Colour dithering using a space filling curve

John W. Buchanan, Oleg Verevka
Department of Computing Science
Technical Report TR95-04
University of Alberta
Edmonton, Alberta.
{juancho,oleg}@cs.ualberta.ca

May 1, 1995

Abstract

Recent results describe how a space-filling curve can be used as an error-diffusion path for half-toning of grey-scale images. The space-filling half-toning method produces clustered approximations to images with less regular artifacts than other available clustering methods. In this paper we present two methods for half-toning colour images using the space-filling technique. The first method processes the red, green, and blue (RGB) (or cyan, magenta, yellow, and black (CMYK)) channels independently, while the second method implements an interpolation between colours that lie on the corners of the RGB colour cube. The first method produces overlapping clusters of red, green, and blue pixels. This technique is useful when the display device allows the three colours to be set independently. For a variety of reasons it is often desirable to reduce the amount of ink used in the display of an image. The second method we present addresses this problem by approximating the regions of pixels with two sets of non overlapping colours. This results in image approximations with a guaranteed maximum pixel coverage of one or two inks (100% or 200% coverage). The quality of image reproduction is severely compromized when this approach is used. Nevertheless, the technique does provide a guarantee of maximum coverage.

These results allow the use of the space-filling technique for dithering colour images. We illustrate the performance of these methods using two images that are printed in various ways on a 300 DPI ink-jet printer. The results are primarily geared to printers, since most modern computer displays allow the selection of specialized palettes, thus making quantization results applicable.

1 Introduction

1.1 Note from the authors

This paper was submitted to Siggraph '95. At the time of submition we were unsure of the significance of this work. The reviewers felt that the work presented in this paper was not worthy of publication in Siggraph. After reading the reviews of this paper, which were by the way very well presented, we felt it best to not pursue this area actively in the immediate future. This technical report details the work that we undertook in the course of this research. One important thing that we learned from this experience is that the Mandrill is a particularly bad image to use as a test image for colour dithering. The reason for this is that it is very forgiving of bad display methods.

1.2 Back to the introduction

It is well known that grey scale images can be approximated by a set of black and white pixels that are appropriately distributed[Ulic87]. When these sets of black and white pixels are viewed from a distance our visual system interprets the black and white pixel map as a grey-scale image. In general this technique also works for the display of colour images. By approximating the image with a series of coloured dots we can display images that appear to be full coloured.

The typical colour printing process starts with colour separation. Colour separation transforms the representation of image colours from screen colour coordinates (RGB) into printer specific coordinates (CMYK). In general, this process is rather complex. In order to ensure a correct colour reproduction it has to incorporate gray balancing and adjustment of contrast, saturation and hue. A detailed description of the process can be found in [Yule67, Ston88, Lind89, Lamm90].

Once an appropriate separation is obtained, all channels are processed independently. This approach is justified since in most colour printers each colour channel is separately added to the final print. Monochrome half-toning techniques are thus directly useful for dithering the separate colour channels. Colour dithering methods must also consider the possibility that there will be some mis-alignment in the position of the inks. Discussions of this problem and proposed solutions can be found in [Yule67, Holl80, Ston88].

2 Overview

There is a large amount of research [Jarv76, Knut87, Ulic87, Ulic88, Geis90, Velh91, Wyvi91, Velh92, Gots93, Zhan93b, Ostr94, Buch95] dealing with the half-toned display of grey-scale images In particular, recent results have detailed how an error dispersion algorithm can be implemented using a space filling curve[Witt82, Velh91, Velh92, Buch95]. The chief advantage of this method is that the colours of the image are approximated by clustered sets of black and white pixels. These clustered sets of pixels introduce less regular artifacts than other clustered dithering techniques [Limb69, Knut87, Velh91, Zhan93a]. On many devices (printers in particular) it is difficult to independently set pixels[1]. In such cases a clustered half-toning method is preferred.

The space-filling method uses sub-sections of a space-filling curve to half-tone the image. Each of these sub-sections of the curve defines a region of the image. The local average of the region determines the number of black and white pixels that must be set to approximate the average. Because the space-filling [2] curve is locally compact the resulting sets of black and white pixels are fairly well clustered. Any quantization error induced by this choice of black and white pixels is propagated to the next curve section. In previous work [Buch95] we introduced two techniques for enhancing the performance of the space-filling technique, Velho and Gomez reported results similar to our first technique in [Velh92].

Having reported these results we were faced with the question: "Can the space-filling curve dithering be used to display coloured images?" and, "if so, how?".

We believe that the answer to the first question is yes. The justification for this affirmative answer is given in the form of two answers to the second question. The two methods that we present process the image along the space-filling curve. Intervals of this curve define regions of the image. For a given region size the position and shape of the regions is independent of the image contents or the dithering method chosen. Thus the two methods that we present differ only in the way they process the pixels of a region. Note also that if we are using edge detection to adaptively size the region size [Buch95], then again, the size and shape of the regions are independent of the dithering method.

The first method processes the three colour channels independently resulting in overlap-

[1] Ink bleeding is the usual cause of this inability to set pixels in isolation.
[2] We use a Hilbert curve

ping red, green, and blue (or cyan, magenta, yellow and black[3]) clusters. The combination of these clusters approximates the average colour of the region.

