Progressive Halftoning by Perona-Malik Error Diffusion and Stochastic Flipping

Jianhong (Jackie) Shen *Member*

Abstract

Halftoning has been a significant topic in image processing due to numerous emerging applications, highly diversified approaches, and enormous challenges in theoretical analysis. Inspired by the wealthy literature on halftoning, as well as the recent PDE (partial differential equations) approach in image processing, the current work proposes a novel progressive halftoning algorithm. It is based upon the celebrated anisotropic diffusion model of Perona and Malik (*IEEE Trans. Pattern Anal. Machine Intell.*, 12:629-639, 1990), and a properly designed stochastic strategy for binary flipping. The halftone outputs from the proposed model are typical samples of some random fields, which share many virtues of successful deterministic halftone algorithms, as well as reveal many interesting features like the blue noise behavior.

Keywords

Halftone, error diffusion, Perona-Malik, edge adaptive, stochastic flipping, blue noise.

I. Introduction: Overview and Contributions

A. Brief Overview: Halftoning and Many Successful Models and Algorithms

Halftoning (for gray scales) refers to the process of converting a continuous-tone (or contone) gray scale image u to a binary one b, so that bi-level printers or other display devices can successfully output such an image. Thanks to enough resolutions and automatic lowpass filtering by the human vision system (HVS), halftone images can well approximate the original contone ones to human observers. We refer to the informative special issue of the *IEEE Signal Processing Magazine* [1] for a comprehensive coverage of the halftoning technology.

The three key ingredients to a typical halftoning process are:

(i) the contone-halftone ($u \to b$) conversion model or algorithm;

(ii) the printer model, which specifies the physical layouts (i.e., the physical dimensions of ink dots) of printed binary values and their superpositions (e.g., [1], [3], [43]);

(iii) the HVS model, which models how the human vision system perceives the distribution of the printed dots (e.g., [1], [27], [32], [40], [49]).

The three are independent but also closely interlaced, often leading to jointly optimal halftoning processes [3], [27], [32]. Algorithm-wise, many works have been mainly focused on the first ingredient, i.e., the contone-halftone conversion algorithm. So will be the primary focus of the current work. Such model-independent study has also been inspired by the profound implications of halftoning in the broader information technology, for instance,

The work has been partially supported by NSF Grant DMS-0202565. **The author is with the University of Minnesota, School of Mathematics, Minneapolis, MN 55455, USA. Tel. (612) 625-3570; Fax. (612) 626-2017; Email: jhshen@math.umn.edu.**

its intrinsic connections to the sigma-delta model for A/D conversions [12], [13], [22], [23], and to information hiding and digital watermarking (e.g., [2], [19]).

In what follows, it shall be assumed that the contone image u has been normalized into the canonical range of $[0,1]$. For convenience, the symbols $\alpha, \beta, \cdots$ shall also be employed to denote general pixels which are commonly expressed by (i,j)'s with $i = 1:n$ and $j = 1:m$ for an image of n by m. Thus for instance, u_α denotes the gray value at a pixel α. Let Ω denote the collection of all pixels.

The core tool in most classical halftoning algorithms is the thresholding operator T:

$$T(a \mid \mu) = \begin{cases} 1 & \text{if } a \geq \mu \\ 0 & \text{if } a < \mu \end{cases},$$

for any $a \in [0,1]$, and a given threshold $\mu \in (0,1)$. For example, for $\mu = 1/2$, $b_\alpha = T(u_\alpha \mid \mu)$ at any pixel $\alpha \in \Omega$ yields a naive halftoning algorithm which is however not quite useful due to the artificial edges and blocks it generates.

There are three major classes of successful halftoning algorithms in the literature.

[a] **Dot Dithering (or Screening).** The blocky effect in the above naive algorithm can be fixed by adopting spatially varying thresholds $\mu = \mu_\alpha$ for $\alpha \in \Omega$ [4], [29], [42]. The threshold field μ is typically made periodic by a suitable tiling of Ω using non-overlapping blocks of a fixed size $I \times J$ (e.g., $I = J = 8$). Then μ is completely determined by its IJ values and their spatial distribution on a template tile.

[b] **Error Diffusion.** In the celebrated error diffusion model of Floyd and Steinberg [18], four important pieces of data determine the algorithm (see Fig. 1).

[b.1] The halftoning path. The path assigns an order $\alpha_1, \cdots, \alpha_{nm}$ to all the pixels, following which the halftoning conversion will be performed from one pixel to the next. The most common choice is the zigzag raster order.

[b.2] The threshold μ. Currently visiting a pixel α_k, one defines the halftone value by $b_{\alpha_k} = T(u^k_{\alpha_k} \mid \mu)$. Here $u^k_{\alpha_k}$ is not the original contone value u_{α_k}, but has been iteratively updated as the path proceeds.

[b.3] The error propagation (or diffusion) front $F(\alpha_k)$. For each pixel α_k, the front $F(\alpha_k)$ is the collection of *unvisited* pixels nearby, into which the current halftone error $e_{\alpha_k} = u^k_{\alpha_k} - b_{\alpha_k}$ is to be diffused. For example, the following shift-invariant front is a familiar choice:

$$F(\alpha) = \{(i, j+1),\ (i+1, j+1),\ (i+1, j),\ (i+1, j-1)\},$$

provided that $\alpha = (i,j)$, and the path is in the raster order from $(i-1,j)$ to (i,j).

[b.4] The diffusion weights $w_{\alpha_k,\beta}$. These are a set of positive weights defined for *each* frontal pixel $\beta \in F(\alpha_k)$, according to which the current halftone error will be diffused and u^k updated to u^{k+1}:

$$u^{k+1}_\beta = u^k_\beta + e_{\alpha_k} \cdot w_{\alpha_k,\beta}.$$

The values at a non-frontal pixel β is not affected, i.e., $w_{\alpha_k,\beta} = 0$.

In summary, Floyd and Steinberg's error-diffusion algorithm relies upon

$$(\text{visiting path},\ \mu,\ F(\alpha)\text{'s},\ w_{\alpha,\beta}\text{'s}).$$

The original choices made by Floyd and Steiberg [18] have witnessed some well known artifacts such as worms and limited cycles [17], [18], [28]. Many remarkable works have since emerged to improve the halftone qualities by inventing and adopting better paths [41], [47], spatially varying thresholds [14], [29], [40], proper frontal masks [29], and adaptive weights [29], [41], [48], [49].

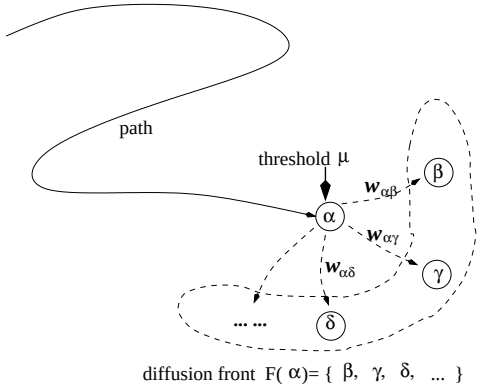

Fig. 1. The four key ingredients of an abstract halftoning model based on error diffusion: the ordered path visiting all the pixels, the thresholds μ's, the error propagation fronts $F(\alpha)$'s, and the weights $w_{\alpha\beta}$'s by which the halftone errors are to be proportionally distributed.

