

An Interactive 3D Integration of Parallel Coordinates and Star Glyphs

Elena Fanea*

Sheelagh Carpendale*

Tobias Isenberg*

Department of Computer Science
University of Calgary, Canada

ABSTRACT

Parallel Coordinates are a powerful method for visualizing multi-dimensional data but, when applied to large data sets, they become cluttered and difficult to read. Star Glyphs, on the other hand, can be used to display either the attributes of a data item or the values across all items for a single attribute. Star Glyphs may readily provide a quick impression; however, since the full data set will require multiple glyphs, overall readings are more difficult. We present Parallel Glyphs, an interactive integration of the visual representations of Parallel Coordinates and Star Glyphs that utilizes the advantages of both representations to offset the disadvantages they have separately. We discuss the role of uniform and stepped colour scales in the visual comparison of non-adjacent items and Star Glyphs. Parallel Glyphs provide capabilities for focus-in-context exploration using two types of lenses and interactions specific to the 3D space.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction Techniques; H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical User Interfaces, Interaction Styles.

Keywords: Parallel Glyphs, parallel coordinates, star glyphs, multi-dimensional data sets, 3D visualization.

1 INTRODUCTION

Parallel Coordinates are a powerful visual representation of multi-dimensional data due to their capacity to display the entire dataset while supporting interactions that lead to identification of relationships among items and among dimensions [6, 7]. Several researchers have been working on methods for improving aspects of their capability to support visual analysis of multi-dimensional data. The main motivation for developing a new visual representation is to address the problem of overlapping polylines. When two or more items in the data have the same values for a subset of the attributes, the polylines drawn to show each item may overlap, depending on the order in which the attributes are plotted. This artifact frequently occurs in Parallel Coordinates when even a very small number of data items are shown. In this case it is hard to read the visualization to determine the values of dimension for a specific data item. In addition, due to the nature of Parallel Coordinates the order in which data items are plotted is lost. In cases where this order is meaningful we lose the ability to compare the development of data values along one dimension or to compare it with another dimension. In contrast, evaluating trends in the order of data values for one dimension is easy with the Star Glyph visualization. However, because the Star Glyph visualization is spatially separate, direct comparisons between two glyphs are error prone.

In this paper we introduce Parallel Glyphs, a new visualization to overcome these problems. Our technique extends 2D Parallel Coordinates into the third dimension and naturally connects them

with Star Glyphs. We show that, in fact, Parallel Coordinates and Star Glyphs are orthogonal to each other and that they belong to a continuum of visualization techniques. We enrich this technique by applying colour scales to the 3D Parallel Glyphs to support comparison and selection tasks in 3D. In addition, we demonstrate that the new interaction methods that are possible in 3D such as a ring ruler enhances comparison of data values and can lead users to entirely new insights into the data.

The remainder of this paper is organized as follows. In the next section we review previous work that is related to our own technique. In the following section we introduce our concept of Parallel Glyphs and discuss the use of colour scales with the new visualization. Section 4 addresses new interaction techniques that are possible with Parallel Glyphs. Then, we present a short case study that illustrates the application of our technique to a specific data set. Finally, we give a conclusion and suggest directions for future work.

2 RELATED WORK

Parallel Coordinates [6, 7] and Star Glyphs [14, 15, 18] are both well recognized multi-dimensional visualizations. In fact, there continues to be considerable research extending them and improving the set of interactions possible with them [3, 12, 13, 18, 19, 25].

Advancements for interacting with Parallel Coordinates include brushing and multi-dimensional reordering to better support the interactive exploration of multi-dimensional data [2, 12, 13, 18, 19]. Brushing can be used to interactively select a range of values from a given dimension. This selected range is useful for filtering [25] and discovering answers for particular queries [1]. In addition, focus + context views can be provided by expanding or contracting the spaces between dimensions [25]. This can be important when the number of dimensions in the data set leads to a particularly dense display [27].

Trends in data can be revealed in Parallel Coordinates using angular brushing [5]. With this interaction technique the user specifies a subset of slopes to mark the corresponding data points as part of the current focus. The subset resulting from this brushing can be used for filtering the initial set. Thus, the user can choose not to display the items or dimensions that are relatively unimportant for a particular task [25].

