

DENOISING NATURAL COLOR PHOTOS IN DIGITAL PHOTOGRAPHY

TONY CHAN, YANG WANG, AND HAOMIN ZHOU

ABSTRACT. As digital photography rapidly replacing the traditional film photography as the photography of choice for all but a few devoted professionals, post image processing of natural color photos such as denoising becomes increasingly an integral part of digital photography. Although many denoising schemes have been designed, almost none specifically target natural color photos. Noise in natural color photos have special characteristics that are substantially different from those that have been added artificially.

In this paper we propose the multiscale total variational method (MTV) for denoising. Standing alone the MTV method is effective in denoising monochromatic images. However, it demonstrates outstanding denoising capabilities for natural color images. Key to the success is the understanding of the characteristics of digital noise in natural color images as well as a non-traditional color space we have introduced specifically for the purpose. An automatic stopping criterion is applied to each channel to prevent over processing.

1. NOISE IN NATURAL COLOR PHOTOS

With the surging popularity of digital cameras, digital photography is rapidly replacing the traditional film photography as the photography of choice for virtually all but a few devoted professionals. In digital photography, post image processing is an integral part for obtaining better images even for the casual picture takers. Post image processing is especially important for people who are willing to go beyond point-and-shoot, and one of the key steps in image processing is denoising.

All digital cameras today take color photos. (Some cameras allow for black-and-white images, but these are converted from color images using in-camera firmwares.) Noise is present in virtually all digital photos, and there are several sources for it. When light (photons) strike the image sensor, electrons are produced. These “photoelectrons” give rise to analog signals which are then converted into digital pixels by an Analog to Digital (A/D)

Key words and phrases. Color images, denoising, wavelet TV, MTV, YCrCb color space.

This paper was completed while the first and second authors are working for the National Science Foundation as IPAs. The views expressed in this paper are those of the authors, and do not necessarily reflect those of the NSF. The third author is supported in part by the National Science Foundation grant DMS-0410062.



FIGURE 1. The original natural color image without artificial noise. Noise is not obvious due to limitation on the size of the display.

Converter. The random nature of photons striking the image sensor is an important source for noise. This type of noise, known as *photon shot noise*, is roughly proportional to the square root of the signal level as a result of the Central Limit Theorem. Thus the lower the signal is the higher the noise becomes relative to the signal. As a result noise in color images can be very pronounced in images shot under low light conditions because the signals must be amplified more. In general, noise level is very low for photos shot outdoor using low ISO (ISO 100 or less). But with most consumer compact cameras noise becomes visible at ISO 200, and it becomes unacceptable at ISO 400 or higher. With more advanced and expensive digital SLRs noise remains low even at ISO 400, and becomes unacceptable at ISO 1600 or higher. Noise is in general much worse under artificial lighting, especially under fluorescent lighting. One of the important characteristics of digital noise is that they are not uniform across all channels. Very often noise is concentrated in the blue channel while the green and the red channels are relatively clean. For photos taken under artificial lighting (without a flash), the blue channel can be so noisy that it is often unrecognizable, see Figures 1 and 2.




FIGURE 2. A zoom-in (upper-left) of the color image in 1 and its RGB channels. Color noise is more evident. The red (upper-right) and green (lower-left) channels are much cleaner than the blue (lower-right) channel.

Another significant source of noise is the so called *leakage current*. Semiconductor image sensors work by converting energy from photons into electrical energy, in the form of a current or voltage signal. Unfortunately thermal energy present in the semiconductor can also generate an electrical signal that is indistinguishable from the optical signal. As temperature increases, so does leakage current in the circuit. The effects of leakage current are most apparent in long exposures in which the light signal is very low.

Modeling noise in digital color photos can be a difficult task. The photon shot noise is clearly signal dependent and thus not uniform from pixel to pixel. Nearly all digital cameras¹ today use the so-called Bayer Pattern in their photo sensors, where half of their pixels are used to capture the green channel and the other half are divided evenly to capture the red and the blue channels. These partial data are then interpolated to complete the RGB channels of a color photo. So unless we access the raw data (most consumer digital cameras do not have this feature) it is clear that noise is not independent from pixel to pixel in any channel. Most digital photos are in JPEG format, which degrade images through quantization and artifacts. Furthermore, all cameras employ proprietary in-camera sharpening, denoising and anti-aliasing. These factors combine to make effective modeling of noise, at least in the images taken by consumer cameras exported in JPEG format, virtually impossible. For this reason, any noise model assuming independent and identically distributed noise from pixel to pixel can be unrealistic.

