

# Dichromatic Steganography

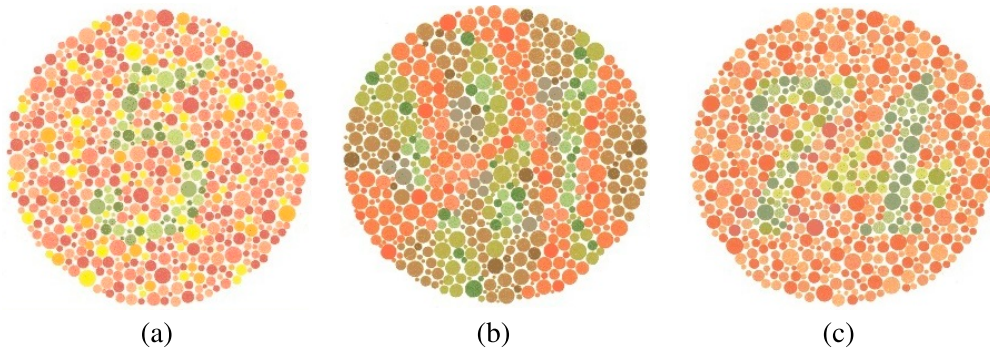
Craig S. Kaplan

School of Computer Science, University of Waterloo, Ontario, Canada; csk@uwaterloo.ca

## Abstract

To a person with colour vision deficiency (colour blindness), some pairs of colours that other people can easily tell apart will be perceived as similar or identical. In the right context, this deficiency can be harnessed as an unusual ability. In particular, some combinations of colours can serve as camouflage, allowing a design to conceal information that is nearly invisible to a viewer with normal colour vision, but comprehensible by someone with the right form of colour vision deficiency. I explore the possibility of using dichromatic vision as a vehicle for a subtle form of steganography, present tools for processing digital images to achieve this effect, and discuss the possibilities of this rather specialized aesthetic domain.

## Introduction



**Figure 1:** Three sample Ishihara plates: a “vanishing” plate (a), a “hidden image” plate (b), and a “transformation” plate (c).

Colour vision deficiency (CVD), also known as colour blindness, is an anomaly of the human visual system causing an individual to be unable to distinguish as many colours as the general population. Most people with CVD have difficulty distinguishing reds and greens, as well as other combinations such as blues and magentas. Less common varieties lead to confusion between other sets of colours, or reduce all vision to shades of grey. CVD affects about 5%–8% of men and 0.5%–1% of women. It is generally regarded as a minor inconvenience, exacerbated by default assumptions about colour vision in technology and design.

The Ishihara Tests [4], including the three examples shown in Figure 1, are the most popular diagnostic tool for CVD, not to mention iconic and attractive works of graphic design. Each one is a pattern of coloured dots, with colours chosen to communicate potentially different information to subjects with and without CVD. In Figure 1(a), for example, trichromats (i.e., people without CVD) see the digit 5, but a person with red-green colour blindness sees a field of monochromatic dots. Many software tools exist to produce images like Figure 1(a) [3], or for simulating CVD in order to reveal potential issues in colour design [2].

Far more interesting to me are Figures 1(b) and (c), which contain hidden information visible *only* to a person with CVD. In both cases, a number is hidden in the dots; a trichromat will see random fields of colour

in (b) and a different number in (c). This hidden information is made possible through a kind of camouflage. The trichromat's eye is overwhelmed by the profusion of colours, but these colours collapse down to a smaller set for a person with CVD, making the hidden number easier to discern. Biologists have speculated that there may be an evolutionary advantage to maintaining a sub-population of a species with CVD—they can alert the group to camouflaged predators [8].

In this paper I explore the possibility of creating art targeted at people with CVD, specifically the form of red-green colour blindness called *deuteranopia*, a form of *dichromacy*. Because generated images contain information concealed from trichromats, I refer to this practice generically as *dichromatic steganography*. Such art is not altogether new. Working by eye, the Hong Kong artist Carol Man created a series of paintings in the spirit of the Ishihara tests, containing hidden content for people with CVD [7]. However, I am not aware of other prior examples of CVD-based steganography, nor of software tools for automating the creation of such images. In this paper I develop a model and set of algorithms, inspired by the Ishihara tests and past analyses of them. After a brief overview of the necessary elements of colour theory, I present techniques for generating vanishing, hidden, and transformation images, corresponding to Figures 1(a), (b), and (c).

