MIT Media Laboratory Perceptual Computing Group Technical Report #336.
Also appears in: Proc. of the Second IEEE Int. Conf. on Image Processing, Washington, DC, Oct. 1995.

COLOR HALFTONING WITH M-LATTICE

Alex Sherstinsky and Rosalind W. Picard

Media Laboratory, E15-383
Massachusetts Institute of Technology, Cambridge, MA 02139
shers@media.mit.edu and picard@media.mit.edu

ABSTRACT

This paper presents new results in the halftoning of color images. The first automatic generation of the Wall Street Journal type halftones of gray-scale images was done using the M-lattice system and reported in [1]. The M-lattice system was derived from the reaction-diffusion model, first proposed by Turing in 1952 in order to explain mammal coat patterns [2]. The M-lattice is a non-linear dynamical system that is well-suited for a variety of applications formulated as constrained non-linear optimization. In particular, it can perform image processing operations that emphasize oriented patterns [3]. The present study uses this property to extend the special-effects halftoning of gray-scale images to that of color images. The quality metric comes from the directionality information extracted by steerable filters [4], [5] from the gray-scale version of the original color image. The binary requirement is stated as an explicit constraint, and all three (red, green, and blue) halftone components are synthesized simultaneously by the M-lattice.

1. INTRODUCTION

The color image halftoning technique discussed in this paper is based on the method of non-linear programming with an orientation-sensitive quality metric [1], [3]. The computational substrate for solving the non-linear program is the M-lattice [6]. This system is rooted in the reaction-diffusion model, first proposed by Turing in 1952 to explain the formation of animal patterns such as zebra stripes and leopard spots. The M-lattice is bounded and has a lot of flexibility in how its variables can interact. In particular, it is well-suited to a variety of applications formulated as constrained non-linear optimization.

Faithful halftoning is the task of tricking the human visual system into seeing exactly the original multi-tone picture in a replica image consisting of only the two extreme intensities. While the faithful halftoning of color images is a mature discipline it is still a challenging problem. The reason is that including the color information makes the concerns encountered in gray-scale halftoning all the more complicated [7]. For example, the non-linear effects due to binarization exacerbate the moire patterns, while the printer imperfections create more visually apparent artifacts. The state of the art techniques for faithful halftoning (and, more generally, quantization) of color images can be found in [8], [9] and references therein. These papers cover a number of central issues in color printing. The concept of utilizing the properties of human visual system for the quantization of color images is analyzed in [8]. Visual models are employed to develop an efficient quantization algorithm in luminescence-chrominance color space, which produces perceptually high-quality quantized images. The issues of printer distortions and how to compensate for them are elaborated in [9].

On the other hand, special-effects halftoning is a relatively new direction [1], [10]. As the name implies, the goal is the automatic synthesis of caricatures that accentuate certain desirable features of the given image. For example, many newspaper portrait styles emphasize lines and curves in the original image.

One distinctive attribute of the special-effects halftones is the fact that the error, instead of being uniformly diffused, is directed into places that exaggerate desirable aspects of the image. Hence, the usually undesirable error is reshaped into a perceptually-pleasant feature.

The main contribution of this paper is the extension of the gray-scale special-effects halftoning technique reported earlier [1] to produce a new algorithm, which performs the automatic synthesis of color caricatures in the style of the Wall Street Journal portraits [1].

In order to make this paper self-contained, a brief review of the M-lattice system(the computational vehicle) and of the orientation extraction (system parameters) appears in Section 2 and Section 3, respectively. Section 4 describes the special-effects color halftoning technique and illustrates it with two examples. Section 5 summarizes the study.

2. BACKGROUND: M-LATTICE SYSTEM

We briefly review the essentials of the M-lattice system [1]. Let $\psi_i(t) \in \Re$ be a state variable as a function of time at each lattice point i, where $i = 1, \ldots, N$. Let $\chi_i(t)$ be an output variable, obtained from $\psi_i(t)$ via $\chi_i(t) = g(\psi_i(t))$. The "warping" function, $g(u)$, is a saturating piece-wise linear non-linearity with an arbitrarily large number of segments. The values of $\chi_i(t)$ will correspond to the intensities of the pixels in the output image at the time when the system has converged. Construct $\vec{\psi}(t)$ and $\vec{\chi}(t)$ by concatenating $\psi_1(t), \ldots, \psi_N(t)$

This research was sponsored by Hewlett-Packard Laboratories, Palo Alto, California.

