

Te,Te,Hi,Hi: Eye Gaze Sequence Analysis for Informing User-Adaptive Information Visualizations

Ben Steichen, Michael M.A. Wu, Dereck Toker,
Cristina Conati, and Giuseppe Carenini

Department of Computer Science, University of British Columbia
2366 Main Mall, Vancouver, BC, V6T1Z4, Canada
{steichen,mikewu,dtoker,conati,carenini}@cs.ubc.ca

Abstract. Information visualization systems have traditionally followed a one-size-fits-all paradigm with respect to their users, i.e., their design is seldom personalized to the specific characteristics of users (e.g. perceptual abilities) or their tasks (e.g. task difficulty). In view of creating information visualization systems that can *adapt* to each individual user and task, this paper provides an analysis of user eye gaze data aimed at identifying behavioral patterns that are specific to certain user and task groups. In particular, the paper leverages the sequential nature of user eye gaze patterns through *differential sequence mining*, and successfully identifies a number of pattern differences that could be leveraged by adaptive information visualization systems in order to automatically identify (and consequently adapt to) different user and task characteristics.

Keywords: Information Visualization, Eye Tracking, Pattern Analysis.

1 Introduction

Information visualization (Infovis) has long been established as a powerful tool to help humans understand and communicate information in a concise and effective manner. Many different types of Infovis techniques have been devised for a wide variety of applications, but they each typically follow a one-size-fits-all paradigm with respect to users, i.e., their design is seldom personalized to the specific characteristics of users (e.g. perceptual abilities) or their tasks (e.g. task difficulty). However, recent research has shown that user differences can significantly influence task performance and satisfaction with Infovis [1][2][3][4], suggesting that *user-adaptive* Infovis could be of significant benefit to individual users.

Such user-adaptive visualizations are the long-term goal of our research. In particular, we aim to design novel Infovis systems that can (i) identify relevant user and task characteristics in real-time (i.e., during interaction); and (ii) adapt to these characteristics in order to improve a user's visualization processing. This paper contributes to the first of these two challenges with an analysis of user gaze data, aiming to uncover specific gaze behaviors that are indicative of different user and task characteristics. These behaviors could then be used to drive a system that automatically identifies and dynamically adapts to new system users.

Previous work already indicates that different visualizations, tasks, and user characteristics impact a user's eye gaze behavior [5][6]. Furthermore, Steichen et al.

[7] have shown that gaze data can be leveraged to predict several user and task characteristics. However, the features used in this related work only include *summative* measures of gaze data (e.g. *total number of gaze fixations*, *mean of fixations durations*, *number of transitions between two areas*, etc.), as opposed to taking into account the *sequential* nature of eye movements. Furthermore, although the accuracies achieved using these summative gaze features were better than the baseline [7], they are arguably too low for reliable real-world application.

Our paper contributes to this line of research by leveraging the sequential nature of user eye gaze patterns through *differential sequence mining*, and our results show that several sequential gaze behaviors are indicative of specific user/task characteristics. Hence, our results complement prior work with valuable additional features that Infovis systems can leverage to more accurately identify and adapt to such characteristics.

In the rest of the paper, we first discuss related work on adaptive information visualization, pattern detection, and eye tracking. Next, we describe the user study that generated the gaze data used in the paper. We then present our method for analyzing gaze sequences, followed by the analysis results, discussion and future work.

2 Related Work

Information visualization research has traditionally maintained a one-size-fits-all approach, typically ignoring an individual user's needs, abilities and preferences. However, recent research has shown that, for example, cognitive measures such as perceptual speed and verbal working memory have an influence on a user's performance and satisfaction when working with visualizations (also depending on task difficulty) [1][2][8]. Likewise, Ziemkiewicz et al. [3], as well as Green and Fisher [9] have shown that the personality trait of locus of control can impact relative visualization performance for different visualizations. These results indicate that there is an opportunity to apply adaptation and personalization to improve usability.

