INVERSE HALFTONING BY POINTWISE SHAPE-ADAPTIVE DCT REGULARIZED DECONVOLUTION

Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian

Institute of Signal Processing, Tampere University of Technology
P.O. BOX 553, 33101 Tampere, Finland
firstname.lastname@tut.fi

ABSTRACT

We present a novel inverse halftoning method based on regularized deconvolution combined with shape-adaptive DCT (SA-DCT) filtering. Using a linear model of the error diffusion halftoning, after performing regularized deconvolution, we construct a continuous-tone estimate by local SA-DCT filtering on adaptive neighborhoods defined by the anisotropic *LPA-ICI*. These neighborhoods are highly adaptive to image details. The true signal encompassed by them is mostly homogeneous and does not contain edges or singularities. Hence, the proposed method faithfully preserves image details and effectively suppresses halftone patterns. The experimental results demonstrate outstanding performance in terms of both visual quality and mean squared error, significantly outperforming the current state-of-the-art inverse halftoning methods.

1. INTRODUCTION

Halftoning is the transformation of a continuous-tone grayscale image into a black and white binary one. In order to take into account the characteristics of the human visual system which acts as a low-pass filter, halftones are generated in such a way that the difference between the halftone image and the original grayscale image is compacted into the high-frequency components. Halftoning techniques include ordered dithering or screening (dispersed-dot and clustered-dot), error diffusion, blue-noise dithering, direct binary search, and others. Ordered dithering and error diffusion are the most computationally efficient and commonly used in applications.

Inverse halftoning is the reconstruction process of a continuous-tone image from a binary halftone. Based on the nature of the halftone process, the inverse halftoning should mimic the human visual system. Thus, roughly speaking, all inverse halftoning techniques should perform a sort of low-pass filtering. Although fixed-kernel low-pass filtering is simple to implement, usually it yields unsatisfactory results (presence of blur and artifacts). Various approaches to the inverse halftoning problem have been proposed, including deconvolution in wavelet domain [1], projection onto convex sets [2], look-up tables (LUT) [3], neural networks [4], and others.

In this work, we consider exclusively the reconstruction of a continuous-tone image from an *error diffusion* halftone. We extend the inverse halftoning method based on anisotropic *LPA-ICI* regularized deconvolution [5] by adopting the shape-adaptive DCT (SA-DCT) [6] regularized deconvolution [7].

This work was supported by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence program [2006 - 2011]).

2. NOTATION AND ERROR DIFFUSION MODEL

Let us denote the true (continuous-tone) image by $y : X \to [0,1] \subset \mathbb{R}$, where $X \subset \mathbb{Z}^2$ is the set of pixel coordinates. The halftone (binary) image $z : X \to \{0,1\} \subset \mathbb{Z}$ is obtained by applying error diffusion halftoning with a pre-defined error filter h^{ed}, such as Floyd-Steinberg's [8] or Jarvis' [9] filters. The standard error diffusion halftoning algorithm [8] is as follows.

- Inputs: y, h^{ed}. Output: z. Initialization: $\widetilde{y} = y$.
- In a raster-scan manner, for each pixel coordinate, i.e. $\forall x \in X$, do the following three steps.
  1. $z(x) = \begin{cases} 1, & \text{if } \widetilde{y}(x) > 0.5 \\ 0, & \text{otherwise.} \end{cases}$
  2. Compute the error $e_x = \widetilde{y}(x) - z(x)$.
  3. Update $\widetilde{y}$ by distributing e_x among unprocessed neighbors of x:

$$\widetilde{y}(x + x') = \widetilde{y}(x + x') + e_x h^{\text{ed}}(x'), \ \forall x'.$$

The Floyd-Steinberg's and Jarvis' error filters that can be used in the above procedure are defined as

$$h^{\text{ed}}_{\text{F-S}} = \frac{1}{16} \begin{bmatrix} 0 & 0^* & 7 \\ 3 & 5 & 1 \end{bmatrix}$$

and

$$h^{\text{ed}}_{\text{J}} = \frac{1}{48} \begin{bmatrix} 0 & 0 & 0^* & 7 & 5 \\ 3 & 5 & 7 & 5 & 3 \\ 1 & 3 & 5 & 3 & 1 \end{bmatrix},$$

respectively, where $*$ denotes the position $x' = \langle 0,0 \rangle$ in step 3 above.

