Mark M. Shovman, Kenneth C. Scott-Brown, Andrea Szymkowiak & James Bown

The Meaning of Graphs: Scatterplots as a Test-Bed for Theories of Object-Hood

Introduction

In natural and social sciences, novel insights are often derived from analysis of data using various visualisation techniques. But what principles underpin the extraction of meaningful content from these visualisations? Abstract data visualisation can be traced at least as far back as William Playfair's work (1801), with the introduction of time-series plots, bar charts and pie charts. Tukey (1977) predicted the change toward the use of computers to draw graphs and advocated the use of graphs in exploratory data analysis. The continued use and development of these and other basic data visualisation techniques has led to a recognised set of heuristic 'best-practice' guidelines (e.g., Bigwood & Spore, 2003). Recently, however, the quantity and complexity of data that require analysis has dramatically increased, making these standard tools and techniques inadequate for the task (Thomas & Cook, 2005). For example, a single typical bar chart visualises only one continuous and two discrete variables. Although, the ubiquity of computing power now enables a vast range of novel visualisations that are rich, multimodal and interactive, the most effective way to exploit this power to support analysis of large, complex data sets is unclear (Chen, 2004).

A recent area of research, Visual Analytics, addresses these issues, focusing on analytical reasoning facilitated by interactive visual interfaces (Thomas & Cook, 2005). In a NIH/NSF report (Johnson et al., 2006) Visual Analytics is described as a vibrant field, drawing inspiration and manpower from mathematics, computer science, arts and psychology and serving many areas of research, industry, finance and defence. Fundamental research in Visual Analytics examines the cognitive mechanisms involved in perception, encoding, memorising, and retrieval of visualisations (Johnson, et al., 2006). There is a substantial amount of data on perception of standard two-dimensional (2D) charts (for a review see Shah & Hoeffner, 2002), and many findings from other areas of fundamental research in cognition are successfully applied to building efficient presentational visualisations (Kosslyn, 2006). However, research exploring the cognitive mechanisms involved in perceiving animated, interactive and three-dimensional visualisations is less advanced (but see Polys, Kim, & Bowman, 2007; Ware & Mitchell,

GESTALT THEORY
© 2012 (ISSN 0170-057 X)
Vol. 34, No.3/4, 381-398

2005). Further fundamental research in comprehension of these stimuli is vital to development of efficient novel Visual Analytics tools and techniques.

Conversely, research in perception and comprehension of data visualisations, can inform study of perception and comprehension. Abstract data visualisations such as bar-charts or scatter-plots are a useful test-bed for object-hood theories, for several reasons. First, visual artefacts in written media and visual analytics are developed and constrained by the visual cognitive capacity of the users and the physical constraints of the media they are presented on. This is in contrast to the objects in nature or the man-made objects we use in modern life, where, to varying degrees, the appearance is constrained and determined by these objects' function, origin, or manufacture considerations. Thus, in an abstract visualisation, the notion of perceptual object is clearly dissociated from that of a physical object, making focusing on the specific features of perceptual objects that are not related to physical properties of real-world objects easier. Second, visual analytics provides stimuli that are rich and yet easily controlled; importantly, systematic modulation of an abstract dataset is much easier than comparable modulation of a natural scene. Third, abstract data visualisation formats such as bar- or pie-charts are an established part of contemporary visual parlance, and thus offer a level of ecological validity higher than the tightly-controlled stimuli and tasks used in many psychophysical experiments. Another advantage of visualisations over conventional psychophysical stimuli such as Gabor patches (Gabor, 1946) is that the visual analysis tasks give themselves more readily to meaningful interpretation, again, due to them being part of the cultural baggage. That answers the call from Koenderink (2010) that highlights the importance of a meaningful task or 'question' (p37) in the research of seeing.

Formally, every visualisation is a mapping from data space into visualisation space (Mackinlay, 1986). Thus, for example, in a yearly precipitation chart, date and precipitation levels are mapped onto horizontal and vertical axes of the chart respectively. The same mapping translates the tasks the analyst must undertake in the data domain to visual tasks: finding the rainiest time of the year is translated to finding a region of the chart with the largest area under it. Likewise, unexpected findings are first perceived as visual objects and then translated back into the data domain: a single sharp peak in a low flat region is seen and then interpreted either as a sudden shower in the middle of a dry season – or an error in data collection. Thus the meaning of the data can be directly related to the structure of the visual scene.

Analysis of visualizations in terms of visual elements that carry meaning is a promising approach: it is generalisable between subject domains; and it can benefit from the wealth of research in visual perception, specifically that regarding the perceptual parsing of the scene (Kubovy & Pomerantz, 1981). For instance, previous findings (Shovman, Szymkowiak, Bown, & Scott-Brown, 2009) indicate that the majority of exploratory data analysis tasks, when expressed in visual terms, consist of either: outlier detection, clustering (group detection) or trend detection. Outliers, clusters, and trends are thus the results of perceptual parsing of the scene, carrying meaning that has no immediate physical analogue. In that sense visual analytics is akin to reading; however it has more visual versatility and relies much more on spatial organisation of constituent elements.

Our previous research (Shovman, et al., 2009) explored outlier detection in a 3D scatterplot. The study a used 'pop-out' paradigm, described by Treisman and Gelade (1980) and since then applied to the study of Gestalts by Pomerantz (2006). It is based on a 'pop-out effect' which can be observed in tasks that consist of detecting an odd-one-out object in an array of distracting objects. When the number of distracting objects is varied, target detection time either remains constant (that being the 'pop-out') or increases linearly with the number of distracters. This distinction is dichotomous and is associated with, respectively, either implicit immediate perception or an explicit and laborious process of serial search. The 'pop-out' effect has been recently connected with the formation of high-level perceptual objects (Hochstein & Ahissar, 2002), making it an appropriate tool to research formation of meaning and insights in abstract visualisations.