One of the only attempts to adapt to individual user differences in visualization systems is presented in [10], where a user's visualization expertise and preferences are dynamically inferred through monitoring visualization selections (e.g. how long it takes a user to decide on which visualization to choose). Using this inferred level of user expertise and preferences, the system then attempts to recommend the most suitable visualizations for subsequent tasks. However, this work does not monitor a user's behavior *during* a task, and thus cannot adapt in real-time to help the user with the current task. By contrast, the system developed by Gotz and Wen [11] is, to the best of our knowledge, the only system that actively monitors real-time user behavior during visualization usage to infer needs for intervention. In their work, a user's interactions through mouse clicks are constantly tracked to detect *suboptimal usage patterns*, i.e., click activities that are of a repetitive (hence inefficient) nature. Each of these suboptimal patterns indicates that an alternative visualization may be more suitable to the current user activity. Once these patterns are detected, the system then triggers adaptation interventions similar to those in [10], namely they recommend alternative Infovis. However, there are a number of shortcomings of this work. First, their usage patterns are determined by experts *a priori*, rather than being based on empirical findings. Secondly, they only utilize explicit visualization interactions, therefore the approach is not suitable if a user is only "looking" at a visualization

without manipulating its controls/data. Thirdly, they do not try to infer properties of the tasks (e.g. easy vs. difficult tasks), nor does their approach try to adapt to any user characteristics.

A solution to the first issue (i.e., requiring an expert's *a priori* identification of patterns) is presented in [12], albeit not applied to information visualization. Specifically, Kinnebrew and Biswas [12] present an approach to identify differences in activity patterns between two predefined groups (in their case 'effective learning' vs. 'non-effective learning') through *differential sequence mining*. By extracting all combinations of interaction sub-patterns for both groups, and then comparing the overall pattern occurrence differences between groups, they are able to identify patterns that are most discriminative. Thus, they propose to monitor these specific patterns in an adaptive system to be able to personalize to the inferred user characteristics. In this paper, we leverage this idea of differential sequence mining.

Regarding the second limitation of [11] (i.e. requiring explicit user interactions), we look at using *eye tracking* as an alternative/complementary source of real-time behavior information, since visual scanning and processing are fundamental components of working with any information visualization (and they are in fact the only components for non-interactive visualizations). Although such technology is currently confined to research environments, the rapid development of affordable, mainstream eye tracking solutions (e.g. using standard webcams) will soon enable the widespread application of such techniques [13]. In the fields of HCI and Infovis, research has already been conducted on identifying user gaze differences for alternative visualizations [6], task types (e.g. reading vs. mathematical reasoning) [14], or individual user differences [5]. Techniques used in these studies typically involve calculating *summative* measures (e.g. total number of gaze fixations, mean of gaze durations, number of transitions between two areas, etc.) [6][5] to identify differences between groups of users. By contrast, our work takes into account the *sequential* nature of gaze data. Related work on *sequential scan path analysis* typically involves comparing (using string distance functions) full gaze traces from a complete interaction [15][16]. Various tools have been developed for such analyses, such as eyePatterns [17], allowing the comparison/clustering of whole gaze path sequences. By contrast, our work consists of identifying individual (short) gaze patterns that could be tracked and leveraged *during* a user's interaction. In particular, we propose to apply the above-mentioned idea of *differential sequence mining* [12] to *eye tracking* data, in order to uncover eye gaze patterns that are indicative of specific user and task characteristics. For this purpose, we leverage an additional functionality provided by eyePatterns [17], namely the counting of frequencies of shorter patterns within the full scan paths.

In terms of actually using eye tracking data for real-time prediction, most research has so far focused on identifying user intentions or activities, for example for predicting user intentions in playing simple puzzle games [18], for recognizing user activities within analytics and e-learning tasks [19], or for predicting user learning [20][21]. By contrast, our work is focused on predicting user and task characteristics during information visualization usage. In particular, in our previous work [7], we have proposed to classify users based on these features using machine learning techniques. While the experiment results in [7] have already shown promise for real-time user/task classification, the achieved accuracies were arguably too low for reliable real-world implementation. However, the features used in [7] only consisted of summative measures

(similar to [18]), thereby leaving ample room for improvement through other features. In this paper, we provide additional features to potentially improve the results from previous work, by taking into account the sequential nature of users' gaze traces.

