# Color Image Processing Using Wavelet-based Multiscale Edges

*Bruce Thomas*
*Mitfarb Consulting, Santa Cruz, California, U.S.A.*

## Abstract

This paper introduces a new approach to color image processing that operates on multiscale edge descriptions of color component images. The multiscale edges are computed in the wavelet domain and recorded using the Wavelet Transform Modulus Maxima representation. The resulting description provides a means for assessing the strength and significance of individual edges, with deference to edges that coincide between color components. This leads to an effective means for suppressing noise in color images. A method of edge sharpening is also demonstrated. Results show the viability of this new form of processing.

## Introduction

The bands of data comprising a color image are uniquely correlated. Perhaps the most apparent similarity is the existence of common image structure. Color component images typically exhibit edges and object outlines that coincide, yet vary in size and strength. These are inherent traits that can be exploited for color image enhancement.

In the following, we present a new framework for processing three-component color image data. The new framework is based on a multiresolution analysis of color image edge information. In many ways, this follows directly from Marr's[7] computational theory of vision in which multiscale edges are used to derive a "primal sketch". According to Marr, the "primal sketch" provides a fundamental description of the underlying scene. We extend Marr's conjecture to color by cross analyzing the multiscale edges of each color component in an effort to correlate the image structure found in each component.

Marr's work on vision employed Laplacian-of-Gaussian (LoG) filters which can be used to locate the multiscale edges of an image. These filters constitute a model of the simple receptive field of the human visual pathway, hence they bear biological relevance. Unfortunately LoG filters provide an analysis-only computational pathway. There is no effective way to reconstruct image data from a Laplacian-of-Gaussian based edge decomposition.

The necessary tools arrived when Mallat and Zhong[5] devised the Wavelet Transform Modulus Maxima (WTMM) representation along with an associated reconstruction algorithm. This representation provides a Marr-like analysis of multiscale edges, but also provides a mechanism for synthesizing image data from a set of multiscale edge information. The ability to analyze and synthesize image data to and from multiscale edgemaps enables the processing strategy outlined Figure 1.

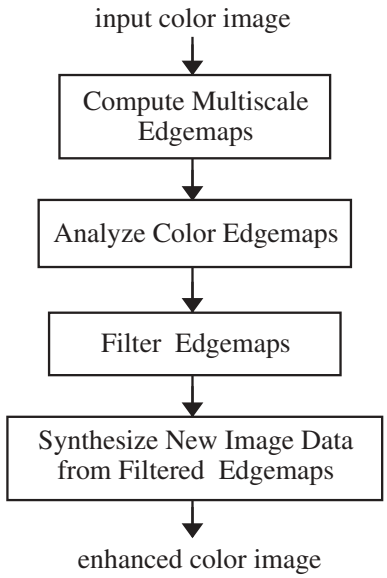


*Figure 1. Outline of multiscale processing strategy.*

The first step is to split the input color image into a constituent set of color component images and then compute the multiscale edgemaps of each component image. The respective edgemaps are then analyzed in an effort to identify physically significant edges. These edges are then filtered to achieve a desired enhancement (e.g. boost for strengthening, diminish for suppression). The final step is to synthesize new color component images from the filtered edgemaps using the WTMM reconstruction algorithm.

It should be noted that aspects of this processing strategy have recently been explored by other investigators. For example, Kotera et al.[3, 4] have developed image sharpening algorithms that assess the strength of edges in a color

image and apply scale-specific filter operators accordingly (i.e. large scale sharpening operators are applied to soft edges, fine scale sharpening operators are applied to hard edges). Likewise, Scheunders[8] has devised an alternate way of analyzing multiscale edges in color images using wavelets, with the express purpose of image filtering and enhancement.

## WTMM Representation

The Wavelet Transform Modulus Maxima (WTMM) representation of an image records the location, magnitude, and orientation of the image's multiscale edges at dyadic scales of analysis. It is constructed from a particular class of wavelets that compute derivatives of smoothed versions of the input signal.[5] For two-dimensional signals, partial derivatives are computed at each scale, $s$, resulting in the following horizontal and vertical wavelet transform subbands

$$W_s^H f(x,y) = s\frac{\partial}{\partial x}[f(x,y) * \theta_s(x,y)] \quad (1)$$

$$W_s^V f(x,y) = s\frac{\partial}{\partial y}[f(x,y) * \theta_s(x,y)] \quad (2)$$

where $\theta_s = (1/s^2)\,\theta(x/s, y/s)$ is a dilated version of a smoothing function $\theta(x,y)$. The partial derivatives (1) and (2) can be combined to form gradient vectors which can then be used to locate edges at scale $s$.