[c] **Dot Diffusion.** Knuth's dot diffusion algorithm [28] adopts both the tiling structure in the dot dithering model and the diffusion mechanism in the error diffusion model. By periodic tiling, error diffusion in the dot diffusion model is confined within each tile. All the IJ pixels in a template tile T of I by J (e.g., $I = J = 8$) are properly ordered or ranked by the class matrix [32] $R = (r_{i,j})_{I\times J} = (r_\alpha)_{\alpha\in T}$ so that $\{r_\alpha : \alpha \in T\} = 1 : (IJ)$, and the pixels are ordered by $\{\alpha_1, \cdots, \alpha_{IJ}\}$ with $r_{\alpha_k} = k$. Within each tile, the halftoning path then follows this specific order. Furthermore, the error at a current pixel α_k only diffuses to those surrounding pixels (out of 8) whose ranks r_β are higher than k, and the weights are assigned according to some fixed rules.

In the work by Mese and Vaidyanathan [32], the class matrix is not chosen *a priori*, but optimally designed subject to proper HVS models. Compared with error diffusion, dot diffusion facilitates parallel implementations and gains both speed and performance.

Besides these models, there have also been in existence numerous other successful halftoning techniques, e.g., via look-up tables, least square minimizations, Markovian random fields, direct binary search (DBS), and fuzzy algorithms [3], [20], [27], [29], [33], [35], [40]. In addition, there are also numerous interesting post-halftoning works including resolution increment [26] and inverse halftoning [1], [32].

Some notable features in all these successful models and algorithms are:

(a) spatial tiling, as in the old JPEG protocol;

(b) *a priori* diffusion patterns for the errors;

(c) the algorithms are sequential, or made periodically parallel by tiling;

(d) single round or pass to complete the contone-halftone conversion;
(e) no straightforward continuum description as the grid sizes converge to zero.
In terms of the competing complexities among modeling, analysis, and computation (see, e.g., [10]), these features are unnecessarily advantageous or disadvantageous. The many interesting works mentioned above have been devoted to the better understanding of all these ingredients.

B. Main Contributions of the Current Work

Inspired by all the above remarkable works, in the current paper we propose a new halftoning model and algorithm, which is built upon the celebrated anisotropic diffusion model of Perona and Malik [36] for adaptive image denoising and enhancement.

The new model is progressive and can be symbolically denoted by (see Fig. 2)

$$b^{n+1} = \mathrm{PMSF}(b^n \mid u, \tau), \quad n = 0, 1, \cdots,$$

where u is a given contone image to be halftoned to some binary image b, and τ a natural positive parameter for the Perona-Malik diffusion. For each n, the model consists of two steps - the PM step, followed by the SF step. The PM step is mainly achieved by the *Perona-Malik* type of diffusion, while the SF step is a suitable rule for *stochastic flipping*.

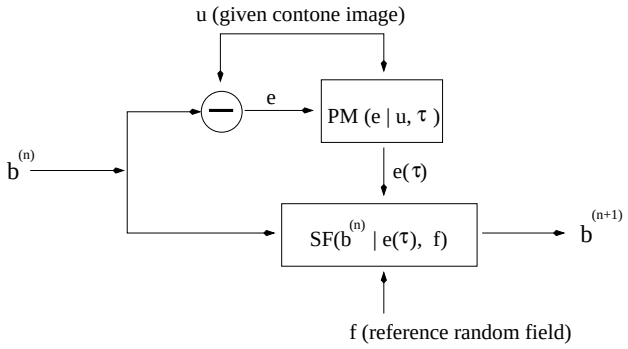

Fig. 2. The flow chart of the proposed new halftoning model.

The main characteristics of the current work can be highlighted as follows.

[1] For the first time, the two major noticeable literatures on diffusions, i.e., error diffusions in halftoning processes and Perona-Malik's celebrated anisotropic diffusions for adaptive image enhancement, are connected and integrated.
[2] Error diffusion via Perona-Malik's model is image adaptive, and does not significantly exchange errors among different objects.
[3] There are no preassigned artificial tiling, error flow patterns, or distribution weights.
[4] The new algorithm improves halftoning qualities *progressively*, instead of only using a single pass, and allows global parallel implementation.
[5] The core technique for making feasible Perona-Malik's adaptive error diffusion is an effective stochastic strategy for on-off flipping.
[6] There are no hard thresholds involved. At each stage it is the diffused error field that offers the likelihood information for progressive flipping.

The next section (Section II) first gives a bird view of the global structures of the new halftoning algorithm that consists of two key steps. Section III explains in details the PM step involving a modified version of Perona-Malik's anisotropic diffusion. Section IV explains the design and analysis of the SF step. Computational examples and their performances are discussed in Section V. The paper ends with a brief conclusion of the current work and a discussion on some future works.

II. Description of the Main Model and Algorithm

In this section, we outline and highlight the main structures of the proposed new halftoning model. More details will be elaborated in the two forthcoming sections.

Let Ω denote the image domain, which could be a continuum like a rectangle, or a discrete lattice. A general point in Ω will be called a pixel, and denoted by α, β, etc. Let u be a given normalized contone image, so that at any pixel α, the image value u_α belongs to $[0, 1]$.

For convenience, for a halftone image b, b_α=1 will be called an "on" pixel and $b_\alpha = 0$ "off." It is also assumed that a printer deposits an ideal black ink dot at a pixel α if and only if $b_\alpha = 0$, i.e., when the pixel is "off."

The proposed new model, referred to as the PM-SF model or algorithm for short, is progressive based upon the iterative procedure described by

$$b^{n+1} = \text{PMSF}(b^n \mid u, \tau), \quad n = 1, 2, \cdots, \tag{1}$$

where,

(i) b^n is the halftone estimation at a stage n,
(ii) u the original contone image to be halftoned, and
(iii) τ a natural parameter to be explained later.

Assuming that the halftone sequence (b^n) converges in some suitable sense, we then define $b^\infty = \lim_{n\to\infty} b^n$ to be the ultimate halftone image of the given contone target u.

At each stage n, the PM-SF model as an operator is constituted of two parts - the PM step and the SF step, i.e., edge-adaptive error diffusion of *Perona and Malik's* type [36], followed by a properly designed decision rule for *stochastic flipping*. The PM step takes the form of

$$e^n(\tau) = \text{PM}(e^n \mid u, \tau), \quad e^n = u - b^n, \tag{2}$$

whose output, the diffused error $e^n(\tau)$, is then fed into the SF step,

$$b^{n+1} = \text{SF}(b^n \mid e^n(\tau), f). \tag{3}$$

In alternating iterations, they accomplish the progressive updating of the halftone image. Here in the SF step, $f = (f_\alpha)$ is random field consisting of independent and identically distributed (i.i.d.) random samples which are uniformly distributed in $[0, 1]$. The PM step is a *deterministic* process, while the SF step is *stochastic* since it reacts to a random reference field f. The *effective* action of the SF step is in fact independent of f *statistically* speaking, as long as f is genuinely random (or *typical* in the information-theoretic sense [11]).

