

Tracing Tuples Across Dimensions: A Comparison of Scatterplots and Parallel Coordinate Plots

X. Kuang¹ and H. Zhang¹ and S. Zhao¹ and M. J. McGuffin²

¹National University of Singapore, Singapore

²École de technologie supérieure, Canada

Abstract

One of the fundamental tasks for analytic activity is retrieving (i.e., reading) the value of a particular quantity in an information visualization. However, few previous studies have compared user performance in such value retrieval tasks for different visualizations. We present an experimental comparison of user performance (time and error distance) across four multivariate data visualizations. Three variants of scatterplot (SCP) visualizations, namely SCPs with common vertical axes (SCP-common), SCPs with a staircase layout (SCP-staircase), and SCPs with rotated axes between neighboring cells (SCP-rotated), and a baseline parallel coordinate plots (PCP) were compared. Results show that the baseline PCP is better than SCP-rotated and SCP-staircase under all conditions, while the difference between SCP-common and PCP depends on the dimensionality and density of the dataset. PCP shows advantages over SCP-common when the dimensionality and density of the dataset are low, but SCP-common eventually outperforms PCP as data dimensionality and density increase. The results suggest guidelines for the use of SCPs and PCPs that can benefit future researchers and practitioners.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing algorithms

1. Introduction

Multivariate data is a commonly encountered type of data (e.g., in relational databases), consisting of a list of points or tuples, each corresponding to a row in a table, whose columns are the attributes or variables of the data. Two widely used visualization techniques for multivariate data are parallel coordinate plots (PCP) and scatterplots (SCP) [Weg90, KD09, TGS04, SS05, AR11]. PCPs display each tuple as a polygonal line intersecting parallel axes, each representing one of the variables, thus providing a continuous view of the multidimensional values of the data tuples [Ins85]. SCPs, on the other hand, show only 2 variables per plot, but can be combined to visualize multivariate data with more than 2 dimensions, such as in a scatterplot matrix [Har75].

Despite the call for rigorous evaluation of experimental visualization techniques over a decade ago [WB97], to date, much still remains unknown about the respective advantages of PCPs and SCPs for different user analytic tasks. To our knowledge, there are only two empirical compar-

isons of these techniques. One [LMvW10] asked users to estimate correlation coefficients using PCPs and SCPs, and another [HvW10] asked users to count clusters; both studies found SCPs to be superior. Given these results, it seems unclear what advantage, if any, PCPs provide. However, the tasks in the two previous studies are just two of many possible tasks. Several other tasks with visualizations have been identified [Shn96, AES05] and have yet to be tested.

We extend previous efforts by comparing SCPs and PCPs for the task of *value retrieval*, a fundamental task that is the first in the taxonomy of analytic tasks by Amar *et al.*'s [AES05] and said to be a building block of other tasks such as finding extrema or sorting [AES05]. As an initial exploration, our study focuses on differences due to the basic visual designs of SCPs and PCPs in their static form. We believe it is important to understand trade-offs due to their basic visual designs before investigating the effects of visual or interactive enhancements. Therefore, brushing, linking as well as additional visual enhancements such as gridlines are not included in this investigation. Furthermore, value re-

retrieval by visual scan is commonly performed in practice since it is an integral component of many higher-level tasks in which explicit clicks would be inappropriate.

We conducted two controlled experiments involving four visualization techniques: three SCP variants (SCP-common, SCP-rotated, and SCP-staircase) and the baseline PCP, on datasets of varied dimensionalities and densities. It was found that SCP-rotated and SCP-staircase are not suitable for value retrieval. PCP and SCP-common yield better performance and are preferred by participants, but each is suited for different scenarios: PCP is better at low dimensionality and low density, while SCP-common is better in the opposite case. Increasing dimensionality seems to only affect performance with PCP, not SCP-common. Increasing density, while affecting both visualizations, has a stronger effect on PCP than SCP-common. Such differences are likely due to the different value retrieval strategies adopted by users and the different visual encodings of data tuples in the two visualization techniques (points versus lines). These results may be used by researchers and practitioners to better understand the differences between PCPs and SCPs, and to promote their appropriate use in the future.

2. Related Work

Two aspects of previous research are related to our study: variants and hybrids involving scatterplots and parallel coordinate plots, and their comparisons.

A single SCP depicts two variables, and is thus insufficient for multivariate data. The scatterplot matrix (SPLOM) [Har75] shows every possible pairing of variables with multiple SCPs. Other variants with multiple SCPs have been proposed [QCX*07, VM CJ10] that show a subset of the SCPs in a SPLOM, arranged with various layouts. Qu *et al.* [QCX*07] showed a row of SCP cells, where consecutive SCPs have an axis in common that is rotated (a technique we call SCP-rotated). These SCPs correspond to cells that are adjacent to the diagonal in a SPLOM. Viau *et al.* [VM CJ10] consider rows of SCPs taken directly from a SPLOM, in which all the SCPs of the row have the same vertical axis (a technique we call SCP-common), privileging the variable along the shared, vertical axis. Viau *et al.* [VM CJ10] also presented a novel “staircase” arrangement (we call SCP-staircase), where adjacent SCPs share a common axis.