We consider the linear model of the error diffusion proposed by Kite et al. in [10]. If we denote the discrete

Fourier transform (DFT) of y and z by $Y = \mathcal{F}(y)$ and $Z = \mathcal{F}(z)$, respectively, the adopted observation model can be expressed in DFT domain as

$$Z = PY + Q\eta, \tag{1}$$

where η is white Gaussian noise of standard deviation $\sigma = 1$, and P and Q are the "signal transfer function" and the "noise transfer function," respectively. These transfer functions are defined by

$$P = \frac{K^{\text{gain}}}{1 + (K^{\text{gain}} - 1) H^{\text{ed}}},$$

and

$$Q = \frac{1 - H^{\text{ed}}}{1 + (K^{\text{gain}} - 1) H^{\text{ed}}},$$

where $K^{\text{gain}} \in \mathbb{R}$ is a scalar constant and $H^{\text{ed}} = \mathcal{F}\left(h^{\text{ed}}\right)$. It was found in [10] that $K^{\text{gain}} = 2.05$ and $K^{\text{gain}} = 4.5$ can be used for most images without large error for, respectively, Floyd-Steinberg's and Jarvis' error filters.

Because H^{ed} is usually a low-pass filter, Q is a high-pass filter that compacts the energy of the additive noise in Equation (1) in the high frequency DFT components.

3. INVERSE HALFTONING BY SHAPE-ADAPTIVE DCT FILTERING

Let us present the main steps of the proposed inverse halftoning method. Given an error-diffusion halftone observation z and both the kernel h^{ed} and the scalar constant K^{gain} which characterize the error diffusion model (1), the method is given by the following steps, which mimic the deblurring approach from [5].

1. Perform a regularized inversion (*RI*) of the product with P in (1), which produces the *RI* estimate z^{RI}.
2. Filter the noise from z^{RI} by hard-thresholding in local SA-DCT domain and as a result obtain the intermediate estimate $\widehat{y}^{RI}$ of the true continuous-tone image.
3. Using $\widehat{y}^{RI}$ as a "pilot" estimate, perform regularized Wiener inverse (*RWI*) on z in order to obtain the *RWI* estimate z^{RWI}, which is a better estimate than z^{RI} because the use of the power spectrum of $\widehat{y}^{RI}$ improves the accuracy of the deconvolution.
4. Filter the noise from z^{RWI} by Wiener filtering (using $\widehat{y}^{RI}$ as a "pilot" estimate) in local SA-DCT domain and compute the final estimate $\widehat{y}$.

For steps 1 and 3, we use the Tikhonov *RI* and *RWI* operators, respectively. Although *regularization* is conventionally used to contain the energy of the noise after inversion, here we use it to compensate for the inevitable inaccuracies associated with the assumed *linear* model (1) of the *nonlinear* error diffusion halftoning. The *nonlinearity* is due to the bi-level quantization in step 1 of the algorithm presented in the previous section.

We apply the SA-DCT on neighborhoods defined adaptively for each processed coordinate $x \in X$ by the 8-directional anisotropic *LPA-ICI* [5, 7]. Each such neighborhood has octagonal shape defined by the polygonal hull of the adaptive *LPA-ICI* kernels' supports. We denote the neighborhood at x by $\tilde{U}_x^+ \subset X$. In the sequel, the subscript $_{|\tilde{U}_x^+}$ denotes the restriction of a signal on $\tilde{U}_x^+$. For example, $z_{|\tilde{U}_x^+} : \tilde{U}_x^+ \to \{0, 1\}$ is defined as

$$z_{|\tilde{U}_x^+}(x') = z(x'), \ \forall x' \in \tilde{U}_x^+.$$

We denote the cardinality of $\tilde{U}_x^+$ (the number of elements in $\tilde{U}_x^+$) by $\left|\tilde{U}_x^+\right|$.

