A Survey, Classification and Analysis of Perceptual Concepts and their Application for the Effective Visualisation of Complex Information

Burkhard Wünsche

Department of Computer Science
University of Auckland,
Private Bag 92019, Auckland, New Zealand
Email: burkhard@cs.auckland.ac.nz

Abstract

Information visualisation has become increasingly important in science, engineering and commerce as a tool to convey and explore complex sets of information. This paper introduces a visualisation schema which uses visual attributes as the principle components of a visualisation. We present a new classification of visual attributes according to information accuracy, information dimension and spatial requirements and obtain values for the information content and information density of each attribute. The classification applies only to the perception of quantitative information and initial results of experiments suggest that it can not be extended to other visual processing tasks such as preattentive target detection.

The classification in combination with additional guidelines given in this paper provide the reader with a useful tool for creating visualisations which convey complex sets of information more effectively.

Keywords: information visualization, visual effectiveness, visual perception, visual attributes, preattentive processing

1 Introduction

Over the past century knowledge has risen exponentially while human capacity to absorb information has stayed constant. As a result it has become increasingly important to maximise the information acquired from data while minimizing the cost of interacting with it. Visualisation is an attempt to achieve this goal by representing complex sets of information by an image (the *visualisation*) which enables the user to better understand and interpret the data.

The biggest challenge in visualizing information is that the visualisation has to be displayed on a two-dimensional screen using color as an additional output dimension. Animating a visualisation introduces an additional output dimension which is frequently reserved for the independent time variable. Using the characteristics of human visual perception it is possible to simulate additional output dimensions. A third spatial dimension is obtained by using stereoscopic techniques or by using pictorial cues to simulate depth perception (see section 3).

2 The Visualisation Process

Traditionally the visualisation process has been represented by a pipeline which performs *data encoding*. We take the perception and interpretation of the data into account and extend the traditional pipeline model by a *data decoding* step as shown in figure 1.

The first stage of the data encoding step is the *data transformation* stage that converts information into a form more suitable for visualisation. This can involve creation of new quantities and subsets, data type changes, and modeling operations (e.g., model a directory structure as a tree). The subsequent *visualisation mapping* converts the transformed data into graphical representations which the *rendering* stage then displays on a screen or by printing. Some authors (e.g., (Chi 2000, Ware, Chi & Gossweiler 2000)) prefer to subdivide the mapping stage further into *visual transformation* (or *data modelling*) and *visual mapping*. However, in many applications these two stages are combined and the parameters of a model (shape, size, colour, texture) represent the encoded information (Wünsche 2003*b*). The data decoding step describes how visual information is perceived and processed and consists of *visual perception* and *cognition*.

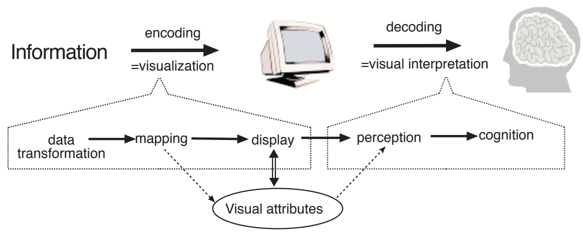

Figure 1: The visualisation process.

The encoding and decoding step of our schema are connected via *visual attributes* such as shape, position, and color, and *textual attributes* such as text and symbols which themselves are represented by simple visual attributes. A visualisation is effective if the decoding can be performed efficiently and correctly. "Correctly" means that perceived data quantities and relationships between data reflect the actual data. "Efficiently" means that a maximum amount of information is perceived in a minimal time.

3 Visual Attributes

The most basic visual attributes are line orientations, color, transparency, position and size (Ferweda 1998, Schiffman 1996, Davidoff 1991). All of these attributes are determined in the initial visual processing step performed by the brain.

While color is a basic visual attribute its perception by the brain is a complex process. Using the concept of *trichromacy* perceived colors can be represented as weighted sum of three *primary colors* red, green, and blue or more intuitively by the perceptual attributes *hue*, *brightness* and *saturation*. The brain's ability to locate colors in this 3D space is limited and even color experts have difficulties separating hue and lightness (Davidoff 1991). Color perception is wavelength dependent with long wavelengths perceived easiest whereas short wavelengths are only identified where luminance is relative high. Some hues such as yellow appear brighter than others such as blue even if they have the same intensities (Schiffman 1996).

Color perception is also influenced by surrounding colors (color illusions). For example, a color patch is perceptually shifted by the color of adjacent patches (*simultaneous contrast*). If colors of different intensities meet, nonexisting intensity changes are perceived (*Mach bands*) (Keller & Keller 1993). Prolonged exposure to a color produces an afterimage of the complementary color which can change perception of subsequent visual impressions. Some black and white patterns can cause color sensations (subjective colors) (Schiffman 1996).

The brain utilises low-level visual attributes for performing more complex visual tasks such as the perception of shape, Gestalt, and depth which together are referred to as *spatial vision* (Ferweda 1998). Other higher order tasks are figure-ground perception and texture perception. The involved visual attributes are called high-level visual attributes.

Texture is perceptually characterised by its spatial frequency, contrast and orientation (Schiffman 1996). Recognition of feature patterns is accomplished using primitive textural features (*textons*) such as length, width and orientation with line segment orientation being particularly important. Pattern detection is orientation dependent and is influenced by adaption (familiarity) (Ferweda 1998).