Spatially sampled versions of the preceding calculations can be computed efficiently using a digital filter bank. This results in horizontal and vertical wavelet subbands at dyadic scales of analysis, $s = 2^j$, for $j = 1, 2, \ldots, \log_2(N)$ where $N$ is the smaller of the input image's pixel dimensions. Spatially sampled versions of edge gradient modulus and angle data are computed from the resulting wavelet coefficients

$$M_{2^j} f(m,n) = \sqrt{|W_{2^j}^H f(m,n)|^2 + |W_{2^j}^V f(m,n)|^2} \quad (3)$$

$$A_{2^j} f(m,n) = \arctan\left|\frac{W_{2^j}^V f(m,n)}{W_{2^j}^H f(m,n)}\right| . \quad (4)$$

A Canny-like edge detector is then used to locate the modulus maxima, i.e. edges. This leads to the WTMM representation which records the position, $(m,n)$, magnitude $M_{2^j} f(m,n)$, and angle $A_{2^j} f(m,n)$ of multiscale edges. The integer $j$ is an index that identifies the scale of analysis (resolution). The WTMM representation also records a low-pass "coarse-image" that is used during reconstruction.

## Analyzing Color Edgemaps

The WTMM representation and reconstruction algorithm developed by Mallat and Zhong provide a means for processing color image data in the multiscale edge domain. The first step is to compute multiscale edgemaps for each component of the input color image. In the following we assume the input is an RGB image.

Let $M_{2^j}R(m,n)$, $M_{2^j}G(m,n)$, and $M_{2^j}B(m,n)$ be "modulus images" from the WTMM representations of the R, G, and B components, respectively. Modulus images have zero value at pixel locations where there is no edge, and edge magnitude values at all other pixel locations. The corresponding angle data are stored in "angle images" $A_{2^j}R(m,n)$, $A_{2^j}G(m,n)$, and $A_{2^j}B(m,n)$. Angle images have undefined pixel values at locations where there is no edge, and edge gradient angle data at all other locations. Three low-pass "coarse images" are also produced. These are labelled $S_{2^J}R(m,n)$, $S_{2^J}G(m,n)$, and $S_{2^J}B(m,n)$.

Once we have generated the WTMM representation for each component image, it becomes possible to analyze corresponding sets of edgemaps, per the second block of the processing strategy outlined in Figure 1. The goal is to identify physically significant edges from the pool of edges found in the multiscale description. A simple way to do this is to sum the red, green, and blue modulus data (for a specific scale of analysis) and determine where the resulting "aggregate" data exceeds a fixed threshold. If all three components exhibit an edge at a certain location, then the resulting aggregate value is expected to be large, which indicates that a physically significant edge exists at that location.

The location of large aggregate values, and hence physically significant edges, are recorded in an "edge likelihood" mask $k_{2^j}(m,n)$. This mask is computed by first summing component modulus data, creating an aggregate modulus image

$$M_{2^j}\Sigma(m,n) = M_{2^j}R(m,n) + M_{2^j}G(m,n) + M_{2^j}B(m,n) \quad (5)$$

then clipping all pixels using a scale-specific threshold, $T_{2^j}$, and normalizing the resulting image so that it has values falling on $(0,1]$. In summary,

$$k_{2^j}(m,n) = \mathrm{g}_{T_{2^j}}(\, M_{2^j}\Sigma(m,n)\,)\,/T_{2^j} \quad (6)$$

where $\mathrm{g}_T(x)$ is a clipping function, e.g. $\mathrm{MIN}(x,T)$.

Figure 2 shows an example of an edge likelihood mask for scale $j = 2$. White pixels have the maximum possible mask value, 1, and correspond to physically significant edge locations. All other edges have values that vary, with respect to the clipping level, in direct proportion to their aggregate strength. This may be regarded as a measure of

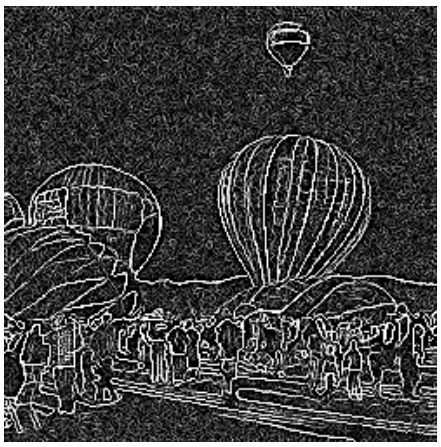

*Figure 2. Example of an "edge likelihood" mask computed from a noisy image at scale $j=2$ using (6).*

the "edge likelihood" of these remaining edges (the likelihood that an edge is structurally significant).

