PERCEPTUALLY OPTIMAL RESTORATION OF IMAGES WITH STACK FILTERS

J.J. Huang and E.J. Coyle

School of Electrical and Computer Engineering
1285 EE Bldg., Purdue University
West Lafayette, IN 47907-1285

ABSTRACT

The present approach to the MAE-based design of stack filters for image restoration does not always produce the desired visual result. Thus, in this paper, a new stack filter design algorithm is developed. It is based upon a Weighted Mean Absolute Error (WMAE) criterion instead of the traditional MAE criterion, which assigns the same weights to all errors. The weights in this WMAE criterion are designed with the aid of the Visible Differences Predictor (VDP), which can estimate the sensitivity of the human visual system to changes in images. Experiments with this WMAE approach show that the stack filters it produces perform significantly better in image processing applications than those designed with the MAE approach.

1. INTRODUCTION

Stack filters are a class of discrete-time, nonlinear filters that have been developed by several authors [1, 2, 3]. The adaptive algorithms [4, 5, 6] for determining a stack filter which minimizes the mean absolute error criterion have been developed and successfully applied to problems such as edge detection and noise reduction in images.

One difficulty with this theory of adaptive stack filtering is that it yields filters that do not always produce the desired visual result in image processing applications. The hypothesis motivating this paper is that this difficulty is due to the error criterion that is used, not to fundamental limitations of stack filters. We claim that the MAE criterion does not assign the same level of significance as the human visual system to certain types of noise artifacts.

A new stack filter design algorithm is therefore proposed. It is based upon a Weighted Mean Absolute Error (WMAE) criterion instead of the MAE criterion, which assigns the same weight to all errors. The weights in the WMAE criterion are designed with the aid of the Visible Differences Predictor (VDP) [7], which can estimate the sensitivity of the human visual system to changes in images.

Experiments with this new approach show that the stack filters it produces perform significantly better in image processing applications than those designed with the MAE approach. They yield a much better tradeoff between noise reduction and detail preservation than all other approaches we have investigated. More significantly, the WMAE criterion produced by the algorithm and the stack filters it yields work well even when the images and noise types to which they are applied are significantly different than those used to train them.

2. REVIEW OF THE VISIBLE DIFFERENCES PREDICATOR (VDP) ALGORITHM

The visible differences predictor described in [7] is an algorithm for the assessment of image fidelity. The goal of the VDP is to determine the degree to which those image distortions become visible differences. Two images, one noise free and the other distorted, and parameters for viewing conditions and calibration are the input to the algorithm. The output image is a map of the probability of detecting the differences between the two images as a function of their location in the images.

In our research, we employ a simplified version of the VDP. The first reason for simplification is that we will compute the VDP many times during the execution of our training algorithm. A fast version is therefore critical. The second, more significant reason for using a simplified VDP algorithm is to obtain better localization of the noise corrupting the images we are considering. The noise types we use include impulsive noise of arbitrary distribution and line drop-out noise. Both have significant high-frequency content. When present they yield a VDP image between the noisy and desired images that exhibits a great deal of ringing and, consequently, is almost useless for locating the noise in the image.

To solve these problems of complexity and localization, we use only the highest spatial-frequency channels in the VDP algorithm. Finally, we need a single number that can be used as a measure of perceptual error when we train the WMAE criterion. The output of the VDP, though, is an array of numbers indicating the probability of detection of differences at each pixel location. We obtain our perceptual error measure from this array by thresholding it and then summing all the entries of the resulting threshold map.

3. MINIMUM WMAE STACK FILTERING AND THE VDP ALGORITHM

One difficulty with the theory of minimum MAE stack filtering is that the filters it produces do not always yield the desired visual result in image processing. The possible remedies for this problem include choosing a different error criterion entirely and modifying the mean absolute error criterion. The latter approach is the one taken in this paper since it still allows the optimization algorithms developed for minimum MAE stack filtering to be used, albeit with some modifications. Whether this is the correct approach can be determined by specifying the modification to be made and determining through experiments whether it produces the desired results.