The second method we present approximates the average colour $\overline{c}$ of a region by finding two colours $\{C_1, C_2\}$ in the corners of the RGB (CMY) cube. These colours are the vertices of the closest line to $\overline{c}$ Having chosen the two colours that will approximate the region we must set the pixels within the region to either C_1 or C_2. We have investigated two different ways of setting the pixel colours in the region. The first way to select the colors is based on the original space-filling work. The average $\overline{c}$ of the region is projected to the line between the two points C_1 and C_2. The position of the projected average determines how many pixels are set to C_1 and how many pixels are set to C_2. The second way to choose the colours of the pixels in the region is similar to colour quantization. The colours of the pixels in the region are set to the closest of the two colours C_1 or C_2. Both of these pixel colour selection schemes results in a colour approximation where at most two inks are set per pixel. Thus the maximum coverage of any pixel is 200%. If C_1 and C_2 are restricted to single ink colours then the maximum coverage is 100%.

Test images

We use the mandrill image to illustrate the performance of the various techniques. The original image of the mandrill is presented in figure 3. In the final image (figure 12) we use another test picture in order to facilitate the comparison of the different techniques.

3 Space-filling curve half-toning

Witten and Neal [Witt82] proposed a dispersed error dot diffusion method for dithering grey-scale images. They used a space-filling curve along which the error was propagated. They showed that the use of the space-filling curve allowed one-dimensional error propagation to be used since the shape of the curve removed or hid some of the artifacts that are associated with one-dimensional error propagation. Velho and Gomez [Velh91] showed that when the curve was processed in regions of pixels (adjacent intervals along the curve) the resulting black and white pixels were clustered. This provided a clustered half-toning

[3]For the sake of clarity we will use red, green, and blue for our discussions. All of the results can be directly applied to cyan, magenta, yellow, and black images.

Figure 1: Sheepherder by Peter Hurd. 128×128 printed at 75 DPI with a maximum region size of 7. Image on the left is the space-filling technique. The centre image uses adaptive region size with a edge threshold of 10. The image on the right uses fixed region sizes, but sorts the cluster prior to setting the colours of the pixels.

method whose artifacts were less regular than the artifacts introduced by other clustered half-toning methods [Limb69, Knut87, Zhan93a, Velh91].

The method that Velho and Gomez presented could produce clusters that straddled edges in the original image. This had the effect of destroying the edges in half-toned image. In a previous paper [Buch95] we presented two methods that ensured that the edges of an image were accurately displayed in the dithered image. The first method uses an edge detection filter to change the size of the regions. By ensuring that no region straddles an edge we can guarantee that no cluster within the regions will cross an edge. Velho and Gomez [Velh92] reported similar results achieved by using a gradient measure along the path.

Edges in monochrome images occur when there is a difference between a local bright area and a local dark area. By setting the pixels in the brightest area of the region to white we preserve the local distribution of light and dark pixels. This has the effect of highlighting edges and other fine details of the images. One way of selecting the bright pixels within a region is to sort the pixels. This is the second method of edge enhancement that we presented in [Buch95]. We present a reduced example from this paper in figure 1. This is a 128×128 image printed at 75 DPI. Notice how clear the edges are in the sorted image.

3.1 Error propagation at edges

When an edge is encountered one has to make a decision about the propagation of the quantization error. If the error is propagated to the next cluster then the edge may not be as clear in the dithered image. We have found that the edges in the image can be enhanced if the error term is set to zero when an edge is encountered. Unfortunately this has the effect of introducing a minor local hue shift. In our implementation, we have allowed the user to choose whether or not the error is to be propagated past the edges in the image.

3.2 Cost of sorting the pixels

The cost of copying and sorting the pixels in the region can be quite high if the region size is large. However, we have found that the optimum region size is between 3 and 11. In fact there is no need to completely sort the pixels in the region, it is sufficient to find the n brightest pixels. The sorting routine we use requires a parameter that indicates how many bright pixels are required.

3.3 Multi channel edge detection

In this paper we present two methods for using the space-filling technique for printing coloured images. In our investigation we found that for both methods benefitted from the use of adaptive region sizes. In our implementation the user selects an edge threshold ϵ. An edge is defined when a difference of more than ϵ is detected along the path in any of the channels. As reported in [Buch95] edge thresholds between[4] 20 and 40 seem to be the most useful.

4 Full-colour approximation

The first implementation of space-filling colour dithering that we describe is in some sense the most straight-forward. Each of the channels are processed independently, thus requiring that the colour separation be done as a pre-processing step.

Given a region of the curve the average intensities for the different channels are computed. Using these averages, the channels are half-toned and the resulting colours are

[4] We are assuming standard 8 bit pixel values.

overlayed. In figure 4 we show the result of this colour dithering method. This image was generated using a cluster size of 7. As expected this exhibits the main problem associated with the space-filling technique, namely that clusters straddle the edges in the scene. As we discussed previously, the display of images using the space-filling technique can be enhanced by using edge detection to adaptively set the size of the regions. Using an edge threshold of 20 results in a much clearer display of the mandrill (see figure 5). Sorting the pixels prior to setting their intensity preserves the local distribution of light and dark areas. In figure 6 we present the mandrill image generated using sorting with a step size of 7. In figure 7 we present the results when edge detection and sorting are combined.

