UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Total Variation Denoising and Enhancement of Color Images Based on the CB and HSV Color Models

Tony F. Chan
Sung Ha Kang
Jianhong Shen

June 2000

CAM Report 00-25

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90095-1555

http://www.math.ucla.edu/applied/cam/index.html

Total Variation Denoising and Enhancement of Color Images Based on the CB and HSV Color Models

Tony F. Chan, Sung Ha Kang and Jianhong Shen *

Abstract

Most denoising and enhancement methods for color images have been formulated on linear color models, namely, the channel-by-channel model and vectorial model. In this paper, we study the total variation (TV) restoration based on the two nonlinear (or non-flat) color models: the Chromaticity-Brightness (CB) model and Hue-Saturation-Value (HSV) model. These models are known to be closer to human perception. Recent works on the variational/PDE method for non-flat features by several authors enable us to denoise the chromaticity and hue components directly. We present both the mathematical theory and digital implementations for the TV method. Comparison to the traditional TV restorations based on linear color models is made through various experiments.

1 Introduction

Any tool that attempts to denoise and enhance digital color images must rely on two dependent ingredients — the representation of colors (i.e. the color model) and the restoration model formulated on that representation. Therefore, we first discuss briefly the color models, followed by an introduction to the variational restoration method.

1.1 RGB, CB, and HSV color models

In image processing, color has been represented or modeled in various ways [8]. In this paper we shall focus on the RGB model and HSV model.

In the RGB representation of color images, at each pixel $p = (x, y)$, the vectorial value $I(p) = (u_1(p), u_2(p), u_3(p))$ represents the intensity of the three primary colors separately. Each monochromatic component u_i is called one

*Research was supported by NSF Grant DMS-9626755 and by ONR under N00014-96-1-0277. The authors are with the Department of Mathematics, UCLA, Los Angeles, CA 90095-1555, USA, ({chan,skang}@math.ucla.edu), and School of Mathematics, University of Minnesota, Minneapolis, MN 55455, USA, (jhshen@math.umn.edu).

channel. The RGB model has inspired the so-called CB model, which decomposes an RGB pixel value $I(p)$ into two components — the brightness component $u(p) = \|I(p)\|$, or the Euclidean length of the vector, and the chromaticity component $f(p) = I(p)/u(p)$. The chromatcity component lives on the unit sphere S^2. Features like this that live on non-linear manifolds are said to be *non-flat* [7, 16]. It has been shown by several authors lately that the CB model is well suited for denoising, edge detection and enhancement, and segmentation [7, 15, 16, 17, 18].

HSV is another color system that is believed to be more natural than the RGB systerm for human perception. The three variables are: hue H, saturation S, and value V. The S and V are linear features and take values in the interval $[0, 1]$. S encodes the "purity" of color: larger S corresponds to purer color. The value V stores the intensity information, so that larger V value means brighter color. The hue variable H, though also takes values in $[0, 1]$, is nevertheless a circular or periodic feature: as H increases from 0 to 1, the color spectrum revolves from red, yellow, green, cyan, blue, to magenta, and eventually back to red. Therefore, the hue lives on the unit circle S^1 in some sense, and is another example of non-flat features [9]. The circularization is easily made by the exponential mapping $H \to \exp(i2\pi H)$.

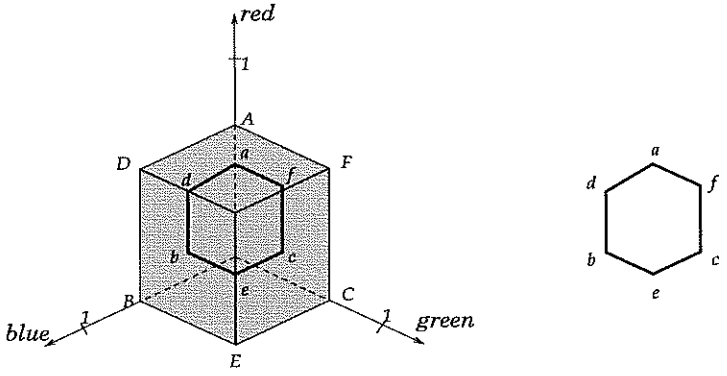

An iso-V surface in the RGB space consists of three 2-D squares.

An iso-S line on the iso-V surface. From a, d, ..., to f, a, H increases from 0 to 1.

Figure 1: An iso-V surface and an iso-S line on it in the RGB space. See Eq. (1,2) for the definition. Notice that an iso-V & S line is a 3-D "hexagon," as shown separately on the right. From $a, d, \cdots$ to f, a, the hue value H increases from 0 to 1.

