BEHAVIOUR & INFORMATION TECHNOLOGY, 1998, VOL. 17, NO. 1, 18 – 26

Visualization using colour: visual presentation of events in particle physics

HANS DREVERMANN† and DAVID TRAVIS‡

†CERN, ECP, 1211 Geneva 23, Geneva, Switzerland

‡System Concepts Ltd, 2 Savoy Court, Strand, London WC2R 0EZ, UK

Abstract. This paper applies a model of colour vision to achieve optimal use of colour in a software system that visualizes the results of experiments in high energy physics. It shows how the elements of the visualization were designed and provides details of why particular colours were chosen. By grounding these findings in psychological research, it is able to show how other computer systems that use colour may profitably apply this methodology.

1. Introduction

In many high technology areas, where data sets are large and complex, increasing reliance is placed on computing systems to carry out data analysis. These data, or part of them, are then typically visualized in a pictorial form to exploit the exquisite pattern recognition of the human visual system.

Because the visual system is a key component in these types of analyses, it is surprising that few practical examples of visualization systems have been built upon our knowledge of human vision.With notable exceptions (Travis 1990, Rogowitz and Treinish 1993, Wickens *et al.* 1994), much of the previous work has been based on visual design heuristics, and some trial and error. This approach has a long history for non-interactive tasks such as map reading, newspaper layout and the interpretation of pictorial instructions, where it works well (Tufte 1990). However, it often requires specialized and talented individuals to design the system and it takes a considerable amount of time. Interactive systems, on the other hand, require a flexible approach since data under analysis need to be manipulated in real time and so may change in ways that cannot be predicted in advance.

One such flexible approach, based on human perception, is to integrate findings from colour science, visual psychophysics and cognitive psychology to derive specific guidance for display design (Travis 1991). There are good reasons for following this approach, because human visual perception is subject to various visual illusions (Gregory 1977). These illusions can affect the perception of graphs, features and objects, suggesting trends that are not present or hiding trends that are (Walraven 1985). For example, the use of a rainbow scale to pseudo-colour continuous data (such as contour data) can hide the very trends that the user wishes to see. In contrast, approaching the problem from the perspective of human vision allows the user to exploit illusory effects to *improve* perception. One such example is provided later in this paper.

1.1. *Problem*

1.1.1. *High energy physics:* High energy physics poses particular problems for data visualization. For example, consider the data collected from ALEPH (see Figure 1), a particle detector in the Large Electron-Positron (LEP) collider located at CERN near Geneva (Decamp *et al.* 1990). LEP has the form of a circle of about 27km in circumference. Located in an underground tunnel, its function is to accelerate electrons and positrons in opposite directions around the tunnel up to energies that can reach almost 100 billion electron volts (100 GeV). The electrons and positrons are made to collide at a rate of several collisions per second, at specific locations along the ring, at energies sufficient to produce the heaviest known elementary particles. One collision every few seconds results in the annihilation of an electron-positron pair. When this happens, energy is concentrated in an extremely small volume, followed by an explosion into many particles, leaving the interaction point at almost the speed of light. The set of data describing all the outgoing particles is defined as an event.

The experiments rely on tracking chambers to provide sample points on the flight paths of the outgoing

0144-929X/98 $12.00

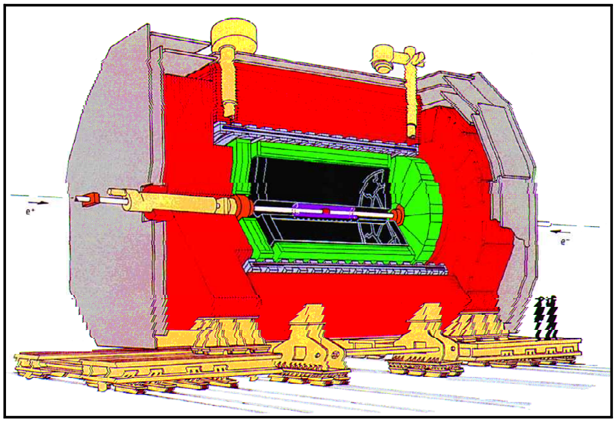

Figure 1. The components of ALEPH shown in a cut-away perspective view. The dimensions of the detector are approximately 12 ×12 ×12m, its weight is about 3000 tons and its cost in the neighbourhood of 70 million Swiss francs.

charged particles, and on calorimeters that measure the energy of all particles, independent of their charge. Only once a full reconstruction of these electronic signatures has been performed can the physics of the underlying collision/annihilation process be studied and identification of particles be performed. The results are then used to test existing theories on elementary particle research and to find new phenomena (Okun 1985).

The subject of our visualization challenge is *the event*, which comprises a digital description of the sample points of charged particles resulting from the collision, and of their energies when they are absorbed. These events are recorded with very high spatial precision (distances, which may be as low as 10 µm, are sampled at 100 000 locations), and hence each event generates a large amount of data. After data processing, an average event will contain the co-ordinates of around 2 500 points (rising to 4000 points for a more complex event). In reality, each of these points is described by an array of about 12 variables. About five million such events have been recorded since the ALEPH experiment began at CERN in 1989.

For data analysis, tracks are reconstructed using powerful mathematical techniques (Buskulic *et al.* 1995, Pusztaszeri 1996) and a small sample of events is then displayed graphically on computer workstations to provide a 'reality check' on the data. These graphical representations are also printed out on paper and acetate for presentation in papers and at scientific conferences and for publicity. The focus of this case study is one such system known as DALI (Drevermann *et al.* 1993). This has been used by hundreds of physicists to analyse the events from the ALEPH experiment.