Scale-specific threshold $T_{2^j}$ is calculated from the histogram of the respective aggregate modulus image data. Figure 3 demonstrates the method used in this report. The desired threshold is computed by doing a linear least squares fit of the right, downward edge of the main lobe of an aggregate modulus image's histogram. The least squares fit is limited to histogram data falling between $90\%$ and $10\%$ of the main lobe's peak value. The x-intercept of the resulting line specifies the threshold.

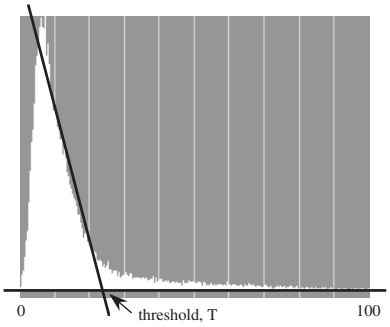


*Figure 3. Sample aggregate modulus histogram with construction lines showing how auto-threshold is computed.*

## Denoising Example

The analysis of color edgemaps leads the way to edgemap filtering operations. As a first example of edgemap filtering, we show how "edge likelihood" masks can be used to suppress noise in color images. The input is assumed to be a noisy RGB image. This is a practical problem that affects digital cameras.[1]

Image denoising is accomplished by recognizing that edge likelihood masks, $k_{2^j}(m,n)$, can be used to form multiplicative masks that attenuate noise-related edges (i.e. edges that are not deemed to be physically significant). This is carried out by multiplying individual component modulus images by the appropriate mask. Previous work[11] discussed the use of edge likelihood masks directly. Improved noise suppression is achieved by using squared versions of the edge likelihood masks, as used herein. The necessary masks are computed and applied to the two finest scales of analysis, i.e. $s=2^1$ and $s=2^2$. We restrict masking to these scales since these are scales for which color component edges exhibit the greatest proportion of coincident edges.[9, 10]

The following filtering operation performs "accelerated" noise suppression

$$\begin{aligned} M_{2^j}R'(m,n) &= |k_{2^j}(m,n)|^2\, M_{2^j}R(m,n) \\ M_{2^j}G'(m,n) &= |k_{2^j}(m,n)|^2\, M_{2^j}G(m,n) \\ M_{2^j}B'(m,n) &= |k_{2^j}(m,n)|^2\, M_{2^j}B(m,n) \\ &\text{for } j=1,2 \end{aligned} \tag{7}$$

where primed quantities represent filtered data. Note that the filtering operation does not affect the angle or coarse image data (these are left unfiltered).

The filtering operation (7) produces three new sets of WTMM data, where each set corresponds to a particular color component. New, denoised image components are then synthesized from each of the new sets of WTMM data using Mallat and Zhong's reconstruction algorithm. The resulting color component images are combined to form a denoised color image. Figure 4 demonstrates the denoising property of this form of processing. The left half of Figure 4 shows the original, noisy image while the right half shows the result obtained from edgemap filtering.

*Figure 4. Demonstration of noise suppression produced by multiscale edgemap filtering (image is originally in color).*

### Experimental Results

In order to assess the effectiveness of the filtering operation (7) we compare it to wavelet thresholding. Wavelet thresholding is a method of image denoising that is known to have near optimal performance (as measured by signal-to-noise ratio).[2] The following tests were performed on images contaminated by additive Gaussian ($\sigma = 20$) noise.

Figure 6 (at end of report) shows noise suppressed outputs obtained from multiscale edgemap filtering (7) and component separable wavelet thresholding, along with the original noisy image from which these outputs were derived. The wavelet thresholding computations were performed using undecimated Daubechies wavelets with a hard threshold of $T = 3.2\sigma_{est}$ where $\sigma_{est}$ is an estimate of the standard deviation of the noise (cf. Mallat[6]).