In short, at each stage n, the proposed PM-SF halftone model bears the following form:

$$b^{n+1} = \text{PMSF}(b^n \mid u, \tau) = \text{SF}(b^n \mid \text{PM}(u - b^n \mid u, \tau), f). \tag{4}$$

The next two sections elaborate in full details on these two main steps.

III. The PM Step: Perona-Malik Type of Error Diffusion

A. The Classical Perona-Malik Model

The celebrated anisotropic diffusion model of Perona and Malik [36] was introduced as a new tool for adaptive image denoising and enhancement. The key characteristic compared with traditional *linear* diffusion models or the *linear* scale-space theory [46] is that the diffusive activity is automatically adapted to *edges*. Ever since David Marr's pioneering works on computer vision and artificial intelligence [30], [31], edges have been long recognized as the most important type of visual cues in image perception.

A general Perona-Malik diffusion model bears the following form [36]:

$$\frac{\partial u}{\partial t} = \nabla \cdot [g(|\nabla u|)\nabla u]\,, \tag{5}$$

for $u = u(x,t)$ on a Lipschitz domain $x = (x_1, x_2) \in \Omega$, and often with Neumann's adiabatic boundary condition. Here the diffusivity coefficient $D = g(|\nabla u|)$ depends upon the image u itself instead of being fixed, making the equation nonlinear and challenging for rigorous mathematical investigation [5], [25]. For image denoising and enhancement, the initial condition of (5) is the given noisy or degraded image $u(x,0) = u_0(x)$.

Another important piece of information needed for achieving ideal image enhancement is the *stopping time* T. Since diffusing processes are all anti-gradient random walks in the microscopic scale, without being stopped in the middle, a diffusion process will invariably converge to an uninteresting constant image (with zero gradients everywhere). There are several ways in the literature dealing with the issue of stopping time, including for example, optimal control or via the connection to the noise level in variational models [44], [45].

The desired action of the Perona-Malik diffusion is furnished by the image dependent diffusivity coefficient $D = g(p)$ with $p = |\nabla u|$. Qualitatively, D should behave like an edge signature function so that

$$D = O(1), \quad \text{away from edges where } p \simeq 0; \quad \text{and} \quad \simeq 0, \quad \text{near edges where } p \simeq \infty. \tag{6}$$

In this manner, the "flux" flow $\boldsymbol{j} = -D\nabla u$ can be effectively confined within each homogeneous patch, and the edge blurring effect is controlled and diminished. Following Perona and Malik [36], some plausible candidates of g for achieving such desirable effects include, the Gaussian

$$g(p) = A\exp(-bp^2), \quad \text{for some two positive constants } A \text{ and } b, \text{ and}$$

the Cauchy function

$$g(p) = \frac{A}{1+bp^2}, \quad \text{for two positive constants } A \text{ and } b.$$

Another useful candidate is motivated by Rudin-Osher-Fatemi's total variation Radon measure [37] as well as the mean curvature motion [16]:

$$g(p) = \frac{1}{p}, \text{ or its regularized version, } \frac{1}{\sqrt{p^2+b^2}},$$

for some positive weight b.

We refer to [5], [25] for further mathematical explorations into the wellposedness of the Perona-Malik diffusion and its regularized variants, as well as to the extensive contributions made by Weickert and his colleagues for efficient numerical computations [44], [45]. Other recent novel applications and extensions of Perona-Malik's nonlinear diffusion mechanism could be found, e.g., [6], [7], [8], [9], [24], [38], [39].

B. The PM-Step of Error Diffusion

In the proposed new halftoning model, the PM step has been closely inspired by the Perona-Malik diffusion briefly covered above. Necessary modification also has to be made to suit the specific application of error diffusion in halftoning.

At step n, if the halftone image b^n already well approximates the given contone image u, then the error field $e^n = u - b^n$ shall more or less look like homogenous oscillations. As a result, the edge information in the original contone image u cannot be clearly read in e^n, and it becomes less appropriate to directly apply the Perona-Malik diffusion for e^n:

$$\frac{\partial e}{\partial t}(x,t) = \nabla \cdot [g(|\nabla e|)\nabla e(x,t)]\,, \quad \text{with the initial condition} \quad e(x,0) = e^n(x).$$

A natural way to dissolve the concern is to turn to the original contone image u, and solve the following *linear* diffusion equation instead:

$$\frac{\partial e}{\partial t}(x,t) = \nabla \cdot [g(|\nabla u|)\nabla e(x,t)]\,, \quad \text{with} \quad e(x,0) = e^n(x) = u(x) - b^n(x),\ x \in \Omega. \tag{7}$$

The Neumann adiabatic boundary condition will be similarly imposed along $\partial\Omega$. Thus to $e(x,t)$, the diffusivity coefficient D is actually provided by the given contone image $u(x)$. For any given stopping time $T = \tau$, we shall then naturally define the diffused error at step n by:

$$e^n(\tau) = \mathrm{PM}(e^n \mid u, \tau) = e(x,\tau), \tag{8}$$

where $e(x,\tau)$ is the solution to (7) at stopping time τ. Recall that in the discrete setting a pixel x is also denoted by α, and $e(x,\tau)$ by $e^n_\alpha(\tau)$.

The key ideas behind this novel design of error diffusion can be highlighted as follows:

[a] Suppose u is a constant contone image, and b its *ideal* halftone version. Then the error field $e(x) = u(x) - b(x)$ shall be maximally homogeneous in space. Consequently, diffusion as a lowpass filtering process shall be able to wipe out the error eventually with sufficiently large τ. For a general contone image, therefore, the *magnitude* of $e(x,\tau)$ or $e^n_\alpha(\tau)$ signifies the degree of *imperfectness* of the current halftone field b^n_α at α.

[b] Noticeable features remaining in the diffused error field $e^n(\tau)$ are often caused by blocky or chunky halftoning, and can thus be employed to improve the halftone quality in the next SF step.

[c] Both conceptually and perceptually in terms of image and vision analysis, it is desirable to confine error diffusion within each homogeneous object in the original image u. This is furnished by the Perona-Malik diffusivity $D = g(|\nabla u|)$ adapted to the given contone image u.

For cartoonish or piecewise constant images, we can further quantify the statement [c] on error confinement.

Theorem 1 *Suppose u is piecewise constant with each piece occupying a Lipschitz domain. Let B denote any of these* generic *patches (assumed to be an open domain). Then under the PM error diffusion (7) and the assumptions that $g(0) = A > 0$ and $g(\infty) = 0$, the net total error on B: $E(t) = \int_B e(x,t)dx$ does not change with time at each fixed step n.*

Proof: The rigorous mathematical statement of the theorem and its proof will inevitably depend upon a mollification process on u (so that the classical derivative ∇u can be well defined, as in [5]). Here we give a heuristic but more straightforward proof.

