

Magic Pictures

6pm Gresham College, 30 October, 2003

Harold Thimbleby
Gresham Professor of Geometry

Mysterious '3D Art,' 'magic eye' or 'holusions' make fantastic poster art. The posters look like meaningless patterns of dots or patterns, but spring into an amazing and convincing three-dimensional image when they are viewed correctly. Technically, these pictures are called *autostereograms*. An autostereogram is a single picture which the eyes can decode as a three dimensional object. (In contrast, a *stereogram* is two pictures, one for each eye, that creates a three dimensional image.)

If you shut one eye, the world does not suddenly appear to go flat: the brain is doing a lot more than relying on two eyes to see in three dimensions — it uses all sorts of cues, such as size and shadows, to build up a sense of distance and depth.

Artists know many ways of getting a depth effect: the size of objects (larger being closer), perspective and lighting give subtle hints. Shadows can be used to make things appear **closer** than the background their shadows fall on. Overlapping occludes (hides or partially hides) objects that are further away. Then there is our use of world knowledge: for example, we know that a human that looks the same size as our outstretched fist must be further away, as humans are known to be higher than hands. And so on.

However, when we look at a solid object with both eyes, we do see something extra that is not possible to see in an ordinary flat picture of a 3D scene, which looks the same to each eye. In three dimensions, looking at a solid, each eye sees it slightly differently because the eyes are in different positions each side of your head. Of course, in an ordinary flat picture each eye sees the same flat thing.


For the brain to build the 3D sensation, each eye must receive the proper image, just as it would have done when looking at a solid. There are several ways to do this. Historically, the stereoscope was very popular. It was invented by the physicist Sir Charles Wheatstone in 1832 (see also my Gresham Lecture on mobile phones, 28 November 2002). The stereoscope is a contraption, somewhat like a large pair of binoculars that presents each eye with a separate image (usually a photograph). Queen Victoria was entranced by them, and they became very popular. Although Victorian photographs were usually monochrome, a stereoscope can produce 3D images in full colour.

Rather than use a stereoscope to get a 3D effect, another way is to use red and green filters, called (wonderfully!) anaglyphoscopes. Here, a single picture is printed in two colours, red and green. Red is used for one eye's image and the green for the other. By printing a picture in two colours and using filters to look at it, the eyes end up with different pictures, as required to get a 3D effect. An anaglyph picture is seen in monochrome, because the colours are needed exclusively to make the filters work. There are some other techniques but they are more difficult to print: polarised light, holography, or tiny prisms. The prism method is sometimes seen on children's toys, and quite often on postcards.

But without separating the eyes' two images more subtle techniques are required. There are many techniques, some involving mirrors (and some smoke), but autostereograms do not require any special equipment, so they are much more accessible. They can be drawn by hand, although, as we shall see, they are ideal for playing with computers.

Conventional stereograms usually have to be viewed with a stereoscope (like a pair of binoculars) or with coloured eye filters. Also, the size of a stereogram is usually limited, because the two pictures cannot be so large that they overlap (although you can use mirrors or lenses to get round this problem). Autostereograms, however, need no special viewing equipment, and can be made as large as desired. (I've made some several metres across for large posters.) An autostereogram works everywhere you look at it, however large it is. Thus, it is possible to make an autostereogram as wide as you like, perhaps even for wallpapering a room. It is possible to arrange for the pattern on the left and right hand sides of an autostereogram to be identical, and if you do this, you really could wallpaper a room — sharing one side of a sheet for its picture as well as for the picture on the next sheet along. There are

easier ways of getting a simple 3D effect from wallpaper, which doesn't let you create arbitrary 3D images, but just different depths for different rows of wallpaper patterns. This is the *wallpaper illusion* — and it is very easy to construct wallpaper illusions in word processors: try going 'wall eyed' with the following pattern (which was typed directly in Microsoft Word, just using different font sizes):



To understand how autostereograms work, imagine looking through a window at something outside. Let's say you can see an apple. If you look with your left eye only, the apple appears a bit to the left of where it looks than if you look at it with your right eye. (This is actually quite confusing, because you may think the window moves to the right instead! But the effect is the same.) Each eye sees the apple differently, and gets a different angle on what you are looking at. Although the apple can be seen in two different places through the sheet of glass looking with the eyes alternately, when you look with both eyes at the same time, the brain cleverly fuses the two apple images into a single image, which it understands as the solid apple beyond the glass.

