Visual Exploration of Hierarchical Data Using Degree-of-Interest Controlled by Eye-Tracking

Nelson Silva[1,2], Lin Shao[1], T. Schreck[1], Eva Eggeling[1,2], and Dieter W. Fellner[1,3]

[1]Institute for Computer Graphics and Knowledge Visualization, Graz University of Technology, Austria
[2]Fraunhofer Austria Research GmbH, Austria
[3]Interactive Graphics Systems Group, Technische Universitaet Darmstadt, and Fraunhofer IGD, Germany

Abstract

Effective visual exploration of large data sets is an important problem. A standard technique for mapping large data sets is to use hierarchical data representations (trees, or dendrograms) that users may navigate. If the data sets get large, so do the hierarchies, and effective methods for the navigation are required. Traditionally, users navigate visual representations using desktop interaction modalities, including mouse interaction. Motivated by recent availability of low-cost eye-tracker systems, we investigate application possibilities to use eye-tracking for controlling the visual-interactive data exploration process. We implemented a proof-of-concept system for visual exploration of hierarchic data, exemplified by scatter plot diagrams which are to be explored for grouping and similarity relationships. The exploration includes usage of degree-of-interest based distortion controlled by user attention read from eye-movement behavior. We present the basic elements of our system, and give an illustrative use case discussion, outlining the application possibilities. We also identify interesting future developments based on the given data views and captured eye-tracking information.

In: W. Aigner, G. Schmiedl, K. Blumenstein, M. Zeppelzauer (eds.): Proceedings of the 9th Forum Media Technology 2016, St. Pölten, Austria, 24-11-2016, published at http://ceur-ws.org

1 Introduction

In this paper, we consider using eye-tracking information to create an adaptive Visual Analytics system. A main idea of Visual Data Analysis is to support analytic reasoning by interactive visual interfaces to data. This typically involves the integration of capabilities of data analysis in terms of visual information exploration, and the computation capabilities of computers to create capable knowledge discovery environments [KMSZ06, ABM07]. The need for effective data analysis solutions is obvious as more and more digital information is being generated and collected in many areas, e.g., in medicine, science, education, or business. Data analysis problems can be diverse, such as the amount and speed of data being generated, while it needs to be processed and analyzed. Other problems are related with the filtering, aggregation and visualization of this same data.

There are different types of data, among which graphs and networks are important data structures to model many relevant data sets [vLKS+11, HMM00]. The sizes of graphs grow quickly in many domains, and these sizes hinder visual exploration of the data, as visualization of large graphs is a challenge. As the above cited surveys show, there are numerous graph visualization methods available, however displaying just a few thousands of nodes effectively remains a problem. Therefore, data reduction, e.g., by clustering/collapsing of graph nodes is a common approach to limit the data complexity. In terms of hierarchies (trees, or dendrograms) these can be reduced to show only a certain depth of the hierarchy, and group together all elements of a sub tree at the sub tree root.

Traditionally, we can use pointing devices to zoom, pan and navigate to other areas of the graph, expand/collapse nodes, or adjust the level of abstraction [GST13]. However, several problems may arise with the mouse navigation, zoom and expand collapse strategies. When the user pans a graph, the mouse click can be anywhere on the graph, including empty parts of it; also, when the user stops panning, the mouse cursor is "parked" somewhere in the visualization, and no useful information can be inferred regarding the user intents.

Existing *eye-tracking* devices allow to track the fixation areas of a user in front of a display. Among others, eye-trackers are often used for evaluation purposes, or to experimentally study human visual attention. In our work we are interested in the question, if eye-tracking information can benefit the visual data exploration process, in addition or as an alternative to standard interaction approaches. Nowadays, we can use affordable eye-tracker devices [SZCJ16], like the EyeTribe system[1] to monitor the behaviors of the users while exploring a visualization. According to our experience, the device allows a useful tracking of the user gazing on specific regions and individual nodes, on a comparably large tree to be explored.

We present a concept and preliminary implementation of an approach to apply eye-tracking for the purpose of supporting visual exploration of large graphs. Our assumption is that the user gaze indicates areas of interest in a tree, and consequently we can use this information to dynamically expand or compress parts of the tree. Furthermore, we can also capture a *visual history* of the exploration process during which a user explores a tree view, with applications for example of 'replaying' analysis sessions or documenting interesting findings done along the way. This paper is our first step towards an experimental system by which we can explore design alternatives for eye-based interaction and visualization, as well as to conduct user studies.