1.1.2. *Model of colour vision:* To solve the problems posed by human perception and to optimize picture quality in DALI, we applied the following model of colour vision because of the abundant use of colour in DALI. The model, based upon research in physiology, psychophysics and cognitive psychology, is essentially a 'zone' model with three broad stages.

The first stage is trichromatic since it is based on the quantum catches in three types of cone photoreceptor (Dartnall *et al.* 1983). These three classes of cone have broad and overlapping spectral sensitivities (Stiles 1939) but may be loosely referred to as long-wave sensitive (L), medium-wave sensitive (M) and short-wave sensitive (S). Each class is effectively colour blind: it is unable to distinguish a change in chromaticity from a change in luminance (Naka and Rushton 1966, Baylor *et al.* 1987). To extract chromatic information from these signals, the outputs of the different classes of cone must be compared.

This happens at a second 'opponent' stage of colour processing (Hurvich and Jameson 1957). Two classes of opponent pathway have been isolated: one, (M-L), that differences the output of M and L cones; and a second, (M+ L)-S, that differences the output of S cones with some combination of M and L cones (Derrington *et al.* 1984). Luminance information, (M+ L), is multiplexed with the chromatic information carried in the M and L opponent pathway (Ingling and Marinez-Uriegas 1983, Lennie and D'Zmura 1988).

In the third stage of colour vision the two chromatic signals and the luminance signal are further analysed by higher-order colour mechanisms to produce the world of colour. Very little is known about these complex, higher-order colour mechanisms, other than that they have their neural basis at the cortical level (Gegenfurtner *et al.* 1994). More complete descriptions of this model, and the research findings that support it, have been published elsewhere (Kaiser and Boynton 1996).

There are many design issues with DALI that are important from the perspective of human vision. Some of these were resolved by empirical testing, but other issues remain. In this paper, the model of colour vision is used to interpret the design of DALI and show that it provides pointers to the design of other systems with extensive use of colour.

2. The anatomy of an event

In this section the elements of event display in DALI (Drevermann *et al.* 1995) that are relevant here are described and related to basic concepts in visual perception, especially colour vision.

2.1. *Detector views*

ALEPH comprises a series of concentric detector cylinders shown in figure 1. The components of ALEPH are three tracking chambers, a superconducting solenoid magnet, two calorimeters, which operate by measuring the deposited energy of the particles they absorb, and muon chambers, built specifically to identify high energetic muons. The three tracking chambers are the VDET (red in figure 1), ITC (purple) and TPC (black). It is mainly the TPC that is used to search for tracks of charged particles. The two calorimeters (green and red) absorb charged and neutral particles which leave behind a characteristic signature. The purpose of DALI is to make explicit the results of the particle collision.

The first difficulty concerns the visualization of the detector. Since the data are inherently three dimensional, it appears logical to present the results as a perspective three-dimensional projection, such as figure 1. For particle collisions, the perspective projection may prove useful to obtain overview (Shneiderman 1994), but an integral part of the physicist's task is to focus on individual tracks and hits. The cut away view in figure 1 is inappropriate to show tracks because the focus of interest is often the very centre of the detector, which in this representation is too small for the task. Additionally, when perceptually reconstructing a three-dimensional image from a two-dimensional picture of an event, some tracks that are actually distant can appear to be very close together. The net result is that the viewer must adopt perceptual hypotheses. With real-world objects, these perceptual hypotheses are usually (although not always) correct. But with data from particle physics, these hypotheses bear no relationship to the underlying event. Supplementing the event data with visual landmarks from the detector (as in figure 2a)[1]does not help and results in a more complex image which hides the very data we are trying to visualize unless the landmarks are reduced to a bare minimum. Wireframe images prove especially problematic at the high magnifications that are necessary to view individual hits. Alternative three dimensional effects can be achieved, for example, with perspective projections, shading and volume rendering, smooth rotations, and stereopsis. But despite numerous attempts to visualize the data from ALEPH in three dimensional realistic pictures (Drevermann *et al.* 1995), it has proved impossible to achieve an appropriate level of magnification without causing the visualization to collapse into