The precise transition from the $[r, g, b]$ variables to $[H, S, V]$ is realized by:

$$V = r \vee g \vee b \tag{1}$$

$$S = 1 - \frac{r \wedge g \wedge b}{r \vee g \vee b}, \tag{2}$$

where for a pair of real numbers a and b: $a \vee b := \max(a, b)$, $a \wedge b := \min(a, b)$. The formula for H appears to be more complicated, but essentially is a simple piecewise linear function along any iso-S & V line (See Fig. 1). In the figure, one observes a typical iso-S & V line, which is a 3-D "hexagon" shown separately on the right side. As one goes from $a, d, \cdots$ to f, a, the H value is defined so that its value increases linearly from 0 to 1. Thus on each one of the six segments, the net increment is 1/6. With this in mind, The formula H can be expressed easily. For examples, along the segment $[a, d]$ (along which r and g are fixed)

$$H = H(b) = \frac{1}{6}\left[\frac{b-g}{r-g}\right];$$

and along $[b, e]$ (along which r and b are fixed)

$$H = H(g) = \frac{2}{6} + \frac{1}{6}\left[\frac{g-r}{b-r}\right].$$

Readers can easily figure out the formulas for the rest.

Currently, many color image processing tools are based on the RGB model, mostly because linear spaces are easy to work with. As a result, despite their similarity to human color perception, the CB and HSV color models have been less favored due to their non-flatness. In this paper, our main interest is to construct restoration models based on these two non-flat color representations.

1.2 Statistics, least squares, and the variational method

For denoising and enhancement of color images, the two classical mathematical tools are statistics and the least square estimation.

Order statistics plays a crucial role in the denoising and enhancement of both gray and color images, and its most successful application is the *median filter* [8, 17, 18]. Median filters perform especially well for eliminating outliers and enhancing edges. Recently median filters have been generalized to the chromaticity feature f on the unit sphere [17, 18]. From the numerical point of view, as in Monte Carlo simulations, the implementation of statistical algorithms is often comparatively inefficient.

The least square estimation is another popular tool throughout image and signal processings. While benefiting from its simplicity in digital implementation and advantages in numerical linear algebra, we suffer seriously from its linear nature. The square norm is notoriously unsuited for jumps in 1-D signals and edges in 2-D images. Denoising models based on the least square estimation inevitably blur sharp edges, which exist universally in gray or color images, and are crucial for the human perception of image structures.

In the past decade, the variational and PDE method has attracted much attention in image processing because of its flexibility in modeling and various advantages in numerical PDEs. Applications can be widely found in image segmentation, denoising, deblurring, enhancement, inpainting, and motion estimation (see the two monographs [10, 19], for examples).

It has been known that by proporly choosing the right cost or energy functionals in the variational formulation or the right diffusion patterns in the PDE formulation, edges in noisy images can be well restored and enhanced. Among the many possible choices, the *Total Variation* (TV) is the simplest but sufficiently efficient measurement for enhancement or denoising. Its PDE form (via the Euler-Lagrange equation) is interestingly connected to anisotropic diffusions and the mean curvature motion [10]. Ever since the first explicit demonstration of its successful performance in image restoration was made by Rudin, Osher and Fatemi [12, 13], the TV model has been studied by many authors (see the References section for examples).

Let us first explain the general framework of the variational approach for denoising and enhancing color images based on the linear RGB color models. The traditional methods can be classified into two categories — the channel-by-channel approach and the vectorial approach.

In the channel-by-channel approach, each channel u_i is assumed to be contaminated by noise η_i so that the observation becomes $u_i^0(x,y) = u_i(x,y) + \eta_i(x,y)$. A typical channel-by-channel denoising model carries the form of

$$\min_{u_i} R_i(u_i) \quad \text{subject to} \quad \frac{1}{|\Omega|}\int_\Omega |u_i - u_i^0|^2 = \sigma_i^2, \tag{3}$$

where Ω is the image domain, $|\Omega|$ its area, R_i the regularity functional, and σ_i the noise level. For example, in the total variational approach, we take

$$R_i(u_i) = \int_\Omega |\nabla u_i| dxdy.$$

Such regularity functional has been proven both mathematically and computationally to be capable of extracting edges snowed by noise [2, 13].

In the vectorial denoising approach, the cost functional is on the vectorial function $I = (u_1, u_2, u_3)$:

$$\min_{I} R_{3D}(I) \quad \text{subject to} \quad \frac{1}{|\Omega|}\int_\Omega \|I - I^0\|^2 = \sigma^2, \tag{4}$$

where R_{3D} is a regularity functional for vector-valued functions. Typically in applications, suppose R_i is a suitable 1-dimensional regularity functional for the i-th channel, then one simply takes $R_{3D} = \sqrt{R_1^2 + R_2^2 + R_3^2}$. In the case when R_{3D} is connected to the total variation functional and anisotropic diffusions, we observe the work of Sapiro and Ringach [14] and Blomgren and Chan [1, 3].

In practice, both of the constrained optimization problems (3) and (4) are replaced by their unconstrained forms. For example, for the vectorial case, we solve instead

$$\min_I \left(R_{3D}(I) + \frac{\lambda}{2}\int_\Omega \|I - I^0\|^2 \right).$$

Here λ is an appropriate Lagrangian multiplier, which depends on the noise level. In practice, it is often estimated or chosen *a priori* [3, 4, 5, 13].

1.3 Variational restoration for non-flat features

Recently, the variational and PDE method has been generalized to non-flat features by Perona [11], Tang, Sapiro and Casselles [15, 16], Chan and Shen [7], and most recently by Kimmel and Sochen [9]. Tang et al. and Chan et al. generalize the total varitional model, while Kimmel et al.'s model profits from the general framework of Polykov action and Beltrami operators for Riemannian manifolds. All these works attempt to successfully restore edges in non-flat features. Chan and Shen [6] also proved that near edges, the models in [7] and [9] are equivalent. But the differential equations based on the total varitation norm are much simpler and thus easier for computation [7, 15, 16].

