ADAPTIVE WEIGHTED HIGHPASS FILTERS USING MULTISCALE ANALYSIS

Robert D. Nowak, Member, IEEE, and Richard G. Baraniuk, Member, IEEE *

Department of Electrical and Computer Engineering
Rice University
6100 South Main Street
Houston, TX 77005–1892
E-mail: nowak@ece.rice.edu, richb@rice.edu
Fax: (713) 524–5237

Submitted to *IEEE Transactions on Image Processing*, February 1996
Revised: December 1996, July 1997

Abstract

In this paper, we propose a general framework for studying a class of weighted highpass filters. Our framework, based on a multiscale signal decomposition, allows us to study a wide class of filters and to assess the merits of each. We derive an automatic procedure to tune a filter to the local structure of the image under consideration. The entire algorithm is fully automatic and requires no parameter specification from the user. Several simulations demonstrate the efficacy of the proposed algorithm.

*This work was supported by the National Science Foundation, grant no. MIP 94–57438, and by the Office of Naval Research, grant no. N00014–95–1–0849.

1 Introduction

Recognition of image features depends on the local level and contrast in the neighborhood of the feature. One of the primary steps in recognition is *edge* or *boundary extraction*. To aid in this task it is often desirable to enhance the image detail and edges using a highpass filtering scheme. Unfortunately, highpass filtering also amplifies noise present in the image.

The local intensity affects the eye's sensitivity to noise in images. Specifically, the visual system is much less sensitive to noise in bright areas of an image than it is in dark areas. This observation is commonly referred to as Weber's Law [2]. In view of Weber's Law, an image enhancement filter can avoid degrading noise amplification by sharpening dark regions of an image less than bright regions. One very simple method to accomplish this is to weight the amount of highpass filtering proportional to the local mean. This gives rise to a class of nonlinear image enhancement filters known as *mean-weighted highpass filters* [4, 9].

Empirical evidence also suggests that the visual system is less sensitive to noise in the edges or highly structured regions of an image. This effect is known as *masking by structure* [7]. The masking effect implies that noise amplification due to highpass filtering is less noticable in highly structured areas of an image. Therefore, a reasonable approach to improve highpass filtering enhancement is to weight the output of the highpass filter proportional to the output of a local edge detector. This idea has led to nonlinear *edge-weighted highpass filters* [1, 8].

One limitation of existing weighted highpass filters is that the filter structure is fixed. This means that the *scale* of the local mean or edge detector is fixed. Hence, the user must specify a local neighborhood for the mean or explicitly define what is meant by a local edge. Also, these algorithms typically require user specified weighting parameters and often threshold the nonlinear highpass image in an ad hoc fashion.

In this paper, we propose a general framework for studying the class of weighted highpass filters based on a multiscale signal decomposition. In order to find the best weighted highpass filter for a given image, we project a linear highpass filtered version of the image onto a subspace of multiscale weighted highpass filtered images. Each weighted highpass filtered image in the subspace provides a degree of enhancement tempered by the suppression of the highpass amplification in dark or homogeneous regions of the original image. Projecting the linear highpass filtered version onto this subspace produces a linear combination of weighted highpass filtered images that match the important image details, while suppressing excessive noise amplification in dark or homogeneous regions. In effect, this design method produces an adaptively weighted highpass filtered image that balances the trade-off between enhancement and noise amplification.

The paper is organized as follows. In Section 2, we review previous work on weighted highpass filters and discuss some of the limitations of existing methods. We also give a brief review of multiscale analysis. Section 3 introduces a novel weighted highpass filter based on multiscale analysis. Several simulations demonstrate the efficacy of the proposed filter in Section 4. Conclusions are drawn in Section 5.

2 Previous Work

2.1 Unsharp Masking: A standard method of image enhancement is *unsharp masking* [2, p. 249]. In unsharp masking, the original image is enhanced by subtracting a signal proportional to a smoother version of the original image. Equivalently, a signal proportional to a highpass filtered version of the original image can be added to the original. Let H denote a linear highpass filter, let $f(x, y)$ be an image, and consider the enhanced image $g = f + Hf$. Adding the highpass filtered image to the original enhances or emphasizes edges and detail in the image. Alternatively, suppose we have a blurred image f and a linear restoration filter R. We may consider the difference between f and the restored Rf as a highpass filter, that is, $Hf = Rf - f$. With this notation, linear deblurring can also be viewed as a form of unsharp masking.

2.2 Weighted Highpass Filters: The enhanced or restored image g may be undesirable if noise in the original image f is amplified by H. Weber's Law and the masking effect [2] suggest the following nonlinear approach to image enhancement. Let L denote a linear filter that is tuned to a specific type of local image feature. By "local" we mean that the output image Lf at the point (x, y) depends only on the local neighborhood of f about (x, y). By "tuned" we mean that $|Lf(x, y)|$ is large if a local image feature, such as an edge or region of high intensity (high local mean), is near (x, y) in f. A weighted highpass filter is defined by the mapping

$$f(x, y) \;\mapsto\; H_{\mathrm{W}}f(x, y) \;\stackrel{\mathrm{def}}{=}\; |Lf(x, y)|^p \cdot Hf(x, y), \quad p \geq 1. \tag{1}$$

Here, $|Lf|^p$ is the image formed by raising every point $Lf(x, y)$ in the image Lf to the p-th power. The image $|Lf|^p$ "weights" the highpass filtered image Hf pointwise according to the strength of the local features associated with L. For instance, if L corresponds to a local mean, then $H_{\mathrm{W}}f$ is roughly proportional the output image obtained by applying H only in regions with high local mean [4, 9]. If L is a local edge-detector, then $H_{\mathrm{W}}f$ is proportional to the output image obtained by applying H only in regions where an edge is detected [1, 8].