In this new approach the filtering of the image takes place in two stages. In the first stage, the goal is to remove noise that is positive-going; the second stage removes negative-going noise. Smaller windows can be used for each of these procedures than would be required if the filter were to remove noise of both signs simultaneously. To achieve the desired visual effects, the weights in the mean absolute error criterion used to design the two filters are modified so the criterion more closely matches a perceptual error criterion.

The modifications made to the error criterion concern the weights assigned to the errors made for each observation vector. In the (unweighted) MAE criterion, all errors are assigned an equal weight of 1.0. There is complete freedom, though, in how these weights are assigned. We have chosen to allow each weight to take any value between 0 and 1, and to select the weights independently for each type of error. Exploiting this freedom does not, however, yield stack filters with significantly better performance. What is also needed is a two-stage approach to filtering. In the first stage, the goal is to remove positive-going noise; in the second, the goal is to remove negative going noise. In each stage, the error weights are modified so that errors in one direction are very heavily penalized.

Consider the first stage filter, whose goal is to suppress positive-going noise. The weights corresponding to positive-going errors are chosen to be less than 1.0, while those for negative-going errors are left equal to 1.0. If the difference between these two weights is large enough, the resulting filter almost completely suppresses positive-going noise impulses. If the difference is too large, though, the filter designed with this error criterion will introduce more negative-going noise.

For example, for the image of Einstein with 20% two-sided, additive impulsive noise of amplitude 200, the choice of 0.9 for the lowered weight does not produce a filter which eliminates all positive-going noise. Choosing the weight to be 0.15 will lead to complete removal of positive-going noise, but also causes more negative-going noise to be generated. We have found that setting the lowered weight to be 0.25 achieves the desired visual result. There thus exists an optimal value for the lowered weight.

Once the first stage filter has been designed, the second stage filter is designed to remove the negative-going noise from the output of the first-stage filter. The technique used to choose the weights in the error criterion for this filter is the same as for the first-stage filter.

By using a two-stage filtering procedure, and modifying the error weights for each stage, we have found that filters of smaller window size can be used while still achieving dramatically better visual results. A cascade of two 4x4 filters designed this new way can be trained and applied to an image in 2 minutes on a Sparc 5, and the resulting filter will outperform one-stage stack filters with the largest window for which training is feasible (5×5). No loss of robustness occurs when the two-stage procedure is used.

The weights in this Weighted Mean Absolute Error (WMAE) criterion are designed with the aid of the Visible Differences Predictor (VDP), which can estimate the sensitivity of the human visual system to changes in images. The VDP algorithm feeds back the visual error to the two-stage stack filtering algorithm to adjust the weights until the filter reached minimizes the VDP based visual error criterion.

4. EXPERIMENTAL RESULTS

In this section, the filtering behavior of the new cascaded filter is examined. The results are obtained by variation of the error weights. Fig. 1(b) shoes the convergence behavior of the visual error metric when the full training algorithm described above is executed. The experimental results show that visual error converges to the global minimum. Note that there are several local maxima in the graphs. They occur when a reduction in the weights in the WMAE causes the filter to destroy more details even though it is removing more noise. This phenomenon can be observed in the sequence of filtered images resulting from the algorithm.

Fig. 1(a) shows the noisy image that will be the target of the filtering algorithm. Fig. 1(c) shows the effects of applying the original algorithm which finds a stack filter that minimizes the unweighted mean absolute error criterion. As expected, many impulsive and strip type errors remain in the output image.

Fig. 1(d) shows the result of applying the new algorithm with the perceptually optimal Weighted Mean Absolute Error (WMAE) criterion. Essentially all of the noise is gone, despite the high probability of its occurrence. The cost is some loss of image detail when compared with the output of the MMAE stack filtering technique. The detail that is lost, though, is very tolerable when compared with the visual effects of impulses remaining in the image.

We have applied the WMAE stack filter trained with the image of Einstein to other images in order to test the robustness of the filter. We found that training with an image which has been made symmetric by concatenating the original image and its left-right and top-bottom flipped versions produces the best – i.e., most robust – filters. When symmetrically trained filters are applied to images for which they were not trained, they usually yield lower MAE and WMAE than filters that were not symmetrically trained.