One of the more common tasks accomplished using scatterplots (Shovman, et al., 2009) is trend detection. This poses the theoretical question: is a trend, then, a perceptual object in its own right? Or, more generally: what are the perceptual objects in a scatterplot? Previous research (Shovman, Szymkowiak, Bown, & Scott-Brown, 2008) suggests that increasing number of points in a 3D scatterplot outlier detection task leads to an increase in search time. Thus according to Feature Integration Theory (Treisman & Gelade, 1980) each and every point is a complex perceptual object. An alternative prediction emerges from Reverse Hierarchy Theory (Hochstein & Ahissar, 2002). In this view, bottom-up scene parsing is guided by task, and thus results in the identification of objects with the highest degree of task relevance; therefore a perceptual object will be a trend of points rather than a single point.

These two theoretical explanations generate conflicting predictions for a 3D trend detection task. Based on the evidence of visual search in 3D scene perception both theories would predict that trends in a 3D scatterplot will not be immediately perceived. However, the empirical prediction following from FIT would be that increasing the number of points will increase the response latency. Conversely, Reverse Hierarchy Theory would predict that increasing the number of point groups would increase the difficulty of the task, and would make no prediction of the effect of increasing the number of points. An experiment was designed to dissociate between these theoretical accounts. Mainly, if the results would support the RHT account, it would indicate that pop-out methodology is a valid tool to assess formation of high-level semantics (i.e. meaning) in abstract visualisations.

Method

The experimental task consists of identifying a non-random positioning of points in a 3D scatterplot. A stimulus scene comprises several groups of points, one being the target group and all others being the distracters. In the target group there is a non-random (i.e. structured) relationship between its constituent points' positions, while points in the distracter groups are positioned randomly. To ensure that the task is not reducible to planar trend detection, the trend in the target group is made inherently three-dimensional. The groups are presented superimposed in space, and distinguished by a non-spatial parameter, namely colour.

The paradigm incorporates interactively rotating stimulus sets that are a combination of a variable number of data sets with a variable number of points in each set. The sets are defined by colour to enable multiple choice coding and responses by the participants. The design used is a within-subject design with five levels of number of points per group, three levels of number of groups, and two levels of depth displacement noise added to the target group (see Stimuli below). The measured dependent variable is the net interaction time for correct answers.

1. Participants

Participants were recruited by convenience sampling and were paid for their time. Informed consent was obtained from all participants. Overall 16 participants (12 male and 4 female) whose ages ranged from 17 to 41 with median at 20 were recruited. Nine participants were students in the School of Computers and Creative Technologies, three from the School of Social and Health Sciences, one from the Dundee Business School and two were members of staff of the Abertay University. Two participants were later excluded from the analysis: one because the average response time for that participant was 3.21 standard deviations slower than the mean; another participant showed evidence of undiagnosed blue-yellow colour blindness.

2. Apparatus

The experiment was run on a Hewlett-Packard PC, (2.4GHz processor & 2Gb RAM). Graphics were rendered with NVidia GeForce 8800 GTS card on a 19" LG Flatron L226WTQ LCD display. The experimental program was custom written in Java[1] using the Java3D[2] package for 3D rendering. Viewing was performed through natural pupils or with participant's optical prescription suffi-

[1] http://www.java.com
[2] https://java3d.dev.java.net/

cient to give normal vision. To improve ecological validity and generalisability of the research results onto common visual analytics practices, experiments were run in a quiet experimental chamber under normal daylight office lighting conditions, using normal office hardware, with the default colour and 3D settings for monitor and PC. Mild head restraint in the form of a chin rest and forehead bar, coupled with the use of a fixed chair ensured constant viewing distance of 57 cm was maintained during data collection. Stimuli generation and subsequent data analysis of mouse logs was achieved using Matlab[3] and SPSS[4].

3. Procedure

The task of the observer in each trial was to identify the colour of the group of points that were positioned non-randomly. A 4AFC paradigm[5] required the observer to indicate for each trial, the colour of the cloud of data points arranged in a structured shape. Keyboard keys 1, 2, 3 and 4 (situated on the top row of the keyboard) had coloured stickers added to them: red, green, blue, and yellow respectively. Before starting with the main experiment task, participants undertook a short training exercise to become familiar with it. The training comprised nine trials of increasing complexity. The first training trial consisted of identifying a square grid structure within a single distracting noise group. More noise groups and more complex 3D structures were gradually added, with the last three trials being of the same complexity as the main task. Participants could repeat the training if they wished.

The experiment itself comprised 120 self-paced trials presented in random order. After every 30 trials, a pause screen was displayed for at least 30 seconds. The pause screen was gray, with black text in 24pt font presented in the middle. For the first 30 seconds, the text read “Block X/4. Please rest your eyes for at least 30 seconds”. During these 30 seconds no input from either keyboard or mouse was processed. After 30 seconds, the text changed to “Please press any key to continue” and any input from the keyboard would close the pause screen and advance the main test to the next trial. During the briefing, participants were specifically encouraged to take a break and rest their eyes between the sets. The duration of the experiment, including explanations and training, was approximately 30 minutes for each participant.