Shape information is directly derived from luminance, motion, binocular disparity, color, and texture, with luminance yielding shadow and subjective (illusory) contour information (Davidoff 1991). Shape perception is dominated by the curvature of the silhouette contour (figure-ground boundary) and 3D surface shading (Humphreys 1992) with diffuse shading being the most important shape cue. Shape perception is highly orientation dependent such that rotated versions of the same form can be perceived as different shapes. Perception can also be dependent on previous stimuli. Familiar shapes and configurations can improve the recognition of a target if it is a part of them (Schiffman 1996).

Depth perception is achieved using binocular vision and visual cues. Binocular vision includes disparity, convergence and motion parallax. Disparity depends on a object viewed by two eyes which are slightly displaced so that the perceived images differ slightly. The displacement of the retinal images of an object is converted by the brain to depth information. Motion parallax is the effect that the relative distance an object moves determines the amount its image moves on the retina. For visualisation purposes binocular vision is achieved by using stereo goggles or VR Head Displays. Independent of this visual cues such as size, brightness, perspective, overlay, texture gradient, and aerial perspective (Humphreys 1992) are used to aid depth perception.

The concept of *Gestalt* originates from the fine arts and expresses the notion that the "whole contains more information than the parts". Perception of Gestalt is influenced by proximity, similarity, continuation, closure, symmetry, and the *law of Prägnanz*, which states that the eyes tend to see the simplest and most stable figure (Schiffman 1996). Context might also play a role in Gestalt perception (Humphreys 1992).

Figure-ground perception describes the observation that an object can be instantly separated perceptually from its background. This is due to physically different attributes of the figure and the background but is also influenced by size, angle, and association with meaningful shapes (Schiffman 1996).

4 Classification of Visual Attributes

In this section we suggest a classification of visual attributes according to representational accuracy, perceptual dimension and spatial requirement. The classification supports the identification of suitable visual attributes for representing a given data set and hence forms the basis for mapping data onto visualization icons.

Not all visual attributes are equally well suited to display quantitative information. For many attributes their perceived scale is a power of the actual scale (Steven's law) (Cleveland 1985). The power is close to one for the perception of length so that length variations can be estimated quite accurately. For area and volume changes the power is smaller than one so that small areas are usually perceived larger than they actually are and vice versa for large areas. In addition perception of visual attributes can be influenced by orientation, e.g., angles with a horizontal bisector are seen larger than angles with a vertical one (Cleveland 1985). Also it has been shown that slope changes influence the perception of vertical distances.

As a consequence the suitability of visual attributes for information encoding differs. We use the term *information accuracy* as a measure of how accurately a human can estimate a quantitative variable represented by that visual attribute. Cleveland shows that such a variable is most accurately represented by a position along a scale, and then in decreasing order of accuracy by interval length, slope or angle, area, volume and color as indicated below (Cleveland 1985).

highest accuracy	↑	position on scale
of representation	│	interval length
	│	slope/angle
	│	area
lowest accuracy	│	volume
of representation	↓	color

The above ranking changes when visualizing ordinal or nominal data which occur frequently in information visualisation (Mackinlay 1986). The suitability of a high level visual attribute for visualizing quantitative data depends on which low level attribute dominates its perception. For example, if a texture consists of colored strokes with a direction and length then the stroke color and direction are usually easier perceived than the length of individual strokes.

We suggest to further differentiate visual attributes by their information dimension and spatial requirements. *Information dimension* refers to the number of dimensions inherent in the visual attribute. Length and slope represent only one dimension but color can be used to represent at least two dimensions. Texture is usually composed of several basic visual attributes such as color and the length and orientation of texture elements. The total information dimension is therefore the sum of the dimensions of the inherent basic attributes. An additional output dimension can be represented by the spatial frequency of a texture. Similarly shape has been shown to represent multiple independent output dimensions.

We define the *spatial requirement* of a visual attribute as the smallest unit of space (i.e., pixels on a screen) necessary to identify a piece of information. Whereas color has a minimal spatial requirement only limited by the resolution of the human visual system a texture requires a much larger space of the output medium to enable the viewer to identify inherent information. For clarity the spatial requirement is described both subjectively (low-high) and then in brackets more objectively by the dimension of the occupied space. For example, a position on a scale is given by a point (dimension 0) whereas volumes and other 3D shapes occupy three-dimensional regions. Color can be

	Information dimension	Information accuracy	Spatial requirement (dimension)	Information content	Information density
Position on scale	1–3	High	Low (0)	High	High
Length	1	High	Medium (1)	Medium	Low
3D Direction	2	Medium	Medium (1)	Medium	Medium
Area	1	Medium	Medium (2)	Low-Medium	Low
Volume	1	Medium	High (3)	Low-Medium	Very Low
Shape	$\geq$3	Low-Medium	High (3)	Medium-High	Medium-High
Texture	$\geq$3	Low-Medium	Medium (1-3)	Medium-High	Medium-High
Color	2	Low	Low ($\geq$0)	Medium	High

Table 1: Classification of common visual attributes.

represented by points, however, color perception is very poor for isolated pixels so that for many applications the space requirement of the color attribute is higher than zero dimensions.

The *information content* of a visual attribute can now be defined as the product of information accuracy and information dimension. The *information density* is given by dividing the information content of a visual attribute by its spatial requirement. A listing of common visual attributes classified using above criteria is shown in table 1.

An alternative classification of visual attributes for information visualization has been suggested by Dastani (Dastani 2002). The author classifies visual attributes into spatial attributes, non-spatial attributes and topological attributes and derives from that different perceptual structures with the aim of finding a structure-preserving mapping from data structures onto perceptual structures.

5 Mapping Information onto Visual Attributes

In order to create effective visualisations information must be mapped to visual attributes in a way that optimises its perception and understanding. The task is difficult since the perception, interpretation and comprehension of visual input is influenced by context, attentional focus, expectations, prior knowledge, past experiences and subjective biases (Healey, Interrante & Rheingans 1999).