3 User Study

In order to investigate the effect of different task and user characteristics on a user's eye gaze behavior, we designed and ran a user eye tracking study with two basic visualization techniques, namely bar graphs and radar graphs. By choosing two different types of visualizations, we aimed to support the generalizability of our results. The study consisted of a set of tasks that required participants to evaluate student grades in eight different courses. The tasks were based on a set of low-level analysis tasks that Amar et al. [22] identified as largely capturing people's activities while employing Infovis. The tasks were chosen so that each of our two target visualizations would be suitable to support them. A first battery of tasks involved 5 questions comparing the grades of one student with the class average for 8 courses, e.g., "*In how many courses is Maria below the class average?*". A second battery of tasks involved 4 questions comparing the performance of two different students along with the class average, e.g., "*Find the courses in which Andrea is below the class average and Diana is above it?*". Each user performed a total of 28 trials, which included 20 tasks from the first battery (10 per visualization) and 8 tasks from the second battery (4 per visualization).

The long-term user traits that we investigated in this study consisted of the following three cognitive abilities (which we found not to be correlated with each other): *perceptual speed* (a measure of speed when performing simple perceptual tasks), *verbal working memory* (a measure of storage and manipulation capacity of verbal information), and *visual working memory* (a measure of storage and manipulation capacity of visual and spatial information). Perceptual speed and visual working memory were selected because they were among the perceptual abilities explored by Velez et al. [2], as well as among the set that Conati and Maclaren [1] found to impact user performance with radar graphs and a Multiscale Dimension Visualizer (MDV). We also chose verbal working memory because we hypothesized that it may affect a user's performance with a visualization in terms of how the user processes its textual components such as labels. In the following sections, we describe our analysis and results regarding the effect that these *three user traits*, as well as *task difficulty* (defined in section 5.4) have on a user's gaze behavior.

4 Eye-Gaze Processing

An eye tracker captures gaze information through fixations (i.e., maintaining gaze at one point on the screen) and saccades (i.e., a quick movement of gaze from one fixation point to another), which can be analyzed to derive a viewer's attention patterns. In this paper, we analyze a user's attention with respect to so-called 'Areas of Interest' (AOI), which relate to specific parts of the Infovis used in the study.

4.1 Areas of Interest

A total of five AOIs were defined for each of our two visualizations, selected in order to capture their distinctive and typical uses. Figure 1 shows how the AOIs map onto bar and radar graph components respectively.

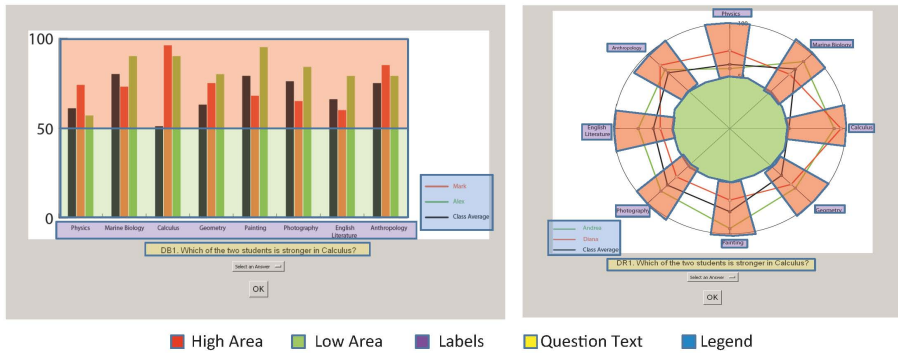


Fig. 1. The five AOI regions defined over bar graph and radar graph

- *High Area (Hi)*: covers the upper half of the data elements of each visualization. This area contains the relevant data values. For the bar graph, it corresponds to a rectangle over the top half of the vertical bars; for the radar graph, it corresponds to the combined area of the 8 trapezoidal regions covering the data points.
- *Low Area (Lo)*: covers the lower half/the inner area of the bar/radar graph.
- *Labels (La)*: covers all the data labels in each graph.
- *Question Text (Te)*: covers the text describing the task to be performed.
- *Legend (Le)*: covers the legend showing the mapping between each student and the color of the visualization elements that represent their performance.