A common interactive technique for reducing clutter and revealing data relationships in Parallel Coordinates is the manual or automatic reordering of the dimensions [13, 18]. While manual reordering may be more tedious than automatic reordering, it is often preferred because it allows the user to establish a superior reordering [25]. Users can take specific aspects of the data into consideration that might be overlooked by an algorithm that reorders based on similarity measures. Dimension reordering is also applicable to Star Glyphs as each spike in a glyph represents a dimension and adjacent spikes can be easier to compare [25].

Another problem with the Parallel Coordinates technique arises when it is used to visualize highly-dimensional data. This typically results in visual clutter since plotted lines overlap extensively, making the visualization difficult to read. One approach to the mitigation of this problem is to use hierarchical clustering, which repeatedly groups dimensions based on similarity measures. Thus,

*e-mail: {fanea | sheelagh | isenberg}@cpssc.ucalgary.ca

this representation provides a multi-resolution view of the data that reveals trends at different degrees of summarization. The same idea of cluster-based hierarchy can be applied for both dimensions [9, 25, 26] and data items [3, 12]. Another technique, polyline averaging, which represents a range of polylines with an averaged polyline, can be used for dynamically summarizing a set of polylines [16]. It reduces the clutter arising from overlapping polylines and is an alternative to more computationally demanding methods such as hierarchical clustering.

For hierarchically clustered representations XMDVTOOL [12, 18] implements n -dimensional brushing for interaction at a chosen level in the hierarchy. Angular brushing can also be done automatically or in relation to the hierarchy, if one is established [5].

Focus + context methods are general techniques developed to address problems of visibility of details and have been applied to Parallel Coordinates [25]. Multi-focus distortion operations allow the user to enlarge or shrink several regions in the same presentation [26]. The space between dimensions can be increased or decreased with a simple click. If, with high-dimensional data, the space between dimensions is too tight to allow an accurate click, a structure-based distortion called INTERRING can be used [27]. In an automatic application of distortion, the spacing parameter is the distance (for Parallel Coordinates) or angle (for Star Glyphs) [25]. Although the spacing between dimensions in Parallel Coordinates and Star Glyphs is typically uniform, non-uniform variations therein can convey additional information about the dimensions [25].

The various types of interactions possible with Parallel Coordinates, Star Glyphs, Scatter Plots, and Dimensional Stacking have been compared by several authors with respect to each of these four visualization techniques [3, 12, 13, 18, 19, 25].

Several systems use Parallel Coordinates along with other visual representations. For example, Wong et al. [23] developed a 2D visualization system that contains Parallel Coordinates and Scatter Plots, each presented in their own regions of the window. The representations are linked so that interactive changes to one are reflected in the other. The two techniques combined better support data exploration.

Tominski et al. [17] suggest organizing Star Glyphs on time based axis. This would provide a linked 3D layout for Star Glyphs with one glyph per data item but was left as a suggestion for future work. Kiviat diagrams [4] are 2D visualizations that resemble Star Glyphs in that radial spikes emerge from a common centre and each spike corresponds to a data item. Kiviat tubes are obtained by displaying Kiviat diagrams along a time axis and rendering the surrounding surface as suggested by Tominski et al. [17]. The surface that connects all the Kiviat diagrams perhaps emphasizes the shape of the tube and possibly conceals information about individual data items. These ideas are similar to our Parallel Glyphs in that they link Star Glyphs but there is no integration with Parallel Coordinates.

Finally, Wernert et al. [21] make an exploration of alternate multi-dimensional data layouts but these representations are not linked either visually, structurally, or interactively.

3 PARALLEL GLYPHS: INTEGRATING PARALLEL COORDINATES WITH GLYPHS

Our technique integrates Parallel Coordinates with Star Glyphs to create Parallel Glyphs. Parallel Coordinates are a 2D visual representation for multi-dimensional data (see Figure 2(a)) that uses a set of vertical axes arranged left to right for plotting the values of data items where each axis corresponds to one dimension. For each data item, a polyline connects the points on each dimensional axis that correspond to the item's value for that dimension.