## Colour Perception and Metamers

Although the emission spectra of real-world objects can conceptually be arbitrary functions, most people experience colour as a three-dimensional phenomenon. The implied reduction in dimensionality is a product of human physiology and psychology, and not of light itself. Humans retinas contain three kinds of *cone cells*, which respond differently to photons of different wavelengths. They are usually called Long, Medium, and Short (or *L*, *M*, and *S*), in reference to the wavelengths of light at which their responses peak. To a first approximation, an object's colour corresponds to the response levels of the three cone types to the spectrum of light emitted by the object.

Each cone type has a *response curve*, giving its normalized excitation level to each wavelength in the visible spectrum. The cone's total response to an incoming spectrum can be computed by integrating the product of the spectrum and the response curve. It follows that these response curves form a basis for a three-dimensional subspace of the space of all emission spectra, and an object's colour is determined by the coordinates of the projection of its spectrum onto that subspace. I refer to this subspace as *LMS space*.

Obviously, many spectra will project down to the same point in LMS space. Such spectra are called *metamers*: distinct functions that are perceived as the same colour. Technology for reproducing colour is ultimately based on synthesizing a weighted mix of a small set of predetermined spectra called *primaries* (e.g., the red, green, and blue components of a pixel) to produce a metamer for the emission spectrum of a given light source. There is no need to reproduce the original spectrum in full: by definition, a metamer provokes the same total response from the viewer's cone cells, and they will perceive the same colour.

LMS space is useful for modelling human vision, but not practical for most colour manipulation tasks. The standard for computers is RGB, in which three numbers determine the brightnesses of a pixel's red, green, and blue primaries. Closely related is sRGB, in which those three numbers are passed through exponential transfer functions to account for the historical nonlinearity in brightness of real-world display devices. The Rosetta Stone for colour representation is CIE XYZ, developed in the early 20th century based on human experiments matching colours to a set of three light sources with known spectra [5, Chap. 28].

Colour vision deficiency usually takes the form of *anomalous trichromacy*, in which one of the three cone types is deficient, or *dichromacy*, in which it is absent altogether. In particular, *protan*, *deutan*, and *tritan* CVD correspond to deficient *L*, *M*, and *S* cones, respectively. In this paper I will focus on deutan CVD, and specifically on *deuteranopia*, the complete absence of *M* cones and the most common form of dichromacy. The effectiveness of my results will diminish for individuals with progressively milder *M* deficiencies (known as *deuteranomaly*). The algorithms presented here adapt readily to protan and tritan CVD as well.

LMS space offers a model in which to understand the experience of dichromatic colour. A person with deuteranopia perceives colours in a space of two dimensions rather than three, based on the responses of their  $L$  and  $S$  cones alone. Two colours in LMS space that differ only in  $M$  will be metamers to a deuteranope but not to a trichromat. More generally, to each combination of  $L$  and  $S$  there corresponds a *confusion line* of colours that ranges over all possible  $M$  values. For example, some shades of red and green lie on one such confusion line, and therefore cannot be distinguished by a deuteranope. Thus deuteranopia is one form of “red-green colour blindness”.