4.2 Sequence Generation and Pattern Frequency Measures

Using the AOI definitions presented above, we converted users’ raw eye gaze data into sequences of AOIs. Specifically, in this paper we define a *sequence as the sequence of AOIs for a complete user scan path of an entire trial*. To generate these sequences, for each trial we mapped each fixation onto one of the 5 AOIs. For example, ‘...Hi-La-Le...’ represents one fixation at the ‘High’ AOI, followed by a fixation at the ‘Label’ AOI, followed by a fixation at the ‘Legend’ AOI. Similarly, ‘...Te-Te-Hi-Hi...’ would represent two consecutive fixations at the ‘Text’ AOI, followed by two fixations at the ‘High’ AOI. In total, we converted the raw eye gaze data from 32 users (each having performed 28 trials), for a total of 725 complete corresponding sequences. Note, some user trials had to be removed due to issues with calibration or other user related matters (e.g., excessive user movements).

Using the eyePatterns tool [17], we then extracted frequencies of *patterns* within these complete sequences. In this paper, *patterns* are defined as *sub-sequences within a sequence*. We extracted patterns in both their “expanded” and “collapsed” form. An expanded pattern comprises all fixations, including consecutive fixations within the same AOI (i.e. *repetitions*). Collapsed patterns lump repetitions into a single AOI.

For each of the different user/task groups to be compared, e.g. high vs. low perceptual speed users (split via a median split, as was done for the other user characteristics), we extracted the occurrence frequencies of expanded and collapsed patterns of length 3 to 7¹. Similar to Kinnebrew et al. [12], we then compared the frequencies of patterns (e.g. ‘Hi-Hi-Hi’, ‘Hi-Hi-La’, ‘Hi-La-La’ etc.) between different groups, in order to see whether some of these patterns are more common for particular groups. More specifically, we compared the following measures:

- *Sequence Support - SS* (s-support in [12]): The number of sequences in a group where the pattern occurs (regardless of how frequently the pattern re-occurs within the sequence) as a proportion of the total number of sequences in the group, i.e.,

$$SS_g(p) = \frac{\text{number of sequences that contain pattern } p \text{ in group } g}{\text{total number of sequences in group } g}$$

For example, if a group (e.g., high perceptual speed users) consists of 300 sequences, and 150 of these sequences contain a particular pattern (e.g. “Hi-La-Te”), then the group’s SS for this pattern would be 50%. The SS measure thereby represents how ‘common’ it is for a pattern to appear for sequences of a particular group (i.e. if a pattern’s SS is very high for a group, this means that most sequences from this group contain this pattern).

- *Average Pattern Frequency - APF* (instance support, or i-support, in [12]): Total number of occurrences of the pattern in all sequences of a group (including reoccurrences within a sequence), divided by the group’s total number of sequences, i.e.,

$$APF_g(p) = \frac{\text{number of occurrences of pattern } p \text{ in group } g}{\text{total number of sequences in group } g}$$

For example, if a group consists of 300 sequences, and a particular pattern appears 600 times in total (including reoccurrences within a sequence), then the group’s APF for this pattern would be 2. The APF thereby represents the ‘reoccurrence frequency’ of a pattern in a particular group (i.e. a pattern with a very high APF for a particular group means this pattern frequently reoccurs in this group).

- *Proportional Pattern Frequency - PPF*: Total number of occurrences of the pattern per group (including reoccurrences) as a proportion of the total number of patterns in the group, i.e.,

$$PPF_g(p) = \frac{\text{number of occurrences of pattern } p \text{ in group } g}{\text{total number of patterns in group } g}$$

For example, if the total number of patterns across all sequences (including reoccurrences) in a group is 10,000 and a particular pattern appears 600 times in total, then the group’s PPF for this pattern would be 6%. The PPF thereby represents the ‘relative frequency’ of a pattern in a particular group, compared to all other patterns from the group (i.e. a pattern with a very high PPF for a particular group means that this pattern is, compared to other patterns, occurring more frequently).

¹ Patterns longer than 7 were very infrequent. Patterns of length 1 and 2 are typically included in simple summative analyses (as presented in Section 2).

While PPF was not included in [12], we believe that it allows an additional dimension of analysis, showing how patterns change in terms of *relative* frequency. In addition, while SS and APF are potentially influenced by overall task length (i.e. longer response times lead to more opportunities for sequences to emerge), this is less of a problem for PPF as it is proportional to the overall number of patterns.