The current paper can be seen as the extension and completion of the works of [7, 15, 16] on the total variation approach for enhancing and denoising non-flat image features. We study in details the TV model and its numerical implementation for restoring color images based on the CB representation and HSV color model. Previous works on general non-flat features enable us to work with the chromaticity feature f (on the sphere S^2) and hue feature H (along the circle S^1) directly. Detailed comparison to the channel-by-channel and vectorial approaches are illustrated through numerical examples. Our results show convincingly the advantages of the CB and HSV color models over those linear ones for the construction of restoration models.

The paper is organized as follows. Section 2 introduces the mathematical models, and Section 3 details their numerical implementations. Numerical experiments and comparison are explained in Section 4.

2 The TV Formulation on CB and HSV Models

In this section, we explain the total variation formulation for the two nonlinear color models.

2.1 TV for the CB model

In the RGB representation, a color image is a mapping

$$I : \Omega \to R^3_+ = \{(r, g, b) : r, g, b > 0\}.$$

I can be separated to the *brightness component* $u = \|I\|$, and the *chromaticity component* $f = I/\|I\| = I/u$. Mathematically speaking, this is the implicit spherical coordinates for the Euclidean space. We shall not introduce the two Eulerian angles (i.e., the latitude and longitude) for f, thanking to our previous work on non-flat features [7].

The brightness u can be treated as a gray image, and thus any scalar denoising model to can be applied. The chromaticity component f stores the major color information, and is "non-flat" since it takes values on the unit sphere S^2. We thereby apply the general framework of non-flat total variational denoising model as studied in details by Chan and Shen [7]. In this chromaticity-brightness

approach, the final restored color image I is assembled from the two restored components:

$$I(p) = u(p) \times f(p).$$

We now detail the mathematics.

Given a noisy image I^0, let $u^0 : \Omega \to R$ be its brightness component. Then the (linear) scalar TV restoration model applied to u^0 is

$$\min_u \quad \int_\Omega |\nabla u| \, dx \, dy + \frac{\lambda}{2} \int_\Omega (u - u^0)^2 \, dx \, dy,$$

where λ is the Lagrange relaxation parameter. The associated Euler-Lagrange equation (in a formal level) of this cost functional is

$$-\nabla \cdot (\frac{\nabla u}{|\nabla u|}) + \lambda(u - u^0) = 0. \tag{5}$$

To avoid possible singularity for $|\nabla u|$ in the denominator, we condition it to $|\nabla u|_a = \sqrt{|\nabla u|^2 + a^2}$ for some small a in implementation. This is equivalent to minimizing directly

$$\int_\Omega |\nabla u|_a \, dx \, dy + \frac{\lambda}{2} \int_\Omega (u - u^0)^2 \, dx \, dy.$$

The total variation as a cost functional legalizes the existence of edges or sharp jumps in the brightness component, which is very important since the sharp boundaries of objects in images can be well restored [13].

Now we discuss how to restore the noisy chromaticity component f^0. Let $f : \Omega \to S^2$ be a general chromaticity feature. Assume that f is smooth. Then $\partial_x f(p)$ and $\partial_y f(p)$ are two tangent vectors in the tangent space $T_{f(p)}S^2$. Let $\|\cdot\|$ be the induced Riemannian norm of $T_{f(p)}S^2$ in R^3. Then the *total variation* for chromaticity is defined to be [7]

$$\mathcal{E}^{TV}(f) = \int_\Omega e(f;p) \, dp = \int_\Omega \sqrt{\|\partial_x f(p)\|^2 + \|\partial_y f(p)\|^2} dp, \quad dp = dx \, dy.$$

Here since the unit sphere S^2 is embedded in R^3 and $f = (f_1, f_2, f_3) \in R^3$, we have

$$e(f;p) = \|\nabla f\| = \sqrt{\|f_x(p)\|^2 + \|f_y(p)\|^2} = \sqrt{\|\nabla f_1\|^2 + \|\nabla f_2\|^2 + \|\nabla f_3\|^2}.$$

Let d be any reasonable distance between two chromaticity points on S^2 [7]. For example, one can choose the *embedded distance* or the cord distance

$$d(f,g) := \|f - g\|_{R^3} = \sqrt{(f-g)^2}, \quad f, \, g \in S^2. \tag{6}$$

Another natural choice would be the *geodesic distance* or arc distance on S^2:

$$d(f,g) = \arccos \langle f, g \rangle,$$

where $\langle f, g \rangle$ denotes the inner product in R^3. Then the TV restoration model becomes

$$\min_f \ \mathcal{E}^{TV}(f) \quad \text{subject to} \quad \frac{1}{|\Omega|}\int_\Omega d^2(f^0, f)dp = \sigma^2.$$