Parallel coordinates [Ins85] lend themselves naturally to multivariate data due to their inherently multidimensional design. Research into PCP variants has examined the use of curves instead of polylines [The00], variations in colors and transparency, and animation for line disambiguation, as surveyed by Holten and van Wijk [HvW10]. Qu *et al.* [QCX*07] have extended PCPs with S-shaped axes to indicate wind direction. Artero *et al.* [AdOL04] proposed an interactive PCP variant.

Hybrid visualizations that combine SCPs and PCPs have

included embedding SCP cells between PCP axes [HvW10], scattering points along curves between PCP axes [YGX*09], the parallel scatterplot matrix [VM CJ10], and highly flexible custom visualizations integrating SCPs and PCPs [CvW11].

In contrast to the many variants and hybrids of SCPs and PCPs, and evaluation within PCP variants [HLKW12], comparisons between these two families of visualizations have been rare. Li *et al.* [LMvW10] found SCPs to be significantly superior to PCPs for judging correlation coefficients. Holten and van Wijk [HvW10] compared cluster identification performance over several PCP variants, and found that the PCP variant with embedded SCPs significantly outperformed other variants, implying that SCPs hold an advantage over PCPs. Our work extends these previous studies by comparing performance in value retrieval with PCPs and three variants of SCPs.

3. Experiment Design

We conducted two controlled experiments to compare SCP and PCP visualizations. The next few subsections first describe aspects common to both experiments.

3.1. Task

We define a “value retrieval” task [AES05] in the context of multivariate data: given the numerical value of one attribute of a data tuple, find the numerical value of another attribute of the same data tuple. Value retrieval is a common, fundamental user analytic task. For example, if a user wants to find the average mileage for a car with 230 horsepower in a multivariate visualization, s/he may first locate the horsepower axis and find a data tuple corresponding to 230 horsepower, and then trace the tuple to the mileage axis and read its value off that axis. In general, it is possible for some axes to correspond to categorical (such as car brands) or ordinal (such as degree of satisfaction) variables, however our study focuses on the most general case: quantitative variables.

3.2. Independent Variables

Our experiments involved three independent variables: visualization technique, data dimensionality, and data density.

3.2.1. Visualization Technique

PCPs have a single straightforward layout (Figure 1:(a)). SCPs, in contrast, afford many different layouts. The aforementioned full SPLOM shows all pairings of variables, so its space utilization is quadratic with the number of variables. PCPs, however, have space requirements linear with the number of variables. A fair comparison requires all techniques occupy the same space. Therefore, we evaluated three SCP variants with linear space requirements:

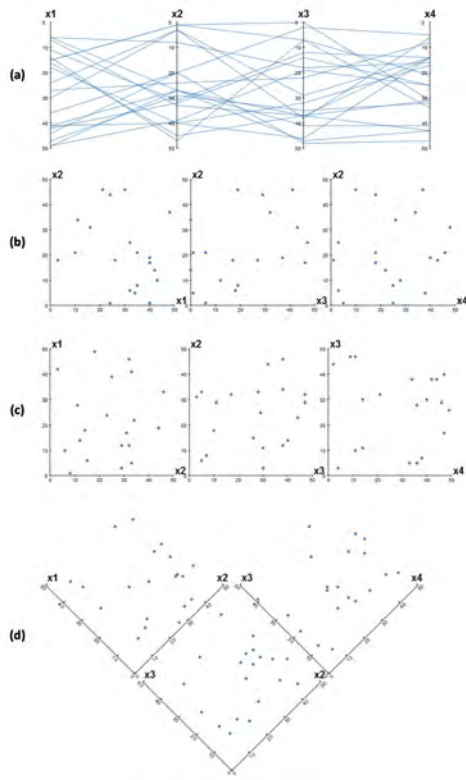


Figure 1: The four evaluated techniques (a). Baseline PCP; (b). SCP-common; (c). SCP-rotated; (d). SCP-staircase.

1. SCP-common: a row of SCPs taken from a standard SPLOM. SCP-common has the advantage of having a common and aligned vertical axis for all its individual cells (Figure 1:(b)).
2. SCP-rotated: a row of SCPs formed from the SCPs adjacent to the diagonal of a SPLOM (Figure 1:(c)).
3. SCP-staircase: adjacent SCPs in this layout have a common and aligned axis (Figure 1:(d)).

As shown previously [VMCJ10], PCPs and each of the above SCP variants require $O(NL^2)$ space, where L is the length of the axes, and N is the number of variables.

3.2.2. Data Dimensionality and Density

Two important characteristics of a multidimensional visualization are the dimensionality of the data, and the density of tuples (number of tuples per unit display area). Since both characteristics may affect the difficulty of value retrieval, they were both varied in our experiments.

