

3D creativity for the Web

Harold Thimbleby
Middlesex University, London
harol d@mdx. ac. uk

Abstract

Autostereograms on the Web make a very appealing combination:

- The World Wide Web brings together art, science and technology, both in the medium and in its content. Autostereograms, likewise: based in science, they are art made possible by computer graphics technology.
- The Web is extremely accessible: it makes multimedia available to more people than any other computer-based medium. Autostereograms, similarly, are easily viewed without any special equipment and give true stereoscopic 3D effects.

This paper discusses stereoscopic 3D effects, and how drawing autostereograms lends itself to the networking constraints and page layout freedom characteristic of the Web. There are wonderful opportunities for making science and art interesting and accessible to wide audiences.

1. Introduction

There are many ways to add depth to our experiences. If we wish to add depth literally, rather than metaphorically, we can do so in all sensory channels. Sound, for example, can be used to give an impression of distance, and even some clues about the environment between the listener and the sound source. Vision, however, is used for creating depth most easily and consistently between different people. Techniques for creating visual depth can be divided conveniently between those requiring two eyes (binocular) and those where only one eye is sufficient (monocular).

Closing one eye hardly makes any difference to our perception of depth, thus proving the power of monocular depth cues, such as size, perspective and texture. Stereoscopic vision (i.e., 3D binocular vision), however, requires each eye to see a different image—by definition, since if they saw the same image, what was seen could be done with monocular vision.

It is easy to have a binocular view *without* any three dimensional percept or impression. The children's illusion of making a hole in the hand (by looking with one eye through a tube next to the hand, and the other eye looking at the hand) is a typical example: the two different images combine in an interesting but stereoscopically flat way.

Although Leonardo da Vinci made comments about perspective, he did not appreciate stereoscopic vision. Stereoscopic three dimensional vision was probably first described by le Père Cherubin of Orleans in 1677—and was then neglected until Charles Wheatstone invented the *stereoscope* in 1832. His gadget attracted the interest of Queen Victoria, and became immensely popular. It also created a lot of scientific interest since it was an instrument that permitted new and interesting psychological experiments, and its appearance coincided with increased instrumental approaches to psychological experimentation.

Surprisingly, many seemingly ‘elementary’ questions about stereopsis—how we see with two eyes—still remain to be explored.

In real life, a solid object is viewed from slightly different angles and therefore each eye gets a slightly different image of it. Crucially for stereoscopic depth viewing, the images are similar (unlike the hand and hole images), and the visual system can *fuse* them, creating the impression of depth. To reproduce a 3D picture, then, requires two sufficiently similar images.¹ One approach is to place the images side by side (Figure 1) and direct each eye’s gaze to one image—this is a popular approach for VR systems where each eye has its own miniature monitor.




Figure 1. Daguerreotype stereoscopic still life by Antoine Claudet (1797-1867), 1850s. Reproduced with permission. Copyright: Museum of the History of Science, University of Oxford, Oxford.

3D, and 3D for the Web

There is a large range of technologies for binocular viewing, which we will not discuss further here (for example, using shutter glasses)—because they all rely on additional gadgetry, to which relatively few Web users will have access. Simple optics (e.g., prismatic lenses) make viewing stereograms easier, but we shall prefer free viewing, without needing further equipment for use on the Web. Without practice, it is quite tricky to see stereograms unaided, because the eyes have to converge at a distance different to the plane of the stereogram; when this is done, the stereogram is at the wrong distance to be in focus. Fortunately, the

¹ Strictly, two different retinal outputs are required, not two different images. If one eye is dark adapted and the image is moving, then the different response times of the eyes lead to a stereoscopic depth effect. This is the Pulfrich Phenomenon. (It is very entertaining to dance when wearing glasses with one dark filter.)

slight effort contributes enormously to the mystique and popularity of free viewing stereoscopic techniques.



Figure 2. Popular virtual image illusion. If the left eye looks at the cage and the right at the bird, a single bird-in-cage image is seen—as well as two ‘spare’ parts to the left and right. Stereoscopically, the bird in cage lies in a different plane from the picture itself (behind it, if viewed as described).



Figure 3. When the eyes’ images have corresponding features, the apparent depth of the combined image can be controlled. Here, the nose will appear closer than the rest of the face (cross-eyed viewing reverses the depth).