As expected, this implementation of the space-filling technique mirrors the results reported in [Velh91, Velh92, Buch95]. Clustered sets of pixels are used to approximate the image. The size of these clusters is controlled by a user defined parameter. The behavior of the algorithm at the edges is also determined by the user. By coupling this method with an appropriate colour separation method the images can be tuned to a chosen printer.

5 Two-colour approximation

Traditional colour reproduction is based on colour separation of the original image into three or four colour plates. Each colour plate is dithered and printed separately. For economical and mechanical reasons it is desirable to control the amount of ink in the final print. Moreover, overlay of more than two inks in one area on a transparent material leads to over-exposed films or transparencies that are too dark.

Stone et. al. [Ston88] reported that the pixel coverage in images can be reduced to 280%, where the overlay of four colours would equal 400%. In this section we present a space-filling colour half-toning technique that guarantees a maximum pixel coverage of 200% or 100%. We refer to this process as the two-colour approximation space-filling method.

The key idea behind this method is to approximate the average colour of a region by using two sets of pixels that are individually set to one of two chosen primary colours C_1 and C_2 . Consider a colour printer with 4 inks (CMYK). If we restrict the maximum dot coverage to 100%, the set of primary colours is exactly the set of inks — cyan, magenta, yellow and black — plus the white colour of the paper. In the case of the maximum coverage of 200%, the additional red, green, and blue primaries are allowed (these are all combinations of two of the CMY colours). The difference between the 100% coverage and the 200% coverage methods lies in the choice of legal colours. Once we have selected the primary colours that are to be used, the techniques are identical.

The selected set of primary colours lies on the vertices of the RGB colour cube. The lines between these primary colours are the set of colours that can be approximated by using pairs of primary colours. Some of these lines intersect the line from white to black. These lines are excluded so that the approximation of grey portions of the image are computed using black and white pixels. The resulting lines are the basis of our approximation technique.

Given the average colour of a region $\overline{c}$ we find the closest primary colour C_1. The second colour C_2 is chosen by examining the lines that have C_1 as one of their vertices. By choosing the closest line to the average colour $\overline{c}$ we define the two primary colours that will be used. The two-colour dithering process will represent the $\overline{c}$ by an appropriate number of C_1 and C_2 coloured pixels in the region.

Given two colours C_1 and C_2 there are two methods for setting the colours of the pixels c_i in the region.

5.1 Projection

The first colour allocation method is based on the space-filling gray scale dithering. Given two approximating colours C_1 and C_2, a colour to be approximated $\overline{c}$, and a region of n pixels: we must determine how many pixels are to be set to C_1 and how many are to be set to C_2. Projecting $\overline{c}$ onto the line $\overline{C_1C_2}$ gives us this information. The number of pixels set to the first primary C_1 is defined by:

$$n_1 = \left\lfloor n\frac{\|c_* - C_1\|}{\|C_1 - C_2\|} + 0.5\right\rfloor,$$

where c_* is the projection of $\overline{c}$ onto the line $\overline{C_1C_2}$. The other $n - n_1$ pixels are set to C_2.

The error propagation and sorting algorithms described in section 3 are directly applicable here. If the c_i colour of the pixel in the region is substituted by $C_{j_i}, j_i \in \{1, 2\}$ then the local quantization error is $\epsilon_i = c_i - C_{j_i}$ and the global error for this region is:

$$\epsilon = \sum_{i=1}^{n}(c_i - C_{j_i})$$

This quantization error is propagated to the next region of the curve.

Sorting can be used to minimize the global error ϵ and to highlight edges inside the region. Pixels of the region can be sorted according to their projections on the line $\overline{C_1C_2}$. The n_1 pixels that are closest to C_1 are set to C_1. An example of this two-colour half-toning method is presentedin figure 8. A 100% coverage of the same image is presented in figure 9.

5.2 Quantization

Another method for selecting the colours of the region pixels is to set the pixels c_i to the closest primary C_1 or C_2. This technique resembles colour quantization without error propagation within one region. The global approximation error for the given region has to be transfered to the next region to ensure the correct colour reproduction of the entire image. In figure 10 we present the mandrill using this colour allocation method, the 100% version is presented in figure 11. We have found that the colours are more accurately displayed using this method even though large clustering artifacts are evident in uniform gradient regions.

6 Examples

All of the images presented here were produced on a HP 300DPI Deskjet 1200C/PS ink jet printer. The half-toning of the image was done on a SGI workstation so that no half-toning is performed by the postscript engine.

In figure 12 we present a set of twelve images printed at 300DPI. These are presented so that the performance of our technique can be evaluated in the context of other dithering techniques. Table 1 describes the parameters used for the display of each of the sample images.