Suppose an artist has one eye shut and draws the apple on the glass, then shuts the other eye and draws the apple again. The artist will have drawn two apples. If you could look at these drawings of apples in the right way, you could fool yourself that there was a real apple behind the glass. All autostereograms do is replace the glass window with a sheet of paper, and draw the apple, or whatever, once for each eye in the appropriate place. But one has to be *very* careful about drawing things that overlap.

To solve the problem about overlap, imagine the apple is painted with black and white dots, as in Figure 2. If we choose the black and white dots very carefully, the dots needed for the right-hand side of the left apple can be *exactly* the same as the left-hand dots needed for the right-hand apple. Figure 3 shows what this looks like, and it is indeed a 3D picture that you could see as a 'real' spherical apple. Choosing the dots properly and positioning them is an interesting problem in geometry and programming.

The second problem is the spare parts of the apple that each eye sees. Consider Figure 1, and imagine the apples have been drawn onto the glass: the right-hand eye certainly sees an apple, but it also sees the whole apple drawn for the left eye, which leaves over a crescent shaped texture. The left eye sees a spare right-hand crescent, and the right eye a spare left-hand crescent. These textures have to be recycled in other parts of the scene, in fact just as the left part of the image was recycled into the right part of the other image. Ultimately the entire picture is made up of dots, recycled from one place to another! This recycling generates a lot of equations that a computer is well placed to solve.

Finally, the amount the apple seemed to move as you looked at it with your eyes alternately depends on its distance away. So by adjusting the distances between dots, it is possible to give the impression that the object is at various distances. In the figures here, we have given the apple a spherical surface, but it could have had any shape, and we could have therefore represented any 3D object. However, there are psychological issues on the geometrical limits that restrict what the brain is able to fuse — if the apple images are too far apart, they cannot be perceived as originating from a single image further away, but they will be seen as two separate images that are nearer.

Autostereograms are certainly interesting art, but they are also interesting from psychological and computing perspectives.

Autostereograms do not (and cannot) make any sense at all when viewed with one eye alone, as you can easily see! Yet when viewed as 3D pictures, they can show any sort of object, including optical illusions. Thus the illusory effects (for the illusions that work in this way) must occur in the brain *after* the eyes' two images are combined. Figure 6 shows an example of the Top Hat illusion, which is a 'top hat' that looks taller than it is wide, but in fact is equally tall as it is wide. That this illusion and many others work in autostereograms, proves that their effect is not simply caused in your eyes or your retinas, because your individual retinas cannot see 'anything' — the illusion must happen after the

brain has worked out how to reconstruct the three dimensional object. Since we know how the optic nerves run through the brain and where they combine, we can pinpoint, as it were, where the illusion occurs.

All autostereograms made from dots or discrete symbols necessarily create rough images of 3D objects. Yet when you look carefully at some pictures, they may appear to be smooth, even in places that aren't there! This effect is particularly noticeable with curved shapes, because the dots are square and cannot possibly show a smooth round object. This optical illusion is closely related to the Kanizsa triangle, shown in Figure 7. You may get this effect in Figure 9, which does not really have layers of neat circular discs: the 'discs' are jagged.