The new algorithms developed in this paper have found important applications in restoration of images captured with defective CCD and CMOS imaging arrays. For example, these filters do an excellent job of compensating for the white-spots produced by device defects in CMOS imaging arrays [8]. In this latter case, we used a one-stage WMAE stack filter because the noise is only positive-going. Because space precludes inclusion of these applications in this paper, the images showing the effectiveness of the WMAE stack filters in these applications can be found at
http://yake.ecn.purdue.edu/~jrjen.

5. CONCLUSION

In this paper, we presented a minimum WMAE stack filtering algorithm that uses the visible differences predictor to achieve better performance in image processing applications than that achieved by the minimum MAE stack filtering approach. The new algorithm retains the iterative nature of the present adaptive minimum MAE algorithm, but it allows the weights of the error criterion to vary during the training process. The variation of these weights is guided by the perceptual error measure that is based on the VDP algorithm. Through the experimental results, we verified that the visual error of the new training algorithm decreases to a minimum, at which point the algorithm produces a stack filter which is optimal in both the WMAE sense and the perceptual sense.

6. REFERENCES

[1] P. D. Wendt, E. J. Coyle, and N. C. Gallagher, Jr. "Stack filters," IEEE Trans. Acoustics, Speech, and Signal Proc., 34:898–911, August 1986.

[2] J. P. Fitch, E. J. Coyle, and N. C. Gallagher, Jr. "Median filtering by threshold decomposition," IEEE Trans. Acoustics, Speech, and Signal Proc., 32:1183–1188, December 1984.

[3] J. P. Fitch, E. J. Coyle, and N. C. Gallagher, Jr. "Threshold decomposition of multidimensional rank order operators," IEEE Trans. Circuits Syst., 32:445–450, May 1985.

[4] J.-H. Lin and E.J. Coyle. "Minimum mean absolute error estimation over the class of generalized stack filters," IEEE Trans. Acoustics, Speech, and Signal Proc., 38:663–678, April 1990.

[5] Y. T. Kim and J.-H. Lin. "Fast training algorithm for stack filters," submitted to the IEEE Trans. Signal Processing.

[6] J. Yoo, K. L. Fong, J.-J. Huang, E. J. Coyle, and G. B. Adam III. "Fast algorithm for designing stack filters," submitted to the IEEE Trans. on Image Processing.

[7] S. Daly. "The visible difference predictor: an algorithm for the assessment of image fidelity," Chapter 14 in Digital Images and Human Vision, ed. by A. B. Watson, MIT Press, 1993.

[8] S. S. Mahant-Shetti and D. A. Martin. "Non-linear filters for whitespot compensation," IEEE/EURASIP Workshop on Nonlinear Signal and Image Processing, 1997.