## Vanishing Images

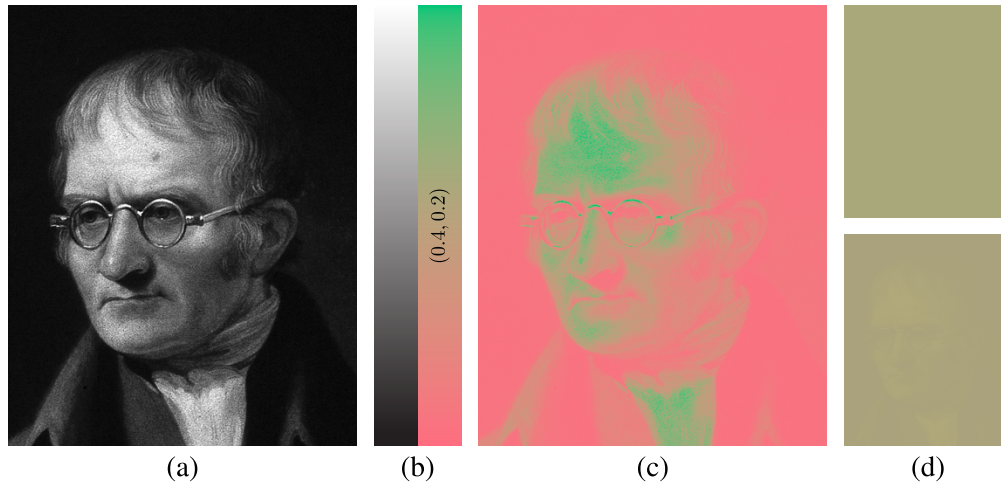
We are now in a position to understand the construction of the subset of the Ishihara tests known as “vanishing” plates, in which a foreground shape or object is visible to trichromats but invisible to dichromats. For a given form of dichromacy, we need only choose a palette of colours that lie on a single confusion line. Then, in an arrangement of dots, we colour the dots that make up the foreground using colours from one end of the confusion line, and colour background dots using colours from the other end. A trichromat will easily distinguish the two families of dots and identify the foreground object; a dichromat will see dots all of a single colour. Figure 1(a) gives an example of an Ishihara vanishing plate, with a foreground object in shades of green and the background in shades of red and yellow.

Ishihara’s work predated standardized colour spaces like CIE XYZ; he developed his plates experimentally, with feedback from human subjects. With modern mathematical and computational tools, many people have automated the construction of images like these [1, 3].

Burrus offers a Python library called DaltonLens [1] for CVD-aware colour management, including a program for generating vanishing plates. His main goal is to provide a simple tool that simulates various forms of dichromacy so that a trichromat can preview whether visual information (say, a user interface or a diagram) is comprehensible to people with CVD [2]. Prior research has shown that deuteranopes perceive a colour space that is well approximated by a planar quadrilateral within LMS space with vertices at black, blue, white, and yellow, hereinafter denoted the “Yellow-Blue quad” or “YB quad”. Burrus’s deuteranopia simulation projects colours along confusion lines onto this quad.

The YB quad is also a useful framework in which to generate vanishing images. For any point on the quad, specified as fractional amounts of blue and yellow between 0 and 1, Burrus constructs a confusion line parallel to the  $M$  axis in LMS space. He then intersects that line with a parallelepiped corresponding to the full RGB colour cube transformed into LMS space. The result is a “confusion segment”, a line segment in LMS space whose points are all deuteranopic metamers, and all convertible to legal RGB colours. The endpoints of a segment  $s$  are colours  $c_1$  and  $c_2$ , each represented as a triple of real numbers in LMS space. We may then express  $s$  parametrically over a real parameter  $t \in [0, 1]$  by defining  $s(t) = (1 - t)c_1 + tc_2$ . In an Ishihara-style vanishing image, we select colours for the foreground object and the background based on  $t$  values in disjoint subintervals of  $[0, 1]$ . Burrus shows how to convert the resulting LMS values first into CIE XYZ space, then into linear RGB, and finally into sRGB, allowing us to create standard raster and vector images.

Figure 2 offers a novel artistic application of vanishing images, based on the observation that a confusion segment’s endpoints usually do not have identical luminances. Given any point in the YB quad, I construct its confusion segment  $s$ , arranging its endpoints so that  $s(1)$  has higher luminance than  $s(0)$ . Now let  $I$  be a greyscale image, such as the portrait of CVD pioneer John Dalton in Figure 2(a). Let  $p \in [0, 255]$  be a pixel of  $I$ . I compute the LMS colour  $s(p/255)$ , convert it into an sRGB colour, and store it in the corresponding pixel of a colour image  $I'$ . The result is a colourful image that is nevertheless able to communicate shading to a trichromat, but that is invisible to a deuteranope.