The mapping between data and visual attributes is usually determined by the intended function of a graphical representation. The following functions are common:

- Display quantitative information
- Draw attention
- Show correlation

Quantitative information is best displayed by length and position and is therefore reflected in the size of a graphical representation. Depending on the required accuracy of the representation, the available space and the number of simultaneously displayed data sets alternative attributes might be more appropriate as shown in table 1.

Figure 2 shows as an example a circular vector field which we visualise with different graphical representations (*visualisation icons*) for the data (Wünsche 2003*b*). The velocity direction is indicated by a Line Integral Convolution texture (LIC) (Cabral & Leedom 1993, Stalling & Hege 1995) and by the direction of the vector arrows. While vector arrows give more precise directional information they can easily be misleading since there is no indication of to which data point they apply. The vector magnitude is represented by the colour of the LIC texture (a poor representation, but continuous), by the length of an arrow (precise, but only available for selected data points) and by a height field which offers the most accurate representation since the boundaries of the domain can be used as a scale. Note that only the height field shows clearly that the vector magnitude increases linearly from the centre of the data set.

Attention can be drawn to a target by using bright or highly saturated colors, sharp boundaries, or movement or

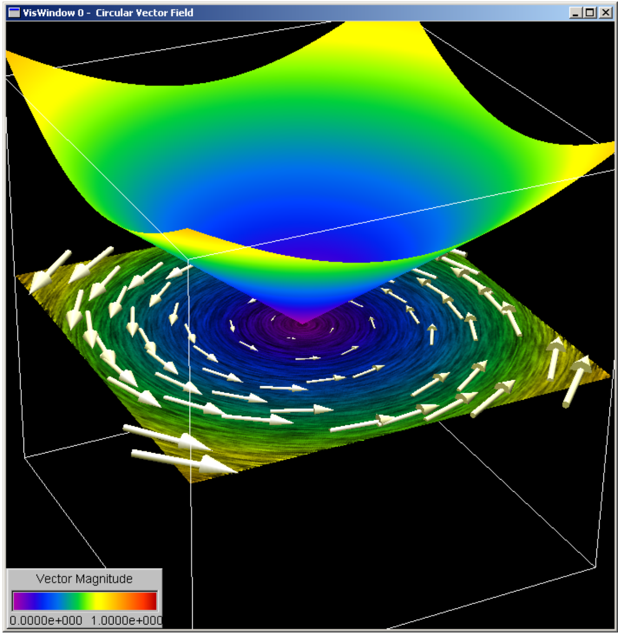

Figure 2: A circular vector field visualised using different visualisation icons.

change (Rheingans & Landreth 1995). Target identification is also influenced by linear separation, color category, and color distance (Healey et al. 1999). If the complexity of the scene allows it instant target identification can be achieved by using preattentive features.

Preattentive features are very simple features that are perceived without conscious attention. An example is the instantaneous perception of a red dot in a cloud of blue ones. The underlying mechanism has been contributed to different sensory dimensions for basic visual attributes such that a unique feature in any dimension is immediately detected (Davidoff 1991). Preattentive vision seems to be dependent on primitive textural features such as length, width and orientation of simple elongated shape as well as their end connections, angle orientations, and intersections (Schiffman 1996). In addition preattentive vision exists for shape, curvature, closure, color (hue), intensity and more complex visual attributes such as texture and depth (Healey et al. 1999).

5.1 Analysis of Visual Attributes

The classification in table 1 was obtained from results previously presented in the literature and from our own experiences and some simple experiments with students. More formal experiments are necessary to verify and quantify the results for practical applications.

Our classification was motivated by the goal to measure the suitability of a visual attribute for displaying quantitative information. An interesting question is whether our classification still holds true when a visual attribute is used for other tasks, such as instant target recognition by preattentive processing.

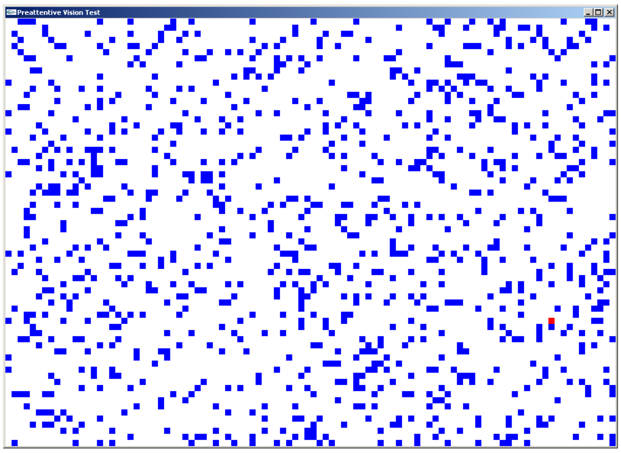

Figure 3: Screen shot of a program to test preattentive target identification.

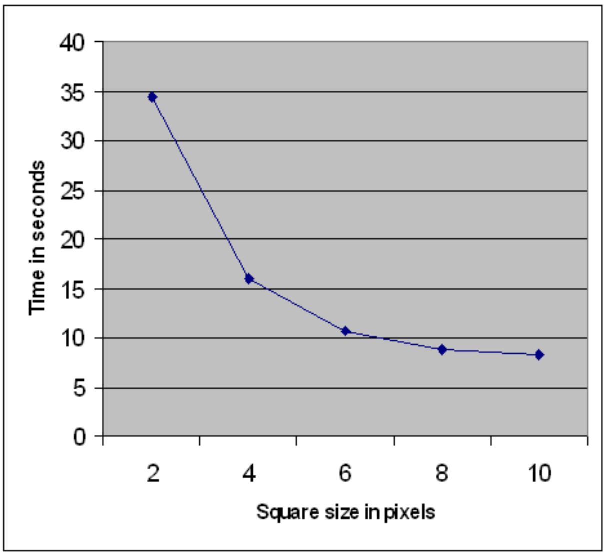

Figure 4: A chart of the time it took to find a red square among 1000 blue squares using different square sizes.