Star Glyphs are also a 2D visual representation for multi-dimensional data (see Figure 1). Each glyph represents either a data

item or a data dimension. Spikes, frequently bounded by a polygon, emanate radially from a central point. When a glyph represents a data item, its spikes vary in length according to the values the item has in each dimension. When, on the other hand, a glyph represents one dimension, the spikes vary in length according to the value each item has for that particular dimension.



Figure 1: Examples for Star Glyphs.

3.1 Adding the Third Dimension

One significant motivation for developing Parallel Glyphs is to address the problem of overlapping polylines. This artifact frequently occurs in Parallel Coordinates even when only a small number of data items are shown. For example, Figure 2(a) shows a dataset that contains only few data items. Looking at the right side of the Parallel Coordinate display in Figure 2(a) one can see four polylines, indicating that there are four data items. However, on the left side only three polylines are visible. The 2D visualization of Parallel Coordinates can be ambiguous. For example, Polyline A is obscured by either Polyline B or C between dimensions 1 and 3. This fact can affect even simple tasks such as revealing trends. By opening the Parallel Coordinates into 3D (see Figure 2(b)) we are able to tell not only that Polyline A follows the path of Polyline B but also that there is in fact another polyline (D) following a different path that was entirely hidden by the other polylines.




Figure 2: Advantage of unfolding Parallel Coordinates into 3D.

As this test dataset illustrates, the problem of ambiguity created by overlapped polylines is likely to appear in real data, especially when the number of data items increases.

3.2 Construction of Parallel Coordinates in 3D

The Parallel Glyphs integrates Parallel Coordinates with Star Glyphs by extending Parallel Coordinates into 3D space and unfolding them around a pivot axis (see Figure 3). In Parallel Coordinates data items are plotted as polylines over a set of dimension axes.

In order to be able to unfold the visualization into 3D we assign each polyline representing a data item its own set of dimension axes and utilize the baseline that connects the bottom of the dimension axes as our pivot axis. Thus, while our 2D view of Parallel Glyphs matches the original Parallel Coordinate view, without any further transformations we can now successively apply rotations around the pivot axis. This results in a visualization in which the data items can be interactively unfolded into three dimensions. We can stop or pause this opening at any chosen angle to change the degree of unfolding, ranging from the original 2D view to a complete 360 degree rotation.

In the new visualization, Star Glyphs emerge during the unfolding process. By connecting the outer end points of all line segments for a given dimension the glyphs appear linked along the pivot axis (see Figure 4). Thus, our 3D visualization naturally integrates both types of visualizations for multidimensional data sets. Both Parallel Coordinates and Star Glyphs can be shown connected to each other and can be explored interactively.

When building the 3D visualization we considered both orthographic and perspective projection. The perspective projection is typically used in applications that convey the sense of depth in a more realistic way. With this projection, the length of the glyphs' spikes varies, closer segments being shown longer than the farther ones. These variations make it difficult for users to accurately visually compare the lengths of the spikes. Therefore, we decided to use orthographic projection in our visual representation thus preserving the lengths of the segments.

Traditionally, Parallel Coordinates have one axis for each dimension of the data set. Consequently, the Star Glyphs emerging in our visualization have one glyph per dimension, each spike corresponding to one data item. This is not the conventional way to map the data in Star Glyphs: most related literature shows each glyph representing all dimensions of a single item. In fact, our visualization can generate this representation as well (see Figure 5) by merely switching the way the data is read from the data table: changing from columns over rows to rows over columns or vice versa. This also results in a modified Parallel Coordinates representation that shows dimensions as polylines and objects as axes (see Figure 6).

A series of questions are raised by this switch. In many situations there are more items than dimensions which would lead to an increased number of glyphs aligned along the pivot line (as seen in Figure 5). This will decrease space between them, raising the question of the maximum reasonable size of data sets. However, by adding some 3D interaction techniques to the new visualization it may be possible to increase the number of glyphs supported and enhance flexibility. For example, the pivot line that connects the glyphs is just an abstract element and not a preferred dimension or a time axis. Therefore, even if the order of glyphs is initially the same as the order of items in the data table, this can be modified to better support each task.