4.3 Statistical Analysis

For our statistical analysis, we followed Kinnebrew et al.'s [12] approach of only considering patterns that are above a certain SS threshold in at least one of the groups to be compared (e.g., at least in either the high or low perceptual speed user group). This ensures that we do not analyze patterns that are too infrequent for any realistic application in an adaptive system. We used the threshold values of 40% for expanded patterns, and 30% for collapsed patterns². To evaluate the statistical significance of SS and PPF differences, we conducted two sets of Pearson's chi-square tests for each user/task characteristic group (our independent measures), separately for expanded and collapsed sequences (with SS and PPF for each pattern being our dependent measures). Since we evaluated multiple patterns for differences within each set, we increased the likelihood of performing a type I error. Thus, we applied the Bonferroni correction to test individual comparisons at significance levels of α/n , where $\alpha = .05$ and $n =$ the number of patterns analyzed in each set (with there being 8 sets in total, i.e. one expanded and one collapsed set per task/user characteristic).

In contrast to SS and PPF, calculating statistical significance for APF differences involves the comparison of means, hence requiring the variances of the pattern frequencies for individual sequences. However, this information is not available in the eyePatterns tool (only overall frequencies are available), and we were hence not able to calculate statistical significance for this measure. We will therefore not discuss APF in the results section below, except for cases where it may provide some complementary insight on intuitively contradictory results.

5 Results

With the goal of finding gaze patterns that may characterize users with specific cognitive characteristics, or are indicative of different task characteristics, we compared the SS and PPF across i) *low vs. user high perceptual speed users*, ii) *low vs. high user verbal working memory users*, iii) *low vs. high user visual working memory users*, and iv) *easy vs. difficult tasks*.

5.1 Perceptual Speed (PS) - Low vs. High

We found a number of patterns that statistically significantly differed in terms of SS and/or PPF between low and high PS users (see Table 1).

² Note that Kinnebrew et al. used 50% as a threshold, however, they allowed for 'gaps' in sequences, which is not supported by the pattern extraction feature in eyePatterns. We therefore slightly lowered the threshold of expanded sequences to 40%. Also, collapsed patterns are generally more infrequent, hence the slightly lowered threshold of 30%.

As shown in the first row of Table 1, patterns involving two fixations at the 'High' AOI (Hi-Hi), followed by a fixation at the 'Label' AOI (La) have a statistically significantly greater PPF for high PS users compared to low PS users. Similarly, the inverted pattern La-Hi-Hi (second row of Table 1) has a statistically significantly greater PPF for high PS users compared to their low PS counterparts. This means that high PS users comparatively make more use of the data labels after/before looking at multiple values displayed in the visualization (i.e. the 'High' AOI). The fact that we did not find a statistically significant difference for this pattern in terms of SS indicates that this pattern is still common for low PS users (i.e. it still occurs in a similar number of sequences), but it does not reoccur as often.

By contrast, 'High' AOI to 'Label' AOI transitions that are broken up by an intermediate fixation at the 'Low' AOI (i.e., Hi-Lo-La) occur more frequently for low PS users. This result was found for both SS and PPF, indicating a strong difference between groups. One possible interpretation for this finding might be that low PS users are less precise when trying to locate the small 'Label' AOIs after visiting one of the visualization values.

Table 1. Pattern differences that occurred between low and high perceptual speed users

Pattern	Sequence Support (SS) <i>Stat. sig. greater for:</i>	Proportional Pattern Frequency (PPF) <i>Stat. sig. greater for:</i>
Hi-Hi-La	-	High PS users
La-Hi-Hi	-	High PS users
Hi-Lo-La *	Low PS users	Low PS users
Te-Te-Lo	Low PS users	-
Te-Te-Te	Low PS users	-
Te (x4)	Low PS users	-
Te (x5)	Low PS users	-

* found for collapsed and expanded patterns

The final set of results in Table 1 shows that it is much more common for low PS users to have some repeated fixations within the 'Text' AOI. In particular, we found that the SS was greater for several 'Text'-related patterns (Te-Te-Te, Te (x4), etc.), suggesting that a single appearance of such a pattern in a user's sequence may indicate low PS. This finding may therefore signify that low PS users generally require more effort to process the larger textual components of visualizations.

