1800, adj=false];
2010-11-02 12:21:53,296 ; **Tooltip shown**;BMI 04.07.07; BMI (kg/m^2); 26,1 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:21:56,218 ; **Tooltip shown**;BMI 04.07.07; BMI (kg/m^2); 26,1 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:21:59,546 ; **Tooltip shown**;BMI 28.01.08; BMI (kg/m^2); 28,7 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:22:07,546 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=649, extent=581, min=0, max=1800, adj=false];
2010-11-02 12:22:13,750 ; **Tooltip shown**;Lipide 28.01.08; Cholesterin (mg/dl); 225 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:22:15,953 ; **Tooltip shown**;Lipide 28.01.08; Cholesterin (mg/dl); 225 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:22:18,281 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=16, extent=581, min=0, max=1800, adj=false];
2010-11-02 12:22:19,343 ; **NewFacet** via toolbar;;
2010-11-02 12:23:16,703 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=1404, extent=581, min=0, max=2000, adj=false];
2010-11-02 12:23:21,890 ; **Tooltip shown**; 09.04.08; Kreatinin (mg/dl); 0,9 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:23:29,343 ; **Tooltip shown**; 09.01.08; Blutdruck diastolisch (mmHg); 80 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:23:31,125 ; **Tooltip shown**; 09.01.08; Blutdruck systolisch (mmHg); 125 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:23:33,546 ; **Tooltip shown**; 09.04.08; Blutdruck systolisch (mmHg); 130 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:23:35,312 ; **Tooltip shown**; 09.04.08; Blutdruck diastolisch (mmHg); 80 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:23:43,437 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=1346, extent=581, min=0, max=2000, adj=false];
2010-11-02 12:23:54,625 ; **Vertical scroll bar moved**; javax.swing.DefaultBoundedRangeModel[value=575, extent=581, min=0, max=2000, adj=false];
2010-11-02 12:24:04,718 ; **Tooltip shown**;BMI 15.10.08; BMI (kg/m^2); 28,1 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:24:05,375 ; **Tooltip shown**;BMI null Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:24:09,031 ; **Tooltip shown**;BMI 15.10.08; BMI (kg/m^2); 28,1 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;
2010-11-02 12:24:13,359 ; **Tooltip shown**;BMI 01.06.06; BMI (kg/m^2); 27,8 Timescale: 2006-02-01 06:12:00 2009-01-12 15:48:00;

Fig. 3 Section of a logfile from VisuExplore showing an interaction pattern (in this case scroll-tooltip-new diagram). In the block on the left the logged activities are colour-coded (scroll - yellow, tooltip - green, new diagram - red) to enhance readability and the recognition of patterns.

common form of behaviour. Explorative behaviour to find single values and explorative behaviour to identify a predictor, therefore, have something in common. In this case, it is a behaviour which does not change the appearance of the screen (as in encode) or the data used for exploration (as in filter). When analysing a single systems, results always depend on the characteristics of this systems. The analysis of two different systems with very different features enables us to detect more general interaction patterns. Using data from two different target groups (physicians and students for VisuExplore, students for Gravi++) is also an advantage in this context because it indicates that patterns are stable across different users. We tested the usability of both Gravi++ and VisuExplore. During the investigation, no serious usability problems occurred, therefore the users could concentrate on solving the tasks.

5.1 VisuExplore: Interaction Patterns

In this section, we describe interaction patterns adopted by the users of the system VisuExplore. The term interaction patterns, in this context, means specific sequences of activities adopted by users of InfoVis systems. These sequences usually consist of only a few steps and occur again and again. Most of the logfiles consist predominantly of the interaction patterns we identified. They are used in a flexible manner and can be slightly adapted to the situation. They can, e.g., be nested (one interaction pattern is part of another interaction pattern). Therefore, it is difficult to count these interaction patterns automatically. It is also difficult to count these patterns manually because of the huge size of some of the logfiles. We identified these patterns through the visual representation of the logfiles (see [18]). When the logfiles are presented visually, the interaction patterns are quite obvious (Fig. 3).

