Multi-Color and Artistic Dithering

Victor Ostromoukhov, Roger D. Hersch
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
{victor, hersch}@di.epfl.ch

Abstract

A multi-color dithering algorithm is proposed, which converts a barycentric combination of color intensities into a multi-color non-overlapping surface coverage. Multi-color dithering is a generalization of standard bi-level dithering. Combined with tetrahedral color separation, multi-color dithering makes it possible to print images made of a set of non-standard inks. In contrast to most previous color halftoning methods, multi-color dithering ensures by construction that the different selected basic colors are printed side by side. Multi-color dithering is applied to generate color images whose screen dots are made of artistic shapes (letters, symbols, ornaments, etc.). Two dither matrix postprocessing techniques are developed, one for enhancing the visibility of screen motives and one for the local equilibration of large dither matrices. The dither matrix equilibration process corrects disturbing local intensity variations by taking dot gain and the human visual system transfer function into account. Thanks to the combination of the presented techniques, high quality images can be produced, which incorporate at the micro level the desired artistic screens and at the macro level the full color image. Applications include designs for advertisements and posters as well as security printing. Multi-color dithering also offers new perspectives for printing with special inks, such as fluorescent and metallic inks.

Keywords: color halftoning, artistic dithering, dither matrix equilibration, non-standard inks, side by side printing

1. Introduction

The reproduction of color images requires in the general case (1) separating the image colors (for example red, green and blue) into the set of available printable colors (for example cyan, magenta, yellow and black), (2) halftoning each of the printable color layer and (3) possibly calibrating the system to ensure that the printed colors are close to the original image colors.

Most existing color halftoning techniques are based on the independent halftoning of each of the contributing color layers. In offset and in many electrographic printers, the color layers are generated independently at angles of 15, 45, 75 degrees in order to avoid interferences between the cyan, magenta and black layers ([24], chapter 13). Ink-jet printers often use error-diffusion to halftone each of the color layers independently.

Existing approaches for color separation with non-standard inks[1], for example the segmentation of the CIE-XYZ color space into a set of tetrahedra having as a common edge the black-white axis [14], the creation of correspondence tables between input colors and combination of output inks [1], or the selection of an optimal subset of inks according to an objective function [19] consider that color layers can be printed independently one from another. This is true as long as the inks are transparent and as long as the dots of each screen are randomly positioned with respect to the dots of the other screens, as assumed by the Neugebauer equations [17].

If these assumptions are not true, halftoning each of the color layers independently may generate color shifts depending on the amount of superposition between screen dots of individual halftone layers. As long as the inks are transparent, the color shifts are small and the reproduction fidelity can be ensured by calibrating the printing device [6], for example, in the case of process colors, by establishing a 3D mapping between CIE-XYZ coordinates and the output Cyan, Magenta, Yellow and Black surface coverage values [5].

Some applications however require that different inks are always printed side by side without overlapping. This is for example the case when printing with opaque inks [9], [10]. Manufacturers of valuable documents, such as banknotes, identity cards and checks often make use of the high registration accuracy of their original printing equipment to create a graphic design which is both visually pleasant and difficult to imitate using standard printing processes ([22], chapter 6).

We propose a multi-color dithering method which automatically enforces *side by side printing* of several color layers. Combined with color separation by tetrahedral interpolation, multi-color dithering makes it possible to print images made of non-standard inks. In addition, we explore ways of creating large dither matrices to produce artistically screened color images, i.e. color images whose screen elements are made of artistic color shapes (letters, symbols, ornaments, etc.).

In section 2, we present the new multi-color dithering method, which is a generalization of standard black and white dithering. We show how multi-color dithering can be used for the color separation of non-standard inks. In section 3, we describe the synthesis of dither matrices improving the visibility of screen motives in areas of high ink coverage. In section 4, we describe an equilibration process which improves the quality of large dither matrices by compensating for local uneven motive distribution and dot gain. In section 5, we apply the previously described dither matrix synthesis methods to generate screens representing freely chosen artistic shapes. We show the results in the form of color pictures which comprise at the micro level the desired artistic screens and at the macro level the global color image. The conclusions and perspectives are presented in section 6.

2. Multi-color dithering and color separation.

Error-diffusion in color space [20] may be used for side by side printing with non-standard colors, i.e. for ensuring that different inks are always printed side by side. Error-diffusion in color space is a straightforward extension of standard error-diffusion [21]: given an input color vector (for example in RGB space) the closest printable color is selected and printed on the current output pixel. The 3D difference vector (error) between input and selected color is distributed among the current pixel's neighbors [7]. A disadvantage of color error diffusion are the classical error-diffusion artifacts and the fixed dot size which in case of dot gain may induce a strongly non-linear tone reproduction behavior [8].

The new multi-color dithering method we introduce here has the same advantages as standard dithering. It enables the creation of dither matrices of the desired size and orientation and the genera-

1. Inks which differ from the standard cyan, magenta, yellow and black process colors

tion of dither threshold levels according to the desired dot growth behavior. Furthermore, small individual dither matrices can be assembled to form large super-matrices incorporating many intensity levels ([6], chapter 9).

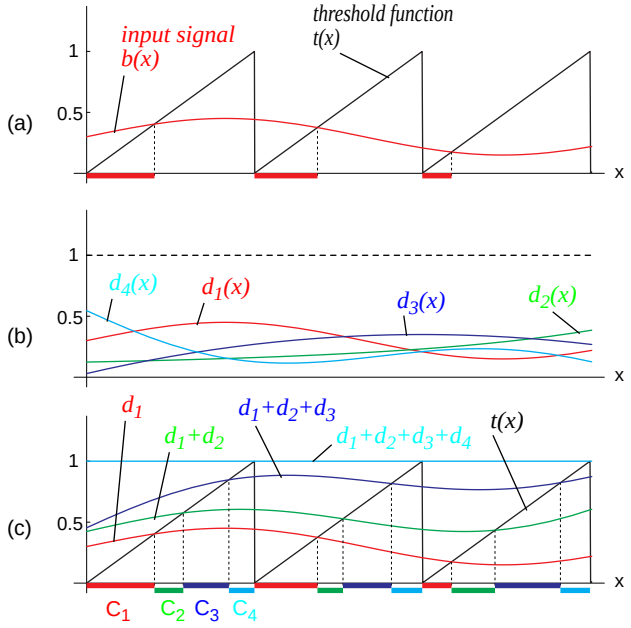

Fig. 1 (a) Black-white dithering of a variable darkness input signal, (b) relative color intensities, (c) multi-color dithering.

