Task-Based Evaluation of Multi-Relational 3D and Standard 2D Parallel Coordinates

Camilla Forsell[a] and Jimmy Johansson[b]

[a]Department of Information Science, Uppsala University, Uppsala, Sweden
[b] Norrköping Visualization and Interaction Studio, Linköping University, Norrköping, Sweden

ABSTRACT

Multivariate data sets exist in a wide variety of fields and parallel coordinates visualizations are commonly used for analysing such data. This paper presents a usability evaluation where we compare three types of parallel coordinates visualization for exploratory analysis of multivariate data. We use a standard parallel coordinates display with manual permutation of axes, a standard parallel coordinates display with automatic permutation of axes, and a multi-relational 3D parallel coordinates display with manual permutation of axes. We investigate whether a 3D layout showing more relations simultaneously, but distorted by perspective effects, is advantageous when compared with a standard 2D layout. The evaluation is accomplished by means of an experiment comparing performance differences for a class of task known to be well-supported by parallel coordinates. Two levels of difficulty of the task are used and both require the user to find relationships between variables in a multivariate data set. Our results show that for the manual exploration of a complex interrelated multivariate data set, the user performance with multi-relational 3D parallel coordinates is significantly faster. In simpler tasks, however, the difference is negligible. The study adds to the body of work examining the utility of 3D representations and what properties of structure in 3D space can be successfully used in 3D representations of multivariate data.

Keywords: Evaluation, parallel coordinates, multi-relational 3D parallel coordinates, multivariate data

1. INTRODUCTION

Information visualization techniques are used to facilitate decision making within many domains. Often the information represented is a large data set having a complex and multivariate nature which requires the user to perform extensive analyses. To aid the exploration process a multitude of tools and methods are offered for use but it seems reasonable to presume that all techniques are not equally suited to all tasks. Empirical comparisons could provide valuable insight into whether, when and why each visualization can provide effective cognitive support. Evaluation also helps us choose promising design ideas for further study.[1,2]

This paper presents an experimental study conducted with the objective of evaluating different parallel coordinates visualizations. We compare standard 2D parallel coordinates[3] and 3D multi-relational parallel coordinates.[4] The study reported here investigates whether the multi-relational 3D parallel coordinates technique is useful for a specific task. Our objective is to identify classes of user task for which each representation of data is best suited, specifically how different axis arrangements influence efficiency. The amount of information represented in one view of the data is crucial, as well as the angle of the view onto the 3D representation.

Parallel coordinates visualizations mainly facilitate two types of examination of a data set: comparisons between data items and comparisons between variables. Our study addresses their support of tasks in which users need to be able to accurately and effectively explore a multivariate data set in order to find relationships between variables. In the case of 3D visualization the possibility of misleading distortions caused by the perspective view has to be taken into consideration. However, our perceptual representations are based on a class of qualitative properties of structure in 3D space, which are easily perceived.[5–7] Since the judgements required by users in this study are based on easily perceived qualitative properties, the 3D multi-relational parallel coordinates technique should not be affected by any distortions of data caused by perspective effects. Theory and earlier findings suggest it should be superior to a 2D layout of data.

Further author information: (Send correspondence to Camilla Forsell)
Camilla Forsell.: E-mail: camilla.forsell@dis.uu.se
Jimmy Johansson.: E-mail: jimjo@itn.liu.se

When examining relationships between variables, the task either involves determining the strength of a predetermined relationship (for example, the amount of correlation) or determining the nature of the relationship. In the present work we focus on parallel coordinates as an aid for a decision maker in determining the nature of a relationship. This is a task that logically precedes determining the strength of a relationship and there are domains where, for example, it would be beneficial to determine if a relation in a data set is mainly linear or not, has a discontinuity or not, or has cyclical components or not.

2. BACKGROUND AND RELATED WORK

The parallel coordinates technique is a commonly used tool for visualizing multivariate data by means of a 2D representation with parallel axes. Recent work has evaluated its efficiency for different tasks in comparison with other multidimensional techniques.[8,9] In an effort to 'link tools to task', Andrienko and Andrienko[10] produced a taxonomy of tasks. One can discern two classes of task claimed to be effectively supported by parallel coordinates visualizations: either examining properties and relationships of objects represented by the set of polylines connecting the axes, or examining the relationships between variables represented by the axes. Some examples of relationships between variables (used in the present study) are illustrated in figure 1.

Despite the popularity of parallel coordinates, the technique is not without limitations. One limitation concerns the number of data items that it is possible to simultaneously display. Visualization of a medium or large data set, meaning one containing thousands of data items and above, often results in visual clutter since plotted lines between axes overlap, making the visualization difficult to interpret. This has, for example, been addressed by Wegman and Luo,[11] Fua et al.[12] and Johansson et al.[13] Several efforts have also been made to extend parallel coordinates into a 3D display. Wegenkittl et al.[14] introduced extruded parallel coordinates and 3D parallel coordinates to visualize dynamical systems. A similar 3D parallel coordinates view, 'the cube', was presented by Falkman[15] for visualizing clinical data. Combining parallel coordinates with star glyphs using a 3D display was recently presented by Fanea et al.[16]

The order in which the axes in the parallel coordinates display are arranged clearly influences the insight gained from one particular view of the visualization. The 2D parallel coordinates configuration only permits analysis of relationships between adjacent axes which allows simultaneous recognition of trends and patterns between different pairs of adjacent axes in the same view. However, it can never permit a simultaneous analysis of all possible combinations between one single axis and all the others in a data set. Manually repositioning the axes facilitates the data analysis but is often time consuming, especially if the data set contains a large number of variables. Wegman[17] showed that a minimal set of orderings of the axes, so that every possible adjacency in the visualization is presented, takes $\frac{N+N \bmod 2}{2}$ permutations. This provides a quick overview of the data and all possible combinations of variables but, since several axes change position simultaneously at each iteration, it is difficult to keep track of individual axes. Another approach regarding repositioning of axes has been developed by Peng et al.[18] who propose using a measure of visual clutter to order the axes so that more information can be gained from one single parallel coordinates view. However, it does not resolve the problem of only being able to examine relationships between adjacent axes. Finally, regardless of using a manual or automatic permutation of axes, data analysis requires a large amount of visual scanning and users need to remember information between different views. Obviously, with a large multivariate data set including many variables, the cognitive load can become very large.