A disadvantage of a traditional 2D Star Glyph representation is that the loss of context can lead to difficulties when conducting comparisons across glyphs. New interactions with the glyphs on the pivot axis can address these issues. The pivot axis naturally aligns the glyphs (see Figure 4), however, rotating the pivot axis can provide either an end view for direct comparison (see Figure 7(a)) or it can be rotated so that the data item of interest is displayed top-most (see Figure 7(b)). The orthographic projection makes comparisons valid and maintaining the polyline connections provided by the Parallel Coordinates can also aid comparison (see Figure 8). Additional 3D interactions are described in Section 4.

3.3 Interacting with Colour for Readability

In Parallel Coordinate visualizations the polylines representing the data items are typically drawn with a uniform colour or with




Figure 3: Several steps of unfolding the 2D Parallel Coordinates into the 3D Parallel Glyphs visualization.




Figure 4: Star Glyphs emerging in the new visualization.




Figure 5: Star Glyphs, each glyph representing one object as opposed to one dimension as shown, e. g., in Figure 4.




Figure 6: Parallel Coordinate representation resulting from Figure 5.




Figure 7: Star Glyphs aligned to coordinate axes to allow comparison of data items.




Figure 8: Setting glyphs into relation to each other by arranging them on the pivot axis and by connecting them with polylines.

proximity-based shades that convey information about the hierarchical structure of the data [3]. For glyphs in general colour is an important attribute that is used to represent data [8, 24] in addition to the size, position, orientation. However, for Star Glyphs in particular the most common way of mapping data attributes is not through the colour but through the spikes' lengths, which represent the values of the corresponding data item. Thus, a Star Glyph's shape is its main visual cue. Unfortunately, comparing spikes becomes increasingly difficult with growing distance between them, especially if one wants to compare spikes across more than one glyph. In 2D, the glyphs' spikes are drawn radially and it can be difficult to accurately compare the lengths of spikes set at wide angles apart from each other. 3D rotations can create further difficulties. Orthographic projection does make these comparisons valid for views that are orthogonal to any two of the x , y , and z -dimensions.

We address this problem by adding circular textures to the glyphs as additional visual cue. We create these textures by adjusting uniform colour scales [10] (see Figure 9) to offer the best readability for glyphs, according to the following three criteria:

- provide a means for comparing pixels situated at a distance,
- follow some natural ordering of colours, and
- be readable by colour-blind people.

While uniform colour scales are discussed extensively in literature [10], these scales are typically based on algorithms that attempt to maximize the number of just noticeable differences (JND) between neighbouring pixels. This conflicts with the first criterion above. Instead of JND, we are interested in recognizable differences across a distance.

The approach we take extends the notion of a uniform colour scale. We use a set of stepped scales with 16, 10 and 8 steps based on the numerical values provided by [11] (see Figure 10). Each sector of a stepped colour scale has a single colour that is obtained by averaging the amount of red, green, and blue of corresponding pixels from the uniform colour scale. This approach allows the colour scale to meet the second criterion above (the natural sense of order induced by a uniform scale) while adding the ability to easily notice ranges of values (first criterion).

According to [10] the heated and magenta scales are recommended because the human eye can better notice variations in luminance along the yellow-orange axis and variations in hue along the magenta axis. Knowing that users have different preferences




Figure 10: Stepped heated colour scales applied to the glyphs to increase comparability between glyphs using 8, 10, and 16 steps, respectively.




Figure 9: Uniform heated colour scale applied to the glyphs using texture mapping.

when choosing the colour scale we provide several other options as well. We also considered colour scales based on opposite colours with the two opposite colours on each end of an axis, as recommended in [20] but we chose not to use them for two important reasons. First, the variation on the red-green channel would be difficult for colour-blind people (contradicting the third criterion above). Second, the variations on the blue-yellow channel are hard to distinguish on very thin sectors [20] as is likely to happen in glyphs with many spikes.

We complement the support offered by a uniform colour scale with a ring, centred about the pivot point for a glyph, that can be interactively sized (see Figure 11). When the user clicks a point inside a glyph, this ring is drawn on each glyph as a reference for comparison. The rings have a radius equal to the distance between the reference point and the centre of its corresponding glyph. These rings act like a ruler and were created to assist in fine comparisons.