2.3 Limitations of Previous Work: One important drawback to the mean-weighted and edge-weighted filters previously studied in [1, 4, 8, 9] is that the filter scale is fixed. Hence, such filters may only be appropriate for image detail at a fixed scale. Our idea is to wed the ideas of multiscale analysis and weighted highpass filters to produce an adaptive filter that automatically adjusts to the local detail of the image at hand. Before discussing our method, we briefly review the multiscale analysis of images.

2.4 Signal Characterization Using Multiscale Edges: The notion of multiscale signal analysis is motivated by the need to detect and characterize the edges of small and large objects alike. In an image, different structures give rise to edges at varying scales — small scales correspond to fine detail and large scales correspond to gross structure. In order to detect all image edges, one must study the image at each scale. Multi-scale image processing tools include scale space, pyramid algorithms, and wavelet transforms.

In this paper, we will follow the approach of Mallat and Zhong, who use the scales of a separable wavelet transform to characterize the important edges in an image (see [3] for more information on the wavelet transform). Consider first the analysis of continuous images. To analyze such images, we employ a smoothing function ϕ, a wavelet function ψ, and an infinite number of scales. The functions ϕ and ψ proposed in [3] are depicted in Fig. 1.

Smoothed versions of the image f are obtained by convolution with ϕ in both x and y directions. Larger scales (smoother images) are obtained by dilating ϕ. Dilation of ϕ by factors of two halves the resolution each time as we move up through scales. We denote the smoothed image at scale 2^j by $S_{2^j} f$. Note that $S_1 f = f$.

Edge and detail information in f is obtained by convolution with ψ. Detail information at larger scales is obtained by dilating ψ. At scale 2^j we have three detail images: $W^h_{2^j} f$, $W^v_{2^j} f$, and $W^d_{2^j} f$, where the superscripts h, v, and d denote the horizontal, vertical, and diagonal (both horizontal and vertical) applications of ψ, respectively.

To analyze discrete images, we use an undecimated two-channel filterbank with discrete analysis filters h and g and a range of scales J limited by the number of pixels in the image. In general, $2^J \leq N$ for an $N \times N$ image. (See [3] for more information on the discrete wavelet transform. In particular, Tables I and II in [3] provide the filters h and g corresponding to ϕ and ψ of Fig. 1.

In [3] it is shown that the modulus maxima of the wavelet transform provide a nearly complete characterization of an image. Mallat and Zhong characterize the image edges at scale 2^j by the local maxima of

$$M_{2^j} f(x,y) = \sqrt{|W^h_{2^j} f(x,y)|^2 + |W^v_{2^j} f(x,y)|^2}. \tag{2}$$

3 AdaptiveWeighted Highpass Filters

In this section, we utilize local edge and local mean information carried by the smooth and detail images at varying scales to develop a class of weighted highpass filters. Our goal is to choose the best weighted highpass filter for a given image.

3.1 Multiscale Mean-Weighted Filters: We can easily formulate the mean-weighted highpass filter in the multiscale framework. Pointwise multiplication of the highpass image Hf with $|S_{2^j} f|^p$ yields a $p+1$st order[1] weighted highpass filter with response strongest in regions where the local mean (at the scale 2^j) is large. Adjusting the scale 2^j is equivalent to adjusting the size of the local neighborhood used to compute the mean. We thus have the following collection of mean-weighted highpass filtered images:

$$\{\, |S_{2^j} f|^p \cdot Hf \;:\; j = 1, \ldots, J, \;\; p = 1, \ldots, P \,\}. \tag{3}$$

The exponent p controls the relative weighting in light and dark regions; increasing p tends to emphasize areas of peak intensity. The scale bound J limits the range of scales used for local feature detection. J acts as a regularization parameter: a small value of J gives maximum regularization by focusing the filters on only very local features, while a large value allows the filters to incorporate more global, gross structure at the expense of less regularization. In practice, the choice of J is problem-dependent, but prior information may suggest a resonable choice depending on which types of features are dominant in the image under study. Experience has shown (see Section 4 below) that reasonable values for J lie in the range $1 \leq J \leq 4$.

3.2 Multiscale Edge-Weighted Filters: We define the detail modulus as

$$|D_{2^j} f(x,y)| = \sqrt{|W^h_{2^j} f(x,y)|^2 + |W^v_{2^j} f(x,y)|^2 + |W^d_{2^j} f(x,y)|^2}. \tag{4}$$

[1] The pointwise product of a linear filtered image and a pth order filtered image is $p+1$st order polynomial in the data.

Our experiments have shown that $|D_{2^j} f|$ provides better results for our application than $M_{2^j} f$ from (2), possibly because it treats edges at different orientations more fairly. Pointwise multiplication of the highpass image Hf with $|D_{2^j} f|^p$ produces a $p+1$st-order weighted highpass filtered image tuned to edges at the scale 2^j — an edge-weighted highpass filtered image. The multiscale analysis produces a set of edge-weighted highpass filters; each is tuned to edges at a prescribed scale:

$$\{\, |D_{2^j} f|^p \cdot Hf \;:\; j = 1, \ldots, J,\;\; p = 1, \ldots, P \,\}. \tag{5}$$

Increasing the exponent p tends to localize the weighting to areas where the detail modulus is large.

3.3 Adaptive Filter Design: Multiscale analysis provides a suite of weighted highpass filters, (3) and (5), suitable for image enhancement. The question now becomes: Which one is best for a given image? Even more generally, we may consider the collection of filtered versions of f

$$\mathcal{C}_f = \left\{ \sum_{j=1}^{J} \sum_{p=1}^{P} \alpha_{j,p} |D_{2^j} f|^p \cdot Hf \;+\; \beta_{j,p} |S_{2^j} f|^p \cdot Hf \right\}, \tag{6}$$

with arbitrary real coefficients $\{\alpha_{j,p}, \beta_{j,p}\}$. The collection $\mathcal{C}_f$ is simply the subspace of weighted highpass filtered images spanned by (3) and (5). The collection $\mathcal{C}_f$ is quite general. In particular, it can model any nonlinear filter scheme involving polynomial combinations of the original image pixels.