Examining the two-colour approximations we notice some step-like effects. These are most evident in the regions of uniform gradient. In figure 2 we illustrate the problem. Two neighboring regions R_a and R_b have similar averages $\{\overline{c}_a, \overline{c}_b\}$ but are approximated by different colours c_{a*} and c_{b*}. This is because the average colours project to different lines on the colour cube, even though they are close to each other. The first colour selection

Floyd-Steinberg	Dispersed Dot ordered (4×4)	Clustered dot ordered (4×4)
Full colour Region = 7 No edge No sorting	Full colour Region = 7 Edge = 20 No sorting	Full colour Region = 7 Edge = 20 Sorting
2 Colour Region = 7 No edge Sorting 200% coverage	2 Colour Region = 7 Edge = 20 No sorting 200% coverage	2 Colour Region = 7 Edge = 20 Sorting 200% coverage
2 Colour (Quantization) Region = 7 Edge 20 200% coverage	2 Colour (Quantization) Region = 3 Edge = 20 200% coverage	2 Colour Region = 7 Edge = 20 Sorting 100% coverage

Table 1: Parameters for images in figure 10

method then tries to approximate the projected average colour of the region. This results in two different colours being approximated in neighboring regions.

This quantization effect is not so apparent when we use the nearest colour selection. This is because a small shift in the average colour of two adjacent regions typically only changes one of the approximating colours, as in this example. Thus a significant number of pixels in neighboring regions are set to the same colour C_1. This is of course not true when large colour shifts are encountered.

7 Conclusion

We have presented two methods for printing coloured images using the space-filling curve dithering method. The first method provides a full colour dithering method. The advantages of this method are that clustered sets of pixels are used to approximate the image. These clusters do not exhibit the same regular artifacts as those produced by other clustering dithering techniques. The second method provides a tool for approximating images with a guaranteed maximum pixel coverage of 100% or 200%. This two-colour dithering method relies on two different colour allocation methods. The first uses a projection of

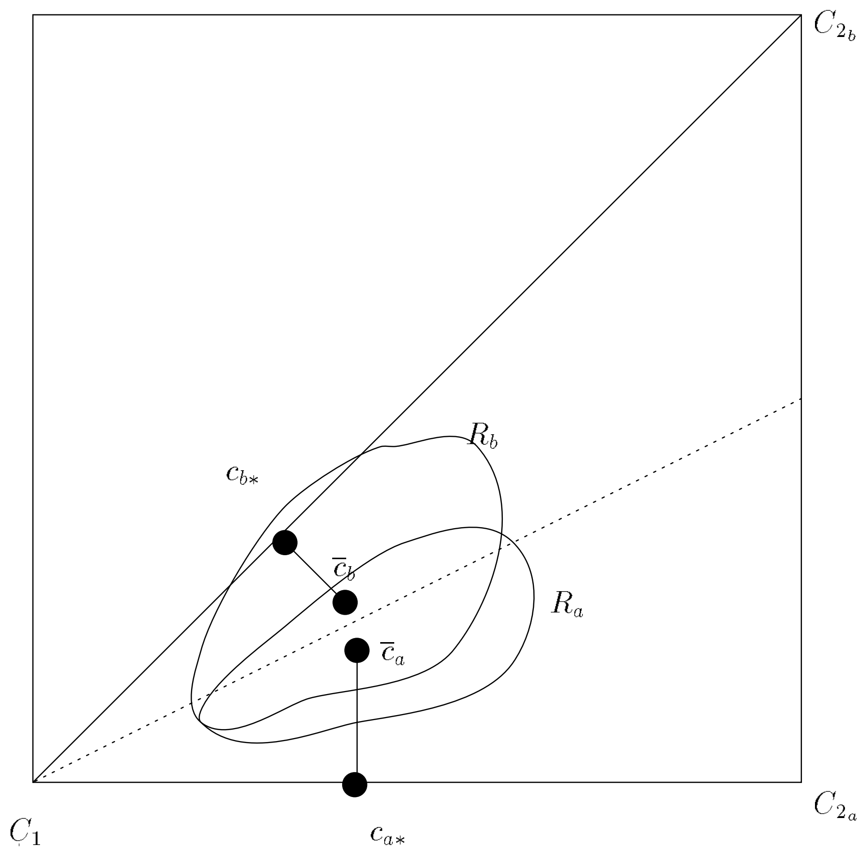

Figure 2: Two regions $\{R_a, R_b\}$, their averages $\{\overline{c}_a, \overline{c}_b\}$, and their projected averages $\{c_{a*}, c_{b*}\}$.

the average colour to determine the proportion of the two approximants to use, the second simply sets the pixels to the closest approximant. Reasonable results can be produced with both of these colour allocation schemes. The projection method seems to produce a better local approximation to the colour but can introduce step-like artifacts in uniform gradient areas. The second colour allocation method produces granier images but does not exhibit the step-like artifacts introduced by the projection method.

The computational requirements of the two techniques are modest. The full colour dithering method only becomes expensive if sorting is used to enhance the display of the edges. The two-colour approximation has a higher overhead associated with the search for the two approximation colours.

If we choose a region size of 1 then both the full colour dithering and the 2 colour dithering produce reasonable results. This amounts to an approximation of each pixel with a one-dimensional error propagation along the space-filling path. If we choose not to propagate the error past the edges then we must have clusters large enough to accommodate the loss of information. In our experience error should be propagated past the edges when the region size is smaller than five. When the regions are larger than five the artifacts introduced by dropping the error term at the edge are not that apparent. We must note that in terms of image correctness, it is always wrong to drop the error term at an edge. However, esthetically we have found that the images produced by dropping the error term at the edges are often better.