None of these results were found using simple summative measures in previous work (e.g. [5]), hence showing that our sequential analysis can indeed reveal complementary features for inferring user characteristics.

5.2 Verbal Working Memory (Verbal WM) - Low vs. High

We found that it is much more common for low verbal working memory users to have *highly repeated fixations at the 'Text' AOI* (i.e., through patterns such as Te (x5, x6, x7)) compared to their high verbal working memory counterparts (see Table 2).

This is in line with previous results found in [5], where the overall proportion of time spent in the text AOI was found to be higher for low verbal working memory users. In fact, our findings further qualify these previous results, indicating that the increased time spent by low verbal working memory users stems from highly repeated transitions within the text AOI, rather than repeatedly coming back to the text AOI after visiting other AOIs. Another result we found was that it is more common for high verbal working memory users to have an increased frequency of the pattern La-Lo-Hi. The interpretation of this result is less intuitive, but nonetheless it may serve as an additional feature for detecting low/high verbal working memory users (as it represents another result not previously found using only summative measures).

Table 2. Pattern differences between low and high verbal working memory users

Pattern	Sequence Support (SS) <i>Stat. sig. greater for:</i>	Proportional Pattern Frequency (PPF) <i>Stat. sig. greater for:</i>
Te (x5)	Low Verbal WM users	-
Te (x6)	Low Verbal WM users	-
Te (x7)	Low Verbal WM users	-
La-Lo-Hi *	High Verbal WM users	-

* found for collapsed and expanded patterns

5.3 Visual Working Memory (Visual WM) - Low vs. High

While the study in [5] found no effect of visual working memory on gaze measures, our sequential analysis reveals that low visual working memory users had increased repetitions in the text AOI (however, only the Te (x6) pattern was statistically significant), as well as two increased patterns involving the 'High', 'Low', and 'Label' AOI (see Table 3). These results are similar to the above findings on perceptual speed and verbal working memory, and represent, to the best of our knowledge, the first results linking visual working memory to eye gaze behavior.

Table 3. Pattern differences for low vs. high visual working memory users

Pattern	Sequence Support (SS) <i>Stat. sig. greater for:</i>	Proportional Pattern Frequency (PPF) <i>Stat. sig. greater for:</i>
Te (x6)	-	Low Visual WM users
Hi-Lo-La *	-	Low Visual WM users
La-Lo-Hi *	-	Low Visual WM users

* found for collapsed and expanded patterns

5.4 Task Difficulty - Easy vs. Difficult

In addition to analyzing the effect of user characteristics, our work also aims to find the impact of characteristics related to a user’s *task*. To this end, we also analyzed pattern differences with respect to the overall ‘difficulty’ of a task. For this measure, we generated, *a posteriori*, an aggregated difficulty value for each task through a

principal component analysis (PCA) using *task completion time* and a user's reported *confidence* on the task (see [5] for a detailed description of this PCA analysis).

The results regarding the pattern frequency differences between *easy* and *difficult* tasks are shown in Table 4. As can be seen from this table, there were many differences regarding repeated fixations of the 'High' AOI (Hi-Hi-Hi, Hi x4, etc.). This difference occurred for both SS and PPF measures, showing that repeated fixations in the 'High' AOI are a strong indicator for a difficult task.

Table 4. Pattern differences for easy vs. difficult tasks

Pattern	Sequence Support (SS) <i>Stat. sig. greater for:</i>	Proportional Pattern Frequency (PPF) <i>Stat. sig. greater for:</i>
Hi-Hi-Hi	Difficult Tasks	-
Hi (x4)	Difficult Tasks	Difficult Tasks
Hi (x5)	Difficult Tasks	Difficult Tasks
Hi (x6)	Difficult Tasks	Difficult Tasks
Hi (x7)	Difficult Tasks	Difficult Tasks
Te-Te-Te	Difficult Tasks	Easy Tasks
Te (x4)	Difficult Tasks	Easy Tasks
Te (x5)	Difficult Tasks	-
Te (x6)	Difficult Tasks	-
Te (x7)	Difficult Tasks	Difficult Tasks
Te-Hi-Hi	Difficult Tasks	-
Te-Te-Hi	Difficult Tasks	-
Hi-Hi-Te	Difficult Tasks	-
Te-Hi-Hi-Hi	Difficult Tasks	-
Te-Te-Hi-Hi	Difficult Tasks	-
Te-Te-Te-Hi	Difficult Tasks	-
Te-Hi (x4)	Difficult Tasks	-