This uneven distribution of noise poses some challenges. Excessively denoising the blue channel can easily lead to color artifacts (such as color bleeding). A good denoising scheme must take this into consideration. One viable solution to this problem is to work in another color space rather than the RGB color space. A standard practice is to work in a color space that separates the luminance and chrominance. Commonly used color spaces are CIELAB, CIELUV and YCrCb. The YCrCb color space has the advantage for being linear. It is the color space used for JPEG and JPEG 2000. One of the innovations in this paper is to design a new color space that effectively takes into account the distribution of noise. This new color space offers superior performance in our tests.

2. MULTISCALE TOTAL VARIATIONAL METHOD FOR DENOISING MONOCHROMATIC IMAGES

Denoising methods for monochromatic images can essentially be classified into four categories: neighborhood filters, frequency domain methods, variational PDE based methods and non-local methods. The most commonly used neighborhood filters include blurring filters, median filter or variations of it, and others such as the filters described in [32, 33]. Neighborhood filters are easy to implement and fast, and in some applications they can be

¹Sigma digital SLRs such as SD9 and SD10, which employ sensors by Foveon, are notable exceptions.

effective, although in general their effectiveness is limited. In frequency domain methods a wavelet, DCT or other type of transformations is first performed. A filter is then applied to the transformation. The most common frequency domain method is the wavelet thresholding. The wavelet thresholding method assumes that noise appear in the wavelet transformation as small nonzero coefficients in the high frequency range and sets them to zero. The wavelet thresholding method is very effective in removing noise, and very fast. It is still perhaps the best “quick and easy” denoising scheme. However, it suffers from Gibbs oscillations at discontinuities. These oscillations can be reduced, although not eliminated, by using soft wavelet thresholding [16, 17] and translational invariant wavelet thresholding [13]. As we shall see, using the right color space the wavelet thresholding can be a very effective method for denoising color images. Methods that are not PDE based and do not fit the description of the other two, including some statistical methods, can be classified as non-local denoising methods, see e.g. [2, 7, 31, 20, 25]. These methods are typically slower, and some of them assume certain statistical properties are known. Given the right images, non-local methods can yield excellent results. For example, the non-local method developed in [2] and its refinement in [20] work very well for images with repeat patterns or large homogeneous areas.

The multiscale total variation (MTV) method we describe in this paper is a variational PDE based denoising method. It is a variation of the wavelet TV. We start with a standard noisy monochromatic image model

$$(2.1) \quad z(x) = u_0(x) + n(x),$$

where $z(x)$, $u_0(x)$ and $n(x)$ are real valued functions defined on $\Omega \subset \mathbb{R}^2$, where Ω is a finite domain such as a rectangle. The function $u_0(x)$ denotes the underlying noise-free image, $z(x)$ the observed image, and $n(x)$ the noise. With a variational PDE based denoising method, the denoised image is the minimizer of certain energy functional $\mathbf{E}(u)$. Typically $\mathbf{E}(u)$ can be written as

$$(2.2) \quad \mathbf{E}(u) = D(u, z) + R(u),$$

where $D(u, z)$ denotes the “distance” between u and observed image z , and $R(u)$ is a regularization term that smoothes out the image. The idea of using variational method described here has been around for some time. In most cases the distance $D(u, z)$ is taken to be the L^2 distance $D(u, z) = \int_{\Omega} (u - z)^2$. Earlier efforts focused on least square based functionals

$R(u)$'s such as $\|\Delta u\|_2^2$, $\|\nabla u\|_2^2$ and others. While noise can be effectively removed, these regularization functionals penalize discontinuity, resulting in soft and smooth reconstructed images, with subtle details lost. This is not acceptable in digital photography, as photographers often place premium emphasis on sharpness. The innovation of the total variational (TV) scheme by Rudin, Osher and Fatemi [26] is to set $R(u)$ to be the total variation $\int_{\Omega} |\nabla u|$ of u . With the total variation regularizer, extensive studies have shown that it does not penalize edges in u , thus it allows for sharper reconstructions, see e.g. [1, 6, 9, 15]. Among all the variational PDE based techniques, the TV minimization scheme offers one of the better combinations of noise removal and feature preservation.