2 Related Work

We briefly provide an overview of possible applications of eye tracking in evaluation, and as an interaction modality. We also discuss visualization of hierarchic data and degree-of-interest techniques.

2.1 Eye-Tracking and Applications

Eye-movement tracking is a method that is used to study, among others, usability issues in Human Computer Interaction (HCI) contexts. Pool and Ball [PB05] give an introduction to the basics of eye-movement technology, and present key aspects and metrics of practical guidance in usability-evaluation studies for capturing user eye movements as an input mechanism to drive system interaction.

Etemadpour et al. [EOL14] address eye-movement tracking on user studies regarding the accomplishment of typical analysis tasks for projected multidimensional data, such as tasks that involve detecting and correlating clusters. The authors examine and draw conclusions on how layout techniques produce certain characteristics that change the visual attention pattern.

In lab-based user experiments using eye-movements tracking, large and complex gaze trajectory data sets are generated. There is work which develops tools to help understand eye-movement patterns [BKR+14]. These should support the definition and exploration of a large number of areas-of-interest (AOIs). Eye-movement tracking data is usually analyzed using different methods [HNA+11] and visualization techniques [BKR+14]. Our work follows an AOI-based approach.

One important question refers to the justification of why should one use eye-movement tracking and not just the typical pointing devices, such as, the mouse. Many past studies debate the correlation between eye-movement tracking and mouse movements. These studies presented values from as high as 84% (in a study from 2001) [CAS01], to 69% [Coo06], to as low as 32% [RFAS08] (in a study from 2008) of correlation between eye and mouse movements. These results are usually dependent on the design of the user interfaces.

Another relevant discussion is centered on the many advantages of using eye-movement tracking analysis and on how to perform a correct eye-movement tracking evaluation [JK03]. Eye-movement tracking allows for a fast and continuous tracking of the interest of the users in real time, allowing the detection of moments of confusion, indecision and high interest regions [GW03]. Also, previous studies discuss an important link between cognitive processes and eye-movements [Hay04]. The accuracy of eye-movement tracking can be kept high by designing a user interface where the size of the areas of interest is big enough and in accordance to the eye-tracker characteristics and the experiment setup. More and more, new eye-trackers are also less affected by negative technical factors, e.g., the users' head movements that usually reduce the accuracy of the eye-tracker device and calibration difficulties.

2.2 Hierarchy Visualization and Focus-and-Context

In this work we consider visualization of hierarchic data. While hierarchies arise in many contexts, one prominent use of hierarchies is in data clustering. Generally, hierarchical graphs together with one of the many clustering techniques [MRS08] can form beneficial tools for the visual exploration of large data sets.

[1] `https://theeyetribe.com` (accessed 09/2016)

This visual exploration can be done in the form of trees (or dendrograms), due to their potential for visual abstraction [CdART12]. In a hierarchical clustering, users may chose the level of detail by which they explore data. Areas more close to the root contain more aggregate information, and areas closer to the leafs include more detail data.

An effective interaction technique for navigating large visualization spaces is to control the level of detail information shown throughout a given visualization. Furnas [Fur86] defines a degree of interest function (DOI-function) where to each node in the graph structure an interest score is defined. This score in turn is used to expand important areas while reducing other less important areas. Lamping et al. [LR96] demonstrated a focus+context (fisheye) scheme for visualizing and manipulating large hierarchies. Generally, the expansion or reduction can operate on different aspects, e.g., on the geometric, semantic or data-oriented level. Previous work [PGB02] was done regarding the dynamic re-scaling of branches of the tree to best fit the available screen space with an optimized camera movement. Concerning the aspect of developing adaptive visualizations, an important survey [CK15] was presented that highlights many techniques for emotion-driven detection, measurement and adaptation, among others. These are very relevant for our concept of adaptation based on degree-of-interest.

