

Structure-Aware Error Diffusion




Figure 1: Halftone images produced with different techniques. Visual quality of our images is comparable with that produced with structure-aware halftoning. At the same time, our method is two to three orders of magnitude faster than structure-aware halftoning (less than one second vs. hundreds of seconds). This difference is even bigger for larger images. **Important notice:** all images in this article have been prepared to be visualized on a graphic display, at 300 dots per inch. To visualize them correctly using Acrobat reader, please set the menu Edit/Preferences/Page Display/Resolution to “Custom resolution 100 pixels/inch”, and visualize using 300% zoom.

Abstract

We present an original error-diffusion method which produces visually pleasant halftone images while preserving fine details and visually identifiable structures present in the original images. Our method is conceptually simple and computationally efficient. The source image is analyzed, and its local frequency content is detected. The main component of the frequency content (main frequency, orientation and contrast) serve as lookup table indices to a pre-calculated database of modifications to a standard error diffusion. The modifications comprise threshold modulation and variation of error-diffusion coefficients. The whole system is calibrated in such a way that the produced halftone images are visually close to the original images (patches of constant intensity, patches containing sinusoidal waves of different frequencies/orientations/contrasts, as well as natural images of different origins).

Our system produces images of visual quality comparable to that presented in [Pang et al. 2008], but much faster. When processing typical images of linear size of several hundreds of pixels, our error-diffusion system is two to three orders of magnitude faster than [Pang et al. 2008]. Thanks to its speed combined with high visual quality, our error-diffusion algorithm can be used in many practical applications which may require digital halftoning: printing, visualization, geometry processing, various sampling techniques, etc.

1 Introduction

Digital halftoning is a technique for rendering continuous-tone images on devices having a limited number of available tone values. A typical application for digital halftoning algorithms is printing. In fact, millions of printers rely on digital halftoning in order to reproduce on paper a wide range of images: photos, business graphics, synthetic images, etc. Image quality and speed are two crucial issues to which any digital halftoning must be very sensitive. The algorithm that we introduce in this paper addresses both issues: it combines the quality of one of the best known techniques with the speed of conventional error-diffusion methods.

According to [Pappas et al. 2003; Lee and Allebach 2007], there are three classes of digital halftoning algorithms:

- Class I algorithms employ point processes, also called dithering [Bayer 1973; Ulichney 1987].
- Class II algorithms use neighborhood processes, such as error diffusion [Floyd and Steinberg 1976].
- Class III algorithms make use of iterative algorithms such as global search or direct binary search (DBS) [Analoui and Allebach 1992; Baqai et al. 2003].

Class I algorithms are the fastest, but their quality is insufficient for many practical applications, such as driving popular ink-jet printers or visualization on displays. Class III algorithms are the best from the quality viewpoint. Nevertheless, their complexity is too large, and they are impractical for most real-world applications. Class II algorithms are the most popular ones: they combine quality and speed. Nevertheless, the quality of images produced with Class II algorithms is far from perfect; it is clearly inferior to that of Class III.

Recently, impressive structure-aware halftoning has been pro-

posed [Pang et al. 2008]. This Class III algorithm performs iterative optimization according to a perceptual texture-sensitive metric. The method produces excellent results, but at very high cost: images of practical interest require minutes of processing time, which is obviously unacceptable for most practical applications.

The goal of the present article is to show that a Class II structure-aware algorithm can be built. The name of our system—structure-aware error diffusion—reflects its essential nature: on one hand, it inherits the essential features of structure-aware halftoning. On the other hand, it is basically a Class II (error-diffusion) algorithm. Consequently, our algorithm benefits from the best features of both: it has quality comparable to that of Class III algorithms, and the speed of Class II ones.

The rest of the article is organized as follows: In Section 2 we give a brief review of the main relevant prior art. In Section 3 we give an overview of the system. In Section 4 the results are presented and discussed. Finally, in Sections 5 and 6 we discuss limitations and draw conclusions.