Figure 3 shows an application developed by us to determine the properties of visual attributes for preattentive processing. After starting the application the screen is initially empty. If the user presses a key 1000 blue squares and one red square are randomly distributed over the screen without overlaps. As soon as the user finds the red square he/she presses another key. The task is repeated 10 times and the total time between each pair of key presses is recorded.

Two experiments were performed using 38 volunteers. All but two of the volunteers were between 20–30 years old and had limited previous experience with visualizations. Two volunteers were older and considerable more experienced.

The first experiment tested how preattentive perception depends on the object size. The users performed the above described task for squares with side lengths of 2,4,6,8,and 10 pixels. The results for one subject are illustrated in figure 4. It can be seen that the time to find the red square increases roughly exponentially with smaller pixel size. An interesting result was that the ratio of the times for performing the five tasks were similar for all 38 subjects, e.g., finding the red square with square sizes of 2×2 pixels took 3-5 times as long as when the square sizes were 10×10 pixels. However, the actual times recorded varied greatly between subjects. The slowest person (who was partially colour blind) took for each task 3-5 times as long as the fastest one. Other factors which might be responsible for these differences but have not been examined in detail are lack visual acuity and other vision defects, reaction time, and mental awareness.

Closer examination showed that the volunteers differed in the way they performed the tasks. Whereas some users kept a constant distance from the screen other users moved closer to the screen with decreasing size of the squares. We found that for large pixel sizes the user must sit relatively far behind the screen in order to find the target without having to scan the screen. In contrast for small squares it seems to be virtually impossible to find the target preattentively with one glance and sitting closer to the screen and looking at different areas of the screen seems to be the most effective way to find the target.

Our initial results indicate that target identification strongly depends on the size of the target and the spatial resolution of the eye. In particular the results might indicate that some properties of visual attributes, e.g., the fact that colour has a low spatial requirement, do not apply to target identification tasks.

In the second experiment we tested how preattentive target identification is affected by the colour dimension of non-target squares. Initially all non-target squares were blue and in the subsequent four tasks we added green, black, yellow and cyan squares, respectively. In all five tasks the number of squares was the same for each colour present in the image. The results of the study are shown in figure 5.

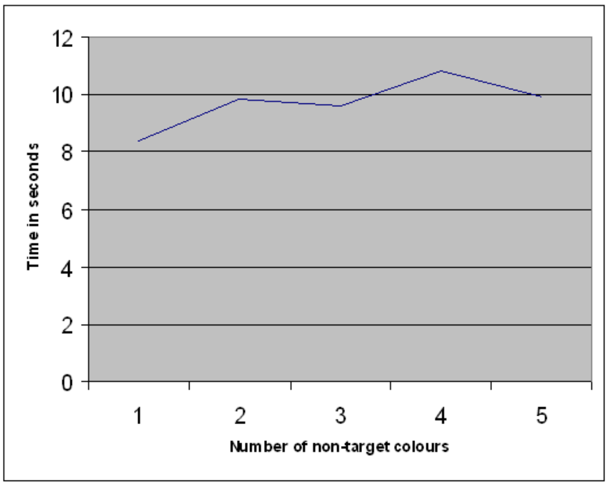

Figure 5: A chart of the time it took to find a red square among 1000 non-target squares with 1, 2, 3, 4 and 5 different colours, respectively.

It can be seen that the time for finding the red square increases with an increasing number of colours, but the increase is not as significant as if the size is reduced. An interesting observation is that for most volunteers the time required to perform the task with 5 differently coloured non-target squares was smaller than if using only 4 different colours. One reason for this might be that yellow is a "hot" colour which draws attention. Using 5 colours rather than 4 reduces the number of yellow coloured squares and as such might make it easier to find the red square.

As a conclusion from these initial experiments we suggest that the classification in table 1 does not extend to preattentive target identification. For example, colour is listed as having a high information density because it has a low spatial requirement and a high information density. However, our experiments show that for preattentive target identification a certain minimum amount of screen space is necessary and that the number of different colours which can be used is limited. We intend to perform similar experiments for other preattentive attributes in order to determine their information content and information density for preattentive target identification. Further experiments are necessary in order to determine the results for other preattentive tasks.

5.2 Combining Visual Attributes

As mentioned previously different visual attributes can be used to create graphical primitives (visualization icons) to represent data. The visualization icons themselves can also be combined in order to yield more information than the sum of the individual icons (Keller & Keller 1993). Additional information may exist in the form of correlation between multiple variables or as higher-order visual information (Gestalt).

Correlation between related data sets is perceived most easily when similar visualisation icons are used (Keller & Keller 1993). Perception can be further improved by using multiple visualisation techniques simultaneously for the same data (Rheingans & Landreth 1995).

Gestalt concepts in visualisation are demonstrated by Laidlaw et al. who use densely-arranged normalised tensor ellipsoids in order to obtain a texture-like representation of a diffusion tensor field which improves the perception of features and field properties (Laidlaw, Ahrens, Kremers & Readhead 1998).