*

These lecture notes give a complete program for doing your own autostereograms, which you can also download from Harold Thimbleby's web site. The notes don't give many deep psychological, geometrical or programming details (which the lecture will cover), but... if you have access to the web or some equally big encyclopaedias, here are some phrases worth looking up:

For historical ideas, start with *Bela Julesz*, one of the first people to look closely at random dot stereograms. He was particularly interested in what random dot images tell us about vision. (Tonight's lecture will give some earlier examples from the 1940s.)

For psychological issues: look up *binocular fusion*.

From a computer programming point of view, the recycling of the dots in the optimal way is a constraint satisfaction problem, and solving it led to a new and efficient algorithm for the *union-find* or *equivalence class* problem. If you are a serious mathematician or computer scientist, my article in *Software—Practice and Experience* gives you good background. There are easier ways to approximately solve the equations, but these result in images that are harder to see; hopefully the lecture will have time to discuss them.

Looking up *autostereogram*, *SIRDS* (single image random dot stereogram), *3D*, and so on will quickly collect a vast amount of information. Some images are animated, and those are great fun.

Further reading

- H. W. Thimbleby & C. Neesham, "How to Play Tricks with Dots," *New Scientist*, 140(1894), pp26–29, 1993. (Popular account, and provides a nice image and some program code to do it yourself.)
- H. W. Thimbleby, I. H. Witten & S. Inglis, "Displaying 3D Images: Algorithms for Single Image Random Dot Stereograms," *IEEE Computer*, 27(10), pp38–48, 1994. (Provides background and a complete explanation of programming autostereograms.)
- H. W. Thimbleby, "An Efficient Equivalence Class Algorithm with Applications to Autostereograms," *Software—Practice & Experience*, 26(3), pp309–325, 1996. (A more formal approach than the 1994 paper.)
- C. W. Tyler & M. B. Clarke, "The Autostereogram," *SPIE Stereoscopic Displays and Applications*, 1258, pp182–196, 1990. (This is the original paper introducing autostereograms.)




Figure 1. An apple seen through a window.




Figure 2. The apple itself, well how it might look if we gave it a random texture of black and white dots.




Figure 3a. The apple of Figure 1 as it would be drawn on the window. If you can see it (e.g., by crossing your eyes) you'll see a three dimensional apple, surrounded by two crescent shaped textures, one on each side. Note how the right hand side of the left circle is the same pattern as the left hand side of the right circle, simply because they overlap. Since the two circles are images of the same object, then the left and right hand sides of the apple must be the same.




Figure 3b. A complete autostereogram, containing the apple image from Figure 3a. If you look carefully you can see common patterns between the Figures, and how they are 'recycled' repeatedly across the picture.