Until the Bela Julesz experiments (1971) with random dot stereograms in the 1960s, there was no experimental evidence to prove that humans have strictly binocular depth perception that works in the absence of any monocular depth cues. Julesz invented random dot stereograms, which have no monocular information. When viewed with two eyes, a binocular percept is created, and it has depth. The apparent depth is compatible with the simple physics for the light paths for corresponding points in the eyes’ visual fields—provided the two images are not too different. (Just what “too different” means is interesting to establish: it depends on what the images are, whether they are moving, and other factors.)

As late as 1967, Ogle and Wakefield argued that there was a possibility of monocular contours in one eye’s random patterns, and they doubted Julesz’s results. It is likely they did not fully understand the computer-based approach to drawing random patterns, which need not leave any artefacts that manual methods would more likely have done.

Like conventional stereoscope pictures, random dot stereograms are made from two separate pictures, like the stereogram of Figure 1. They are constructed starting from a single random dot field, with the second image copied from the first but laterally moving some dots left or right depending on the 3D image depth required. As can be appreciated from Figure 1, it is not possible to make the images much larger since they would overlap.

Alternative optical arrangements could be made to avoid overlap; Wheatstone’s early stereoscope used mirrors to allow the images to be larger but to lie in different planes (1838). Ideally, for our purposes where we have no

control over the viewer's equipment (it could be anywhere in the world), we would rather stereograms that did not require special viewing equipment.

Chris Tyler (1983) recognised that random dot stereograms can be effectively overlaid, so creating *single image* random dot stereograms (abbreviated SIRDS, also called autostereograms). Effectively, the right hand side of the left image is designed to carry on into the left hand side of the right image. Tyler devised an algorithm (discussed by Stork & Rocca, 1989) that left no visible joint. Moreover, the Tyler algorithm allowed a stereogram to be viewed from any position, permitting stereograms to be made arbitrarily large. (Previously, not overlapping the eyes' separate images required either small stereo pairs, or optical devices: either choice leading to restricted viewing conditions.)

Online discussion of autostereograms can be found in al t. 3d and Inglis (1996); printed references include Thimbleby & Neesham (1993); Horibuchi (1994); Richardson (1994); Watkins & Mallette (1996). Tyler & Clark's (1990) very creative paper is particularly recommended for further reading. A recent review of the literature is Davis & Hodges (1995).

2. Wallpaper effect

Brewster described the wallpaper effect in the last century (Brewster, 1844) A horizontally repeating pattern (such as Figure 4) can be viewed with the eyes incorrectly converging (for the distance to the wallpaper), so that the eye images are out of phase. The pattern row being viewed therefore appears at a distance other than its physical distance.




Figure 4. A simple wallpaper pattern (showing one complete tile, which could be the pattern for one vertical strip of paper).




Figure 5. Using the image from Figure 3, it is possible to create an indefinitely repeating pattern. Two tiles are shown. Alternate images will have their noses in and out, though for psychological reasons this is extremely difficult to see. Pictures like Figure 3 are therefore interesting to use as a Web backgrounds, and can be tiled across the entire page.

By adjusting the repetition rate, the apparent distance of the pattern can be controlled. Sophisticated wallpaper effects can be achieved when the width of the vertical strips (i.e., the wallpaper itself) is the lowest common multiple of all horizontal repetition periods.

Figure 4 shows a simple wallpaper pattern, with variation of depth vertically up the page (i.e., each horizontal row is seen at the same depth). Figure 5 shows an

example wallpaper pattern, used to give variation of depth laterally as well as vertically.

3. Autostereograms

Consider a horizontal plane passing through a viewer's eyes and the autostereogram image. We assume the autostereogram is vertical² and parallel to the line between the eye's optic centres. Normally the eyes would converge onto the plane of the autostereogram, and each point on the autostereogram would be imaged in corresponding points in each eye's retina. The autostereogram would appear flat (notwithstanding monocular depth interference). However, suppose the eyes diverge so that they are separately fixated on matching points on the autostereogram texture. Since the retinal images correspond, the brain fuses a virtual image beyond the autostereogram. The distance of this image can be varied by altering the geometry of the autostereogram: placing the corresponding points closer makes them appear nearer in depth, and *vice versa*.

It is also possible to see an autostereogram with hyper-converged eyes. The virtual image is then between the viewer and the autostereogram, rather than behind it. An autostereogram designed for one style of viewing will be depth-reversed when viewed by the other method. The slight difference in geometry, apart from the inversion, is generally not noticeable, since variation due to distance and parallax are visually more salient.