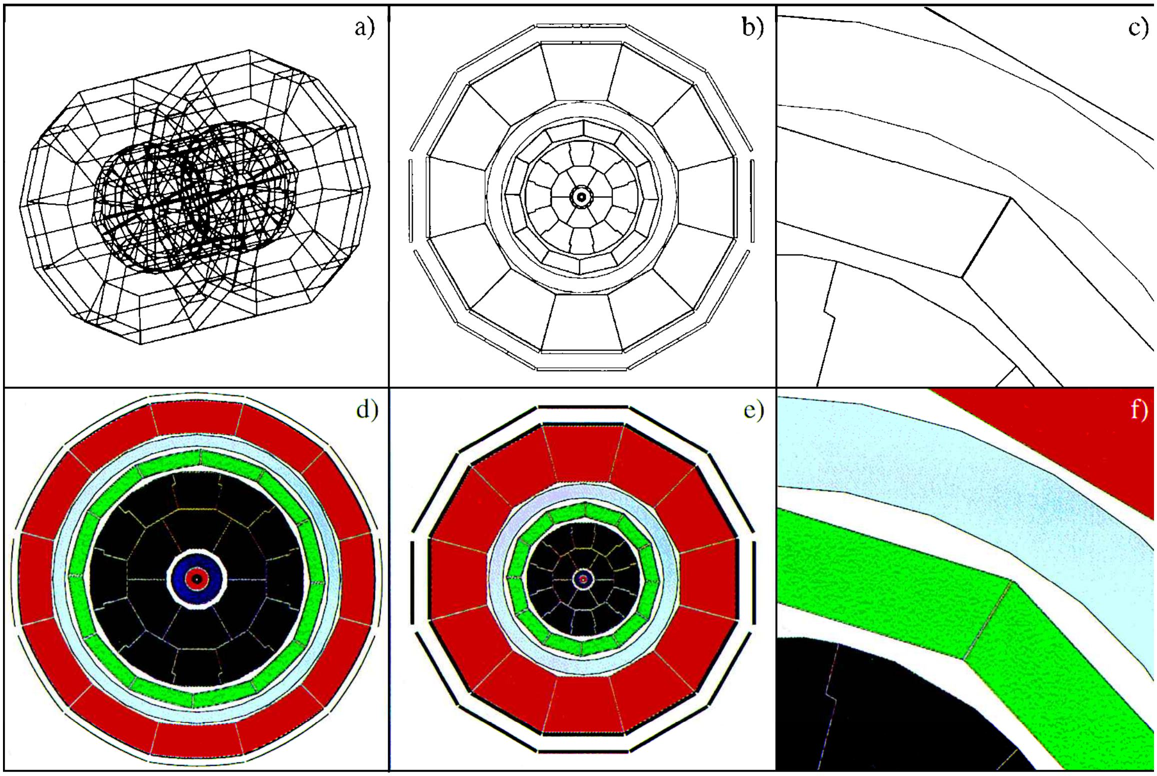

Figure 2. Six schematic views of the ALEPH detector: (a) wireframe perspective view, (b) wireframe cross-section. (c) segment through wireframe cross-section, (d) coloured cross section with 'fish eye' transformation, (e) coloured cross section, (f) segment of coloured cross section.

noise. Techniques from virtual reality (VR) prove to be useless: lines representing tracks dominate our recognition of points (hits) in both two- and three- dimensions. Therefore, it is vital to be able to identify tracks by looking at points only. This is much more difficult, time consuming and visually tiring with a VR projection than relying on the use of well chosen, two-dimensional, schematic and abstract projections.

Hence, DALI shows 2-dimensional cross-sections through the detector (one type of cross-section is shown in figure 2b). This view shows six layers of subdetectors which show up as six concentric rings. In figure 2e the visualization is further enhanced by colour coding the various elements using the same code as in figure 1. The white areas between the rings show regions where no detectors are present (in ALEPH these areas contain cabling, electronic devices and mechanical support structures). The grey ring is a superconducting solenoid magnet. The use of colour gives an improvement compared to the wireframe-only picture (figure 2b), especially if a section of the detector is shown (compare figures 2c and 2f).

One disadvantage with this view is that it provides equal emphasis to the outer subdetectors as the inner subdetectors. What is needed is a transformation that decreases the scale with increasing radius, so the outer subdetectors appear shrunk. This emphasizes the commonly used construction principle of detectors, namely that precision and sampling distance decrease when moving from the inner to the outer detectors.

This view was therefore transformed (figure 2e) to a schematic 'fish eye' picture (figure 2d), calculated in the following way. From the Cartesian co-ordinates X and Y the polar co-ordinates ρ and φ are derived. These are transformed to ρ_F and φ_F by:

$$\rho_F = \frac{\rho}{1 + \alpha\rho} \text{ and } \varphi_\mathrm{F} = \varphi$$

From ρ_F and φ_F the Cartesian co-ordinates X_F and Y_F are recalculated and drawn with a suitable linear scale to conserve the total picture size. The factor α is chosen interactively. Further justification for the fish eye view is given in the discussion of track visualization. The elements of the detector show up more clearly on screen if the background around and between them is shaded or coloured because colour helps the viewer segregate the various elements (see figure 3, which shows an actual event).

2.2. *Visualizing hits*

The second issue concerns the representation of 'hits', the small data points that show the paths of the various particles. Within the innermost three subdetector layers the hits are recorded with very high precision. To illustrate this for the TPC, consider figure 4 which shows six different two-dimensional transformations applied to better understand the structure of the underlying data. The length of the line in figure 4b corresponds to 70 mm in the detector itself. As the measuring precision within this area of the detector is about 0.1 mm, the detector is able to sample this range with an accuracy of one unit in 700. Yet, on a high resolution display, this line is just 100 pixels long and hence the display can sample the line at just one unit in 100. So here a single pixel represents seven error units. Moreover, on a high resolution display, a single pixel is often below contrast threshold, so the hits must be represented by at least a matrix of 2 × 2 pixels (in practice, a 3 × 3 matrix is used). This means that (for the current display) the size at which we can represent hits is limited by the human visual system, and not by the measuring precision of the detector.