The user interface in this experiment consisted of a button positioned on a horizontal bar, in the bottom of the screen. Clicking the mouse on the button and sliding it along the horizontal bar rotated the scene around the vertical axis. To

[3] http://www.mathworks.com/products/matlab/
[4] http://www.spss.com/
[5] In a four alternative forced choice paradigm (4AFC) a participant is given four response options in every trial; the absence of null response option (e.g. ‘I do not know’) requires the participant to make a perceptual decision every trial.

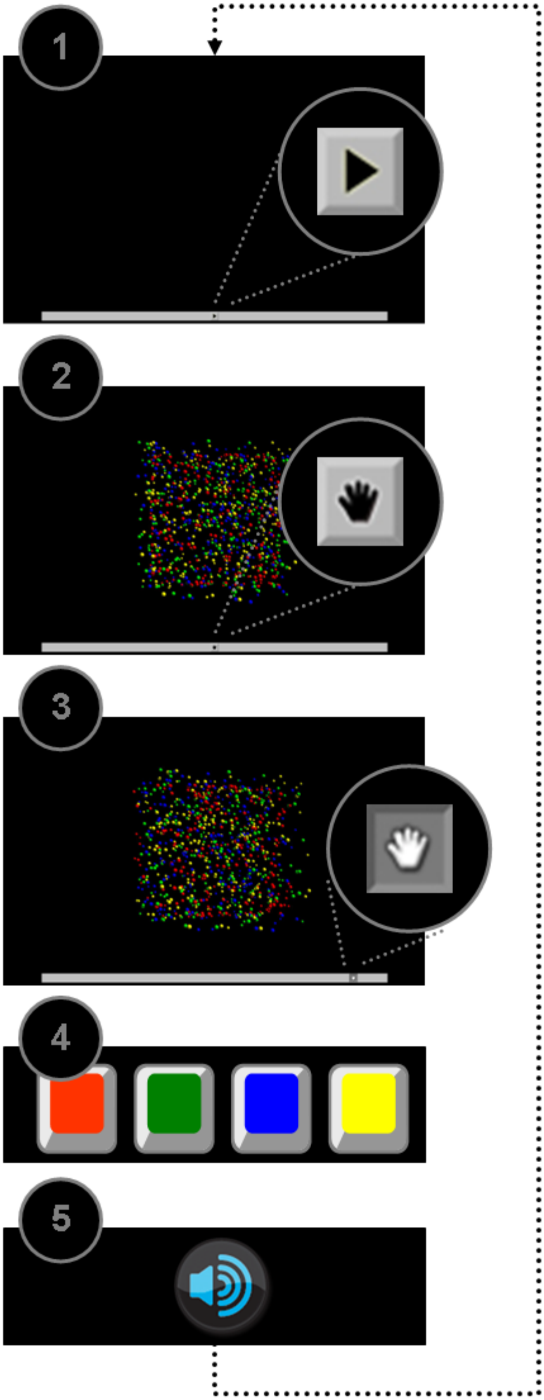

Fig. 1 – single trial timeline. 1: waiting for the participant to press the 'next trial' button; no stimulus is presented. 2: stimulus presented in a frontal view. 3: interaction phase; participant is rotating the view. 4: response: participant presses one of four coloured keys on the keyboard. 5: auditory feedback; harmonious sound for correct response, or discordant sound for an incorrect one.

indicate the participant's response and thus finish a trial, participants were instructed to press one of the four colour-coded buttons (keys 1, 2, 3 and 4 on the top row of keyboard). This response was time-stamped and recorded together with all the mouse movements and actions (see Figure 1 for the timeline of a trial).

To preserve attention and minimise the possibility of lapses of attention during a repetitive task (Norman & Shallice, 2000) over a period of time, performance feedback was incorporated to responses. Auditory stimuli easily elicit emotional responses (LeDoux, 1996) and, since the main experimental task is purely visuomotor, an auditory feedback can be expected to only minimally interfere with it. Thus, two short sound samples, one positive (harmonious tone) and one negative (discordant tone), were prepared and used as feedback throughout all the experiments. They were sounded after each trial, indicating successful or unsuccessful performance. A recording of a training session can be seen at the website.

4. Stimuli

The stimulus in a trial consisted of a composition of two, three, or four datasets displayed simultaneously. The datasets were presented in the centre of the screen, but in different colours. One was the target data set, and the rest were distracter data sets. In each trial, all sets had the same number of points, however in different trials the number of points in each set varied from 64 to 100, 144, 196 and 256 points (i.e. $4n^2$ for n in the range of 4 to 8), for a linear increase in the planar density of points. The manipulation of the number of datasets and number of points per set was fully balanced in the experimental design, resulting in 15 options for total number of data points in a stimulus, from 128 (2 sets of 64 points each) to 1024 (4 sets of 256 points each). The angular subtense of the stimulus array was 18 degrees of visual angle.

In distracter datasets, the spatial position of each point was randomly drawn from a uniform distribution, within a cubic volume. In target datasets, planar position of points was randomly generated in the same way as in the distracter sets, while the depth position of each point was defined by a formula $z = \boldsymbol{a} \cdot max(x^4 - \boldsymbol{b}x^2, y^4 - \boldsymbol{b}y^2)$. In half of the target sets, $\boldsymbol{a}$ was 1; in the other half $\boldsymbol{a}$ was -1; likewise, $\boldsymbol{b}$ was either ½ or ¾ (see Figure 2). In half of the sets a random 20% noise factor was added to the depth position. After that the data point depths were adjusted so that the dataset would fit the same cubic volume as the distracter sets. The combination of two **a** values, two **b** values and two noise conditions gave eight different types of target data sets, which, combined with five data set sizes resulted in 40 different target sets overall.