**Figure 2:** A vanishing image based on a portrait of John Dalton (a). Grey levels are mapped to a confusion segment (b) at point  $(0.4, 0.2)$  on the YB quad, yielding a final image (c). We can simulate the appearance of the result to a dichromat (d) using Burrus’s software (top) and Photoshop (bottom). The simulators produce slightly different outputs because of low-level differences in their colour models.

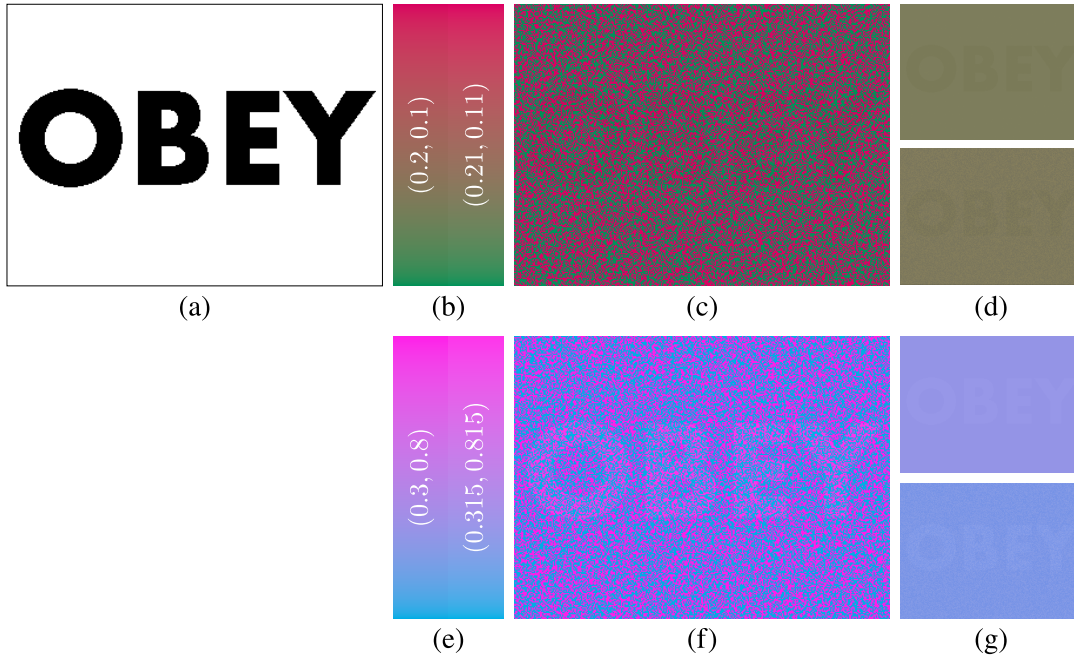
### Hidden Images

While the rendering in Figure 2 is mathematically interesting, and a good starting point for exploration of CVD-aware image processing, I find it to be aesthetically disappointing. After all, CVD already makes many artworks difficult or impossible to interpret; the world does not really need more artifacts that behave that way by design. Far more interesting is to create images with the opposite behaviour: content that is visible *only* to a viewer with CVD.

Ishihara’s “hidden digit” plates, like the one shown in Figure 1(b), demonstrate that images of this type are possible in principle. In a hidden digit plate, a viewer with the correct form of dichromacy will be able to discern a faint digit in the pattern of dots, one that is invisible (or nearly so) to a trichromat. Lakowski carefully measured the trichromatic colours of hidden digit plates and discussed how those colours are organized [6]. As he reports, the trick is to select the dot colours that make up the foreground digit and background from two nearby confusion segments in LMS space. To a trichromat, the colour variation *along* the two segments should dominate their perception, masking the variation *between* the two segments. A deuteranope sees only two colours projected onto the YB quad and can therefore more easily distinguish the digit from the background.