2 Previous work

The error-diffusion algorithm, first introduced by Floyd and Steinberg [1976], has been extensively studied in the literature [Ulichney 1987; Kang 1999; Sharma 2002]. We consider that this research area is well-established and well-documented. For this reason, we explore in this section only the bibliography directly related to our work. We refer interested readers to excellent and very complete textbooks mentioned hereabout.

The main advantage of the initial version of the algorithm is its simplicity combined with fairly good overall visual quality of the produced binary images. The basic algorithm has a number of inherent drawbacks: first, the algorithm produces identifiable visually harmful artifacts in highlights and in dark areas (worm artifacts). Moreover, at certain intensity levels, close to $1/2$, $1/3$, $2/3$, etc., patches of regular structure may appear. These problems have been addressed in [Ostromoukhov 2001] by introducing intensity-dependent variable diffusion coefficients. This technique has been further improved in [Zhou and Fang 2003], by introducing intensity-dependent noise, in order to break regularities near critical intensity levels such as $1/2$, $1/3$, $2/3$, etc. Both papers use very similar optimization techniques for finding near-optimal intensity-dependent diffusion coefficients. This approach became a popular standard base (see [Pang et al. 2008]). For this reason, we shall designate by the term “standard error diffusion” a generic implementation of the method described in [Ostromoukhov 2001] and improved in [Zhou and Fang 2003]. In fact, we build our system on top of this “standard error diffusion”: we effectively use intensity-dependent diffusion coefficients. In addition to this “standard” component, we use a set of modified diffusion filters, as explained in Section 3.2. Modified diffusion filters depend on local spectral content of the image (local frequency/orientation/contrast); the proportion between “standard” and modified diffusion filters is adjusted during the calibration phase, as explained in Section 3.3.

Eschbach and Knox [1991] improved the basic error diffusion mechanism by modulating the threshold process with the edge response, in order to improve the standard error diffusion in the textured area. They used the inverse of the input image as the threshold. Their approach (also known as edge-enhancement) has been further improved and generalized by [Hwang et al. 2004; Kwak et al. 2006]. The latter considers both local luminance average and variation. However, even these improved versions of edge-enhancement contain a major drawback: the strength of the enhancement is controlled by a unique global coefficient. Consequently, the method is sensitive to only one particular sub-range

of frequencies and contrasts. In complex images with mixed frequency/contrast content this method can fail. This particular drawback has been explained in [Pang et al. 2008]. In accordance with Pang et al. [2008], we give more evidence of fundamental insufficiency of the edge-enhancement technique, in the most advanced version described in [Kwak et al. 2006]. Basic and generalized edge-enhancement techniques show one additional drawback when repetitive structures are present in the image. In fact, harmful undesirable Moiré-like patterns may appear. These harmful artifacts are clearly visible in Figure 2, no matter whether strong or weak enhancement coefficients were used. Once again, our method does not suffer from this harmful effect, as shown in Figure 2.

In the present paper, we further improve the ideas of [Kwak et al. 2006]. We do use image-dependent threshold modulation, but we replace the unsharp filter by a frequency- and orientation-dependent Gabor filter as explained in Section 3.1. Moreover, we modulate the strength of the threshold modulation by an amount determined during the calibration phase, as explained in Section 3.3. Our method provides a consistent improvement over [Kwak et al. 2006], even though visual difference between the results produced with both methods may appear rather subtle in some cases.

More recently, [Pang et al. 2008] introduced a new, more sophisticated structure-aware halftoning method. This method preserves both the global tone and the texture information, by using, within a complex optimization process, a combined metric which takes into account tonal and structural fidelity. This method generates results of excellent quality, one of the best known to date. However, as we mentioned before, this method belongs to Class III, and is therefore very slow. We use [Pang et al. 2008] as the target reference for quality for our Class II structure-aware error diffusion.

2.1 Local frequency analysis

Classical Fourier analysis [Bracewell 1980] represents an image as a sum of sine waves of different orientations, periods, amplitudes, and phases. Usually, a Discrete Fourier Transform (DFT) [Bracewell 1980] is performed on images represented by discrete sets of intensities. Frequencies describe the images’ essential nature: high frequencies correspond to fine details whereas low frequencies represent large-scale variations. Natural images may have heterogeneous frequency content.