It is easy to show that the TV minimization scheme leads to solving the PDE

$$(2.3) \quad \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) - \lambda(u - z) = 0.$$

But in practice, one introduces the time variable t and solve for $u(x, t)$ by time-marching the equation

$$(2.4) \quad u_t = \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) - \lambda(u - z) = 0, \quad u(x, 0) = z(x).$$

The end result $u(x, T)$, if T is large enough, will have noise removed or reduced. In essence the time-marching by (2.4) is to use gradient flow to minimize the energy $\mathbf{E}(u)$. An important attribute of the TV minimization scheme is that it takes the geometric information of the original images into account, in that it does not penalize edges. On the contrary, significant edges are sharpened. This is similar to the anisotropic diffusion methods, see [28] and the references therein.

However, the TV minimization scheme has its own weaknesses. It is well known that when (2.4) is left running for too long a denoised image will tend to become a cartoon-like piecewise constant image, wiping out all subtle details [22, 24]. With a more gentle run, one may not remove enough noise. For optimal results it is important to have an automatic stopping criterion. This is difficult to do. Although there are some attempts in this direction [19, 23, 30], some properties of the noise (such as the variance) are assumed to be known, which is not entirely realistic for natural color photos.

A modified version of the TV denoising scheme based on wavelets was introduced in Chan and Zhou [12]. Using this scheme, Wang and Zhou [29] devised an automatic stopping criterion for the time-marching process (2.4). This criterion works surprisingly well in

experiments, see [29] for a detailed discussion. In this paper we propose a new method based on the wavelet TV denoising scheme. The cartoon-like tendency in the standard TV method is a result of excessive diffusion of low frequency features in an images. The new method, which we call *Multiscale Total Variation* (MTV) method, will concentrate the diffusion on high frequency features where the noise resides. Our tests illustrate that the MTV method is highly effective, particularly for denoising natural color images. Comparing to the standard TV or wavelet TV method, the MTV method requires fewer number of iterations to obtain comparable or better results. Figure 3 shows some comparisons.

To see how MTV scheme works, let $\{\psi_j : j \in I\}$ be an orthonormal or biorthonormal wavelet basis for $L^2(\Omega)$ such as those found in [14, 27]. So we may expand any $f(x)$ as

$$f(x) = \sum_{j \in I} c_j \psi_j(x),$$

for some real (c_j) . In practice, we have always used the biorthonormal 7-9 wavelet basis as our basis $\{\psi_j\}$. (The conventional notation uses two sub-indices to denote a wavelet basis. Here we use only one for brevity. There should not be any confusion.) Now expand the observed image function $z(x)$ using the basis $\{\psi_j(x)\}$,

$$z(x) = \sum_{j \in I} \alpha_j \psi_j(x).$$

Let

$$(2.5) \quad u(x, \boldsymbol{\beta}) := \sum_{j \in I} \beta_j \psi_j(x)$$

where $\boldsymbol{\beta} = (\beta_j)$. Now we set the distance functional $D(u, z)$ to be

$$(2.6) \quad D(u, z) := \sum_{j \in I} \lambda_j (\beta_j - \alpha_j)^2,$$

where $\lambda_j > 0$. The key feature is that λ_j *decreases* as the scale becomes more localized. More precisely, we use smaller values λ_j for high frequency terms and larger values for lower frequency terms. The term $R(u)$ remains to be the total variation. Thus the MTV method is to find the minimizer of the energy functional

$$(2.7) \quad \mathbf{E}_{MTV}(u, z) := \sum_{j \in I} \lambda_j (\beta_j - \alpha_j)^2 + \int_{\mathbb{R}^2} |\nabla_x u(x, \boldsymbol{\beta})| dx.$$

where $u = u(x, \boldsymbol{\beta})$. The idea is that since λ_j are smaller for high frequency terms the smoothing is done mostly on high frequency features. The goal of denoising is to minimize



FIGURE 3. Some comparisons using artificial noise. The original image (upper-left) is added with Gaussian white noise (upper-right). The standard TV denoised image (lower-left) has more noisy residual than the wavelet TV denoised image (lower-right). Both are obtained with same number of iterations.