One way to overcome the axes order limitation of the parallel coordinates display is to change the axis configuration. Instead of arranging the axes equidistantly, a circular (multi-relational) arrangement can be used to allow one variable to be simultaneously compared with all other variables in the data set. This is achieved by positioning one axis, the focus axis, at the centre of a circle and arranging the other axes so that they are positioned, equally spaced, on the circumference of the circle. This can be done using a 2D display[19] or a 3D display,[4,20] see figure 2 for an illustration of the 3D display.

These multi-relational parallel coordinates techniques have the advantage that the $N-1$ relationships between the focus variable and all other variables can immediately be investigated. Compared to the standard 2D parallel coordinates display these configurations reduce the interaction needed in the data analysis. In general, within N views all variables can be put into the centre and all possible pairwise relationships can easily be examined. In fact all possible combinations are actually seen using $N-1$ visualizations but this would require the user to memorize more information.

The addition of the third dimension to the display has inherent problems, however, and the developer has to be aware of the problem of misleading distortions in perspective displays.[21] Literature concerning the general usefulness of 3D representations for different tasks have concluded that the geometry relating actual space and perceived space is best

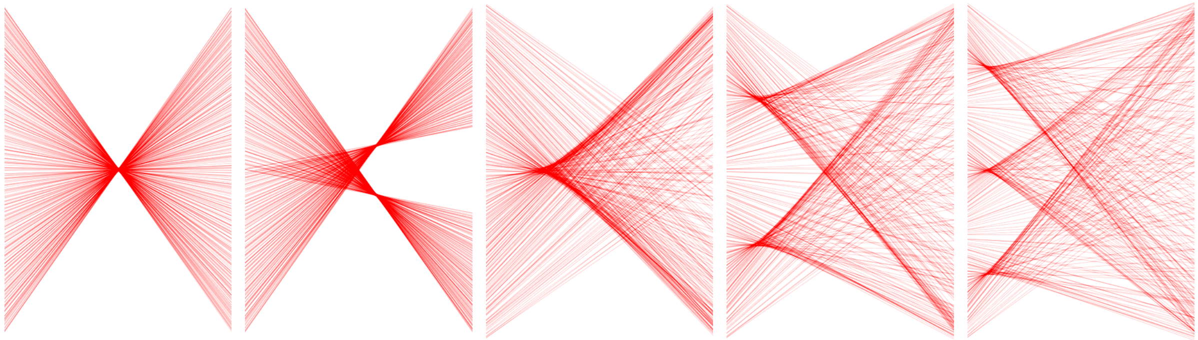

(a) v_1, v_2 (b) v_1, v_3 (c) v_1, v_4 (d) v_1, v_5 (e) v_1, v_6

Figure 1. The five mathematical relationships used in our study displayed using parallel coordinates. v_1, v_2 means that the relationship is found between variables v_1 and v_2. (a) Negative correlation. (b) Negative correlation with a discontinuity. (c) Sine function with one period. (d) Sine function with two periods. (e) Sine function with three periods.

described as being affine.[6,7] This means that distances along the depth direction (extending directly away from the eye) cannot be related systematically to distances in a direction perpendicular to the depth direction (the frontoparallel direction). Hence, the location of objects along the line of sight into a computer display (or depth axis) would be ambiguous. A series of empirical studies,[22–24] have shown that human observers make large errors when asked to estimate metric properties of patterns slanted or tilted in depth. They do not make such errors with slanted stimuli, however, when asked to judge qualitative properties, defined as properties invariant over affine transformations. When examining multi-relational 3D parallel coordinates it is easy to see that a number of relationships between axes will form patterns that are of this easy to perceive, qualitative nature. Therefore, this visualization method would, for theoretical reasons, not be affected by the general difficulties often found when examining the usability of 3D presentations.[25] Instead its benefits in terms of being able to depict more inter-relationships between axes, compared with standard 2D parallel coordinates, should be directly useful to users in tasks where such inter-relationships are important.

3. EXPERIMENT

The aim of the present experiment was to examine whether the obvious differences in visual representation (axis arrangement) and means of interaction (changing the position of axes) between the different visualizations described above would translate into actual performance differences on different tasks. We presumed they would. We compared two parallel coordinates visualizations with (two) different arrangements of axes and (two) different types of direct manipulation by mouse actions:

- A standard 2D parallel coordinates display (2Dm) with manual permutation of axes by drag-and-drop.
- A multi-relational 3D parallel coordinates display (3Dm) with manual permutation of axes by mouse click.
- As a control condition we also included a 2D parallel coordinates display (2Da) with automatic permutation of axes.

Evaluating performance requires clearly defined answers, hence a key component in this experiment was the creation of the test data set. The data set must contain a well-defined set of relationships that the test participants can locate without interference from other factors in the data set, see Keim et al.[26] for an in depth discussion. In order to have full control of the relationships present we constructed a synthetic data set (described in detail in section, 3.1.1) which included five pairwise mathematical relationships. These relationships, or patterns as they will also be referred to from here on, were used as stimuli. All other patterns appearing in the visualization were to be considered as distractors. The five stimuli patterns depicted the following mathematical relationships; a negative correlation, a negative correlation with a discontinuity, and three instances of a sine function having one, two and three periods, respectively (figure 1). In choosing these mathematical relationships, we emphasized the usefulness of parallel coordinates in aiding a decision maker to determine the nature of a

