Reinstating Floyd-Steinberg: Improved Metrics for Quality Assessment of Error Diffusion Algorithms

Sam Hocevar[1] and Gary Niger[2]

[1] Laboratoire d'Imagerie Bureautique et de Conception Artistique
14 rue de Plaisance, Paris, France
[2] 143 Rolloffle Avenue, Tarzana, California 91356
sam@hocevar.net, gary_niger@gnaa.us

Abstract. In this contribution we introduce a little-known property of error diffusion halftoning algorithms which we call *error diffusion displacement*. By accounting for the inherent sub-pixel displacement caused by the error propagation, we correct an important flaw in most metrics used to assess the quality of resulting halftones. We find these metrics to usually highly underestimate the quality of error diffusion in comparison to more modern algorithms such as direct binary search. Using empirical observation, we give a method for creating computationally efficient, image-independent, model-based metrics for this quality assessment. Finally, we use the properties of error diffusion displacement to justify Floyd and Steinberg's well-known choice of algorithm coefficients.
Keywords: halftoning, error diffusion, image quality, human visual system, color quantization

1 Introduction

Image dithering is the process of reducing continuous-tone images to images with a limited number of available colours. Applications vary tremendously, from laser and ink-jet printing to display on small devices such as cellphones, or even the design of banknotes.

Countless methods have been published for the last 40 years that try to best address the problem of colour reduction. Comparing two algorithms in terms of speed or memory usage is often straightforward, but how exactly a halftoning algorithm performs quality-wise is a far more complex issue, as it highly depends on the display device and the inner workings of the human eye.

Though this document focuses on the particular case of bilevel halftoning, most of our results can be directly adapted to the more generic problem of colour reduction.

2 Halftoning algorithms

The most ancient halftoning method is probably classical screening. This highly parallelisible algorithm consists in tiling a dither matrix over the image and using its elements as threshold values. Classical screening is known for its structural artifacts such as the cross-hatch patterns caused by Bayer ordered dither matrices [1]. However, modern techniques such as the void-and-cluster method [2], [3] allow to generate screens yielding visually pleasing results.

Error diffusion dithering, introduced in 1976 by Floyd and Steinberg [4], tries to compensate for the thresholding error through the use of feedback. Typically applied in raster scan order, it uses an error diffusion matrix such as the following one, where x denotes the pixel being processed:

$$\frac{1}{16}\begin{vmatrix} - & x & 7 \\ 3 & 5 & 1 \end{vmatrix}$$

Though efforts have been made to make error diffusion parallelisable [5], it is generally considered more computationally expensive than screening, but carefully chosen coefficients yield good visual results [6].

Model-based halftoning is the third important algorithm category. It relies on a model of the human visual system (HVS) and attempts to minimise an error value based on that model. One such algorithm is direct binary seach (DBS) [10], also referred to as least-squares model-based halftoning (LSMB) [16].

HVS models are usually low-pass filters. Nasanen [9], Analoui and Allebach found that using Gaussian models gave visually pleasing results, an observation confirmed by independent visual perception studies [11].

DBS yields halftones of impressive quality. However, despite efforts to make it more efficient [12], it suffers from its large computational requirements and error diffusion remains a more widely used technique.

3 Error diffusion displacement

Most error diffusion implementations parse the image in raster scan order. Boustrophedonic (serpentine) scanning has been shown to cause fewer visual artifacts [7], but other, more complex processing paths such as Hilbert curves [8] are seldom used as they do not improve the image quality significantly.

Intuitively, as the error is always propagated to the bottom-left or bottom-right of each pixel (Fig. 1), one may expect the resulting image to be slightly translated. This expectation is confirmed visually when rapidly switching between an error diffused image and the corresponding DBS halftone.

This small translation is visually innocuous but we found that it means a lot in terms of error computation. A common way to compute the error between an image $h_{i,j}$ and the corresponding binary halftone $b_{i,j}$ is to compute the mean square error between modified versions of the images, in the form:

$$E(h,b) = \frac{(||v * h_{i,j} - v * b_{i,j}||_2)^2}{wh} \tag{1}$$

where w and h are the image dimensions, $*$ denotes the convolution and v is a model for the human visual system.