Since the image u is piecewise constant and B is a *generic* open patch, the size of the gradient $p = |\nabla u|$ is exactly zero on B, and infinity in the interior boundary of B (i.e., $\partial B \cap \Omega$). Thus in terms of the diffusivity coefficient, one has

$$D = g(p) \equiv A, \quad \text{on B}; \quad \text{while} \quad \equiv 0, \quad \text{along } \partial B \cap \Omega. \tag{9}$$

Therefore, by the divergence theorem on the Lipschitz patch B,

$$\begin{aligned} \frac{dE(t)}{dt} &= \frac{d}{dt}\int_B e(x,t)dx = \int_B e_t(x,t)dx \\ &= \int_B \nabla \cdot [D\nabla e]dx = \int_{\partial B} D\frac{\partial e}{\partial \vec{n}}ds \\ &= \int_{\partial B\cap\Omega} D\frac{\partial e}{\partial \vec{n}}ds + \int_{\partial B\cap\partial\Omega} D\frac{\partial e}{\partial \vec{n}}ds = 0, \end{aligned} \tag{10}$$

where $\vec{n}$ denotes the outer normal of ∂B and ds its arc length. The last equality follows from (9) and the Neumann condition $\partial e/\partial \vec{n} = 0$ along $\partial\Omega$. This completes the proof. ■

As emphasized in the proof, a more rigorous argument should rely upon suitable mollification techniques, and the above proof must be considered heuristic since the transition from $p = 0$ to $p = \infty$ is in terms of Dirac's delta function or the total variation Radon measure [15], [21]. The above proof does have captured sufficiently well the essence of error confinement at each individual step.

Next we discuss how to quantitatively employ the information hidden in $e^n(\tau)$ to polish the halftone quality (in the SF step).

IV. The SF Step: Stochastic Flipping of On-Off's

A. The Intuition Behind the SF-Step

We first explain qualitatively the ideas that have motivated the SF step.

At each stage n, after the PM diffusion, the halftone error $e^n = u - b^n$ is mollified to $e^n(\tau)$. The SF step fully exploits the information contents in $e^n(\tau)$.

Suppose at a given pixel α, $e^n_\alpha(\tau)$ is positive and noticeably large. It means that *on average* near α, there are not enough "on" pixels (or too many black dots). One should then employ this information to turn on more pixels near α.

Similarly, if at a given pixel α, $e^n_\alpha(\tau)$ is negative with noticeably large magnitude, one should then carry out the opposite, i.e., turning *off* more pixels near α since negativity implies that *on average* there are too many "on" pixels.

Accordingly, it seems natural to adopt the following decision rule. First copy b^{n+1} from b^n, and then,

$$\begin{aligned} &\text{if } e^n_\alpha(\tau) > 0, \text{ turn on } b^{n+1}_\alpha \text{ if } b^n_\alpha \text{ was off; else} \\ &\text{if } e^n_\alpha(\tau) < 0, \text{ turn off } b^{n+1}_\alpha \text{ if } b^n_\alpha \text{ was on.} \end{aligned} \tag{11}$$

This updating is to be applied to each pixel $\alpha \in \Omega$.

A closer check shows a problem with such a *deterministic* decision rule. Consider for example a constant contone image $u_\alpha \equiv 0.5$. Suppose one starts with an initial halftone guess $b^0_\alpha \equiv 0$, i.e., with all pixels off. Then

$$e^0_\alpha = u_\alpha - b^0_\alpha \equiv 0.5, \quad \text{and} \quad e^0_\alpha(\tau) \equiv 0.5 > 0.$$

Thus by the updating rule (11), all the pixels shall be turned on and the new halftone image becomes $b^1_\alpha \equiv 1$. Repeating the processing, one ends up with a sequence of halftone approximations $(b^n)_{n=0}^\infty$ which alternate between purely off-tone images and purely on-tone images. In particular, they would never converge to any reasonable halftone estimation b^∞.

This motivates the following stochastic flipping (SF) rule, which is the second important ingredient of the proposed new halftone model.

B. The Algorithm of Stochastic Flipping (SF)

Instead of the deterministic binary decision rule in (11), the SF algorithm adopts the stochastic (or fuzzy) approach, and strictly speaking outputs a random field instead of a deterministic single halftone image.

As before, at step n, first copy b^{n+1} from b^n. Then if $e^n_\alpha(\tau) \geq 0$ at a pixel α (i.e., *undershooting*), one applies the following stochastic rule:

$$\text{with probability } e^n_\alpha(\tau), \text{ turn on } b^n_\alpha \text{ to } b^{n+1}_\alpha = 1 \text{ if it has been off;} \tag{12}$$

otherwise when $e^n_\alpha(\tau) < 0$ (i.e., *overshooting*), one then instead applies:

$$\text{with probability } 1 + e^n_\alpha(\tau), \text{ turn off } b^n_\alpha \text{ to } b^{n+1}_\alpha = 0 \text{ if it has been on.} \tag{13}$$

Since the decision is pixelwise decoupled, this stochastic rule clearly admits parallel implementation and computation.

To effectively carry out the flipping task, a convenient way is to first generate a *reference* random field f on the image domain Ω, so that

$$(f_\alpha \mid \alpha \in \Omega) \text{ is a field of i.i.d. random variables uniformly distributed in } [0,1]. \tag{14}$$

Then algorithmically, the SF step described by (12) and (13) is to be realized by

$$b^{n+1} = \mathrm{SF}(b^n \mid e^n(\tau), f), \qquad \text{as follows.}$$

First copy b^{n+1} from b^n. Then (12) is *numerically* computed by: if $e^n_\alpha(\tau) \geq 0$,

$$\text{turn on } b^n_\alpha \text{ to } b^{n+1}_\alpha = 1 \text{ if } f_\alpha \leq e^n_\alpha(\tau) \text{ (no action otherwise);} \tag{15}$$

and similarly for (13), if $e^n_\alpha(\tau) < 0$,

$$\text{turn off } b^n_\alpha \text{ to } b^{n+1}_\alpha = 0 \text{ if } f_\alpha \geq 1 + e^n_\alpha(\tau) \text{ (no action otherwise).} \tag{16}$$

Notice that if α is already on in the former case, or off in the latter, the above action does not change the current on-off status.

Coupled with the PM step in the preceding section, this completes the description of the entire new halftoning model. Next we develop some theoretical analysis for the SF step.

Theorem 2 (Equivalence) *Indeed (15) is equivalent to (12), and (16) to* (13).

Proof: When $e^n_\alpha(\tau) \geq 0$, since f_α is uniformly distributed in $[0, 1]$, (15) implies that

$$\begin{aligned} &\text{Prob}(\ b^{n+1}_\alpha = 1 \mid b^n_\alpha = 0 \text{ and } e^n_\alpha(\tau) \geq 0) \\ =&\text{Prob}(\ f_\alpha \leq e^n_\alpha(\tau) \mid b^n_\alpha = 0 \text{ and } e^n_\alpha(\tau) \geq 0) \\ =&e^n_\alpha(\tau), \end{aligned} \tag{17}$$

which is precisely (12). The equivalence of the other pair can be similarly established. Here we have used the notation of conditional probability. ■

To help understand the convergence of the SF step, define at each step n the *updating profile* d to be

$$d^n_\alpha = |b^{n+1}_\alpha - b^n_\alpha|, \quad \alpha \in \Omega. \tag{18}$$

Then d^n is binary and $d^n_\alpha = 1$ if and only if the status of α is updated to the opposite. Thus we further define

$$R^n = \frac{1}{|\Omega|} \sum_{\alpha \in \Omega} d^n_\alpha \qquad \text{to be the flipping rate per pixel (frpp) at step } n, \tag{19}$$

where $|\Omega|$ denotes the total number of pixels in Ω. For example, $R^n = 5\%$ would mean that at each pixel α, the likelihood that b^{n+1}_α differs from b^n_α is less than 5%.