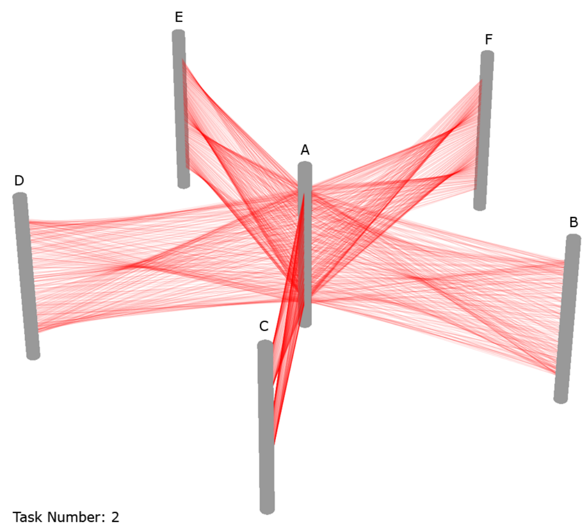

Figure 2. A screen shot from a 3D multi-relational parallel coordinates stimulus display.

relation (such as whether it is mainly linear or not, or has a discontinuity or not, or has cyclical components or not) rather than in determining the strength of a predetermined relation (the amount of correlation).

Each visualization had a six-axes arrangement, in circular or in parallel visualizing six variables labelled A–F (figures 2 and 3). The participants were not required to learn anything about the identity of, and logic behind the variables. They only had to recognize and discriminate between the patterns. Participants were instructed that some of the patterns in the visualization, the distractors, could be similar to a stimulus pattern but that only one was identical in all qualitative details and thus a correct answer (figure 4). Also, depending on how the axes were positioned with respect to each other, a stimulus pattern could be mirrored. These two different patterns (the stimulus pattern and its mirror reflection) were to be regarded as one and the same pattern. Their qualitative properties are identical and they differ only in orientation (figure 4).

To measure the actual results when comparing how the different visualizations would assist the participants, we created a judgement situation to be performed as a search task. We let participants perform two different tasks, both requiring them to identify and locate relationships between variables in a multivariate data set. We chose these tasks to be representative of typical operations that people perform on multivariate data sets. One of the tasks, referred to as task B, had a high complexity. The participants were asked to, in any individual trial, explore the visualization and search for all of the five possible stimulus patterns. However, in any visualization only four out of these five were present. The participants' task was to identify and locate these four and also to identify the one not present. The data set was constructed such that one, and only one, of the six variables, could form all of these stimulus patterns when positioned adjacent to the other variables (hence, in any visualization each of the patterns could appear only once). So the task implied the identification of the axis onto which that variable was mapped, a fact explicitly pointed out to the participants. As an aid for identification the participants had instruction materials where the five patterns were illustrated (figure 5). The patterns were labelled from 1–5 and these numbers were used when giving responses. The other task, A, had a lower complexity. Here the task was to search the visualization in order to find a specific one of the five possible patterns. All five patterns were present in each visualization and the target pattern could appear between any of the axis pairs. The target for each separate search was presented to the participant on a target-card (figure 6).

The experimental tasks required participants to modify the visualization to change the position of the attribute axes in order to search the data set. In all three conditions all modifications were carried out by mouse actions and the results were interactively updated. In the standard 2D parallel coordinates visualization with manual permutation of axes (2Dm), the participants changed an axis position by first selecting it with the middle button and then dragging the selected axis horizontally, with the button pressed, and dropping it beside the axis where it should be repositioned. As an axis was selected it turned blue, as feedback, and stayed blue while dragged until repositioned. In the multi-relational 3D parallel

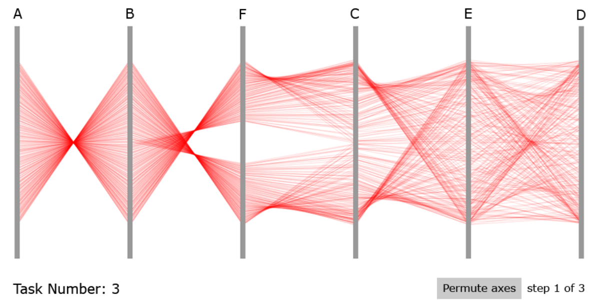

Figure 3. A screen shot from a standard 2D parallel coordinates stimulus display. The button to the right was only present in the 2D parallel coordinates stimulus display with automatic permutations where it was used for reposition of axes.

coordinates visualization (3Dm), the focus axis was changed by means of a single mouse click. A participant pointed to one of the outer axes as the new focus axis. By means of a middle button mouse click the axis was swapped with the previous focus axis. To rotate the visualization participants moved the mouse horizontally or vertically with the left mouse button pressed. In order to prevent disorientation the rotation was constrained such that the visualization could not be turned upside down. In the 2D parallel coordinates visualization with automatic permutations (2Da), the user alternated between the permutations by means of a mouse click on a button. Using the algorithm for six axes meant that in three views of the visualization, including the first displayed, all axes had been positioned adjacent to each other and all possible relationships had been displayed once. Hence, with two clicks the user had explored the entire data space. The button was placed down to the right of the visualization as illustrated in figure 3. It was labelled, "step 1 out of 3" (at the onset of trial display), "step 2 out of 3" (after one click) and "step 3 out of 3" (two clicks) respectively. A third mouse click again showed the first iteration.

Before the experiment we predicted:

H1: In task A (simple) there would be no significant performance differences in terms of search times between the 2Dm visualization and the 3Dm visualization.

H2: In task B (complex task) performance with the 3Dm visualization would be faster than 2Dm since the need for many comparisons between axes would be more difficult and time consuming using the 2Dm visualization.

H3: 2Da would be faster than 2Dm and 3Dm in both tasks. In terms of interaction it is the easiest and most straightforward visualization. Not being able to control the permutations and decide where to reposition an axis is constraining to the user but, with this visualization having only six variables, we did not expect that to slow down performance in these two tasks.

H4: Regarding accuracy we did not predict any differences between the visualizations.