While DFT provides a global analysis, many image processing applications need local analysis. Some existing techniques have been designed to that end [Turner 1986; Simoncelli and Freeman 1995]. Our error-diffusion method uses local frequency content (dominant frequency/orientation/contrast within a small window) in order to get better sensitivity to structures present in the image. In presence of a complex frequency content, our method takes into account the most prominent frequency component only.

Local spectral analysis is not part of our contribution. Any analysis tool can be used as long as it detects the main frequency content. Local frequency content analysis described in Chapter 13 of [Jahne 2004] is well-documented and satisfies our needs. Let us briefly summarize their technique. A gradient vector Δg is computed for each pixel of the image. Finding the dominant orientation within a window is done by maximizing the term $\sum_{\mathcal{N}} w(\Delta g \cdot \vec{k}_0)^2$ over the neighborhood, where w is a windowing function and \vec{k}_0 is the local orientation vector. Then, the DFT of the neighborhood window is computed and the dominant frequency and contrast are extracted, using known dominant orientation. For more details, refer to [Jahne 2004].




Figure 2: Patches of sine waves. In each set of 6x6 patches, frequency changes along the x axis; contrast – along the y axis. Standard error diffusion spoils low-contrast patches, especially for high frequencies. Edge-enhancement is unable to reproduce faithfully all frequencies and contrasts. Please notice the excellent visual coherence of the results produced with our method: the progression of relative contrast is visually close to that of the original.




Figure 3: Schematic representation of the architecture of our structure-aware error diffusion. Please refer to Sections 3.1 and 3.2 for the meaning of all terms.

3 Structure-aware error diffusion

In this section, we provide an overview of our structure-aware error-diffusion algorithm.

The conventional error-diffusion algorithm processes each pixel of an input image in a particular sequential order (typically scanline or serpentine). The intensity of the pixel is compared with a thresh-

old (constant or position-dependent). The result of that comparison determines the intensity of the output pixel, 0 or 1; the quantization error is distributed to the unprocessed neighboring pixels according to a particular diffusion filter.

In our method, uniform thresholding is adaptively replaced by thresholding from a parametrized anisotropic Gabor filter, and standard diffusion filter is adaptively replaced by a parametrized Gaussian filter. The adaptive transition from standard to structure-driven error diffusion is controlled by a weight parameter w .

During the offline calibration step, we select a perceptually optimal set of parameters for each of a large number of combinations of frequency/orientation/contrast (typically, 6^3 combinations). In the run-time part of our method, analyzed frequency content of each pixel then serves as index to the calibration table, and error diffusion is performed in accordance with the retrieved parameters. Figure 3 gives a schematic representation of our run-time architecture; pseudo-code is given in Appendix A.

3.1 Threshold modulation

Threshold modulation is a powerful tool for edge enhancement in error diffusion [Knox 1989; Eschbach and Knox 1991; Hwang et al. 2004; Kwak et al. 2006]. In this paper, we bring two important

modifications to previous threshold modulation schemes. First, our modulation depends on the local frequency content of the image (local dominant frequency/orientation/contrast). Second, we calibrate the amount of threshold modulation in such a way that for each specific combination of dominant frequency/orientation/contrast, our error-diffusion system produces halftone output which is *perceptually* close to the source image.

[Kwak et al. 2006] applied an isotropic unsharp mask filter on the input image for calculating the threshold matrix. This improves the original edge enhancement [Knox 1989; Eschbach and Knox 1991], but, as [Pang et al. 2008] pointed out, the method suffers from blurriness because important and unimportant frequencies are equally amplified.

More recently, [Kang et al. 2009] remarked that feature enhancement using isotropic filters such as difference-of-Gaussians does not preserve the sense of “directedness”, making aggregates of light edge pixels look less like lines.