$F(u, z)$ and find the minimizer $u^* := u(x, \beta^*)$ such that

$$(2.8) \quad \mathbf{E}_{MTV}(u^*, z) = \min_{\beta} \mathbf{E}_{MTV}(u, z).$$

One can use simple calculus of variation to obtain the derivative of the objective functional (2.7). For $u = u(x, \beta)$ where $\beta = (\beta_j)$,

$$\begin{aligned} \frac{\partial \mathbf{E}_{MTV}(u, z)}{\partial \beta_j} &= \int_{\mathbb{R}^2} \frac{\nabla_x u}{|\nabla_x u|} \cdot \nabla_x \psi_j dx + 2\lambda_j(\beta_j - \alpha_j) \\ &= - \int_{\mathbb{R}^2} \nabla_x \cdot \left[\frac{\nabla_x u}{|\nabla_x u|} \right] \psi_j dx + 2\lambda_j(\beta_j - \alpha_j). \end{aligned}$$

Then the Euler-Lagrange equation for the model is

$$(2.9) \quad - \int_{\mathbb{R}^2} \nabla_x \cdot \left(\frac{\nabla_x u}{|\nabla_x u|} \right) \psi_j(x) dx + 2\lambda_j(\beta_j - \alpha_j) = 0.$$

In practice, rather than solving the Euler-Lagrange equation (2.9) directly for denoising, we introduce an artificial time parameter t and time-march the image using gradient flow. More precisely, we set $\beta = \beta(t) = (\beta_j(t))$ and solve the following time evolution equation,

$$(2.10) \quad \frac{\partial \beta_j}{\partial t} = \int_{\mathbb{R}^2} \nabla_x \cdot \left(\frac{\nabla_x u}{|\nabla_x u|} \right) \psi_j(x) dx - 2\lambda_j(\beta_j - \alpha_j), \quad (\beta_j(0)) = (\alpha_j).$$

The minimizer of the TV wavelet model is the steady state of the above equation. However, we often stop the process before the actual minimizer is attained because, depending on the parameter λ_j the actual minimizer can be overly smoothed while noise might be effectively removed long before that. The automatic stopping criterion in [29] for the wavelet TV method is easily applicable for the MTV method.

3. COLOR SPACE AND AUTOMATIC STOPPING CRITERION

One may argue that color images are no different from three monochromatic images once we consider the three channels separately, and therefore to denoise a color photo one only needs to denoise the three monochromatic channels separately. This view, however, misses some important subtle characteristics in naturally captured color images that, when fully utilized, yield superior results. To denoise color photos we must first understand the nature of the noise in these images, and take full advantages of all available informations.

The most commonly used color space is the RGB color space. In the RGB color space we denoise each of the three channels to complete the denoising of the color image. This approach yields unsatisfactory results, particularly for images taken under artificial lighting in which the blue channel is excessively noisy. A better way is to separate the luminance from the chrominance. There are a few ways one can achieve this. One way is to use the




FIGURE 4. The Y (left) and Cr (right) channels of the natural color image shown in Figure 2

LAB color space, which is nonlinear against the RGB color space. Another choice is the YCrCb color space. Given that YCrCb is a linear transformation of RGB it is widely used in applications such as color video and JPEG compression of color images. The advantage of separating luminance from chrominance is that human vision is typically less sensitive to diffusions in chrominance. This is illustrated in Figures 4, 5 and 6. Figure 4 and the left on Figure 5 show the YCrCb channels of the natural color image displayed in Figure 2. We then performed a rather destructive wavelet thresholding on the chrominance channels Cr (on the right of Figure 5) and Cb (on the left of Figure 6). The right on Figure 6 shows the re-composed color image. As one can see, there is very little discernable difference between the two color images. This robustness against diffusion in the chrominance does not extend, however, to the luminance channel Y. In fact, even a tiny blurring in the luminance channel will be immediately visible in the re-composed color image. Given these characteristics of the luminance-chrominance decomposition, we would want to be more aggressive in denoising the chrominance channels while less so in denoising the luminance channel.