The system VisuExplore was tested with physicians and students. We looked at these two groups separately, because while physicians were mainly asked to solve three tasks, students were also asked to look for usability problems and to test the system. This influenced the way students interacted with the system. They used more features than physicians and in addition it is sometimes difficult to determine when students finished with a task and started to test the system. We had to consider this when interpreting students' logfiles.

Table 2 shows the strategies physicians used predominantly. All physicians used the basic strategy scroll-tooltip. Additionally, some moved the diagrams, looked at the exact values in the table, or created a new diagram. It should be noted, that none used more than one additional feature of the system (that is, more than three activities in a sequences).

Students on the other hand did not use scrolling and tooltip as much as physicians. This might be due to the fact that they have no medical training and therefore exact values are not as meaningful to them as to the physicians. Students tended to use one feature after the other (new diagram, measure, zoom/pan, move diagram, etc.).

5.2 Gravi++: Interaction Patterns

In [18] we already looked at strategies subjects used to solve tasks in Gravi++. Table 2 shows the commonly used interaction patterns we found.

As in VisuExplore subjects had a preference for certain interaction patterns. They added persons' and questionnaires' icons, decided on a visualisation early on and then, looked at the development over time and used the highlight function. Tooltip is, like in VisuExplore, an often used function, indicating an interest to look at exact values.

There are some differences between the interaction patterns for Gravi++ and VisuExplore. One of these differences is that in Gravi++ scrolling was not possible. Another one is that Gravi++ apparently encouraged the users more to use select. When taking these differences into account, which are due to the specific interface design, many similarities can be detected. This analysis is based on the Yi et al. [29] categorisation which makes similarities visible. Hover-drag (1. Gravi++ pattern) is very similar to scroll-tooltip-move (2. VisuExplore pattern). Time-hover-drag (3. Gravi++ pattern) is also similar to the 2nd VisuExplore pattern. These patterns are based on the basic sequence explore-abstract/elaborate. Activity pattern 4 and 5 in Gravi++ are similar to scroll-tooltip (1.VisuExplore pattern) if we take into account that selection is more common in Gravi++. Interaction pattern 2 in Gravi++ and pattern 4 in VisuExplore also have some similarities, although the categorization does not make this visible. What is in one case encode is in the other case categorised as filter. Both categories reflect activities which introduce new data to the visualisation or remove existing data. We would like to point out that there is one basic sequence which is very often adopted in both systems (explore-abstract/elaborate). This sequences is often combined with other activities (reconfigure, filter, encode). Users tend to prefer activities which do not change the appearance of the screen and which are exploratory in nature. Other activities (like reconfigure, filter, encode) are not adopted as often. A possible reason for this is that users do not want to change their mental model too often, but first try to adapt the existing model to the data. We would also like to point out that activity sequences are usually not longer than three or four activities. These sequences occur again and again. Users of Gravi++, for example, went through time-hover(-drag) sequences again and again to make sense of data from different points in time.

5.3 Transition Probabilities

The quantitative analysis of logfile data is a difficult problem. In the analysis of eye-tracking data, different forms of sequence analysis have been used (see e.g. [4]), but it is an open questions whether these methods can be used to identify small sequences of interactions based on logfile data.

Another method of describing interaction patterns is the computation of transition probabilities. Transition probabilities express whether an activity follows another (or the same) activity frequently or not. They are also an indication of certain sequences of activities. Transition probabilities as indicators for specific strategies were also used by Ratwani et al. [22]. They based the computation of these probabilities on more high-level processes (utterances in verbal protocols), therefore, their results cannot easily be compared to ours. Nevertheless, Ratwani et al. mention a few important facts, e.g., that the processes they observed are cyclical. They especially point out that it is still unclear what cognitive processes underlie the users' activities.

We looked closer at the transition probabilities – which activity follows another activity. To get the numbers in Table 3 and Table 4, we counted all transitions in the logfiles (from one logfile entry to the next) and, based on that, computed the relative frequencies for the transitions from one activity to another (or the same activity, because users might do the same activity again after a short break). We also categorized all activities according to the schema of Yi et al. [29] to be able to compare the results from the two different systems.

