Halftoning With *M*-Lattice

Alex Sherstinsky
Media Laboratory, Massachusetts Institute of Technology
Cambridge, Massachusetts

Abstract

This paper summarizes the halftoning techniques that employ the M-lattice, a non-linear dynamical system recently introduced to the signal processing community. The *M*-lattice system was derived from the reaction-diffusion model, first proposed by Turing in 1952 in order to explain mammal coat patterns. The *M*-lattice system is closely related to the analog Hopfield network and the cellular neural network, but has more flexibility in how its variables interact. In particular, this model is well-suited for a variety of applications formulated as constrained non-linear optimization. The present overview demonstrates the use of this model for three different image halftoning examples. The first example synthesizes halftones free of correlated artifacts; it illustrates the noise-shaping capability of the *M*-lattice system. The second example synthesizes halftones in the creatively "hand-drawn" style of the Wall Street Journal portraits; it illustrates how a more flexible quality metric can be used when the binary requirement is stated as an explicit constraint. The third example extends this monochrome "special-effects" halftoning method to allow color images; all three (red, green, and blue) halftone components are synthesized simultaneously by the *M*-lattice.

Introduction

The present research has originated in the investigation of the usefulness of reaction-diffusion systems for modeling natural textures. A reaction-diffusion system is a set of heat equations coupled by, typically non-linear, reaction terms. The reaction-diffusion model was first proposed by Turing in 1952 in order to explain mammal coat patterns, such as zebra stripes and leopard spots. Until recently, reaction-diffusion systems have been researched predominantly by mathematical biologists working on theories of natural pattern formation and by chemists working on modeling the dynamics of complex chemical reactions[1]. However, the past three years have seen a significant surge in interest in reaction-diffusion systems, primarily for exploiting them in the areas of computer graphics and image processing[2,3,4].

In order to form patterns a valid reaction-diffusion system must exhibit local instability to small random perturbations. That notwithstanding, the system should be stable in the large-signal sense for practical reasons. A major difficulty associated with the reaction-diffusion paradigm in its standard form is that the system is stable only for a restricted class of non-linear reaction functions. This drawback narrows the scope of the model's engineering applications, due to numerical overflow.

A common approach aimed at preventing numerical overflow from plaguing the simulations of reaction-diffusion systems on the digital computer has been to clip the magnitudes of the state variables by adding an "if" statement to the numerical method (*e.g.*, Forward Euler) used for solving the system of differential equations[3]. However, this technique does not guarantee that the system will reach equilibrium; moreover, it destroys the mathematical integrity of the original dynamical system.

By using a warping function to facilitate stability, the *M*-lattice system allows more flexible non-linear interactions than the reaction-diffusion system. Three of the capabilities of this model are illustrated in an application to digital halftoning of images.

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[5]. 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[6,7] 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 [6]. 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 [7].

On the other hand, special-effects halftoning is a relatively new direction[8,9]. 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 color image halftoning technique discussed in this paper is based on the method of non-linear programming with an orientation-sensitive quality metric[8,10]. The computational substrate for solving the non-linear program is the *M*-lattice[11]. 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.

The present color halftoning algorithm the extension of the gray-scale special-effects halftoning technique reported earlier[8] to produce a new algorithm, which performs the automatic synthesis of color caricatures in the style of the Wall Street Journal portraits.

Background: *M*-Lattice System

We briefly review the essentials of the *M*-lattice system[8]. 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 $\psi_1(t), \ldots, \psi_N(t)$ and $\chi_1(t), \ldots, \chi_N(t)$, respectively into column vectors.

Definition *Suppose that a given function,* $\Phi(\vec{\chi}(t))$, *is continuous, twice-differentiable, and bounded above. Let the matrix* **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*[†] *is the following non-linear dynamical system*:

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

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 [11]. 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 $\vec{\chi}$.[11] Thus, in many situations it is advantageous to use the *M*-lattice system for non-linear optimization. In the examples described in **Halftoning**, the directionality information defines the quality metric for the non-linear program, solved by the 3-lattice system.