3 Concept for Visual Exploration of Hierarchical Data Guided by Eye Tracking

Next, we introduce our concept (Figure 1) for exploring hierarchical data based on degree-of-interest, guided by eye tracking. We will exemplify our concept by using scatter plots as the leaf elements in the hierarchy, which are to be explored by a user. The hierarchy is created by a hierarchical clustering algorithm using feature-based similarity between scatter plots. Our approach relies on eye-tracking to determine the degree-of-interest, which in turn distorts (i.e., magnifies/compresses) the hierarchy display.

3.1 Considered Application: Hierarchy of Scatter Plots

For the exploration of complex data sets, target visualization techniques such as scatter plots, parallel coordinates or glyph representations can be used to discover interesting findings in the data. In our approach, we rely on a set of scatter plot visualizations to represent all pairwise combinations of a high-dimensional data set. To explore a potentially large set of scatter plots hierarchically, we apply hierarchic clustering. Input to the clustering is a distance matrix between the set of scatter plots. The latter is obtained making use of image features, which have been shown to work well for the comparison of scatter plots [SvLS12]. More precisely, we compute a 25-dimensional intensity histogram for each plot. Then, we use the Euclidean distance between histograms to compute the distance (average linkage) of each plot. Using these visual features, the scatter plots can be arranged hierarchically (e.g., in a tree or circular layout) and the exploration for visually similar plots becomes more efficient.

3.2 Hierarchical Layout of Scatter Plots

Typically, there are a large number of scatter plot views for a high-dimensional data set, these views grows quadratically with the number of data dimensions. Specifically, an n-dimensional numeric data set can be represented in $\frac{n\times(n-1)}{2}$ distinctive views using two distinctive dimensions. To facilitate the exploration, we take the computed feature vectors of the scatter plots and apply a hierarchical clustering to structure the plots based on their visual similarities. Thus, we receive a structured representation of the space of scatter plots that arranges similar scatter plots spatially close. To create the hierarchy structure, we compute the average distance (average linkage) of each scatter plot and build a dendrogram tree, which contains all scatter plots on the leaf node level, see Figure 3. As usual, the dendrogram height describes the similarity (histogram distance) of the scatter plots.

3.3 Degree of Interest for Navigation of the Hierarchy

The above described dendrogram provides a useful spatial organization of the input space (scatter plots). Yet, the tree may still be large and complex, especially if we have a large number of leaf and internal nodes to inspect and compare. Hence, we introduce *spatial distortion* to enlarge parts of the tree currently being looked at by the user, while visually aggregating the remainder of the tree. To this end, we apply eye tracking using an EyeTribe (see Section 1) setup to track user gazes. Specifically, we measure the user attention on the tree nodes to compute a degree-of-interest (DOI) score for the elements of the dendrogram. Initially, we show the overall dendrogram using semantic zoom to fit the whole hierarchy onto the display space. From there, the user starts the graph exploration from any point in the view space. While the user navigates through the view space, the eye tracker captures the gaze path. When the user explores specific branches of the tree or local nodes, the eye gaze path and eye-fixation durations are recorded for each link and node of the tree. Therefore, besides a measure of interest based on similarity between scatter plots, we can now update interest metrics based on time and number of visits to a node.

Potentially, this recording can also be done for local parts of the scatter plots, i.e., tracking if the user is dedicating more viewing time to certain local areas in a plot. Such analysis may be useful to detect e.g., correlations, dense areas or clusters in a given plot. Each scatter plot involves the representation of variables (for x and y axes respectively), the interest of the user on these variables (axis) can also be tracked. In the next section, we apply the current eye gaze location in order to focus the display using semantic zoom. Conceptually, more applications are possible (see also Section 5).

3.4 Degree-of-Interest Visualization Using Eye-Tracking

We apply eye-movement information to allow the user to navigate through a hierarchy of clusters of scatter plots using semantic zoom. We define an *eye-tracking mode*, which if enabled, controls the expansion and collapsing of sub-trees in the display based on eye fixation. Specifically, the area where the user looks at is visually expanded, revealing the scatter plots under the sub-tree. The neighboring (remaining) sub-trees are represented using just node and link symbols. While they do not show particular scatter plots, this reduced representation is still indicating basic data properties like number of scatter plots represented, or structure of the similarity relationships within the dendrogram.