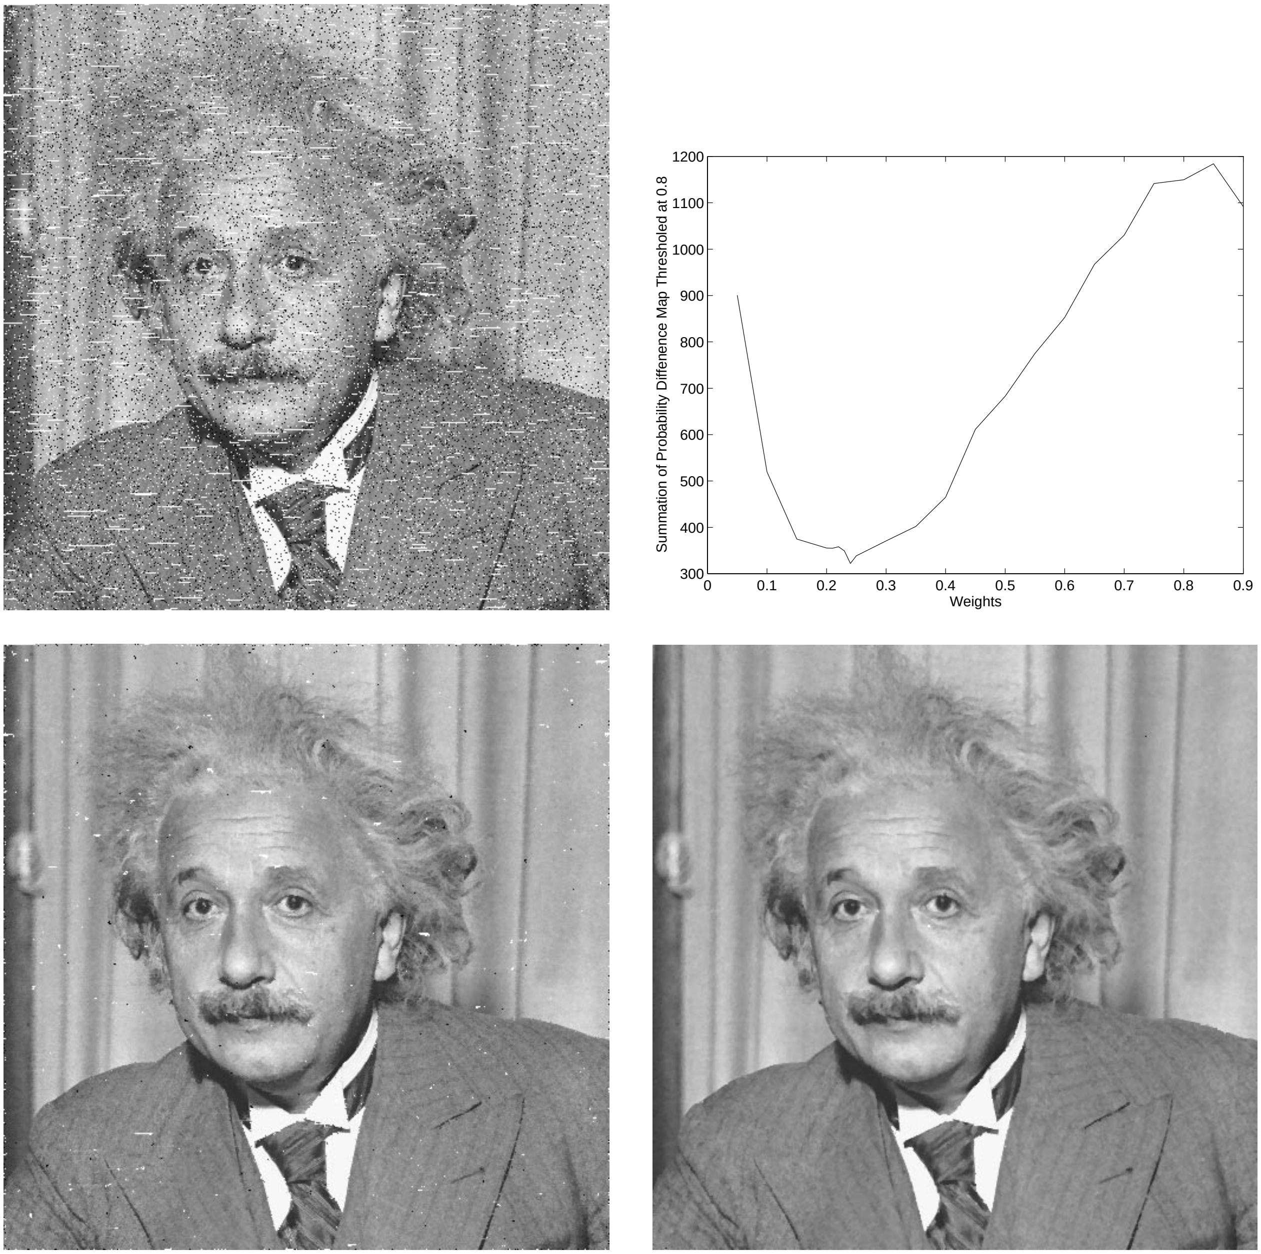

(a)	(b)
(c)	(d)

Figure 1: (a) Einstein plus 10% two-sided, additive impulsive noise with amplitude 200, and with line drop-outs of average length 5 and occurrence probability .005. (b) Convergence behavior of visual error metric of 4 x 4weighted stack filter with VDP algorithm applied to image (a). (c) Output of stack filter designed using original MMAE method. (d) Output of stack filtering operation designed with new Weighted MAE method.