Following Li *et al.* [LMvW10], we conducted a pilot study with five participants (1 female, 4 males) to identify the feasible range of data density for each visualization technique. We fixed the size of cells in the visualizations to be

49 cm². For each visualization technique, participants tried to finish value retrieval tasks with datasets of increasing densities, starting at 5 tuples per cell with increments of 5 tuples, until they found it too difficult to complete the trials and gave up. We recorded the number of tuples each participant completed just before giving up as the maximum tolerance.

On average, the maximum tolerance for density was 50 tuples for SCP-common, 35 tuples for SCP-rotated, 30 tuples for SCP-staircase, and 45 tuples for PCP, suggesting that users are less frustrated with SCP-common and PCP when they are dealing with dense datasets.

Since reported densities in different studies may have different units, we must normalize to standard units (tuples/cm²) to allow for comparisons. The densities of 10, 20, 30, and 40 tuples in a cell of 49 cm² correspond to 0.20 tuples/cm², 0.41 tuples/cm², 0.61 tuples/cm², and 0.81 tuples/cm², respectively. In previous work [LMvW10], the densities used were 10, 40, 160 tuples displayed in a 24 cm × 26cm area, which is equivalent to 0.016 tuples/cm², 0.064 tuples/cm², and 0.26 tuples/cm², respectively. In comparison with Li *et al.*'s [LMvW10], our densities are higher in terms of tuples/cm², but lower in terms of total number of tuples. Because our experiment displays multiple plots to participants, we cannot have the same number of tuples as Li *et al.*'s [LMvW10] previous experiment without increasing their screen density even more, thus our chosen values are a compromise. For convenience, we use “tuples” instead of tuples/cm² to refer to density in the rest of the paper.

3.3. Dependent Variables

Two dependent variables, completion time and error distance, were used to measure user performance.

Completion time, in milliseconds, is measured from the appearance of the task stimuli on the screen to the moment the user hit a key to indicate that s/he has found the answer. Note that the time spent on typing in the exact numerical value is not counted in completion time. This is because we intended to prevent the input time from contaminating the raw result for value retrieval.

Error distance is measured as the absolute difference between the actual target value and participant's input. For example, if the actual value for a tuple on the target axis is 15, but the participant keys in 10, the error distance would be $Abs(10 - 15) = 5$. The smaller the error distance, the better the accuracy it is. We chose the continuous scale of error distance instead of a Boolean category of hit and miss to measure the errors, in favor of its added level of details.

3.4. Apparatus

Two iMac11,3 computers with 2.7GHz quad-core Intel Core i5 processors running on OS X Lion were used for the experiment. Each computer was equipped with a standard mouse




Figure 2: The stimuli used in the experiment. 1) Basic experimental information: trial number, time spent on the current trial, and task description; 2) the red \times indicating the value for the tuple of interest; 3) the highlighted target axis.

and keyboard. The display size was 27 inches (597.73mm by 336.22mm), 2560 by 1440 pixels, corresponding to a pixel pitch of 0.233mm. The experiment software was implemented in JavaScript with Protovis[†] and run in Firefox browser version 8.0.1 in full screen mode.

3.5. Stimuli

Figure 2 illustrates an example of stimuli used in the experiment. The top of the screen displays information about the current trial: number, time spent, and task description (e.g. for an N dimensional dataset, it shows "with the highlighted X_1 value, what's the corresponding X_N value?"). Just below this is the main experimental area in which the data and the visualization techniques are displayed.

Cell size: To fully utilize the screen estate while allowing the participants to simultaneously view the maximum number of dimensions without scrolling, each plot cell has a fixed length of 70mm, which translates to 300 pixels in our display configuration. This allows a maximum of 8 dimensions to be comfortably displayed (e.g. $300 \times 7 = 2100$ pixels for the 7 scatterplots + $50 \times 6 = 300$ pixels for the 6 visible gaps of 50 pixels each between adjacent scatterplots + spaces before and after the first and last scatterplot).

Tuple size and color: The data tuples in SCP are visualized using points of 4-pixel radius; for PCP, each data tuple is represented using a line of 1 pixel in width, both rendered with anti-aliasing. Based on our observation, these are the minimum data tuple sizes for participants to comfortably recognize under the current screen resolution. All data tuples are displayed in blue. All axes, numeric labels, and tick marks on the axes are in black. The value for the target data tuple is highlighted in red on the corresponding axis.

Stimuli generation: Data tuples are generated randomly with uniform distribution along each dimension according to the density requirement. The numeric values of all data

tuples are integers between 0 and 50. This range is fixed for all axes across all conditions and techniques so that it can serve as a constant. To avoid possible ambiguity of multiple data tuples having the same value as the highlighted tuple (in which case the users are unable to determine the tuple to trace from), when choosing the target tuple, we purposely avoid those with neighbors that are closer than 8 pixels or equivalently 1.9mm on all dimensions.