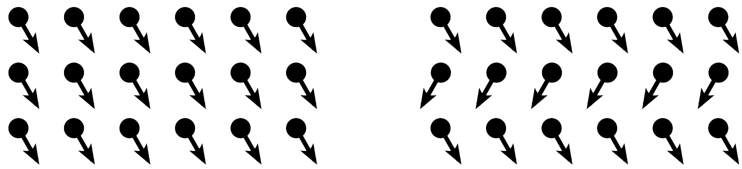

Fig. 1. Floyd-Steinberg error diffusion direction in raster scan (left) and serpentine scan (right).

To compensate for the slight translation observed in the halftone, we use the following error metric instead:

$$E_{dx,dy}(h,b) = \frac{(||v * h_{i,j} - v * t_{dx,dy} * b_{i,j}||_2)^2}{wh} \tag{2}$$

where $t_{dx,dy}$ is an operator which translates the image along the (dx, dy) vector. By design, $E_{0,0} = E$.

A simple example can be given using a Gaussian HVS model:

$$v(x,y) = e^{\frac{x^2+y^2}{2\sigma^2}} \tag{3}$$

Finding the second filter is then straightforward:

$$(v * t_{dx,dy})(x,y) = e^{\frac{(x-dx)^2+(y-dy)^2}{2\sigma^2}} \tag{4}$$

Experiments show that for a given image and a given corresponding halftone, $E_{dx,dy}$ has a local minimum almost always away from $(dx, dy) = (0, 0)$ (Fig. 2). Let E be an error metric where this remains true. We call the local minimum E_{min}:

$$E_{min}(h,b) = \min_{dx,dy} E_{dx,dy}(h,b) \tag{5}$$

For instance, a Floyd-Steinberg dither of *Lena* with $\sigma = 1.2$ yields a per-pixel mean square error of 3.67×10^{-4}. However, when taking the displacement into account, the error becomes 3.06×10^{-4} for $(dx, dy) = (0.165, 0.293)$. The new, corrected error is significantly smaller, with the exact same input and output images.

Experiments show that the corrected error is always noticeably smaller except in the case of images that are already mostly pure black and white. The experiment was performed on a database of 10,000 images from common computer vision sets and from the image board *4chan*, providing a representative sampling of the photographs, digital art and business graphics widely exchanged on the Internet nowadays [13].

In addition to the classical Floyd-Steinberg and Jarvis-Judice-Ninke kernels, we tested two serpentine error diffusion algorithms: Ostromoukhov's simple error

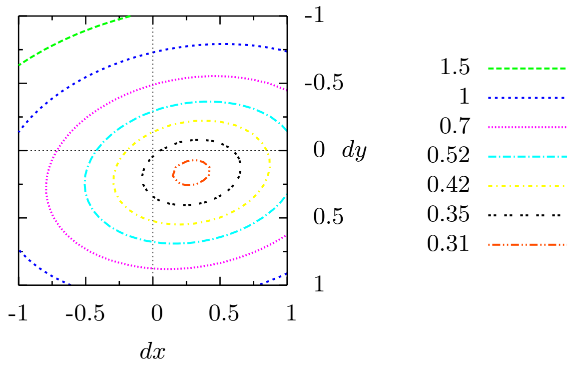

Fig. 2. Mean square error for the *Lena* image ($\times 10^4$). v is a simple 11×11 Gaussian convolution kernel with $\sigma = 1.2$ and (dx, dy) vary in $[-1, 1] \times [-1, 1]$.

diffusion [15], which uses a variable coefficient kernel, and Wong and Allebach's optimum error diffusion kernel [14]:

	$E \times 10^4$	$E_{min} \times 10^4$
raster Floyd-Steinberg	3.7902	3.1914
raster Ja-Ju-Ni	9.7013	6.6349
Ostromoukhov	4.6892	4.4783
optimum kernel	7.5209	6.5772

We clearly see that usual metrics underestimate the quality of error-diffused halftones, especially in raster scan. Algorithms such as direct binary search, on the other hand, do not suffer from this bias since they are designed to minimise the very error induced by the HVS model.