3.1. Regularized inverse (*RI*)

We compute the *RI* of the convolution in (1) by multiplying the spectrum Z of the halftone observation with the *RI* operator

$$T^{RI} = \frac{\overline{P}}{|P|^2 + \varepsilon_1^2 |Q|^2},$$

where $\varepsilon_1 \in \mathbb{R}$ is a fixed regularization parameter and $\overline{P}$ is the complex conjugate of P. Hence, the *RI* estimate is

$$\begin{aligned} z^{RI} &= \mathcal{F}^{-1}\left(T^{RI} Z\right) \\ &= \mathcal{F}^{-1}\left(\frac{|P|^2}{|P|^2 + \varepsilon_1^2 |Q|^2} Y\right) + \qquad (2a) \\ &\quad + \mathcal{F}^{-1}\left(\frac{\overline{P} Q}{|P|^2 + \varepsilon_1^2 |Q|^2} \eta\right) \qquad (2b) \end{aligned}$$

The first addend (2a) is a biased estimate of the true image and the second addend (2b) is colored noise.

3.2. Filtering of the *RI* estimate

We filter locally the *RI* estimate z^{RI} in SA-DCT domain in order to attenuate the additive colored noise (2b). For each processed coordinate $x \in X$, we do this in the following manner. First, use *LPA-ICI* to obtain an adaptive neighborhood $\tilde{U}_x^+$. Then, compute the SA-DCT of $z^{RI}_{|\tilde{U}_x^+}$. The SA-DCT coefficients are given by

$$\varphi_{z^{RI},x} = \mathcal{T}_{\tilde{U}_x^+}\left(z^{RI}_{|\tilde{U}_x^+} - m_{\tilde{U}_x^+}\left(z^{RI}\right)\right), \tag{3}$$

where $\mathcal{T}_{\tilde{U}_x^+}$ is the orthonormal SA-DCT transform and $m_{\tilde{U}_x^+}\left(z^{RI}\right)$ is the mean value of $z^{RI}_{|\tilde{U}_x^+}$. The mean subtraction is done in order to avoid the "mean weighting effect" [11] of the orthonormal SA-DCT.

Let the SA-DCT coefficients be indexed as $\varphi_{z^{RI},x}(i) \in \mathbb{R}$, where $i = 1, \ldots, \left|\tilde{U}_x^+\right|$. In addition, let us designate the SA-DCT basis element corresponding to $\varphi_{z^{RI},x}(i)$ as $\psi^{(i)}_{\tilde{U}_x^+}$, whose domain is zero-extended from $\tilde{U}_x^+$ to X, i.e. $\psi^{(i)}_{\tilde{U}_x^+}(x') = 0$ for all $x' \in X \setminus \tilde{U}_x^+$. The variance of each SA-DCT coefficient can be obtained as

$$\sigma^2_{\varphi_{z^{RI},x}}(i) = \sigma^2 \left\| T^{RI} Q \mathcal{F}\left(\psi^{(i)}_{\tilde{U}_x^+}\right)\right\|_2^2 \tag{4}$$

and the variance of the mean as

$$\sigma^2_{m_{z^{RI},x}} = \frac{\sigma^2}{\left|\tilde{U}^+_x\right|^2}\left\|T^{RI}Q\mathcal{F}\left(\chi_{\tilde{U}^+_x}\right)\right\|_2^2, \quad (5)$$

where $\chi_{\tilde{U}^+_x} : X \to \{0,1\}$ is the characteristic function of $\tilde{U}^+_x$.

The colored noise in $z^{RI}_{\left|\tilde{U}^+_x\right.}$ is then attenuated by hard-thresholding,

$$\varphi^{thr}_{z^{RI},x}(i) = \begin{cases} \varphi_{z^{RI},x}(i), & \text{if } \left|\varphi_{z^{RI},x}(i)\right| < \varphi_{thr}(i) \\ 0, & \text{otherwise,} \end{cases}$$

for all $i = 1, \ldots, \left|\tilde{U}^+_x\right|$, where the threshold

$$\varphi_{thr}(i) = \sigma_{\varphi_{z^{RI},x}}(i)\,\lambda_{thr}\sqrt{2\ln\left(\left|\tilde{U}^+_x\right|+1\right)}$$

depends on the variance (4) of each SA-DCT coefficient. The inverse SA-DCT yields the local estimate of the true signal on $\tilde{U}^+_x$,

$$\widehat{y}^{RI}_{\tilde{U}^+_x} = \mathcal{T}^{-1}_{\tilde{U}^+_x}\left(\varphi^{thr}_{z^{RI},x}\right) + m_{\tilde{U}^+_x}\left(z^{RI}\right). \quad (6)$$