Geometrically, hyper-converged viewing can be used at greater distances, but monocular cues have greater precedence and viewing is quite difficult (i.e., the autostereogram is manifestly a flat sheet in the distance). Convergent viewing is appropriate for computer use, since converging beyond a monitor screen is easy—provided it is slightly reflective—simply converge at one's own reflection which lies behind the monitor screen.

4. Autostereograms on the web

Although some people argue that there is better 3D in the world than there ever will be on the Web (and they beg, please leave the Web bandwidth to them)! 3D is fun and one of the obvious ways the Web can engage and entertain people around the world.

Patch High School, Stuttgart, are using autostereograms both as a science project in its own right, and as a Web project (Dawson, 1996). Patch School is using the Web so that visitors can add their comments and suggestions: this will not only make their Web site more interesting, but will also encourage the school's pupils as they see how their work is attracting interest from around the world.

Obviously the Web can be used to publish autostereograms, and there are some nice ones out there. But the Web has constraints. We want to minimise use of bandwidth to make pictures arrive faster.

Autostereograms can be drawn at very low resolution, for example using characters in place of pixels (cf. Figures 4 and 6).

² We ignore vertical parallax in the explanation.



Figure 6. Autostereogram made at low resolution using characters rather than pixels.

But if we want graphical realism, unfortunately, autostereograms don't have any useful properties that lend themselves to better compression. We could send a texture and a 3D model, compressed in many ways, but we would lose some of the advantages of autostereograms—we would need to assume the client side had appropriate software, and we would be trading bandwidth for the user spending time doing the 3D rendering. We lose the main reason to being attracted to autostereograms: they can be viewed without special equipment.

However, by using backgrounds the client (browser) can tile a full page by repeating a smaller image. The HTML code `<BODY BACKGROUND="file">` does this nicely for an image in the file *file*. This speed-up can make autostereograms much more attractive for frequent use. We therefore require wallpaper autostereograms—and an algorithm that can tile properly.

Tiling autostereograms

Tyler's paper (1983) raised the possibility of creating wallpaper, where the right side of a single image stereogram could be continued with its own left edge. However, his algorithm was based on a unidirectional scan, so it could not make the left and right sides match, except in the ingenious case of using alternate 3D models adjacent to complimentary models, which then cancel out their lateral shift on the dot pattern. His paper gives a good example of an autostereogram tile made in this way.

Tyler pointed out that discontinuous autostereogram tiling causes visual effects at the joints (e.g., lustre); artists may wish to use these interesting effects as an integral part of their design.

The Tyler algorithm takes a strip of texture and copies it left to right (or right to left) across the picture. The width of the texture is stretched or shrunk according to the 3D depth model. The result is a continuously varying apparent depth, according to the depth model. Unfortunately, the right hand side of the

picture—because of its stretching and shrinking—has no relationship to the left hand side.

Thimbleby, Witten and Inglis argued the Tyler algorithm could be improved with a different approach (Thimbleby, Witten & Inglis, 1994; Thimbleby, 1996). Their approach lends itself to constraining the left and right hand sides of an image to be equal, and therefore is one way to create wallpaper autostereograms. The geometry requires certain pairs of pixels to be equally coloured.³ The algorithm finds the implied equivalence classes. To extend to tiling, we add the constraints that certain right pixels equal certain left pixels, as if the tiled plane extended indefinitely in each direction. That is geometrically sufficient to cater for horizontal tiling; unfortunately the smaller the tile (so the faster it can be downloaded) the harder the constraints are to solve and leave any artistic freedom in the content of the image.

Vertical tiling is trivial: merely requiring the top and bottom of the autostereogram texture to correspond, and this can always be achieved provided the 3D model corresponds along the line of the joint (i.e., equal depth and curvature at the edges), since this implies the constraints are identical at the joint.

Drawing autostereograms

An autostereogram is a texture that when viewed in the right way reveals a 3D depth percept.




Figure 7. Example components of an autostereogram.

Although the autostereogram below (Figure 8) looks like a conventional stereogram, the background texture is continuous across the picture. The background is an autostereogram, which happens to represent the flat surface on which the butterfly is resting. The image could be extended left and right indefinitely, perhaps to represent additional 3D objects. If it did so, then the two monocular butterfly images would have to be ‘recycled’ in the rendering of the other objects.⁴

³ Equal luminance is sufficient.