We also found that patterns involving repeated fixations in the 'Text' AOI had a larger SS for more difficult tasks. Interestingly, however, we found that PPF for two of these 'Text' AOI patterns was statistically significantly greater for easier tasks. In order to investigate this counterintuitive result, we also looked at the APF measure, which revealed that the average re-occurrence of repeated 'Text' AOI patterns actually increases for more difficult tasks. However, because the aforementioned 'High' AOI patterns had increased by a much greater extend (up to threefold), the *proportional* occurrence of 'Text' AOI patterns within the difficult task group (as measured by PPF) had actually decreased. While this analysis makes use of the APF measure that we could not check for statistical significance, these very high numbers for the 'High' AOI pattern increase therefore seem to be the most plausible explanation for this seemingly contradictory finding.

In addition to these patterns involving repeated fixations in either only the 'High' or 'Text' AOIs, we found that patterns involving the combination of the two (including intermediate fixations at the 'Low' AOI) also increased for difficult tasks.

None of these results have been found previously using summative measures, which confirms that our sequential pattern analysis was able to find many new discriminatory features.

6 Conclusion, Discussion, and Future Work

In conclusion, our analysis has found a number of gaze behavior differences between different user/task groups during Infovis usage. While some results confirm previous findings that were discovered using simple summative measures (e.g. increased fixations in the text AOI for high verbal working memory users), our novel application of *differential sequence mining* was able to uncover many additional results, including new results for perceptual speed, verbal working memory, visual working memory, and task difficulty. In particular, many results were found in terms of repeated transitions within a single AOI, such as frequent repetitions within the “high” or “text” AOIs being indicative of low values for certain user characteristics. Also, given that these results were found using rather simple tasks and visualizations, it is possible that more complex scenarios would lead to even stronger differences.

In view of building Infovis systems that can adapt to each individual user and task characteristics, these findings provide important indicators as to which particular patterns could be monitored for predicting and adapting to the various characteristics. For example, observing a user’s frequent exhibition of a pattern that is common for a certain user/task group may indicate that the user belongs to this group, and that she may therefore benefit from a specific type of adaptive support. Carenini et al. [24] have already presented a number of possibilities for providing such support in Infovis systems, and have also highlighted that the effect of interventions indeed depends on various user/task characteristics (e.g., a user’s subjective rating of different highlighting mechanisms is shown to be affected by visual WM).

While we found a number of patterns to be indicative for more than one user/task characteristic group (e.g. repeated patterns in the text AOI was found for all 4 characteristics), it is worth noting again that our pattern difference results can be complemented with features from i) other eye gaze features (e.g. non-sequential features, pupil dilation features), and/or ii) other interaction features (e.g. mouse clicks). To this end, the next steps of our research consist of complementing the results from this paper with previously reported summative features [5] to build a combined machine learning model for automatically inferring user/task characteristics.

References

1. Conati, C., Maclaren, H.: Exploring the role of individual differences in information visualization. In: Proc. of the Working Conf. on Advanced Visual Interfaces, pp. 199–206 (2008)
2. Velez, M.C., Silver, D., Tremaine, M.: Understanding visualization through spatial ability differences. In: IEEE Visualization, VIS 2005, pp. 511–518 (2005)
3. Ziemkiewicz, C., Crouser, R.J., Yauilla, A.R., Su, S.L., Ribarsky, W., Chang, R.: How locus of control influences compatibility with visualization style. In: 2011 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 81–90 (2011)
4. Dillon, A.: Spatial-semantics: How users derive shape from information space. *J. Am. Soc. Inf. Sci.* 51, 521–528 (2000)
5. Toker, D., Conati, C., Steichen, B., Carenini, G.: Individual user characteristics and information visualization: connecting the dots through eye tracking. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 295–304 (2013)