In contrast to correlated variables unrelated variables are best displayed using orthogonal (independent) visual attributes such as shape, colour, movement, and texture. Many visualisation icons utilise multiple visual attributes so extra care has to be taken when combining such icons. In general it has been shown that the brain can handle a maximum of about seven unrelated elements (Keller & Keller 1993). Note that different visualisation icons can also be used to display the same data in order to reinforce information or to highlight different aspects of the data (explicit redundancy).

Further guidelines for combining visualisation icons are obtained from research on graphing data. For example, visualisation icons which overlap should be visually distinguishable (Cleveland 1985). When visualisation icons with similar shape are used colour can be used to discriminate between them (Cleveland 1985). Mackinlay (Mackinlay 1986) extends work from Bertin (Bertin 1983) and classifies graphical encoding techniques into marks (points, lines, areas), positional (1D, 2D, 3D), temporal (animation), retinal (colour, shape, size, saturation, texture and orientation), maps, connections (tree, network) and others. The author uses this classification to create a composition algebra which specifies whether two encoding techniques can be used for the same task. While the work was developed for graph design many results apply to the field of scientific visualisation. For example, when size and shape are composed together small sizes must be avoided since the shapes of the small objects may be hard to distinguish.

This observation is reflected in our classification of visual attributes given in table 1. For example, length has a medium spatial requirement but shape has a high spatial requirement. Consequently an icon using both length and shape for encoding data is restricted in its usage by the spatial requirements for shape encoding. In general the spatial requirement of an icon is given by the largest spatial requirement of any of its visual attributes.

The effectiveness of visualisation icons is also influenced by the chosen background. The background can be used to highlight and support features in the image and can be used to provide supplementary information and 3D perspective (Keller & Keller 1993). Keller and Keller recommend that the background of a visualisation should have a neutral (unsaturated) colour with a good contrast to the foreground. The authors further recommend the use of a horizontal (landscape) view since it corresponds to the normal field of vision. 3D scenes should be oriented in such a way that important features are in the foreground and not covered by other scene components (Keller & Keller 1993).

6 Increasing the Effectiveness of a Visualisation

A general approach for the creation of effective visualisations is given by the “Natural Scene Paradigm” which is based on our ability to immediately perceive complex information in natural scenes (Robertson 1991). Implementing this paradigm involves clear 3D structures and the association of data with recognizable properties of objects.

In many cases no natural association between data and icons exist and the understanding of a visualisation is dependent on the target audience having a priori knowledge about it. In particular familiarity with the data set and the particular visualisation techniques is often required.

6.1 Lighting

Since our environment is illuminated by a wide variety of natural and artificial light sources, the human brain is well adapted to perceive geometric information from shading. As mentioned in section 3 the single most important clue in shape recognition is diffuse illumination. Hence lighting is essential if using icons which encode information by shape (such as height fields, isosurfaces, and tensor ellipsoids). On the other hand, illuminating an object changes its perceived colour so that lighting should be disabled if colour is the primary visual attribute. For example, in figure 2 lighting is disabled for the rendering of flat or nearly flat (“shape-less”) surfaces such as the textured ground plane and the heightfield but lighting is enabled for the vector arrows.

Some authors suggest that diffuse shading is the most important shape cue and that adding specularity does not significantly improve perception of shape differences (Rodger & Browse 2000). Therefore specular material properties should be avoided for icons which use colour as a secondary visual attribute. On the other hand specular highlights can help to distinguish object details, such as the radius of a rounded edge, so that adding specularity to objects with low colour variations (e.g., isosurfaces) improves the amount of perceived information.

The use of shadows can further improve the perception of the 3D geometry of an object. For example, shadows have been successfully employed for visualizing 3D vectors over 2D slices (Klassen & Harrington 1991). Shadows can also be used to indicate the distance of an object from a background plane and help to indicate the spatial order of objects. It has also been shown that shadowing increases the accuracy (but not speed) of object positioning (Hubona, Wheeler, Shirah & Brandt 1999). However, Hubona et al. (Hubona et al. 1999) show that using multiple shadowing light sources decreases user performance for positioning and resizing tasks, which indicates that the perception and interpretation of scientific visualisations might also suffer if more than one light source is used for shadow creation.

6.2 Perceptual Clues

In general it is difficult to design a visualisation using the natural scene paradigm. More concrete techniques for improving perception and understanding are

- Shape clues
- Contextual clues
- Annotations

Shape clues are used to improve the perception of the 3D geometry of a scene. Two major classes of shape clues exist: illumination (explained in the previous subsection) and explicit redundancies. Techniques based on explicit redundancies include emphasizing of silhouette curves (figure-ground boundary) and contour curves (depth discontinuities) (Saito & Takahashi 1990) and the use of mirrors. Projections of coloured shadows on the 3 coordinate planes have also been used (Allen B. Tucker 1997).

Contextual clues improve perception by enabling the brain to relate abstract visualisation icons to familiar objects or properties. Examples of contextual clues inherent in a data set are coastlines, bounding boxes, and model outlines which improve the perception of positional information. Motion blur can be used to indicate velocities. Additional contextual clues to make data more readable include numbered scales, grid lines, and abstract objects to suggest value and relationships (see (Tufte 1983)). Object recognition can be increased by comparing an object with similar ones familiar to the user. This can be achieved by using multiple windows with the same view, by using split-screen techniques or by using overlay techniques.