We now propose an automatic procedure for choosing approximating the optimal filtered image in $\mathcal{C}_f$ for a given image. The idea is very straightforward. By design, all of the filtered images in $\mathcal{C}_f$ are highpass enhanced yet also suppress noise in smooth or low intensity regions. However, each of these enhanced images was obtained using filters tuned to structures at a different scale. Weighted highpass filters at one scale may be preferable to others depending on the signal and noise structure. More generally, a combination of weighted highpass filtered images may be preferable to any one.

We would like to choose the "best" weighted highpass filtered image from all possibilities. Ideally, the best weighted highpass filter provides the same level of enhancement as the linear highpass filter in regions of high intensity or in regions around a local edge, while reducing noise amplification in other areas. Hence, our objective is to preserve as much signal detail as possible in the weighted highpass filtered image. However, due the conflicting requirements of enhancement and noise suppression, different weighted highpass filters provide varying degrees of enhancement.

We advocate finding the weighted highpass filtered image that is closest to the linear highpass filtered image. The underlying principle behind this approach is that, by design, none of the weighted highpass filtered images can "match" the amplified noise component of the linear highpass filtered image. However, there is a best weighted highpass filtered image that comes close to matching the desired enhancement of true image detail.

Mathematically, we justify our approach as follows. Consider the optimisation programme

$$H^\star f \;=\; \arg \min_{H_{\mathrm{W}} f \in \mathcal{C}_f} \|H_{\mathrm{W}} f - Hf\|_F^2, \tag{7}$$

where $\|\cdot\|_F$ denotes Frobenius matrix-norm. The solution $H^\star f$ is the weighted highpass filtered image in $\mathcal{C}_f$ closest in norm to the linear highpass filtered image, or equivalently, the projection of the linear highpass filtered image onto

the subspace spanned by the set of weighted highpass filtered images. We can compute $H^\star f$ by adjusting the filter parameters $\{\alpha_{j,p}, \beta_{j,p}\}$ in (6).

Let us consider this minimisation more carefully. The error $\|H_{\rm w} f - Hf\|_F^2$ can be decomposed into two components. Let $e_1(H_{\rm w} f, Hf)$ denote the sum of squared errors in pixels of high intensity and/or those pixels near an edge. Let $e_2(H_{\rm w} f, Hf)$ denote the sum of squared errors in the remaining pixels. Thus we have $\|H_{\rm w} f - Hf\|_F^2 = e_1(H_{\rm w} f, Hf) + e_2(H_{\rm w} f, Hf)$. Since, by design, all weighted highpass filters have a very small gain except in bright or edgy regions, the error $e_2(H_{\rm w} f, Hf)$ can be well approximated by

$$e_2(H_{\rm w} f, Hf) \approx e_2(0, Hf). \tag{8}$$

That is, $e_2(H_{\rm w} f, Hf)$ is dominated by the contribution due to the linear highpass filter. Hence, this component of the overall error is approximately independent of the choice of weighted highpass filter. Programme (7) is thus approximately equivalent to the desirable minimisation

$$H^\star f \;\approx\; \arg \min_{H_{\rm w} f \in \mathcal{C}_f} e_1(H_{\rm w} f, Hf). \tag{9}$$

Therefore, $H^\star f$ is approximately equal to the weighted highpass filtered image that is closest to the linear highpass filtered image in bright regions and/or near edges. By design, $H^\star f$ is very small in other areas of the image.

Remark 1: $H^\star f$ is not necessarily the weighted highpass filtered image that provides the greatest reduction in noise amplification in bright regions and/or near edges. In fact if two highpass filters have equal e_1 errors, then the minimisation in (7) produces the filtered image with the smaller e_2 error and hence the lesser reduction of noise amplification. However, equation (8) shows that this suboptimal reduction of noise is negligible in comparison to the noise in the linear highpass filtered image, and equation (9) shows that $H^\star f$ very close to the optimal solution.

Remark 2: The minimization (9) would require prior knowledge of the location of edges and/or bright regions in the noise-free image. In problems of practical interest, we only have the noisy image to work with and therefore such prior knowledge is unavailable. The minimization in (7) represents a practical alternative to the desired optimisation. The examples in Section 4 demonstrate the excellent performance of the proposed weighted highpass filters.

Remark 3: Note that we may pose the minimization over any subspace spanned by a subset of the edge-weighted and/or mean-weighted filtered images.

Remark 4: The filter $H^\star$ is unique and can be computed in a simple fashion. First let

$$\mathbf{d}_{j,p} = \text{vec}(|D_{2^j} f|^p \cdot Hf), \quad \mathbf{s}_{j,p} = \text{vec}(|S_{2^j} f|^p \cdot Hf), \quad \mathbf{h} = \text{vec}(Hf), \tag{10}$$

where the operator "vec" forms a column vector from a matrix by stacking its columns. Since the Frobenius norm coincides with the vector 2-norm, (7) can be rewritten as

$$\mathbf{h}^\star \;=\; \arg \min_{N \in \mathcal{C}_f} \Big\| \sum_{j=1}^{J} \sum_{p=1}^{P} \alpha_{j,p} \mathbf{d}_{j,p} + \beta_{j,p} \mathbf{s}_{j,p} - \mathbf{h} \Big\|_2^2. \tag{11}$$

It is clear that the filter is specified by the $2JP$ parameters $\{\alpha_{j,p}\}$ and $\{\beta_{j,p}\}$. Now define the matrix $\mathbf{X} = [\mathbf{d}_{1,1}, \ldots, \mathbf{d}_{J,P}, \mathbf{s}_{1,1}, \ldots, \mathbf{s}_{J,P}]$ and the parameter vector $\boldsymbol{\gamma} = [\alpha_{1,1}, \ldots, \alpha_{J,P}, \beta_{1,1}, \ldots, \beta_{J,P}]^T$. The filter parameters

are given by

$$\boldsymbol{\gamma}^\star \;=\; \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}^{2JP}} \|\mathbf{X}\boldsymbol{\gamma} - \mathbf{h}\|_2^2. \tag{12}$$

The adaptive weighted highpass image, in vectorized form, is given by

$$\mathbf{h}^\star = \mathbf{X}\boldsymbol{\gamma}^\star. \tag{13}$$

Recall that, once we have $H^\star$ in hand, we form the enhanced image g as the unsharp mask $g = f + H^\star f$.