A comparison of Figures 6b and 6c reveals that the edgemap filtering approach of (7) does a better job of maintaining edges and image details than the wavelet thresholding approach. This is particularly evident in the magnified portions of each image shown in Figures 6e and 6f. Multiscale edgemap filtering is also less prone to smudging artifacts. This was supported by tests in which wavelet thresholding was performed using reduced size thresholds. Smudging artifacts continued to be evident for $T = 2.4\sigma_{est}$, and gave way to noise and wrinkle artifacts at $T = 1.6\sigma_{est}$. The edgemap filtering approach produced no such smudging as thresholds $T_{2^1}$ and $T_{2^2}$ were reduced.

Additional tests were carried out using luminance-only processing. These were carried out by denoising the Y-component of a YIQ transformation of the RGB input image (using edgemaps and wavelet information computed strictly from the Y-component). The denoised luminance image, Y', was then used to derive a corresponding R'G'B' image using Y'/Y scaling ratios (taking care to clip Y' in order to avoid out-of-gamut R'G'B' values). The edgemap filtering method was labelled 'lummsk' while the wavelet thresholding method was labelled 'lumthr'. The images produced by both of these methods exhibited color speckling artifacts which gave them an unsatisfactory appearance.

Denoising performance was also measured quantitatively. Figure 5 shows the results in graphical form. Performance was quantified using signal-to-noise ratio improvement (SNRI) which measures $\mathrm{SNR}out - \mathrm{SNR}in$.

The SNRI data graphed in Figure 5 reveal that component separable wavelet thresholding ('septhr') consistently outperformed the other approaches in terms of SNRI. This is consistent with wavelet thresholding's theoretical properties, but bears little witness to the blurry visual quality of wavelet thresholding outputs (cf. Figures 6c and 6f). The data also shows that multiscale edgemap filtering ('clrmsk') delivered strong SNRI performance in all of the images that were tested.

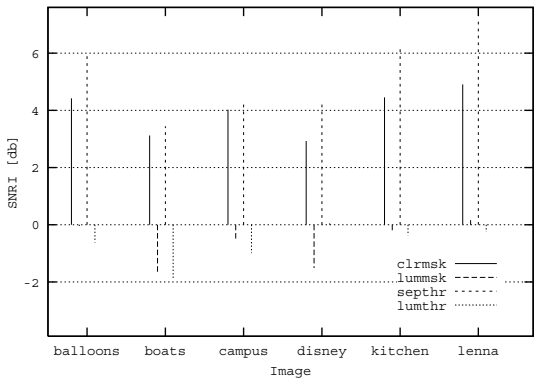


*Figure 5. Graphical depiction of $\sigma = 20$ denoising results.*

Perhaps the most startling aspect of Figure 5 is the fact is that all of the luminance-based tests exhibited negative or near-zero SNRI values.† This suggests that luminance-based denoising algorithms generally introduce more noise than they suppress. When it comes to suppressing strong additive Gaussian noise from color images, one is better off using algorithms that process all of the components of a color image.

## Edge Sharpening Example

As a second example of edge map filtering, we show how "edge likelihood" masks can be used to sharpen color image edges. The sharpening of color image edges has also been addressed in a multiscale setting by Kotera et al.[3, 4]

Here, instead of using edge likelihood masks to attenuate noise related edges, we use the masks to amplify edges that are considered to be physically significant. This is accomplished by once again employing edge likelihood masks, $k_{2^j}(m,n)$, as multiplicative masks. This time however, we incorporate a scale factor, $\alpha$, which specifies the amount by which edges are to be boosted.

Multiscale edge sharpening is accomplished by the following filtering operation

$$\begin{aligned} M_{2^j}R'(m,n) &= \alpha\, k_{2^j}(m,n)\, M_{2^j}R(m,n) \\ M_{2^j}G'(m,n) &= \alpha\, k_{2^j}(m,n)\, M_{2^j}G(m,n) \\ M_{2^j}B'(m,n) &= \alpha\, k_{2^j}(m,n)\, M_{2^j}B(m,n) \\ &\text{for } j = 1,2 \quad (\text{with } \alpha > 1.0)\ . \end{aligned} \tag{8}$$

As before, the filtering is restricted to the two finest scales of analysis and does not affect angle or coarse image data (these are left unfiltered). The processing also uses edge likelihood masks directly, without squaring, so that all edges are boosted evenly. An edge enhanced color image is obtained by substituting the filtered modulus edgemaps into

† n.b. previously reported[11] lummsk data are anomalous.

the color image's original WTMM data, then synthesizing new color component images from the filtered WTMM data.