znwjt kmeoyrkgcjuoznwjtkmeoyrkgcjuoznwjtkmeoyrkgcjuoznwjtkmeoyrkgcjuoznwjtkmeoyra
 qwumqmftztxyadhlljqwumqmftztxyadhlljqwumqmftztxyadhlljqwumqmftztxyadhlljqwumqmftztxo
 vouuwanwebjnawgkkvouuwanwebjnawgkkvouuwanwebjnawgkkvouuwanwebjnawgkkvouuwanwebje
 ydmhczknicuiddxwqydmhczknicuiddxwqydmhczknicuiddxwqydmhczknicuiddxwqydmhczknicus
 dpqcqahaqkwxlxsdlpqcqahaqkwxlxsdlpqcqahaqkwxlxsdlpqcqahaqkwxlxsdlpqcqahaqkws
 ffwmgxmbfltyuhtuwffwmgxmbfltyuhtuwffwmgxmbfltyuhtuwffwmgxmbfltyuhtuwffwmgxmbflta
 idzvitpeeexrlckduidzvitpeeexrlckduidzvitpeeexrlckduidzfvitpeeexrlckduidzfvitpeeexry
 letlyxquulmzbkiljletlyxquulmzbkiljletlyxquulmzbkiljletdlyxquulmzbkiljletdlyxquulmzw
 rzcpbrezzaypppydrzcpbrezzaypppydrzcpbrezzaypppydrzchpbrezaypppydrzchpbrezaypk
 urioynbbbtsrvkcpvurioynbbbtsrvkcpvurioynbbbtsrvkcpvurinoybbbtsrvkcpvurinoybbbtsrs
 bykacdxbbizlbnrmabykacdxbbizlbnrmabykacdxbbizlbnrmabyksacdxbbizlbnrmabyksacdxbizls
 hmofcjzggavcndjukhmofcjzggavcndjukhmofcjzggavcndjukhmodfcjzggavcndjukhmodfcjzggavcx
 behrpooyijjatmvefbehrpooyijjatmvefbehrpooyijjatmvefbehgrpooyijjatmvefbehgrpooyijjal
 jkavdbowwzjjhbzpfjkavdbowwzjjhbzpfjkavdbowwzjjhbzpfjkagvdbowzjjhbzpfjkagvdbowzjjv
 ixceuiyhhenctcgloixceuiyhhenctcgloixceuiyhhenctcgloixcreuiyhhenctcgloixcreuiyhencf
 jghvxfaelyshnsqjghvxfaelyshnsqjghvxfaelyshnsqjghbvxfaelysqjghbvxfaelysy
 khhdwvqqdvnwyervkhhdwvqqdvnwyervkhhdwvqqdvnwyervkhhdwvqqdvnwyervkhhdwvqqdvnq
 uvbsojjpshdivyjonuvbsojjpshdivyjonuvbsojjpshdivyjonuvbsoejpshdivyjonuvbsoejpshdip
 yofocpparjrbyrtguyofocpparjrbyrtguyofocparjrbyrtguyofocaparjrbyrtguyofocaparjrbh
 xxvmllektfplastwkxxvmllektfplastwkxxvmllektfplastwkxxvmllektfplastwkxxvmllektfplf
 nxiosvjjxcnmsufpgnxiosvjjxcnmsufpgnxiosvjjxcnmsufpgnxioslvjjxcnmsufpgnxioslvjjxcnmu
 fifnxggsbxgbbhggjffifnxggsbxgbbhggjffifnxggsbxgbbhggjffifnxggsbxgbbhggjffifnxggsbxgf
 oarojbpamrrzzucbtoarojbpamrrzzucbtoarojbpamrrzzucbtoarojbpamrrzzucbtoarojbpamrrs
 vwaiedseqqxfzvhpvcvwaiedseqqxfzvhpvcvwaiedseqqxfzvhpvcvwaiedseqqxfzvhpvcvwaiedseqqxy
 qqsgyqpotqipugeafqqsgyqpotqipugeafqqsgyqpotqipugeafqqsgyqpotqipugeafqqsgyqpotqir
 hjfdompiskrrgenhxhjjfdompiskrrgenhxhjjfdompiskrrgenhxhjjfdompiskrrgenhxhjjfdompiskrw
 lpnklzbdctjivyewnlpnklzbdctjivyewnlpnklzbdctjivyewnlpnklzbdctjivyewnlpnklzbdctji
 lihkrwoqcvsolkqrclihkrwoqcvsolkqrclihkrwoqcvsolkqrclihkrwoqcvsolkqrclihkrwoqcvsi
 cjuoznwjtkmztxokgcjuoznwjtkmz-|-kgcjuoznwjtkmz-|-kgcjuoznwjtkmztxokgcjuoznwjtkme
 dhlljqwumqmfefjeyadhlljqwumqmfefjeyadhlljqwumqmfefjeyadhlljqwumqmfefjeyadhlljqwumqmf

Figure 4. Instead of using dots or colours, here we've used letters (which are available on all computers). The autostereogram shows a 'Top Hat' floating above the paper that seems taller than it is wide, which would prove the height illusion occurred in your brain **after** the 3D image was reconstructed — except that in this case the aspect ratio of the text has not been corrected! That is, the lines of text here are taller than the characters are wide, so the top hat *really* looks taller, because it is. However, when you do correct for height, the illusion still works.