4 An image-independent corrected quality metric for error-diffused halftones

We have seen that for a given image, $E_{min}(h, b)$ is a better and fairer visual error measurement than $E(h, b)$. However, its major drawback is that it is highly computationally expensive: for each image, the new (dx, dy) values need to be calculated to minimise the error value.

Fortunately, we found that for a given raster or serpentine scan error diffusion algorithm, there was often very little variation in the optimal (dx, dy) values (Fig. 3 and 4).

For each algorithm, we choose the (dx, dy) values at the histogram peak and we refer to them as the *algorithm's displacement*, as opposed to the *image's displacement* for a given algorithm. We call $E_{fast}(h, b)$ the error computed at

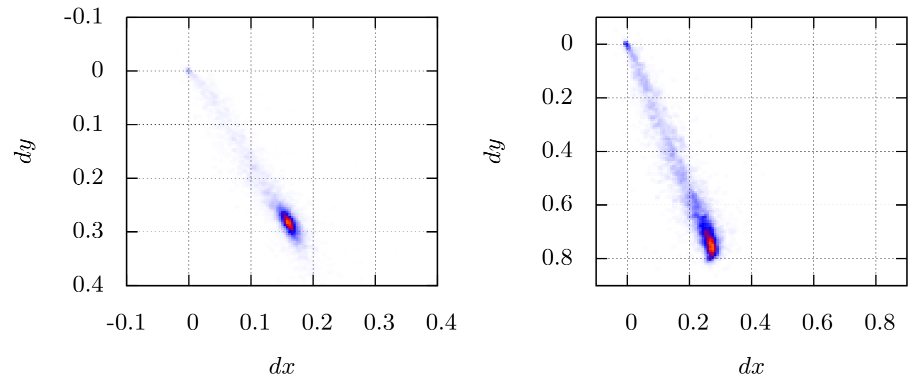

Fig. 3. error diffusion displacement histograms for the raster Floyd-Steinberg (left) and raster Jarvis, Judis and Ninke (right) algorithms applied to a corpus of 10,000 images

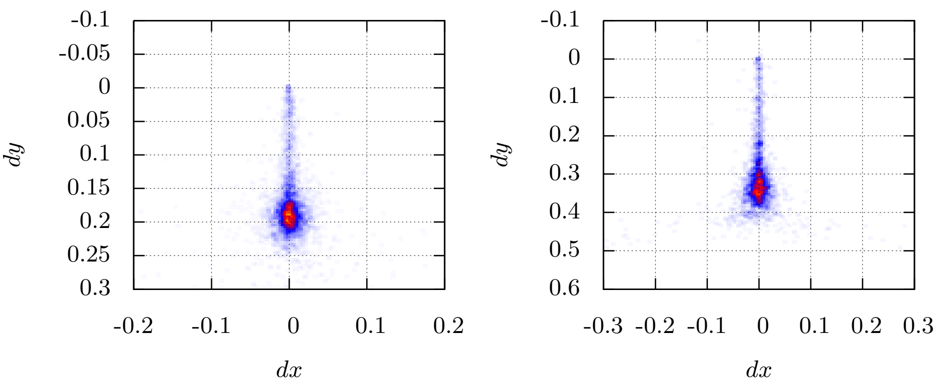

Fig. 4. error diffusion displacement histograms for the Ostromoukhov (left) and optimum kernel (right) algorithms applied to a corpus of 10,000 images

(dx, dy). As E_{fast} does not depend on the image, it is a lot faster to compute than E_{min}, and as it is statistically closer to E_{min}, we can expect it to be a better error estimation than E:

	$E \times 10^4$	$E_{min} \times 10^4$	dx	dy	$E_{fast} \times 10^4$
raster Floyd-Steinberg	3.7902	3.1914	0.16	0.28	3.3447
raster Ja-Ju-Ni	9.7013	6.6349	0.26	0.76	7.5891
Ostromoukhov	4.6892	4.4783	0.00	0.19	4.6117
optimum kernel	7.5209	6.5772	0.00	0.34	6.8233

5 Using error diffusion displacement for optimum kernel design

We believe that our higher quality E_{min} error metric may be useful in kernel design, because it is the very same error that admittedly superior yet computationally expensive algorithms such as DBS try to minimise.