3.6. Procedure

Prior to the experiment, each participant was introduced to the visualization techniques and the value retrieval task. They were also instructed to finish the trials as quickly and accurately as possible while not using any visual aid (mouse cursor, finger, ruler, pen tip, etc.) other than their eyes. They were also informed that there is no ambiguity in the highlighted value.

A training session familiarized the participants with the techniques. They were instructed to continue practicing until they were fully comfortable with the value retrieval tasks with each technique before starting the main experiment.

For each trial in the main experiment, upon determination of the numeric value on the target axis, the participants were expected to hit the space bar, after which the timer is stopped and the visual stimuli is masked. The participants are then required to take their time to key in the numeric value in the provided input box. The visual stimuli were masked to prevent participants' visual residue from affecting their responses, which should not change after hitting the space bar.

Considering the switch between different techniques may result in relative longer response time to readapt, a pop-up window is shown whenever there is a change in techniques between the trials to remind the participants and to facilitate mental adjustment between different techniques. Upon finishing all the trials in the official session, the participants were invited to a brief interview to collect their subjective opinions. Their responses were audio recorded with their consent.

3.7. Result Analysis Method

Both experiments used the within-subject design involving three independent variables: technique, density, and dimensionality. Data were analyzed using factorial Repeated-Measures ANOVA, with significance level of $\alpha = .05$. Mauchly's test was used to verify the assumption of sphericity. Pairwise comparisons for the main effects of different variables were corrected using Bonferroni adjustments.

4. Experiment 1

This first experiment is to provide an overall understanding of the performance differences among the four techniques and to identify the winning techniques.

[†] <http://mbostock.github.com/protovis/>

Participants: 12 participants, 5 females and 7 males aged 20 to 25 years, from the university community, volunteered for the experiment. All participants had seen and used 2D SCP before, but none had experience with either PCPs or one of the SCP variants on multivariate data.

Experiment setup: Techniques were counterbalanced using balanced Latin Square. Participants were randomly assigned to four groups of three participants each.

For each technique, participants perform 3 trials in each of the three different data densities: 10, 20, and 30 tuples.

Within each technique and dimension combination, participants perform the trials in three different dimensions (2D, 4D, 6D). Presentation order of the dimensions and densities is both from easy to hard, (i.e., 2D, 4D, 6D for dimensions, and 10-tuple, 20-tuple, 30-tuple for densities) to allow participants to ease gradually to more difficult conditions. Note that since the main purpose of this experiment is to obtain an overall picture for the performance differences among the four techniques, we only counterbalanced the main factor, technique, in this first experiment.

After training, each participant performed the entire experiment in one sitting, including breaks, and post questionnaires in approximately 1 hour. In summary, the design was as follows (excluding trainings): 12 participants \times 4 visualization techniques (PCP, SCP-common, SCP-rotated, SCP-standard) \times 3 levels of data dimension (2D, 4D, 6D) \times 3 levels of data density (10 tuples, 20 tuples, 30 tuples) \times 3 repetitions of trails = 1296 trials in total.

4.1. Results

For experiment 1, we focus on revealing the overall performance for the four techniques. With regards to the main effect of the techniques, Mauchly's test verified the assumption of sphericity has been met in both error distance ($p = .119$) and completion time ($p = .057$) analysis.

Error Distance

Figure 3:(a) shows the average error distance of each technique. Repeated-measures ANOVA tests suggest that there is a significant main effect of the technique ($F_{(3,33)} = 22.34$, $p < .001$, $\eta^2 = .67$, observed power = 1.0).

For reporting the results in pairwise comparison among these four techniques, we use "[]" to enclose techniques with comparable performance ($p > .05$) and ">" to indicate the technique on the left of the operator is significantly better than the technique on the right side ($p < .05$). The relative accuracy performance relationship among the four technique is [PCP (0.98), SCP-common (2.7)] > [SCP-rotated (5.04)] > [SCP-staircase (7.31)].

Completion Time

Figure 3:(b) shows the average completion time with standard errors for the four techniques. Similar with the error




Figure 3: The average error distance (left) and completion time (right) with standard error bars among four techniques

distance, it shows that both PCP and SCP-common techniques are better than the SCP-rotate and SCP-staircase ($p < .05$).

Repeated-measures ANOVA tests suggest that the four techniques have significant difference in the completion time of tracing tuples across dimensions ($F_{(3,33)} = 27.83$, $p < .001$, $\eta^2 = .717$, with an observed power = 1.0). Post hoc tests further indicate the differences and ordering among the four techniques as [PCP (8.99s), SCP-common (12.02s)] > [SCP-rotated (18.58s), SCP-staircase (17.93s)].

Experiment 1 Summary

Comparing error distance and completion time among the four techniques, PCP and SCP-common are clearly the two better techniques. Both SCP-rotated and SCP-staircase are not suitable for value retrieval tasks, taking significantly longer time and are more error-prone.