Figure 5 Put the bird in its cage.




Figure 6. Autostereogram of a top hat. The 3D image here is a top hat. The hat is in fact as wide as it is high, but it does not appear to be so — which is a standard optical illusion. You cannot see anything at all when you look at the picture with one eye, so the illusion you see must happen after the images from your two eyes have been combined.




Figure 7. A Kanizna triangle. A light white `cake slice' with a smooth and continuous outline appears to float in front of the picture. Kanizna triangles work in autostereograms, though, curiously, the 3D sensation appears to be very much smoother than they really are, with your brain filling in the missing bits, just as it has filled in the missing parts of the white triangle.

Program it yourself

A complete program for constructing autostereograms is shown below, which is written as a JavaScript program for a web page — you should be able to run it on any computer with a web browser if you

store it in an HTML file. This program constructs autostereograms out of coloured letters (a, b, c...), which can be drawn on any computer (any symbols can be used, with the proviso that they are the same width — that's why the <tt> font is used here). See Figures 4 and 9 for examples of typical results.

```

<html><head><title>Text autostereograms</title></head><body>
<h1>An autostereogram in JavaScript...</h1>
<tt><script>
// Draw an autostereogram, using letters
// © Harold Thimbleby, 1993-2003.

var    maxE = 15,      // separation of eyes, in characters. About 2.5 inches
        minE = 10,
        maxX = 80,    // width of image, in characters
        maxY = 30;    // height of image, in lines

var aspect = 2; // there are 6 characters/cm horizontally, 3/cm vertically; so aspect ratio is
2
function image(x, y) // return object depth at (x,y), a number minE to maxE
{
    var a = (x-maxX/2)/aspect, b = y-maxY/2;
    var r = Math.sqrt(a*a+b*b); // radius from centre
    r = r/(maxY/2-1);
    if( r > 1 ) return maxE;
    return Math.round(minE+r*(maxE-minE)); // so: r = 0 -> minE, r = 1 -> maxE
}

// colours and letters are allocated in turn, so it looks nicer!var colourcount = 0,
lettercount = 0;
var colours = ["maroon", "red", "fuchsia", "green", "lime", "yellow",
    "aqua", "teal", "blue", "olive", "gray", "black"];
var letters = ["H","a","r","o","l","d","T","h","i","m","b","l","e","b","y"];

function randomChar() // return a letter and colour, randomly (using HTML codes as desired)
{
    if( colourcount >= colours.length ) colourcount = 0;
    if( lettercount >= letters.length ) lettercount = 0;
    return "<font color="+colours[colourcount++]+>"+(letters[lettercount++])+"</font>";
}

var same = new Array(maxX), colour = new Array(maxX);

for( var y = 0; y < maxY; y++ )
{
    for( var x = 0; x < maxX; x++ )
        same[x] = x;
    for( var x = 0; x < maxX; x++ )
    {
        var s = image(x, y), i = Math.round(x-s/2), j = i+s;
        if( 0 <= i && j < maxX )
        {
            var pi = same[i];
            while( pi != i && pi != j )
            {
                if( pi < j ) i = pi;
                else { i = j; j = pi }
                pi = same[i];
            }
            same[i] = j;
        }
    }
    for( var x = maxX-1; x >= 0; x-- )
        document.write(colour[x] = same[x] == x? randomChar(): colour[same[x]]);
    document.write("<br>");
}

var leftSpot = Math.round((maxX-maxE)/2), rightSpot = leftSpot+maxE;
for( var x = maxX-1; x >= 0; x-- ) // draw the convergence guides...
    document.write(x == leftSpot || x == rightSpot? "|": "&nbsp;");

</script></tt><p>
To view, look at the two vertical lines at the bottom of the image and stare into the distance,
so that the two lines merge. (It may take a while to do this!) You will then actually see
<i>three</i> lines: a merged central line, and two lines on each side of it. When you've done
this, your eyes are at the correct angle to view the background of the autostereogram. With
care, you will be able to move your gaze around the whole picture, and you will see the three
dimensional image in it. The picture here is a 'pyramid' made of five discs, coming out towards
you (or away from you if you cross your eyes).
</body></html>