The Euler-Lagrange equation for its unconstrained variational formulation is given by:

$$-\partial_x^*\left(\frac{\partial_x f}{\|\nabla f\|}\right) - \partial_y^*\left(\frac{\partial_y f}{\|\nabla f\|}\right) + \frac{\lambda}{2}\text{grad}_f\ d^2(f^0, f) = 0,$$

where $\text{grad}_f d^2(f^0, f)$ denotes the gradient vector of the scalar function $d^2(f^0, f)$ on S^2 and ∂_x^* and ∂_y^* the covariant derivatives acting on vector fields on the sphere. According to [7],

$$\partial_x^*\left(\frac{\partial_x f}{\|\nabla f\|}\right) + \partial_y^*\left(\frac{\partial_y f}{\|\nabla f\|}\right) = \nabla \cdot \left(\frac{\nabla f}{\|\nabla f\|}\right) + \|\nabla f\| f.$$

Furthermore, if we choose the embedded distance (6), then

$$\frac{1}{2}\text{grad}_f d^2(f^0, f) = \Pi_f(f - f^0) = -\Pi_f f^0.$$

Here Π_f is the orthogonal projection from $T_f R^3$ onto the tangent plane $T_f S^2$:

$$\Pi_f g = g - \langle g, f \rangle f, \quad \text{for any vector} \quad g \in T_f R^3.$$

Eventually, the restoration equation for the TV model is

$$f_t \ \ (\text{or } \ 0) \ \ = \nabla \cdot \left(\frac{\nabla f}{\|\nabla f\|}\right) + \|\nabla f\| f + \lambda \Pi_f f^0. \tag{7}$$

We shall explain in Section 3 the numerical implementation or the digital version of the non-linear restoration equations (5) and (7).

2.2 TV for the HSV model

In the HSV color model, the hue feature H is non-flat and lives on the unit circle S^1, while the other two are both linear or flat ones. Therefore, for the saturation S and value V, we simply apply the scalar (flat) TV denoising model (5). For the circular feature H, one copies the equation for chromaticity (7) based on the straightforward modification from S^2 to S^1:

$$H = (H_1, H_2) \in S^1 \subset R^2 \quad \text{and} \quad \|\nabla H\| = \sqrt{\|H_x\|^2 + \|H_y\|^2}.$$

Our experiment (Fig. 7) shows that denoising separately the hue component H and the saturation component S does not produce satisfactory visual results. Our explanation for this deficiency is that to human vision, the two components

are highly correlated. Therefore, a better approach is to restore the combination of H and S, which in the classical literature of color image processing, is also called chromaticity. In this paper, to distinguish it from the chromaticity appearing in the CB representation, we name it the "HS-chromaticity."

As we see from Fig. 1, in the RGB color space, for a fixed value V, the other two variables H and S span the iso-V surface which consists of three 2-D squares. The surface is not smooth since there are a corner and three folding lines. Working with such irregular surfaces is inconvenient. Thus, we invent the so-called *disk transform* to "straighten" this folded surface as follows. For each point on the iso-V surface with the hue and saturation (H, S), define a complex number

$$Z = S \times \exp(i2\pi H). \tag{8}$$

Since both H and S take all values in $[0, 1]$, the image of each iso-V surface is the unit disk in the complex plane (see Fig. 2). Under the disk transform, each iso-S line on an iso-V surface (i.e., the 3-D hexagon on the right side of Fig. 1) is mapped onto a circle centered at the origin (see Fig. 2).

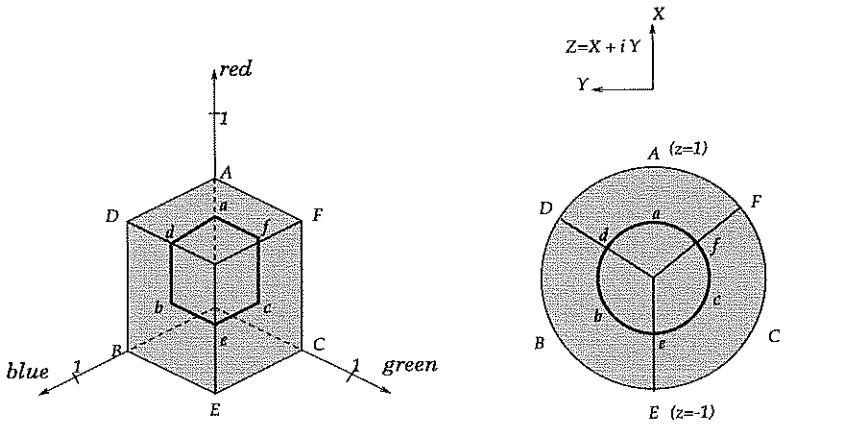

Figure 2: The disk transform of an iso-V surface: $Z = S\exp(i2\pi H)$.

The unit disk is a smooth and convex domain and thus much easier to work with. To restore the combination of H^0 and S^0 in the HSV color model, is now equivalent to restoring a complex-valued *scalar* feature Z^0. More precisely, for any given noisy image under the HSV model

$$(H^0(x,y), S^0(x,y), V^0(x,y)), \quad (x,y) \in \Omega,$$

it is sufficient to denoise two *scalar* components — the real-valued signal $V^0(x,y)$, and the complex-valued signal $Z^0(x,y)$. Since both are *scalar* functions instead of vectorial or non-flat ones, we can apply the same scalar TV model (5). The

numerical results shall convince us the advantage of working with the combination of H and S, or their disk transform.