⁴ Line artifacts in the image are due to it being drawn at too low a resolution for the printer—it looked OK on the screen at 72DPI! (The artifacts that arise in printing autostereograms can be exploited to reduce copyright problems but not diminish peoples’ appreciation of autostereograms on the Web!)




Figure 8. Small example autostereogram, derived from data shown in Figure 7.

Unfortunately, the constraints implied by representing *both* a 3D model and doing so on wallpaper are severe: the design of an image that ‘works’ is quite tricky. There are various ways of reducing the constraints without significantly reducing the impact of the autostereogram image. For example, by introducing noise, some constraints can be dropped—this results in fewer constraints but a blurring of the 3D image. Another method is to insert points at infinity into the 3D model. Interestingly, although this makes the 3D image harder to find, some people report even greater fascination with the effect. Maybe ordinary autostereograms have got too easy for them?

5. Other techniques

Autostereograms, being a combination of 2D and 3D images, lend themselves to combine several layers of illusions and depth effects.

Simple illusions

Autostereograms create a variety of straight forward optical surprises. View an autostereogram so that the 3D image is clearly seen. Rotate the paper or head about the axis going through the head and perpendicular to the image: the 3D image will vanish. The converse of this surprise explains why some people cannot see autostereograms: if the autostereogram is not held with its patterns parallel to the line of eyes, then nothing can be seen anyway.

Rotation of an autostereogram about axes in the plane of the image (e.g., like starting to turn a page in a book over) create fascinating parallax effects. Normally when a 3D object is rotated, the eyes obtain changing views of it. In the autostereogram, although the image is rotated, the views do not change as they would for a solid object. The 3D image seems to follow the observer. A hand pointing directly out of the paper is quite dramatic!

Copy an autostereogram onto a transparent sheet, and view it until you really appreciate the solidity of the image. Now turn the sheet over and look again. You expect the 3D image to be viewed from behind, because you are looking at the back of the image: but the 3D object has ‘turned around’ and is still the same way out.

Low resolution autostereograms necessarily create a noticeably quantised 3D image. Prolonged viewing of a naturalistic shape (such as a face) results in the image becoming smoother. This and other illusory contour illusions (such as 3D versions of the Kanizsa triangle) are easy to reproduce.

Combined textures and autostereograms

The artist has relatively free choice in the use of the texture in an autostereogram. Historically—what a word for a subject less than a decade old!—

three main sources have been used: random fields, interesting textures (e.g., naturalistic textures, like leaves), and textures based on rendering the 3D models being represented in the autostereogram itself. It is also possible to use additional 3D effects in the texture itself, such as narrowing or other distortions for perspective effects.

Of course, an autostereogram can be drawn using a texture that is a naturalistic scene: for instance with sharp foreground and hazed distance, and many have been done like this, particularly with the autostereogram image hidden in various features, such as stony beaches or in forest leaves.

An interesting possibility is to use textures that are themselves illusions, such as the smiley face shown earlier, or Hermann grids.

Anaglyphs

Anaglyphs are common 3D pictures, making use of colour filters to separate the two eyes' images from a single picture. Typically, anaglyphs use red and green filters in spectacles (anaglyphoscopes): the red-filtered eye seeing any green in the picture as black, and *vice versa* for the other eye. Anaglyphs can be used on the Web like any other picture, but they require the use of special (though widely available) spectacles.

The algorithms we have discussed for autostereograms also work for anaglyphs, which results in an autostereogram that can be viewed easily with spectacles, and can also be viewed *without* spectacles! Hence people having difficulty viewing autostereograms can start with the aids. Moreover, auto-anaglyphs can be tiled, and therefore used efficiently in Web backgrounds.

Other illusions

Autostereograms can be made more interesting by using them to show optical illusions such as the Muller-Lyer illusion (arrows that look different lengths).

A single image from a random dot stereogram contains no depth information (since the depth information is solely in the correlation with the other image in the pair). Therefore if a 3D model of an optical illusion is made, and is still visible as an illusion when seen in stereo, one can be certain that the illusory effect does not invoke single eye processing effects. This remains true for autostereograms (although Thimbleby *et al.* showed that a simple 3D model can be recovered by monocular processing, they used a process that does not correspond to any known physiological process).

The Necker Cube (Figure 9) is a simple image that can be made into an autostereogram (e.g., Figure 10). The normal Necker Cube is an ambiguous 2D projection of a 3D object, and it can 'flip' as you see it in different ways. Curiously, to see a Necker Cube in an autostereogram requires stereovision, which would therefore represent it unambiguously, but to most people it still has depth ambiguity. These flat-but-3D models of Necker cubes evoke great interest in viewers. Evidently the construction of the 'solid' percept is created independently in the visual pathway of the brain and subsequent to the 3D reconstruction of the image.