Zhang et al. (Zhang, Curry, Morris & Laidlaw 2000) apply contextual clues to the field of medical imaging and use easily identified anatomical features to improve the understanding of the 3D geometry of visualised nerve fiber structures. An example from our own work (Wünsche & Lobb 2001, Wünsche 2003*b*) is shown in figure 6. The tube-like structures indicate nerve fiber tracts, whereas the green and red isosurfaces represent the eyes and the ventricles, respectively. The latter two objects are *anatomical landmarks* which indicate the orientation of the data set (the eyes are in the front of the head), improve the perception of the position of the nerve fibers inside the head, and clarify the perception of the 3D geometry since the fibers tracts furthest away from the view point are occluded by the ventricles.

Finally annotations can be used to identify features and to explain relationships. Examples are legends, labels, and markers. Legends should be comprehensive, informative and draw attention to important features in the data set (Cleveland 1985). Care has to be taken that the annotations do not distract from the actual goal of the visualisation (Tufte 1983).

6.3 Exploration Techniques

The perception of a visualised data set is further improved by enabling the user to interact with the data. Common types of interaction are rotation, translation (pan) and zoom. Walk through and fly through features are also popular. An example of the resulting improvement in perception is explained in (Wünsche 2003*a*, Wünsche & Young 2003): If large numbers of icons are distributed over a 3D domain, rotating the model around its axis enables the brain to differentiate icons in the foreground and the background. Animating the interaction, e.g., using continuous rotations or automatic fly-throughs, can help the user to concentrate on the data.

Other common interaction techniques are fish-eye views and cut-away (clipping) techniques. A generalisation of clipping is the *sectioning tool* (Wünsche 2003*b*).

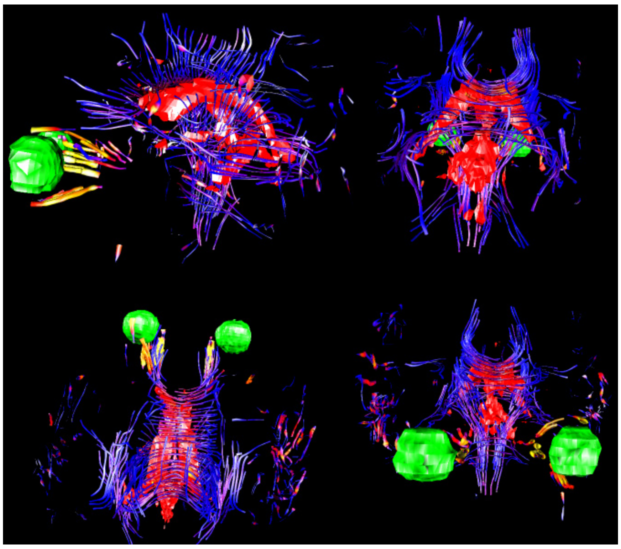

Figure 6: A visualisation of the nerve fiber structure in the brain. Perception is improved by inserting familiar anatomical structures such as the eyes (green) and the ventricles (red).

The tool slices a data set into sections and arranges them regularly in 3D. Inner structure is revealed but the global structure can still be perceived due to the brain's ability to visually interpolate slice data.

Region-of-interest techniques allow the user to extract a region of interest from the data volume. In most cases regions of interest are interactively defined by placing a box or a sphere into the volume. In some cases simple shapes are not sufficient to extract an interesting region such as an anatomical abnormality in a medical data set. Ney and Fishman (Ney & Fishman 1991) present a tool for interactively creating shapes suitable to define arbitrarily shaped regions of interest in a data set.

Fuhrmann and Gröller (Fuhrmann & Gröller 1998) introduce magic lenses and magic boxes as a tool to improve 3D interaction. Magic lenses are planar usually circular objects which magnify the scene behind the lens and remove the volume in front of it. Magic volumes are explicit focus volumes in which a detailed representation of the visualisation is displayed. The back faces of the box are opaque in order to reduce distraction.

Large data sets can be efficiently explored by using multiscale visualisations which use different levels of abstraction for detail views and overviews of the data. An example is given by Stolte et al. (Stolte, Tang & Hanrahan 2003).

Data brushing, originally developed for multivariate data, can also be utilised in scientific visualisation by selecting or highlighting icons shown in one view in all other views of the data. The technique can be used to increase rendering speed, to highlight or to extract features and to facilitate the discovery of relationships between subsets of the data (Wong & Bergeron 1996, Wong & Bergeron 1997).

Finally image graphs (Ma 1999, Ma 2000) and visualisation spreadsheets (Chi, Riedl, Barry & Konstan 1998, Jankun-Kelly & Ma 2001) allow the user to interactively change visualisation parameters while seeing their effects on the final image.

An increasingly popular approach to dealing with extremely large data sets is the use of immersive environments such as virtual reality (VR) workbenches (Bryson & Levit 1991, Bryson 1996) and CAVE theatres (Jaswal 1997). Interaction with data is achieved using data gloves (Bryson & Levit 1991, Bryson 1996, Fröhlich, Barrass, Zehner, Plate & Göbel 1999) or natural interaction techniques such as speech and hand gestures (Sharma, Zeller, Pavlovic, Huang, Lo, Chu, Zhao, Phillips & Schulten 2000).

For large data sets exploration results might be improved by employing a collaborative visualisation in which research teams collectively analyze data (Wood, Wright & Brodlie 1997). Fuhrmann et al. suggest that collaborative exploration is facilitated by using augmented reality which combines familiar physical surroundings with synthetic data (Fuhrmann, Löffelmann, Schmalstieg & Gervautz 1998). Issues relating to collaborative control are discussed in (Bresnahan, Insley & Papka 2000). Recently collaborative visualisation over the Internet has been proposed as an effective learning tool (Pea 2002).

Direct interaction with the data can be replaced or supplemented by a presentation simulating an interaction. For example, a sequence of successively magnified images reveals structure whereas a simultaneous display of images using different techniques shows multiple aspects of a data set (Keller & Keller 1993).