Background: Estimating Local Orientation

We employ the computation-saving "steerable" set of basis filters described in [12]. Steerable filters have been shown to give a good match to orientation perception by humans[13]. 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 V_i and the rotation matrix by Θ_i:

$$V_i = \begin{bmatrix} \sigma_{i,x}^2 & 0 \\ 0 & \sigma_{i,y}^2 \end{bmatrix}, \Theta_i = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix} \tag{2}$$

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 $\vec{n} \in Z^2$ be the pixel position. Then the (unnormalized) oriented low-pass filter is given by:

$$h_i^0(\vec{n}) = \exp\{-\vec{n}^T \Theta_i^T V_i \Theta_i \vec{n}\}. \tag{4}$$

Halftoning

Suppose $\vec{n} \in Z^2; s(\vec{n}) \in [-1,1]$ is the continuous-tone (or 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. The halftoning method must yield an image which appears perceptually similar to the original gray-scale image. Least-squares halftoning approaches receive continual attention, because they can employ explicit models of the human visual system and of the printing device[14].

Noise-Shaping Least-Squares Halftoning

It is generally agreed that error diffusion produces the best results in terms of artifacts[15]. However, the causality of the algorithm prevents it from making sharp transitions and tracking edges properly[16]. In contrast, the least-squares halftoning techniques render edges well, but suffer from granular artifacts. We show that the *M*-lattice system naturally combines noise shaping with least-squares optimization, thereby offering the benefits of both.

Given a least-squares halftoning technique, set $\Phi(\vec{\chi})$ to the negative of the distortion measure. For example, if

$$\Phi(\vec{\chi}) = (\mathbf{H}^T \vec{s})^T \vec{\chi} - \frac{1}{2} \vec{\chi}^\gamma \mathbf{H}^T \mathbf{H} \vec{\chi}, \tag{5}$$

then at an equilibrium (1) yields:

$$-\mathbf{A}\vec{\psi} = \mathbf{H}^T \vec{s} - \mathbf{H}^T \mathbf{H} \vec{\chi}, \text{ or} \tag{6}$$

$$-(\mathbf{A} - \mathbf{H}^T \mathbf{H})\vec{\psi} = \mathbf{H}^T \vec{s} - \mathbf{H}^T \mathbf{H}(\vec{\chi} - \vec{\psi}). \tag{7}$$

Now set $-(\mathbf{A}-\mathbf{H}^T\mathbf{H})=\mathbf{I}$, and let $\vec{q} \overset{\text{def}}{=} \vec{\chi} - \vec{\psi}$ be the quantization error (or the quantization noise)[17]. Then (6) and (7) become:

$$\vec{\psi} = \mathbf{H}^T \vec{s} - \mathbf{H}^T \mathbf{H} \vec{q} \tag{8}$$

Thus, according to (8), the *M*-lattice system performs non-causal error diffusion in the steady-state limit. For perceptual reasons, it is desirable to minimize the low-frequency content of the quantization error. Since $\mathbf{H}^T\mathbf{H}$ is a smoothing filter, $\mathbf{H}_Q \overset{\text{def}}{=} \mathbf{I}\text{-}\mathbf{H}^T\mathbf{H}$ becomes a high-pass filter. Then it follows that A $= -\mathbf{H}_Q$. The action of the high-pass noise shaping filter, $\mathbf{H}_Q$, gives the quantization noise the perceptually pleasant "blue" character[18]. We exploit the fact that **A** can have off-diagonal elements by making it act as a perceptually-based filter. Therefore, the resulting images correspond to local minima that are visually more pleasant than those produced using a diagonal **A** matrix.

Starting with the equation for error diffusion, (8), and reversing the above steps leads to (6), the equation for the *M*-lattice system in steady state. Error diffusion has been modeled as a Hopfield network that uses $q(\psi_i)$ in place of $g(\psi_i)$.[16] However, the non-monotonicity of $q(\psi_i)$ causes instability. In contrast, slightly perturbing **A** so as to make it negative-definite guarantees that (1) will be stable for binary outputs. Hence, the *M*-lattice system is a more suitable model for non-causal error diffusion.

For the sake of simplicity we programmed the *M*-lattice system with the symmetric version of the original Floyd & Steinberg error filter. Figure 1 shows the magnified version of a test image and the result of halftoning it by the *M*-lattice system. The new method provides accurate detail rendition without introducing correlated texture. However, some perceptual artifacts still occur, because the filter coefficients have not yet been optimized after the causality constraint was lifted.