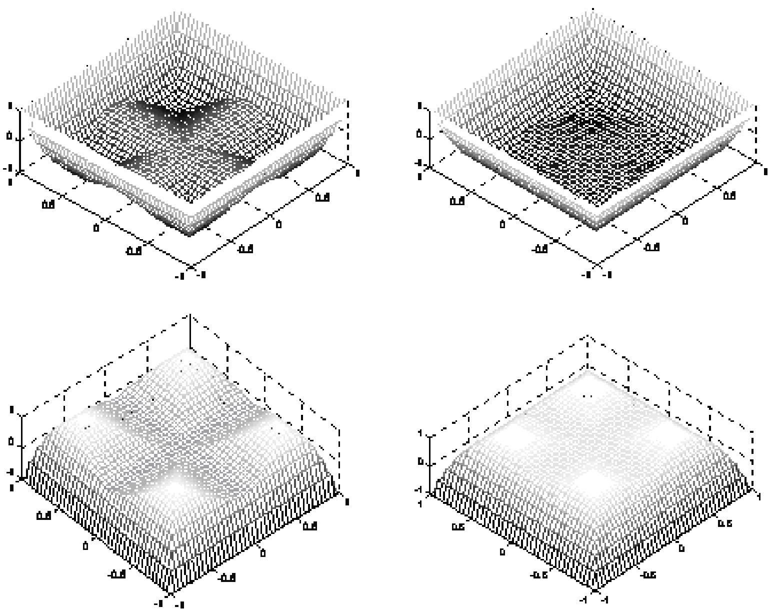

Fig. 2 - 3D structures used as targets in the experiment. Vertical dimension in the figure corresponds to depth dimension in the experiment, 'up' being 'closer to the viewer'. Top row: positive ***a*** coefficient, concave structures; bottom row: negative ***a*** coefficient, convex structure. Left column: ***b*** coefficient ½, convoluted structure; right column: ***b*** coefficient ¾, smoother structure.

The size of each point was randomly assigned from a uniform distribution, subtending between 0.1 and 0.4 degrees of visual angle when rendered at zero depth. The choice of colour for the data sets, as well as the choice of specific distracter datasets, was random for each run, aiming to reduce the irrelevant biases caused by specific colours or distracter datasets. The colours were system-default red, green, blue and yellow. A sample of resulting stimuli (without rotation) is shown in Figure 3.

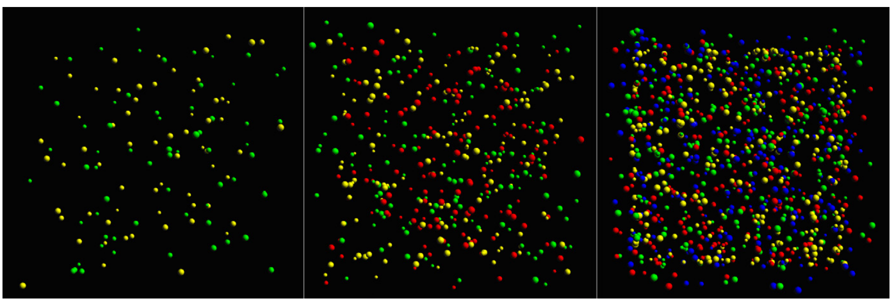

Fig. 3 – sample stimulus arrays. Left panel: two datasets of 64 points. Middle panel: three datasets of 144 points. Right panel: four datasets of 256 points.

5. Mouse Log Analysis

All mouse and keyboard responses in the experiment were time-stamped and logged. Post-experiment, these logs were processed to produce more detailed parsing of trial time course. Trial time was divided into three stages: before the first rotation; from the first to the last rotation; and from the last rotation till the keyboard press that was the response. The rotation time was further refined by taking into account only the time actually spent rotating the scene, i.e. dragging the button icon along the length of the slider bar, and discarding the times when no rotation occurred. The sum of dragging periods resulted in a measure of net rotation time. In addition, mouse movements were parsed into side-to-side sweeps, and for each sweep, time and range were calculated. This resulted in two derived variables for each trial: average rotation range and average rotation speed.

Results

Participants' accuracy ranged from 40.8% to almost perfect performance at 97.5%. The chance rate in this experiment design is as follows: a third of the trials are 2AFC, a third 3AFC, and a third 4AFC; thus chance rate is (1/2+ 1/3 + 1/4)/3 = 36.1%. Thus all participants performed above chance level. There was also evidence of improvement in accuracy during the experiment: every trial is on average 0.14% more accurate than the previous one. For response time analyses, only the trials with correct responses were included. In 66 trials (3.9% of all trials), a correct response was made with no rotation of the stimulus. To ensure that main analyses were performed only on the results of tasks that were completed using 3D scene reconstruction, all the trials with zero interaction times were excluded.

Usually, experiments of pop-out effect (Hershler & Hochstein, 2005; Hochstein & Ahissar, 2002; Pomerantz, 2006; Pomerantz, Sager, & Stoever, 1977; Treisman, 1982; Treisman & Gelade, 1980) analyse raw response times. However, several studies have demonstrated that response times are not distributed normally (Ratcliff, 1979; Zaitsev & Skorik, 2002). Characteristic features of human response time distributions are positive-only values, high negative skew, cut-off to the left and a long trail of outliers to the right of the mean (Zaitsev & Skorik, 2002). This violates the normalcy assumption of most standard statistical analyses (Langdridge & Hagger-Johnson, 2009), specifically of the linear regression used in the pop-out effect paradigm. Several theoretical distributions that have these characteristic features have been proposed (e.g., Zaitsev & Skorik, 2002), the simplest of them being a log-normal distribution (Limpert, Stahel, & Abbt, 2001). In the analyses of the experimental data below, the distribution of response times fits the lognormal distribution better than the normal distribution, and so a natural logarithm link function was applied to the data prior to subsequent analysis.