3.1. Method

3.1.1. Stimuli

As described above each stimulus display was comprised of a visualization with six axes in either a parallel or a circular arrangement. Between each pair of axes a pattern was depicted. There were five specific stimulus patterns which were constructed in a synthetic data set containing 300 six-dimensional data items (six variables each having 300 samples). In the two variations of 2D parallel coordinates visualizations, this gave 300 polylines intersecting all six axes. The 3D multi-relational parallel coordinates visualization gave 300 instances of five lines, where each of these five lines connects a point on the focus axis with one of the five outer axes. Given x_i, where $i = 1, \ldots, 300$, equally distributed on the interval $[0,1]$, the sampled data set containing the six variables $v_1, \ldots, v_6$ were constructed as follows:

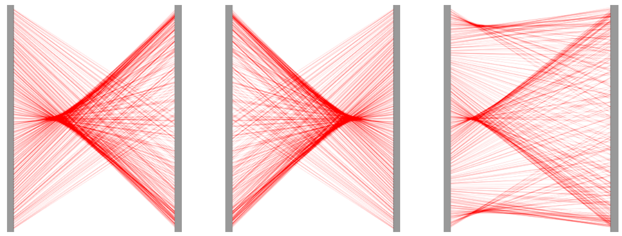

(a) (b) (c)

Figure 4. If the target of search, when exploring the visualization in a stimulus display, was pattern (a) then pattern (b) would be the correct answer. That pattern is identical in every aspect except for the orientation. Pattern (c), however, would be an incorrect answer, even though it is similar in several aspects.

$$v_1(x_i) = x_i + e \tag{1}$$

$$v_2(x_i) = -x_i + e \tag{2}$$

$$v_3(x_i) = \begin{cases} -x_i + e, & \text{if } 0 \leq x_i \leq 0.5 \\ -x_i - \gamma + e, & \text{if } 0.5 < x_i \leq 1 \end{cases} \tag{3}$$

$$v_4(x_i) = \sin(2\pi x_i) + e \tag{4}$$

$$v_5(x_i) = \sin(4\pi x_i) + e \tag{5}$$

$$v_6(x_i) = \sin(6\pi x_i) + e \tag{6}$$

where e is a small random noise and γ is an offset, in this experiment set to 0.25, creating a discontinuity. After construction, all variables were normalized to the range $[0,1]$. For a multivariate data set containing N variables, there are $_NC_2$ pair-wise combinations that can occur. Our data set with six variables, yields 15 possible combinations. The five combinations giving the relationships that the participants in our study examined are illustrated in figure 1. These relationships are obtained from the combinations between the first variable, v_1, and the other five variables $v_2, \ldots, v_6$.

3.1.2. Apparatus and Viewing Condition

The experimental sessions took place in a large room with two long tables facing each other. On each of the two tables five stationary Windows workstations were arranged side by side. All tests were carried out on a Pentium III computer with a Dell 19" monitor set to a resolution of 1280 $\times$ 1024 pixels, 85 Hz refresh rate and 32-bit colour depth. The computers were equipped with Nvidia TNT graphics cards. The visualizations were produced through a program that showed the visualization in a rectangular display of size 1200 $\times$ 800 pixels. The viewing distance was approximately 57 cm and the viewing position was unrestricted. The participants interacted with the visualization by means of a three button mouse and a keyboard placed in front of each workstation.

3.1.3. Experimental Design

The study was performed as a four-factor mixed design with two within-subjects factors: task type A (simple) vs. B (complex), and block of trials. The two between-subject factors were: type of visualization (2Dm vs. 3Dm vs. 2Da) and sequence of presentation of task type. Having type of visualization as a between-subject factor was chosen to prevent carry-over effects (e.g., learning or fatigue) from one visualization to the next. Participants were randomly assigned to one of the three visualizations. An equal number of males and females participated in each visualization group and sequence of task. Each participant completed both tasks with one of the three visualization types, the sequence of presentation of the two tasks was counterbalanced. Each experiment was performed over two separate timed trial sessions. Half of the participants in each viewing condition (type of visualization) started with task A and half of the participants started

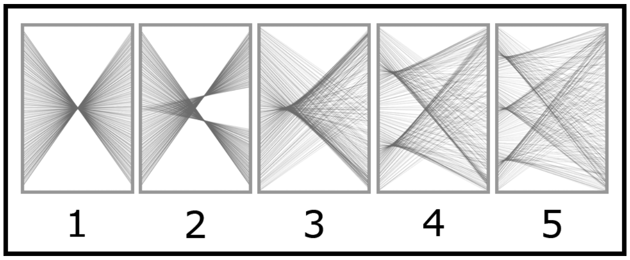

Figure 5. The five patterns as they were presented to the participants in their instruction material. They were equally spaced on an A4 page, each pattern having a size of 6.5×5.5cm

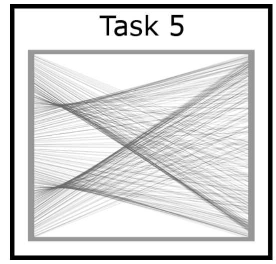

Figure 6. The target for search was presented on a target card, size 10×8cm (pattern 6.5×5.5cm).

with task B. Participants completed all trials in the first task before beginning the second task. The session in which the participants performed task A consisted of two blocks of 20 trials. In each block the five patterns appeared four times each in a randomized order. Hence each participant searched for a total of 40 patterns (5×4×2). The other session, where the participants performed task B, consisted of 15 trials. These trials were not blocked, thus the four patterns present out of the five possible (one pattern being absent) varied randomly from trial to trial. This design yielded a total of 55 trials per participant.