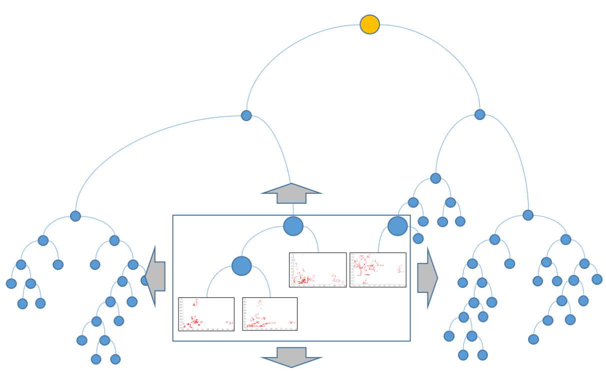

Figure 1: Concept: Exploration of large scatter plot spaces. The user eye-gaze is detected, leading to an expansion of the focused sub-tree (center rectangle). The remaining data is shown using a node-link representation (context, outside center rectangle).

When the user stops the *eye-tracking mode*, the application goes back to a state where it tracks only the user interest on each specific node, i.e. eye-gaze duration on each node (no pan control). We also show an overview of dendrogram areas visited so far (see Figure 2 for a gaze history view). This view allows to keep track of visited and unvisited areas, and constitutes input for further data analysis (see also Section 5).

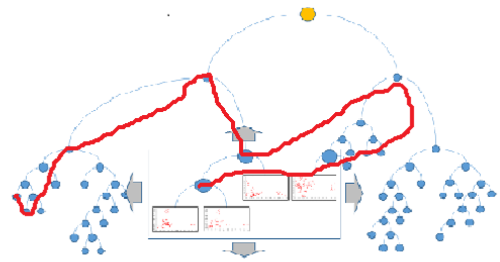

Figure 2: Gaze History Mock-up: It can be activated in the navigation panel. The user can track tree areas explored so far by an overlaid trace path (red line). It serves as a global map of explored/unexplored areas, and it is used for further analysis (see Section 5).

3.5 Benefits of Our Approach

We did informal, preliminary tests of our proposed navigation with 10 users. The feedback so far was positive, both to the semantic zoom mode and the gaze history view. The navigation was considered as rather smooth, and users can navigate without larger difficulties. Just by looking slightly away in the tree view, the corresponding movement is initiated in a very intuitive way. This facilitates the entire process of exploring the data in the tree, i.e., the view panning is synchronized with the field of view and the eye-movements of the user. When the user stops using the eye-based navigation, attention information starts to be collected again (eye-gaze duration on each area-of-interest) and it is the basis for the analysis of user interest detection and possible subsequent recommendation of interesting views. Note that in our concept we consider only gaze-based navigation. Of course, we can rely in addition on mouse/keyboard input to facilitate navigation, e.g., for labeling, saving views/bookmarking, etc.

4 Implementation and Application

To test our approach, we developed a proof-of-concept allowing the exploration of a large tree (dendrogram).

4.1 System Implementation

In our tree visualization, the leaf nodes are composed of visual representations of the data, i.e., scatter plots with pairs of data attributes. We use our own modified version of the JUNG system [OFWB03] for the tree visualization. We made changes on the adjustment of the lens size in the view space and positioning of the lens, now they are controlled by user eye-movements and updated in real-time. We created a customized tree layout to display color-coded nodes according to the computed similarity distance measure (darker color = lower similarity), and also the ability to display visual representations of the data on the leaf nodes, i.e, scatter plots.

The initial preset for DOI specification is the calculated similarity distance between scatter plot images. For this calculation, we make usage of a basic descriptor from the Java Image Processing Cookbook [2] that is based on the average calculation of 25 color triples for each image. After performing the comparison between each image using the descriptors, we create a distance matrix with the computed distances between all images. This matrix is handed out to an agglomerative hierarchical clustering algorithm [Beh16].

We tested our system with several hierarchical trees. Here, we illustrate the application of a hierarchical tree exploration of our data set. At the root and top subtrees we can find information about clusters of similar scatter plots, and at the leafs we find individual scatter plots. For this proof-of-concept we use a dendrogram comprising 269 scatter plots (leaf nodes) and 536 edges. Figure 3 shows a zoomed in view of the tree, the color-coded nodes according to distance similarity of the scatter plots, and the circular zooming lens (in gray color). In the navigation panel (top-right corner), we can get an overview of the entire tree size and respective available view space. The current view size (depending on the zoom level) and location is denoted by a white rectangle. The lens can be used to perform a close zoom into the scatter plot image, and it can be used to activate the display of a different visual representation of the data.