3.4 Local Adaptive Weighting: The adaptive weighted highpass filter described in the previous section is tuned to each individual image. This tuning is a global optimisation over the entire image. However, the scale of local structure may differ within the image itself. Consequently, no single scale, nor weighted highpass filter, is locally optimal at all points in the image.

In this section, we briefly describe an adaptive filter that adjusts its weighting coefficients at each point in the image. A related idea is considered in [5] to improve the performance of the weighted highpass filter proposed in [1]. The adaptive algorithm is more computationally intensive, but can provide significant improvements over the global adaptive weighted highpass filter described in the previous section. The local adaptive algorithm computes the filter at the point (x, y) by considering the error between weighted and linear highpass filtered images only in a local neighborhood around (x, y) rather than by considering the total error over the entire image. The procedure for the local adaptive filter algorithm is straightforward; we give a brief description below.

We again consider the collection of weighted highpass filters $\mathcal{C}_f$ defined in (6). However, rather than computing the globally adaptive filter according to (7), at each point (x, y) in the image we compute a local adaptive filter as follows. First, let $B(x, y)$ denote a local neighborhood about (x, y) in the image and let $\|\cdot\|_{B(x,y)}$ denote the Frobenius matrix norm restricted to the neighborhood $B(x, y)$. That is, for images f and g

$$\|f - g\|_{B(x,y)}^2 = \sum_{(i,j) \in B(x,y)} \left[f(i, j) - g(i, j)\right]^2 . \tag{14}$$

We will use the same notation for the vector 2-norm when working with vectorized versions of the images. Now at each point in the image, the adaptive filter parameters are obtained by solving

$$\boldsymbol{\gamma}^\star(x, y) \;=\; \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}^{2JP}} \|\mathbf{X}\boldsymbol{\gamma} - \mathbf{h}\|_{B(x,y)}^2. \tag{15}$$

Using these parameters, the output of the local adaptive weighted highpass filter at the point (x, y) is given by

$$\mathbf{h}^\star(x, y) = \mathbf{X}(x, y)\boldsymbol{\gamma}^\star(x, y), \tag{16}$$

with $\mathbf{X}(x, y)$ the row vector

$$\left[|D_1 f(x, y)|, \ldots, |D_{2^J} f(x, y)|^P,\, S_1 f(x, y)|, \ldots, |S_{2^J} f(x, y)|^P\right] . \tag{17}$$

4 Simulations

In this section we present several examples to illustrate the performance and flexibility of the both the global and adaptive multiscale adaptive weighted highpass filters.

4.1 Edge-Weighted Enhancement: We consider two examples of image enhancement. The original images are shown in Fig. 3(a) (blurred Lenna image), and Fig. 4(a) [PET (Positron Emission Tomography) brain image].[2] The 256×256 Lenna image was blurred through convolution with the kernel $b_a = \begin{bmatrix} 1 & 0 & 1 \\ 0 & a & 0 \\ 1 & 0 & 1 \end{bmatrix}$ with $a = 8$. The 128×128 PET image was blurred also through b_a, with $a = 4$.

The key feature in these examples is that both images are processed by exactly the same adaptive weighted highpass filter algorithm — with no tweaking of parameters to handle the drastically differing image structures. First, the images are enhanced using a linear highpass filter H whose convolution mask is given by

$$\mathbf{H} \;=\; 0.5 \begin{bmatrix} -1 & 0 & -1 \\ 0 & 4 & 0 \\ -1 & 0 & -1 \end{bmatrix}. \tag{18}$$

The linearly enhanced images, shown in Figs. 3(b) and 4(b), are computed as $g = f + Hf$.

The space of edge-weighted highpass filtered images considered in both cases is

$$\text{Span}\{\, |D_{2^j} f| \cdot Hf \;:\; j = 1, \ldots, 4 \,\}. \tag{19}$$

The adaptive weighted filter parameters for the Lenna image are $\gamma^\star = 0.017 \times [0.32, -0.03, -0.09, 1.0]^T$, where the ordering is $[|D_{2^j} f|\, Hf,\;\; j = 1, \ldots, 4]$. The globally adaptive nonlinear enhancement $f + H^\star f$ of the Lenna is shown in Fig. 3(c). Note that the adaptive combination of the weighted highpass filtered images involves negative coefficients — parts of the image are both "built up" and "chipped away" by the component filters in order to optimise the enhancement.

The adaptive filter parameters for the PET image are $\gamma^\star = 0.033 \times [0.67, -0.56, 0.25, 1.0]^T$. The globally adaptive nonlinear enhancement $f + H^\star f$ of the PET image is shown in Fig. 4(c).

Note that the adaptive filters are quite different for the two images. However, in both cases the resulting nonlinear filter enhances the detail of the image while reducing the noise amplified by the linear highpass filter.

4.2 Locally Adaptive Edge-Weighted Enhancement: We compare the global adaptive weighted filter (13) to the adaptive adaptive weighted filter (16). The image of Fig. 5(a) was obtained by first convolving the 256×256 bridge image with b_a as above (with $a = 4$) and then adding a small amount of white Gaussian noise. The image was enhanced using the linear highpass filter given in (18). The linear highpass filter enhancement is shown in Fig. 5(b).