For each new trial the mapping between the variables and the axes was randomized. Depending on the task and visualization type, the probability that the sought pattern (or patterns) would appear directly in the initial view differs. For task A (simple) the probability is $\frac{2}{N} = \frac{1}{3}$ for 3Dm and $\frac{2(N-1)(N-2)!}{N!} = \frac{1}{3}$ for both 2Dm and 2Da. For task B (complex), the probability is $\frac{1}{N} = \frac{1}{6}$ for 3Dm and 0 for both 2Dm and 2Da. It is clear that this difference is a feature of the two different representations (2D vs. 3D) which may have a small effect on the user performance and so must be taken into account in the analysis of the response times.

3.1.4. Procedure

The experiment was carried out in groups of between six and nine participants per session with each participant carrying out the test individually using a separate workstation. Applying this method during evaluation is somewhat unusual but there were no reasons to expect any effect on the results due to having the participants perform their experimental tasks in the same room at the same time. Therefore, we chose this experimental setup in order to be time efficient. A session was divided into two parts: an introduction for familiarizing the participants with the situation and the timed trial. In order to prevent any disturbance the participants were instructed that when the timed trial parts of the experimental session had begun questions would not be allowed. They were also asked to remain quiet and seated when they had completed all of the tasks until the administrator determined the test to be over. A time limit of one and a half hours was used. Any participant not finished within this time was to be excluded from the analysis (the latter, however, was not made explicit to the participants only that they did not have to complete the remaining tasks). This did not occur during this experiment, however.

First, background information was obtained from each participant concerning their age, sex, occupation and experience with information visualizations. Participants then reviewed written instruction material and completed a block of practice trials to learn the concepts and usage of the visualization and the two types of task to be performed. The administrator helped participants when needed during the training session to resolve any confusion about how to operate the software and carry out the tasks before the timed trials began. Participants were instructed that, although being timed they should try to be as accurate as possible in solving each task. Also, feedback was given on practice trials to ensure they would know the correct answers when identifying the stimuli patterns. The participants' task in any individual trial was to search the visualization and find the correct pattern (task A) or patterns (task B) as specified. The sequence of trials was self-paced. Each trial was composed of a stimulus display and a response display. To switch from the stimulus display to the response display, the space key was pressed. Response times were measured from onset of the stimulus display until pressing space, thus not including the time to record the response. The display used for response in task B had five buttons numbered from 1–5. Each number corresponded to one of the five stimuli patterns creating a logical mapping between the images on the

Table 1. Summary of the error data. For task A (simple) data was obtained for 10 persons and 20 trials (second block) per visualization type. The largest amount of error for any participant was 3 out of 20 trials. For task B (complex) data was obtained for 10 persons and 15 trials per visualization type. The largest amount of error for any participant was 2 out of 15 trials.

	2Dm	3Dm	2Da
Total errors in task A (200 trials)	4	5	5
Total errors in task B (150 trials)	4	5	2

instruction material and the response display. As described earlier, the correct response in this task was the number of the pattern not present in the visualization. In task A the correct response was the axis pair between which the target pattern appeared. There were 16 buttons each labelled with one possible axis combination (e.g, B-C, C-B, A-E). To initiate a new trial after giving a response, the participant clicked another button on the response display labelled "Continue" which was followed by a new button labelled "Start". A new visualization was then presented and the experiment continued. This procedure was used to ensure that the participants would not start a new timed trial until ready. This was of great importance in the task A session where the participants had each target for search presented on a separate target-card. These were numbered and arranged upside-down in a pile of 40 cards. Participants were instructed to turn the next card before clicking the "Start" button to initiate a new trial. Errors and reaction times for each trial were recorded. No feedback was provided.

After finishing all trials the participants then completed a subjective satisfaction questionnaire concerning the visualization. They responded to 16 statements and rated their satisfaction with the visualization using a Likert-scale ranging from 1 (strongly disagree) to 5 (strongly agree). To summarize, each participant completed 10 practice trials (five in each task type) and 55 experimental trials. A full experimental session lasted seventy minutes (including the introductory part).

3.1.5. Participants

30 participants took part in the experiment, 15 male and 15 female. They were all undergraduate or graduate students at Uppsala University, aged between 21 and 39 years. All participants had normal or corrected-to-normal vision. None of the participants were familiar with the parallel coordinates visualization and they had no prior knowledge of the purpose of the experiment or the specific hypotheses. They received a small compensation for taking part in the experiment.

3.2. Results

Data obtained from all 30 participants were analysed. First we analysed the error data. Errors were scarce. The results are summarized in table 1. No significant differences between visualization types in terms of error rates, for either of the two tasks, were found by means of resampling methods.[27] The probability for the difference in the group mean error rates in the 2Da condition (1.3%) and the 3Dm condition (3.3%) in task B to occur by chance was estimated to be approximately 0.5. Also, the few trials where participants made errors were often not fastest and the accurate ones often not slowest. This indicates that no speed-accuracy trade-off was present. We therefore conclude that, for our participants, the search times are a valid representation of performance.

Thereafter, we turned to the analysis of the time data. First, we analysed the data from the 15 trials in task B, the complex task. Before statistical testing we employed a logarithmic transformation since reaction time data are typically not normally distributed. Group mean values were calculated and a between-subject ANOVA was carried out using a decision criterion of 0.05 (the same criterion was also used in all subsequent testing). Factors were visualization type (2Dm vs. 3Dm vs. 2Da) and sequence of task type. There was a significant effect of visualization type $F(2,24) = 5.528$, $p < 0.01$) as well as of sequence ($F(1,24) = 4.642$, $p < 0.041$). No interaction effects were found. In order to specifically test our hypothesis that participants using the 3Dm visualization would perform better than participants using the 2Dm visualization a t-test (for independent samples) was performed. The response times for the 3Dm visualization were significantly faster than for the 2Dm visualization ($t = 2.4891$, $n = 10$, $p < 0.0228$) thus our hypothesis was not rejected. The group mean value for the search times in the 3Dm condition was 23.9 seconds with a standard deviation of 1.35 while in the 2Dm condition it was 37.2 seconds with a standard deviation of 1.61. We also performed the same test comparing 3Dm and 2Da (group mean value in the 2Da condition was 20.0 seconds, standard deviation was 1.70). No significant difference was obtained ($t = -0.9339$, $n = 10$, $p < 0.3627$).

