Development of a Texture Model Based Adaptive Halftone Algorithm

Jennifer Alford
University of Iowa
and Theophano Mitsa
GE Medical Systems

Abstract

We develop an adaptive halftoning technique that incorporates a multichannel texture model. A filter bank is used to define a feature vector for each pixel. The feature vector is used to select an error constraint that will preserve neighborhood features. Results indicate that perceptually significant features can be enhanced while preserving smoothness.

Introduction

It is generally accepted that the best quality halftones are those that result in an error spectrum that greatly attenuates low frequency error and allows uniformly distributed error in the high frequencies. This is described as a blue noise distribution.[1] Examples of algorithms that achieve this vary from Floyd and Steinberg's error diffusion,[2] to dithering with the blue noise mask,[3] to iterative techniques that explicitly shape the error with a visual model based error constraint.[4] These all provide a pleasing balance between grayscale reproduction and edge rendering. Unfortunately, allowing unconstrained errors in the high frequency can compromise image quality because, for some image content, it is the high frequency information that is of most significance to the observer. For example, in order to discern if a house exterior is of wood or shingles, it is necessary to see the cracks separating the shingles. Several researchers have tried to resolve this by incorporating image dependent schemes that enhance edges without sacrificing the pleasing blue noise properties.[5,6] The improvements achieved with these approaches suggest that adapting halftoning techniques to local image content is a promising direction of research.

The work presented here seeks to improve halftone quality by identifying and preserving distinctive regions of texture in an image. The assumption is that, for some image content, the high frequency information is the most perceptually significant and that constraining the high frequency error will improve perceived image quality. Previous work used the image error spectrum to adapt error constraints for images of uniform texture.[7] This work extends the concept of adapting error constraints to general images. This is achieved by a three stage process. First a preprocessing stage uses a multichannel model of human texture perception to define feature vectors for each pixel. The feature vectors contain information about the perceptually significant component at each pixel. Second, an evaluation stage interprets the feature vector to determine an error metric that will preserve the significant perceptual components at each pixel. Third, an iterative halftoning stage minimizes the error constraint that has been defined for each pixel. An overview of the system can be seen in Figure 1.

Texture Model

Much work supports the use of a multichannel filter bank as the mechanism by which texture is perceived by humans.[8,9] Multichannel filter banks have been used for many texture discrimination problems in image processing. The specific analytical form of the model varies among implementations but they all have several features in common. They all have a set of bandpass filters that are tuned to different radial frequencies and orientations. In all cases, the output of the filters provides feature information that is used to discriminate texture. How the channel output is used varies.

For this study, we choose the multichannel filters to be a set of M filters used by Coggins and Jain for texture segmentation.[10] These are a bank of radially symmetric filters tuned to different frequencies and a bank of four angular filters. The radial filters, defined by Ginsburg,[11]
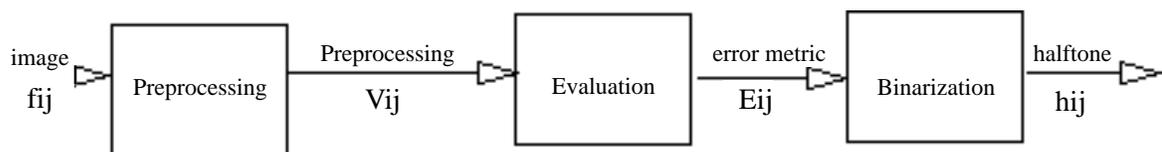

Figure 1. Overview of the three stages of processing. Preprocessing incorporates a multichannel texture model embedded in a contrast normalization model to detect perceptually significant information. The evaluation stage uses the feature vector at each pixel to define an error constraint adapted to neighborhood information. The binarization stage uses an iterative procedure to produce a halftone.

are Gaussian function with center frequencies one octave apart on a log scale. The extent of the filters are between 1 and 2 octaves. The number of filters is M = $\log_2 N + 1$ for a 2N × 2N region. Specifically, the *m*th channel filter is defined in the frequency domain as

$$H_k(f_r) = e^{-0.5\frac{(\ln(f_r)-\ln(u_k))^2}{\sigma^2}} \qquad (1)$$

where f_r is radial frequency, $\sigma = 0.275$, and $u_k = 2^{k-1}$. The angular filter is defined as

$$G_\theta(u,v) = e^{-0.5\frac{A(\theta,u,v)^2}{\sigma^2}} \qquad (2)$$

where $\sigma = 17.8533$ and $A(\theta,u,v) = \min[|\theta - \arctan(v/u)|, |\theta - 180) - \arctan(v/u)|]$, $\theta = 0°,45°,90°,135°$.