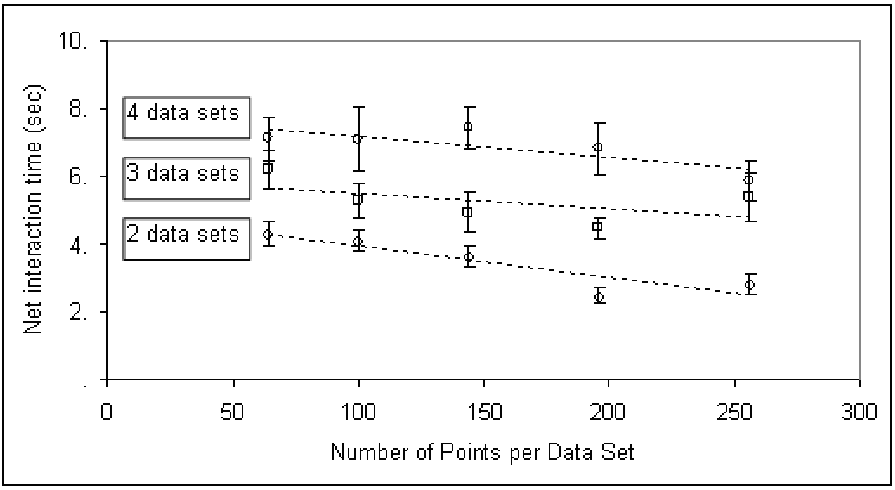

Fig. 4 – summary graph showing net interaction time, in seconds, as a function of number of points per stimulus data set, for the three conditions (two, three, or four data sets). Error bars at 1 SEM.

A summary of the relation between net interaction time, number of data point sets, and number of data points per set is shown in Figure 4. The analysis was conducted using a multiple linear regression model with criterion of net interaction time and two key predictors: number of data sets and number of points per data set. The analysis (R^2=0.14; F(2,1018)=84.2; p<0.001) has shown that both main predictors are significantly associated with the interaction time. Interaction time increases with more data sets (b=0.34; β=0.34; p<0.001), but decreases with more data points per set (b= -0.002; β=-0.17; p<0.001).

In order to account for remaining variability in the data, an additional linear regression was run, in which the model allowed stepwise inclusion of other parameters that could be expected to affect interaction time. These parameters were: coefficients ***a*** and ***b*** of the target data set formula; the amount of depth noise added to the target data; serial position of the trial in the experimental sequence; and two variables derived from mouse log analysis: average mouse sweep time, and average mouse sweep range. Similarly to the distribution of net interaction time, average mouse sweep time is better fitted by a lognormal distribution; therefore a logarithm link function was used.

Contrary to predictions, coefficient ***a*** (concavity or convexity of the target structure) does not show a significant effect on interaction time. However, all the other parameters passed the inclusion criteria of the stepwise algorithm. A summary of the resulting regression analysis (R^2=0.48; F(7,1013)=133.3; p<0.001) is presented in Table 1.

Predictor	b	β	p
Number of data points per data set	0.000	-0.069	0.003
Number of data sets	0.201	0.200	<0.001
ln(Average mouse sweep time)	0.855	0.525	<0.001
Trial serial position order	-0.002	-0.092	<0.001
Average mouse sweep range	0.000	0.090	<0.001
Amount of depth noise	0.012	0.076	0.001
Target structure coefficient ***b***	0.379	0.058	0.011

Table 1 – linear regression coefficients for the augmented main analysis. b: unstandardized coefficient; β: standardized coefficient; p: significance.

In this enhanced analysis, the relation of the interaction time to the key parameters: number of data sets and the number of data points per set, remains significant. In addition, there are positive associations of interaction time with the amount of depth noise; with coefficient b (smoothness) of the target structure; and with both mouse sweep time and mouse sweep range. The associations of interaction time with trial presentation order is negative: subsequent trials take less time.

Discussion

This paper focuses the central question of this special issue, namely "What is a perceptual object?" on the area of Visual Analytics, an area where it is less frequently asked. The experiment presented shows how the context of graph perception offers a unique pathway to the view of meaning in abstract structures, and, more generally, to the structure and form of human perceptions in situations where the perceptual behaviour is natural and purposive, yet the sensory environment is relatively abstract and removed from the natural world and the environments in which humans evolved to see and act. This separation between the physical world and the search behaviours associated with it gives a highly specialised paradigm space to address questions of interest to form, shape structure and object-hood since the meaning of a pattern can be manipulated with a data structure. The issues at stake addressed in this paper can be subdivided as follows: (a) how useful can visual analytics stimuli and tasks be in the conduct of rigorous experimentation in the area of interplay of visual perception, semantics, interactivity, etc. and (b) although pop-out is not a low-level perception tool, can it be used to objectively assess high-level semantic constructs?