2.1 Standard black-white dithering

Before explaining multi-color dithering, let us describe the basics of standard dithering for black-white (two colors). To simplify the explanations we assume that an input grayscale image with normalized darkness values between *0* (white) and *1* (black) is dithered by comparing at each output location corresponding input darkness and dither threshold values[1]. If the darkness *b(x)* is higher than the dither threshold value *t(x)*, then the output location is marked as black, else it is marked as white (Fig. 1a).

Standard dithering converts a darkness value into a surface coverage. Conceptually, we can look at a given darkness value *b* as a percentage *b* of black and as percentage of *(1-b)* of white. The dithering process converts an input signal of darkness *b* to a surface coverage *b* of black and *(1-b)* of white.

2.2 Multi-color dithering

Let us now extend dithering to color. Suppose that we would like to print with 4 different color inks C_1, C_2, C_3, C_4 (called basic colors). At each pixel of the output pixmap, the color separation we use gives us the relative percentages of each of the basic colors (Fig. 1b), for example d_1 of color C_1, d_2 of color C_2, d_3 of color C_3 and d_4 of color C_4. One of the basic colors, for example C_4, may be white.

Extending dithering to multiple colors consists in intersecting the relative cumulative amounts of colors d_1, d_1+d_2, and $d_1+d_2+d_3$ with the dither function *t* (Fig. 1c). In the interval where d_1 (x) > *t(x)*, the output location will be printed with basic color C_1 (Fig. 1c). In the interval where $d_1(x)+d_2(x) > t(x)$ and $d_1(x) \le t(x)$, the output location will be printed with basic color C_2. In the interval where $d_1(x)+d_2(x)+d_3(x) > t(x)$ and $d_1(x)+d_2(x) \le t(x)$, the output location will be printed with basic color C_3. In the remaining interval where $d_1(x)+d_2(x)+d_3(x)+d_4(x) > t(x)$ and $d_1(x)+d_2(x)+d_3(x) \le t(x)$, the output locations are printed with basic color C_4. Multi-color dithering therefore converts the relative amounts d_1, d_2, d_3, d_4 of basic colors C_1, C_2, C_3, C_4 into relative coverage percentages and ensures by construction that the contributing colors are printed side by side.

Color dithering is generally applied with 4 basic colors, since a point in 3D color space within the printer's gamut can be described by a barycentric combination of 4 colors.

2.3 Color Separation

When given a set of basic colors (inks), each specified by its tri-stimulus value in a given 3D color space (RGB or CIE-XYZ space), the 3D volume covered by these basic colors, called the printable gamut, can be segmented into a set of mutually adjacent tetrahedra [4]. The vertices of each tetrahedron correspond to four neighboring basic colors. Many tetrahedrizations of a point set in 3D exist; there is however only a single tetrahedrization which ensures that the enclosing sphere of a tetrahedron does not include another tetrahedron. Properties of tetrahedrizations and methods of construction are well described in the literature ([12], chapter 20).

Color separation of an input tristimulus value (RGB or CIE-XYZ) is obtained by locating in the selected color space the tetrahedron enclosing the given tri-stimulus value and by finding the barycentric coefficients d_1, d_2, d_3, d_4 used to express the input tristimulus value as a linear combination of the tetrahedron's vertices. These barycentric coefficients give the relative amounts of basic colors C_1, C_2, C_3 and C_4 used to reproduce the input tristimulus value.

As an illustration for tetrahedral decomposition and interpolation, let us consider an RGB color cube whose vertices correspond to the basic colors black, red, green, blue, cyan, magenta, yellow and white (Fig. 2a). A color wedge (Fig. 2b) with vertices close to cyan, blue, red and yellow is reproduced by multi-color dithering using the available set of basic colors. This color wedge, a planar slice in RGB space (Fig. 2a), intersects all tetrahedra into which the RGB cube is decomposed. In each tetrahedron, the corresponding color wedge part is reproduced using the 4 basic colors associated with the tetrahedron's vertices.

The dither matrix used for producing Fig. 2b has been obtained by discretizing and renumbering an egg crate function ([6], section 9.3.1, Fig. 9.6a). The resulting dithered color wedge incorporates ring shaped screen elements. All contributing basic colors are printed side by side.

If color reproduction fidelity is an issue, a calibrated input device (for example a scanner) is needed which provides a mapping from RGB input device values to CIE-XYZ device-independent values. Color separation in CIE-XYZ space is possible by tetrahedral decomposition of the volume formed by the measured CIE-XYZ values of the basic colors (inks + paper white). Colors close to the basic colors will be reproduced correctly. However, colors requiring the combination of several basic colors may deviate from their desired CIE-XYZ value, depending on various parameters such as printer registration accuracy, dot gain and ink density distribution of printed screen dots. The present contribution does not deal with printer color calibration. However, a possible printer calibration may be achieved by printing a large number of samples covering the printer's gamut and by measuring their CIE-XYZ values in order to build a 3D calibration table providing a mapping between device independent input CIE-XYZ values and output space "predicted" CIE-XYZ values (CIE-XYZ values predicted by linear interpolation within each tetrahedron).

When printing with transparent inks, we may want to use the superposition of one or sevral pairs of selected inks as additional basic colors. This is easily done by measuring for each pair its CIE-XYZ tristimulus values and incorporating the superposition of the two inks as a new color C_j into the set of available basic colors.

In this paper, we generally assume that input color values are within the range of printable colors. If the input color values are not located within the range of printable colors, a gamut mapping method must be applied. Several gamut mapping methods are known to produce convenient results [19].

Appendix I gives an example of a color separation with a set of non-standard offset inks (see single color page insert).

1. The darkness of an image pixel is one minus its normalized intensity value (in the graphic arts, a maximal intensity value is white and a minimal intensity value is black).