Our method relies on the hypothesis that anisotropic filters should perform better than their isotropic counterpart in presence of anisotropic image structure. We chose the Gabor filter [Gabor 1946] to accomplish this task for its very simple spectral properties. Since anisotropy is a property that varies throughout the image, the parameters of the Gabor filter change locally. We checked our hypothesis, and found that our method produced very satisfactory results.

Let us describe the threshold modulation part of our structure-aware error diffusion. Standard error diffusion uses constant threshold. We use threshold modulation determined as a linear interpolation between two terms: (1) frequency-content-dependent filtered version F_G of input image Img , and (2) noise modulation m , as in [Zhou and Fang 2003].

The term F_G is the result of a local convolution with a suitable Gabor kernel and input image Img modulated by a variable strength parameter β :

$$F_G(x, y) = \beta \cdot \sum_i \sum_j Img(x - i, y - j) \cdot Gabor_{\theta, f}(i, j)$$

where the function $Gabor_{\theta, f}$ is defined from local frequency f and orientation θ given by image analysis:

$$Gabor_{\theta, f}(x, y) = \exp\left(-\frac{x'^2 + y'^2}{2\sigma_G^2}\right) \cos(fx') + c,$$

$$x' = x \cos \theta + y \sin \theta,$$

$$y' = -x \sin \theta + y \cos \theta,$$

and c is a constant that nulls the sum of the filter coefficients. In our current implementation, a fixed value $\sigma_G = 1.6$ was used.

The amount of threshold modulation β constitutes the main balancing factor for sensitivity to different natures of structural content. It is controlled carefully at the calibration step for the whole range of frequencies and contrasts. Since the image analysis described in Section 2.1 provides a finite number of combinations of frequency/orientation, all Gabor kernels were pre-calculated for fast run-time access (see Appendix A).

3.2 Modified diffusion filter

The method described in [Ostromoukhov 2001] provides a collection of 3-coefficient error-diffusion filters that produce best results on regions of constant tone. However, this method was conceived

for low-frequency images and may blur high-frequency details. Notice that the midtone-improved version of the standard error diffusion [Zhou and Fang 2003] suffers from the same drawback, as shown in Figures 1 and 7. Ulichney mentioned in his book [Ulichney 1987] that larger filter size can allow sharper feature preservation. Also, [Marcu and Abe 1996] explored the use of analytical coefficient distribution, in particular that of a Gaussian kernel.

According to our hypothesis, diffusion coefficients following a distribution based on frequency, orientation, and contrast should preserve structure sensitivity better than standard coefficients in presence of strong frequency content. We checked this hypothesis, and found that best results were obtained from a 12-coefficient Gaussian distribution with adjustable variance σ^2 and anisotropy a .

Both σ and a are controlled by calibration. In the absence of local frequency content, standard coefficients H_a are used. Transition from standard to structure-driven diffusion coefficients H_b is controlled by linear interpolation with the calibration parameter w . Gaussian kernels are pre-calculated for fast run-time access (see Appendix A).

3.3 Calibration

Perhaps the most crucial step of our method is its calibration, the process by which we ensure that any local frequency information is associated with proper threshold modulation and error diffusion coefficients.

Our method relies on one fundamental hypothesis: by obtaining optimal halftoning results on an artificial image containing a single characteristic frequency, we should obtain near-optimal halftoning results on a region of a natural image that shows the same local frequency.

Automatic calibration was used by [Ostromoukhov 2001] and [Zhou and Fang 2003] for patches of uniform pixel luminance. Even though such an approach could be applied to our case, we chose to rely instead on the well-established method of side-by-side comparisons used, for example, in the classical color matching experiment [Wyszecki and Stiles 2000]. In our method, a uniform reference patch is shown side-by-side with a halftoned calibration patch whose parameters are adjusted by the observer until the best possible match between the two patches is reached. The experiment is repeated for several values of frequency/orientation/contrast of the reference patch.

Figure 4 illustrates the calibration process. To calibrate a given frequency, orientation, and contrast, we create a reference square patch with that frequency and orientation displaying two close but different contrasts: the central part, delimited by the red circle, shows the contrast to calibrate; the peripheral area shows a contrast value which has already been successfully calibrated.