The addition of the mean $m_{\tilde{U}^+_x}\left(z^{RI}\right)$ compensates for its subtraction in (3). The above procedure, i.e. Equations (3-6), is repeated for each processed coordinate $x \in X$.

Because the neighborhoods corresponding to neighboring spatial coordinates can (and usually do) overlap, the global filtered *RI* estimate is computed as a weighted average of the overlapping local ones,

$$\widehat{y}^{RI} = \frac{\sum_{x\in X} w_x \widehat{y}^{RI}_{\tilde{U}^+_x}}{\sum_{x\in X} w_x \chi_{\tilde{U}^+_x}}, \quad (7)$$

where the adaptive weight $w_x \in \mathbb{R}$ is defined as

$$w_x = \frac{1}{\left(\left|\tilde{U}^+_x\right|\sigma^2_{m_{z^{RI},x}} + \sum_{i\in S}\sigma^2_{\varphi_{z^{RI},x}}(i)\right)\left|\tilde{U}^+_x\right|},$$

where the set S contains the indices of the SA-DCT coefficients that are retained (i.e. not equal to zero) after hard-thresholding. Let us remark that the local estimates $\widehat{y}^{RI}_{\tilde{U}^+_x}$ are zero-extended to X in Equation (7).

3.3. Regularized Wiener inverse (*RWI*)

Provided the intermediate estimate $\widehat{y}^{RI}$, we can apply a *RWI* on (1) by multiplying it with the *RWI* operator

$$T^{RWI} = \frac{\overline{P}\left|\widehat{Y}^{RI}\right|^2}{\left|P\widehat{Y}^{RI}\right|^2 + \varepsilon_2^2|Q|^2},$$

where $\varepsilon_2 \in \mathbb{R}$ is a fixed regularization parameter. Hence, the *RWI* estimate is

$$\begin{aligned} z^{RWI} &= \mathcal{F}^{-1}\left(T^{RWI}Z\right) \\ &= \mathcal{F}^{-1}\left(\frac{\left|P\widehat{Y}^{RI}\right|^2}{\left|P\widehat{Y}^{RI}\right|^2 + \varepsilon_2^2|Q|^2}Y\right) + \quad (8a) \\ &\quad + \mathcal{F}^{-1}\left(\frac{\overline{P}\left|\widehat{Y}^{RI}\right|^2 Q}{\left|P\widehat{Y}^{RI}\right|^2 + \varepsilon_2^2|Q|^2}\eta\right) \quad (8b) \end{aligned}$$

Similarly to (2a) and (2b), the terms (8a) and (8b) are a biased estimate of the true image and additive colored noise, respectively.

3.4. Filtering of the *RWI* estimate

We filter the colored noise (8b) from z^{RWI} by empirical Wiener filtering in local SA-DCT domain.

For each neighborhood $\tilde{U}^+_x$, we compute the SA-DCT coefficients of the noisy z^{RWI} and the SA-DCT coefficients of intermediate estimate $\widehat{y}^{RI}$ as

$$\varphi_{z^{RWI},x} = \mathcal{T}_{\tilde{U}^+_x}\left(z^{RWI}_{\left|\tilde{U}^+_x\right.} - m_{\tilde{U}^+_x}\left(z^{RWI}\right)\right)$$

and

$$\varphi_{\widehat{y}^{RI},x} = \mathcal{T}_{\tilde{U}^+_x}\left(\widehat{y}^{RI}_{\left|\tilde{U}^+_x\right.} - m_{\tilde{U}^+_x}\left(\widehat{y}^{RI}\right)\right),$$

respectively. The variances of $\varphi_{z^{RWI},x}$ and the variance of $m_{\tilde{U}^+_x}\left(z^{RWI}\right)$ can be computed by Equations (4) and (5), respectively, after replacing T^{RI} with T^{RWI}.