The current experiment was aimed at deciding between two theoretical models of hierarchical perception, the Feature Integration Theory (FIT) (Treisman & Gelade, 1980) and the Reverse Hierarchy Theory (RHT) (Hochstein & Ahissar, 2002). In a multifactorial design, the number of points and the number of groups of points was manipulated to assess 3D shape reconstruction via interactive scene rotation. In half of the target groups, the points aligned perfectly to the structure; in the other half a random factor has been added to the points' positioning. With this design, FIT predicts that the time required to detect the structure would increase with the number of low-level objects, i.e. the individual points. RHT, on the other hand, would predict that it would increase with the number of high-level objects, i.e. the point groups. Both theories predict that the randomness of position would disrupt performance to a detectable degree. Structure detection times increase with number of points per group and decreased with number of points per group; a pattern of data consistent with RHT. Additional, stepwise regression analysis ruled out other parameters of the stimulus and participant behaviour as contaminating factors.

To address point (a) above, from an applied, visual analytics perspective, the results imply that a 3D scatterplot can be an efficient tool for trivariate trend detection, especially when the number of individual data points is high, but there are not many options to choose from. Ideally, a 3D scatterplot with a single data set that can either be structured or unstructured would be the optimal format for trend detection tasks, however for exploratory data analysis a more open ended choice of different types of pattern detection would also be appropriate. The possibility of improvement in this task has been demonstrated both in the short-term, in learning during the course of the experiment, and, anecdotally, in the long-term, in the exceptional performance of participants who are experts in manipulating and interpreting computer-generated 3D scenes.

The relevance of 3-D scatterplots to Visual Analytics can easily be overlooked in favour of more traditional 2-D modes of presentation. It should be emphasised that standard statistical tests used in the General Linear Model tradition cannot detect the patterns used in the current study. A typical Visual Analytics task is one where the shape of the data is not known beforehand. In these cases it would not be appropriate to prejudge the test used to detect it (Keehner, Hegarty, Cohen, Khooshabeh, & Montello, 2008). The traditional reliance on the General Linear Model or linear regression analyses (e.g., Keehner, et al., 2008) to establish statistical relationships favours the detection of relationships between pairs of variables rather than across ranges of variables; it is also restricted to finding linear relationships only. To detect an unknown nonlinear relationship between two variables, statistics textbooks (e.g., Tabachnik & Fidell, 2006) suggest plotting the variables as a 2D scatterplot. Similarly, a 3D scatterplot can be used to detect relationships between three variables, making trends in data visible to a human observer as coherent shapes in space. 3D scatterplots have only become available recently, with the widespread use of computers capable of rendering interactive graphics. A static, 2D, representation of a 3D scatterplot (e.g. in a printed medium) is almost impossible to perceive, since it lacks the crucial depth cue of relative motion (Regan, 2000).

The use of motion in graph perception poses a methodological question for the role of 2D resolvable structures in scatterplot-based experimental paradigms. Should essentially 2D structures be included in a 3D structure perception task? If they are, then this effectively increases the search space for the participant, since more options exist, but it also makes some trials much easier. This type of increase of variety of target data set would at the very least engender greater ecological validity in the task. It would also allow more scope for top-down cognitive constraints. Removing a priori constraints on the task should help engender realism in the task. A systematic exploration of the relevant dimensions of colour space would also help set the limits on resolvable differences in colour groupings in such a task and define the maximum number of data sets usable in such visualisations.

In a wider context, this experiment provides theoretical evidence supporting the connection between visual search and formation of high-level, task-relevant scene semantics which answers point (b) above. In principle, the pop-out paradigm can address questions of scene meaning when situated in the appropriate context. This argues that such a paradigm can be used to address the fundamental theories of the workings of perceptual system (Fechner, 1860/1966; Gibson, 1979; Koffka, 1935; Marr, 1982; Steinman, Pizlo, & Pizlo, 2000) and connections between sensation, perception, comprehension and action (e.g., Finday & Gilchrist, 2005; Land & Tatler, 2009). Rather than relying on bottom up explanations such as retinal image analysis or optic flow, the performance in a visual analytics task involving complex interactive analysis can be better envisaged and addressed in paradigms extending beyond the traditional Fechnerian approach. Given the need for impression formation and structure perception in exploratory data analysis (Tukey, 1977) and the rejection of reliance on strict adherence to null hypothesis significance testing (Keehner, et al., 2008), a top-down view of the processing is justified.

Koenderink (2010) highlights the importance of the task in the act of perception since without a question 'there can be no answer' to see (p.37). The search for meaning in data visualisation is particularly relevant when the analysts are inspecting data for a pattern or an outlier, but they don't know what pattern or outlier they are looking for. The point made by a post-Gestalt based research perspective is that whether the search on a surface is for footprints in the sand in the real world, or evidence of activity in abstract data structures, the perceiving necessarily requires a prior position from the perceiver. Progressing beyond pure gestalt towards meaning, Pinna and Albertazzi (2010) point out that as well as grouping and figure ground the principles of 'figureality' (p. 288) of a percept must be considered. From the perspective of abstract data visualisations, this would comprise the colour and volume of an object in 3D space (Pinna & Reeves, 2006). In particular we must pay attention to the distinctions between stationary and non-stationary events. 3D Graph visualisation sits at the junction between 2-D assays of visual perception and the temporal component of structure perception evident in 3D graphical paradigms.

Graphs have evolved to conform to the restrictions of human perception and to present an efficient interface between the data and human comprehension. They can be contrasted to, on the one hand, the written word, which has few physical constraints, but also less variety of spatial organisation, and on the other – to physical objects, both natural and man-made, whose form and function has to conform to certain physical constraints depending on their origin. As such, graphs provide a test-bed for theories of meaning that is both ecologically valid and controllable, presenting an excellent opportunity for research into processes of translation of form into meaning. The next level to address is the transfer from visual insight to the insight in terms of the data domain. It is the capacity to vary the data domain context that provides 3D graph based visual analytics research with the promise of a rich experimental paradigm to address fundamental questions of meaning.