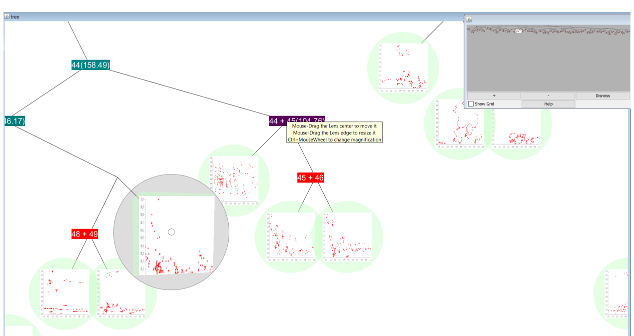

Figure 3: Zoomed-in view of the hierarchical tree. The mouse wheel can be used to: increase/decrease the lens magnifying ratio; increase/decrease the size of the circular zooming lens (gray circle) by clicking and dragging its border. A navigation panel (top-right corner) gives an overview of the actual position in the tree.

[2]Java Image Processing Cookbook (`http://goo.gl/FBXbjp`)

4.2 Application

Our data set is retrieved from the *Eurostat*[3] data repository, which provides a collection of data sets containing information on EU related topics (e.g., economy, population and industry). We use a preprocessed data set from preliminary work [SSB+15], which contains 27 statistical attributes from 28 EU countries showing temporal changes over time.

All navigation actions presented in the following examples illustrate a typical usage of our navigation system. Figure 4 shows an example of an ideal view over a small data portion, where the user is able to see the majority of the data. Practically, for larger data sets users will often have a more narrow view over the entire tree, depicted in Figure 5 with the 3 narrow views.

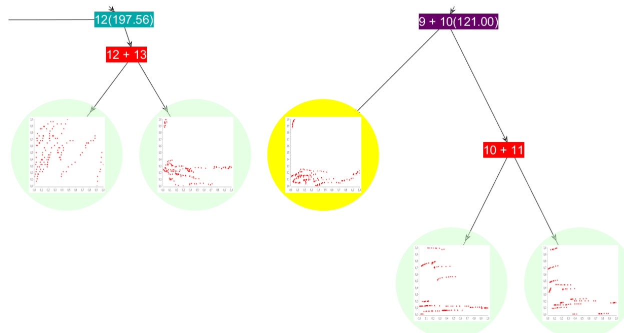

Figure 4: Ideal Case: Users can view several clusters of related scatter plots at once. Due to limitations of display space, this is often not possible, hence the need for adaptive visualization for navigation.

The navigation order (view sequence) followed by the user on these narrowed views can be random, it might just follow the similarity distance measures (depicted by the color-coded node rectangles). Figure 5, shows that the user first moved to view V1, where a group of interesting clustered scatter plots is visible (Figure 6). In view V1, the system detects a high gaze duration and infers that the interest is on one of the scatter plots (marked with "*"). After an in-depth inspection of this area, the user navigates to a view V2 (Figure 7) over another group of clustered scatter plots. The next most interesting and similar scatter plot (yellow color) is occluded in view V2 and it is only visible in view V3 (Figure 8).

It might take time until an interesting scatter plot (view V3) is spotted by the user. Also, there might exist other interesting scatter plots in another part of the tree, in a more far, and yet hidden location, e.g., Figure 9.

[3]Statistical Office of the European Union (`http://ec.europa.eu/eurostat`). Accessed 09/2016.

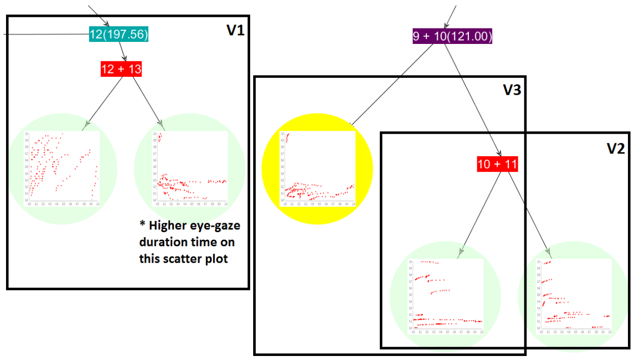

Figure 5: In practice, users may have limited views over a large space that must be explored. Therefore, the views (V1, V2, V3) might be limiting and not following an ideal sequence of exploration that would lead to finding interesting factors and to the creation of a useful mental model while exploring the data set. We take the eye-gaze duration time in account to infer about the interest of the user.