This filterbank is part of a larger model of contrast normalization proposed by Heeger and shown in Figure 2, taken from [12]. This model suggests normalizing the output of each channel at a single location with pooled neighborhood information. The N × N input image is filtered by the bank of M filters. In the spatial domain, the magnitude of the deviation about the mean at a single pixel indicates the relative contribution of that energy band at that point relative to the other spatial locations in the image. Each channel output is standardized to zero mean and unit variance. The transformation of each channel effectively boosts the high frequencies. Then the output of each channel is modified using pooled information from all channels at an individual pixel. We deviate from the contrast normalization model by dividing by the standard deviation of the values rather then dividing by the total energy.
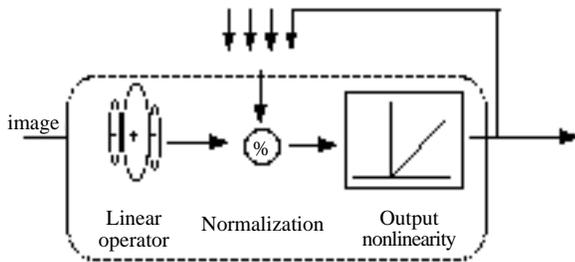

Figure 2. Model of contrast normalization that produces the final feature vector.

Evaluation of the Feature Vector

The input to the evaluation stage is the set of N × N feature vectors, V_{ij}, each with M elements, defined by preprocessing. The goal is to define an error metric at each pixel that will preserve high frequency information without sacrificing the desirable attenuation of low frequencies achieved with a global low-pass visual model. In order to avoid introduction of grainy artifacts in lower frequencies, we decided on a simple approach of defining a set of *k* possible error constraints, denoted e_k, that are all high pass in nature but that vary in dropoff

frequency and slope. Shifting the drop off in this way allows suppression of high frequency errors without sacrificing blue noise characteristics. Which of the *k* error constraints to apply is determined by comparing the values of the feature vector to a predefind threshold *t*.

We found that comparing the $V_{ij}[1]$ to *t* was sufficient to identify areas of fine detail that needed to be enhanced from other parts of the image that were best rendered with a smoother constraint. The interpretation is that because each feature vector is standardized within itself, each element contains information about the distribution of the other elements relative to itself. $V_{ij}[1]$ represents low frequency information. When its value falls below the threshold, its energy is relatively insignificant compared to the distribution of energy in the other, higher frequency channels and the higher frequencies should be preserved. When it is above the threshold, its energy is more perceptually significant than the distribution of energy in the other, higher frequency channels and thus more error in the high frequencies can be tolerated. This is illustrated with an example in Figure 3. Figure 3a shows the values of $V_{ij}[1]$ for every pixel in the image. Figure 3b shows which pixels for which e_1 was selected. Constraint e_2 was applied to all other pixels.

Figure 3. (a) First element of the feature vector for all pixels. (b) Pixels selected from error constraint e_1 with t = 0.5.

Binarization

The final stage takes as input the original grayscale image and the N × N error constraints to produce a halftone. This is implemented as an iterative procedure based heavily on the work of Pappas[4] because of its simplicity and its sensitivity to the error metric. Specifically, the algorithm developed will test the effect of changing only one pixel at a time and it will compute the error in the spatial domain over a neighborhood of the candidate pixel. The algorithm that will generate a halftone according to the error metric developed in this study.

1. Initialize binary image with a simple thresholding method.
2. For each pixel (*i*,*j*) in the image
   Select binary value that minimizes the error metric at pixel (*i*,*j*)
3. Repeat (2) until convergence

Let E(h(*i*,*j*)) designate the error metric computed at pixel (*i*,*j*) and let h((i,j) indicate the possible values of

the output pixel h(i,j). Then E(0) represents the error metric computed when h(i,j)=0 and E(1) represents the error metric computed when h(i,j)=1. The algorithm will select for output the value that minimizes E(h(i,j)).

Results

The algorithm described above was implemented in C on an HP 715/80 workstation. The multichannel filterbank was implemented with parameters $s = 0.275$ and without angular sensitivity. The set of error constraints were computed as low-pass modifications to a contrast sensitivity function described in previous work.[3] Constraint e_1 was generated with viewing distance of 36 inches and a print resolution of 150dpi. This filter produces smoothing grayscale reproduction. Constraint e_2 was generated with viewing distance of 36 inches and a print resolution of 300dpi. This filter produces edge enhancing effects The constraints were realized as 19th order FIR filters. Data was a set of 256×256 8bpp grayscale images with varying subjects including people, landscapes, and textures. The images were printed at 150dpi.

Figure 4. a) input grayscale, (b)upper right, e1 applied globally, (c) lower left, e_2 applied globally, (d) adaptive error constraints.