3 Numerical Implementation

To digitally implement the above non-linear differential equations, we have applied the approach of *digital TV filtering* as proposed in [5]. The digital TV filter can been seen as the finite difference realization of the differential equations, but is simpler and more self-contained [5].

Let Ω be a discrete digital domain or a graph. Pixels in Ω are denoted by $\alpha, \beta, \cdots$. In the conventional rectangular setting, it is also denoted by $\alpha = (ij)$. Denote N_α the neighbors of the pixel α. For instance, in the rectangular setting,

$$N_\alpha = N_{(ij)} = \{\ (i+1,j),\ (i-1,j),\ (i,j-1),\ (i,j+1)\ \}.$$

Of course, there is much freedom in defining the neighbors. For example, one may include $(i \pm 1, j \pm 1)$ in $N_{(ij)}$. If $\beta \in N_\alpha$, we also write $\beta \sim \alpha$.

Let $u : \Omega \to R$ be the brightness component. We define its local variation to be $|\nabla_\alpha u| = \sqrt{\sum_{\beta\sim\alpha}(u_\beta - u_\alpha)^2}$, and its conditioned form is $|\nabla_\alpha u|_a = \sqrt{\|\nabla_\alpha u\|^2 + a^2}$ for some small a.

By introducing the weights $\omega_{\alpha\beta}(u) = \frac{1}{|\nabla_\alpha u|_a} + \frac{1}{|\nabla_\beta u|_a}$, Chan, Osher, and Shen [5] showed that the Euler-Lagrange equation (5) is replaced by a system of nonlinear algebraic equations:

$$\sum_{\beta\sim\alpha} \omega_{\alpha\beta}(u_\beta - u_\alpha) + \lambda(u - u^0) = 0, \quad \alpha \in \Omega.$$

The techniques of linearization and iteration scheme lead to the so-called digital TV filter [5]. The digital TV filter $\mathcal{F}$ is nonlinear data dependent filter $\mathcal{F} : u \to v$ once the noisy image u^0 is given. At any pixel $\alpha \in \Omega$,

$$v_\alpha = \mathcal{F}_\alpha(u) = \sum_{\beta\sim\alpha} h_{\alpha\beta}(u)u_\beta + h_{\alpha\alpha}(u)u^0_\beta, \tag{9}$$

where the low-pass filter coefficients are

$$h_{\alpha\beta}(u) = \frac{\omega_{\alpha\beta}(u)}{\lambda + \sum_{\gamma\sim\alpha}\omega_{\alpha\gamma}(u)}, \qquad h_{\alpha\alpha}(u) = \frac{\lambda}{\lambda + \sum_{\gamma\sim\alpha}\omega_{\alpha\gamma}(u)}.$$

The digital TV filter is applied in an iterative fashion. To restore and denoise the brightness component u^0, one starts with a random guess $u^{(0)}$ (for example, $u^{(0)} = u^0$, conveniently though unnecessarily), and then generate $u^{(n)} = \mathcal{F}(u^{n-1})$ for $n = 1, 2, \cdots$. $u^{(n)}$ converges to the optimal restoration u.

We now discuss how to apply the digital TV filter to the chromaticity component. Let $f^0 : \Omega \to S^2$ be the noisy chromaticity. Taking the embedded

distance $d(f,g) = \sqrt{(f-g)^2}$ for any $f, g \in S^2$, we define the local variation at a pixel α of a chromaticity feature $f : \Omega \to S^2$ to be:

$$e(f;\alpha) = \left[\sum_{\beta \sim \alpha} d^2(f_\beta, f_\alpha) \right]^{\frac{1}{2}}.$$

The digitized total variation plus the fitting constraint becomes [5]

$$\mathcal{E}^{TV}(f,\lambda) = \sum_{\alpha \in \Omega_n} e(f;\alpha) + \lambda \sum_{\alpha \in \Omega_n} \frac{1}{2} d^2(f_o, f),$$

which provides the target cost function for optimization. The digital Euler-Lagrange equation is shown to be [7]:

$$0 = \sum_{\beta \sim \alpha} \Pi_{f_\alpha}(f_\beta) \left(\frac{1}{e(f;\alpha)} + \frac{1}{e(f;\beta)} \right) + \lambda \Pi_{f_\alpha}(f^0_\alpha), \quad \alpha \in \Omega. \tag{10}$$

Here Π_f is the orthogonal projection defined in the previous section. By setting

$$\omega_{\alpha\beta}(f) = \frac{1}{e(f;\alpha)} + \frac{1}{e(f;\beta)},$$

we rewrite the restoration equation (10) as

$$\Pi_{f_\alpha}\left(\sum_{\beta \sim \alpha} \omega_{\alpha\beta} f_\beta + \lambda f^0_\alpha\right) = 0, \quad \alpha \in \Omega.$$

Chan and Shen [7] showed that this equation on the unknown optimal restoration $f : \Omega \to S^2$ can be similarly solved by the digital TV filter $\mathcal{F}$. Unlike the brightness component, since the feature lives on the unit sphere, $g = \mathcal{F}(f)$ now needs an extra step of projection:

$$\begin{aligned} \tilde{g}_\alpha &= \textstyle\sum_{\beta \sim \alpha} h_{\alpha\beta} f_\beta + h_{\alpha\alpha} f^0_\alpha; \\ g_\alpha &= \tilde{g}_\alpha / \|\tilde{g}_\alpha\|. \end{aligned}$$

One starts the iterative filtering process with an initial guess $f^{(0)}$, then generate $f^{(n)} = \mathcal{F}(f^{(n-1)})$. Chan and Shen [7] showed that the limit of $f^{(n)}$ indeed solves the restoration equation (10).