```

Figure 9 shows what that program shows (the real output is in rainbow colours):



Figure 9. A stack of discs, output copied from the JavaScript program run in a web browser. The 3D image is very deep, and may be quite hard to find, particularly because of all the steps in distance. Once you’ve worked out what the image is, you will be able to keep it in view more easily as you look around it. Have a close look at the edges of the discs, and try and work out what shape you *think* you see. (NB The image above may have been reproduced in black and white; the program generates an autosterogram in colour, which looks prettier. Colour *per se* does not affect perception of the 3D image.)

What does the program do? It draws each row of the image separately, driven by the geometrical model provided by the function `image`. (Here, `image` is pretty trivial.) The `image` routine specifies how far apart matching pixels or letters should be to form part of the 3D image. If at coordinate (x, y) , there are two pixels each side at distance $\pm s/2$ that should match, then the program essentially solves a system of many very simple equations of the form $\text{colour}_i = \text{colour}_j$, where i and j are $x-s/2$ and $x+s/2$ respectively. In the example, each row of the picture has about 70 such equations. The value of s can be found by simple geometry, depending on the viewing conditions and the depth of the image required at the point (x, y) . The solution is not fully determined, so arbitrary colours and shapes can be assigned to some points along the line: that is the role of the function `randomChar`, which in the example gives successive letters from my name. On some lines, the constraints of the equations allow my name to appear more or less readably!

This is a very brief discussion (you should try drawing some pictures to see what is going on) — but the lecture should make all clear.

Viewing autostereograms

It may take some time to see anything in autostereograms! You have to look *at* the picture (the window) to see the dots in focus, yet you have to deconverge your eyes as if looking *through* the window to see the reconstructed image. The key problem is that we are not accustomed to focusing and accommodating at different distances: you have to focus on the image, but converge our eyes as if the image was further away.

You can practice converging your eyes with Figure 5, which is the familiar trick of getting the bird in the bird cage. When your eyes are deconverged, you see the bird inside the bird cage. Achieving this feat is exactly the same problem as seeing the 3D posters, though it is very much easier as you know what you are doing!

As you deconverge your eyes, you tend to focus into the distance, because that is where you would normally be looking. Consciously bringing the picture back into focus probably ruins the attempt at looking beyond it. If you wear glasses to correct for short sight, you may have an advantage if you remove them: your natural focus will be closer than the corresponding distance of the patterns. You will therefore have little difficulty diverging your eyes and keeping the dots in focus at the same time. (Of course, your glasses may also correct for other sight defects so removing them may not help much.)

Another method is to photocopy the pictures onto transparent sheets. Hold the sheet half way between your eyes and a plain white wall. The 3D picture should be seen 'on the wall.' By viewing in bright light you reduce the differences in focusing needed for the sheet and the wall.

Using a computer program for drawing autostereograms you can repeatedly display the same picture using different random numbers. This makes the dots shimmer, but does not affect the geometry of the image. Thus, there is no stable image for your eyes to converge onto, and deconverging to see the depth effect is very much easier. If you have the equipment to do it, you can easily animate autostereograms to view 3D movies, though you need good random numbers to stop them creating distracting artefacts from frame to frame.

Note that the geometry only works if the autostereogram you are trying to look at is upright with respect to your head. The image's horizontal lines must be in the same planes as your two eyes.

If you still can't see autostereograms, try anaglyphs (red/green pictures). These are much easier to see. If you can't see 3D depth in anaglyphs, it is possible that you have a binocular vision problem.