Because the transformations in figure 4 affect the spatial geometry of the tracks, a visual cue is required so that the physicist can compare hits and tracks across the various views. Colour is an ideal solution to this problem since it is a proven technique for grouping and coding data (Christ 1977) and for visual search

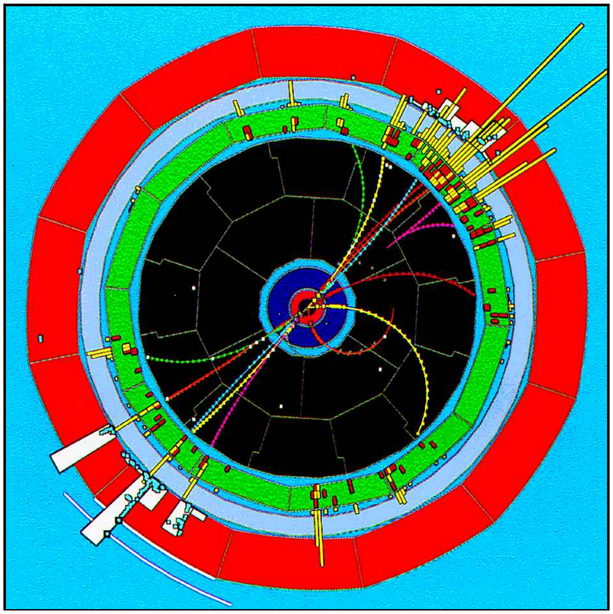

Figure 3. A typical event measured in the ALEPH detector at CERN. One sees clearly that there is a two-jet structure (at approximately 2 o'clock and 8 o'clock). The histogram peak at 6 o'clock with the absence of any track pointing to it is the typical signature of a neutral particle (identified as a photon as it has no hits in the outer, red, calorimeter). The bifurcating pink track at about 2 o'clock consists of two tracks originating from the decay of a neutral particle.

(D'Zmura 1991). The reader can judge the effectiveness of this by choosing a sequence of hits in figure 4a and then identifying the same sequence in the remaining projections (figures 4b–f). This task is virtually impossible without the use of colour.

However, the choice of colour needs to be optimized: hits on the same track should be easy to identify on the black background, but remain distinct from hits on different tracks. Various algorithms have been proposed to maximize colour difference on displays (Carter and Carter 1982, DeCorte 1990), but a significant disadvantage with them is that they require lengthy calibration procedures to define colours within CIE colour space. DALI is used on so many different colour workstations, at CERN and across the world, that to ensure colour calibration of all of these displays is simply not practical.

An alternative approach is to use basic colours. Psychologists and anthropologists have shown that people are quite happy to classify virtually all colours using the basic colour names white, grey, black, red, green, yellow, blue, pink, brown, orange and purple. Basic colours are easy to name (Berlin and Kay 1969), well separated in colour space (Smallman and Boynton 1990), and simple to produce on a display (Travis and Johns 1994). For DALI we chose to use basic colours to code the hits.

Hits may be categorized into three broad classes. The first class are unused by the analysis software; these are considered unimportant and are coded grey. A second class are unassigned; these can be potentially very interesting and are coded white. The third class are assigned to a particular track, and these are given one of the basic colours yellow, red, orange, pink, green and blue. The precise chromaticities of the colours assigned to the hits are given in table 1.

A complication here is that, although small pixels are necessary for spatial precision, they mitigate against accurate colour perception. The colour of the small points (around 4' of visual angle) is often very hard to identify unambiguously: it is difficult to separate yellow and white, green and blue, and red, orange and pink (see table 1 for the precise chromaticity co-ordinates of these colours). Figure 5 for example shows how a change in background colour alters our ability to detect lines and hits. This figure uses eight different colours for the lines and hits, and the colours are the same across columns: all that changes is the colour of the background. It is clear that, due to the effects of colour contrast, the