The colour scales can also be applied to the polylines, having one main dimension selected as a criterion for choosing the colour (see Figure 12). This is especially useful when using the 2D Parallel Coordinates visualization to observe, for example, data trends. However, when using a stepped colour scale based on a uniform one, the steps may not be sufficiently distinguishable for other tasks such as selection.

To mitigate this problem we provide a set of stepped colour




Figure 11: Using rings as rulers to compare values between glyphs. In this example it is now possible to tell that the big top spike on the middle glyph is slightly shorter than the corresponding spike on the glyph to its left and that it represents the second biggest value for this dimension.

scales based on a stepped rainbow scale (see Figure 13) that has a distinct colour for each step. With this colour scale it is easy to identify and select polylines having the values corresponding to the main dimension in the same range. This is valid even if the polylines are not adjacent or if the glyphs are folded up or only slightly rotated. Of course, the rainbow scale will not meet the second criterion above (natural order of colours) and so a legend for the colour scale as an aid to the visualization reader is recommended.

4 INTERACTION

Data exploration using Parallel Glyphs is enhanced by several interactions, some of which are specific to either the 2D or the 3D visualization. We extend the common techniques possible with 2D Parallel Coordinates such as brushing the polylines and dimensions, filtering tuples or dimensions, and reordering dimensions with new interactions that originate from the 3D nature of our visualization. It is now possible to unfold the Parallel Glyphs using various angles as described in Section 3.2 and apply 3D manipulations to the resulting visualization. For example, we are able to rotate the visualization in all three directions around the pivot axis. A particularly




Figure 12: Brush-selecting a range of polylines with the help of a stepped colour scale applied to the lines based on a selected dimension. The colours have been assigned according to the middle dimension. Note that even with stepping applied to the colour scale it is hard to make out differences.




Figure 13: Using a stepped rainbow scale helps to increase the contrast between adjacent colour ranges to better support brushing. In addition, trends are more visible for the entire dataset compared to Figure 12.

interesting situation is when the glyphs are unfolded and the entire model is rotated such that the glyphs directly overlap each other. Translucent rendering of glyph interiors supports direct comparison of all glyphs (see Figure 7(a)). In addition to these interactions, we implemented several new techniques described in the following.

4.1 Tuple Rearrangement

Having transformed Parallel Coordinates to three dimensions, the order in which tuples are arranged now also has a meaning in addition to the order of dimensions. Thus, we provide the possibility to interactively change this order. The user can select one (through clicking) or several (through brushing) polyline(s) and then to rotate them around the pivot axis to a different position in the sequence. This allows one to adapt the sequence of polylines in the data set to directly compare polylines that were previously far apart. Thus, the user can avoid one of the main issues of a Star Glyphs visualization, the difficult comparison of non-adjacent spikes. Although in Parallel Coordinates the values of different data items corresponding to

the same dimension are plotted along the same axis, their comparison is not necessarily easy because they might be overlapped by other polylines. Parallel Glyphs offer the best of both worlds, supporting the interactive adjustment of overlapping polylines and the possibility of placing items adjacently for a given task.

Besides facilitating the comparison of data items, the polyline reordering can lead to identification of patterns or trends in the data set, as dimension reordering has been already shown to be useful for this type of task. Our technique provides interactive item reordering as the user might have knowledge about the data that would elude an algorithm. However, an automatic operation equivalent to the automatic dimension reordering would be interesting to develop.

The fact that the angle between two consecutive spikes of the glyphs is no longer constant opens up new paths for data exploration. In Parallel Coordinates, in order to analyze a subset of the data independently from the rest of the set, the items not important for the task are either made invisible or faded out. In our technique, the items involved in a task can be maintained in the top side of the screen while the remaining subset can be rotated to the bottom. Thus, the entire set is visible and available for selection at any moment, giving the user full overview of the data and making it easy to switch between items from one category of interest to another. Both the top and bottom subsets of the data items benefit independently from different angles of unfolding, while the entire model can be rotated in 3D, in order to find the best view (see Figure 14).




Figure 14: Rearrangement of selected data items in 3D for comparison without losing context information.