Two examples typical of the results are shown in Figures 4. Figures 4a depicts the input grayscale image. For comparison with non-adaptive constraints, halftones were generated in which the same error constraint was applied at every location in the image. Figure 4b is the halftone that is generated when a visual model based error metric e_1 is used for every location in the image. This represents the widely accepted balance between smooth grayscale reproduction and edge rendering that is produced by minimizing the error with respect to a visual model. This is a good quality image that is very

smooth but that has dulled details. Observe the smooth rendition of the leaf in the lower left corner. Figure 4c is the halftone that is generated when error constraint e_2 is applied at every location in the image. This is also a good quality image that is very sharp but that is somewhat grainy. This is best observed by noting the enhanced detail of the ribbing in the dark petals and the graininess in the bright petals. Figure 4d is the product of the adaptive technique that detects the areas that need enhancement but that leaves the other regions smooth. Note, particularly, that the detailed ribbing in the petals is now apparent without introduction of graininess in the bright petals.

Figure 5 similarly represents another picture. Figure 5a is the input grayscale. Figures 5b and 5c are halftones generated when constraints e_1 and e_2, respectively, are applied globally. Figure 5d is the result of the adaptive halftone algorithm described here. Note the sharpness in the feathers and hair from Figure 5c are apparent with the smoothness of the face from Figure 5b.
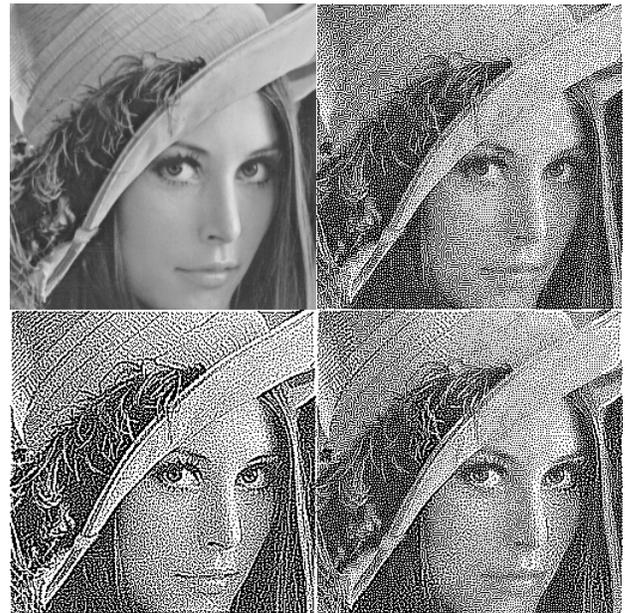

Figure 5. a) input grayscale, (b)upper right, e_1 applied globally, (c) lower left, e_2 applied globally, (d) adaptive error constraints.

Discussion

Adaptive halftone constraints based on image content can improve halfone quality. The texture model used here to identify significant high frequency distributions in neighborhoods of the image clearly helps select the portions of the image that benefit from edge enhancement while allowing the rest of the image to be rendered more smoothly. The implementation described here was a much simpler model than could be used. Future work will include adding angular sensitivity to expand the bank of filters. It will also include a larger set of error constraints that will refine the optimum constraint. Another direction of experimentation will be to vary the method

of evaluation to test weighted combinations of channel outputs against a threshold. There is much to be studied in order to tune the model for optimal performance. However, the results presented here indicate the contrast normalization model to be very promising as a means to adapt halftone error constraints.

Acknowledgements

References

1. R. Ulichney, *Digital Halftoning* (MIT Press, Cambridge, Massachusets, 1987).
2. R. W. Floyd and L. Steinberg, *Proceedings of the SID* **17**, 75–77 (1976).
3. T. Mitsa, K. Varkur, and J. R. Alford, in *SPIE Conference on Electronic Imaging* (1993).
4. T. N. Pappas and D. L. Neuhoff, in *SPIE/IS&T Symposium on Electronic Imaging Science and Technology, Human Vision, Visual Proc., and Digital Display III* 165–176 (San Jose, California, 1992).
5. J. Sullivan, R. Miller, and G. Pios, *Journal of the Optical Society of America A* **10**, 1714–1724 (1993).
6. R. Eschbach and K. Knox, *Journal of the Optical Society of America A* **8**, 1844–1850 (1991).
7. J. R. Alford and T. Mitsa, in *ICIP* (Washington, D. C., 1995).
8. B. A. Wandell, *Foundations of Vision* 1-476 (Sinauer Associates, Inc., Sunderland, Mass, 1995).
9. Reed, T. R. & Du Buf, J. M. H. CVGIP: *Image Understanding* **57**, 359–372 (1993).
10. J. M. Coggins and A. K. Jain, *Pattern Recognition Letters* **3**, 195–203 (1985).
11. A. P. Ginsburg, (University of Cambridge, 1977).
12. D. J. Heeger, E. P. Simoncelli, and J. A. Movshon, in *Proceedings of the National Academy of Science* 623–627 (Irvine, CA, 1996).

❋ Previously published in *IS&T/OSA Optics & Imaging in the Information Age*, pp. 195–198, 1996.