The images that we have printed using these techniques lead us to believe that space-filling can be effectively used for the printing of high-quality coloured images. It is clear that printing high-quality images is a challenge. Printing these high quality images on a given printer requires the tailoring the colour-separation process and the dithering process to the printer. We believe that the techniques presented here provide a number of degrees of freedom to make this tailoring easier.

The latter two techniques (100% and 200%) coverage are clearly not good techniques. They are presented here in the hope that:

1. No one makes the same mistake again.

2. They challenge others to think about the issue of displaying images with reduced ink coverage.

References

[Buch95] J. W. Buchanan and O. Verevka. "Edge preservation with space-filling curve half-toning". *To be appear in Graphics Interface*, May 1995.

[Geis90] Robert Geist and Robert Reynolds. "Colored Noise Inversion in Digital Halftoning". *Proceedings of Graphics Interface '90*, pp. 31–38, May 1990.

[Gots93] Craig Gotsman. "Halftoning of image sequences". *The Visual Computer*, Vol. 9, No. 5, pp. 255–266, March 1993.

[Holl80] T.M. Holladay. "An optimum algorithm for halftone generation for displays and hard copies". *Proc. SID*, Vol. 21, pp. 185–192, 1980.

[Jarv76] J.F. Jarvis and C.S. Roberts. "A new technique for displaying continuous tone pictures on bilevel displays". *IEEE Transactions on Communications*, Vol. COM-24, pp. 891–898, 1976.

[Knut87] D. E. Knuth. "Digital halftones by dot diffusion". *ACM Transactions on Graphics*, Vol. 6, No. 4, pp. 245–273, October 1987.

[Lamm90] Michael G. Lamming and Warren L. Rhodes. "A Simple Method for Improved Color Printing of Monitor Images". *ACM Transactions on Graphics*, Vol. 9, No. 4, pp. 345–375, October 1990.

[Limb69] J. O. Limb. "Design of Dither Waveforms for Quantified Visual Signals". *Bell Systems Technical Journal*, Vol. 48, No. 7, pp. 2555–2582, 1969.

[Lind89] Bruce J. Lindbloom. "Accurate Color Reproduction for Computer Graphics Applications". *Computer Graphics (SIGGRAPH '89 Proceedings)*, Vol. 23, No. 3, pp. 117–126, July 1989.

[Ostr94] V. Ostromoukhov, R. D. Hersch, and I. Amidror. "Rotated Dispersed Dither: A new Technique for Digital Halftoning". *Computer Graphics (SIGGRAPH '94 Proceedings)*, pp. 123–130, July 1994.

[Ston88] Maureen C. Stone, William B. Cowan, and John C. Beatty. "Color gamut mapping and the printing of digital color images". *ACM Transactions on Graphics*, Vol. 7, No. 4, pp. 249–292, October 1988.

[Ulic87] R. Ulichney. *Ditital Halftoning.* MIT Press, Cambridge, Massachusetts, 1987.

[Ulic88] R. A. Ulichney. "Dithering with blue noise". *Proc. IEEE*, Vol. 76, No. 1, January 1988.

[Velh91] Luiz Velho and Jonas de Miranda Gomes. "Digital halftoning with space filling curves". *Computer Graphics (SIGGRAPH '91 Proceedings)*, Vol. 25, No. 4, pp. 81–90, July 1991.

[Velh92] L. Velho and J. M. Gomes. "Space Filling Curve Dithering with Adaptive Clustering". *Proceedings of SIBGRAPI'92 – V Simposium of Computer Graphics and Image Processing*, pp. 1–9, 1992.

[Witt82] I.H. Witten and R.M. Neal. "Using Peano curves for bilevel display of continuous-tone images". *IEEE Computer Graphics and Applications*, Vol. 2, No. 3, pp. 47–52, May 1982.

[Wyvi91] Geoff Wyvill and Craig McNaughton. "Three plus five makes eight: A simplified approach to haltoning". *Scientific Visualization of Physical Phenomena (Proceedings of CG International '91)*, pp. 379–392, 1991.

[Yule67] J. A. Yule. *Principles of Color Reproduction.* Wiley, 1967.

[Zhan93a] H. Zhang and B. Wyvill. "The taxonomy of behavioural simulation in computer graphics — A perspective from the calculus of communicating systems". *Proceedings of the 1993 Western Computer Graphics Symposium*, pp. 1–2, March 1993.

[Zhan93b] Y. Zhang and R. E. Webber. "Space diffusion: An improved parallel halftoning technique using space-filling curves". *Computer Graphics (SIGGRAPH '93 Proceedings)*, pp. 305–312, August 1993.