Our first experiment was a study of the Floyd-Steinberg-like 4-block error diffusion kernels. According to the original authors, the coefficients were found "mostly by trial and error" [4]. With our improved metric, we now have the tools to confirm or infirm Floyd and Steinberg's initial choice.

We chose to do an exhaustive study of every $\frac{1}{16}\{a, b, c, d\}$ integer combination. We deliberately chose positive integers whose sum was 16: error diffusion coefficients smaller than zero or adding up to more than 1 are known to be unstable [17], and diffusing less than 100% of the error causes important loss of detail in the shadow and highlight areas of the image.

We studied all possible coefficients on a pool of 3,000 images with an error metric E based on a standard Gaussian HVS model. E_{min} is only given here as an indication and only E was used to elect the best coefficients:

rank	coefficients	$E \times 10^4$	$E_{min} \times 10^4$
1	7 3 6 0	4.65512	3.94217
2	8 3 5 0	4.65834	4.03699
5	7 3 5 1	4.68588	3.79556
18	6 3 5 2	4.91020	3.70465
...	...	...	...

The exact same operation using E_{min} as the decision variable yields very different results. Similarly, E is only given here as an indication:

rank	coefficients	$E_{min} \times 10^4$	$E \times 10^4$
1	6 3 5 2	3.70465	4.91020
2	7 3 5 1	3.79556	4.68588
15	7 3 6 0	3.94217	4.65512
22	8 3 5 0	4.03699	4.65834
...	...	...	...

Our improved metric allowed us to confirm that the original Floyd-Steinberg coefficients were indeed amongst the best possible for raster scan. More importantly, using E as the decision variable may have elected $\frac{1}{16}\{7, 3, 6, 0\}$ or $\frac{1}{16}\{8, 3, 5, 0\}$, which are in fact poor choices.

For serpentine scan, however, our experiment suggests that $\frac{1}{16}\{7, 4, 5, 0\}$ is a better choice than the Floyd-Steinberg coefficients that have nonetheless been widely in use so far (Fig. 5).

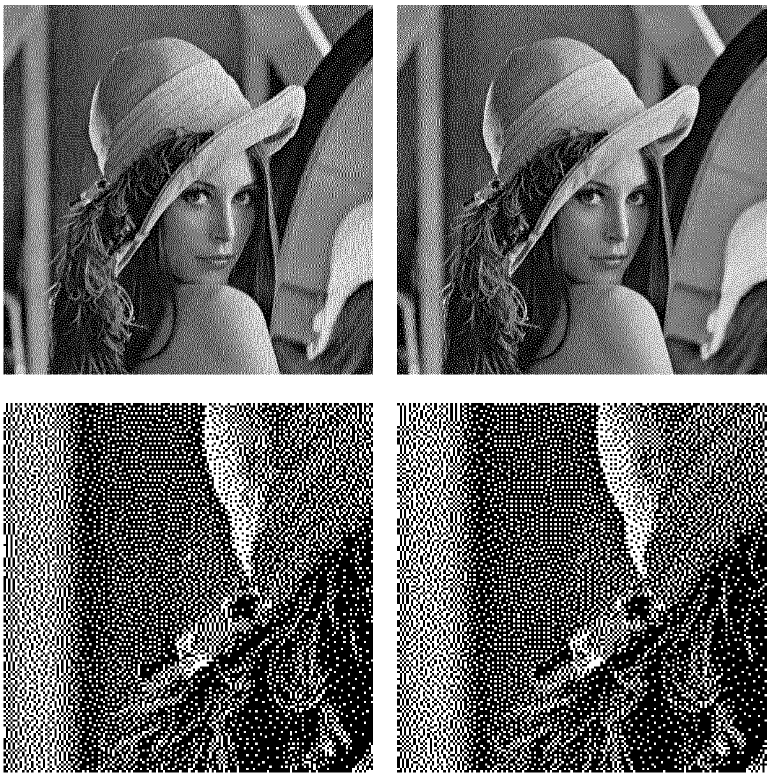

Fig. 5. halftone of *Lena* using serpentine error diffusion (*left*) and the optimum coefficients $\frac{1}{16}\{7, 4, 5, 0\}$ (*right*) that improve on the standard Floyd-Steinberg coefficients in terms of visual quality for the HVS model used in section 3. The detailed area (*bottom*) shows fewer structure artifacts in the regions with low contrast.