The problem with the standard luminance-chrominance decomposition, such as LAB and YCrCb, is that the luminance is “contaminated” by the blue channel, where noise concentrates as we have pointed out earlier. As a result the luminance channel can be somewhat noisy, and therefore substantial denoising will often have to be performed on




FIGURE 5. The Cb channel (left) of the natural color image shown in Figure 2. Using wavelet thresholding to severely blur the chrominance Cr channel (right), in which most of the details are removed.




FIGURE 6. The other chrominance channel Cb (left) is also severely blurred by wavelet thresholding. While the re-composed color image (right) after blurring the CrCb channels seems to be a reasonable approximation to the original one.

it. This can adversely affect the quality of the denoised color image. To get around this problem we introduce a new color space, the *modified YCrCb* color space (m-YCrCb color

space). In the m-YCrCb space, the ‘‘luminance’’ channel is a linear combination of only the green and the red channels. More precisely, the m-YCrCb color space is obtained via the following linear transform from the RGB color space:

$$\begin{aligned} Y_m &= 0.666G + 0.334R, \\ Cr_m &= (R - Y)/1.6, \\ Cb_m &= (B - Y)/2. \end{aligned}$$

To denoise a color image we perform the MTV denoising scheme on each of the three channels Y_m , Cr_m , and Cb_m . An automatic stopping criterion is applied in the MTV scheme. Since the luminance channel in the m-YCrCb color space no longer contains any part of the blue channel it is usually much cleaner. The automatic stopping criterion, which we describe in details below, will stop the process for the Y_m channel after only a few iterations. The denoising of the Cr_m channel also takes only a few iterations for the very same reason. The process takes much longer in general for the Cb_m channel, in which noise is concentrated.

We now describe our automatic stopping criterion for the MTV denoising scheme with the wavelet basis $\{\psi_j : j \in I\}$. Again we remark that the conventional notation for wavelet bases use two or more indices, such as $\{\psi_{jk}\}$. In this paper we only use one index for conciseness, and there should not be any confusion. Like in the wavelet hard thresholding scheme, we first choose a threshold $\rho > 0$. Let $J_\rho = \{j \in I_D : |\beta_j(0)| = |\alpha_j| \leq \rho\}$, where $I_D \subset I$ is the index set corresponding to the diagonal portion of the highest frequency wavelet coefficients. Intuitively speaking, as in the wavelet hard thresholding scheme, the coefficients $\{\beta_j(0) : j \in J_\rho\}$ will indicate how noisy the image is. In a noise-free image these wavelet coefficients will mostly be very close to 0. But in a noisy image they will be mostly not close to zero. Define $\mu(t) = \frac{1}{|J_\rho|} \sum_{j \in J_\rho} |\beta_j(t)|$. So $\mu(t)$ measures the noise in the image at time t . The key idea is that an automatic stopping criterion of the time evolution can be designed by measuring the reduction in the value $\mu(t)$ from the original value $\mu(0)$.

In [29] we have described two different automatic stopping criteria, the *relative criterion* and the *absolute criterion*. Both of these can be adopted to the MTV scheme. For the relative automatic stopping criterion, we consider $\mu(t)/\mu(0)$. We will stop the time evolution whenever this value goes below a threshold b . For example, we may set $b = 0.1$. This threshold intuitively says that we stop the time evolution when we have reduced noise by

90%. For the absolute automatic stopping criterion, we stop the time evolution if $\mu(t)$ drops below a threshold c . Since in a noise-free image we expect $\mu(t)$ to be very close to zero, it is reasonable to set an absolute threshold for $\mu(t)$ to achieve a desired denoising effect.