Summary

The burgeoning field of visual analytics offers a novel area for research in principles of Gestalt, specifically in the definitions of object-hood. Abstract visualisations are parsed into perceptual objects that are rich in meaning, the meaning that is based on spatial configuration of these objects and their elements, yet completely devoid of physical restrictions and connotations. The 'Pop-out' paradigm offers an objective testing methodology for the processes of visual scene parsing and onto formation of high-level semantics and pragmatics, i.e. meaning, of the scene. An experiment is presented that applies the 'pop-out' paradigm to trend detection in a 3D scatterplot. The experiment addresses the question: "what constitutes an object in a scatterplot?", and in doing so dissociates between two conflicting theories of perceptual organisation: the Feature Integration Theory and the Reverse Hierarchy Theory, in favour of the Reverse Hierarchy Theory.
Keywords: Visual analytics, Reverse Hierarchy Theory, Feature Integration Theory, perceptual objects, shape perception, 3D visualisation, pop-out.

Zusammenfassung

Visuelle Analytik ist ein neuartiges Feld zur Erforschung von Gestalt-Prinzipien, insbesondere hinsichtlich Definitionen von Objektheit. Abstrakte Visualisierungen werden in bedeutungsreichhaltige perzeptuelle Objekte zergliedert, basierend auf der räumlichen Anordnung dieser Objekte und ihrer Bestandteile, dabei jedoch vollständig ohne physikalische Begrenzungen und Konnotationen. Das `pop-out´ Paradigma stellt eine objektive Testmethode der Prozesse zur visuellen Szenen-Aufteilung und der Entstehung

der Bedeutung der Szene dar. Das `pop-out´ Paradigma wurde in einem Experiment zur Trend-Entdeckung in einer 3D Punktwolke angewendet und untersuchte die Frage „Was konstituiert ein Objekt in einer Punktwolke?". Dabei wurden zwei gegensätzliche Theorien perzeptueller Organisation, `Feature Integration Theory ´und `Reverse Hierarchy Theory gegenübergestellt, und zugunsten der `Reverse Hierarchy Theory´ entschieden.
Schlüsselwörter: Visuelle Analytik, Theorie umgekehrter Hierarchien, Feature Integration Theory, Wahrnehmungsobjekte, Formenwahrnehmung, 3D Visualisierung, pop-out.

Acknowledgements

We are grateful to M Swanston and D Heeley for comments on the early development of this research question, and to Z Pizlo for comments on a previous draft. Preliminary results from this study were presented at ECVP 2008.

References

Bigwood, S. & Spore, M. (2003): *Presenting Numbers, Tables, and Charts*. Oxford: Oxford University Press.
Chen, C. (2004): *Information Visualization: Beyond the Horizon*. London; New York: Springer.
Fechner, G. T. (1860/1966): *Elements of psychophysics*. New York: Holt, Rinehart & Winston.
Finday, J. M. & Gilchrist, I. D. (2005): *Active Vision: The Psychology of Looking and Seeing*: Oxford University Press.
Gabor, D. (1946): Theory of communication. Part 1: The analysis of information. *Electrical Engineers - Part III: Radio and Communication Engineering, Journal of the Institution of, 93*(26), 429-441.
Gibson, J. J. (1979): *The ecological approach to visual perception*. Boston: Houghton Mifflin.
Hershler, O. & Hochstein, S. (2005): At first sight: A high-level pop out effect for faces. *Vision Research, 45*(13), 1707-1724. doi: 10.1016/j.visres.2004.12.021
Hochstein, S. & Ahissar, M. (2002): View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron, 36*(5), 791-804. doi: 10.1016/s0896-6273(02)01091-7
Johnson, C. R., Moorhead, R., Munzner, T., Pfister, H., Rheingans, P. & Yoo, T. S. (2006): *NIH-NSF Visualization Research Challenges Report*: IEEE Press.
Keehner, M., Hegarty, M., Cohen, C., Khooshabeh, P. & Montello, D. R. (2008): Spatial reasoning with external visualizations: What matters is what you see, not whether you interact. *Cognitive Science, 32*(7), 1099-1132. doi: 10.1080/03640210801898177
Koenderink, J. J. (2010): Vision and Information. In L. Albertazzi, G. J. van Tonder & D. Vishwanath (Eds.): *Perception beyond Inference* (pp. 27-58). Cambridge, Mass.: MIT Press.
Koffka, K. (1935): *Principles of Gestalt psychology*. New York: Harcourt, Brace.
Kosslyn, S. M. (2006): *Graph Design for the Eye and Mind*. Oxford: Oxford University Press.
Kubovy, M. & Pomerantz, J. R. (1981): *Perceptual Organization*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
Land, M. F. & Tatler, B. W. (2009): *Looking and Acting: Vision and eye movements in natural behaviour*: Oxford University Press.
Langdridge, D. & Hagger-Johnson, G. (2009): *Introduction to research methods and data analysis* (2nd ed.): Pearson Education (US).
LeDoux, J. (1996): *The emotional brain.* New York: Simon & Schuster.
Limpert, E., Stahel, W. A. & Abbt, M. (2001): Log-normal distributions across the sciences: keys and clues. *Bioscience, 51*(5), 341-352. doi: 10.1641/0006-3568
Mackinlay, J. (1986): Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics (TOG), 5*(2), 110-141. doi: 10.1145/22949.22950
Marr, D. (1982): *Vision: A computational investigation into the human representation and processing of visual information*: Henry Holt and Co., Inc. New York, NY, USA.
Norman, D. A. & Shallice, T. (2000): Attention to action: Willed and automatic control of behavior. *Cognitive neuroscience: a reader*, 376-390.
Pinna, B. & Albertazzi, L. (2010): From Grouping to Visual Meaning: A New Theory of Perceptual Organi-