If the algorithm converges in certain sense, one should be able to expect very small frpp R^n in the end. But unlike any deterministic algorithms, due to the stochastic nature of the SF step, d^n is also a random field and the frpp R^n a random number. We should thus analyze the properties of both d^n and R^n in the probabilistic framework. Let $\mathrm{E}[X]$ denote the expectation or average of a given random variable X.

Theorem 3 *At each step n, given b^n and $e^n(\tau)$, the average frpp is given by*

$$\mathrm{E}[R^n] = \frac{1}{|\Omega|} \sum_{\alpha \in \Omega} \chi\left(b^n_\alpha + \mathrm{sign}(e^n_\alpha(\tau))\right)\ |e^n_\alpha(\tau)|. \tag{20}$$

Here sign *is the standard sign function, and $\chi(w)$ denotes the indicator function of the* closed *unit interval* $[0, 1]$*:*

$$\chi(w) = 1, \quad w \in [0, 1]; \qquad 0, \quad \textit{otherwise.}$$

Proof: By definition, it suffices to show that at each pixel α,

$$\mathrm{E}[d^n_\alpha] = \chi(b^n_\alpha + \mathrm{sign}(e^n_\alpha(\tau)))\ |e^n_\alpha(\tau)|. \tag{21}$$

We then consider separately two cases. First suppose $b^n_\alpha = 0$. When $e^n_\alpha(\tau) < 0$, no flip can occur and $b^{n+1}_\alpha = 0$, which implies that $d^n_\alpha \equiv 0$ regardless of the random reference value f_α and $\mathrm{E}[d^n_\alpha] = 0$. When $e^n_\alpha(\tau) \geq 0$, one has

$$\begin{aligned} \mathrm{E}[d^n_\alpha] &= \mathrm{Prob}(d^n_\alpha = 1) \\ &= \mathrm{Prob}(b^{n+1}_\alpha = 1 \mid b^n_\alpha = 0, e^n_\alpha(\tau) \geq 0) \\ &= e^n_\alpha(\tau), \qquad \text{by (17)}. \end{aligned} \tag{22}$$

Secondly, suppose $b^n_\alpha = 1$. Then no flip can occur when $e^n_\alpha(\tau) \geq 0$, implying that $d^n_\alpha \equiv 0$ and $\mathrm{E}[d^n_\alpha] = 0$. When $e^n_\alpha(\tau) < 0$, one then similarly has $\mathrm{E}[d^n_\alpha] = -e^n_\alpha(\tau)$. In combination, we have established (21), and thus the theorem as well. ■

In particular, it is useful to have the following upper bound solely based on the diffused error from the PM step.

Theorem 4 *Given $e^n(\tau)$ after the PM diffusion at step n, the upper bound for the average frpp is given by*

$$\mathrm{E}[R^n] \leq \frac{1}{|\Omega|} \sum_{\alpha \in \Omega} |e^n_\alpha(\tau)| = \frac{\|e^n_\cdot(\tau)\|_{l^1}}{|\Omega|}, \tag{23}$$

where the vectorial l^1-norm has been applied.

In the next section, we shall apply the results of both Theorems 3 and 4 to partially characterize the computational performance of the PM-SF halftoning algorithm. For instance, an average frpp with $E[R^n] \leq 2\%$ would imply that on average there are only less than 2% of all pixels whose halftone values in b^n are flipped to their opposites in b^{n+1}. Thus acceptable convergence has been reached both statistically and practically.

V. Computational Examples and Discussions

In this section, we demonstrate the numerical performance of the proposed PM-SF halftoning algorithm through some generic examples. Since there have been no universal objective criteria for fairly rating or comparing various halftoning models, we shall mainly discuss the blue noise characteristics of the current model in the end [41], [42].

Fig. 3 shows the halftone ramp image by the PM-SF model (the bottom panel), and its comparison with the output from Floyd-Steinberg's (F.-S.) original error diffusion model [18] (the top panel). The PM-SF model avoids all the structured artifacts (often called "worms" or "tearing") that emerge from the F.-S. model, and clearly shows the characteristics of an *isotropic* random field. (Many remarkable works mentioned in the Introduction have also improved the original F.-S. model by introducing more adaptivity, as explained earlier.)

Fig. 4 shows the halftone images of a stair ramp defined by:

$$u(x,y) = u(x) = \begin{cases} 1 - x/2 & x \in [0, 1/2) \\ 1/2 - x/2 & x \in [1/2, 1], \end{cases}$$

with a stair jump of 1/2 at $x = 1/2$. Due to the *localized* diffusion mask in the F.-S. model (see Fig. 1), the perceived edge sharpness is automatically preserved in the halftone image (the top panel). On the other hand, although the Perona-Malik diffusion has been run

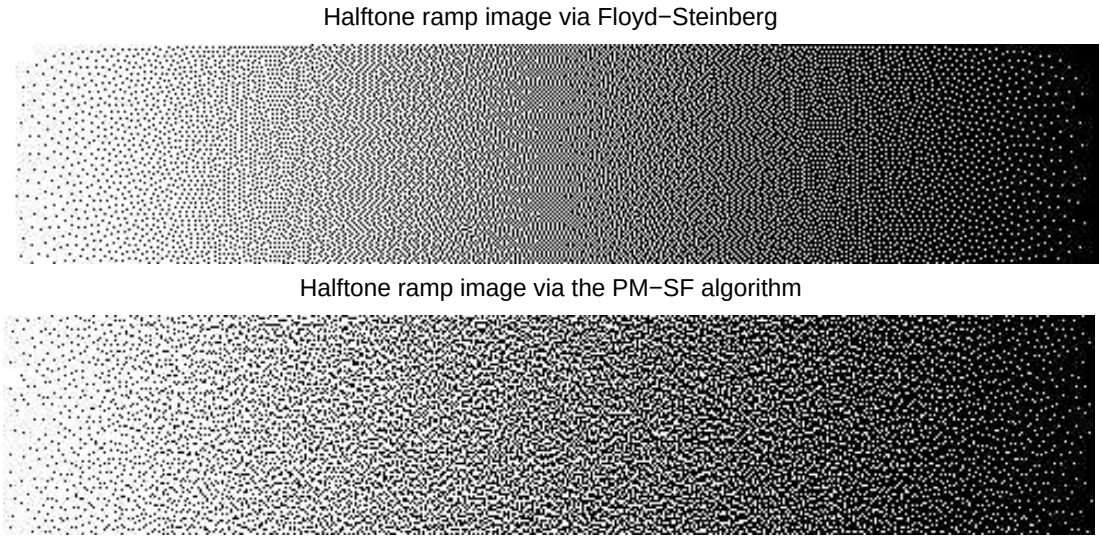

Fig. 3. The halftoned ramp image by the PM-SF algorithm. (Note: due to PDF compressions as well as screen or printer distortions, the displayed or printed halftone images in this section may not be the original MATLAB outputs. Proper zooming rates (like 100%, 150%, or 200% for Acrobat Reader 5.0 or 7.0) should be accordingly adjusted.)

globally on the entire image domain, due to its built-in edge adaptivity, the halftone image also remarkably preserves edge sharpness (the bottom panel).