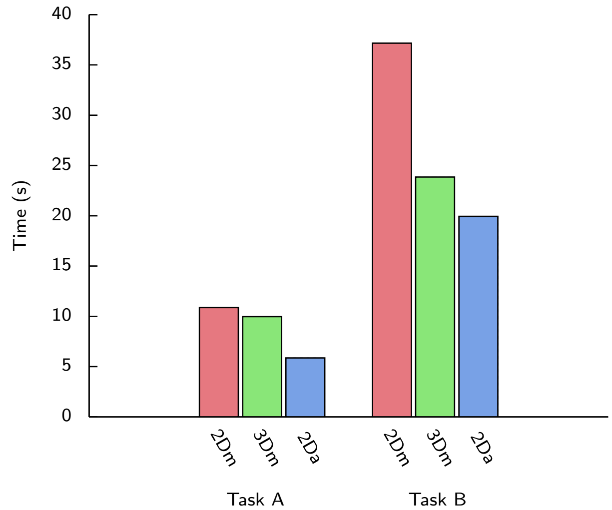

Figure 7. The main findings from both tasks in terms of search time. Manual exploration using 2D parallel coordinates (2Dm) produced the highest group mean completion time in both tasks. Manual exploration using 3D parallel coordinates (3Dm) is 36 percent faster than 2Dm for task B (complex) and the difference is statistically significant. In task A (simple) the difference is negligible.

Next, we turned to task A. The response times were transformed and calculated in the same way as described earlier. A mixed ANOVA was carried out with visualization type (2Dm vs. 3Dm vs. 2Da) and sequence of task type as between-subject factors and block of trials as within-subject factor. Again there was a significant effect for visualization type $F(2,24) = 7.493$, $p < 0.0031$) however, no effect of sequence was obtained. There was an effect of learning, the differences in response times between the two blocks were significant $F(1,24) = 33.543$, $p < 0.00001$). In this task we did not predict any significant difference in performance between participants using the 3Dm visualization and those using the 2Dm visualization. A t-test (independent samples) revealed no significant difference between conditions ($t = 0.6307$, $n = 10$, $p < 0.536$). Thus, the hypothesis could not be rejected. The group mean value for the response times was 10.0 seconds in the 3Dm condition with a standard deviation of 1.29, and 10.9 seconds in the 2Dm condition with a standard deviation of 1.42. A comparison between the 3Dm visualization and the 2Da visualization (in the 2Da condition, the group mean value was 5.9 seconds, standard deviation of 1.50) revealed that there was a significant difference between these types of visualizations as well ($t = -3.5200$, $n = 10$, $p < 0.002$). Between the 2Dm and 2Da visualizations the difference was also significant ($t = -3.6434$, $n = 10$, $p < 0.001$.) The main findings from the analysis of both tasks are summarized in figure 7.

Finally, we analysed the subjective measures obtained from the post-test questionnaire. Figure 8 illustrates ratings for four out of the 16 statements answered, raw response values as well as group mean response values of the five point Likert scale (1 strongly disagree to 5 strongly agree). The statements are one regarding the visualizations overall usability; "I thought the visualization was easy to use", and three regarding their ease of use; "It was easy to perform task A (simple)" "It was easy to perform task B (complex)", and finally, "It was easy to comprehend how the movement of an axis influenced the visualization". Please note that the word "difficult" was used as often as the word "easy" in the questionnaire. However we chose these four statements to be representative of the subjective measures. The obtained differences between the three visualization types in terms of task completion time for both tasks are reflected in the participants' subjective satisfaction ratings of the different visualizations. The standard 2D parallel coordinates visualization with manual permutation obtained the lowest group mean value and the one with automatic permutation the highest. However, on the statement "It was easy to comprehend how the movement of an axis influenced the visualization" the 2Da visualization obtained the lowest group mean value (rating). In spite of their fast performance with (and otherwise favourable rating of) the 2Da visualization, participants seemed to get confused by the automatic permutation of axes. This visualization will most certainly be less supportive in the case of even more complex tasks requiring many comparisons between adjacent axes. As illustrated, the ratings are in favour of the 3Dm visualization.

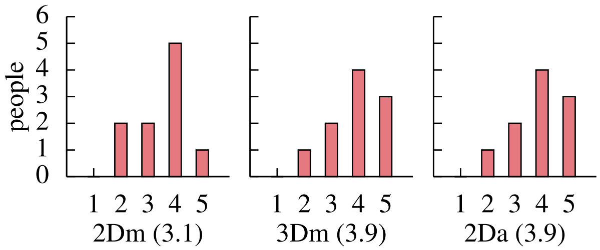

(a) Statement: "I thought the visualization was easy to use". When rating the overall usability there was no difference in opinion between participants using the 3Dm visualization and the participants using the 2Da visualization (in terms of group mean value).

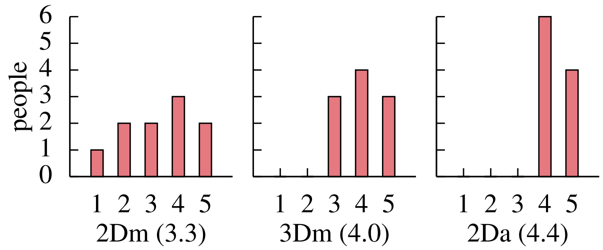

(b) Statement: "It was easy to perform task A (simple)". The participants' responses reflect the differences between the visualizations in terms of mean search time performance. Performance was slowest with 2Dm, slightly faster with 3Dm and faster still using the 2Da visualization.