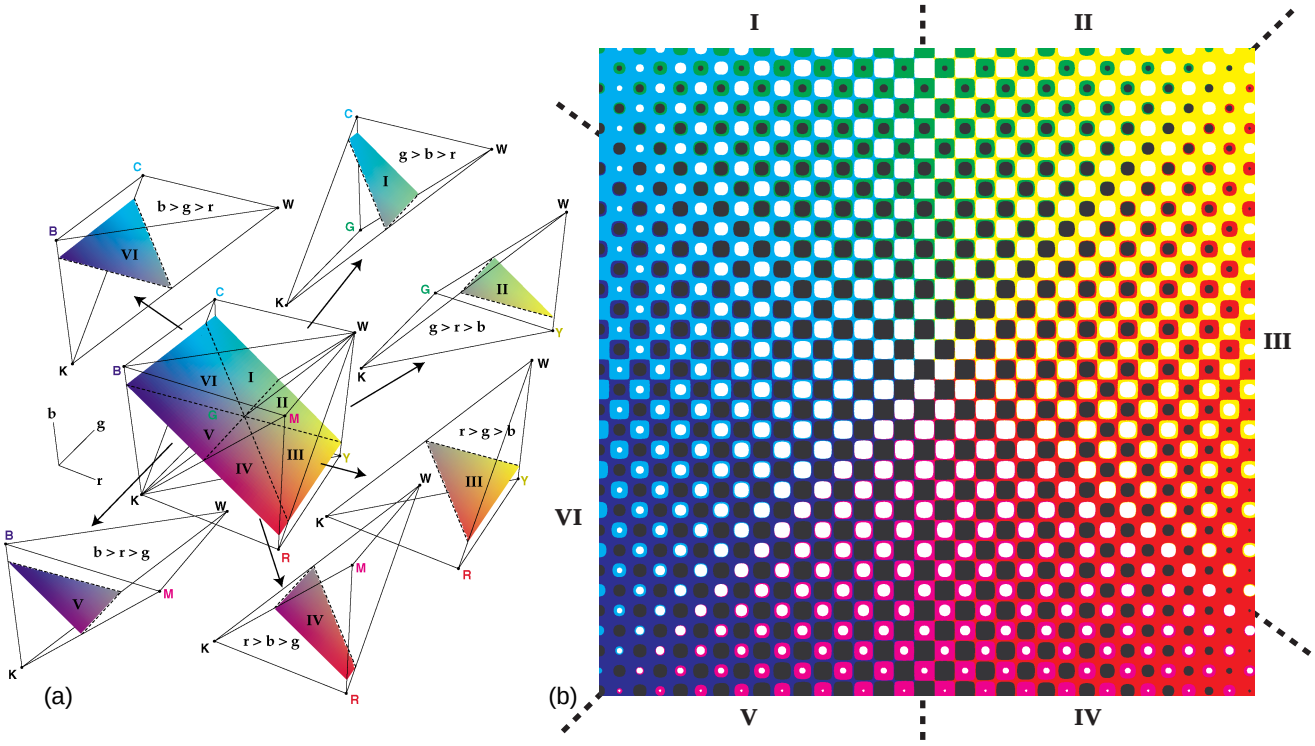

Fig. 2 (a) Color separation by tetrahedral decomposition of the RGB cube, (b) example of color dithered cyan - blue - red- yellow wedge with all basic colors printed side by side.

3. Synthesis of partially continuous, partially random dither matrices

For the purpose of artistic color screening, we need to generate dither matrices incorporating visually appealing symbols and ornamental motives. At low intensities however, one process color (for example black) may become very dominant and the corresponding screen motive may become too large to be recognized.

In order to solve this problem, we propose to generate dither matrices which produce partially ordered and partially random screen dots. To avoid enlarging the screen motive too much when strongly increasing the color surface coverage, we allow the background to contribute to the corresponding color. In the case of a black screen motive, from a certain darkness level, the background becomes successively darker. Thus instead of enlarging the motive shape, its contrast to the background becomes successively less pronounced, until it finally vanishes (Fig. 4b). One may observe that with a partially random screen, the desired screen motive remains visible over a larger intensity range than with a standard clustered screen. The final result however depends on dot gain: if dot gain is small, this effect is clearly visible. When dot gain is important, this effect is counterbalanced by the decreased contrast due to dot gain at high ink surface coverages.

It is relatively easy to create dither matrices generating partially clustered, partially random screens. One may first generate a continuous screen function representing the clustered part of the screen dot and sample it at each dither matrix cell (Fig. 3a). Then a very small amplitude noise[1] can be applied either to all dither matrix cells or to dither matrix cells having low dither values (Fig. 3b). Dither matrix cells are successively numbered according to their respective dither values (histogram equalization). Their ordinal numbers represent, after normalization, their dither threshold levels (Fig. 3c).

A color image dithered with such a mixed dither matrix preserves the screen motive better than with a standard dither matrix, see for example the still visible screen motive in the dark hair of the girl (Fig. 7b)

1. For the sake of simplicity, we apply white noise.

4. Equilibration of large dither matrices

In order to create artistically screened color images incorporating sophisticated screen shapes, we need to synthesize large dither matrices which may cover surfaces of several square millimeters. Since the motives incorporated into such dither matrices may not be well balanced, i.e. specific regions within a large dither matrix may not have a flat histogram of threshold values, alternations of dark and light within the produced screen dots may generate light and dark strips over the resulting dithered color image. This phenomenon is accentuated by dot gain since middle and dark tones tend to become darker.

The dither matrix equilibration method we propose compensates the uneven local surface coverage of the screen motive. It takes into account the dot gain and the human visual system transfer function ([23], chapter 7). It is related to model-based halftoning methods [16], but instead of modifying the final halftone image, it modifies the dither matrix used to produce the final image.

We assume that dot gain is similar for all contributing inks. We therefore equilibrate the dither matrix for a single ink, the black ink. Dither matrix equilibration is an iterative process. We start with the initial dither matrix containing the desired screen motive. We define a number of uniform gray patches equally spread throughout the available intensity range, for example patches at 16 representative intensity levels g_1 to g_{16}. These grayscale patches undergo a transformation comprising dithering, application of dot gain, and application of the human visual system transfer function. The transformation includes the following detailed steps:

(1) With the dither matrix, generate the halftone patches corresponding to the desired representative grayscale values

(2) For each generated halftone patch, take dot gain into account by adding into each pixel a darkness value representing the dot gain of the black neighboring pixels. In our model and for our target electrographic printer we consider that, due to dot gain, each black pixel adds 20% of blackness to its vertical and horizontal neighbors, and 5% of blackness to its diagonal neighbors [18]. We limit the maximal blackness of any pixel to 100%.