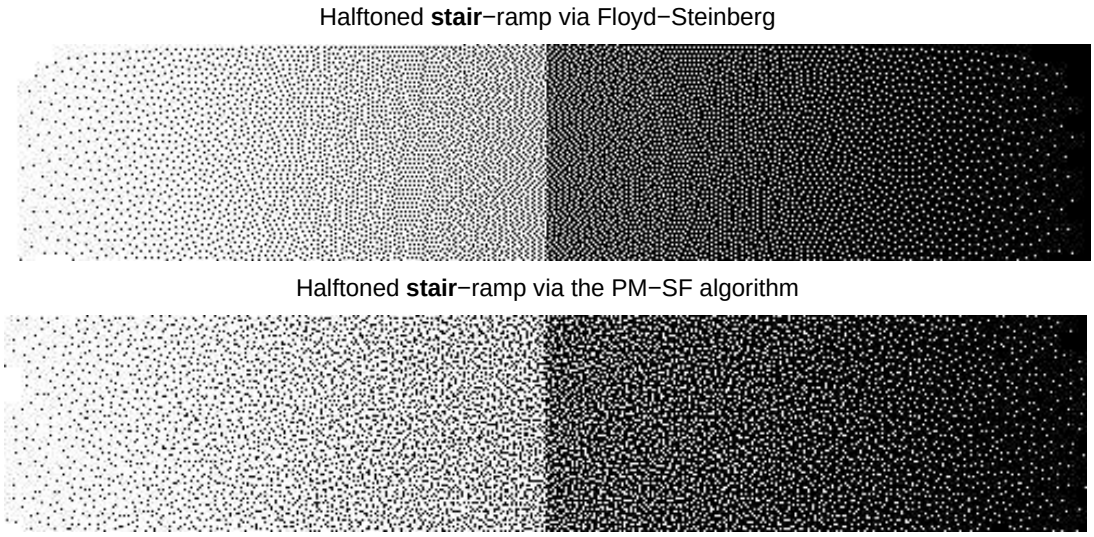

Fig. 4. The halftoned images of a stair ramp via the F.-S. and PM-SF models.

In Fig. 5 and 6, we have shown the halftone images for two popular test images by the PM-SF algorithm. In Fig. 6, in addition, we have also plotted the average frpp (*flip rate per pixel*, defined in the preceding section) versus the progressive step number. It is evident that statistical convergence has been reached when the average frpp drops down to about only 2% and stabilizes.

Finally, we investigate the blue-noise characteristics of the halftone images by the proposed PM-SF algorithm. Blue noises generally refer to signals or images whose power spec-

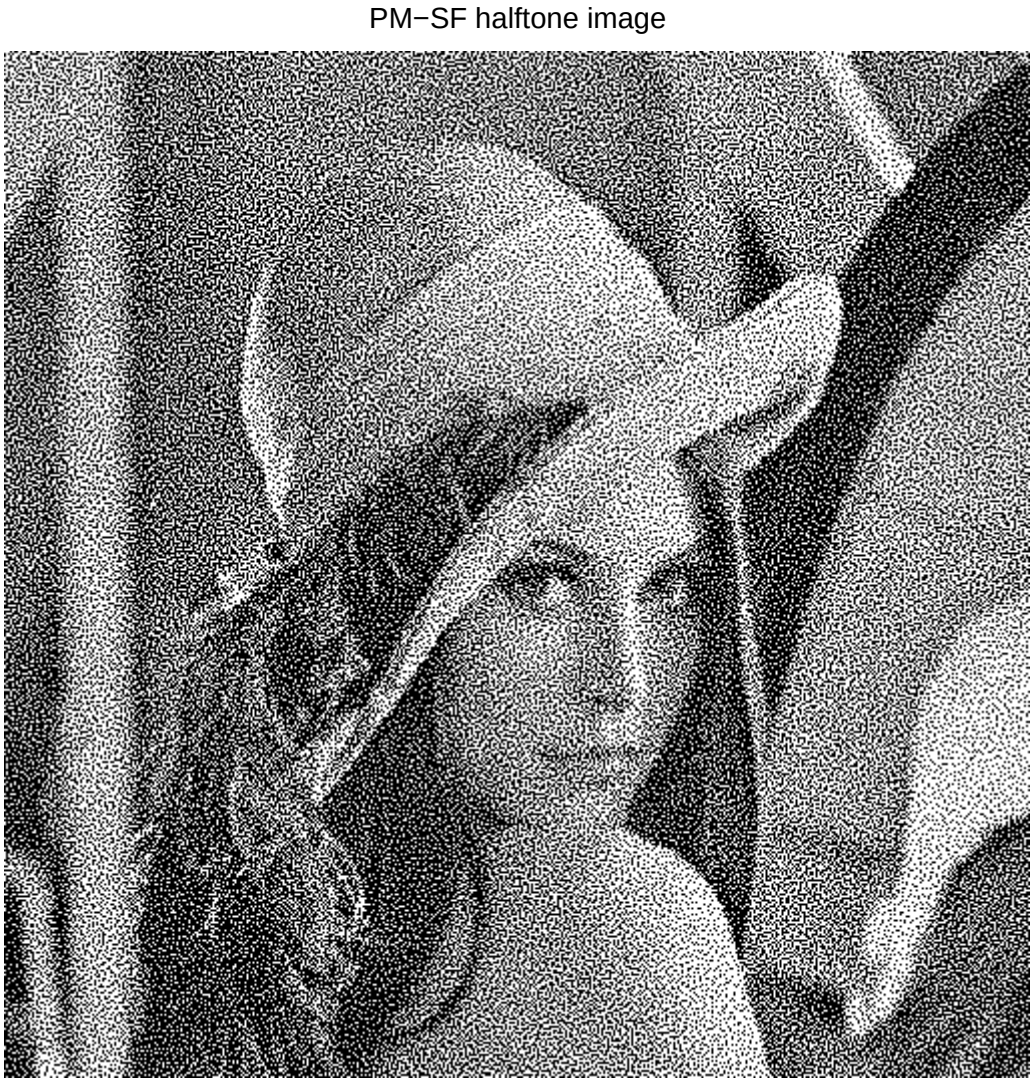

Fig. 5. Halftoned Lenna by the PM-SF algorithm.

tra are only concentrated on the high-frequency end. For instance, a 1-D discrete signal is a blue noise when its power spectrum is given by $P(f) = 1_{[f_p,1]}(|f|)$ for some $f_p > 0$, where f is the normalized frequency variable so that $f \in [-1, 1]$. It has been long acknowledged that good halftoning algorithms should be able to generate blue-noise binary images [41], [42].