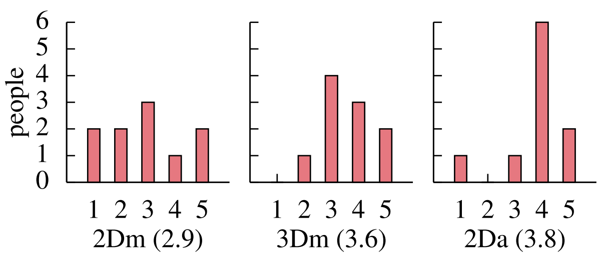

(c) Statement: "It was easy to perform task B (complex)". Again, the participants' responses reflect the differences in terms of obtained search performance.

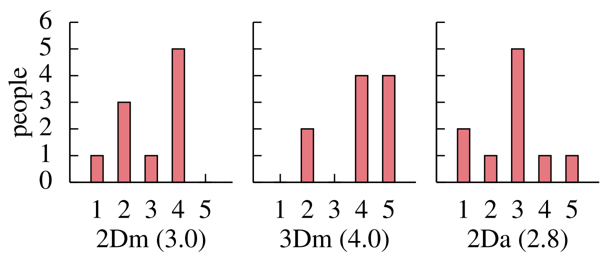

(d) Statement: "It was easy to comprehend how the movement of an axis influenced the visualization". In spite of the ease of performance with (and otherwise favourable rating of) the 2Da visualization, participants seemed to get confused by the automatic permutation of axes. This visualization will certainly be less supportive in more complex tasks. As illustrated, the ratings are quite strongly in favour of 3Dm.

Figure 8. A summarization of the participants' subjective ratings of four Likert scale statements (1 strongly disagree to 5 strongly agree). Group mean values within brackets.

4. DISCUSSION AND CONCLUSIONS

The multi-relational 3D parallel coordinates display has an obvious advantage over standard 2D parallel coordinates for the tasks investigated here in that it simultaneously displays $N-1$ relations of interest as compared to two in the standard 2D case. However, this comes at the cost of having the patterns of lines between the variables distorted by perspective effects. The main aim of this study was therefore to examine whether the combined effect of having more relations simultaneously visible, but having them distorted by perspective effects, was advantageous compared to a standard 2D parallel coordinates display. From perceptual theory, supported by a large number of experiments, we hypothesized that this combined effect would indeed be advantageous. The patterns in the visualizations used in our tasks have properties that are invariant over affine transformations. Therefore distortions in the patterns introduced by the perspective effects should not be a hindrance to a viewer; the basic patterns should all be easily picked up by the human visual system regardless of the perspective.[22–24] The results from the study support this. In the complex task, where simultaneously visualizing five patterns, the 3Dm display led to better performance than the 2Dm display. An illustration of this can be seen in figure 2, the pattern defining the relationship between axes A and C is the pattern shown in figure 1(b). The pattern is easily perceived in the visualization despite the extremely narrow perspective. More surprising was that performance in the 2Da condition was not significantly faster than in the 3Dm condition for the more complex task. Users performing real tasks with real multivariate data sets will usually need to take into consideration several more variables than six. An automatic permutation of axes will be even more confusing in such cases and it is fair to expect the 2Da visualization to be less useful as dimensionality grows. Something that could help explain the result is that we did not use animation which is known to help a viewer to follow transitions.

However, results from the Multiple Object Tracking paradigm (MOT) show that visually tracking more than five objects at a time is difficult.[28,29] Together with the subjective data (figure 8(d)), it therefore seems reasonable to primarily propose the use of the 3Dm visualization for real tasks similar to the ones we employed except in very simple cases.

As described in the methods section the probability, in the more complex task, for all of the four target patterns to occur in the initial display of each trial was $\frac{1}{6}$ in the 3Dm condition and 0 in the 2Dm condition. From a pure experimental point of view this could be considered a confounding factor, jeopardizing the validity of our results. However, the mapping of variables to axes was always randomly determined in all trials and in all conditions in this task. Therefore the difference in probability described above merely reflects inherent properties of the different visualization methods and are therefore representative of their behaviour in real situations. In considering this study as a usability evaluation trying to present representative tasks to the participants, this difference in probability cannot be considered a confounding factor. Moreover, an examination of the log files revealed that the trials in the 3Dm condition where all four target patterns were present in the default view were not systematically performed faster by the subjects than trials where this was not the case.

The ability to perform interactive operations on a visualization is generally considered of great importance. It greatly facilitates data analysis and must be regarded as a prerequisite when dealing with large data sets. The different amount of interaction required in the visualizations we have evaluated implies different costs in terms of time to get new information into the visual field. This time can reliably be predicted by performing a GOMS[30] analysis. We have not performed such an analysis but the click-and-drag rotation and the point-and-click axis selection of the 3Dm visualization should be more costly in terms of time than the simple button click interaction required in the 2Da visualization. Even so there was no significant difference in time performance between these two. This further supports the opinion that the visual analysis is more efficiently supported by the 3Dm visualization.

Finally, even if the multi-relational display worked well for these synthetically generated patterns, we cannot generalize to say it will work for all other data sets. However, it seems fair to argue that as long as the patterns formed between the variables are invariant over affine transformations, our results should generalize.

Overall, we propose that standard 2D parallel coordinates is best suited to tasks concerning relationships between data items while multi-relational 3D parallel coordinates could be efficiently used to establish relationships between variables. Because of their respective usage it is also likely that combining the two views into a single display would provide superior performance for many situations. How to achieve this will be addressed in future user studies where different solutions to the integration problem will be considered and evaluated. Also, further studies with empirical validation on real tasks performed on real data sets with feedback from domain expert users will provide valuable insight.

Future studies are also planned that will further investigate the 3D representation. For instance, to examine how slanted in depth a pattern can be before it is no longer accurately perceived and so to establish a perspective threshold. This is an important aspect affecting both the amount of manipulation needed to interpret the display, and the number of dimensions that can be simultaneously displayed. Such studies are aimed at establishing how many variables can meaningfully be shown in a multi-relational 3D visualization. Other possible future experiments include keeping the same relationships and varying the level of random noise within each pattern to establish perceptual threshold values, and also to investigate other types of relationships than the ones used here.