### Experimental Results

Figure 7 shows the effect that multiscale edge sharpening (8) has on noise free color images. Both results were produced using $\alpha = 2$. The enhanced images exhibit strong contrast with pronounced object outlines in a fashion that resembles unsharp masking.

Luminance-based processing is also possible, and offers computational savings. Tests of a luminance-based version of equation (8) showed that it delivered results that were similar to those obtained using (8) directly. Thus, luminance-based processing appears to have merit in the case of multiscale edge sharpening.

The filtering strategy of (8) provides edge sharpening along with a modest degree of noise filtering. Tests on noisy color images revealed that it is not entirely effective at performing both tasks at once. This remains an area warranting further investigation.

## Conclusion

The components of a color image exhibit unique correspondences that can be exploited for color image enhancement. This report has presented a framework for taking advantage of edge related correspondences.

The proposed approach is based on a multiscale edge analysis and utilizes the interdependencies that exist between color component edgemaps. These interdependencies provide a relatively simple means of locating and gauging physically significant edges at specific scales of analysis, as stored in so called "edge likelihood" masks.

Two examples have demonstrated the capabilities of the proposed framework. The first example showed how multiscale edge filtering can be used to suppress noise in color images. This approach does a good job of maintaining image details, and offers SNR performance that is comparable to wavelet thresholding. This is significant for wavelet thresholding is known to deliver nearly optimum SNR performance but tends to produce blurred results. A second example showed how multiscale edge scaling can be used to sharpen color image edges.

The examples presented here use strictly edge magnitude information and cross-component interdependencies. Future work can build upon these methods by using edge angle information, exploiting across-scale interdependencies, and examining alternate color spaces as well.

## References


1. J. Adams and K. Spaulding, Noise cleaning digital camera images to improve color fidelity capabilities, *Proc. of IS&T/SID 7th Color Imaging Conference*, pp. 197-199 (1999).

2. D. Donoho, I. Johnstone, G. Kerkyacharian, and D. Picard, Wavelet shrinkage: asymptopia?, *J. Royal Statistical Society B*, vol. 57, pp. 301–369 (1995).

3. H. Kotera, Y. Yamada, and K. Shimo, Sharpness improvement adaptive to edge strength of color image, *Proc. of IS&T/SID 8th Color Imaging Conference*, pp. 149–154 (2000).

4. H. Kotera and W. Hui, Multi-scale image sharpening with background noise suppression, *Proc. of IS&T/SID 10th Color Imaging Conference*, pp. 196–201 (2002).

5. S. Mallat and S. Zhong, Characterization of signals from multiscale edges, *IEEE Trans. Patt. Analysis and Machine Intell.*, vol. 14, pp. 710–732 (1992).

6. S. Mallat, *A Wavelet Tour of Signal Processing, 2nd ed.*, Chesnut Hill, MA: Academic Press, 1999.

7. D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman, 1982.

8. P. Scheunders, Wavelet-based enhancement and denoising using multiscale structure tensor, Proc. Intl. Conf. Image Processing, vol.3, pp. 569–572 (2002).

9. B. Thomas and R. Strickland, A study of scene structure in the saturation component of color images, in *Recent Progress in Color Science*, IS&T, Springfield, VA, 1997, pp. 274–280.

10. B. Thomas, *New Aspects of Digital Color Image Enhancement*. Ph.D. dissertation, University of Arizona, 1999.

11. B. Thomas, Wavelet-based color image denoising, Proc. Intl. Conf. Image Processing, vol.2, pp. 804–807 (2000).


## Biography

Bruce Thomas holds degrees in electrical engineering from Cal Poly, San Luis Obispo (B.S.E.L.), Carnegie Mellon (M.S.E.E.) and the University of Arizona (Ph.D.) with an emphasis on optics and color image processing. He has a lifelong interest in photography, visual perception, and color, and has participated in the development of color pre-press and digital camera/imaging systems at DuPont, SuperMac Technology and Sierra Imaging. He is currently an industry consultant.

(a) original image (b) 'clrmsk' output (c) 'septhr' output

(d) original image (e) 'clrmsk' output (f) 'septhr' output

*Figure 6. Examples of color image noise suppression. 'clrmsk' designates multiscale edgemap filtering per eqn. (7). 'septhr' designates component separable wavelet thresholding. (all images are originally in color).*

(a) original image (b) edge enhanced image

(c) original image (d) edge enhanced image

*Figure 7. Examples of color image edge sharpening produced by multiscale edge scaling, eqn. (8). (all images are originally in color).*