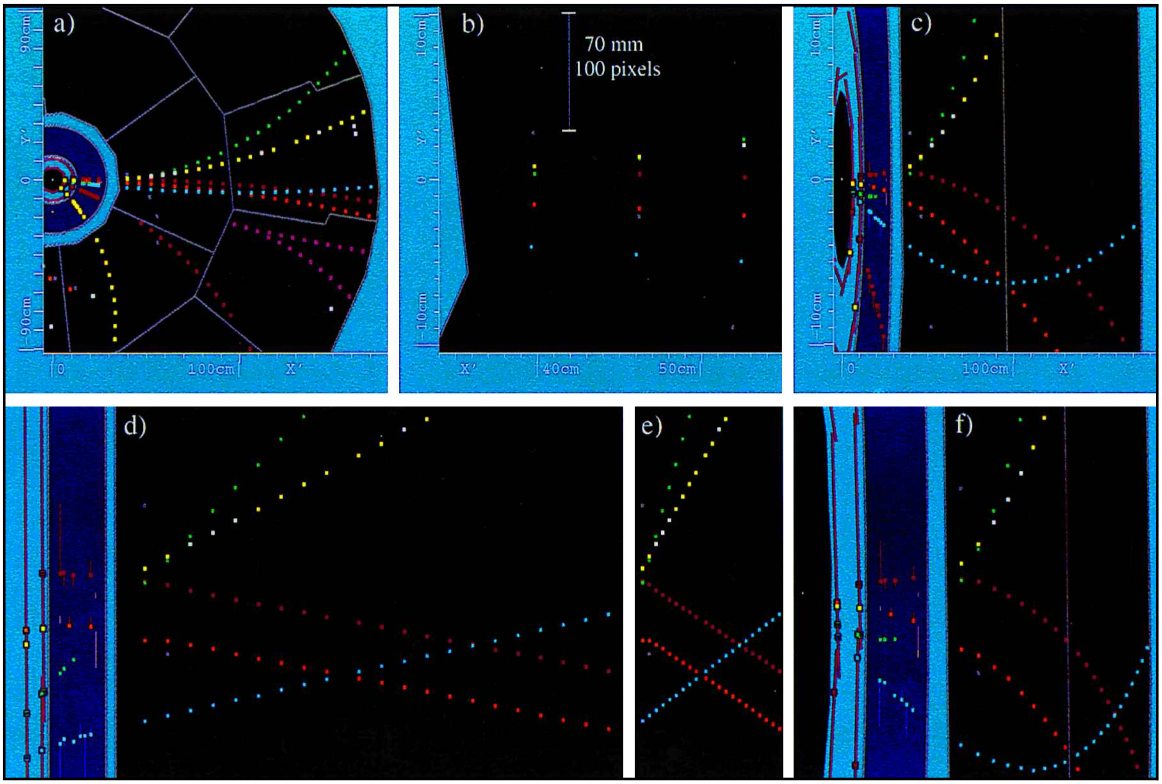

Figure 4. Six views of the same event. This figure shows how the use of colour helps the user compare hits and tracks across the various views. The reader is encouraged to identify a sequence of hits in (a) and identify this same sequence in the subsequent figures. Note also the improved perceptual grouping in (e) compared with (b) (see also figure 6). See text for more details.

background colour has a pronounced effect on the perceived colour of the lines and hits.

The main reason why such small images are hard to *discriminate* is because the very centre of the fovea of the human eye lacks S-cone receptors (Williams *et al.* 1981) making us essentially tritanopic for very small fields. Consider, for example, the yellow hit (the fourth square from the top in figure 5) on the white background. It is not possible to see this hit at a normal viewing distance. The reason for this is because the signals from the L and M cones are about the same for the yellow hit and the white background. It is the S receptors that are so valuable in making the yellow discrimination on a white background because the S signal differs significantly between these two colours. However, because there is a lack of S cones in the central fovea they cannot help us make the discrimination. The net result is that the hit is almost indistinguishable from the background.

Table 1. CIE chromaticity co-ordinates of the basic colours used in DALI (for background on the CIE method of colour specification, see Travis (1991). *Y* values are given in lux. Values were measured with a Minolta Chromameter II. For accurate measurements, the calibration patches used to measure these values (147 × 152 pixels) were larger than the 3 × 3 pixel matrix used to show hits.

Name	*Y*	*x*	*y*
Red	4.800	0.606	0.354
Green	19.100	0.277	0.616
White	24.900	0.272	0.329
Grey	10.200	0.268	0.320
Yellow	24.000	0.376	0.542
Orange	8.800	0.494	0.447
Pink	6.700	0.268	0.139
Blue	14.900	0.194	0.270

The main reason why such small images are hard to *detect* is because of insufficient luminance contrast between the hit and the background. Therefore black was chosen as a background for the central area of the detector. A black background also reduces dazzle effects from the screen that are noticed when a high luminance background, such as white, is used.

2.3. *Visualizing tracks*

A typical task faced by the physicist is to check track recognition in the TPC (the black region in figure 4) and then extrapolate these tracks to the inner two tracking chambers (purple and red). The standard projection (see figure 2d) shown in figure 4a provides a useful overview, but in order to discriminate the tracks within the TPC, the hits must be magnified considerably. The conventional approach to zooming is to use a linear magnification (such as in figure 4b) where the aspect ratio is held constant. However, this restricts the view so severely that track detection is no longer possible. An alternative approach is to compress figure

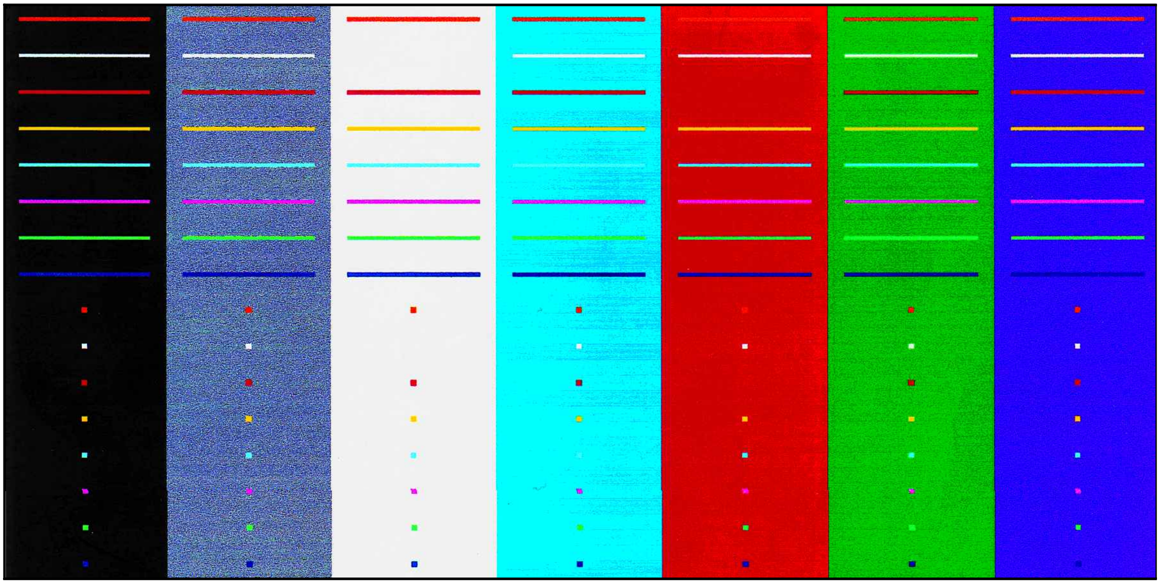

Figure 5. Demonstration of the effects of background colour on the perceived colour of lines (top) and hits (bottom). The physical colour of the lines and the hits is the same on each background, but the perceived colour varies. The line width (2 pixels) and hit size (3 × 3 pixels) used to produce this figure are those most commonly used in DALI. To demonstrate the effects of small-field tritanopia, move the figure away until the yellow hit disappears from the white background (third column). Note that the hit is still distinguishable on the other backgrounds.