Some issues are quite obvious. The highest transition probabilities occur when an activity is followed by the same activity again (e.g., explore-explore). This also indicates that explorative behaviour is predominant (explore-explore in VisuExplore and time-time in

Gravi++). The numbers in Table 3 and 4 also support the idea of interaction patterns described in the last section. We argued that explore-abstract/elaborate is a basic interaction pattern. In table 3, this pattern has the highest probability apart from the transition probabilities for the same activities. In this context, one problem of using transition probabilities can be seen here. This method does not "know" whether an interaction sequences starts or finishes. Therefore, the probability of explore-abstract/elaborate is very similar to the probability of abstract/elaborate-explore. The equivalent of explore-abstract/elaborate in Gravi++ is time/explore-abstract/elaborate. Some indication for this sequence can be seen in the high number for the time-time sequence. In table 4, we can also see that the probability for abstract/elaborate-reconfigure and for filter-abstract/elaborate are quite high. This also supports the interaction patterns described in the last section (Gravi++ interaction patterns 1 and 2).

It is quite obvious that some transitions are significantly more likely than others, but we still conducted a chi-squared test. For this test, we did not take the transition probabilities as they are presented in table 3 and 4. These tables present the data on the transition probabilities in more detail than the Yi et al. [29] categories would provide, to give the readers a comprehensive overview of the users' behaviour. It should also be pointed out that in the tables 3 and 4 transitions with lower probabilities are not shown. For the data from Gravi++ we had to exclude all categories with transition frequencies lower than 5 because the results of chi-squared test become unreliable when too many such values are included in the analysis. For Gravi++ we compared the values for the Yi et al. categories. For the data from VisuExplore we also computed the numbers for the original Yi et al. [29] categories. This means, e.g., that we added the values for explore-explore and for explore(scroll)-explore(scroll). There were no transition frequencies below 5, therefore we did not exclude any of these numbers. For Gravi++, we computed the following chi-squared value: 6562,35. The corresponding chi-squared value for df = 22 and significance level alpha = 0,01 is 40,29. For VisuExplore we computed the following chi-squared value: 18205,52. The corresponding chi-squared value for df = 24 and significance level alpha = 0,01 is 42,98. Both values are highly significant. The chi-squared values are very large because the chi-squared analysis method emphasises larger differences between the expected and the observed values.

In this section, we would also like to compare data coming from physicians and students who were tested when they worked with the VisuExplore system (in contrast to the previous section). This comparison indicates how the tasks assigned by the investigator may change activity patterns. Comparing physicians and students in VisuExplore showed some differences in how they worked. Again this is partly due to the students having been asked to look at the usability which motivated them to use the available interaction methods more extensively. Table 3 shows the main differences between students and physicians. Students scroll and use tooltips less than physicians.

Table 3 VisuExplore: Transition probabilities (in percent)

Transition	Physicians	Students
explore-explore	17,16	13,38
explore(scroll)-explore(scroll)	16,46	4,75
abstr./elab.-abstr./elab.	16,25	13,14
abstr./elab.-explore(scroll)	7,72	2,40
explore(scroll)-abstr./elab.	6,71	2,47
abstr./elab.(zoom)-abstr./elab.(zoom)	4,14	12,49
select-reconfigure	1,41	4,65
explore(scroll)-encode	1,06	1,47
explore(scroll)-reconfigure	0,91	1,09
reconfigure-reconfigure	0,61	4,99
abstr./elab.-encode	0,40	0,93
reconfigure-select	0,35	3,34
abstr./elab.-reconfigure	0,05	0,40

Table 4 Gravi++: Transition probabilities (in percent)

Transition	Probability
time-time	29,38
abstr./elab.-abstr./elab.	14,60
filter-abstr./elab.	9,66
abstr./elab.-reconfigure	8,26
reconfigure-abstr./elab.	6,69

On the other hand, they are more likely to use other features the software has to offer; move, close, expand diagrams or the zoom/pan functions. We think that transition probabilities can provide us with interesting insights, but it is clear that there are still many methodological problems to be overcome.