ACKNOWLEDGMENTS

This work has been funded by the Swedish Foundation for Strategic Research, grant A3 02:116. We would like to thank Matthew Cooper and Mats Lind for advice and feedback.

REFERENCES

1. R. Kosara, C. Healey, V. Interrante, D. Laidlaw, and C. Ware, “Visualization viewpoints,” *IEEE Computer Graphics & Applications* **23**(4), pp. 20–25, 2003.
2. M. Tory and T. Möller, “Human factors in visualization research,” *IEEE Transactions on Information Visualization and Computer Graphics* **10**(1), pp. 72–84, 2004.
3. A. Inselberg, “The plane with parallel coordinates,” *The Visual Computer* **1**, pp. 69–91, 1985.
4. J. Johansson, M. Cooper, and M. Jern, “3-dimensional display for clustered multi-relational parallel coordinates,” in *9th IEEE International Conference on Information Visualization*, pp. 188–193, 2005.

5. C. Forsell, S. Seipel, and M. Lind, “Simple 3D glyphs for spatial multivariate data,” in *IEEE Symposium on Information Visualization*, pp. 119–124, 2005.
6. J. T. Todd and J. F. Norman, “The visual perception of 3-D shape from multiple cues: Are observers capable of perceiving metric structure?,” *Perception and Psychophysics* **65**, pp. 31–47, 2003.
7. J. T. Todd, “The visual perception of 3-D shape,” *Trends in Cognitive Science* **8**(3), pp. 115–121, 2004.
8. M. Lanzenberger, S. Miksch, and M. Pohl, “Exploring highly structured data: a comparative study of stardinates and parallel coordinates,” in *IEEE International Conference on Information Visualisation*, pp. 312–320, 2005.
9. R. M. Pillat, E. R. A. Valiati, and C. M. D. S. Freitas, “Experimental study on evaluation of multidimensional information visualization techniques,” in *Latin American Conference on Human-Computer Interaction*, pp. 20–30, 2005.
10. G. Andrienko and N. Andrienko, “Constructing parallel coordinates plot for problem solving,” in *1st International Symposium on Smart Graphics*, pp. 9–14, 2001.
11. E. J. Wegman and Q. Luo, “High dimensional clustering using parallel coordinates and the grand tour,” *Computing Science and Statistic* **28**, pp. 352–360, 1997.
12. Y.-H. Fua, M. O. Ward, and E. A. Rundensteiner, “Hierarchical parallel coordinates for exploration of large datasets,” in *IEEE Visualization*, pp. 43–50, 1999.
13. J. Johansson, P. Ljung, M. Jern, and M. Cooper, “Revealing structure within clustered parallel coordinates displays,” in *IEEE Symposium on Information Visualization*, pp. 125–132, 2005.
14. R. Wegenkittl, H. Löffelmann, and E. Gröller, “Visualizing the behaviour of higher dimensional dynamical systems,” in *IEEE Visualization*, pp. 119–125, 533, 1997.
15. G. Falkman, “The use of a uniform declarative model in 3D visualisation for case-based reasoning,” in *6th European Conference on Advances in Case-Based Reasoning*, pp. 103–117, 2002.
16. E. Fanea, S. Carpendale, and T. Isenberg, “An interactive 3D integration of parallel coordinates and star glyphs,” in *IEEE Symposium on Information Visualization*, pp. 149–156, 2005.
17. E. J. Wegman, “Hyperdimensional data analysis using parallel coordinates,” *Journal of the American Statistical Association* **85**(411), pp. 664–675, 1990.
18. W. Peng, M. O. Ward, and E. A. Rundensteiner, “Clutter reduction in multi-dimensional data visualization using dimension reordering,” in *IEEE Symposium on Information Visualization*, pp. 89–96, 2004.
19. C. Tominski, J. Abello, and H. Schumann, “Axes-based visualizations with radial layouts,” in *ACM Symposium on Applied Computing*, pp. 1242–1247, 2004.
20. C. Tominski, J. Abello, and H. Schumann, “Interactive poster: 3D axes-based visualizations for time series data,” in *IEEE Symposium on Information Visualization, Poster Proceedings Compendium*, pp. 49–50, 2005.
21. S. R. Ellis, “On the design of perspective displays,” in *44th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 411–414, 2000.
22. J. T. Todd, J. S. Tittle, and J. F. Norman, “Distortions of three-dimensional space in the perceptual analysis of motion and stereo,” *Perception* **24**, pp. 75–86, 1995.
23. J. F. Norman, J. T. Todd, V. J. Perotti, and J. S. Tittle, “The visual perception of three-dimensional length,” *Journal of Experimental Psychology: Human Perception and Performance* **22**(1), pp. 173–186, 1996.
24. M. Lind, G. P. Bingham, and C. Forsell, “Metric 3D structure in visualization,” *Information Visualization* **8**(3), pp. 51–57, 2003.
25. J. S. Tittle, D. W. Woods, M. Howard, and F. Phillips, “The role of 2-D and 3-D task performance in the design of and use of visual displays,” in *45th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 1215–11219, 2001.
26. D. A. Keim, D. Bergeron, and R. Pickett, “Test data sets for evaluating data visualization techniques,” *Perceptual Issues in Visualization* , pp. 9–22, 1994.
27. C. E. Lunneborg, *Data Analysis by Resampling: Concepts and Applications*, Duxbury Press, 1999.
28. P. Cavanagh, “Attention: Exporting vision to the mind,” in *Neuronal Basis and Psychological Aspects of Consciousness*, pp. 129–143, 1999.
29. S. Yantis, “Multielement visual tracking: Attention and perceptual organization,” *Cognitive Psychology* **24**(3), pp. 295–340, 1992.
30. S. K. Card, T. P. A. Moran, and A. Newell, *The Psychology of Human-Computer Interaction*, Lawrence Erlbaum associates, 1993.