6. Goldberg, J., Helfman, J.: Eye tracking for visualization evaluation: reading values on linear versus radial graphs. *Inf. Vis.* 10, 182–195 (2011)
7. Steichen, B., Carenini, G., Conati, C.: User-adaptive information visualization: using eye gaze data to infer visualization tasks and user cognitive abilities. In: Proceedings of the 2013 International Conference on Intelligent User Interfaces, pp. 317–328 (2013)
8. Toker, D., Conati, C., Carenini, G., Haraty, M.: Towards adaptive information visualization: On the influence of user characteristics. In: Masthoff, J., Mobasher, B., Desmarais, M.C., Nkambou, R. (eds.) UMAP 2012. LNCS, vol. 7379, pp. 274–285. Springer, Heidelberg (2012)
9. Green, T.M., Fisher, B.: Towards the Personal Equation of Interaction: The impact of personality factors on visual analytics interface interaction. In: 2010 IEEE Symposium on Visual Analytics Science and Technology (VAST), pp. 203–210 (2010)
10. Grawemeyer, B.: Evaluation of ERST – an external representation selection tutor. In: Barker-Plummer, D., Cox, R., Swoboda, N. (eds.) Diagrams 2006. LNCS (LNAI), vol. 4045, pp. 154–167. Springer, Heidelberg (2006)
11. Gotz, D., Wen, Z.: Behavior-driven visualization recommendation. In: Proceedings of the 14th International Conference on Intelligent User Interfaces, pp. 315–324 (2009)
12. Kinnebrew, J.S., Biswas, G.: Identifying Learning Behaviors by Contextualizing Differential Sequence Mining with Action Features and Performance Evolution. In: Proc. of EDM, 5th Int. Conf. on Educational Data Mining, pp. 57–64 (2012)
13. Sesma, L., Villanueva, A., Cabeza, R.: Evaluation of pupil center-eye corner vector for gaze estimation using a web cam. In: Proceedings of the Symposium on Eye-Tracking Research & Applications, pp. 217–220 (2012)
14. Iqbal, S.T., Bailey, B.P.: Using eye gaze patterns to identify user tasks. Presented at the The Grace Hopper Celebration of Women in Computing (2004)
15. Duchowski, A.T., Driver, J., Jolaoso, S., Tan, W., Ramey, B.N., Robbins, A.: Scanpath Comparison Revisited. In: Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications, pp. 219–226 (2010)
16. Madsen, A., Larson, A., Loschky, L., Rebello, N.S.: Using ScanMatch Scores to Understand Differences in Eye Movements Between Correct and Incorrect Solvers on Physics Problems. In: Proc. of Symp. on Eye Tracking Research & Applications, pp. 193–196 (2012)
17. West, J.M., Haake, A.R., Rozanski, E.P., Karn, K.S.: eyePatterns: software for identifying patterns and similarities across fixation sequences. In: Proceedings of the 2006 Symposium on Eye Tracking Research & Applications, pp. 149–154 (2006)
18. Eivazi, S., Bednarik, R.: Predicting Problem-Solving Behavior and Performance Levels from Visual Attention Data. In: 2nd Workshop on Eye Gaze in Intelligent Human Machine Interaction at IUI 2011 (2011)
19. Courtemanche, F., Aïmeur, E., Dufresne, A., Najjar, M., Mpondo, F.: Activity recognition using eye-gaze movements and traditional interactions. *Interac. Comp.* 23, 202–213 (2011)
20. Kardan, S., Conati, C.: Exploring gaze data for determining user learning with an interactive simulation. In: Masthoff, J., Mobasher, B., Desmarais, M.C., Nkambou, R. (eds.) UMAP 2012. LNCS, vol. 7379, pp. 126–138. Springer, Heidelberg (2012)
21. Bondareva, D., Conati, C., Feyzi-Behnagh, R., Harley, J.M., Azevedo, R., Bouchet, F.: Inferring Learning from Gaze Data during Interaction with an Environment to Support Self-Regulated Learning. In: Lane, H.C., Yacef, K., Mostow, J., Pavlik, P. (eds.) AIED 2013. LNCS, vol. 7926, pp. 229–238. Springer, Heidelberg (2013)
22. Amar, R., Egan, J., Stasko, J.: Low-Level Components of Analytic Activity in Information Visualization. In: Proc. of 2005 Symp. on Information Visualization, pp. 15–21 (2005)