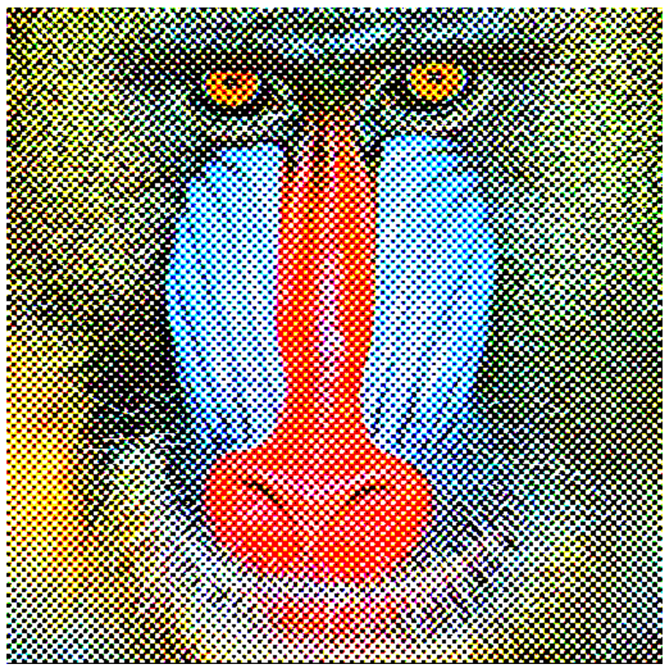

Figure 3: Mandrill used as an example. 512×512 image printed at 150 DPI using clustered dot ordered dithering of order 4.

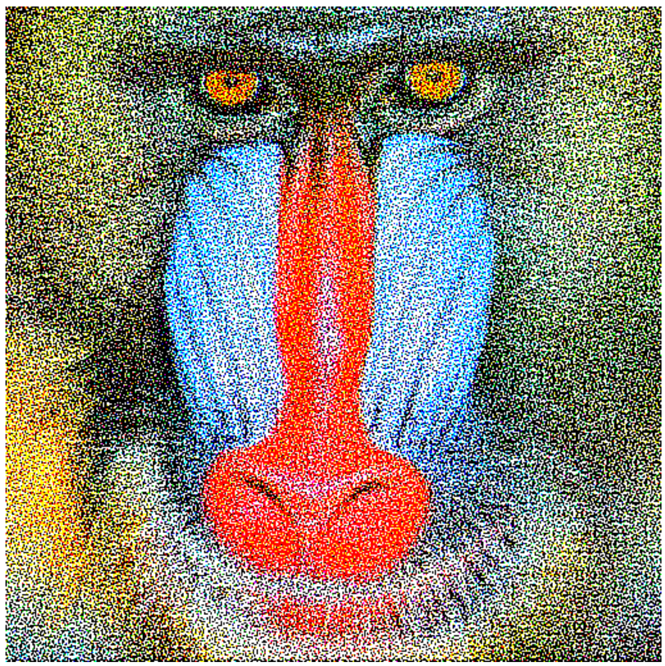

Figure 4: Mandrill printed using space-filling dithering applied to each channel independently. Step size = 7, DPI = 150.

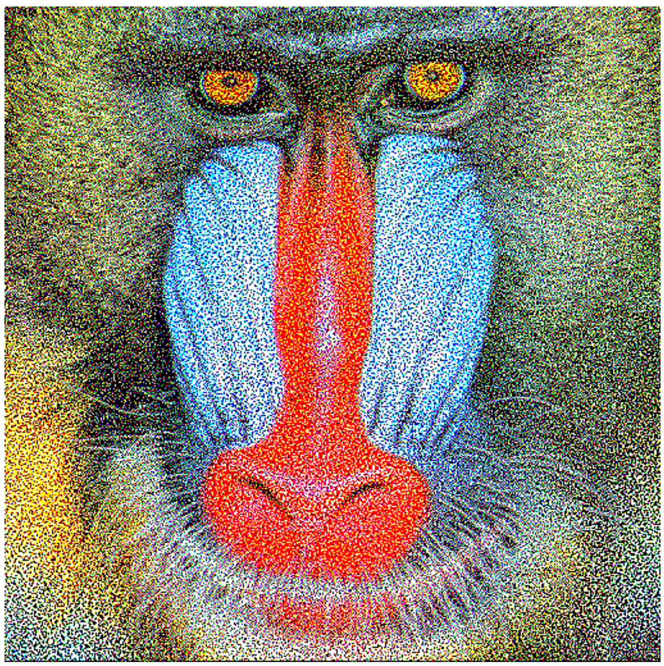

Figure 5: Mandrill printed using space-filling dithering applied to each channel independently. Step size = 7, Edge tolerance = 20, DPI = 150

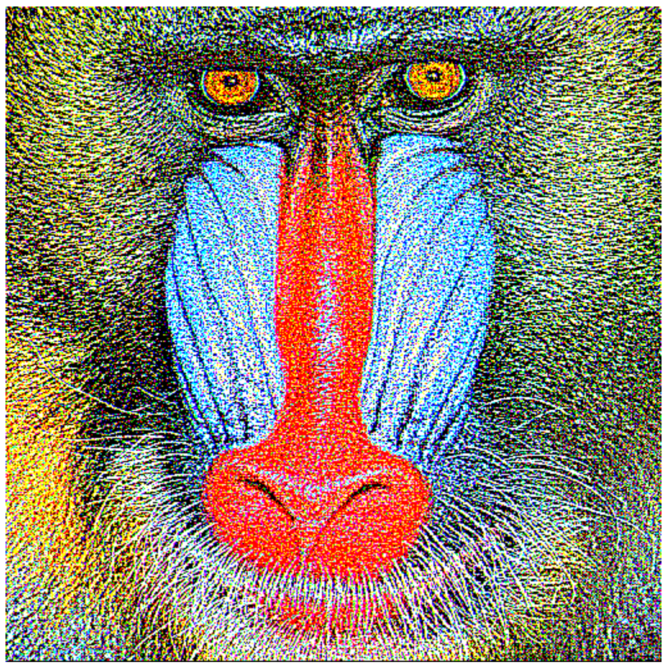

Figure 6: Mandrill printed using sorted space-filling dithering applied to each channel independently. Step size = 7, DPI = 150.