For the TV restoration model based on the HSV color system, we apply the scalar TV filter (9) to both the real scalar function V^0, and the disk transform Z^0 of H^0 and S^0, which is a complex scalar function. Since both the unit interval and unit disk are convex domains, the *maximum principle* of the TV filter [5] guarantees that the restored V takes values in $[0,1]$, and Z on the unit disk. Thus, the restored H and S can be well recovered from Z.

4 Numerical Experiments and Comparison

This section summarizes the performance of the TV models based on the CB and HSV nonlinear color models.

4.1 TV restoration based on the CB model

In Fig. 3, we demonstrate the result of TV restoration applied to the chromaticity-brightness representation. In the middle, we have plotted the image with chromaticity restored only, and at the bottom, the image with both chromaticity and brightness restored. The result is quite successful. The visible noisy red and green dots have been swept out. The eyes and dark lines resume their original black color, and the nose and lips become smoothly red as they should be.

4.2 Comparison of TV restorations based on the CB and linear color models

In Fig. 4, we compare the restortation results by the CB based TV and linear TV's. The two columns on the right show the details of the first column. Compared to the channel-by-channel TV and vectorial TV, the chromaticity-brightness TV restoration seems to give better color control. This example convinces us the advantage of working with the chromaticity and brightness components, and also the fact that the CB representation is closer to human color perception.

Fig. 5 and 6 show another example for the comparison between the linear TV's and the CB based TV. The column on the right in Fig. 5 zooms into that on the left. We see that the author's stamp has been best restored by the CB TV. The linear TV's have somehow blurred the stamp, mostly due to their inefficiency in dealing with the chromaticity component. Fig. 6 shows a 1-dimensional slice of the stamp, from which we clearly see that the CB based TV restoration leads to the best coordination among the three channels: the original stamp is purely gray (i.e. having the idential channels), and thus any ideal restoration should restore such purity.

4.3 TV restoration directly based on the HSV model

In Fig. 7, we show the restoration result from the direct application of the TV models on the three components in the HSV color representation. The zoom-in image clearly shows the unsatisfactory behavior of such an approach. The underlying reason, we believe, is the high correlation between the hue component H and the saturation S for human perception.

4.4 TV restoration based on the HSV model and its disk transform

As discussed in the preceding sections and seen from the previous numerical example, we are led to considering the combination Z of H and S via the disk transform (8). In Fig. 8, we display the restoration result from the scalar TV on both the real function V and the complex function Z. From the zoom-in image, it is clear that such combination is much closer to human perception and thus yields better color restoration.

4.5 TV on CB and HSV under the disk transform: similar performance

The disk transform Z in the HSV representation encodes both of the two types of color information—hue and saturation, and thus is similar to the chromaticity information in the CB representation. Meanwhile, the "value" component V in the HSV system apparently plays the same role as the brightness component in the CB model. As a result, to no one's surprise, the performance of these two approaches should be also very close. Such understanding is better fleshed in Fig. 9, where the restoration results of both methods are tested on the standard "Peppers" image.

Acknowledgments

The third author wishes to thank Professor Stanley Osher for his constant encouragement and support.

References

[1] P. V. Blomgren. Total variation methods for restoration of vector valued images, (Ph.D. thesis). Also Technical report, UCLA Dept. of Math., CAM 98-30, June 1998.

[2] A. Chambolle and P. L. Lions. Image recovery via Total Variational minimization and related problems. *Numer. Math.*, 76:167–188, 1997.

[3] T. F. Chan and P. V. Blomgren. Color TV: Total variation methods for restoration of vector valued images. Technical report, UCLA Dept. of Math., CAM 96-5, 1996.

[4] T. F. Chan and P. V. Blomgren. Modular solvers for constrained image restoration problems. Technical report, UCLA Dept. of Math., CAM 97-52, 1997.

[5] T. F. Chan, S. Osher, and J. Shen. The digital TV filter and nonlinear denoising. *IEEE Trans. Image Processing,* to appear, 2000.

[6] T. F. Chan and J. Shen. Rotational invariance for prior models in image analysis. *Let. Applied Math.,* submitted, 2000.

[7] T. F. Chan and J. Shen. Variational restoration of non-flat image features: Models and algorithms. *SIAM Journal of Applied Mathematics,* to appear, 2000.

[8] R. Gonzalez and R. Wood. *Digital Image Processing.* Addison-Wesley, 1992.