[1] This paper focuses exclusively on the special-effects halftoning, even though the M-lattice system has been used to synthesize high-quality faithful halftones as well [1].

and $\chi_1(t),\ \ldots,\ \chi_N(t)$, respectively into column vectors.

Definition 2.1 *Suppose that a given function, $\Phi(\vec{\chi}(t))$, is continuous, twice-differentiable, and bounded above. Let the matrix $\mathbf{A}$ be real, symmetric, and negative-definite:* $\mathbf{A} \in \Re^{N \times N}$, $\mathbf{A} = [a_{ij}]$, $\mathbf{A} = \mathbf{A}^T$, *and* $\forall\ i\ \lambda_i\,[\mathbf{A}] < 0$. *Then the M-lattice system* [2] *is the following non-linear dynamical system:*

$$\frac{d\vec{\psi}(t)}{dt} = \mathbf{A}\vec{\psi}(t) + \vec{\nabla}_{\vec{\chi}}\Phi(\vec{\chi}(t)). \tag{1}$$

Notice the right-hand side contains two components – a linear function of the state variables and the gradient of a typically non-linear function of the warped state variables. The convergence and stability properties of the M-lattice system are analyzed in [6]. For the present (halftoning) applications, (1) has exhibited convergence (in computer simulation) to fixed points of the form $\vec{\chi} \in \{-1, 1\}^N$ regardless of the initial conditions.

In non-linear optimization, $\Phi(\vec{\chi})$ is the objective function to be maximized. For certain types of objective functions, the M-lattice system converges to the (appropriately defined) local maxima of $\Phi(\vec{\chi})$ with respect to $\vec{\chi}$ [6]. Thus, in many situations it is advantageous to use the M-lattice system for non-linear optimization. In the examples described in Section 4, the directionality information defines the quality metric for the non-linear program, solved by the 3-lattice system.

3. BACKGROUND: ESTIMATING LOCAL ORIENTATION

We employ the computation-saving "steerable" set of basis filters described in [4]. Steerable filters have been shown to give a good match to orientation perception by humans [5]. The output of the steerable filters at each pixel i of the gray-scale version of the original color image gives the angle, $\theta_i \in [-\pi, \pi]$, and relative strength (or magnitude), $m_i \in [0, 1]$, of the dominant orientation present at that pixel.

In the Wall Street Journal type halftoning, we use the orientation to guide the action of the M-lattice system. For example, to design a low-pass adaptive filter that rotates to the dominant orientation, denote the diagonal matrix of variances by $\mathbf{V}_i$ and the rotation matrix by $\mathbf{\Theta}_i$:

$$\mathbf{V}_i = \begin{bmatrix} \sigma_{i,x}^2 & 0 \\ 0 & \sigma_{i,y}^2 \end{bmatrix},\ \mathbf{\Theta}_i = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix}. \tag{2}$$

The relative sizes of $\sigma_{i,x}^2$ and $\sigma_{i,y}^2$ depend on m_i and determine the skewness of filters with respect to the dominant orientation:

$$\sigma_{i,y}^2 = \frac{L}{2}(1 - m_i),\ \ \sigma_{i,x}^2 = L - \sigma_{i,y}^2, \tag{3}$$

where $L \times L$ is the size of the filter mask in pixels. Let $\vec{n} \in \mathcal{Z}^2$ be the pixel position. Then the (unnormalized) oriented low-pass filter is given by:

$$h_i^u(\vec{n}) = \exp\left\{-\vec{n}^T \mathbf{\Theta}_i^T \mathbf{V}_i \mathbf{\Theta}_i \vec{n}\right\}. \tag{4}$$

[2]This is the definition adapted for the present paper. The general M-lattice system is defined in [6].

4. ORIENTATION-DEPENDENT COLOR HALFTONING AS NON-LINEAR PROGRAM

We now consider the problem of synthesizing – for each RGB component – a binary caricature that brings out the directional content of the original color image. The resulting halftoning method must yield a composite halftone RGB image that appears similar to the original color image in the sense of preserving orientations. A least-squares halftoning approach is appropriate for this task, because it can employ an explicit model of perception and printer distortions as the measure of performance [11], [9]. Here we show how to implement such an approach using the M-lattice system and point out the additional benefits brought by using the M-lattice.