Furthermore, the subjective feedback of both SCP-rotated and SCP-staircase is consistent with the quantitative results: 6 out of 12 participants ranked the SCP-staircase as the least preferred technique while the other half ranked SCP-rotated as the least preferred one. The reported reason for disliking SCP-staircase is the difficulty in tracing tuples across non-horizontal lines. The 45-degree tilted cells require the participant to "tilt the head to see (through imagined projection) the correct value". This is not only "more tiring", but also "more difficult to judge whether two points are on the same level". To many participants, such combined difficulties are so discouraging that they "gave up after a while".

While fatigue and perceptual difficulties caused by tilting are the main reasons for participants to dislike SCP-staircase, the difficulty in using SCP-rotated was reported to have a different reason. In SCP-rotated, to trace a tuple from one cell to another, it requires the following set of actions: find the target data tuple based on the value marked on the first axis, read the value of that tuple on the second axis, remember that value and locate that value on the same axis in the adjacent cell, and find the tuple in the adjacent cell with that value. As reported by one participant, "you have to always find and remember the value on the axis to move to the next cell (plot). This is too much work when the number of dimensions increases".




Figure 4: The average completion time among four techniques under varied data dimensions and densities.

While no significant overall performance differences are found between PCP and SCP-common, a further breakdown of the results (Figure 4) struck us with several interesting phenomena.

It is observed that in the 2D case, the performance difference between PCP and SCP-common is small. As the number of dimensions increases to 4, PCP seems to have advantages over SCP-common in all three densities. As the number of dimension increases to 6D, we found that PCP seems to have an advantage over SCP-common in the 10-tuple density case, but becomes inferior to SCP-common in the 30-tuple density case.

While PCP seems to have comparable overall performance with SCP-common, fine-grained investigation revealed that there are differences under different conditions. PCP seems to have advantages over SCP-common when density and dimension are low, but this advantage diminishes as dimension and density increase, indicating the strategy and cost for retrieving values for the two techniques are likely to be different.

5. Experiment 2

In experiment 1, we identified PCP and SCP-common as the two winning techniques for value retrieval. In experiment 2, we attempt to further investigate the influence of dimensionality and density on these two techniques. While not counterbalancing dimensionality and densities were less of a concern in experiment 1, proper counterbalancing is needed for both factors in this experiment as they become the focus of the study. Furthermore, in experiment 1, we learned that both techniques have similar performance in the 2D condition, but as the dimensionality and density increase, greater performance differences seem to emerge. This motivated us to use both higher dimensionality and density conditions in the second experiment.

Participants: 18 participants, 7 females and 11 males, aged between 20 to 30 years, from the university community, volunteered for the experiment. None had participated in experiment 1. All participants had seen and used 2D SCP before,

but none had experience with either PCPs or one of the SCP variants with multivariate data.

Experiment setup: Similar to experiment 1, a within-subject design was used. However, instead of only counterbalancing the technique, all three factors (technique, dimensionality, and density) are counterbalanced. The technique, with only two levels (PCP and SCP-common), is fully counterbalanced. The dimensionality and density both have three levels (4D, 6D, 8D for dimensionality and 20-tuple, 30-tuple, 40-tuple for density), were counterbalanced using Latin Square.

Combining the 2 techniques with 3 different order sequences in dimensions and with 3 different order sequences in density leads to 18 arrangements of the three factors ($2 \times 3 \times 3 = 18$). Participants were randomly assigned to one of the 18 experiment arrangements. For each of the technique, dimensionality, and density combination, participants were asked to perform 5 randomly generated trials.

The flow of the experiment procedure is exactly the same as experiment 1. Each experiment session took approximately 1 hour. The design of experiment 2 can be summarized as follows (excluding trainings):

18 participants \times 2 techniques (PCP, SCP-common) \times 3 dimensions (4D, 6D, 8D) \times 3 densities (20 tuples, 30 tuples, 40 tuples) \times 5 trials for each technique, dimension, density combination = 1620 trials in total.

5.1. Results

For experiment 2, we counterbalanced all three independent variables (e.g. technique, density, and dimensionality). Mauchly's tests verified that the assumption of sphericity have been met for the main effects and interaction effects of these variables we mentioned as follows ($p > .05$)[‡]. The observed power for all significant effects were above .80.

Error Distance

Overall, the repeated-measures ANOVA tests revealed no significant differences between techniques ($p = .436$). However, there were significant main effects in both dimensionality ($F_{(2,34)} = 6.124, p < .01, \eta^2 = .265$), and density ($F_{(2,34)} = 10.637, p < .001, \eta^2 = .385$).

Furthermore, Post-hoc comparison (Bonferroni correction) on dimensionality showed the ordering and differences among the dimension and density conditions to be [4D (1.44)] > [6D (2.85), 8D (2.48)] and [20 tuples (1.18)] > [30 tuples (2.59), 40 tuples (2.99)], respectively. These results are less surprising as the error distance is likely to increase as the dimensionality and density increase (as the task becomes more difficult).