The attenuation of the noise (8b) is done by multiplying the SA-DCT coefficients $\varphi_{z^{RWI},x}$ and the subtracted mean $m_{\tilde{U}^+_x}\left(z^{RWI}\right)$ with the Wiener coefficients

$$\omega_x = \frac{\varphi^2_{\widehat{y}^{RI},x}}{\varphi^2_{\widehat{y}^{RI},x} + \sigma^2_{\varphi_{z^{RWI},x}}} \quad (9)$$

and

$$\varpi_x = \frac{m_{\tilde{U}^+_x}\left(\widehat{y}^{RI}\right)^2}{m_{\tilde{U}^+_x}(\widehat{y}^{RI})^2 + \left|\tilde{U}^+_x\right|^{-1}\sigma^2_{m_{z^{RWI},x}}}, \quad (10)$$

respectively. Hence, the filtered *RWI* estimate on $\tilde{U}^+_x$ is given by

$$\widehat{y}^{RWI}_{\tilde{U}^+_x} = \mathcal{T}^{-1}_{\tilde{U}^+_x}\left(\omega_x\varphi_{z^{RWI},x}\right) + \varpi_x m_{\tilde{U}^+_x}\left(z^{RWI}\right).$$

In analogy to (7), after computing the local estimates of each processed $x \in X$, the global *RWI* estimate (and also final estimate) is computed as a weighted average of the overlapping local ones,

$$\widehat{y} = \frac{\sum_{x\in X} w_x \widehat{y}^{RWI}_{\tilde{U}^+_x}}{\sum_{x\in X} w_x \chi_{\tilde{U}^+_x}}, \quad (11)$$

where

$$w_x = \frac{1}{\left(\left|\tilde{U}^+_x\right|\varpi_x^2\sigma^2_{m_{z^{RWI},x}} + \sum_{i=1}^{\left|\tilde{U}^+_x\right|}\omega_x^2(i)\sigma^2_{\varphi_{z^{RI},x}}(i)\right)\left|\tilde{U}^+_x\right|}.$$

3.5. Fast computation of SA-DCT coefficients' variances

The calculation of the variances of the SA-DCT coefficients as given in Equations (4) and (5) is computationally very demanding because the 2D DFT with the size of the input image should be applied for each transform coefficient of each unique shape. In order to reduce complexity, we approximate the variances by decimation of the 2D DFT power spectra corresponding to the term $T^{RI}Q$ in Equation (4) (and also to the term $T^{RWI}Q$ in the case of *RWI*). These approximations allow us to use 2D DFT of smaller fixed size (we used 32×32 in our experiments), regardless of the size of the input image. More details on this operation can be found in [7].

	Error diffusion filter	
	Floyd-Steinberg	*Jarvis*
Lena	32.99	33.54
Peppers	32.22	31.98
Boats	30.11	30.75
Barbara	26.99	26.10
Couple	29.79	30.36
Hill	30.19	30.55

Table 1. Output PSNR results of the proposed method.

Method	Image	
	Lena	*Peppers*
Proposed	**32.99**	**32.22**
Anisotropic *LPA-ICI* [5]	32.40	31.60
WinHD (Neelamani et al.) [1]	32.10	31.20
Edge-based LUT (Chung et al.) [3]	31.34	31.04
POCS-Wavelet (Unal et al.) [2]	32.18	30.90
RBF+MLP (Huang et al.) [4]	31.32	27.85

Table 2. Comparison of the PSNR results of the proposed method with the results of the current state-of-the-art techniques. The error diffusion halftones were created with the Floyd-Steinberg's error kernel.

4. RESULTS AND CONCLUSIONS

We conducted experiments on error diffusion halftones of the standard test images *Lena*, *Boats*, *Peppers*, *Couple*, *Hill*, and *Barbara*, all of size 512×512. An implementation of the algorithm that produced the results presented in this section can be found at `http://www.cs.tut.fi/~foi/SA-DCT/`.