First, the observer adjusts viewing distance until sine waves in the peripheral area are just imperceptible. At that distance, sine waves in the continuous-tone side of the central area should be at the threshold of visibility. Then, sliders are used to adjust the parameters for uncalibrated contrast until sine waves in the halftone side of the central area are also barely perceptible. The patches are dynamically updated, permitting the observer to act interactively.

Four parameters are required to control the behavior of our halftoning process: β is the strength of the applied threshold modulation, which is the amplitude of the Gabor convolution filter; σ and a are the parameters of the Gaussian error diffusion filter; w is a linear blending factor between no-contrast standard coefficients and noise-modulated threshold, and parametrized Gaussian coefficients and Gabor-filtered threshold.




Figure 4: The calibration patches. Left side: continuous-tone sine waves of fixed frequency and angle. Right side: halftone images produced using adjustable set of parameters. The peripheral area delimited by the red circle serves to determine the appropriate observation distance. For the purpose of illustration, we exaggerated the contrast difference between inner and outer areas.

Once the calibration is done, a 3D lookup table is filled by interpolation of the calibration parameters for all frequency/orientation/contrast. They are used at run-time for a lookup table search of an arbitrary required frequency/orientation/contrast parameter.

All images shown in this paper use the same calibration. In our implementation, we use a total of around 1.3MB of lookup data split among five tables: interpolated calibration parameters for all frequency/orientation/contrast; Gabor filters for a variety of frequencies and orientations; Gaussian diffusion filters for a variety of σ and a ; no-contrast standard diffusion filters for all tones; mid-tone improvement noise weights for all tones [Zhou and Fang 2003]. The pseudo-code in Appendix A indicates how each table is used.

3.4 Kernel size analysis

In Section 2.1, local frequency content is extracted from a pixel’s neighborhood. The quality and computation time of the analysis depend on the size of the neighborhood window. To obtain the best quality/computation-time tradeoff, a systematic study was done on the choice of kernel size.

Table 1 shows how kernel size relates to computation time. Since local frequency analysis is parallel by nature, we implemented it both on GPU and CPU.

Figure 5 illustrates how results produced by our method can be degraded due to a poor frequency analysis. Some important structure details were lost with the 8×8 kernel, even more with the 4×4 one. Best results were observed for the 32×32 kernel, but the quality difference between the 16×16 and the 32×32 kernels is very small. The best tradeoff between speed and quality was reached by the 16×16 kernel on our hardware.

| Kernel size | CPU | GPU |
|----------------|-------|-------|
| 4×4 | 1.18 | 0.04 |
| 8×8 | 2.27 | 0.23 |
| 16×16 | 6.74 | 0.75 |
| 32×32 | 26.31 | 15.09 |

Table 1: Frequency analysis time (in seconds) for a typical 512×512 image. CPU: AMD Athlon 64 Dual Core 6000+ @ 3.01GHz 4GB. GPU: nVidia Corporation GeForce 9800 GTX 512MB.

4 Results

We evaluate the results produced with our method using both subjective (visual) comparisons and objective measurements.

| | Our method | Struct.-aware halfton. | Edge enhanc. | Std. error-diff. | Ordered dither |
|--------|------------|------------------------|--------------|------------------|----------------|
| Ribbon | 43.14 | 41.04 | 36.27 | 48.56 | 30.16 |
| Arm | 38.57 | 36.78 | 28.11 | 43.11 | 29.46 |
| Knee | 35.40 | 34.63 | 27.66 | 40.91 | 28.71 |
| Desert | 39.92 | 38.21 | 40.37 | 43.91 | 28.88 |
| Cat | 39.82 | 37.51 | 35.51 | 43.02 | 28.39 |
| Tree | 40.66 | 37.67 | 34.95 | 44.20 | 28.35 |
| Snail | 38.75 | 36.92 | 35.67 | 43.29 | 30.36 |

Table 2: PSNR comparison.