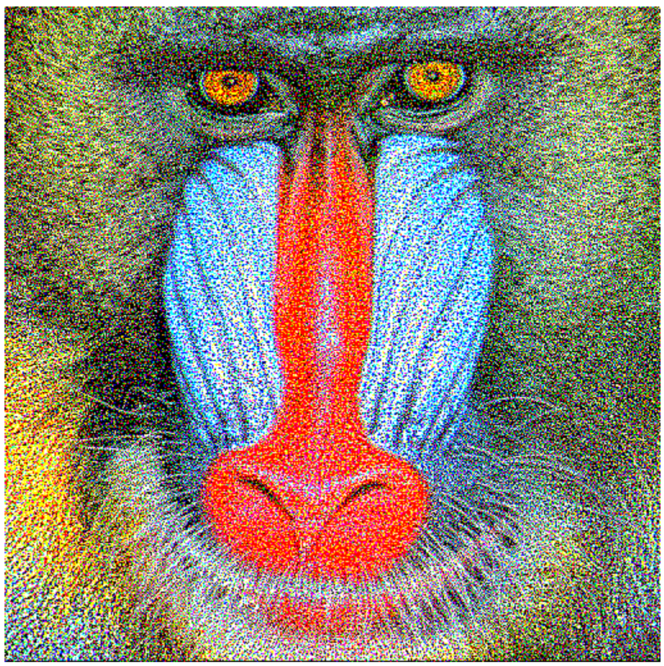

Figure 7: Mandrill printed using sorted space-filling dithering applied to each channel independently. Step size = 7, Edge tolerance = 20, DPI = 150

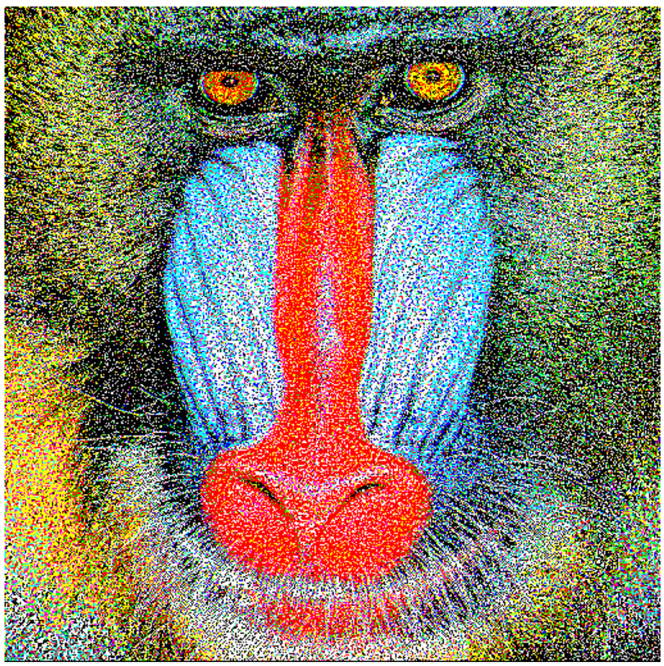

Figure 8: Two colour approximation with 200% coverage. 150 DPI, size = 7, edge threshold = 20, sorted. Colours selected by projecting average onto the closest line.

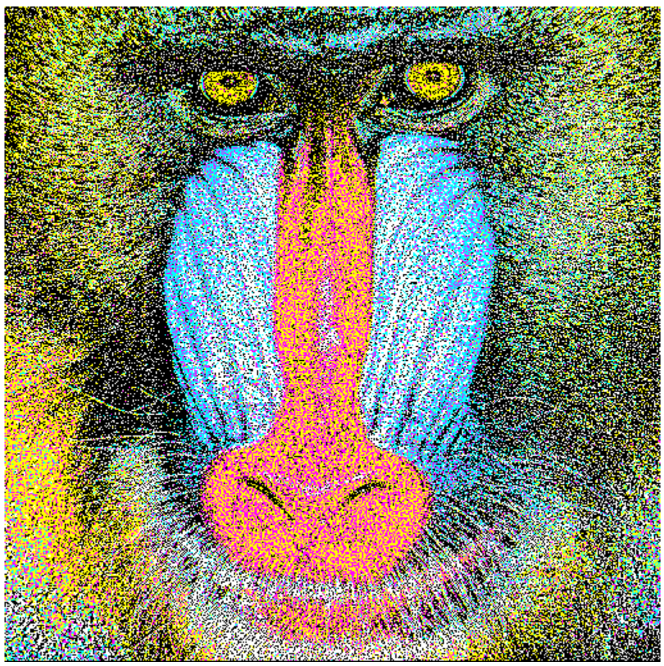

Figure 9: Two colour approximation with 100% coverage. 150 DPI, cluster size = 7, edge threshold = 20, sorted. Colours selected by projecting average onto the closest line.