Suppose $\vec{n} \in \mathcal{Z}^2$; $s(\vec{n}) \in [-1, 1]$ is the finely quantized original input image signal; $y(\vec{n}) \in \{-1, 1\}$ is the output halftone image; and $h(\vec{n})$ is a 2-D filter (not necessarily the same as $h(\vec{n})$ in the previous section). Let $\mathbf{B} = \mathbf{H}^T\mathbf{H}$, where $\mathbf{H}$ is a circulant matrix with $h(\vec{n})$ in the first row. The problem of halftoning can be stated as a non-linear program:

$$\min_{\vec{y}} \frac{1}{2}\vec{y}^T\mathbf{B}\vec{y} - (\mathbf{B}\vec{s})^T\vec{y} \tag{5}$$

$$\text{subject to constraints: } y_i^2 - 1 \geq 0, \tag{6}$$

where the vectors are the standard concatenations of the corresponding sequences. The particular form of constraints, (6), forces each pixel to assume binary values.

In order to solve this problem using the M-lattice system, we combine the objective function to be minimized, (5), with the N constraints, (6), into the Lagrangian cost functional with the help of the Karush-Kuhn-Tucker conditions [12] simultaneously for all three color components:

$$\min_{\vec{y}} \mathcal{L}(\vec{y}), \quad \text{where}$$

$$\mathcal{L}(\vec{y}) = \frac{1}{2}\vec{y}^T\mathbf{B}\vec{y} - (\mathbf{B}\vec{s})^T\vec{y} + \frac{1}{2}\sum_i p_i(y_i^2 - 1), \tag{7}$$

$$p_i \leq 0, \quad p_i(y_i^2 - 1) = 0. \tag{8}$$

The Lagrange multipliers, p_i, are the varying penalty terms that enforce the constraints according to (8). As a result, the unconstrained minimization of $\mathcal{L}(\vec{y})$ in (7) produces the optimal halftone image.

The optimization problem, (7), is "programmed" onto the M-lattice system, (1), by setting $\vec{y}$ equal to $\vec{\chi}$, $\Phi(\vec{\chi})$ to $-\mathcal{L}(\vec{y})$, and taking partial derivatives. This yields:

$$\frac{d\vec{\psi}(t)}{dt} = \mathbf{A}\vec{\psi}(t) + \mathbf{B}\vec{s} - \mathbf{B}\vec{\chi}(t) - \mathbf{P}\vec{\chi}(t), \tag{9}$$

where $\mathbf{P} = \text{Diag}\ \{p_1,\ \ldots,\ p_N\}$. The elements of $a(\vec{n})$ are chosen so as to guide the system towards an optimum corresponding to a perceptually-pleasant halftone. It has been shown that $\mathbf{A} = \mathbf{B} - \mathbf{I}$ is a good choice, because it filters out objectionable correlated spatial patterns [1].

While the gray-scale special-effects halftoning algorithm is implemented as a 1-lattice [1], the RGB color scheme used in the present study is organized as a 3-lattice. This organization can be advantageous in case the quality metric requires the interaction of the color components [3].

Treating halftoning as a non-linear programming problem and solving it with the M-lattice system offers considerable flexibility in the choice of the quality metric and in the functional form of constraints. In order to demonstrate this flexibility, we incorporated orientation detection into the halftoning quality metric. The adaptive filter matrix, $\mathbf{H}$, was designed using (4) so as to include the information about the dominant orientation at each pixel of the gray-scale version of the original color image, shown in Figure 1(a). Figure 1(b) displays the result, which exhibits more of the line and curve features found in hand-drawn "halftones" (such as the Wall Street Journal portraits) [6]. Another example of this technique appears in Figure 2.

According to the poster provided by the Wall Street Journal Classroom Edition program, the entire process of creating a monochrome halftone is done by hand and takes an artist from three to five hours [13]. In contrast, the simulation of the M-lattice system implementation of the color halftoning algorithm on the CM-2 takes 6000 iterations at the time step of 0.01 sec for the total time of approximately 20 minutes including the system time and the I/O.

5. SUMMARY

We have presented a method for halftoning color images automatically in the would-be style of the color Wall Street Journal portraits. As with gray-scale images, the RGB halftones are made optimally close to the original color components in the sense of preserving the dominant directions in the image. And the M-lattice system is the non-linear dynamical system employed to compute these optimal RGB components simultaneously.

Acknowledgements

The authors would like to thank Thrasyvoulos N. Pappas for helpful suggestions.

6. REFERENCES

[1] A. Sherstinsky and R. W. Picard, "M-Lattice: A novel non-linear dynamical system and its application to halftoning," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 2, (Adelaide, Australia), pp. 565–568, Apr. 1994.