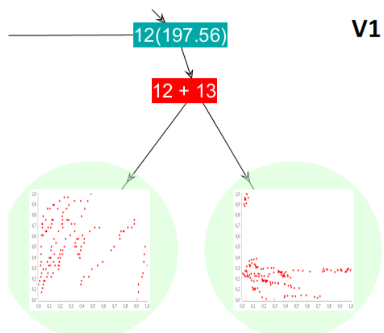

Figure 6: View 1: Realistic view of a first set of scatter plots. The user is focused on a set of scatter plots and unaware of other interesting locations of the tree.

In summary, our example merely demonstrates some of the challenges associated with the exploration of large graphs. The duration of the gazes can be used to expand or collapse sub-trees and hence provide a more organized (less cluttered) overview of the data, reducing the risk of getting lost in the exploration of large dendrogram trees. However, our measures computed directly from the location of the gazes are only a first step to control the views. We plan to collect data about the eye-movement scan paths and respective eye-gaze durations, as well as recurrences to develop more adaptive hierarchy views.

5 Discussion and Extensions

We implemented a proof-of-concept system for which we see numerous extension possibilities. First, our solution allows not only to adjust the amount of visual information presented, but also to capture longer sequences of visual exploration. The analysis of such captured data presents manifold opportunities to enhance the analysis process. For example, similar to previous work on navigation recommendations for exploring hierarchical graphs [GST13], approaches could be developed to suggest what new parts of the tree should be explored next by the user.

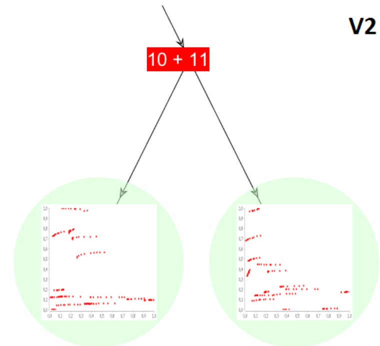

Figure 7: View 2: The user moves to a new location, but misses an interesting scatter plot that is occluded on a top location.

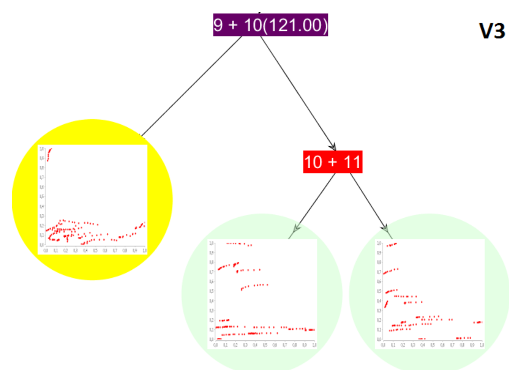

Figure 8: View 3: User navigates to this location and finds an interesting scatter plot related to view V1.

For now, our user focus model considers all elements of the tree (inner nodes and scatter plots). Given sufficient tracking resolution, we may apply the degree-of-interest concept also *locally* within a focused scatter plot. There are numerous ways to heuristically compute interest measures from eye-gaze fixations, fixation sequences, and gaze recurrences. Examples include learning relevance of local patterns, or deducing data groups of interest to a given user or analysis session. In the future, we hope to leverage such information and detect important aspects of the analysis problem at hand (e.g., whether there is exploratory or confirmatory analysis going on in a given session), or detect the level of user expertise. Depending on this infor-

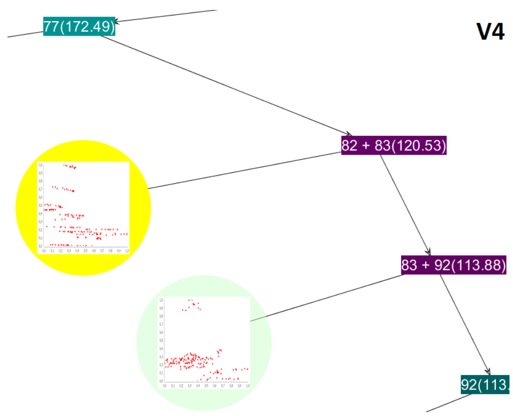

Figure 9: View 4: In a more far location (from V1, V2 and V3) the user notices that there is another interesting scatter plot worth investigation.