In Fig. 7 and 8, we compare the power spectra of the halftoned constant images $u \equiv c$ for both the F.-S. error diffusion model and the proposed PM-SF model. The power spectra have been estimated by the FFT and moving averages (see, e.g., [34]). When $c = 0.35$ (Fig. 7), the power spectra of both models demonstrate the blue-noise behavior. In addition, the isolated bright lines in the left panel shows that the F.-S. halftoning model introduces directional structures, while the rotational invariance in the right panel shows that the PM-SF halftoning model does not. When $c = 0.5$ (Fig. 8), it is well known that the F.-S. algorithm produces the regular check-board pattern, so that $b(i, j) = \cos(i\pi)\cos(j\pi)$, and the power spectrum is concentrated on the single component of the highest frequency (see the arrow in the left panel). In the right panel, the power spectrum of the PM-SF halftone

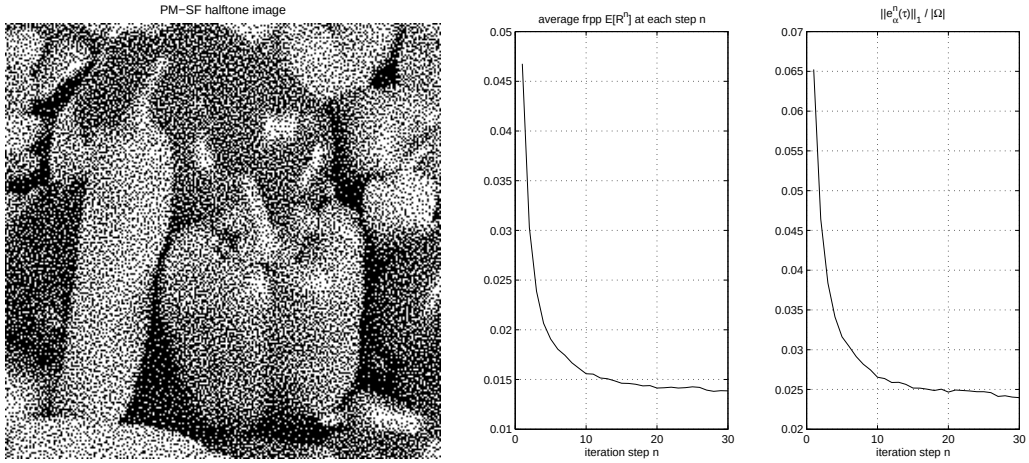

Fig. 6. Left panel: the halftone pepper image from the PM-SF algorithm. Right panel: the average frpp $\mathrm{E}[R^n]$ and its upper bound $\|e^n_\bullet(\tau)\|_1/|\Omega|$ as halftoning progresses (see Theorems 3 and 4). Notice that $\mathrm{E}[R^n] = 1\%$ means that on average there is only one pixel to be flipped in every 10 by 10 patch at step n. Thus the two curves clearly demonstrate the stability as well as the stochastic convergence of the PM-SF algorithm.

image does not differ much from the preceding case of $c = 0.35$, and shows its universality in halftoning different continuous tones.

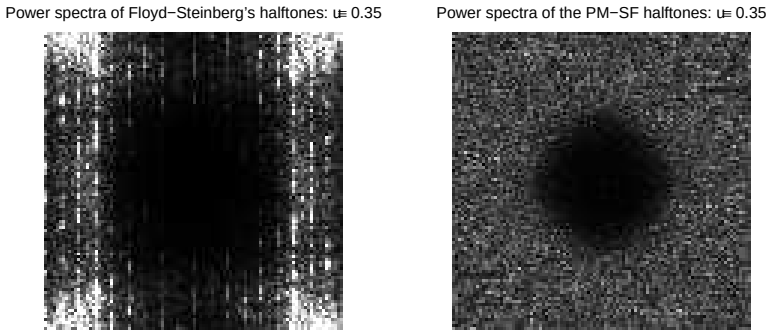

Fig. 7. Comparison of the blue-noise characteristic (I): $u \equiv 0.35$.

VI. Conclusion and Discussions

In this paper, we have introduced a novel progressive halftoning algorithm based on the celebrated edge-adaptive diffusion model of Perona and Malik [36] (PM), and a properly designed stochastic flipping (SF) strategy inspired by the probabilistic interpretation of the diffused errors. Motivations and implementations of these two key models have been explained in details. Generic computational examples further highlight the main characteristics of the proposed PM-SF halftoning model: anti-worming and isotropicity on regions with smoothly varying continuous tones, no significant error "leakage" across edges, and outputing a binary random field instead of a single deterministic one.

Our new model has enriched the contemporary literature on halftoning. Together with numerous other successful halftoning models, the current work contributes to the further advancement of the exciting halftoning technology.

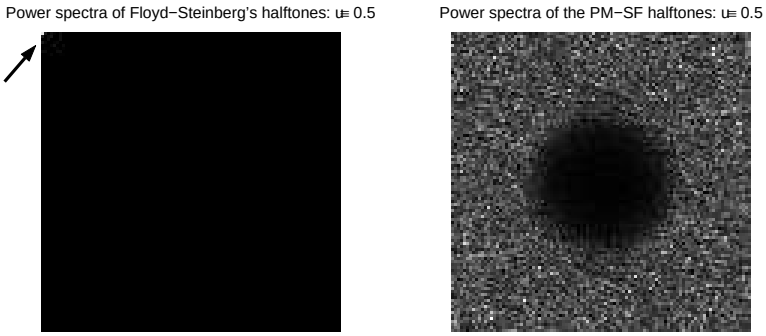

Fig. 8. Comparison of the blue-noise characteristic (II): $u \equiv 0.5$.

Our future works will be focused on (1) combining the PM-SF halftoning model with given printer models and/or HVS models, (2) generalizing the PM-SF model for color contone images, and (3) further characterizing the mathematical properties of the PM-SF model.

Acknowledgments

The author thanks Professors Gil Strang and Truong Nguyen for their constant inspirations on digital signal and image processing. The author also wishes to thank Professor Ingrid Daubechies for persistently advocating the importance of studying and refining interesting engineering problems in contemporary applied mathematics, as well as for timely introducing to me a good mathematical conference for presenting the current work. The author is particularly grateful for the inspiring teaching by and discussion with my two friends: Dr. Chai Wah Wu from the IBM Research, and Professor Sinan Güntürk from the Courant Institute of Mathematical Sciences. Finally, the author, as an applied mathematician working on image processing and computer vision, would feel terribly indebted if without acknowledging the numerous secret inspirations from reading all the remarkable works on halftoning mentioned throughout the work.

Jianhong (Jackie) Shen received the Ph.D degree in Applied Mathematics from the Massachusetts Institute of Technology in 1998, and the B.S. degree from the University of Science and Technology of China (USTC) in 1994. He was a CAM (Computational and Applied Mathematics) Assistant Professor at UCLA from 1998 to 2000. He is currently an Assistant Professor of Applied Mathematics in the University of Minnesota, MN, USA. His current research interests include image, signal, and information processing, vision modeling and computation, as well as multiscale and stochastic modeling in medical and biological sciences. His new book: *Image Processing and Analysis - variational, PDE, wavelets, and stochastic methods*, coauthored with Prof. Tony F. Chan (Dean of Physical Sciences, UCLA), will be published by the SIAM (*Soc. Ind. Appl. Math.*) Publisher in July 2005. Most of his research and teaching activities could be found at www.math.umn.edu/~jhshen.

References

[1] J. P. Allebach and T. Pappas. Digital halftoning: Advances in algorithms and in printing. *IEEE Signal Processing Magazine (Special Issue)*, 20(4), 2003.