The space of edge-weighted highpass filters considered in this case is

$$\text{Span}\{\, |D_{2^j} f|^2 \cdot Hf \;:\; j = 1, \ldots, 3 \,\}. \tag{20}$$

The globally adaptive edge-weighted highpass filter enhancement is shown in Fig. 5(c). The adaptive, local adaptive enhancement based on a 16×16 adaptation region is pictured in Fig. 5(d). Note that the adaptive algorithm is better

[2] Courtesy of Col. Brian W. Murphy, Center For Positron Emission Tomography, State University of New York at Buffalo.

able to adjust to the local structure within the image, yet still reduces the noise that is amplified by the linear highpass filter in homogeneous regions of the image.

4.3 Adaptive Weighted Restoration: In [6], we consider the adaptive weighted restoration of a degraded image with a known blurring function. The results presented there show that our method performs better than conventional linear restoration in both a visual and squared error sense.

5 Conclusions

We have developed a family of adaptive weighted highpass filters based on multiscale analysis. Two significant features distinguish our method from previous work. First, the filters do not have a fixed form like previously proposed filters. Therefore, the filters are capable of matching the structure of the image at hand. Secondly, the design of the adaptive filter is fully automatic. Previously proposed filters have required user specified parameters and/or ad hoc thresholding schemes. We have also derived an adaptive filter that automatically adjusts to varying structure within an image itself. Simulations have demonstrated that the proposed filter provides very good results for images with differing local structure.

There are many possible avenues for future work in this area. For example, multiscale analyses other than that of [3] may produce better results in certain cases. Also, it may be advantageous to decompose the linear highpass filtered image Hf at different scales as well. A deeper understanding of the nonlinear filtering concepts presented here may be gained by noting that weighted highpass filters belong to the class of nonlinear filters known as *Volterra filters*. The theory of Volterra filters should provide insight into the analysis, implementation, and design of nonlinear enhancement filters. On a final, more ambitious note, adaptive weighted highpass filters could provide a plausible model for studying masking phenomena in the human visual system.

References

[1] P. Fontanot and G. Ramponi, "A polynomial filter for the preprocessing of mail address images," *Proc. 1993 IEEE Workshop on Nonlinear Digital Signal Processing*, session 2.1, pp. 6.1-6.6, January, 1993.

[2] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.

[3] S. Mallat and S. Zhong, "Characterization of signals from multiscale edges," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 7, pp. 710-732, July, 1992.

[4] S. K. Mitra, S. Thurnhofer, M. Lightstone, and Norbert Strobel, "Two-dimensional Teager operators and their image processing applications," *Proc. IEEE Wksp. Nonlinear Signal Image Proc.*, pp. 959-962, June, 1995.

[5] S. Mo and V. J. Mathews, "Adaptive binarization of document images," *Proc. 1995 IEEE Workshop on Nonlinear Signal and Image Processing*, pp. 967-970, June, 1995.

[6] R. D. Nowak and R. G. Baraniuk, "Optimally weighted highpass filters using multiscale analysis," IEEE Southwest Symposium on Image Analysis and Interpretation, San Antonio, TX, 1996.

[7] L. A. Olzak and J. P. Thomas, "Seeing spatial patterns," in *Handbook of Perception and Human Performance*, (K.R. Boff, L. Kaufman, and J.P. Thomas, eds.), ch. 7, New York: John Wiley and Sons, 1986. pp. 963-966, June, 1995.

[8] G. Ramponi, "A simple cubic operator for sharpening an image," *Proc. 1995 IEEE Workshop on Nonlinear Signal and Image Processing*, pp. 963-966, June, 1995.

[9] S. Thurnhofer, *Quadratic Volterra Filters for Edge Enhancement and Their Applications in Image Processing*, Ph.D. dissertation, University of California-Santa Barbara, 1994.

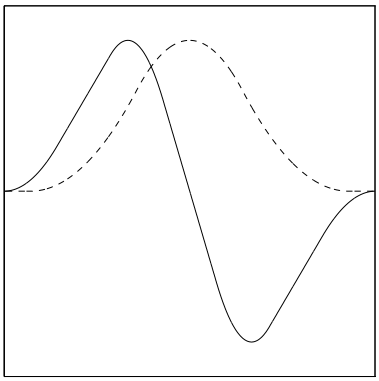

Figure 1: *Smoothing function ϕ (dashed) and wavelet ψ (solid) employed in the multiscale decomposition.*

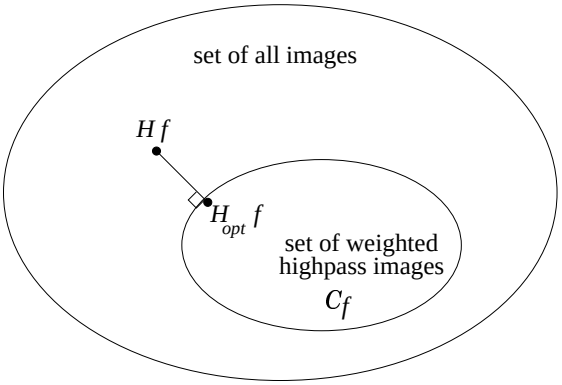

Figure 2: *The projection of the highpass filtered image Hf onto the set of all weighted highpass filtered images C_f defined in (6) yields the optimal weighted highpass filtered image $H_{opt}f$.*