4b horizontally (figure 4c). In this transformation, hits can be resolved and the tracks are recognized and seen over their total length. However track extrapolation to the innermost tracking chamber (VDET) still remains problematic.

For this purpose either the fish-eye view (see figure 2e) can be applied (figure 4f) or the image can be projected in polar co-ordinates (the φ/ρ projection, figure 4d, see also Drevermann *et al*. (1996).

In a typical session with DALI, some of the pictures are displayed side by side, but rarely one after the other. The projections shown in figures 4a, 4d, and 4e in particular are usually examined together.

It is clear from figure 4 that when a number of tracks overlap, they become difficult to distinguish from one another. This difficulty arises from the inability to perceptually group adjacent hits on the same track. The precise visual angle at which this occurs depends on the number of overlapping tracks, but for DALI the informal results suggest that it is when adjacent hits are separated by more than 2° of visual angle. The effect is particularly noticeable when comparing figure 4b with figure 4e and is represented schematically in figure 6. The left-hand sides of figures 6a and 6b show seven overlapping tracks, with and without connecting lines. It is difficult to perceive the horizontal tracks in figure 6b (left half of figure) unless the tracks are compressed (right half of the figure). The reason for this is the strong vertical perceptual grouping, related to the Gestalt 'Law of Proximity' (Koffka 1935, Bruce and Green 1990): elements that are closer together are perceived as a coherent object. This unwanted vertical structure results from the placement of the detectors which sample at discrete points. Adding colour (figure 6c and figure 6d) to encourage the perception of horizontal structure provides some help, but even with as few as two colours (figure 6d) the vertical spatial grouping is still more dominant than the colour grouping. Unwanted vertical grouping can be prevented by applying the Law of Proximity in favour of the horizontal structure, that is by compressing the tracks (right half of the figure, and compare figure 4b with figure 4e). Adding connecting lines (see figures 3 and 6a) also provides a solution. The latter, however, may be dangerous: connecting points with lines provides such a strong grouping cue that it may prevent the physicist from detecting track recognition errors introduced by the analysis programme.

2.4. *Visualizing energy deposit histograms*

Energy deposit histograms are used to identify particles. The histograms can be seen clearly in figure 3 (the yellow and white histograms at the top right of the figure). These histograms may overlap, and the design problem is to allow the physicist to switch attention between the different histograms.

A first step is to draw the frame of the underlying histogram on top of the upper histogram. Secondly, transparency (Reynolds 1994) is used (see figure 7, a magnified region of figure 3). Transparency is achieved

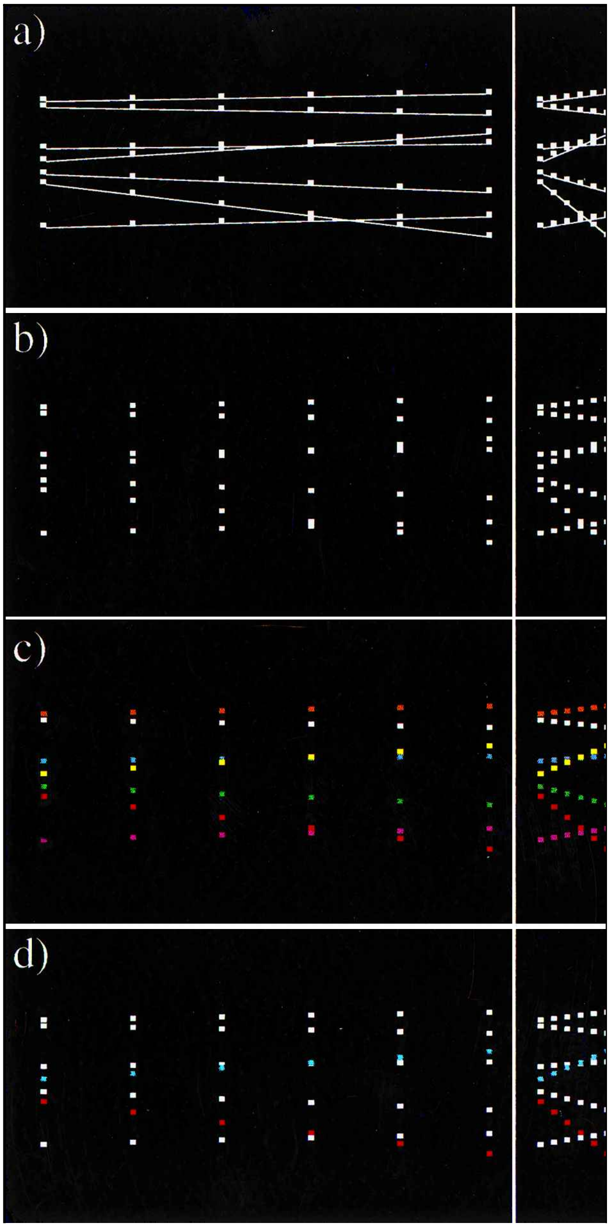

Figure 6. Demonstration of vertical grouping effects that hinder perception of horizontal tracks. Vertical spatial grouping is more salient than horizontal spatial grouping (b), left panel) because of the Gestalt Law of Proximity. Adding colour (c) and (d), left panel) only partially counteracts this tendency. Compression (right panels) and joining points with lines (a) left panel) regains the horizontal structure. See also figure 4.