7 Conclusion

We suggested a classification of visual attributes according to information accuracy, information dimension and spatial requirement. From this classification we determined the information content and the information density of a visual attribute. The classification can be used to select and to develop graphical representations for a given visualisation task by considering the visual attributes most suitable for that task. An initial analysis of visual attributes for preattentive target identification suggests that our results do not extend to this task. Further research is necessary to quantify our classification and to extend it to other visual processing tasks. An important observation which might be useful for the development of user interfaces for data exploration is that the effectiveness of a visualization can not be measured in absolute terms but depends on how it is viewed.

We presented additional guidelines for creating more effective visualisations and motivated them using the "Natural Scene Paradigm". An important result is that lighting should not be applied globally to a scene but only to objects where illumination aids the shape recognition. If the color of an objects encodes important information it might be better to render it without shading. While unnatural this rule optimises the perception of information from shape and colour attributes.

We hope that this work will enable the reader to create visualizations that represent complex sets of information more effectively.

References

Allen B. Tucker, J., ed. (1997), *The Computer Science and Engineering Handbook*, CRC Press Inc., 2000 Corporate Blvd., N.W., Boca Raton, Florida 33431.

Bertin, J. (1983), *Semiology of graphics: diagrams, networks, maps*, University of Wisconsin Press, Madison, Wis. Translated by William J. Berg.

Bresnahan, J., Insley, J. & Papka, M. E. (2000), Interacting with scientific visualizations: User-interface tools within spatially immersive displays, Technical Report ANL/MCS-P789-0100, Argonne National Laboratory, Argonne, Illinois, USA.

Bryson, S. (1996), 'Virtual reality in scientific visualization', *Communications of ACM* **39**(5), 62–71.

Bryson, S. & Levit, C. (1991), The virtual windtunnel: An environment for the exploration of three-dimensional unsteady flows, *in* G. M. Nielson & L. Rosenblum, eds, 'Proceedings of Visualization '91', IEEE, Computer Society Press, Los Alamitos, California, pp. 17 – 24.

Cabral, B. & Leedom, L. C. (1993), Imaging vector fields using line integral convolution, *in* J. T. Kajiya, ed., 'Computer Graphics (SIGGRAPH '93 Proceedings)', Vol. 27, ACM SIGGRAPH, Addison Wesley, pp. 263–272.

Chi, E. H. (2000), A taxonomy of visualization techniques using the data state reference model, *in* 'Proceedings of the Symposium on Information Visualization (InfoVis '00)', IEEE Press, October 9–10, Salt Lake City, Utah, pp. 69–75. URL: `http://www-users.cs.umn.edu/~echi/papers/infovis00/Chi-TaxonomyVisualization.pdf`.

Chi, E. H., Riedl, J., Barry, P. & Konstan, J. (1998), 'Principles for information visualization spreadsheets', *IEEE Computer Graphics and Applications (Special Issue on Visualization)* pp. 30–38. URL: `http://www-users.cs.umn.edu/~echi/papers/cga98/cga-spreadsheet-principle.pdf`.

Cleveland, W. S. (1985), *The elements of graphing data*, Murray Hill, N.J. : AT&T Bell Laboratories.

Dastani, M. (2002), 'The role of visual perception in data visualization', *Journal of Visual Languages and Computing* **13**(6), 601–622.

Davidoff, J. B. (1991), *Cognition through color*, MIT Press, Cambridge, Mass.

Ferweda, J. A. (1998), 'Fundamentals of spatial vision'. Applications of visual perception in computer graphics, Course #32, SIGGRAPH '98, 1-27.

Fröhlich, B., Barrass, S., Zehner, B., Plate, J. & Göbel, M. (1999), Exploring geo-scientific data in virtual environments, *in* D. Ebert, M. Gross & B. Hamann, eds, 'Proceedings of Visualization '99', IEEE, pp. 169 – 173.

Fuhrmann, A. & Gröller, E. (1998), Real-time techniques for 3D flow visualization, *in* D. Ebert, H. Hagen & H. Rushmeier, eds, 'Proceedings of Visualization '98', IEEE, Computer Society Press, pp. 305 – 312.

Fuhrmann, A., Löffelmann, H., Schmalstieg, D. & Gervautz, M. (1998), 'Collaborative visualization in augmented reality', *IEEE Computer Graphics and Applications* **18**(4), 54–59.

Healey, C., Interrante, V. & Rheingans, P. (1999), 'Fundamental issues of visual perception for effective image generation'. Course notes #6, SIGGRAPH 1999.

Hubona, G. S., Wheeler, P. N., Shirah, G. W. & Brandt, M. (1999), 'The relative contributions of stereo, lighting and background scenes in promoting 3D depth visualizaqtion', *ACM Transactions on Computer-Human Interface* **6**(3), 214 – 242.

Humphreys, G. W., ed. (1992), *Understanding vision : an interdisciplinary perspective*, Blackwell, Oxford, UK ; Cambridge, USA.

Jankun-Kelly, T. J. & Ma, K.-L. (2001), 'Visualization exploration and encapsulation via a spreadsheet-like interface', *IEEE Transactions on Visualization and Computer Graphics* **7**(3), 275–287.

Jaswal, V. (1997), Cavevis: Distributed real-time visualization of time-varying scalar and vector fields using the cave virtual reality theater, *in* R. Yagel & H. Hagen, eds, 'Proceedings of Visualization '97', IEEE, pp. 301 – 308.