4.2 Lens Techniques

Depending on the data set, the visual density of the representation can make it extremely difficult to perform visual selection of polylines that satisfy a given criterion or to identify overlapped data items. We address this problem by integrating focus + context capabilities using two types of lenses.

The first one, EDGELENS [22] facilitates the exploration of the data set in Parallel Coordinates by adjusting the segments that fall into its area of influence into Beziér curves (see Figure 15(a)). These curves are shuffled at gradually decreasing distances from the centre of the lens, thus expanding the space available between polylines and providing more insight into the dataset. The centre

of the lens can be easily moved until the desired presentation is obtained. The desired polylines can then be brush selected.

For situations where the desired data items have a particular value or range of values for a specific dimension, a second type of lens may be used. For ease of selection purposes the EPSLENS produces a non-linear expansion of the axis being in its immediate vicinity (see Figure 15(b)). Thus, the points plotted along this axis are separated by increased distances, allowing the identification and selection of polylines required by a particular visual task.




Figure 15: Lens interaction with Parallel Coordinates.

5 CASE STUDY

As a concrete example, we apply Parallel Glyphs to a set of 100 generations of plants generated by a genetic algorithm, each having five attributes: number of flowers, the dimensions of the plant in 3D (width, depth, and height), and the fitness factor. These are the five dimensions used in the visual representation, with each generation represented by one polyline. In this section we look at how the integration of Parallel Coordinates and Star Glyphs along with the available colour scales support the user in solving a series of tasks.

The first task is to compare the variation of flowers and fitnesses along the 100 generations. In Parallel Coordinates it is straightforward to compare the range of values plotted along the first and last axis. However, this representation does not provide information about how these values vary from one generation to another. This task is easily solved with our technique by unfolding the glyphs and by bringing the two desired glyphs next to each other (see Figure 16). In combination with the interactive ring ruler and with rotations in any of the three dimensions, the variation between these two attributes is easily compared, revealing a dependence relationship between the number of flowers and the fitness factor (the two Star Glyphs in front in Figure 16).

In the second task, one wants to understand how the shape of plant changes over a selected number of generations for the plants that have a fitness factor of between 70%–80%. First, to visually select the desired items, the user applies the rainbow scale with 10 steps to the whole set of polylines, using the values of last dimension as a criterion for choosing the colour. Thus, each distinct colour of the colour scale corresponds to a range of 10% of the maximum fitness. Next, polylines that have a colour corresponding to the third sector of the colour scale (green) must be selected, but because the lines are too close to each other, accurate selection is very difficult. Using the EPSLENS the user expands the space between those lines and then brushes only the green ones (see Figure 17(a)). The unselected polylines can be either made invisible (as in the Figure) or rotated to the bottom of the screen to keep them available for further tasks. By choosing to keep only the selected data items the user is free to completely unfold the glyphs formed with the remaining items (see Figure 17(b)). The glyphs corresponding to width and height (second and third from left) have similar shapes but different sizes. This means the variation of these two attributes is similar. However, the glyph corresponding to depth (fourth from




Figure 16: Task 1: Comparison of a sequence of data values using repositioned Star Glyphs.

left) has peaks where the other two have local minima. This means that in generations where the width was small the height was also small but the depth was inverted. This display suggests that the plants in each generation varied their dimensions to try to maintain a constant fitness and volume.




Figure 17: Task 2: examination of a selected subset of the data.

6 CONCLUSION AND FUTURE WORK

In this paper we have introduced Parallel Glyphs which are a three-dimensional extension of Parallel Coordinates integrated with Star Glyphs. Our work demonstrates that Parallel Coordinates and Star Glyphs are not only closely related and in the same family of visualizations but that they are, in fact, orthogonal to each other. We showed that the traditional view of Parallel Coordinates—objects over dimensions—leads to one of two versions of Star Glyphs—one glyph representing all values for one dimension. The other version of Star Glyphs—one glyph representing the values of one data object—can be derived from a second version of Parallel Coordinates in which dimensions are plotted over objects.