mation, the system may adapt its presentation and functionality accordingly. Also, a view recommender module may prevent the user from repetitively going to already examined areas in the display, suggesting instead, previously unseen parts of the view space, based on analysis of the eye-movement (scan) path. To that end, it is interesting to ask how one can do suggestions of other scatter plots to be explored, e.g., based on visual or data-driven similarities.

Another idea is to choose adaptively different visual representations on the nodes based on advanced interest measure computation. For example, in our current visualization approach (scatter plots) we might want to change dynamically between scatter plots, table representation, parallel coordinates visualization, or regression or clusters models computed for a given scatter plot. Also, an interest function should be adaptive also regarding time, e.g., during a longer or repeated analysis cycle. Such interest functions might take into account different objectives, e.g., the user might want to explore the most dissimilar clusters of scatter plots, or explore all scatter plots that have a certain shape. Depending on these objectives the interest function might need to be adjusted.

We also mention that our gaze path visualization could be enhanced to work as an overview tool to represent explored/unexplored regions. A gaze path might eventually serve as a visual history of a whole exploration process. Therefore, we may extend a given gaze path by annotating certain views visited (e.g., scatter plot thumbnails) at certain points in the gaze path (e.g., exceptionally long or short fixation times). Appropriate visual design might communicate a whole analysis session in a single image, which would be a valuable tool for reproducibility and communication of analysis sessions.

6 Conclusion

We presented a concept for visual exploration of hierarchically organized data, that relies on eye-tracking to steer the level of resolution shown. We assume that long gaze fixation times indicates user interest and hence can be used as a proxy to control the visual display of large data. Our effort extends previous work by a new user interface, allowing the navigation and the setting of the degree-of-interest to be determined by eye-movements, and it can be applied on both desktop screens or larger displays (e.g., using wearable eye-trackers). We applied this idea to the specific problem of comparing scatter plot diagrams, and hence support a type of meta-visualization: the elements in the tree are complex objects (visualizations). To this end we applied dendrogram computation based on image features, an approach which may help to overview large amounts of data views by grouping these for similarity. We have shown illustrative use cases for how eye-tracking can enhance a hierarchical data visualization, by mapping eye-gazes to degree-of-interest representations. Yet, our work is in an early stage and we see ample areas for future work. Future work includes high gaze-tracking precision on each node, refinement of interaction operations, view recommending, adaptive visual representations, and analysis provenance visualization. Finally, evaluation of our approach should be done in comparison to non-eye-tracking controls to qualitatively or quantitatively assess strengths and weaknesses of the approach.

References

[ABM07] Wolfgang Aigner, Alessio Bertone, and Silvia Miksch. *Tutorial: Introduction to Visual Analytics*, pages 453–456. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[Beh16] Lars Behnke. Implementation of an agglomerative hierarchical clustering algorithm in java., 09 2016.

[BKR+14] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, and T. Ertl. State-of-the-Art of Visualization for Eye Tracking Data. In R. Borgo, R. Maciejewski, and I. Viola, editors, *EuroVis - STARs*. The Eurographics Association, 2014.

[CAS01] Mon Chu Chen, John R. Anderson, and Myeong Ho Sohn. What can a mouse cursor tell us more?: Correlation of eye/mouse movements on web browsing. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '01, pages 281–282, New York, NY, USA, 2001. ACM.

[CdART12] Stéphan Clémençon, Héctor de Arazoza, Fabrice Rossi, and Viet-Chi Tran. Hierarchical clustering for graph visualization. *CoRR*, abs/1210.5693, 2012.

[CK15] Daniel Cernea and Andreas Kerren. A survey of technologies on the rise for emotion-enhanced interaction. *Journal of Visual Languages and Computing*, 31(Part A):70–86, 2015.

[Coo06] Lynne Cooke. Is the mouse a" poor man's eye tracker"? *Annual Conference-Society for Technical Communication*, 53:252, 2006.