4.1 Visual comparison

For visual evaluation, we use the set of test images already shown in previous articles, namely in [Pang et al. 2008]. This set comprises a wide range of important features: small details of various contrasts, smooth tone gradation, visually identifiable structures (leaves, hair, tissue or brush structure), etc.

Images in Figures 1 and 7 show a clear and undeniable advantage of our method over standard error diffusion: small details are better preserved; our images appear sharper and globally nicer. Visual comparison of our method with structure-aware halftoning [Pang et al. 2008] shows that both methods produce very comparable quality.

We added comparisons using a few artificial images: a grayscale ramp in Figure 6 and patches of sine waves of different frequencies and contrasts in Figure 2. The visual quality of the ramps is not a surprise: in absence of strong frequency content, our error diffusion becomes a mid-tone improved standard error diffusion. The patches in Figure 2 show a clear advantage of our method over all others, except the structure-aware halftoning. Compared to standard error diffusion, our method better reproduces the range of contrasts and frequencies. Our method performs much better than edge enhancement, too. Notice excellent visual coherence of the results produced with our method: the progression of relative contrast is visually close to that of the original.

4.2 Objective quality evaluation

Objective metrics were applied to all test images to confirm our method’s capacity to preserve tones and structures.

The peak-to-signal noise ratio (PSNR) is a standard measure of reconstruction quality for lossy image compression. [Pang et al. 2008] used this metric on blurred versions of grayscale and halftone images to evaluate tone preservation. PSNR for Gaussian blurred ($\sigma = 2$) grayscale and halftone image pairs are given in Table 2. Results show that PSNR for our method remains consistently close to that of standard error diffusion and structure-aware halftoning.

The mean structure similarity measure (MSSIM) is a visual quality metric first introduced by [Wang et al. 2004] and used by [Pang et al. 2008] as a central criterion for their convergence algorithm. From the results shown in Table 3, we find that our method is close to structure-aware halftoning, which confirms visual observation.

4.3 Computation time

Table 4 shows computation time for our structure-aware error diffusion compared to previous methods. Although the quality of our results is close to that of [Pang et al. 2008], our method is dramatically faster. In fact, execution time under one second makes our method suitable for real-world applications.




Figure 5: Structure-aware error diffusion using analysis tool of different kernel sizes.




Figure 6: Gray ramps halftoned with different methods.

| | Our method | Struct.-aware halfton. | Edge enhanc. | Std. error-diff. | Ordered dither |
|--------|------------|------------------------|--------------|------------------|----------------|
| Ribbon | 26.93 | 25.36 | 25.11 | 20.49 | 19.27 |
| Arm | 46.51 | 47.34 | 45.85 | 35.66 | 32.78 |
| Knee | 37.12 | 39.29 | 36.62 | 27.01 | 24.54 |
| Desert | 6.39 | 7.00 | 4.02 | 2.75 | 2.22 |
| Cat | 14.64 | 17.62 | 11.65 | 7.42 | 5.53 |
| Tree | 13.37 | 15.35 | 12.58 | 5.53 | 4.50 |
| Snail | 37.26 | 37.87 | 34.62 | 31.23 | 30.80 |

Table 3: MSSIM comparison.

| Our method | Our method + Analysis | Structure-aware halftoning | Edge-enhanc. | Std. error-diff. |
|------------|-----------------------|----------------------------|--------------|------------------|
| 0.22 | 0.97 | 135.41 | 0.10 | 0.04 |

Table 4: Computation time (in seconds) for a typical 512×512 image

Note that our code is not optimal, especially its analysis part. With a relatively modest effort, our execution time could be further improved.

5 Limitations and future work

Although our method produces very satisfactory results in most cases, there are several issues which could be improved. First, our method works very well in presence of one clearly dominant local frequency. In rare cases when two or more equally important frequencies are present, our method is presently misguided by the analysis. We plan to carry out research which would incorporate treatment of multiple equally important frequencies in the framework described in this article. A better, specific analysis tool should be developed for this purpose.