(3) Apply to the resulting grayscale images a Gaussian low-pass filter approximating to some extent the low pass behavior of the human transfer function. Based on the estimation of about 30 cycles per degree for the cutoff frequency of the human visual system ([13], chapter 7), we approximate the human visual system

transfer function by the Gaussian function $F(q) = Exp(-\pi q^2)$, where the unit on the frequency axis (q-axis) corresponds to the cutoff frequency of 30 cycles per degrees. The corresponding impulse response, i.e. the inverse Fourier Transform of $F(q)$, is also a Gaussian function, $f(r)=Exp(-\pi r^2)$, whose unit (r-axis) corresponds to 1/30 degree of visual angle. Considering a large observation distance (25", twice the normal observation distance), where details of the screen motive should disappear, 1/30 degree corresponds, at 1200 pixels/inch, to 17.45 pixels. To produce the discrete convolution kernel, this Gaussian impulse response function is sampled on a 5σ x 5σ grid, where $\sigma=1/Sqrt(2\pi)$ corresponds on our pixel grid to 7 pixels. Please note that for different printing resolutions, as well as for different observation distances (e.g. for posters to be observed from far away) the discrete convolution kernel needs to be recomputed accordingly.

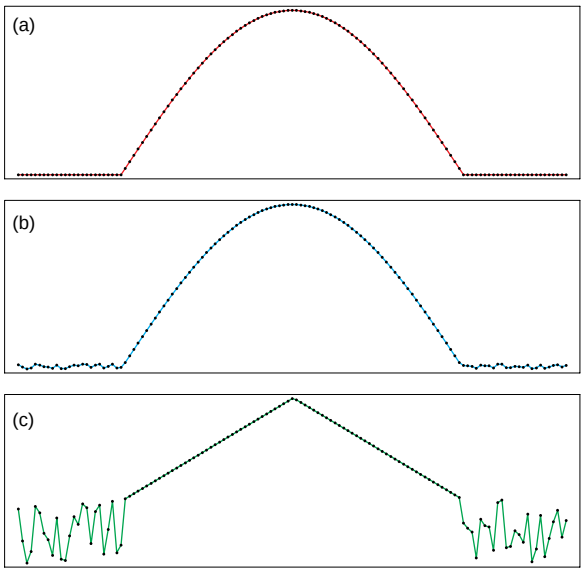

Fig. 3 Producing a dither matrix comprising partially continuous, partially random threshold levels.

The resulting transformed gray patches h_1 to h_{16} correspond to the original uniform grayscale images perceived by a human observer under given conditions (dot gain, resolution, distance).

The difference in gray values between the original uniform grayscale patches and the resulting halftoned patches specify the modification to be applied to the original dither matrix $d(x,y)$ in order to produce a modified dither matrix $d'(x,y)$ which compensates for the locally uneven motive distribution, i.e. the locally non uniform distribution of dither levels. The compensated dither matrix $d'(x,y)$ is obtained by successively compensating the differences associated with the different representative grayscale levels.

In a first step, a subset of threshold levels of the compensated dither matrix $d'(x,y)$ is computed by scaling the threshold levels of dither cells between 0 and the first representative grayscale value g_1 according to a value proportional to the difference between the desired grayscale level g_1 and grayscale level h_1 obtained by the previously described transformation:

```
if d(x,y) <= g1 then d'(x,y) = d(x,y) + k (h1(x,y)-g1)
```

where $d(x,y)$ represents the original dither threshold matrix, and k a scaling factor, chosen between 1/2 and 1.

This compensation is repeated in further steps for all successive representative grayscale levels and finally all dither cells of the original threshold matrix are corrected to yield the compensated dither matrix $d'(x,y)$.

The resulting algorithm can formulated as follows in pseudo-code:

```
for all successive representative intensity levels gi
  for all cells (x,y)
     if gi-1 < d(x,y) <= gi) then
         d'(x,y) = d(x,y) + k (hi(x,y)-gi);
  endfor;
endfor;
```

Fig. 5 graphically shows the application of the equilibration process to a one dimensional dither function $d(x)$.

Further equilibration passes, starting from the resulting corrected dither function $d'(x,y)$ may bring further improvements. The convergence is fast and a small number of passes is sufficient.

The effect of the equilibration process on the produced dither matrix can be observed in Fig. 6b and Fig. 6c. Small surfaces of low dither levels (white) surrounded by large surfaces of high dither levels tend to become larger. Correspondingly, the large surfaces with the high dither levels (black) tend to shrink. As an example, look at the internal cavities within the Allah motive (الله)

Fig. 7a shows a color image produced without equilibration (dither matrix of Fig. 6b) and Fig. 7b a color image with equilibration (dither matrix of Fig. 6c). The equilibration process considerably improves the quality of the resulting image and enables the generation of large dither matrices incorporating visually significant ornamental patterns without explicitly taking care of the pattern's local distribution of dither threshold levels (ideally, i.e. without taking into account dot gain, dither threshold levels over a given surface should be uniformly distributed).

Fig. 4 Grayscale wedge produced (a) with a standard continuous dither matrix and (b) with a mixed dither matrix made of partially continuous, partially random threshold levels.

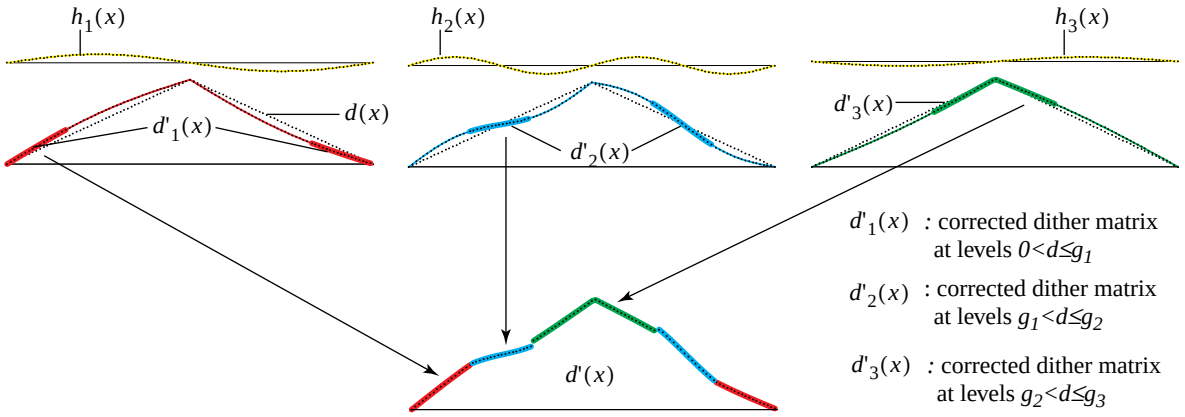

Fig. 5 Equilibrating a 1D dither function *d(x)* to produce the equilibrated dither function *d'(x)*