We have enriched the new 3D visualization by applying colour scales as textures to the 3D Star Glyphs to support the task of comparing spatially separated glyphs. We discussed both uniform and stepped colour scales and compared their merits. We also showed

how these colour scales can lead to colouring the data items' polylines in order to identify trends and support brush-selection. To further support visualization tasks we discussed various interaction possibilities. In addition to well-known techniques for 2D Parallel Coordinates, the unfolding of 2D Parallel Coordinates into the new 3D representation, and 3D transformations such as rotation around coordinate axes we introduced a ring ruler interaction that enables determining even small differences between Star Glyphs. Furthermore, we discussed the newly possible data item rearrangement that enables users to directly compare selected data items with each other without losing the context of the rest of the data set. Finally, we discussed enhancements to interactions that are applied to the traditional 2D Parallel Coordinates visualization such as lens techniques for focus + context data exploration.

In the future we will examine the limits for the supported size of data set for the technique to be usable. In addition, we envision adding shading and other 3D rendering techniques to improve the visualization technique. Also, we will explore 3D shape analysis methods to extract more information from the data set and research automatic reordering of data items.

ACKNOWLEDGEMENTS

We gratefully thank our funding provider Natural Sciences and Engineering Research Council of Canada (NSERC). We would also like to thank researchers from the Interactions Lab at the University of Calgary for their insightful comments on this work and, in particular, Michael Boyle for his help in editing this paper.