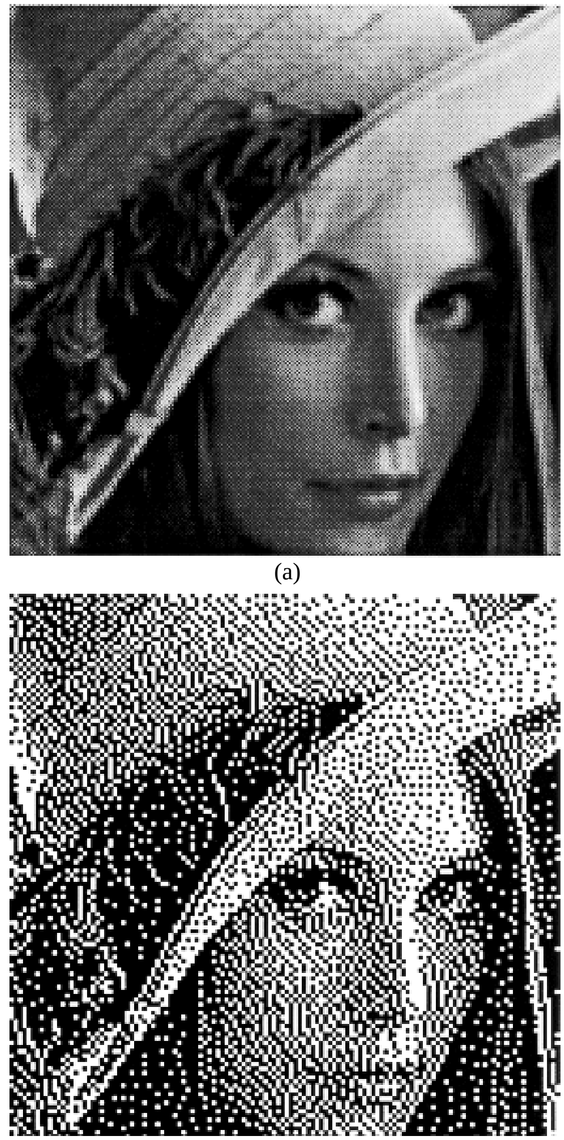

(a)

(b)

Figure 1. Noise-Shaping Least-Squares Halftoning. (a) a portion of the original "Lena" image (magnification is ×2 on a side); (b) the image in (a) halftoned by the M-lattice system.

Halftoning As Non-Linear Program

Suppose $\vec{n} \in Z^2$; $s(\vec{n}) \in [-1,1]$ 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 **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: The problem of halftoning can be stated as a non-linear program:

$$\min_{g} \frac{1}{2} \vec{y}^T \mathbf{B} \vec{y} - (\mathbf{B}\vec{s})^T \vec{y} \tag{9}$$

subject to constraints:

$$y_i^2 - 1 \geq 0, \tag{10}$$

where the vectors are the standard concatenations of the corresponding sequences, $\mathbf{B} = \mathbf{H}^T\mathbf{H}$, and **H** is a circulant matrix with $h(\vec{n})$ in the first row. The particular form of constraints, (10), 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, (9), with the *N* constraints, (10), into the Lagrangian cost functional with the help of the Karush-Kuhn-Tucker conditions[19]:

$$\min_{\vec{y}} \pounds(\vec{y}), \text{ where}$$
$$\pounds(\vec{y}) = \frac{1}{2} \vec{y}^T \mathbf{B} \vec{y} - (\mathbf{B}\vec{s})^T \vec{y} + \sum_i p_i (y_i^2 - 1), \tag{11}$$
$$p_i \leq 0, p_i (y_i^2 - 1) = 0. \tag{12}$$

The Lagrange multipliers, p_i, are the varying penalty terms that enforce the constraints according to (12). As a result, the unconstrained minimization of $\pounds(\vec{y})$ in (11) produces the optimal halftone image.

The optimization problem, (11), is "programmed" onto the *M*-lattice system, (1), by setting $\vec{y}$ equal to $\vec{\chi}, \Phi(\vec{\chi})$ to $-\pounds(\vec{y})$ and taking partial derivatives. This yields:

$$\frac{d\vec{\psi}(t)}{dt} = \mathbf{A}\vec{\psi}(t) + \mathbf{B}\vec{s} - \mathbf{B}\vec{\chi}(t) - \mathbf{P}\vec{\chi}(t), \tag{13}$$

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 **A = B-I** is a good 3-{πce, because it filters out objectionable correlated spatial patterns[8].