5. Artistically dithered images

By creating large size artistic dither matrices and by equilibrating them appropriately, one may create artistic color screens which enable the generation of color images incorporating two layers of information: a macro layer, i.e. the global image, and a micro layer formed by the artistic color screens. As a characteristic design example, we would like to render an image using the square Kufi arab script style for “Allah”, which can be found on the mosque of Badra, Azerbaijan ([3], page 182). The first step consists in designing a continuous screen made of several functions: paraboloid cylinders for the horizontal and vertical segments and half-spheres at segments ends (Fig. 6a). This continuous screen is sampled at the center of dither matrix cells and converted into an array of dither values. A small random value (noise) is added. The consecutive renumbering of cells according to their dither values yields a dither matrix (Fig. 6b), to which the equilibration process is applied (Fig. 6c). The resulting dither matrix is used to produce an output image (Fig. 7b) with red, green, blue, cyan, magenta, yellow, black and white basic colors. One can observe from a certain distance that the produced halftoned image is smooth despite the large size of the screen element and the uneven surface coverage behavior of the original screen function.

The next example shows the generation of a screen element of high aesthetic value laid out according to the beautiful *Cairo tessellation* ([11], page 119). Fig. 9a to Fig. 9e show how this threshold matrix is constructed. First, we define a simple analytical bell function on a unit square (Fig. 9a). Then, we take five halves of this bell function (Fig. 9b) that we project, by a set of linear transformations on a unit square as shown in Fig. 9c. We obtain a five-petals flower pattern. We take twelve such patterns, and, by applying twelve different linear transformations according to [11], we superpose them on a unit square thus obtaining the final analytical function shown in Fig. 9d. This analytical function is converted as in the previous examples into a set of discrete threshold values - our final dither matrix represented in Fig. 9e.

The example shown in Fig. 8 consists of an image rendered with a parallelogram dither array whose dither values form the shape of a celtic spiral motive. The thickness of the spiral varies from very thin in light tones to relatively thick in dark tones. Beyond a certain darkness level, the motive doesn't become thicker, but thanks to a partially ordered, partially random dither matrix, the background of the motive becomes darker. Please note the perfectly smooth transitions between the red spirals in light tones and the black spirals in dark tones.

The final example shown in Appendix I illustrates a real example of a graphic design, dithered either with typographic character shapes or with an oriental motive and printed with non-standard inks on an offset press.

At a first glance, the produced artistically dithered color images have some resemblance with images produced by black-white artistic screening [15]. The applied techniques however are completely different and significant differences between images produced by black-white artistic screening and by artistic color dithering exist. In dark tones, artistically screened black-white images have motives which nearly fill up the screen element space, whereas with color dithering, the size of the visible screen motive grows up to a certain limit; after that limit, the contrast between background and foreground is successively reduced until it vanishes.

In addition, artistic color dithering is able to render any color image and therefore has a much larger application range than artistic screening, which is limited to black-white or duo-tone.

Regarding implementation issues, artistic dither matrices can be designed in different ways, either by programming continuous 2D functions and sampling them into discrete threshold values or by starting from bi-level shape designs, transforming them into gray-scale intensity images (with PhotoShop for example) and applying histogram equalization to generate the dither threshold levels. Once the initial dither matrix is produced, the remaining part of the work is rather automatic. Dither matrix equilibration is slow, but is done only once. After having prepared the final dither matrix, reproducing the image requires color separation by tetrahedral interpolation and dithering. To improve image rendition speed, one can build a 3D look-up table establishing a pre-computed mapping between input tri-stimulus values and output color percentages of the contributing basic colors. With such a look-up table, we expect that multi-color dithering of images using pre-computed dither matrices can be made as fast as other color reproduction techniques (color error diffusion).

6. Conclusions and perspectives

Multi-color dithering is a generalization of standard bi-level dithering. Combined with tetrahedral color separation, multi-color dithering makes it possible to print images made of a set of non-standard inks. In contrast to most previous color halftoning methods, multi-color dithering ensures by construction that the different selected basic colors are printed side by side.

In this contribution, we extended bi-level dithering to multi-color dithering and explored multi-color dithering in the context of artistic color screening. To generate high-quality dithered color images incorporating artistic screen motives, we developed two dither matrix postprocessing techniques, one for enhancing the visibility of screen motives and one for the local equilibration of large dither matrices. By combining within the same dither matrix continuous threshold levels for the screen motive and randomly distributed dither threshold levels for the background, we enhance the visibility of the generated screen shapes at high ink saturation levels. The dither matrix equilibration process we propose to avoid disturbing local intensity variations takes both the physical behavior of the printer, i.e. the dot gain and the human visual system modulation transfer function into account.

Thanks to the combination of the presented techniques, high quality images can be produced, which incorporate at the micro level the desired artistic screens and at the macro level the full color image. Possible applications include innovative designs for publicity and posters.

Multi-color dithering clears the way for further color reproduction applications. It is known that some electrographic printing processes [2] require that different toners be placed beside one another (no overlap allowed). Security printing may make use of multi-

color dithering in order to print side by side with non-standard inks at a high registration accuracy. In that context, many new issues arise, such as the selection of the inks and the design of screen motives making the original very difficult to replicate, both by professional craftsmen and by simple color photocopying. Multi-color dithering also offers new perspectives for printing with special inks, such as fluorescent and metallic inks.

7. Acknowledgement

We would like to thank Orell Füssli Security Printing Ltd, Zürich, Switzerland, for collaborating with us on this project. Thanks also to David Salesin, Eric Stollnitz and Nicolas Rudaz for helpful discussions. This work was partly supported by the Swiss CTI (Grant 3776.1) and by the Swiss National Science Foundation (Grant 21-54127.98).

8. References

[1] H. Boll, "A Color to Colorant Transformation for a Seven Ink Process", in *Device Independent Color Imaging* (Ed. E. Walowit), SPIE Proceedings, Volume 2170, pages 108-118, 1994.

[2] J. Geraedts, S. Lenczowski, "Océ's productive colour solution based on the Direct Imaging Technology", Proceedings IS&T International Conf. on Digital Printing Technologies (NIP-13), pages 728-733, 1997.