[EOL14] R. Etemadpour, B. Olk, and L. Linsen. Eye-tracking investigation during visual analysis of projected multidimensional data with 2d scatterplots. In *Information Visualization Theory and Applications (IVAPP), 2014*, pages 233–246, Jan 2014.

[Fur86] G. W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '86, pages 16–23, New York, USA, 1986. ACM.

[GST13] Stefan Gladisch, Heidrun Schumann, and Christian Tominski. *Navigation Recommendations for Exploring Hierarchical Graphs*, pages 36–47. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[GW03] Joseph H. Goldberg and Anna M. Wichansky. Chapter 23 - eye tracking in usability evaluation: A practitioner's guide. In J. Hyn, R. Radach, and H. Deubel, editors, *The Mind's Eye*, pages 493 – 516. North-Holland, Amsterdam, 2003.

[Hay04] Mary M. Hayhoe. Advances in relating eye movements and cognition. *Infancy*, 6(2):267–274, 2004.

[HMM00] Ivan Herman, Guy Melançon, and M. Scott Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, January 2000.

[HNA+11] Kenneth Holmqvist, Marcus Nystrom, Richard Andersson, Richard Dewhurst, Halszka Jarodzka, and Joost van de Weijer. *Eye Tracking. A comprehensive guide to methods and measures.* Oxford University Press, 2011.

[JK03] Robert J.K. Jacob and Keith S. Karn. Commentary on section 4 - eye tracking in human-computer interaction and usability research: Ready to deliver the promises. In J. Hyn, R. Radach, and H. Deubel, editors, *The Mind's Eye*, pages 573 – 605. North-Holland, Amsterdam, 2003.

[KMSZ06] Daniel A. Keim, Florian Mansmann, Jorn Schneidewind, and Hartmut Ziegler. Challenges in visual data analysis. In *Proceedings of the Conference on Information Visualization*, IV '06, pages 9–16, Washington, DC, USA, 2006. IEEE Computer Society.

[LR96] John Lamping and Ramana Rao. Visualizing large trees using the hyperbolic browser. In *Conference Companion on Human Factors in Computing Systems*, CHI '96, pages 388–389, New York, NY, USA, 1996. ACM.

[MRS08] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.

[OFWB03] J. O'Madadhain, D. Fisher, S. White, and Y. Boey. The JUNG (Java Universal Network/Graph) Framework. Technical report, UCI-ICS, October 2003.

[PB05] Alex Poole and Linden J. Ball. Eye tracking in human-computer interaction and usability research: Current status and future. In *Prospects, Chapter in C. Ghaoui (Ed.): Encyclopedia of Human-Computer Interaction. Pennsylvania: Idea Group, Inc*, 2005.

[PGB02] Catherine Plaisant, Jesse Grosjean, and Benjamin B. Bederson. Spacetree: Supporting exploration in large node link tree, design evolution and empirical evaluation. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis'02)*, pages 57–, Washington, DC, USA, 2002. IEEE Computer Society.

[RFAS08] Kerry Rodden, Xin Fu, Anne Aula, and Ian Spiro. Eye-mouse coordination patterns on web search results pages. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '08, pages 2997–3002, New York, NY, USA, 2008. ACM.

[SSB+15] Lin Shao, Timo Schleicher, Michael Behrisch, Tobias Schreck, Ivan Sipiran, and Daniel A. Keim. Guiding the Exploration of Scatter Plot Data Using Motif-based Interest Measures. *IEEE Int. Symposium on Big Data Visual Analytics*, September 2015.

[SvLS12] Maximilian Scherer, Tatiana von Landesberger, and Tobias Schreck. A Benchmark for Content-Based Retrieval in Bivariate Data Collections. In *Proc. Int. Conference on Theory and Practice of Digital Libraries*, 2012.

[SZCJ16] Popelka Stanislav, Stachon Zdenek, Sasinka Cenek, and Dolezalova Jitka. Eyetribe tracker data accuracy evaluation and its interconnection with hypothesis software for cartographic purposes, 2016.

[vLKS+11] T. von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. van Wijk, J.-D. Fekete, and D. Fellner. Visual analysis of large graphs: State-of-the-art and future research challenges. *Wiley-Blackwell Computer Graphics Forum*, 30 (6):1719–1749, 2011.