6 Conclusion

This study presents results that indicate that users of InfoVis probably adopt specific interaction patterns to reach their conclusions. We investigated two different systems and compared them to find out whether interaction patterns might be valid across systems. We present first results indicating that this is the case. We used two different kinds of investigation approaches. On the one hand, we identified interaction patterns in a qualitative way from logfiles. On the other hand, we computed transition probabilities. The results from the transition probabilities study also support the interaction patterns identified in the qualitative study to a certain extent. The interaction patterns were adopted by two different user groups (physicians, students). The interaction sequences are fairly short (3 or 4 activities) and are executed again and again. In our study, we also found that users prefer exploratory activities and avoid changing the visualisation on the screen very often.

These are preliminary results and more research is still necessary with other systems and other approaches. When testing the VisuExplore system, we also recorded thinking aloud protocols. We want to study these protocols and compare the results with the logfile data to find out whether a more high-level approach yields similar results. We also intend to test InfoVis systems from other areas than medicine. In addition, we want to compute transition probabilities for longer sequences of activities, not just two steps. In the long run, we want to relate the results to theories from the psychology of reasoning and visuo-spatial cognition.

Acknowledgements 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. Furthermore, this work was supported by the Bridge program of the Austrian Research Promotion Agency (project no. 814316) and conducted in cooperation with Danube University Krems, Vienna University of Technology, NÖ Landeskliniken-Holding, Landesklinikum Krems, NÖGUS, systema Human Information Systems.

References

1. Cowley P, Nowell L, Scholtz J (2005) Glass Box: an instrumented infrastructure for supporting human interaction with information. In: Proceedings of the 38th annual Hawaii international conference on system sciences (HICSS)
2. Dix A, Pohl M, Ellis G (2010) Perception and cognitive aspects. In: Keim D, Kohlhammer J, Ellis G, Mansmann F (eds) Mastering the information age solving problems with visual analytics. Goslar: Eurographics Association, pp 109-130
3. Dou W, Jeong DH, Stukes F, Ribarsky W, Lipford HR, Chang R (2009) Recovering reasoning processes from user interactions. IEEE Comput. Graph. & Appl., 29, 3, pp 52-61
4. Fabrikant, SI, Rebich-Hespanha, S, Andrienko, N, Andrienko, G, Montello, DR (2008) Novel Method to Measure Inference Affordance in Static Small-Multiple Map Displays Representing Dynamic Processes. Cartogr. J. 45 (3), 201-215
5. Gotz D, Zhou MX (2009) Characterizing users' visual analytic activity for insight provenance. Inf Vis 8, 1: 42-55
6. Hinum K, Miksch S, Aigner W, Ohmann S, Popow C, Pohl M, Rester M (2005) Gravi++: interactive information visualization to explore highly structured temporal data. In: J. Univers. Comput. Sci., 11, pp 1792-1805
7. Hollan J, Hutchins E, Kirsh D (2000) Distributed cognition: toward a new foundation for human-computer interaction research. ACM Trans. Comput.-Hum. Interact., 7, 2:174196
8. Ivory MY, Hearst MA (2001) The state of the art in automating usability evaluation of user interfaces. ACM Comput. Surv., 33, 4, pp 470-516
9. Kang Y, Görg C, Stasko J (2011) How can visual analytics assist investigative analysis? Design implications from an evaluation. IEEE Trans. Vis. & Comput. Graph., 17, 5:570-583
10. Kerren A, Ebert A, Meyer J (2007) Human-centered visualization environments. Springer, Berlin, Heidelberg
11. Kuniavsky, M (2003) Observing the user experience. Morgan Kauffmann Publishers, San Francisco, San Diego, New York
12. Liu Z, Nersessian N, Stasko J (2008) Distributed cognition as a theoretical framework for information visualization. IEEE Trans. Vis. & Comput. Graph., 14, 6:11731180
13. Liu Z, Stasko, J (2010) Mental models, visual reasoning and interaction in information visualization. IEEE Trans. Vis. & Comput. Graph., 16, 6:999-1008
14. Mirel B (2004) Interaction design for complex problem solving: developing useful and usable software. Morgan Kaufmann Publishers, Amsterdam, Boston, Heidelberg
15. Patel VL, Arocha JF, Zhang J (2007) Thinking and reasoning in medicine. In: Holyoak KJ, Morrison RG (eds) The Cambridge handbook of thinking and reasoning. 2nd reprint, Cambridge University Press, Cambridge, New York, Melbourne, pp 727-750
16. Pillat RM, Valiati ERA, Freitas CMDS (2005) Experimental study on evaluation of multidimensional visualization techniques. In Proceedings of the CLIHC '05 conference, pp 20-30
17. Pike WA, Stasko J, Chang R, O'Connell TA (2009) The science of interaction. Inf Vis 8, 4: 263-274
18. Pohl M, Wiltner S, Miksch S, Rester M, Hinum K, Popow C, Ohmann S (2010) Exploring information visualization - describing different interaction patterns, beyond time and errors: novel evaluation methods for Information Visualization. In: BELIV'10, a workshop at ACM conference on human factors in computing systems (CHI 2010)
19. Pohl M, Wiltner S, Rind A, Aigner W, Miksch S, Turic T, Drexler F (2011) Patient Development at a Glance: An Evaluation of a Medical Data Visualization. In: Campos P, Graham N, Jorge J, Nunes N, Palanque P, Winckler M. (eds) INTERACT 2011, Part IV, LNCS 6949, 2011 pp.292-299
20. Preece J, Rogers Y, Sharp H, Benyon D, Holland S, Carey T (1994) Human-Computer Interaction. Addison-Wesley, Wokingham, England, Reading, Mass., Menlo Park, California
21. Preece J, Rogers Y, Sharp H (2002) Interaction Design. John Wiley, New York
22. Ratwani, RM, Trafton, JG, Boehm-Davis, DA (2008) Thinking Graphically: Connecting Vision and Cognition During Graph Comprehension. J Exp Psychol: Appl, 14, 1: 36-49
23. Rind A, Aigner W, Miksch S, Wiltner S, Pohl M, Turic T, Drexler F (2011) Visual Exploration of Time-oriented Patient Data for Chronic Diseases: Design Study and Evaluation. In: Holzinger A, Simonic K.-M (eds) USAB 2011: Information Quality in e-Health. Springer, Heidelberg, LNCS 7058, pp 301-320
24. Shah P, Freedman EG, Vekiri I (2005) The Comprehension of Quantitative Information in Graphical Displays. In:

Shah P, Miyake A (eds) The Cambridge Handbook of Visuospatial Thinking. Cambridge University Press, Cambridge, New York, Melbourne, pp 426-476
25. Shrinivasan, YB, van Wijk, JJ (2009) Supporting Exploration Awareness in Information Visualization. IEEE Comput. Graph. Appl. 29, 5, pp 34 - 43
26. Tversky B (2007) Visuospatial Reasoning. In: Holyoak KJ, Morrison RG (eds) The Cambridge Handbook of Thinking and Reasoning. 2nd reprint, Cambridge University Press, Cambridge, New York, Melbourne, pp 209-240
27. Wehrend SC, Lewis C (1990) A problem-oriented classification of visualization techniques. In: Kaufman AE et al (eds) Proceedings of the First IEEE Conference on Visualization, San Francisco, CA, October 23 - 26 1990. IEEE Computer Soc. Press, Los Alamitos, CA, pp 139-143
28. Winograd T (2002) Interaction Spaces for Twenty-First-Century Computing. In: Carroll J (ed) Human-Computer Interaction in the New Millenium. Addison-Wesley, Boston, San Francisco, New York, pp 259-276
29. Yi JS, Kang Y, Stasko JT, Jacko J (2007) Toward a Deeper Understanding of the Role of Interaction in Information Visualization. IEEE Trans. Vis. & Comput. Graph., 13, 6, pp.1224-1231
30. Zhou MX, Feiner SK (1998) Visual Task Characterization for Automated Visual Discourse Synthesis. In: Proceedings of the CHI '98 conference, pp 393-399