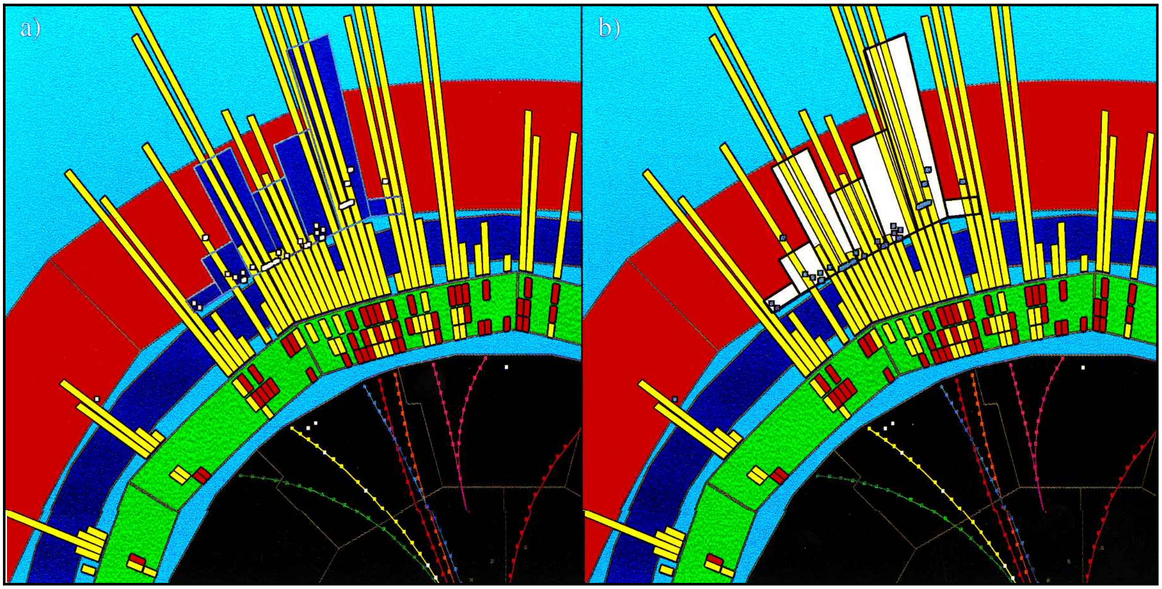

Figure 7. Use of colour assimilation to produce colour transparency. See text for more details.

by desaturating overlapping colours. For this application, rather than compute the region of colour overlap, transparency is achieved by using colour assimilation. Figure 7b shows the effect. Where the yellow histogram overlaps the white histogram it appears noticeably lighter due to the effects of colour assimilation. In reality, the physical colour of the histogram does not change. This is an example of how visual illusions can be used to *improve* our perception. Figure 7a shows that this effect is highly sensitive to the background colour: in this figure, there is no assimilation effect, and hence no effect of transparency. This makes it more difficult to switch attention between the two distributions.

3. Conclusions

The general approach used here, where design decisions are based on a knowledge of human vision, can be applied in almost any system intended to be viewed by end users. Typical applications include network management, air-traffic control systems, weather displays, and financial analysis applications (such as stock market predictors). The approach can also be profitably applied to user interface designs in general, including interfaces for the World-Wide Web.

Acknowledgements

The authors would like to thank the ALEPH Collaboration for the construction and running of the ALEPH detector, which provided the events for analysis. Special thanks are owed to Dieter Schlatter for his comments on the manuscript and Björn Nilsson for his important role in software development.

Note

[1] For clarity, the same event (i.e. the same particle collision) is used in the figures in this paper. Note also that the colour figures in this paper have been reproduced to simulate as close as possible the colour images viewed on screen. Inevitably, some differences between the printed figures and the real images have occurred. Versions of these figures suitable for inspection on screen can be viewed on the World-Wide Web at URL:http//alephwww.cern.ch/DALI/

References

Baylor, D. A., Nunn, B. J. and Schnapf, J. L. 1987, Spectral sensitivity of cones of the monkey *Macaca fascicularis*. *Journal of Physiology*, **390,** 145 – 160.

Berlin, B. and Kay, P. 1969, *Basic Color Terms* (Berkeley: University of California Press).

Bruce, V. and Green, P. R. 1990, *Visual Perception: Physiology, Psychology and Ecology* (2nd Edition). (Hove: Erlbaum).

Buskulic, D. *et al.* 1995, Performance of the ALEPH detector at LEP. *Nuclear Instruments & Methods in Physics Research A.*, **360,** 481 – 506.

Carter, R. C. and Carter, E. C. 1982, High-contrast sets of colors. *Applied Optics*, **21,** 2936 – 2939.

Christ, R. E. 1977, Review and analysis of color coding research for visual displays. *Human Factors*, **17,** 542 – 570.

Dartnall, H. J. A., Bowmaker, J. K. and Mollon, J. D. 1983, Human visual pigments: microspectrophotometric results from the eyes of seven persons. *Proceedings of the Royal Society of London Series B*, **220,** 115 – 130.