[‡] Note that technique has only 2 levels, so the Mauchly's test does not apply.




Figure 5: The interaction effect for technique × density (left) and technique × dimension (right) in terms of error distance.

However, we found a number of significant interaction effects. There were a significant Technique × Density interaction ($F_{(2,34)} = 7.05, p < .01, \eta^2 = .293$), and a Technique × Dimension interaction ($F_{(2,34)} = 10.81, p < .001, \eta^2 = .389$). These interaction effects contain key information for us to reveal the relationship among these factors. Figure 5 shows the interaction effects for Technique × Density (left) and Technique × Dimensionality (right).

We found that for SCP-common, the error distance is relatively stable as dimensionality and density changes, but in PCP, the error distance dramatically increases as the dimension or density increases.

Completion Time

Overall, there are significant main effects on techniques ($F_{(1,17)} = 9.79, p < .01, \eta^2 = .365$), dimensionality ($F_{(2,34)} = 64.17, p < .001, \eta^2 = .791$), and density ($F_{(2,34)} = 42.98, p < .001, \eta^2 = .717$).

Post-hoc (Bonferroni correction) comparison on dimensionality and density finds the following relationship among different levels: [4D (13.62s) > [6D (18.37s), 8D (19.30s)]] for dimensionality and [20 tuples (13.50s)] > [30 tuples (17.50s)] > [40 tuples (20.29s)] for density. Just like the observations we made with error distance, the significant effects found in dimensionality and density are expected as the completion time is likely to increase as the dimensionality and density increase.

However, the significant effect found in technique is somewhat surprising as it differs from what we got from experiment 1. In experiment 1, we found that the completion time is comparable ($p > .05$) between the two techniques with PCP (8.99s) being slightly quicker than that of SCP (12.02s), but experiment 2 tells an almost opposite story, as PCP-common is significantly slower than SCP-common. To understand the reason behind this phenomenon, we need to further analyze the interaction effects below.

Similar to the results found with error distance, we found two significant interaction effects. There were a significant Technique × Density interaction ($F_{(2,34)} = 8.74, p < .01, \eta^2$




Figure 6: The interaction effect for technique × density (left) and technique × dimension (right) in terms of completion time.




Figure 7: The average completion time for SCP-common and PCP under varied dimensions and densities.

= .340), and a Technique × Dimension interaction ($F_{(2,34)} = 73.46, p < .001, \eta^2 = .812$). Figure 6 shows the interaction effects for Technique × Density (left) and Technique × Dimensionality (right).

We found that the increase of dimensionality has almost no effect on SCP-common, but causes the significant performance degradation to that of PCP. On the other hand, the Technique x Density interaction showed that both PCP and SCP-common are affected by increased density. However, the increase in density seems to cause more damage to PCP than that of SCP-common (i.e., at the density of 20 tuples, PCP has almost equal performance with SCP-common, but when the density is increased to 30 or 40 tuples, PCP is much slower than SCP-common, and the performance gap between the techniques increases with number of dimension).

This effect is further elaborated in Figure 7, in which the effects of all three factors on completion time are simultaneously displayed. Overall, it shows that PCP has advantages over SCP-common when dimension and density are low.

Under each particular (density, dimensionality) condition, we applied Pairwise T-test to compare these two techniques. The results show the advantage of PCP over SCP-common in two low dimension and density cases (4D, 20 tuples; 4D, 30 tuples) (both $p < .05$). A single step increment in ei-

ther dimension or density renders PCP comparable to SCP-common, as proven by pairwise T-test in these two cases (6D, 20 tuples; 4D, 40 tuples) (both $p > .05$). Finally, further increase in either dimension or density will make PCP inferior to SCP-common, as demonstrated by pairwise T-test on the rest of the 5 conditions. (8D, 20 tuples; 6D, 30 tuples; 8D, 30 tuples; 6D, 40 tuples; 8D, 40 tuples) (all $p < .05$).

6. Discussion

Experiments 1 and 2 revealed the following relationships between the four techniques:

1. SCP-rotated and SCP-staircase yielded poor performance and users found them difficult to use. This seems to be because tilting the axes 45° makes the task more difficult, and requiring users to remember the value from axis to axis also increases difficulty.
2. PCP and SCP-common performed better and were preferred by participants. However, these two techniques seem suited for different scenarios: PCP is better at low dimensionality and low density, and SCP-common is better when these are higher.
3. The performance of PCP is dependent on dimensionality, while the performance of SCP-common seems roughly independent of dimensionality.
4. Increasing density affects the performance of PCP more than it affects SCP-common.

We now offer theoretical explanations of the observed differences between PCP and SCP, partly to guide future design of visualization techniques.