6 Conclusion

We have disclosed an interesting property of error diffusion algorithms allowing to more precisely measure the quality of such halftoning methods. Having showed that such quality is often underestimated by usual metrics, we hope to see even more development in simple error diffusion methods.

Confirming Floyd and Steinberg's 30-year old "trial-and-error" result with our work is only the beginning: future work may cover more complex HVS models, for instance by taking into account the angular dependance of the human eye [18]. We plan to use our new metric to improve all error diffusion methods that may require fine-tuning of their propagation coefficients.

References

[1] B. Bayer, *Color imaging array.* U.S. patent 3,971,065 (1976)

[2] R.A. Ulichney (Digital Equipment Corporation), *Void and cluster apparatus and method for generating dither templates.* U.S. patent 5,535,020 (1992)

[3] H. Ancin, A. Bhattacharjya and J. Shu (Seiko Epson Corporation), *Void-and-cluster dither-matrix generation for better half-tone uniformity.* U.S. patent 6,088,512 (1997)

[4] R.W. Floyd, L. Steinberg, *An adaptive algorithm for spatial grey scale.* Proceedings of the Society of Information Display 17, (1976) 75–77

[5] P. Metaxas, *Optimal Parallel Error-Diffusion Dithering.* Color Imaging: Device-Indep. Color, Color Hardcopy, and Graphic Arts IV, Proc. SPIE 3648, 485–494 (1999)

[6] T. D. Kite, *Design and Quality Assessment of Forward and Inverse Error-Diffusion Halftoning Algorithms.* PhD thesis, Dept. of ECE, The University of Texas at Austin, Austin, TX, Aug. 1998

[7] R. Ulichney, *Digital Halftoning.* MIT Press, 1987

[8] L. Velho and J. Gomes, *Digital halftoning with space-filling curves.* Computer Graphics (Proceedings of SIGGRAPH 91), 25(4):81–90, 1991

[9] R. Nasanen, *Visibility of halftone dot textures.* IEEE Trans. Syst. Man. Cyb., vol. 14, no. 6, pp. 920–924, 1984

[10] M. Analoui and J. P. Allebach, *Model-based halftoning using direct binary search.* Proc. of SPIE/IS&T Symp. on Electronic Imaging Science and Tech., February 1992, San Jose, CA, pp. 96–108

[11] Ann McNamara, *Visual Perception in Realistic Image Synthesis.* Computer Graphics Forum, vol. 20, no. 4, pp. 211–224, 2001

[12] Bhatt *et al.*, *Direct Binary Search with Adaptive Search and Swap.* http://www.ima.umn.edu/2004-2005/MM8.1-10.05/activities/Wu-Chai/halftone.pdf

[13] moot, http://www.4chan.org/

[14] P. W. Wong and J. P. Allebach, *Optimum error-diffusion kernel design.* Proc. SPIE Vol. 3018, p. 236–242, 1997

[15] Victor Ostromoukhov, *A Simple and Efficient Error-Diffusion Algorithm.* in Proceedings of SIGGRAPH 2001, in ACM Computer Graphics, Annual Conference Series, pp. 567–572, 2001

[16] T. N. Pappas and D. L. Neuhoff, *Least-squares model-based halftoning.* in Proc. SPIE, Human Vision, Visual Proc., and Digital Display III, San Jose, CA, Feb. 1992, vol. 1666, pp. 165–176

[17] R. Eschbach, Z. Fan, K. T. Knox and G. Marcu, *Threshold Modulation and Stability in Error Diffusion.* in Signal Processing Magazine, IEEE, July 2003, vol. 20, issue 4, pp. 39–50

[18] J. Sullivan, R. Miller and G. Pios, *Image halftoning using a visual model in error diffusion.* J. Opt. Soc. Am. A, vol. 10, pp. 1714–1724, Aug. 1993