The output PSNR results of the proposed method are given in Table 1 and a comparison with the results of the current best-performing, to our knowledge, techniques is given in Table 2.

An illustration of *Lena* reconstructed by the proposed method from its Floyd-Steinberg halftone is presented in Figure 1. The estimate has very few artifacts and the image details are well reconstructed. This can be confirmed by observing the enlarged fragments in Figure 2, which demonstrate good preservation of smooth areas (e.g. the cheeks of *Lena*) and sharp edges (e.g. around the poles in *Boats*). Similar conclusions can be drawn form the illustrations given for *Boats* and *Peppers* in Figures 3 and 4, confirming that sharp edges and uniform areas are well reconstructed.

The results show that the proposed method based on deconvolution combined with local SA-DCT filtering is highly effective. Despite being computationally more demanding than the LUT approaches for example, it delivers unprecedented results in terms of both subjective visual quality and PSNR, significantly outperforming the current state-of-the-art inverse halftoning methods.

5. ACKNOWLEDGMENTS

We are thankful to Arsen Melkonyan for optimizing the parameters of the proposed algorithm.

6. REFERENCES

[1] R. Neelamani, R. Nowak, and R. Baraniuk, "WInHD: Wavelet-based Inverse Halftoning via Deconvolution," *IEEE Trans. on Image Process.*, October 2002. Submitted.

[2] G. B. Unal and A. E. Cetin, "Restoration of error-diffused images using projection onto convex sets," *IEEE Trans. on Image Process.*, vol. 10, no. 12, pp. 1836–1841, 2001.

[3] K.-L. Chung and S.-T. Wu, "Inverse halftoning algorithm using edge-based lookup table approach," *IEEE Trans. on Image Process.*, vol. 14, no. 10, pp. 1583–1589, 2005.

[4] W.-B. Huang, W.-C. Chang, Y.-W. Lu, A. Su, and Y.-H. Kuo, "Halftone/contone conversion using neural networks," in *Proc. Int. Conf. Image Process.*, vol. 5, (Singapore), pp. 3547–3550, Oct. 2004.

[5] A. Foi, V. Katkovnik, K. Egiazarian, and J. Astola, "Inverse halftoning based on the anisotropic LPA-ICI deconvolution," in *Proc. Int. TICSP Workshop Spectral Methods Multirate Signal Processing*, (Vienna, Austria), pp. 49–56, Sept. 2004.

[6] T. Sikora and B. Makai, "Shape-adaptive DCT for generic coding of video," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 5, pp. 59–62, Feb. 1995.

[7] A. Foi, K. Dabov, V. Katkovnik, and K. Egiazarian, "Shape-adaptive DCT for denoising and image reconstruction," in *Proc. SPIE Electronic Imaging: Algorithms and Systems V*, vol. 6064, (San Jose, CA, USA), Jan. 2006.

[8] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial greyscale," *Proc. Society for Information Display*, vol. 17, no. 2, pp. 75–77, 1976.

[9] J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A survey of techniques for the display of continuous tone pictures on bilevel displays," *Computer Graphics and Image Process.*, vol. 5, pp. 13–40, 1976.

[10] T. D. Kite, B. L. Evans, T. Sculley, and A. C. Bovik, "Digital halftoning as 2-D Delta-Sigma modulation," in *Proc. Int. Conf. Image Process.*, (Santa Barbara, CA, USA), pp. 799–802, Oct. 1997.

[11] P. Kauff and K. Schüür, "Shape-adaptive DCT with block-based DC separation and ΔDC correction," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 8, no. 3, pp. 237–242, 1998.

Figure 1. On the left is the Floyd-Steinberg halftone of *Lena* and on the right is the continuous-tone image reconstructed by the proposed method (PSNR 32.99 dB).

Figure 2. Illustration of fragments from Floyd-Steinberg halftones (top row) and from the corresponding reconstructed images (bottom row) of *Lena* (left and center) and *Boats* (right).

Figure 3. On the left is the Floyd-Steinberg halftone of *Boats* and on the right is the continuous-tone image reconstructed by the proposed method (PSNR 30.11 dB).