[2] Z. Baharav and D. Shaked. Watermarking of dither halftoned images. *Proc. SPIE*, pages 307–313, 1999.

[3] F. A. Baqai and J. P. Allebach. Halftoning via direct binary search using analytical and stochastic printer models. *IEEE Trans. Image Processing*, 12(1):1–15, 2003.

[4] B. E. Bayer. An optimum method for two level rendition of continuous-tone pictures. *Proc. IEEE Int'l Conf. Communication*, 26:11–15, 1973.

[5] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, 29:182–193, 1992.

[6] T. F. Chan, S. Osher, and J. Shen. The digital TV filter and nonlinear denoising. *IEEE Trans. Image Process.*, 10(2):231–241, 2001.

[7] T. F. Chan and J. Shen. Variational restoration of non-flat image features: models and algorithms. *SIAM J. Appl. Math.*, 61(4):1338–1361, 2000.

[8] T. F. Chan and J. Shen. Nontexture inpainting by curvature driven diffusions (CDD). *J. Visual Comm. Image Rep.*, 12(4):436–449, 2001.

[9] T. F. Chan and J. Shen. *Image Analysis and Processing: variational, PDE, wavelets, and stochastic methods*. SIAM Publisher, Philadelphia, 2005.

[10] T.-C. Chang and J. P. Allebach. Memory efficient error diffusion. *IEEE Trans. Image Processing*, 12(11):1352–1366, 2003.

[11] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., New York, 1991.

[12] I. Daubechies and R. DeVore. Approximating a bandlimited function using very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order. *Annals of Math.*, 158(2):679–710, 2003.

[13] I. Daubechies, R. DeVore, C. S. Güntürk, and V.A. Vaishampayan. Beta expansions: A new approach to digitally correct A/D conversion. *Proc. IEEE Int'l Symp. Circ. Sys. (ISCAS'2002), Scottsdale, Arizona*, pages 26–29, 2002.

[14] R. Eschbach. Error diffusion algorithm with homogeneous response in highlight and shadow areas. *J. Electronic Imaging*, 6:348–356, 1997.

[15] L. C. Evans and R. F. Gariepy. *Measure Theory and Fine Properties of Functions*. CRC Press, Inc., 1992.

[16] L. C. Evans and J. Spruck. Motion of level sets by mean curvature. *J. Diff. Geom.*, 33(3):635–681, 1991.

[17] Z. Fan and R. Eschbach. Limit cycle behavior of error diffusion. *Proc. IEEE Int. Conf. Image Processing*, 2:1041–1045, 1994.

[18] R. Floyd and L. Steinberg. An adaptive algorithm for spatial grey scale. *Proc. Soc. Info. Disp.*, 17(2):75–77, 1976.

[19] M. S. Fu and O. C. Au. Data hiding watermarking for halftone images. *IEEE Trans. Image Processing*, 11(4):477–484, 2002.

[20] R. Geist, R. Reynolds, and D. Suggs. A Markovian framework for digital halftoning. *ACM Trans. Graphics*, 12(2):136–159, 1993.

[21] E. Giusti. *Minimal Surfaces and Functions of Bounded Variation*. Birkhäuser, Boston, 1984.

[22] C. S. Güntürk. *Harmonic Analysis of Two Problems in Signal Quantization and Compression*. PhD thesis, Princeton University, 2000.

[23] C. S. Güntürk and N. T. Thao. Refined error analysis in second-order Sigma-Delta modulation with constant inputs. *IEEE Trans. Inform. Theory*, 50(5):839–860, 2004.

[24] S.-H. Kang and J. Shen. Video dejittering by bake and shake. *UCLA Math. Tech. Report 04-60*, 2004.

[25] S. Kichenassamy. The Perona-Malik paradox. *SIAM J. Appl. Math.*, 57(5):1328–1342, 1997.

[26] H. Y. Kim. Binary halftone image resolution increasing by decision tree learning. *IEEE Trans. Image Processing*, 13(8):1136–1146, 2004.

[27] S. H. Kim and J. P. Allebach. Impact of HVS models on model-based halftoning. *IEEE Trans. Image Processing*, 11(3):258–269, 2002.

[28] D. E. Knuth. Digital halftones by dot diffusion. *ACM Trans. Graph.*, 6:245–273, 1987.

[29] P. Li and J. P. Allebach. Tone-dependent error diffusion. *IEEE Trans. Image Processing*, 13(2):201–215, 2004.

[30] D. Marr. *Vision*. Freeman, San Francisco, 1980.

[31] D. Marr and E. Hildreth. Theory of edge detection. *Proc. Royal Soc. London*, B:207: 187–217, 1980.

[32] M. Mese and P. P. Vaidyanathan. Optimized halftoning using dot diffusion and methods for inverse halftoning. *IEEE Trans. Image Processing*, 9(4):691–709, 2000.

[33] D. L. Neuhoff, T. N. Pappas, and N. Seshadri. One-dimensional least-squares model-based halftoning. *Proc. ICASSP-92, San Francisco, CA*, pages 189–192, 1992.

[34] A. V. Oppenheim and R. W. Schafer. *Discrete-Time Signal Processing*. Prentice Hall Inc., New Jersey, 1989.

[35] D. Ozdemir and L. Akarun. Fuzzy algorithms for combined quantization and dithering. *IEEE Trans. Image Processing*, 10(6):923–931, 2001.

[36] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.*, 12:629–639, 1990.

[37] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[38] J. Shen. On the foundations of vision modeling I. Weber's law and Weberized TV restoration. *Physica D: Nonlinear Phenomena*, 175:241–251, 2003.

[39] J. Shen. Bayesian video dejittering by BV image model. *SIAM J. Appl. Math.*, 64(5):1691–1708, 2004.

[40] J. Sullivan, R. Miller, and G. Pios. Image halftoning using a visual model in error diffusion. *J. Opt. Soc. Amer. A*, 10(8):1714–1724, 1993.

[41] R. Ulichney. *Digital Halftoning*. MIT Press, Cambridge, Massachusetts, 1987.

[42] R. Ulichney. Dithering with blue noise. *Proc. IEEE*, 76:56–79, 1988.

[43] A. Vongkunghae, J. Yi, and R. B. Wells. A printer model using signal processing techniques. *IEEE Trans. Image Processing*, 12(7):776–783, 2003.

[44] J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner-Verlag, Stuttgart, Germany, 1998.

[45] J. Weickert, B.M. ter Haar Romeny, and M.A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Trans. Image Processing*, 7:398–410, 1998.

[46] A. Witkin. Scale space filtering - a new approach to multi-scale description. In R. Ullmann, editor, *Image Understanding*, pages 79–95, Ablex, New Yersey, USA, 1984.

[47] I. H. Witten and R. M. Neal. Using Peano curves for bilevel display of continuous-tone images. *IEEE Comput. Graph. Applicat.*, 2:47–52, 1982.

[48] P. W. Wong. Adaptive error diffusion and its application in multiresolution rendering. *IEEE Trans. Image Processing*, 5:1184–1196, 1996.

[49] P. W. Wong and J. P. Allebach. Optimum error diffusion kernel design. *Proc. SPIE*, 3018:236–242, 1997.