[3] I. Hargittai, M. Hargittai, *Symmetry, A Unifying Concept*, Shelter Publ., 1994.

[4] P.C. Hung, Colorimetric calibration in electronic imaging devices using a look-up-table model and interpolations, *Journal of Elecronic Imaging*, 2 (1), pages 53-61, 1993.

[5] International Color Consortium. Specification ICC.1:1998-09. http://www.color.org.

[6] H.R. Kang, *Color Technology for Electronic Imaging Devices*, SPIE Publication, 1997.

[7] R.V. Klassen, R. Eschbach, K. Bharat, Vector Error Diffusion in a Distorted Colour Space, Proc. of IS&T 47th Annual Conference, 1994, Reprinted in *Recent Progress in Digital Halftoning*, (Ed. R. Eschbach), IS&T Publication, pages 63-65, 1994.

[8] K.T. Knox, Printing with Error Diffusion, in *Recent Progress in Digital Halftoning* (Ed. R. Eschbach), IS&T Publication, pages 1-5, 1994.

[9] H. Küppers. *Die Farbenlehre der Fernseh-, Foto- und Drucktechnik: Farbentheorie der visuellen Kommunikationsmedien.* DuMont Buchverlag, Köln, 1985.

[10] US Patent 4,812,899, issued March 14, 1989, filed Dec 19, 1986, Inventor: H. Kueppers.

[11] G.E. Martin, *Transformation Geometry. An Introduction to Symmetry.* Springer-Verlag, 1982.

[12] G.M. Nielson, H. Hagen, H. Müller, *Scientific Visualization*, IEEE Computer Society, 1997.

[13] L. Olzak, J.P. Thomas, Seeing spatial patterns, in *Handbook of perception and human performance, (*Eds. K. R. Boff, L. Kaufman, J. P. Thomas), Chapter 7, J. Wiley, pages 7-1 to 7-55, 1986.

[14] V. Ostromoukhov, Chromaticity Gamut Enhancement by Heptatone Multi-Color Printing, in *Device-Independent Color Imaging and Color Imaging Systems Integration*, Proc. SPIE, Vol. 1909, pages 139-151, 1993

[15] V. Ostromoukhov, R.D. Hersch, Artistic Screening, *Proceedings of SIGGRAPH 95*, Annual Conference Series, pages 219-228, 1995.

[16] T. N. Pappas, Model-based halftoning of color images, IS&T 8th International Congress on Advanced in Non-Impact Printing Technologies, 1992, reproduced in *Recent Progress in Digital Halftoning* (Ed. R. Eschbach), IS&T Publication, pages 144-149, 1994.

[17] G.L. Rogers, "Neugebauer Revisited: Random Dots in Halftone Screening", *Color Research and Applications*, 23 (2), pages 104-113, 1998.

[18] C.J. Rosenberg, "Measurement based verification of an electrophotographic printer dot model for halftone algorithm tone correction", IS&T 8th International Congress in Non-Impact Printing Technologies, Oct. 25-30, 1992, Reprinted in *Recent Progress in Digital Halftoning*, (Ed. R. Eschbach), IS&T Publication, pages 159-163, 1994.

[19] E. J. Stollnitz, V. Ostromoukhov, D. H. Salesin, Reproducing Color Images Using Custom Inks, *Proceedings of SIGGRAPH 98*, Annual Conference Series, pages 267-274, 1998.

[20] J. R. Sullivan, R. L. Miller, T. J. Wetzel, *Color digital halftoning with vector error diffusion*, US Patent 5,070,413, 1991.

[21] R. Ulichney, *Digital Halftoning*, The MIT Press, Cambridge, Mass., 1987.

[22] R. L. Van Renesse (Ed.), *Optical Document Security*, Artech House, 1998.

[23] B.A. Wandell, *Foundations of Vision*, Sinauer Associates, Inc. Publishers, Sunderland, Mass., 1995.

[24] J.A.C. Yule. *Principles of Color Reproduction.* John Wiley & Sons, New York, 1967.

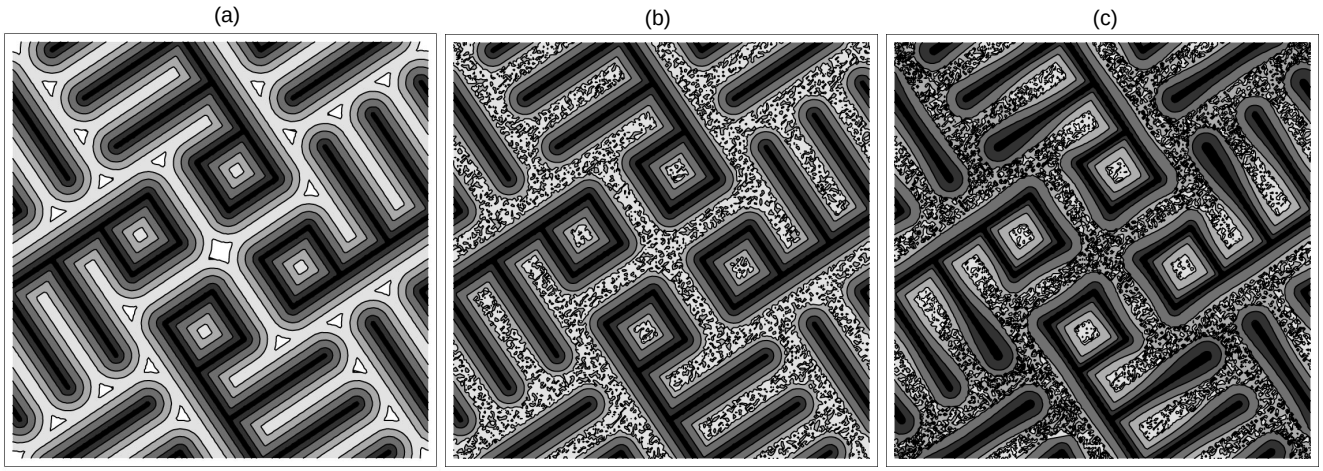

Fig. 6 "Allah" motive, from Mosque of Badra, Azerbaijan: (a) the initial dither matrix, (b) the mixed dither matrix incorporating continuous and random dither levels and (c) the equilibrated mixed dither matrix.