Halftoning with the Hopfield network requires setting $b_{ii} \geq 0$; otherwise, the optimal values of y_i will not be binary[20,17]. However, treating halftoning as a non-linear programming problem and solving it with the *M*-lattice system Gffers considerable flexibility in the choice of the quality metric and in the functional form of constraints.

In order to demonstrate this flexibility, we encorporated orientation detection into the halftoning quality metric. The adaptive filter matrix, **H**, was designed so as to include the information about the dominant orientation at each pixel of the original image, shown in Figure 2(a)[12].

Since no effort is made to design **H** in a way that would result in $b_{ii} \geq 0$, the non-linear constraints provide the only mechanism for driving the output pixels to the limits of the gray scale. Figure 2 displays the result, which exhibits more of the line and curve features found in hand-drawn "halftones" (such as the Wall Street Journal portraits).

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[14,7]. 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.

In order to solve this color halftoning problem using the *M*-lattice system, we combine the objective function to be minimized, (9), with the *N* constraints, (10), into the Lagrangian cost functional with the help of the Karush-Kuhn-Tucker conditions[19] simultaneously for all three color components. The Lagrange multipliers, p_i, are the varying penalty terms that enforce the constraints according to (12). As a result, the unconstrained minimization of $\pounds(\bar{y})$ in (11) produces the optimal color halftone image.

While the gray-scale special-effects halftoning algorithm is implemented as a 1-lattice[8], 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[‡].

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, **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 3(a). Figure 3(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)[11]. Another

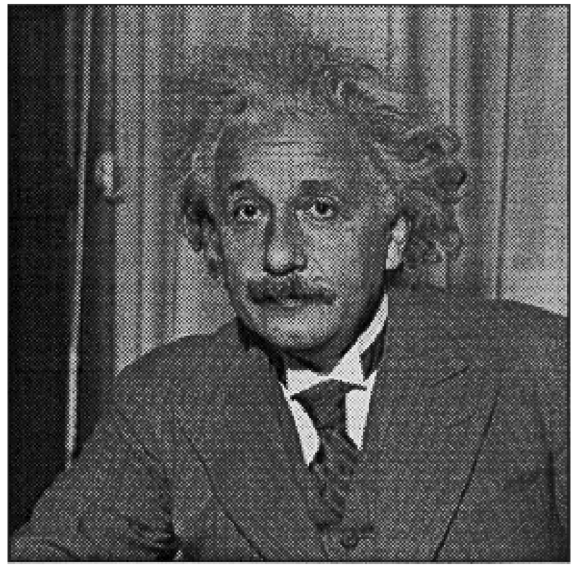

(a)

(b)

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

example of this technique appears in Figure 4.

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[21]. 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.

Summary

We have reviewed the *M*-lattice system and applied it to digital halftoning of images. As a non-linear programming technique, the *M*-lattice system is capable of solving constrained optimization problems with flexible

(a)

(b)

Figure 3. Orientation-sensitive color halftoning. (a) the original "Marty" image; (b) the "Marty" image adaptively halftoned using orientation information at each pixel of the original.

objective functions. Orientation-sensitive halftoning makes use of this property. When the objective function is a quadratic form, the *M*-lattice system can be designed to perform blue noise filtering. This implies that the resulting halftone images can be made not only optimal in the least-squares sense, but also perceptually pleasant.

We have also 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.

† This is the definition adapted for the present paper. The general M-lattice system is defined in [11].

‡ The present study treats the RGB components independently for simplicity.

Acknowledgment

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

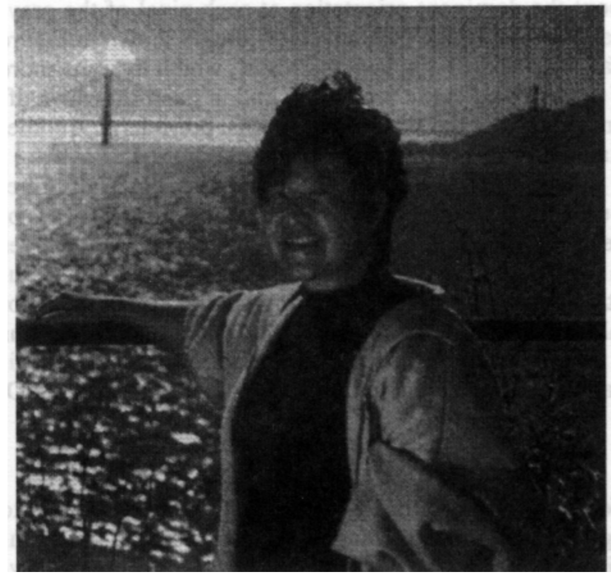

(a)