[9] R. Kimmel and N. Sochen. Orientation diffusion or how to comb a porcupine ? *J. Visual Comm. Image Representation,* to appear, 2000.

[10] J.-M. Morel and S. Solimini. *Variational Methods in Image Segmentation,* volume 14 of *Progress in Nonlinear Differential Equations and Their Applications.* Birkhäuser, Boston, 1995.

[11] P. Perona. Orientation diffusion. *IEEE Trans. Image Process.*, 7(3):457–467, 1998.

[12] L. Rudin and S. Osher. Total variation based image restoration with free local constraints. *Proc. 1st IEEE ICIP*, 1:31–35, 1994.

[13] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithm. *Physica D*, 60:259–268, 1992.

[14] G. Sapiro and D. Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Processing*, 5:1582–1586, 1996.

[15] B. Tang, G. Sapiro, and V. Caselles. Color image enhancement via chromaticity diffusion. Technical report, ECE University of Minnesota, 1999.

[16] B. Tang, G. Sapiro, and V. Caselles. Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *Int. Journal Computer Vision*, 36(2):149–161, 2000.

[17] P. E. Trahanias, D. Karako, and A. N. Venetsanopoulos. Directional processing of color images: theory and experimental results. *IEEE Trans. Image Process.*, 5(6):868–880, 1996.

[18] P. E. Trahanias and A. N. Venetsanopoulos. Vector directional filters — a new class of multichannel image processing filters. *IEEE Trans. Image Processing*, 2(4):528–534, 1993.

[19] J. Weickert. *Anisotropic Diffusion in Image Processing.* Teubner-Verlag, Stuttgart, Germany, 1998.