Figure 9. Necker Cube, normal 2D representation. Where is the ball? You may be able to ‘move’ it from the far face to the front face—or anywhere in between.




Figure 10. Autostereogram of a Necker Cube.

The Ebbinghaus illusion⁵ is shown in Figure 11. Some people report that one inner circle looks larger than the other (they are geometrically of equal size, as can be easily confirmed by measurement). When converted to a floating image in an autostereogram, they appear to be the same size to many people. Some very interesting research by Aglioti, DeSouza and Goodale (1995) showed that, although people may see the inner circles as different sizes, when asked to grasp them, their fingers move to grip their actual physical sizes. These experiments would be interesting to repeat where there is no monocular information, as in the autostereogram.⁶

⁵ In the US it is also known under the name Titchener. (Titchener and Ebbinghaus were both students of Wundt. I'm grateful for M. A. Goodale for pointing this out.)

⁶ The distance to the 3D image will depend on whether the (auto)stereogram is viewed by converging or diverging the eyes, which will affect the apparent size of the ‘object’ one attempts to grasp.




Figure 11. Ebbinghaus illusion.

It may be that the Ebbinghaus illusion ‘fails’ in autostereograms because of the modified effect of size constancy scaling. Size constancy scaling is a perceptual compensation for perspective: distant objects do not look as small as purely geometric considerations would lead one to think. If, for some reason, the visual system thinks the circles are at different distances, then their size would be scaled. In the flat stereoscopic case, perhaps the visual system does not apply scaling since the circles are manifestly at the same distance. Whereas, in Figure 5, we saw that knowledge of human faces overrode what the visual system should have been seeing: it was as if, the 3D image processing said “this is 3D and the nose on the face sticks in” and then knowledge said, “if it’s a nose and in 3D it sticks out of the face.” Disentangling all the influences (using appropriate experimental controls) is one of the great challenges.

Whatever the answers, it is likely that there will be interesting cultural variation in the experience of the illusion—normal or stereo—for example, it is known that people living in open, flat regions have noticeably different concepts of perspective than people living in geometric spaces such as cities like New York. Such questions would be ideal to experiment with given the international and cross-cultural reach of the Web.

6. Further work

Being able to draw autostereograms and being able to draw ones that look artistically pleasing are different skills! The constraints of wallpapering autostereograms makes artistry even more elusive. We may choose to think of this as a drawback, or as a way of enhancing the value of the image.

At a recent presentation of some of the ideas in this chapter, the audience raised many interesting questions that I do not know the answer to. For example: Does autism affect stereo vision? How do different forms of colour blindness affect it? Can anyone see anaglyphs but not autostereograms (whose fusion requires monocular contours)?

There is a lot of scope for experimentation, and new discoveries. Autostereograms bring the edge of science into reach of everyone; it is still a young subject. Moreover, the Web is the largest collaboration of humans the world has ever seen. It must therefore be the largest laboratory available—and the natural interest in autostereograms guarantees a large sample of subjects! So if you design autostereograms (or other experiments) on the Web, register your pages with search engines (so people know they are there) and *do* include ways to collect feedback (using email or CGI scripts).

7. Conclusions

Autostereograms are popular art, ideal for Web posters and Web page backgrounds. We showed how autostereograms can be drawn for backgrounds using the net bandwidth effectively. Even for people who are stereo-blind, who cannot appreciate the hidden 3D images themselves, autostereogram backgrounds can still be attractive art. For those who can see autostereograms, they are like sculpture: they look attractive from a distance, and the effort of viewing them is like the 'effort' of running one's hands over a physical object. The delay in seeing an autostereogram, the effort—whether for artists or viewers—even their 'exclusivity,' contribute to their addictive effect.