In the actual implementation the value ρ does not seem to affect the automatic stopping time sensitively. We usually take $\rho = \frac{2}{|I_D|} \sum_{j \in I_D} |\alpha_j|$. Both the relative criterion and the absolute criterion work well, although we typically use the relative criterion. For an image with moderate noise we set the threshold b to be between 0.05 and 0.1. In more noisy cases such as the one in Figure 2, we use smaller threshold b around 0.03. It should be pointed out that we have tested the automatic stopping time criterion on a number of noisy color images as well as monochromatic medical images. The thresholds for optimal performance stayed remarkably consistent. This is an important attribute for batch processing medical images.

4. EXAMPLES

In this section we compare various denoising schemes using the noisy color image shown in Figure 2. These schemes include the wavelet hard thresholding, the wavelet soft thresholding, and MTV. We first compare the denoising on the RGB color space in Figures 7 and 8. As one can see, given the severity of the noise in the blue channel none of them has performed well, at least not well enough to be a serious tool practically in digital photography. We next show the denoising results on the m-YCrCb space (Figure 9). Clearly in the new color space there is a rather substantial improvement in performance. The reconstructed blue channel is perhaps the best indicator of the effectiveness of the denoising. The MTV method yields a superb reconstructed blue channel that is essentially noise free with all details maintained.

REFERENCES

- [1] R. Acar and C. R. Vogel. Analysis of total variation penalty methods for ill-posed problems. *Inverse Prob.*, 10:1217–1229, 1994.
- [2] A. Buades, B. Coll and J. M. Morel, A review of image denoising algorithms, with a new one, *Multiscale Model. Simul.*, 4 (2005), no. 2, 490–530 (electronic).
- [3] E. Candès and D. Donoho, *Curvelets - A Surprisingly Effective Nonadaptive Representation for Objects with Edges, Curves and Surfaces*, L. L. Schumaker et al. Eds, Vanderbilt University Press, Nashville, TN. 1999.
- [4] E. J. Candès, and F. Guo. *Edge-preserving Image Reconstruction from Noisy Radon Data*, (Invited Special Issue of the *Journal of Signal Processing on Image and Video Coding Beyond Standards.*), 2001.



FIGURE 7. The denoised images by wavelet hard (left) and soft (right) thresholdings in the RGB space. Either noticeable noise still exists due to high noise in blue channel, or the image is excessively smeared.



FIGURE 8. The denoised image by MTV in the RGB space (left) and its blue channel (right). Noticeable noise still exists due to high noise in blue channel even the recomposed image has most of the noise removed

- [5] A. Chambolle, R. A. DeVore, N.-Y. Lee, and B. J. Lucier. *Nonlinear wavelet image processing: variational problems, compression and noise removal through wavelet shrinkage*. IEEE Trans. Image Processing, 7(3):319–335, 1998.
- [6] A. Chambolle and P. L. Lions. Image recovery via Total Variational minimization and related problems. *Numer. Math.*, 76:167–188, 1997.



FIGURE 9. The denoised image by MTV in m-YCrCb color space (left) and its blue channel (right) which is much cleaner than the original one.