In this paper, we limited our consideration to one-channel images. Extension to multi-channel (3-channel RGB or 4-channel CMYK images) may require additional work. A naïve straightforward extension such as mutually independent per-channel halftoning works rather well. But, in order to get even better quality for color images,

a more sophisticated collaborative multi-channel error-diffusion scheme may be required. We plan to carry out research in this direction too.

The images shown in this paper have been prepared using a calibration with respect to a “dummy” LCD display. It would be quite worthwhile to integrate the calibration process into the framework defined by ICC [Sharma 2002], to be used with a concrete printing device having strong hardware constraints (e.g. strong non-linear dot gain, depending on the nature of the printer: electrostatic, ink-jet, etc.).

6 Conclusion

The error-diffusion method presented in this article combines the advantages of several halftoning methods known to date. On one hand, the overall visual quality of the images produced with our method is comparable with that of the best known error-diffusion algorithms, such as [Ostromoukhov 2001; Zhou and Fang 2003]. On the other hand, our algorithm is sensitive to structures present in the source images, in a wide range of frequencies and contrasts. From this point of view, our algorithm approaches the quality of [Pang et al. 2008], without being prohibitively slow. We have tested our algorithm with a wide variety of artificial and natural images, producing very satisfactory results. Our algorithm is conceptually simple and easy to implement.

Thus, we have shown that a structure-aware digital halftoning algorithm of Class II mentioned in Section 1 can be efficiently built.

We think that thanks to its unique combination of visual quality, sensitivity to structures, and speed, our algorithm may be used in a wide range of graphical applications in which digital halftoning is required. Typical applications may be printing, visualization, or various sampling techniques.

Acknowledgements

We would like to thank François Duranleau, Neil Stewart, Jean Vaucher, and Pierre Poulin for their precious help. We are grateful to the authors of [Pang et al. 2008] for insightful discussions. The funding of this work was partially provided by the Natural Sciences and Engineering Research Council of Canada.