We can adapt the analysis above into an algorithm for generating a hidden image based on a given foreground object. Let an object be represented by a “mask image”, a raster image that is black in the interior of the object and white in the intended background. Choose two nearby points on the YB quad and construct foreground and background confusion segments  $s_f$  and  $s_b$ , as discussed above. Now construct a new output image with the same resolution as the mask. For every black pixel in the mask, set the corresponding pixel of the output image to  $s_f(0)$  or  $s_f(1)$ , chosen at random. Similarly, for every white pixel choose one of  $s_b(0)$  or  $s_b(1)$  at random. The resulting image will look like noise to a trichromat, but the foreground object will be faintly visible to a deuteranope. Figure 3 gives an example of a raster image with a hidden object.

Raster-based hidden images are necessarily noisy: every pixel is assigned one of four colours in a random pattern. By moving instead to vector illustrations, we can generate more structured hidden images. Here the input to the algorithm is a set of closed geometric paths, such as polygons. Each path is given one of two



**Figure 3:** A raster-based hidden image. A mask (a) selects pixel colours from one of two confusion segments (b). The resulting image (c) looks like noise to a trichromat, but a deuteranope can discern a faint copy of the mask object, as visualized using simulators by Burrus (d, top) and Photoshop (d, bottom). The bottom row (e–g) repeats the experiment in a different region of the YB quad. The hidden message is a tribute to the cult sci-fi movie *They Live*.


symbolic labels, in order to distinguish the intended foreground and background regions. As before, we assign colours from  $s_f$  to foreground regions, and colours from  $s_b$  to background regions. Figure 4 demonstrates a hidden image generated this way. The composition is somewhat calmer than that of Figure 3, because it is made up of larger blocks of uniform colour. It more closely resembles a traditional Ishihara test.

Figure 5 shows a similar style of image constructed in a hybrid manner. Geometric paths are provided as before, but the distinction between foreground and background is incorporated via a separate black and white raster mask. For each coloured path shown on the left, we choose whether to use  $s_b$  or  $s_f$  by sampling the mask image at the same location. Black paths are carried over unchanged, allowing universally visible features to be incorporated into the resulting illustration.


### Transformation Images

Ishihara’s “transformation” plates are the most sophisticated of his collection. A transformation plate contains one object that is visible to trichromats, and a second, hidden object visible to individuals with different forms of dichromacy. In the example of Figure 1(c), a trichromat sees the number 74, and a deuteranope sees 21.

The good news is that mathematically, transformation images can be generated by merging the techniques used for vanishing images and hidden images. A vanishing image uses a single confusion segment  $s$ , and chooses colours  $s(t)$  for varying values of  $t$  to modulate the information shown to a trichromat without affecting the image’s appearance to a deuteranope. A hidden image distinguishes foreground from background using two confusion segments, but samples those segments randomly to present noise to a trichromat. A transformation image can therefore be computed from *two* auxiliary images: a greyscale foreground image  $F$  and a black and white hidden image mask  $H$ . As before, we choose two nearby confusion segments  $s_b$  and



**Figure 4:** A vector illustration containing a hidden image. The paths (left) are generated from a collection of overlapping background smiley faces, with a single foreground frowny face indicated in bold. The frowny face is faintly visible to a deuteranope in the final rendering (right).




**Figure 5:** A hybrid hidden image inspired by Lewis Carroll's Cheshire Cat, based on colouring unlabelled geometric paths (left) as foreground or background using a separate mask image (centre). In the final image (right), the black trunk and branches are copied over without recolouring.

$s_f$ . Now, for a given position  $(x, y)$  in the output image, we choose a segment  $s = s_b$  or  $s_f$  depending on whether  $H$  is white or black at  $(x, y)$ , and then set the output pixel to the sRGB value of  $s(p/255)$ , where  $p$  is the grey level of  $F$  at  $(x, y)$ .

Turning this construction around, we see that a vanishing image is a special case of a transformation image in which  $s_f = s_b$ , and a hidden image is a special case in which the pixels of  $F$  are random.

As in the previous sections, we can also use this construction to control colours assigned to geometric paths instead of pixels. Indeed, disjoint paths are a better choice than pixels in this case. If we produce a raster output, then in regions where  $F$  changes slowly, the boundary between black and white in  $H$  will be visible to a trichomat as a faint edge between similar colours. By isolating colours into disjoint shapes, like the dots in an Ishihara plate, the viewer loses the ability to evaluate contrasting colours in direct juxtaposition, helping to obscure the hidden image. Figure 6 gives an example of a transformation plate with an Ishihara-like dot pattern, which uses a pure black and white foreground image to create a two-colour composition. Figure 7 shows a more elaborate colouring in which the foreground image changes continuously.