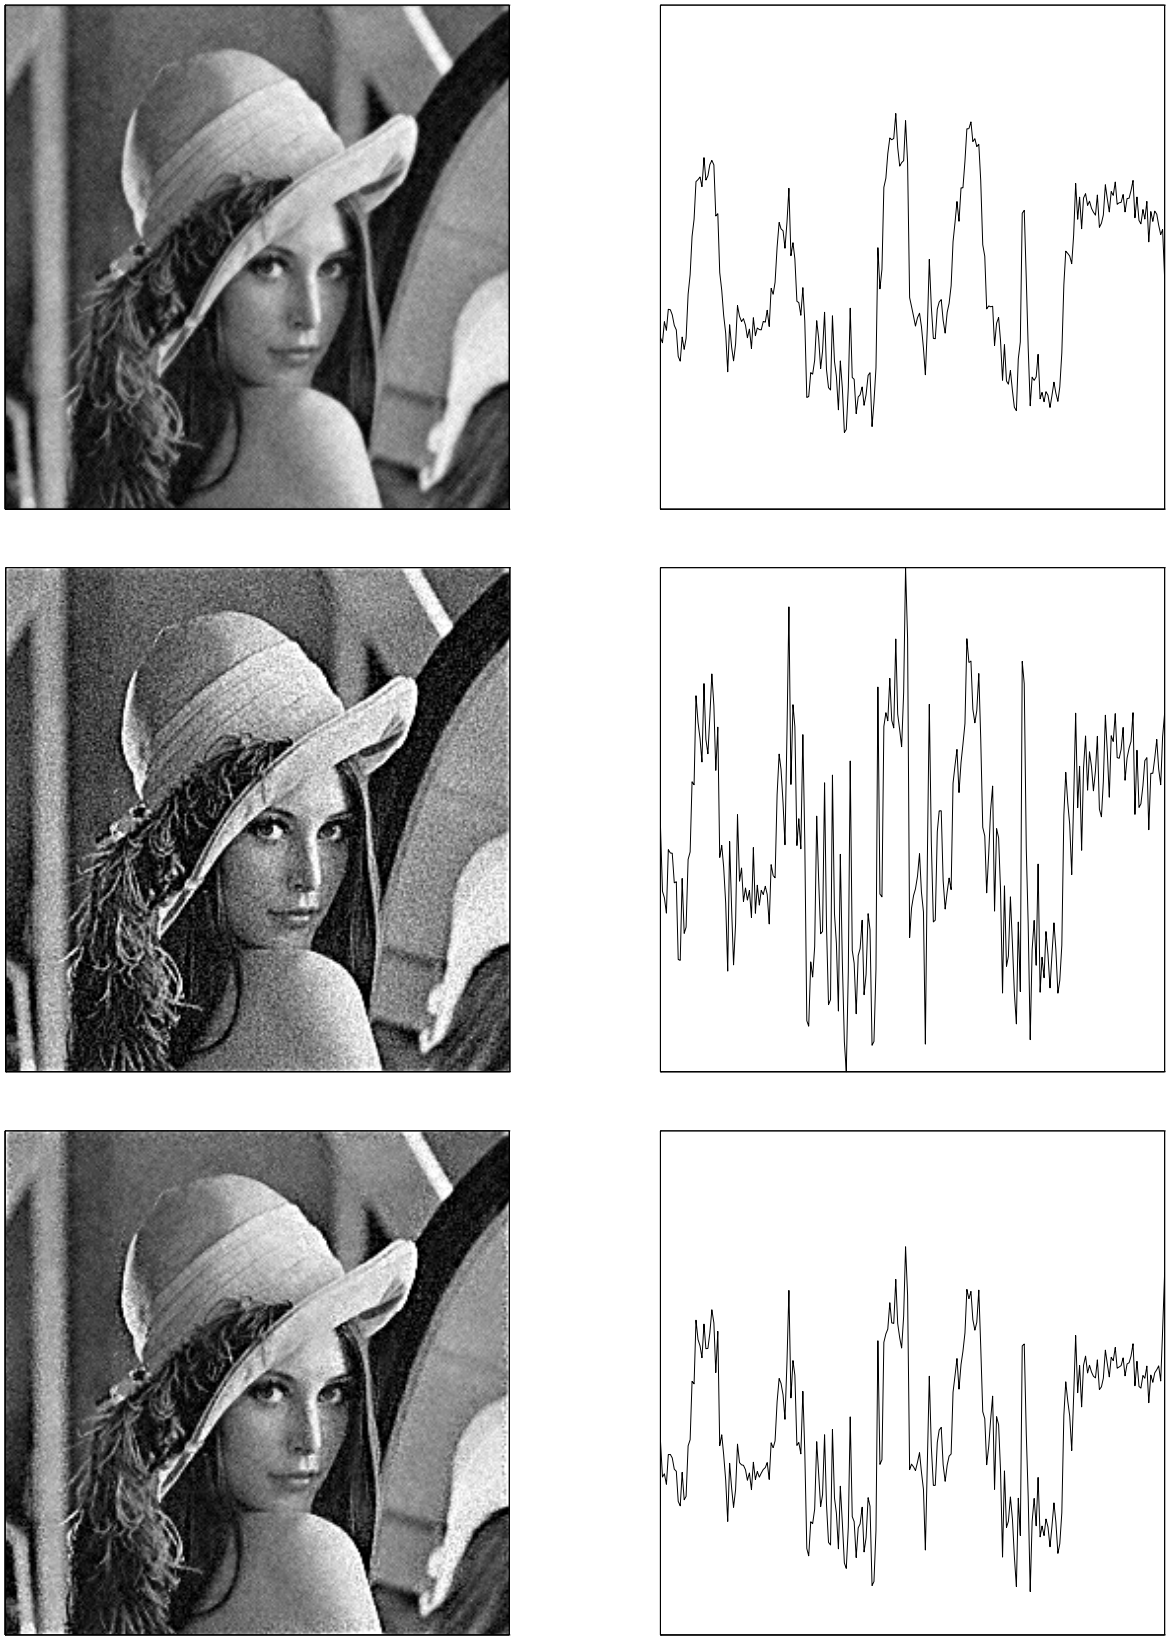

Figure 3: *Optimally weighted enhancement. (a) Original image (blurred Lenna). (b) Image enhanced using linear highpass filter. (c) Image enhanced using optimal edge-weighted highpass filter. At left, we show the image; at right, we show a vertical cross-section through the center of the image.*