[2] A. M. Turing, "The chemical basis of morphogenesis," *Philosophical Transactions of Royal Society of London*, vol. 237, no. B, pp. 37–72, 1952.

[3] A. Sherstinsky and R. W. Picard, "Orientation-sensitive image processing with M-Lattice – a novel non-linear dynamical system," in *Proceedings of First IEEE International Conference on Image Processing*, vol. 3, (Austin, TX), pp. 152–156, Nov. 1994.

[3] The present study treats the RGB components independently for simplicity.

[4] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. on Pat. Anal. and Mach. Intell.*, vol. PAMI-13, pp. 891–906, Sept. 1991.

[5] R. W. Picard and M. Gorkani, "Finding perceptually dominant orientations in natural textures," *Spatial Vision, special Julesz birthday issue*, vol. 8, no. 2, pp. 221–253, 1994.

[6] A. S. Sherstinsky, *M-Lattice: A System For Signal Synthesis And Processing Based On Reaction-Diffusion*. ScD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1994.

[7] "Private communications with Thrasyvoulos N. Pappas, AT&T," Mar. 1995.

[8] R. Balasubramanian, C. A. Bouman, and J. P. Allebach, "Sequential scalar quantization of color images," in *IS & T Final Program and Advance Printing of Paper Summaries IS & T's 46th Annual Conference*, pp. 97–101, 1994.

[9] T. N. Pappas, "Model-based halftoning of color images," in *Proc. IS & T's Eigth International Congress on Non-Impact Printing Technologies*, pp. 270–275, 1992.

[10] Y. Pnueli and A. M. Bruckstein, "Gridless halftoning: A reincarnation of the old method," CIS Report #9323, Center For Intelligent Systems, Technion – Israel Institute of Technology, 1993.

[11] T. N. Pappas and D. L. Neuhoff, "Least-squares model-based halftoning," in *Proc. SPIE Conf. on Elect. Imaging Sci. and Tech.*, (San Jose, CA), Feb. 1992.

[12] M. S. Bazaraa and C. M. Shetty, *Nonlinear Programming*. New York: John Wiley and Sons, 1979.

[13] "Private communications with Richard Kelman, Dow Jones," May 1994.

(a)

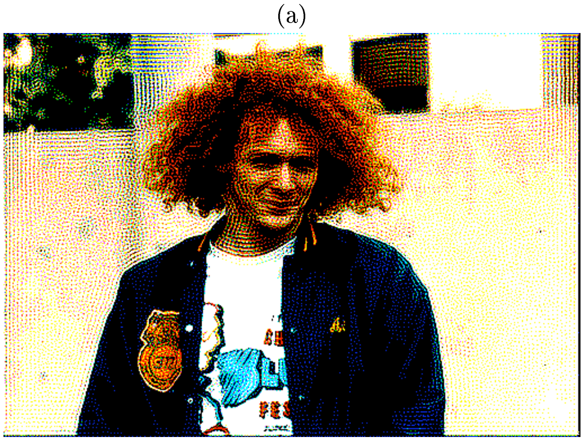

(b)

Figure 1: Orientation-sensitive color halftoning. (a) the original "Marty" image; (b) the "Marty" image adaptively halftoned using orientation information at each pixel of the original.

(a)

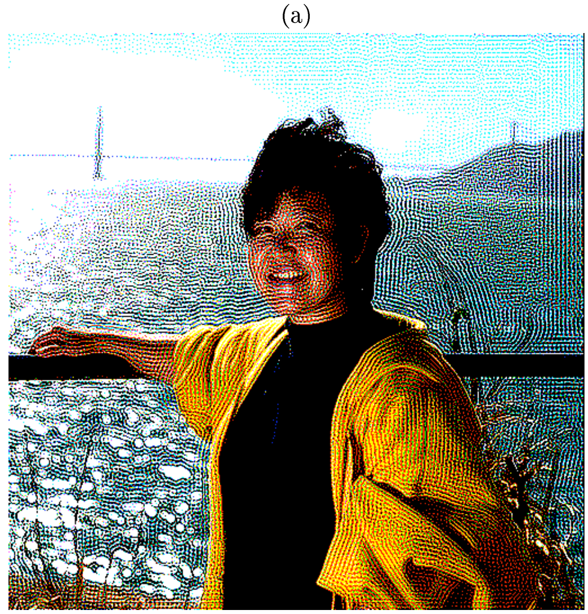

(b)

Figure 2: Orientation-sensitive color halftoning. (a) the original "Betty" image; (b) the "Betty" image adaptively halftoned using orientation information at each pixel of the original.