8. References

- S. Aglioti, J. F. X. DeSouza & M. A. Goodale (1995) "Size-contrast illusions deceive the eye but not the hand," *Current Biology*, **5**, pp679–685.
- D. Brewster (1844) "On the knowledge of distance given by binocular vision," *Transactions of the Royal Society of Edinburgh*, **15**, pp663–674.
- E. T. Davis & L. F. Hodges (1995) "Human Stereopsis, Fusion, and Stereoscopic Virtual Environments," in W. Barfield & T. A. Furness III, editors, *Virtual Environments and Advanced Interface Design*, pp145–174, Oxford.
- M. Dawson (1996) Patch High School, Vaihingen, Stuttgart, Germany. Email: mdawson@patch-ahs.dsi.net, Math, Art, and Music URL: <http://192.253.114.31/Projects/stereograms/stereogram.html>
- S. Frank (1995) *Kai's Power Tools 3.0*, MetaTools Inc.
- S. Horibuchi, ed. (1994) *Stereogram*, Boxtree Ltd.
- S. Inglis (1996) Waikato University, New Zealand, Stereograms FAQ, URL: <http://www.cs.waikato.ac.nz/~singlis/sirds.html>
- B. Julesz (1971) *Foundations of Cyclopean Perception*, University of Chicago Press.
- D. Richardson (1994) *Create stereograms on your PC*, Waite Group Press.
- D. G. Stork & C. Rocca (1989) "Software for Generating Auto-random-dot Stereograms," *Behaviour Research Methods, Instruments & Computers*, **5**, pp525–534.
- H. W. Thimbleby, S. Inglis & I. H. Witten (1994) "Displaying 3D Images: Algorithms for Single Image Random Dot Stereograms," *IEEE Computer*, **27**(10), pp38–48.
- H. W. Thimbleby & C. Neesham (1993) "How to Play Tricks with Dots," *New Scientist*, **140**(1894), pp26–29.
- H. W. Thimbleby (1996) "An equivalence class algorithm for drawing autostereograms," *Software—Practice & Experience*, **26**(3), 309–325.
- C. W. Tyler (1983) "Sensory processing of binocular disparity," in C. M. Schor & K. J. Ciuffreda, editors, *Vergence eye movements: Basic and clinical aspects*, pp199–295.
- C. W. Tyler & M. B. Clarke (1990) "The Autostereogram," in J. O. Merritt & S. S. Fisher, editors, *Stereoscopic Displays and Applications, Proceedings International Society for Optical Engineering*, **1256**, pp182–196.
- C. D. Watkins & V. Mallette (1996) *Stereogram Programming Techniques*, Charles River Media.
- C. Wheatstone (1838) "Contributions to the physiology of vision. Part the first: On some remarkable, and hitherto unobserved, phenomena of binocular vision," *Philosophical Transactions of the Royal Society*, **128**, pp371–394.

9. Appendix: Adobe Photoshop filters

A Photoshop filter has been written by the author to construct autostereograms. Working within Photoshop, it acquires all the professional image manipulation features of Photoshop (and incidentally becomes much easier to program, since the filter does not have any baggage such as file handling or image format conversion to do itself). The filter is available from the author.

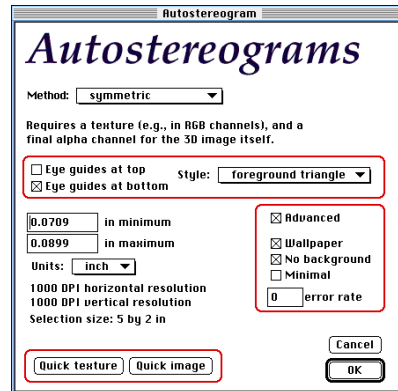


Figure 9. Photoshop filter control panel.

Kai's Power Tools, a wide selection of commercial Photoshop filters (Frank, 1995), provide a simple autostereogram filter, *3D Stereo Noise*. It is very limited: the texture has to be random dots (either monochrome or colour). Apart from the eye separation and 3D image, there is no creative freedom. However, KPT's tool has an educational advantage over ours: it provides a fast preview that you can look at while changing the eye guide separation. This helps an artist get an intuitive feel for the constraints on autostereograms, and, in particular, the way in which eye guide separation can affect texture streaking.

There is a vast amount of shareware for drawing autostereograms. The reader is recommended to do a search on the Web, perhaps starting from Inglis's Stereograms FAQ Web site (see references).

Author's biographical notes

Harold Thimbleby is Professor of Computing Research at Middlesex University, London. He became interested in autostereograms when he was shown one by a Japanese friend. He then developed a symmetric algorithm when experiments showed that asymmetry caused problems for some viewers. His visit to Prof Ian Witten at Waikato University, New Zealand, in 1994 coincided with growing world-wide interest in autostereograms. He can be found on the Web at <http://www.cs.mdx.ac.uk/harold>.