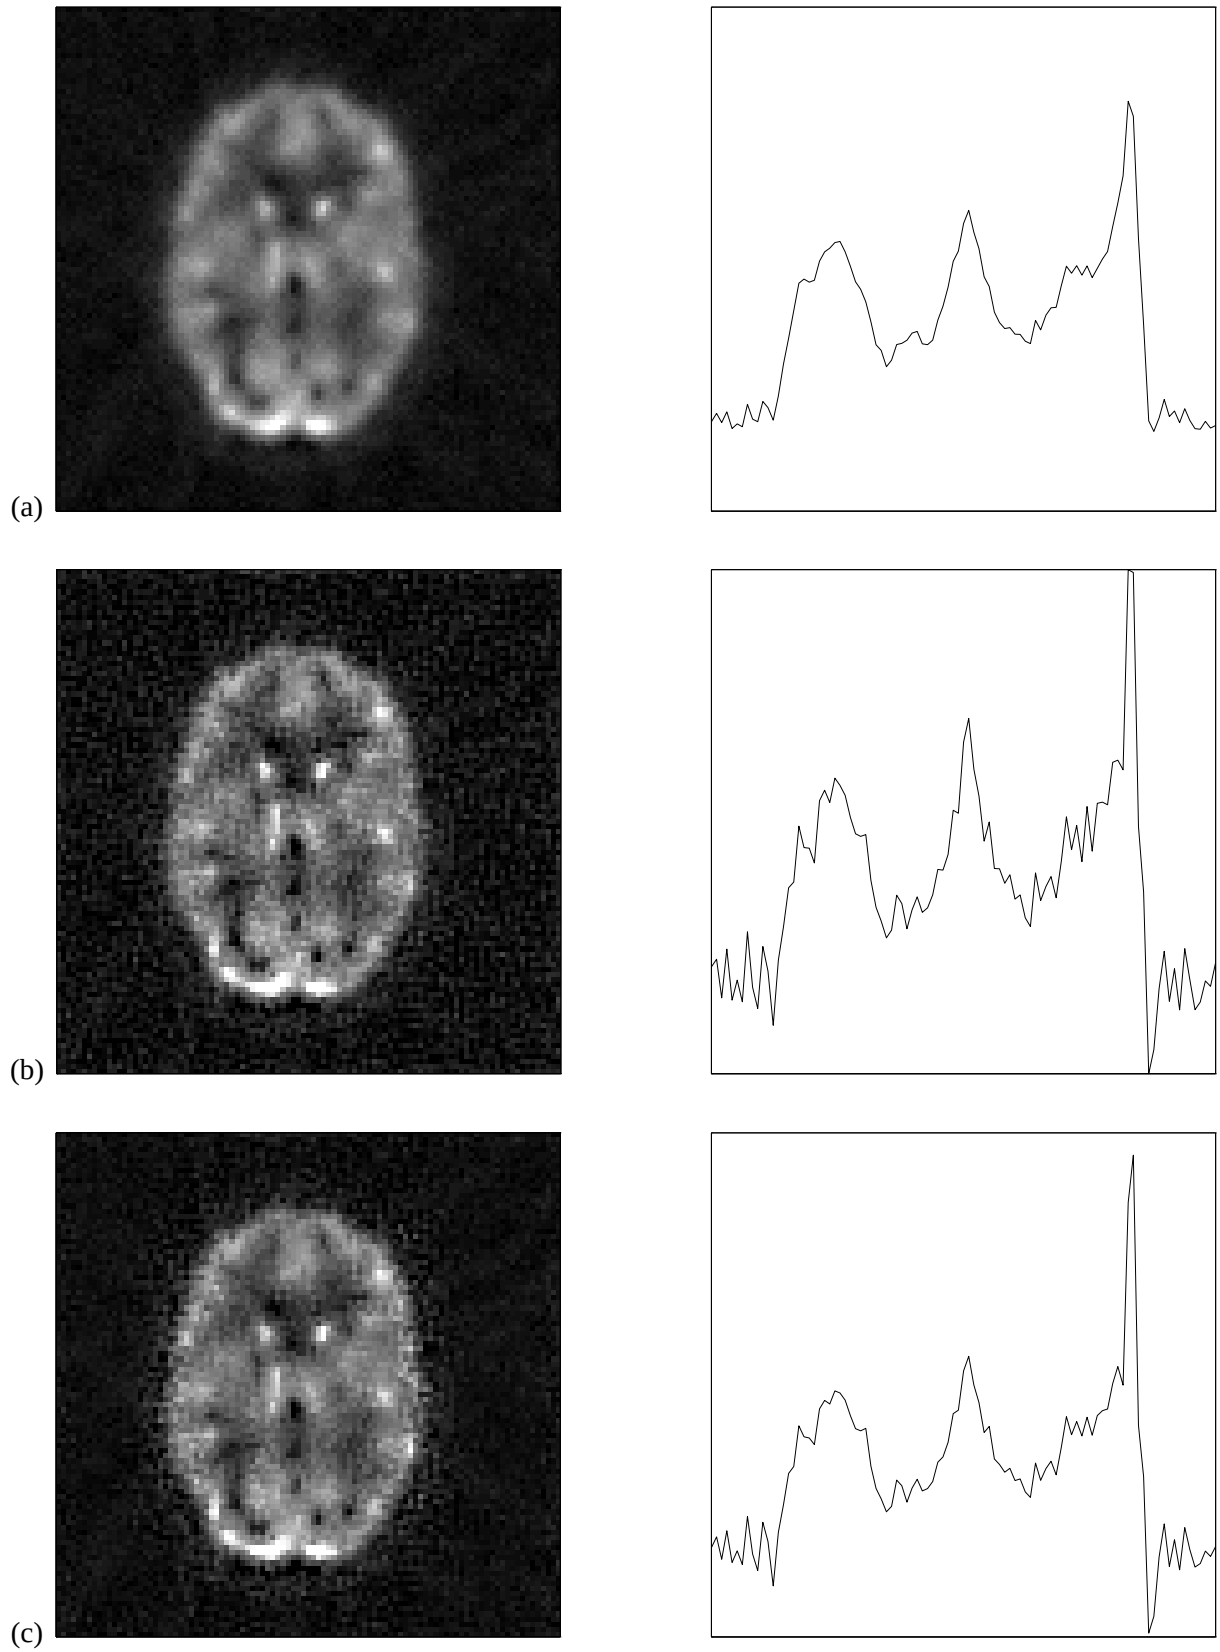

Figure 4: *Optimally weighted enhancement. (a) Original image (PET reconstruction). (b) Image enhanced using linear highpass filter. (c) Image enhanced using optimal edge-weighted highpass filter. Note that the nonlinear filtering algorithm employed here is identical to that used in Figure 3.*