User strategy for value retrieval

To inspect the values of a data tuple across multiple dimensions, one needs to trace it from one cell to another. There are different hypothetical strategies which users may use (Figure 8):

- **"Remember-value"**: The user memorizes the position or numeric value along the axis common to the two cells, which is invariant to the arrangement and alignment of cells.
- **"Count-point"**: The user memorizes the ordinal position of the tuple within a local interval, e.g. a tuple can be identified as "the tuple with second largest X2 attribute value, among all those having X2 values between 0 and 10".
- **"Trace-line"**: This strategy can be used with SCP-common: the user imagines the horizontal line passing through all the points of a tuple, and follows this imaginary line to the point above the target horizontal axis. (Note that SCP-staircase also allows tracing along imaginary lines perpendicular to the shared axes, but this must be repeated for each pair of adjacent scatterplots, rather than done once globally as in SCP-common.)

Actually, users' choice of strategy may be affected by the visualization technique and the specific instance of the trial.




Figure 8: The possible strategies for users to perform when tracing tuples across dimensions

With SCP-common, when data points are sparse, users may trace along an imaginary line, because "it is easier than estimating and remembering the value on the axis", given a sparse neighborhood around the point of question. With increased density, however, tracing along an imaginary line surrounded by many distracters can become difficult. Especially, one participant commented that "without grid lines, the virtual line is quite misleading when there are many neighbors". In such case, the user may prefer one of the other two strategies, neither of which should be hindered by increase in dimensionality. Indeed, for SCP-common, the experimental results found that dimension has no significant effect on completion time and error distance, for a dataset denser than 20 tuples.

For PCP, the only strategy we can think of is to visually follow the polygonal line representing a tuple. Since this operation takes more time with increased dimensionality, we expect dimension should have an effect on performance with PCP, and this is indeed found in our experimental results.

Clutter problem in SCP and PCP

The most fundamental difference between SCP and PCP is the tuple's visual representation: points (SCP) versus a polygonal line (PCP) (or, in some variants of PCP, a smooth curve [HvW10]). The tradeoff between a point and a line may explain why performance with PCP is more sensitive to density than SCP. Points are more space-efficient than lines: adding more points introduces less clutter than adding more polygonal lines.

However, when the screen is not cluttered, a line that intersects with the associated axis allows the user to directly read the numerical value, without the need to visually project the data tuple to the axis through imagination. At low density, tracing along visual lines (as in PCP) may be easier than tracing along an imaginary line or memorizing positions or numeric values (as in SCP). Therefore, tracing tuples across dimensions will be easier with PCP as compared to SCP when




Figure 9: Examples of PCP and SCP with 10 tuples (top) and 40 tuples (bottom). It is apparent that clutter increases faster in PCP.

the screen is not cluttered. Figure 9 shows the examples of equal-density dataset with both PCP and SCP, in which it can be seen that, when the number of tuples increases from 10 to 40, the increased difficulty of task for PCP is clearly greater than for SCP.

Guideline for user

Based on the experimental results, we have come up with a table that can guide users in choosing which visualization technique to use for value retrieving task in respect to dimensionalities and densities.

Figure 10 shows the recommendation of techniques for value retrieval under varied dimensionality and densities. Cells with "PCP" or "SCP-common" means PCP or SCP-common has significant better performance. Cells with "~" means PCP and SCP-common have comparable performance. And, cells with question mark is the condition which we have not covered in this study.

It can be seen that, PCP is recommended for cells in top left corner, which represent multivariate data with lower dimensionalities and densities. The cells in the bottom right corner represent multivariate data with high dimensionality and densities, and SCP-common is preferred. Cells on the diagonal line can use either of the two approaches, which offers users a choice depending on other considerations.

7. Limitation and Future Work

For practical reasons of experimental design, the testing conditions in our study only involved datasets with relatively low dimensions and density. In practice, visual analysts often face datasets presented with much higher on-screen density and dimensions. Future studies may want to further validate

Dimension \ Density	2D	4D	6D	8D
10 tuples	~	PCP	PCP	?
20 tuples	PCP	PCP	~	~
30 tuples	~	~	SCP-common	SCP-common
40 tuples	?	~	SCP-common	SCP-common

Figure 10: The recommendation of techniques based on experimental results. The dotted red rectangle highlights the conditions in experiment 1; the solid green rectangle highlights the conditions in experiment 2.

our experimental results with such scenarios. In addition, many possible PCP and SCP variants have been proposed in the literature in which our study has only investigated a few. Future studies can involve other interesting variants, such as the Radar plot [CCKT83], to further our investigation. Lastly, PCP and SCP are only two of the vast number of visualization techniques presented in the literature. The value retrieval task is also one of the many visual analytical tasks. The InfoVis research community has a long way to go before being able to fully understand the design tradeoffs of the different visualization techniques in different tasks.