REFERENCES

- [1] Almir Olivette Artero, Maria Cristina Ferreira de Oliveira, and Haim Levkowitz. Uncovering Clusters in Crowded Parallel Coordinates Visualizations. In *Proceedings of the IEEE Symposium on Information Visualization 2004 (InfoVis 2004)*, pages 81–88, Los Alamitos, CA, 2004. IEEE Computer Society Press.
- [2] Hong Chen. Compound Brushing. In *Proceedings of the IEEE Symposium on Information Visualization 2003 (InfoVis 2003)*, pages 181–188, Los Alamitos, CA, 2003. IEEE Computer Society Press.
- [3] Ying-Huey Fua, Matthew O. Ward, and Elke A. Rundensteiner. Hierarchical Parallel Coordinates for Exploration of Large Datasets. In *Proceedings of IEEE Visualization 1999 (VIS 1999)*, pages 43–50, Los Alamitos, CA, 1999. IEEE Computer Society Press.
- [4] Steven T. Hackstadt and Allen D. Malony. Visualizing Parallel Programs and Performance. *IEEE Computer Graphics and Applications*, 15(4):12–14, July 1995.
- [5] Helwig Hauser, Florian Ledermann, and Helmut Doleisch. Angular Brushing of Extended Parallel Coordinates. In *Proceedings of the IEEE Symposium on Information Visualization 2002 (InfoVis 2002)*, pages 127–130, Los Alamitos, CA, 2002. IEEE Computer Society Press.
- [6] Alfred Inselberg and Bernard Dimsdale. Parallel Coordinates: A Tool for Visualizing Multi-Dimensional Geometry. In *Proceedings of IEEE Visualization 1990 (VIS 1990)*, pages 361–378, Los Alamitos, CA, 1990. IEEE Computer Society Press.
- [7] Alfred Inselberg. The Plane with Parallel Coordinates. *The Visual Computer*, 1(2):69–91, October 1985.
- [8] Daniel A. Keim. Pixel-Oriented Database Visualizations. *ACM SIGMOD Record*, 25(4):35–39, December 1996.
- [9] Daniel A. Keim. Information Visualization and Visual Data Mining. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):1–8, January 2002.
- [10] Haim Levkowitz and Gabor T. Herman. Color Scales for Image Data. *IEEE Computer Graphics and Applications*, 12(1):72–80, January 1992.
- [11] Haim Levkowitz. ColorCenterColormaps. Web site at <http://www.cs.uml.edu/~haim/ColorCenter/ColorCenterColormaps.htm>. visited April 19, 2005.
- [12] Allen R. Martin and Matthew O. Ward. High Dimensional Brushing for Interactive Exploration of Multivariate Data. In *Proceedings of IEEE Visualization 1995 (VIS 1995)*, pages 271–278, Los Alamitos, CA, 1995. IEEE Computer Society Press.
- [13] Wei Peng, Matthew O. Ward, and Elke A. Rundensteiner. Clutter Reduction in Multi-Dimensional Data Visualization Using Dimension Reordering. In *Proceedings of the IEEE Symposium on Information Visualization 2004 (InfoVis 2004)*, pages 89–96, Los Alamitos, CA, 2004. IEEE Computer Society Press.
- [14] William Ribarsky, Eric Ayers, John Eble, and Sougata Mukherjee. Glyphmaker: Creating Customized Visualizations of Complex Data. *Computer*, 27(7):57–64, July 1994.
- [15] Christopher D. Shaw, James A. Hall, Christine Blahut, David S. Erbert, and D. Aaron Roberts. Using Shape to Visualize Multivariate Data. In *Proceedings of the 1999 Workshop on New Paradigms in Information Visualization and Manipulation*, pages 17–20, New York, 1999. ACM Press.
- [16] Harri Siirtola. Direct Manipulation of Parallel Coordinates. In *Proceedings of IEEE Visualization 2000 (VIS 2000)*, pages 373–378, Los Alamitos, CA, 2000. IEEE Computer Society Press.
- [17] Christian Tominski, James Abello, and Heidrun Schumann. Axes-Based Visualizations with Radial Layouts. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, pages 1242–1247, New York, 2004. ACM Press.
- [18] Matthew O. Ward. XmdvTool: Integrating Multiple Methods for Visualizing Multivariate Data. In *Proceedings of IEEE Visualization 1994 (VIS 1994)*, pages 326–333, Los Alamitos, CA, 1994. IEEE Computer Society Press.
- [19] Matthew O. Ward. Creating and Manipulating N-Dimensional Brushes. In *Proceedings of Joint Statistical Meeting 1997*, pages 6–14, 1997.
- [20] Colin Ware. *Information Visualization*. Morgan Kaufmann Publishers, San Francisco, 2000.
- [21] Eric A. Wernert, Donald K. Berry, John N. Huffman, and Craig A. Stewart. Tree3d-A System for Temporal and Comparative Analysis of Phylogenetic Trees. In *Proceedings of the IEEE InfoVis 2003 Poster Compendium*, pages 114–115, Los Alamitos, CA, 2003. IEEE Computer Society Press.
- [22] Nelson Wong, Sheelagh Carpendale, and Saul Greenberg. EdgeLens: An Interactive Method for Managing Edge Congestion in Graphs. In *Proceedings of the IEEE Symposium on Information Visualization 2003 (InfoVis 2003)*, pages 51–58, Los Alamitos, CA, 2003. IEEE Computer Society Press.
- [23] Pak Chung Wong and R. Daniel Bergeron. Multivariate Visualization Using Metric Scaling. In *Proceedings of IEEE Visualization 1997 (VIS 1997)*, pages 111–118, Los Alamitos, CA, 1997. IEEE Computer Society Press.
- [24] Jing Yang, Anilkumar Patro, Shiping Huang, Nishant Mehta, Matthew O. Ward, and Elke A. Rundensteiner. Value and Relation Display for Interactive Exploration of High Dimensional Datasets. In *Proceedings of the IEEE Symposium on Information Visualization 2004 (InfoVis 2004)*, pages 73–80, Los Alamitos, CA, 2004. IEEE Computer Society Press.
- [25] Jing Yang, Wei Peng, Matthew O. Ward, and Elke A. Rundensteiner. Interactive Hierarchical Dimension Ordering, Spacing and Filtering for Exploration of High Dimensional Datasets. In *Proceedings of the IEEE Symposium on Information Visualization 2003 (InfoVis 2003)*, pages 105–112, Los Alamitos, CA, 2003. IEEE Computer Society Press.
- [26] Jing Yang, Matthew O. Ward, Elke A. Rundensteiner, and S. Huang. Visual Hierarchical Dimension Reduction for Exploration of High Dimensional Datasets. In *Proceedings of the ACM Symposium on Data Visualisation 2003*, pages 19–28, 282, New York, 2003. ACM Press.
- [27] Jing Yang, Matthew O. Ward, and Elke A. Rundensteiner. InterRing: An Interactive Tool for Visually Navigating and Manipulating Hierarchical Structures. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis 2002)*, pages 77–84, Los Alamitos, CA, 2002. IEEE Computer Society Press.