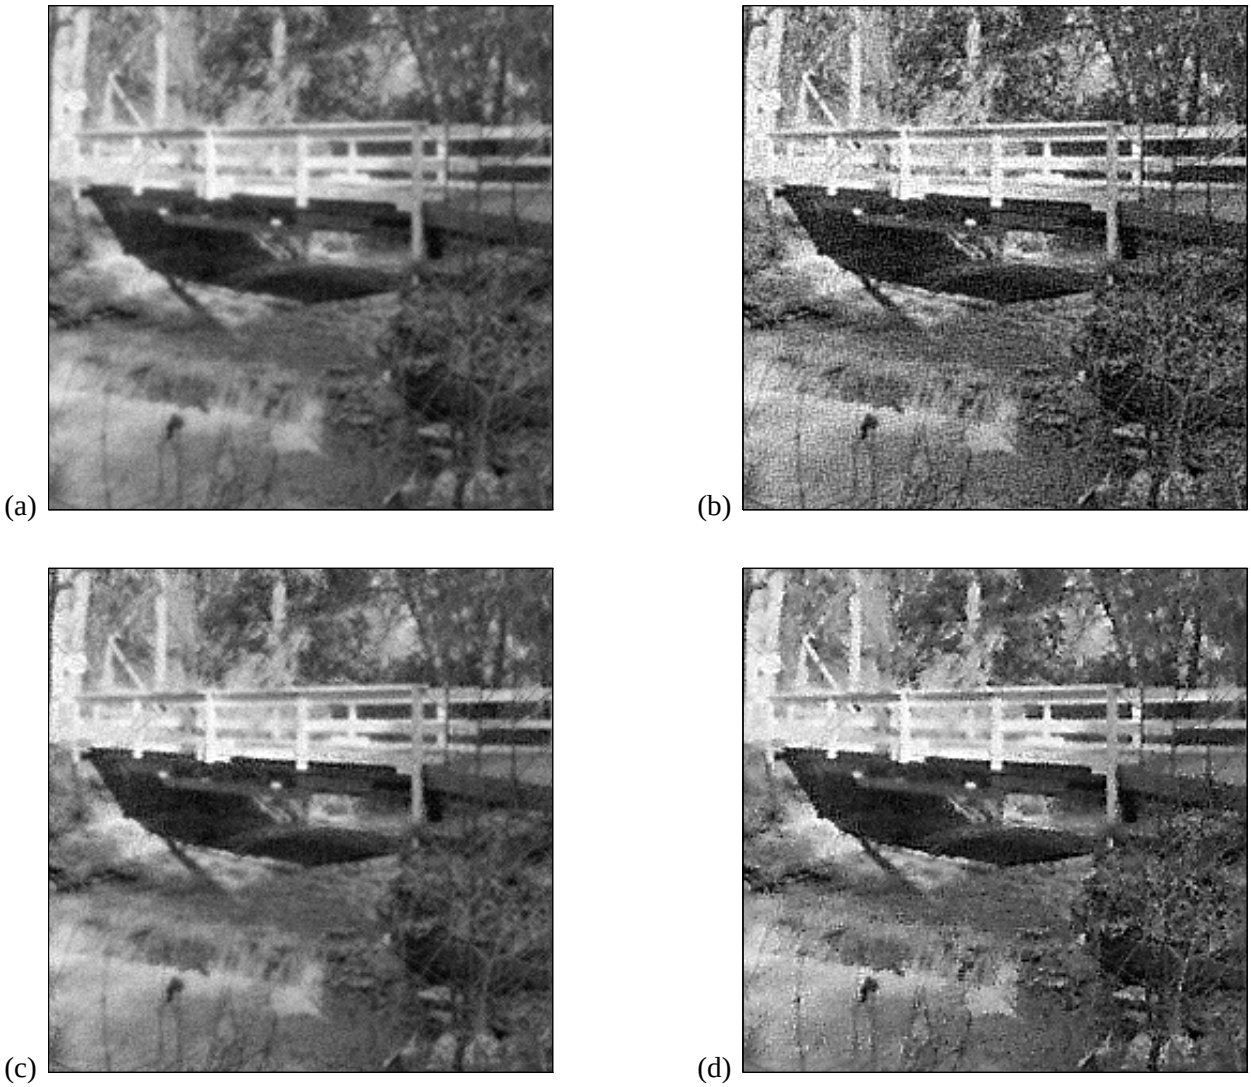

Figure 5: *Optimal adaptive weighted image restoration. (a) Blurred, noisy image of bridge. (b) Image restored using linear highpass filter. (c) Image restored using globally optimal edge-weighted highpass filter. (d) Image restored using locally optimal adaptive edge-weighted highpass.*