- [7] R. Chan, T. Chan, L. Shen, and Z. Shen, Wavelet algorithms for high-resolution image reconstruction, *SIAM J. Sci. Comput.*, 24 (2003), no. 4, 1408–1432.
- [8] T. F. Chan, S. H. Kang and J. Shen, *Euler's Elastica and Curvature Based inpainting*, *SIAM J. Appl. Math.*, 63(2) (2002), 564-592.
- [9] T. F. Chan, S. Osher, and J. Shen, *The Digital TV Filter and Nonlinear Denoising*, *IEEE Trans. Image Process.*, 10(2), pp. 231-241, 2001.
- [10] R. Chan, S. Riemenschneider, L. Shen and Z. Shen. *Tight Frame: the efficient way for high-resolution image reconstruction*. *Applied and Computational Harmonic Analysis*, 17(1) (2004), pp91-115.
- [11] T. F. Chan and H. M. Zhou, *Total Variation Wavelet Thresholding*, to appear in *J. Scientific Computing*.
- [12] T. F. Chan and H. M. Zhou, *Optimal Constructions of Wavelet Coefficients Using Total Variation Regularization in Image Compression*, CAM Report, No. 00-27, Dept. of Math., UCLA, July 2000.
- [13] R. Coifman and D. Donoho, Translation-invariant de-noising, in *Wavelets and Statistics*, 125-150, Springer Verlag, 1995.
- [14] I. Daubechies. *Ten lectures on wavelets*. SIAM, Philadelphia, 1992.
- [15] D. Dobson and C.R. Vogel, *Convergence of An Iterative Method for Total Variation Denoising*, *SIAM Journal on Numerical Analysis*, 34 (1997), pp. 1779-1971.
- [16] D. L. Donoho. *De-noising by soft-thresholding*. *IEEE Trans. Information Theory*, 41(3):613–627, 1995.
- [17] D. L. Donoho and I. M. Johnstone. *Ideal spacial adaption by wavelet shrinkage*. *Biometrika*, 81:425–455, 1994.
- [18] S. Durand and J. Froment, *Artifact Free Signal Denoising with Wavelets*, in *Proceedings of ICASSP'01*, volume 6, 2001, pp. 3685-3688.
- [19] G. Gilboa, N. Sochen and Y. Zeevi, Estimation of optimal PDE-based denoising in the SNR sense, *UCLA CAM Report 05-48*, August, 2005.
- [20] M. Mahmoudi and G. Sapiro, Fast image and video denoising via nonlocal means of similar neighborhoods, *IEEE Signal Processing Letter*, 12 (2005), 839-842.
- [21] F. Malgouyres, *Mathematical Analysis of a Model Which Combines Total Variation and Wavelet for Image Restoration*, *Journal of information processes*, 2:1, 2002, pp 1-10.

- [22] Y. Meyer. *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*, volume 22 of *University Lecture Series*. AMS, Providence, 2001.
- [23] P. Mrazek and M. Navara, Selection of optimal stopping time for nonlinear diffusion filtering, *IJCV*, v.52, no. 2/3, 2003, pp189-203.
- [24] S. Osher, A. Sole, and L. Vese, *Image decomposition and restoration using total variation minimization and H_{-1} Norm*, *Multiscale Modeling and Simulation* 1(3), pp349-370, 2003.
- [25] J Portilla, V Strela, M Wainwright, E P Simoncelli. *Image Denoising using Scale Mixtures of Gaussians in the Wavelet Domain*. *IEEE Transactions on Image Processing*. vol 12, no. 11, pp. 1338-1351.
- [26] L. Rudin, S. Osher and E. Fatemi, *Nonlinear Total Variation Based Noise Removal Algorithms*, *Physica D*, Vol 60(1992), pp. 259-268.
- [27] G. Strang and T. Nguyen. *Wavelets and Filter Banks*. Wellesley-Cambridge Press, Wellesley, MA, 1996.
- [28] Y. Sun, P. Wu, G.W. Wei and G. Wang. *Evolution-Operator-Based Single-Step Method for Image Processing*. *International Journal of Biomedical Imaging*, Vol. 2006, pp 1-28.
- [29] Y. Wang and H.-M. Zhou, Total Variation Wavelet Based Medical Image Denoising, *Int. J. Biomedical Imaging*, Vol 2006 (2006), 1-6.
- [30] J. Weickert, Coherence-enhancing diffusion of colour images, *IVC*, Vol 17, 1999, pp 201-212.
- [31] T. Weissman, M. Weinberg, E. Ordentlich, G. Seroussi and S. Verdú, A discrete universal denoiser and its application to binary images, *Proc. IEEE ICIP*, 1 (2003), 117-120.
- [32] L. P. Yaroslavsky, *Digital Picture Processing – An Introduction*, Springer Verlag, 1985.
- [33] L. P. Yaroslavsky and M. Eden, *Fundamentals of Digital Optics*, Birkhauser, Boston, 1996.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA AT LOS ANGELES, LOS ANGELES, CA 90095-1555

E-mail address: `tfchan@nsf.gov`

SCHOOL OF MATHEMATICS, GEORGIA INSTITUTE OF TECHNOLOGY, ATLANTA, GEORGIA 30332, USA.

E-mail address: `wang@math.gatech.edu`

SCHOOL OF MATHEMATICS, GEORGIA INSTITUTE OF TECHNOLOGY, ATLANTA, GEORGIA 30332, USA.

E-mail address: `hmzhou@math.gatech.edu`