Keller, P. R. & Keller, M. M. (1993), *Visual Cues - Practical Data Visualization*, IEEE Computer Society Press, Los Alamitos, CA.

Klassen, R. V. & Harrington, S. J. (1991), Shadowed hedgehogs: A technique for visualizing 2D slices of 3D vector fields, *in* G. M. Nielson & L. Rosenblum, eds, 'Proceedings of Visualization '91', IEEE, Computer Society Press, Los Alamitos, California, pp. 148–153.

Laidlaw, D. H., Ahrens, E. T., Kremers, D. & Readhead, C. (1998), Visualizing diffusion tensor images of the mouse spinal cord, *in* D. Ebert, H. Hagen & H. Rushmeier, eds, 'Proceedings of Visualization '98', IEEE, Computer Society Press, pp. 127 – 134.

Ma, K.-L. (1999), Image graphs - a novel approach to visual data exploration, *in* D. Ebert, M. Gross & B. Hamann, eds, 'Proceedings of Visualization '99', IEEE, pp. 81–88. ISBN 0-7803-5897-X.

Ma, K.-L. (2000), 'Visualizing visualizations', *IEEE Computer Graphics and Applications* **20**(5), 16 – 19.

Mackinlay, J. D. (1986), 'Automating the design of graphical presentations of relational information', *ACM Transactions on Graphics* **5**(2), 110–141.

Ney, D. R. & Fishman, E. K. (1991), 'Editing tools for 3d medical imaging', *IEEE Computer Graphics and Applications* **11**(6), 63–71.

Pea, R. D. (2002), 'Learning science through collaborative visualization over the internet'. URL: `http://www.nobel.se/nobel/nobel-foundation/symposia/interdisciplinary/ns120/lectures/pea.pdf`.

Rheingans, P. & Landreth, C. (1995), Perceptual principles of visualization, *in* G. G. Grinstein & H. Levkowitz, eds, 'Perceptual Issues in Visualization', Springer Verlag, Berlin, New York, pp. 59 – 73.

Robertson, P. K. (1991), 'A methodology for choosing data representation', *IEEE Computer Graphics and Applications* **11**(3), 56 – 67.

Rodger, J. C. & Browse, R. A. (2000), 'Choosing rendering parameters for effective communication of 3D shape', *IEEE Computer Graphics and Applications* **20**(2), 20–28.

Saito, T. & Takahashi, T. (1990), Comprehensible rendering of 3-D shapes, *in* F. Baskett, ed., 'Computer Graphics (SIGGRAPH '90 Proceedings)', Vol. 24, pp. 197–206.

Schiffman, H. R. (1996), *Sensation and Perception: An Integrated Approach*, 4th edition edn, John Wiley & Sons.

Sharma, R., Zeller, M., Pavlovic, V. I., Huang, T. S., Lo, Z., Chu, S., Zhao, Y., Phillips, J. C. & Schulten, K. (2000), 'Speech/gesture interface to a visual-computing environment', *IEEE Computer Graphics and Applications* **20**(2), 29–37.

Stalling, D. & Hege, H.-C. (1995), Fast and resolution independent line integral convolution, *in* 'Computer Graphics Proceedings (SIGGRAPH '95)', pp. 249 – 256.

Stolte, C., Tang, D. & Hanrahan, P. (2003), 'Multiscale visualization using data cubes', *IEEE Transactions on Visualization and Computer Graphics* **9**(2), 176 – 187.

Tufte, E. R. (1983), *The Visual Display of Quantitative Information*, The Graphics Press, Cheshire, Connecticut.

Ware, C., Chi, E. H. & Gossweiler, R. (2000), Visual perception for data visualization, *in* 'Proceedings of the Human Factor in Computing Systems Conference (CHI 2000)', Amsterdam, Netherlands. Tutorial for CHI 2000, URL: `http://www-users.cs.umn.edu/~echi/tutorial/perception2000/`.

Wong, P. C. & Bergeron, R. D. (1996), Multiresolution multidimensional wavelet brushing, *in* R. Yagel & G. M. Nielson, eds, 'Proceedings of Visualization '96', IEEE, pp. 141 – 148.

Wong, P. C. & Bergeron, R. D. (1997), Brushing techniques for exploring volumetric datasets, *in* R. Yagel & H. Hagen, eds, 'Proceedings of Visualization '97', IEEE, pp. 429 – 432.

Wood, J. D., Wright, H. & Brodlie, K. W. (1997), Collaborative visualization, *in* R. Yagel & H. Hagen, eds, 'Proceedings of Visualization '97', IEEE, pp. 253 – 260. ISBN 0-58113-011-2.

Wünsche, B. C. (2003*a*), A toolkit for visualizing biomedical data sets, *in* 'Proceedings of GRAPHITE 2003', ANZGRAPH and SEAGRAPH, pp. 167–174.

Wünsche, B. C. (2003*b*), The Visualization of Tensor Fields in Biological Tissue, PhD thesis, University of Auckland. (To be published).

Wünsche, B. C. & Lobb, R. (2001), The visualization of diffusion tensor fields in the brain, *in* 'Proceedings of METMBS'01', CSREA Press, pp. 498 – 504.

Wünsche, B. C. & Young, A. A. (2003), 'The visualization and measurement of left ventricular deformation using finite element models', *Journal of Visual Languages and Computing - Special Issue on Biomedical Visualization for Bioinformatics* **14**(4), 299–326.

Zhang, S., Curry, C., Morris, D. S. & Laidlaw, D. H. (2000), Visualizing diffusion tensor MR images using streamtubes and streamsurfaces, URL: `http://www.cs.brown.edu/~dhl/pdf/hbp00dti.pdf`.