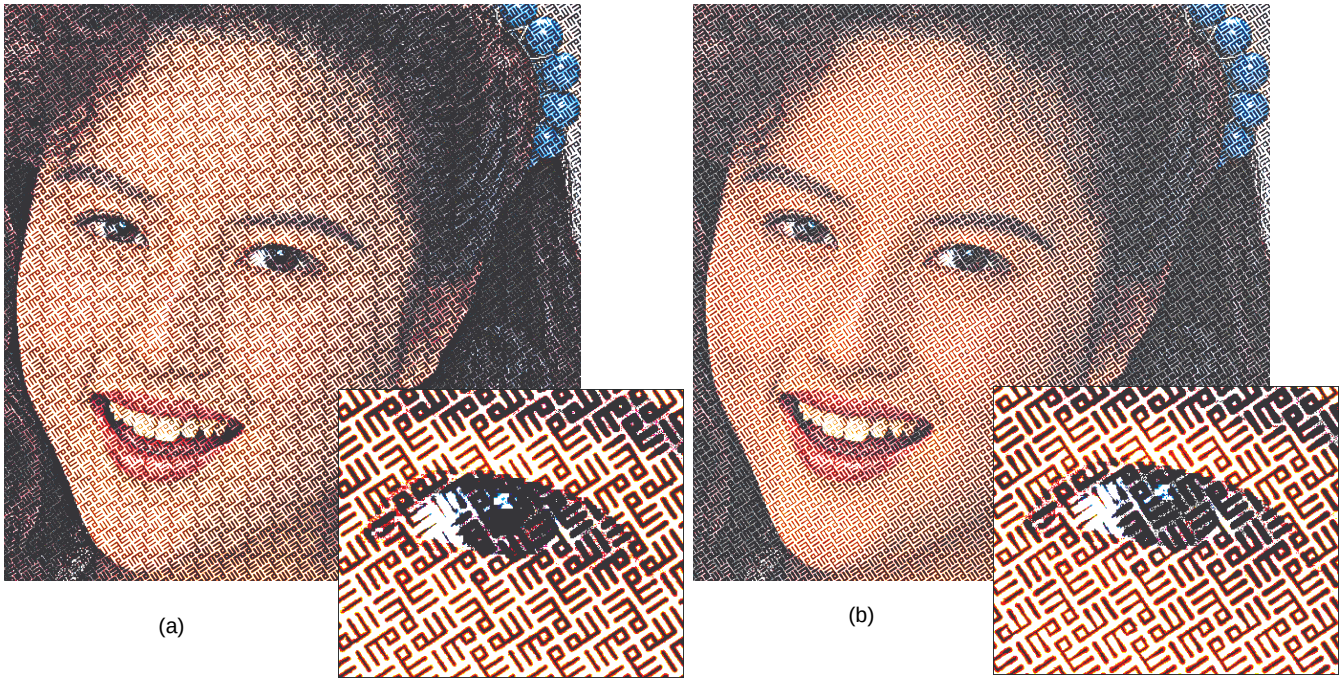

Fig. 7 (a) Color image produced without and (b) with the equilibration process

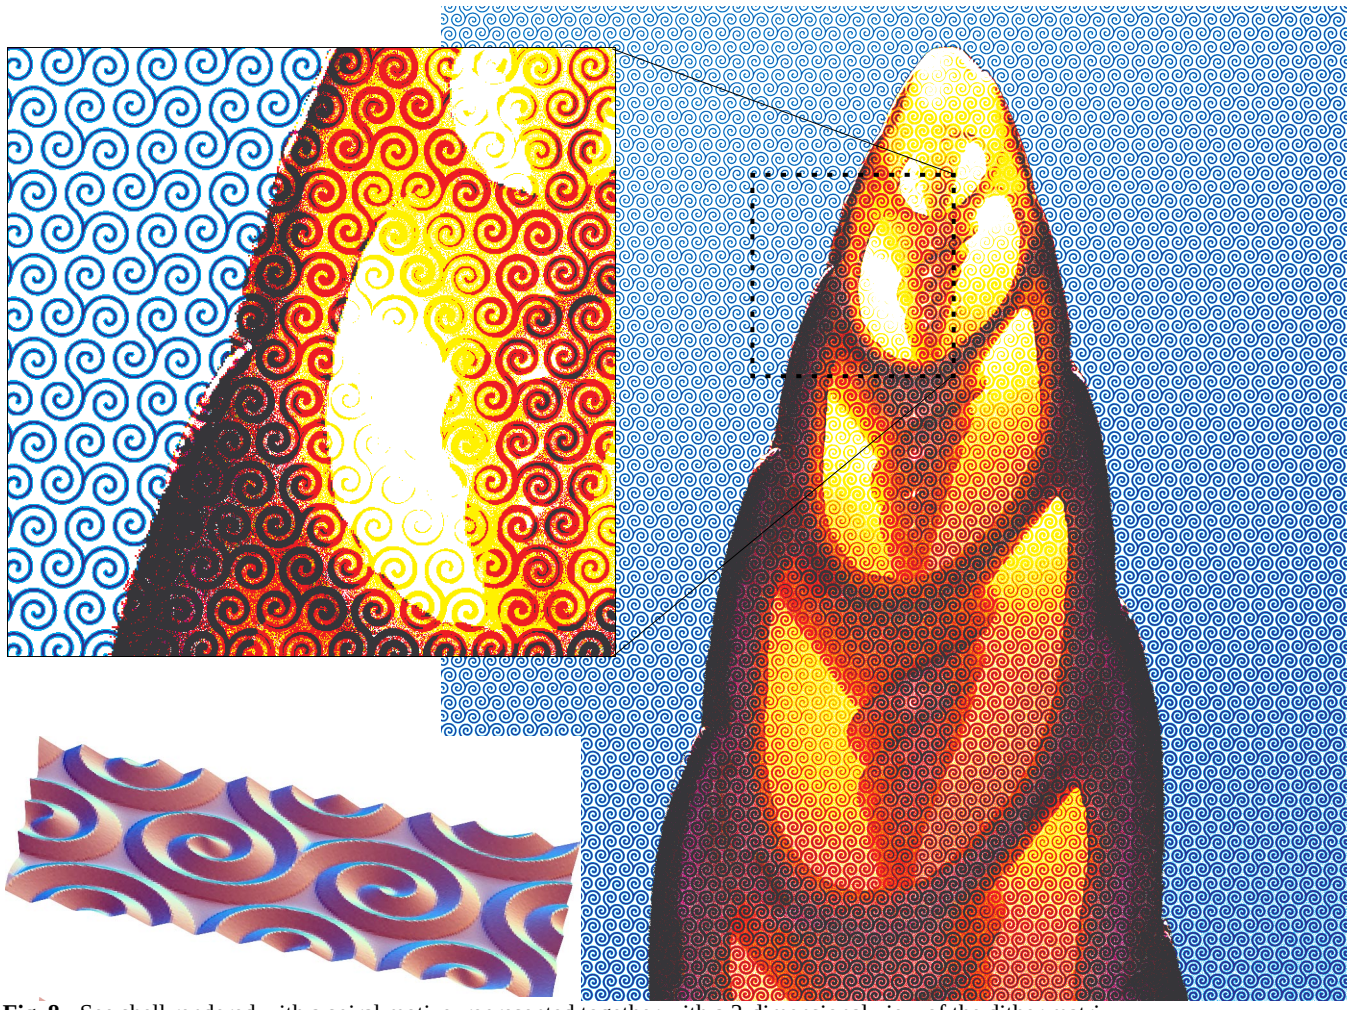

Fig. 8 See shell rendered with a spiral motive, represented together with a 3-dimensional view of the dither matrix.

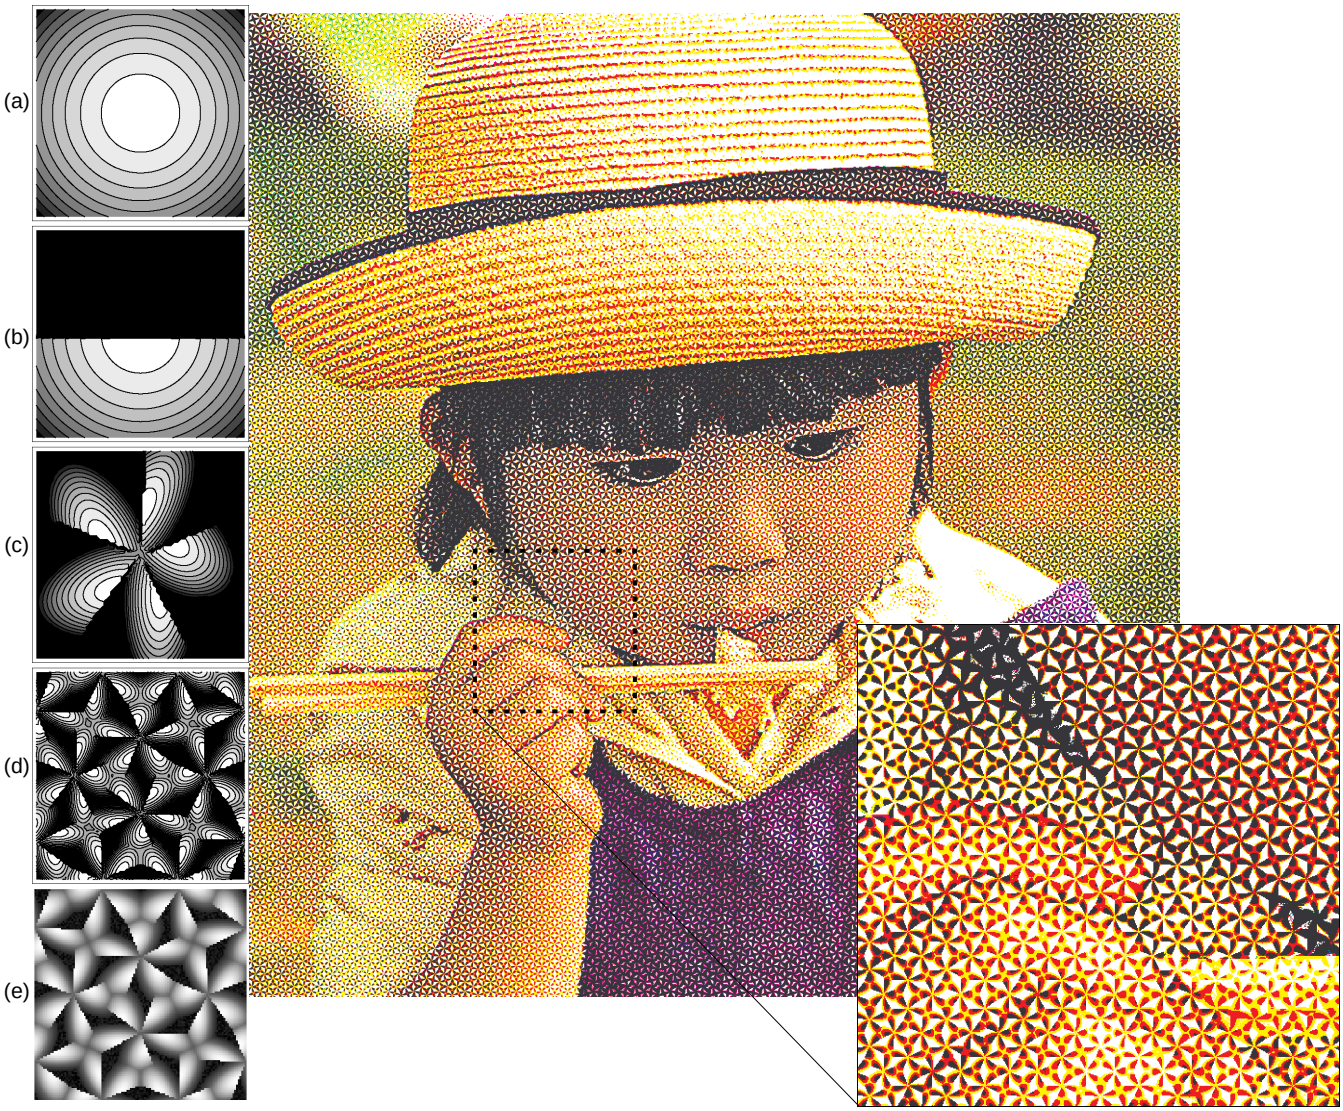

Fig. 9 A sample image produced using a threshold matrix inspired by the *Cairo tessellation.* Figures (a)-(e) show the building process of the threshold matrix.

Appendix I

Images produced on a high-registration offset press, at resolution YYY with non-standard inks XXXX, courtesy Orell Füssli Security Printing Ltd, Zürich, Switzerland.