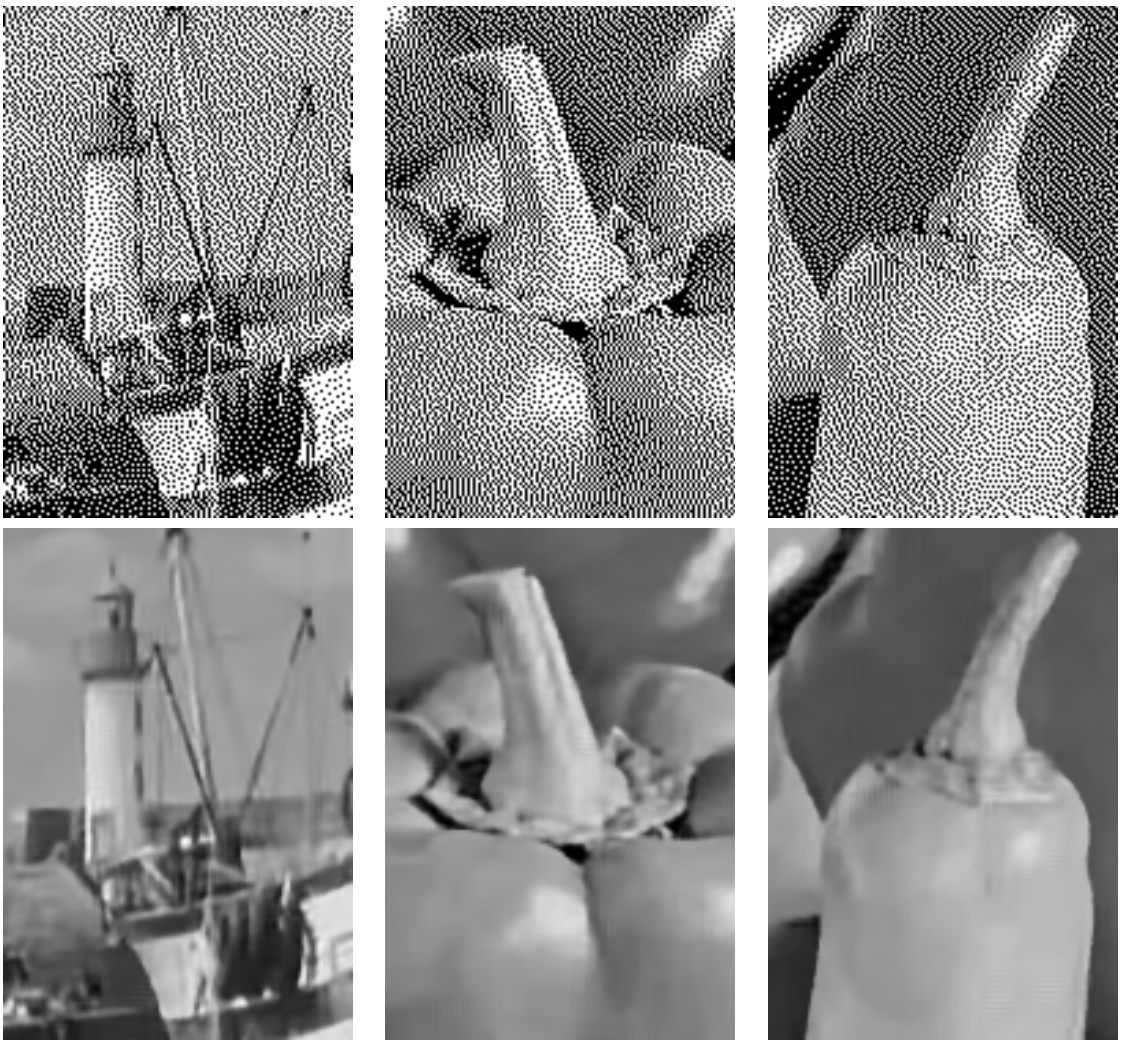

Figure 4. Illustration of fragments from Floyd-Steinberg halftones (top row) and from the corresponding reconstructed images (bottom row) of *Boats* (left) and *Peppers* (center and right).