saion. In L. Albertazzi, G. J. van Tonder & D. Vishwanath (Eds.): *Perception beyond Inference* (pp. 287-344). Cambridge, Mass.: MIT Press.
Pinna, B. & Reeves, A. (2006): Lighting, backlighting and watercolor illusions and the laws of figurality. *Spatial Vision, 19*, 341-373.
Playfair, W. (1801): *The Statistical Breviary; shewing, on a principle entirely new, the resources of every state and kingdom in Europe; illustrated with stained copper-plate charts, representing the physical powers of each distinct nation with ease and perspicuity*. London: Wallis.
Polys, N. F., Kim, S. & Bowman, D. A. (2007): Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. *Computer Animation and Virtual Worlds, 18*(1), 19-38. doi: 10.1002/cav.159
Pomerantz, J. R. (2006): Colour as a Gestalt: Pop out with basic features and with conjunctions. *Visual Cognition, 14*(4), 619-628. doi: 10.1080/13506280500195052
Pomerantz, J. R., Sager, L. C. & Stoever, R. J. (1977): Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance, 3*(3), 422-435. doi: 10.1037/0096-1523.3.3.422
Ratcliff, R. (1979): Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin, 86*(3), 446–461. doi: 10.1037/0033-2909.86.3.446
Regan, D. (2000): *Human perception of objects: Early visual processing of spatial form defined by luminance, color, texture, motion, and binocular disparity*: Sinauer Associates, Sunderland, Mass.
Shah, P. & Hoeffner, J. (2002): Review of graph comprehension research: implications for instruction. *Educational Psychology Review, 14*(1), 47-69. doi: 10.1023/A:1013180410169
Shovman, M. M., Szymkowiak, A., Bown, J. L. & Scott-Brown, K. C. (2008). Use of 'pop-out'paradigm to test graph comprehension in a three-dimensional scatter plot. *Perception, 37* (ECVP Abstract Supplement), 79.
Shovman, M. M., Szymkowiak, A., Bown, J. L. & Scott-Brown, K. C. (2009). *Changing the View: towards the theory of visualisation comprehension.* Paper presented at the IEEE conference on Information Visualisation, Barcelona.
Steinman, R. M., Pizlo, Z. & Pizlo, F. J. (2000): Phi is not beta, and why Wertheimer's discovery launched the Gestalt revolution. *Vision Research, 40*(17), 2257-2264. doi: 10.1016/S0042-6989(00)00086-9
Tabachnik, B. & Fidell, L. S. (2006): *Using Multivariate Statistics*. New York: Allyn & Bacon.
Thomas, J. J. & Cook, K. A. (2005): *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. Los Alamitos, CA: IEEE Computer Society Press.
Treisman, A. M. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance, 8*(2), 194-214. doi: 10.1037/0096-1523.8.2.194
Treisman, A. M. & Gelade, G. (1980): A feature-integration theory of attention. *Cognitive Psychology, 12*(1), 97-136. doi: 10.1016/0010-0285(80)90005-5
Tukey, J. W. (1977): *Exploratory data analysis*. Reading, MA: Addison-Wesley.
Ware, C. & Mitchell, P. (2005): *Reevaluating stereo and motion cues for visualizing graphs in three dimensions* Paper presented at the 2nd symposium on Applied Perception in Graphics and Visualization A Corona, Spain http://doi.acm.org/10.1145/1080402.1080411
Zaitsev, A. V. & Skorik, Y. A. (2002): Mathematical description of sensorimotor reaction time distribution. *Human Physiology, 28*(4), 494-497. doi: 10.1023/A:1016550503294

Mark M. Shovman is an interdisciplinary researcher, studying the perception and cognition aspects of information visualisations, computer games, and immersive virtual reality. Mark was born in 1976 in Georgia, USSR, moving to Jerusalem, Israel in 1990. He is a lecturer in the Institute of Arts, Media and Computer Games, University of Abertay, Dundee.

Kenneth C. Scott-Brown (corresponding author) is a lecturer in Psychology at Abertay University. Born in Scotland in 1970 he gained an MA Hons and then a PhD in Psychology from St Andrews University. Post-docs in Vision Science and Visual Cognition at Glasgow Caledonian University and Nottingham University have fed his wide ranging interests in multidisciplinary research with both theoretical and applied perspectives.

Address: Dr. Kenneth C. Scott-Brown, Division of Psychology, School of Social and Health Sciences, University of Abertay Dundee, Bell Street, Dundee DD1 1HG
E-mail: k.scott-brown@abertay.ac.uk

Andrea Szymkowiak's research interests are in HCI, interaction design and cognitive psychology to investigate the interface between humans and technological systems for various user groups. Born in 1964 in Germany, she obtained a Degree in Psychology at the RWTH Aachen and a PhD in Industrial Engineering in the USA. She is currently a lecturer in Psychology at Abertay University Dundee.

James Bown is Professor of Systems Biology. Born in 1971 in England, he completed his PhD in ecological modelling of plant and fungal community dynamics in 2000. Over the last decade he has developed computational models of complex systems generally and has a particular focus on cancer systems biology. He is interested in developing effective information visualisation techniques to support model development and statistical analysis of results.