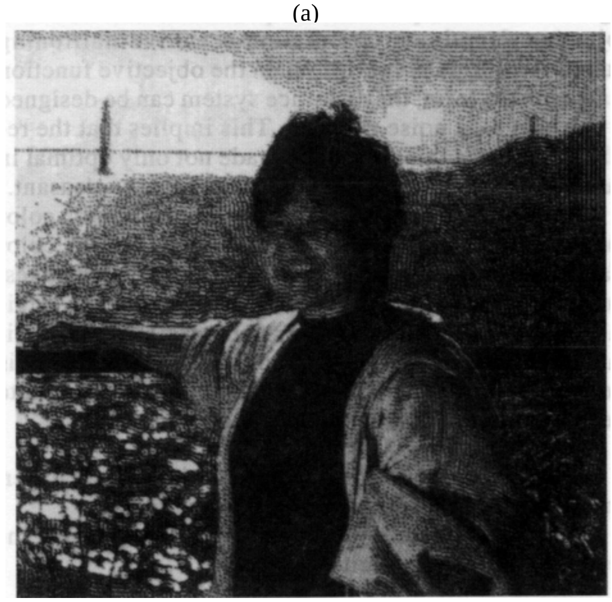

(b)

Figure 4. Orientation-sensitive color halftoning. (a) the original "Betty" image; (b) the "Betty" image adaptively halftoned using orientation information at each pixel of the original.

References

1. A. M. Turing, "The chemical basis of morphogenesis," *Philosophical Transactions of Royal Society of London*, vol. **237**, no. B, pp. 37-72, 1952.
2. A. Witkin and M. Kass, "Reaction-diffusion textures," *Computer Graphics*, vol. **25**, pp. 299-308, July 1991.
3. G. Turk, "Generating textures on arbitrary surfaces using reaction-diffusion," *Computer Graphics*, vol. **25**, pp. 289-298, July 1991.
4. C. B. Price, P. Wambacq, and A. Oosterlinck, "Applications of reaction-diffusion equations to image processing," in *Third International Conference on Image Processing and its Applications*, (Warwick, UK), July 1989.

5. "Private communications with Thrasyvoulos N. Pappas, AT&T," Mar. 1995.
6. 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.
7. 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; ***see Recent Progress in Digital Halftoning, Vol. I, pg. 144.***
8. 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.
9. 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.
10. A. Sherstinsky and R. W. Picard, "Orientationsensitive 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.
11. 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.
12. W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. **PAMI-13**, pp. 891-906, Sept. 1991.
13. R. W. Picard and M. Gorkani, "Finding perceptually dominant orientations in natural textures," *Spatial Vision*, special Juleaz birthday issue, vol. **8**, no. 2, pp. 221-253, 1994. Also avail. as Percep. Comp. TR #229, M.I.T. Media Lab.
14. T. N. Pappas and D. L. Neuhoff, "Least-squares model based halftoning," in *Proc. SPIE Conf. on Electronic Imaging Science and Technology*, (San Jose, CA), Feb. 1992.
15. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grey scale," in *SID Int. Sym. Dig. Tech. Papers*, pp. 36-37, 1975.
16. D. Anastassiou, "Error diffusion coding for A/D conversion," *IEEE Transactions on Circuits and System*s, vol. **36**, pp. 1175-1186, Sept. 1989.
17. A. Sherstinsky and R. W. Picard, "M-Lattice system: Theory and potential applications," Tech. Rep. (In Preparation), M.l.T. Media Lab Perceptual Computing Group, 1994.
18. R. A. Ulichney, *Digital Halftoning*. Cambridge, MA: MIT Press, 1987.
19. M. S. Bazaraa and C. M. Shetty, *Nonlinear Programming*. New York: John Wiley and Sons, 1979.
20. J. J. Hopfield, "Neurons with graded response have collective computational properties like those of twostate neurons," *Proceeding of National Academy of Scienc*e, USA, vol. **81**, pp. 3088-3092, May 1984.
21. "Private communications with Richard Kelman, Dow Jones," May 1994.

❋ Previously published in *IS&T's NIP 11*, 1995, pp. 415–420.