8. Conclusion

In this paper, two controlled experiments compared user performance in value retrieval tasks between four visualization techniques: three SCP variants (SCP-common, SCP-rotated, and SCP-staircase) and the baseline PCP, while varying dimensionality and data density. Results indicate PCP and SCP-common outperform the other two techniques. Furthermore, PCP shows advantages in low dimensionality and low density dataset, while SCP-common outperforms PCP in higher dimensionality and density dataset. We also proposed a guideline for choosing a technique based on the dataset properties. This is the first study we know of that empirically compares PCP and SCP for this task, and also the first study that has found an advantage for PCP over SCP for any task in any conditions. We believe the experimental results, the analysis and reasoning we formulated on the observed phenomena, and the proposed guideline of usage can be valuable for both researchers and practitioners to better understand and utilize PCP and SCP for more effective information visualization.

9. Acknowledgments

This research is supported by the National University of Singapore Academic Research Fund R-252-000-375-133 and by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the IDM Programme Office.

References

- [AdOL04] ARTERO A., DE OLIVEIRA M., LEVKOWITZ H.: Uncovering clusters in crowded parallel coordinates visualizations. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on (0-0 2004)*, pp. 81–88. 2
- [AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *Proceedings of IEEE Symposium on Information Visualization (InfoVis)* (2005), pp. 111–117. 1, 2
- [AR11] AZHAR S., RISSANEN M.: Evaluation of parallel coordinates for interactive alarm filtering. In *Information Visualisation (IV), 2011 15th International Conference on* (july 2011), pp. 102–109. 1
- [CCKT83] CHAMBERS J. M., CLEVELAND W. S., KLEINER B., TUKEY P. A.: *Graphical Methods for Data Analysis*. Wadsworth & Brooks/Cole Publishing Company, New Jersey., 1983. 9
- [CvW11] CLAESSEN J. H. T., VAN WIJK J. J.: Flexible linked axes for multivariate data visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 17, 12 (2011), 2310–2316. 2
- [Har75] HARTIGAN J. A.: Printer graphics for clustering. *Journal of Statistical Computation and Simulation* 4, 3 (1975), 187–213. 1, 2
- [HLKW12] HEINRICH J., LUO Y., KIRKPATRICK A. E., WEISKOPF D.: Evaluation of a bundling technique for parallel coordinates. In *Proceedings of International Conference on Information Visualization Theory and Applications* (2012). 2
- [HvW10] HOLTEN D., VAN WIJK J. J.: Evaluation of cluster identification performance for different PCP variants. In *Proceedings of Eurographics/IEEE-VGTC Symposium on Visualization (EuroVis)* (2010). 1, 2, 8
- [Ins85] INSELBERG A.: The plane with parallel coordinates. *Visual Computer* 1 (1985), 69–91. 1, 2
- [KD09] KINCAID R., DEJGAARD K.: Massvis: Visual analysis of protein complexes using mass spectrometry. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on* (oct. 2009), pp. 163–170. 1
- [LMvW10] LI J., MARTENS J.-B., VAN WIJK J. J.: Judging correlation from scatterplots and parallel coordinate plots. *Information Visualization* 9 (2010), 13–30. 1, 2, 3
- [QCX*07] QU H., CHAN W.-Y., XU A., CHUNG K.-L., LAU K.-H., GUO P.: Visual analysis of the air pollution problem in Hong Kong. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 13, 6 (2007), 1408–1415. 2
- [Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of IEEE Symposium on Visual Languages (VL)* (1996), pp. 336–343. 1
- [SS05] SEO J., SHNEIDERMAN B.: A knowledge integration framework for information visualization. In *From Integrated Publication and Information Systems to Information and Knowledge Environments*, Hemmje M., Niederée C., Risse T., (Eds.), vol. 3379 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2005, pp. 207–220. 1
- [TGS04] TYMAN J., GRUETZMACHER G., STASKO J.: Infovisexplorer. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on* (oct. 2004), p. r7. 1
- [The00] THEISEL H.: Higher order parallel coordinates. In *Proc. Vision, Modeling and Visualization (VMV)* (Saarbrücken, 2000), Girod B., Greiner G., Niemann H., (Ed.) H.-P. S., (Eds.), pp. 119–125. 2
- [VMCJ10] VIAU C., MCGUFFIN M. J., CHIRICOTA Y., JURISICA I.: The FlowVizMenu and parallel scatterplot matrix: Hybrid multidimensional visualizations for network exploration. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 16, 6 (2010), 1100–1108. 2, 3
- [WB97] WONG P. C., BERGERON R. D.: 30 years of multidimensional multivariate visualization, 1997. Chapter 1 (pp. 3–33) of Gregory M. Nielson, Hans Hagen, and Heinrich Müller, editors, *Scientific Visualization: Overviews, Methodologies, and Techniques*, IEEE Computer Society. 1
- [Weg90] WEGMAN E. J.: Hyperdimensional data analysis using parallel coordinates. *J. of the American Statistical Association* 85, 411 (1990), 664–675. 1
- [YGX*09] YUAN X., GUO P., XIAO H., ZHOU H., QU H.: Scattering points in parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 15, 6 (2009), 1001–1008. 2