Decamp, D. *et al.* 1990, ALEPH: a detector for electron-positron annihilations at LEP. *Nuclear Instruments & Methods in Physics Research A*, **294,** 121.

DeCorte, W. 1990, Recent developments in the computation of ergonomically optimal contrast sets of CRT colours. *Displays*, **11,** 123 – 128.

Derrington, A. M., Krauskopf, J. and Lennie, P. 1984, Chromatic mechanisms in laternal geniculate nucleus of macaque. *Journal of Physiology*, **357,** 241 – 265.

Drevermann, H., Kuhn, D. and Nilsson, B. S. 1993, Non-conventional methods for the visualisation of events from high energy physics, in G. M. Nielson and D. Bergson (eds), *Proceedings of Visualisation '93*, San Jose, California, 386 – 390. IEEE Computer Society Press, Los Alamitos, California.

Drevermann, H., Kuhn, D. and Nilsson, B. S. 1995, Event display: can we see what we want to see? *Proceedings of the 1992 CERN School of Computing*, L'Aquila, Italy, CERN Yellow Report 95-05 and CERN/ECP 95-25. This report is also available on the Internet as http://www.cern.ch/CERN/Divisions/ECP/IPT/Papers/CSC95/EDisplay/G.E. Vandoni (Ed).

Drevermann, H., Kuhn, D. and Nilsson, B. S. 1996, Skewplot — a momentum sensitive visualisation of two-dimensional hits from very high multiplicity events. *Proceedings on the HEPVIS96 Workshop on Visualisation in High Energy Physics*. L. Taylor and G. E. Vandoni (Eds) CERN, Geneva, Switzerland, CERN Yellow Report 97-01.

D'zmura, M. 1991, Color in visual search. *Vision Research*, **31,** 951 – 966.

Gegenfurtner, K. R., Kiper, D. C., Beusmans, J. M., Carandini, M., Zaidi, Q. and Movshon, J. A. 1994, Chromatic properties of neurons in macaque MT. *Visual Neuroscience*, **11,** 455 – 466.

Gregory, R. L. 1977, *Eye and Brain: The Psychology of Seeing* (3rd Edition) (London: Weidenfeld and Nicolson).

Hurvich, L. M. and Jameson, D. 1957, An opponent-process theory of colour vision. *Psychological Review*, **64,** 384 – 404.

Ingling, C. R. and Martinez-Uriegas, E. 1983, The spatiochromatic signal of the r-g channel, in J. D. Mollon and L. T. Sharpe (eds) *Colour Vision: Physiology and Psychophysics* (London: Academic Press), 433 – 444.

Kaiser, P. K. and Boynton, R. M., 1996 *Human Color Vision.*

Koffka, K. 1935, *Principles of Gestalt Psychology (New York: Harcourt, Brace and Company).*

Lennie, P. and D'Zmura, M. 1988, Mechanisms of color vision. *CRC Critical Reviews in Neurobiology*, **3,** 333 – 400.

Naka, K. I. and Rushton, W. A. H. 1966, S-potentials from colour units in the retina of fish (*Cyprinidae*). *Journal of Physiology*, **185,** 536 – 555.

Okun, L. B. 1985, *Particle Physics*, (New York: Harwood Academic Publishers).

Pusztaszeri, J.-F. 1996, Combinatorial algorithms for pattern recognition in composite tracking chambers. Thèse No. 1530, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (in English).

Reynolds, L. 1994, Color for air traffic control displays. *Displays*, **15,** 215 – 225.

Rogowitz, B. E. and Treinish, L. A. 1993, An architecture for rule-based visualisation, in G. M. Nielson and D. Bergson (eds) *Proceedings of Visualisation '93*, San Jose, California, 236 – 243.

Shneiderman, B. 1994, Dynamic queries for visual information seeking. *IEEE Software 11*, **6,** 70 – 77.

Smallman, H. S. and Boynton, R. M. 1990, Segregation of basic colors in an information display. *Journal of the Optical Society of America*, **A7,** 1985 – 1994.

Stiles, W. S. 1939, The directional sensitivity of the retina and the spectral sensitivities of the rods and cones. *Proceedings of the Royal Society of London Series B*, **127,** 64 – 105. Reprinted in Stiles, W.S. 1978, *Mechanisms of Colour Vision*, (London: Academic Press).

Travis, D. S. 1990, Applying visual psychophysics to user-interface design. *Behaviour & Information Technology*, **9,** 425 – 438.

Travis, D. S. 1991, *Effective Color Displays: Theory and Practice* (London: Academic Press).

Travis D. S. and Johns, A. 1994, Back to basics. *Proceedings of the Society for Information Display*, June 1994, San Jose).

Tufte, E. R. 1990, *Envisioning Information* (Connecticut: Graphics Press).

Walraven, J. 1985, The colours are not on the display: a survey of non-veridical perceptions that may turn up on a colour display. *Displays*, **6,** 35 – 42.

Wickens, C. D., Merwin, D. H. and Lin, E. 1994, The human factors implications of graphic enhancements for the visualisation of scientific data: dimensional integrality, stereopsis, motion and mesh. *Human Factors*, **36,** 44 – 61.

Williams, D. R., MacLeod, D. I. A. and Hayhoe, M. 1981, Foveal tritanopia. *Vision Research*, **21,** 1341 – 1356.