**Figure 6:** *A transformation image. The foreground image  $F$  (a) controls the colours visible to a trichromat. The hidden mask  $H$  (b) controls the choice of confusion segment. An Ishihara-like arrangement of dots (c) is coloured to produce a design with two opposing messages (d).*

## Discussion


The method presented in this paper is mathematically sound and physiologically well motivated, but it is also unavoidably fragile. After all, any two colours that are distinct to a deuteranope are also distinct to a trichromat. There is no way, then, to achieve perfect steganography, in which the dichromatic content is truly invisible. At best we can hope to camouflage the hidden image behind a dazzle of colour, making it difficult to discern. Even then we must walk a fine line, choosing points on the YB quad whose confusion segments are close enough to be conflated by a trichromat, but different enough to be distinguished by a deuteranope. I chose the colours in my results by trial and error, using simulators to estimate what someone with CVD might see. More systematic experiments could reveal which regions of the YB quad are best suited to these constructions.

Furthermore, these results are always approximations. They are based on imperfect empirical models of human vision, filtered through multiple layers of approximate mathematical transformations, and ultimately displayed on unknown hardware in an unknown lighting environment. It is remarkable that they ever work at all. While it would be ideal to control for all of these variables in a gallery setting, it may be more satisfying to encounter images like these in the wild, where a trichromat might overlook the hint of a secret message implied by subtle colour variations.

I find that these images offer a fascinating glimpse into the psychological grouping principles that guide human vision, much like real-world camouflage. In this context, the tricks of colour and geometry played in hidden and transformation images are devices to guide the trichromat's eye to group the composition in one particular way, obscuring the lower-order grouping possible for a deuteranope.

## Acknowledgements

Thanks to Nicolas Burrus for helpful guidance on the correct use of his DaltonLens library. This work is dedicated to the memory of Linda Carson, and is diminished by the absence of her potential contribution.



**Figure 7:** A transformation image in which the foreground image  $F$  has continuous grey levels (here arranged in a diagonal gradient). The cells of a traditional Islamic geometric pattern are given a range of colours controlled by the gradient, while also concealing a hidden traditional pattern in  $H$ , made up of four copies of the word “ali” in square Kufic calligraphy.

## References

- [1] N. Burrus. “DaltonLens-Python.” <https://github.com/DaltonLens/DaltonLens-Python>. Accessed April 19th, 2023.
- [2] N. Burrus. “Understanding LMS-based Color Blindness Simulations.” <https://daltonlens.org/understanding-cvd-simulation>. Accessed February 8th, 2023.
- [3] F. Couzo. “Ishihara Plate Generator.” <https://franciscouzo.github.io/ishihara/>. Accessed April 19th, 2023.
- [4] D. Flück. “Ishihara’s Test for Colour Deficiency: 38 Plates Edition.” <https://www.color-blindness.com/ishiharas-test-for-colour-deficiency-38-plates-edition/>. Accessed April 19th, 2023.
- [5] J. F. Hughes, A. van Dam, M. McGuire, D. F. Sklar, J. D. Foley, S. K. Feiner, and K. Akeley. *Computer graphics: principles and practice (3rd ed.)*. Boston, MA, USA: Addison-Wesley Professional, July 2013.
- [6] R. Lakowski. “Theory and practice of colour vision testing: A review. Part 2.” *Brit. J. industr. Med.*, vol. 26, no. 4, October 1969, pp. 265–288.
- [7] C. Man. “You See, I Hear You.” <https://carolxman.com/2020/08/22/exhibition-you-see-i-hear-you/>. Accessed April 19th, 2023.
- [8] M. Morgan, A. Adam, and J. Mollon. “Dichromats detect colour-camouflaged objects that are not detected by trichromats.” *Proc. Biol. Sci.*, vol. 248, no. 1323, June 1992, pp. 291–295.