Figure 10: Two colour approximation with 200% coverage. 150 DPI, region size = 7, edge threshold = 20. Colours selected by the nearest colour.

Figure 11: Two colour approximation with 100% coverage. 150 DPI, region size = 7, edge threshold = 20. Colours selected by the nearest colour.

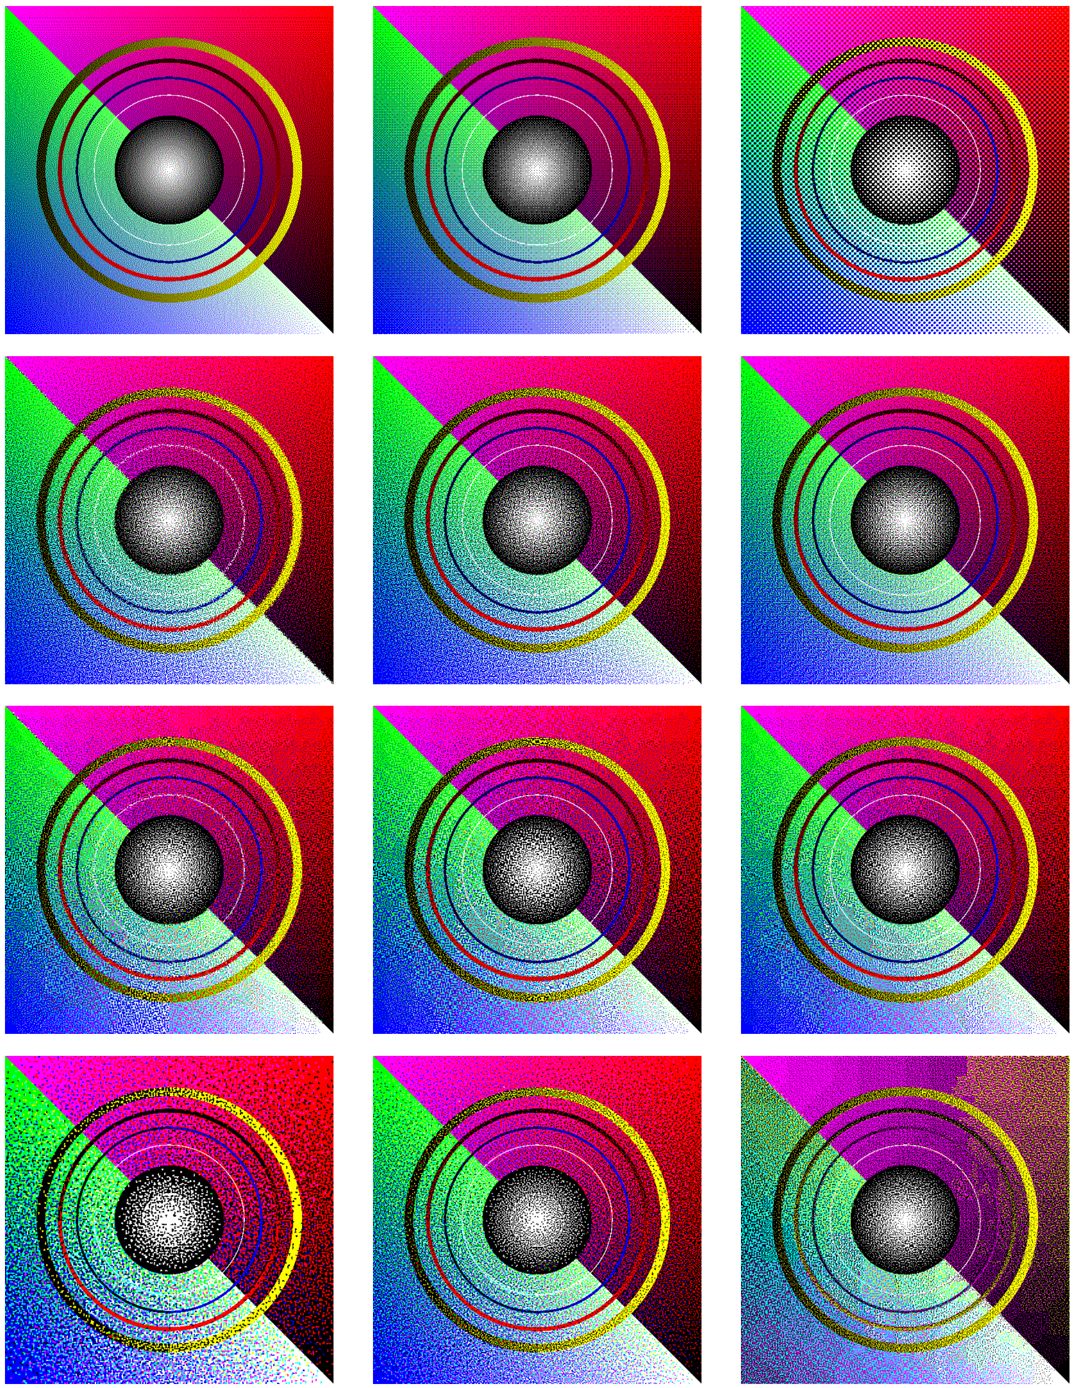

Figure 12: Different versions of test image at printed at 300 DPI. (See Table 1 for the parameters used in this figure)