References

- ANALOUI, M., AND ALLEBACH, J. P. 1992. Model based halftoning using direct binary search. *SPIE 1666*, 96–108.
- BAQAI, F. A., TAYLOR, C. C., AND ALLEBACH, J. P. 2003. Halftoning via direct binary search using analytical and stochastic printer models. *IEEE Trans. Image Processing 12*, 1–15.
- BAYER, B. E. 1973. An optimum method for two-level rendition of continuous-tone pictures. *IEEE Intl. Conf. on Communications 1*, 2611–2615.
- BRACEWELL, R. N. 1980. *Fourier Transform and Its Applications*. McGraw-Hill Education.
- ESCHBACH, R., AND KNOX, K. T. 1991. Error-diffusion algorithm with edge enhancement. *JOSA (A) 8*, 12, 1844–1850.
- FLOYD, R. W., AND STEINBERG, L. 1976. An adaptive algorithm for spatial grey scale. *Proc. Soc. Inf. Display 17*, 75–77.
- GABOR, D. 1946. Theory of communication. *JIEE 93*, 3, 429–459.
- HWANG, B.-W., KANG, T.-H., AND LEE, T.-S. 2004. Improved edge enhanced error diffusion based on first-order gradient shaping filter. In *IEA/AIE'2004: Proceedings of the 17th international conference on Innovations in applied artificial intelligence*, Springer Springer Verlag Inc, 473–482.
- JAHNE, B. 2004. *Practical Handbook on Image Processing for Scientific and Technical Applications, Second Edition*. CRC Press, Inc., Boca Raton, FL, USA.
- KANG, H., LEE, S., AND CHUI, C. K. 2009. Flow-based image abstraction. *IEEE Transactions on Visualization and Computer Graphics 15*, 1, 62–76.
- KANG, H. 1999. *Digital Color Halftoning*. SPIE Press.
- KNOX, K. T. 1989. Edge enhancement in error diffusion. In *Advance Printing of Paper Summaries, SPSE's 42nd Annual Conference*, 310–313.
- KWAK, N.-J., RYU, S.-P., AND AHN, J.-H. 2006. Edge-enhanced error diffusion halftoning using human visual properties. In *ICHIT '06: Proceedings of the 2006 International Conference on Hybrid Information Technology*, IEEE Computer Society, Washington, DC, USA, 499–504.
- LEE, C., AND ALLEBACH, J. P. 2007. The hybrid screen: improving the breed. *SPIE*, R. Eschbach and G. G. Marcu, Eds., vol. 6493, 649318.
- MARCU, G., AND ABE, S. 1996. Halftoning by back error compensation. *IS&T NIP 12*, 132–135.
- OSTROMOUKHOV, V. 2001. A simple and efficient error-diffusion algorithm. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 567–572.
- PANG, W.-M., QU, Y., WONG, T.-T., COHEN-OR, D., AND HENG, P.-A. 2008. Structure-aware halftoning. *ACM Transactions on Graphics 27*, 3 (Aug.), 89:1–89:8.
- PAPPAS, T. N., ALLEBACH, J. P., AND NEUHOF, D. L. 2003. Model-Based Digital Halftoning. *IEEE Signal Processing Magazine 20*, 4, 14–27.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann.
- SHARMA, G. 2002. *Digital Color Imaging Handbook*. CRC Press, Inc., Boca Raton, FL, USA.
- SIMONCELLI, E. P., AND FREEMAN, W. T. 1995. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *Proc 2nd IEEE Int'l Conf on Image Proc.*, IEEE Sig Proc Society, Washington, DC, vol. III, 444–447.
- TURNER, M. R. 1986. Texture discrimination by gabor functions. *Biol. Cybern.* 55, 2-3, 71–82.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press.
- WANG, Z., BOVIK, A., SHEIKH, H., AND SIMONCELLI, E. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4, 600–612.
- WYSZECKI, G., AND STILES, W. S. 2000. *Color Science: Concepts and Methods, Quantitative Data and Formulae (Wiley Series in Pure and Applied Optics)*, 2 ed. Wiley-Interscience.
- ZHOU, B., AND FANG, X. 2003. Improving mid-tone quality of variable-coefficient error diffusion using threshold modulation. *ACM Transactions on Graphics 22*, 3, 437–444.

APPENDIX A: Pseudo-code of the Structure-Aware Error Diffusion

```

STRUCTUREAWAREERRORDIFFUSION(Img, FreqContent)
1  Err ← NIL    ▷ Array of quantization error values
2  for each processed pixel (i, j)
3    ▷ Retrieve threshold and diffusion filter from lookup tables
4    β, σ, a, w ← GETCALIBRATIONPARAMS(FreqContenti,j)
5    Gabor ← GETGABORFILTER(FreqContenti,j)
6    m ← Noise()GETNOISEINTENSITY(Imgi,j)
7    FG ← β · Conv(Img, Gabor)
8    Hb ← GETGAUSSIANFILTER(σ, a)
9    Hd ← GETSTANDARDFILTER(Imgi,j)
10
11   modulatedThreshold ← Lerp(FGi,j, m, w)
12   modifiedDiffusionFilter ← Lerp(Hb, Hd, w)
13
14   ▷ Perform error-diffusion
15   Imgi,j ← Imgi,j + Erri,j
16   Outi,j ← QUANTIZE(Imgi,j, modulatedThreshold)
17   Erri,j ← Imgi,j - Outi,j
18   DISTRIBUTEERROR(i, j, Err, modifiedDiffusionFilter)
19  return Out

```

Lerp(v_1, v_2, t) is a linear interpolation function, which calculates $(1 - t)v_1 + tv_2$ [Pharr and Humphreys 2004].

GETXXX() functions access corresponding lookup tables.

CONV() is the local convolution described in section 3.1

QUANTIZE() compares pixel luminance to threshold and returns BLACK if below, WHITE otherwise.

DISTRIBUTEERROR() distributes quantization errors over unprocessed neighbor pixels according to the *modifiedDiffusionFilter* values.




Figure 7: A few sample images produced with different halftoning techniques.