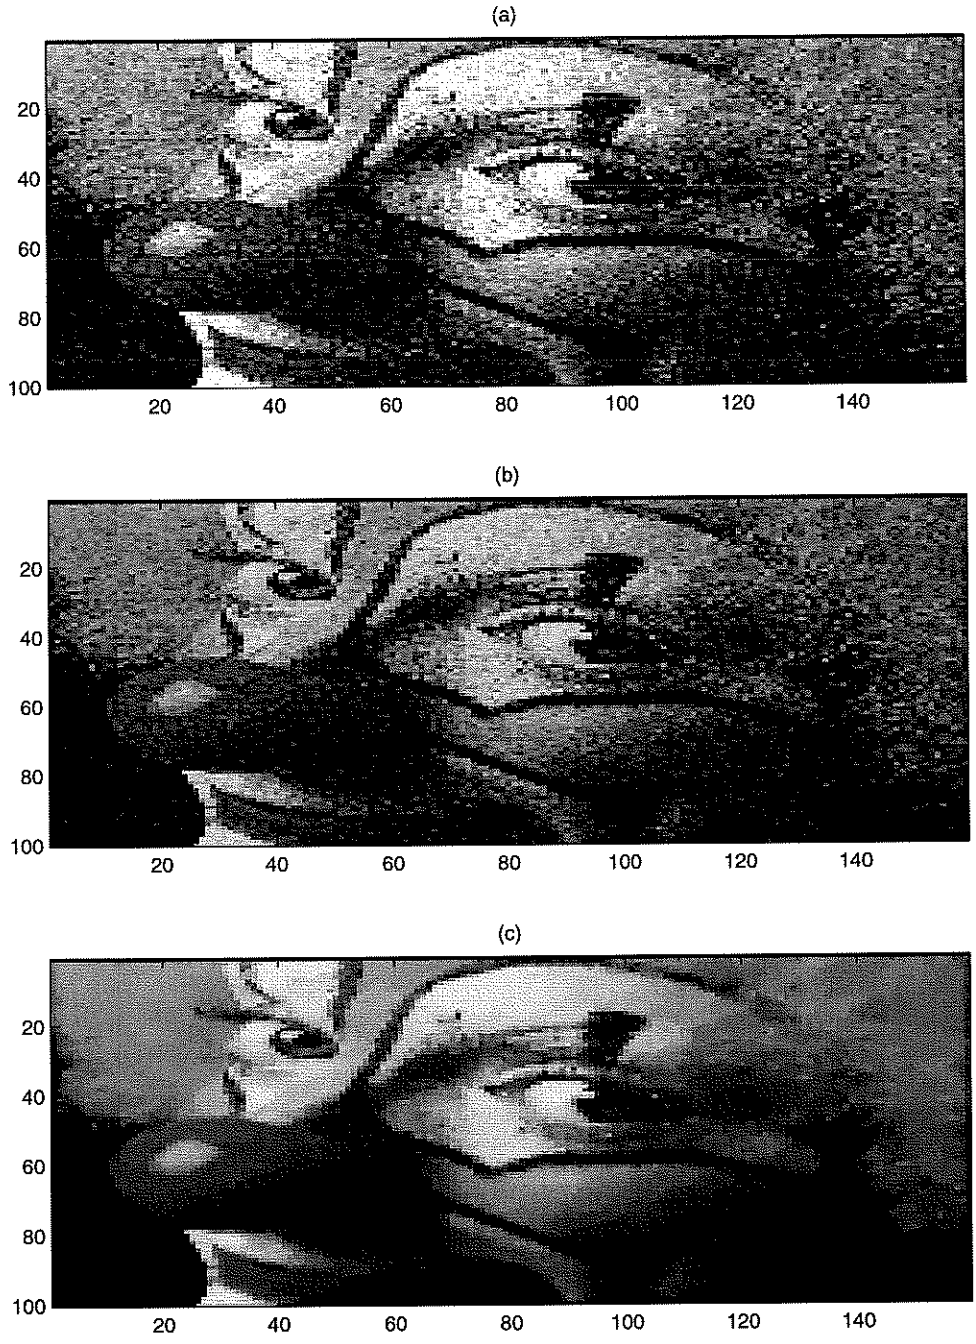

Figure 3: Spherical TV on the chromaticity component and scalar TV on the brightness component. In (b), only the chromaticity is restored, while in (c), both the chromaticity and brightness components are restored.

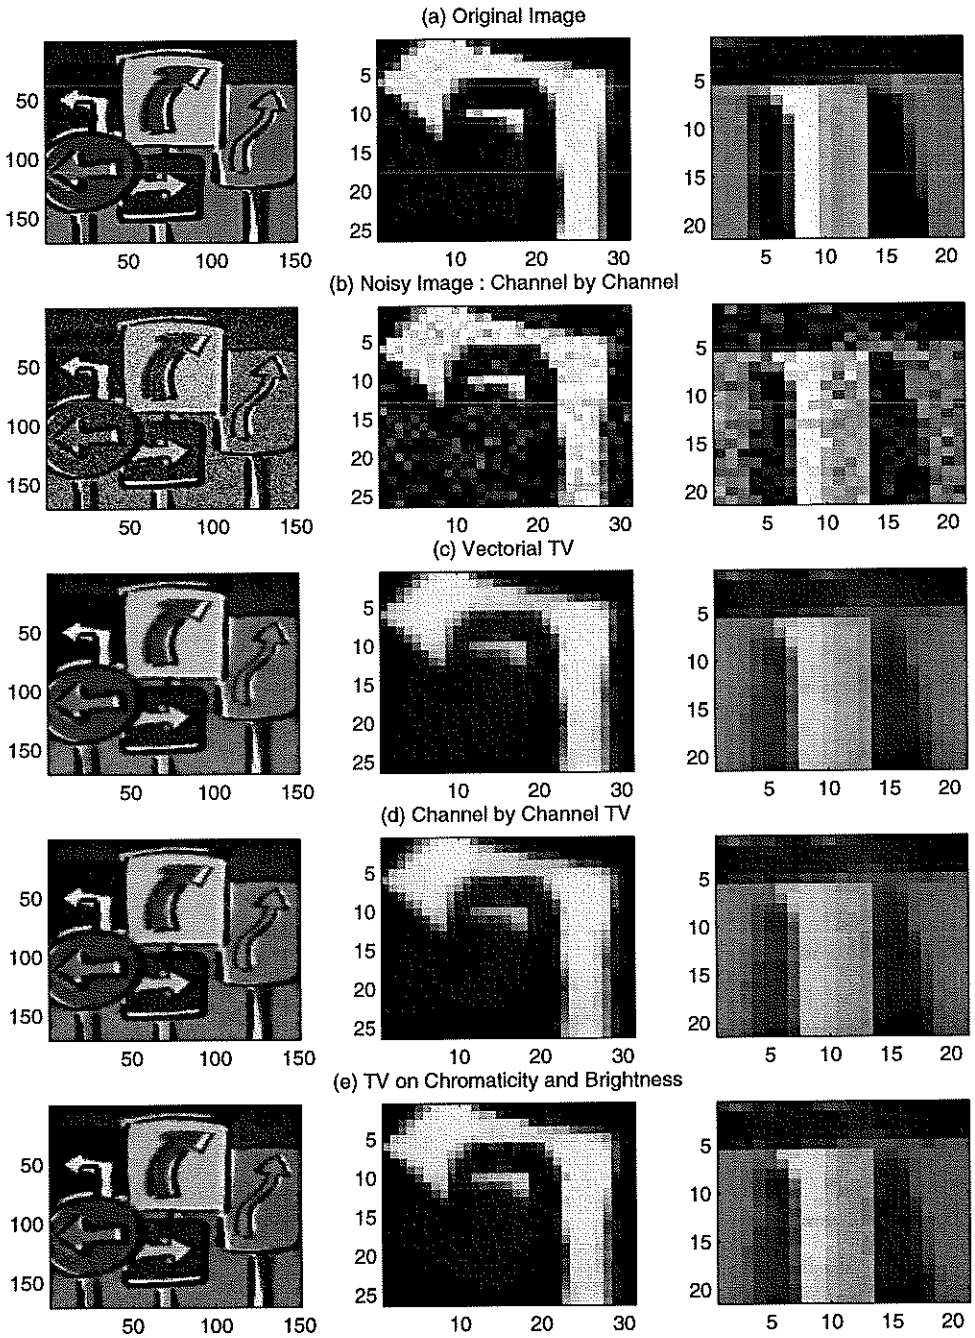

Figure 4: Comparison of the CB based TV and linear TV's (I): CB leads to better color control.

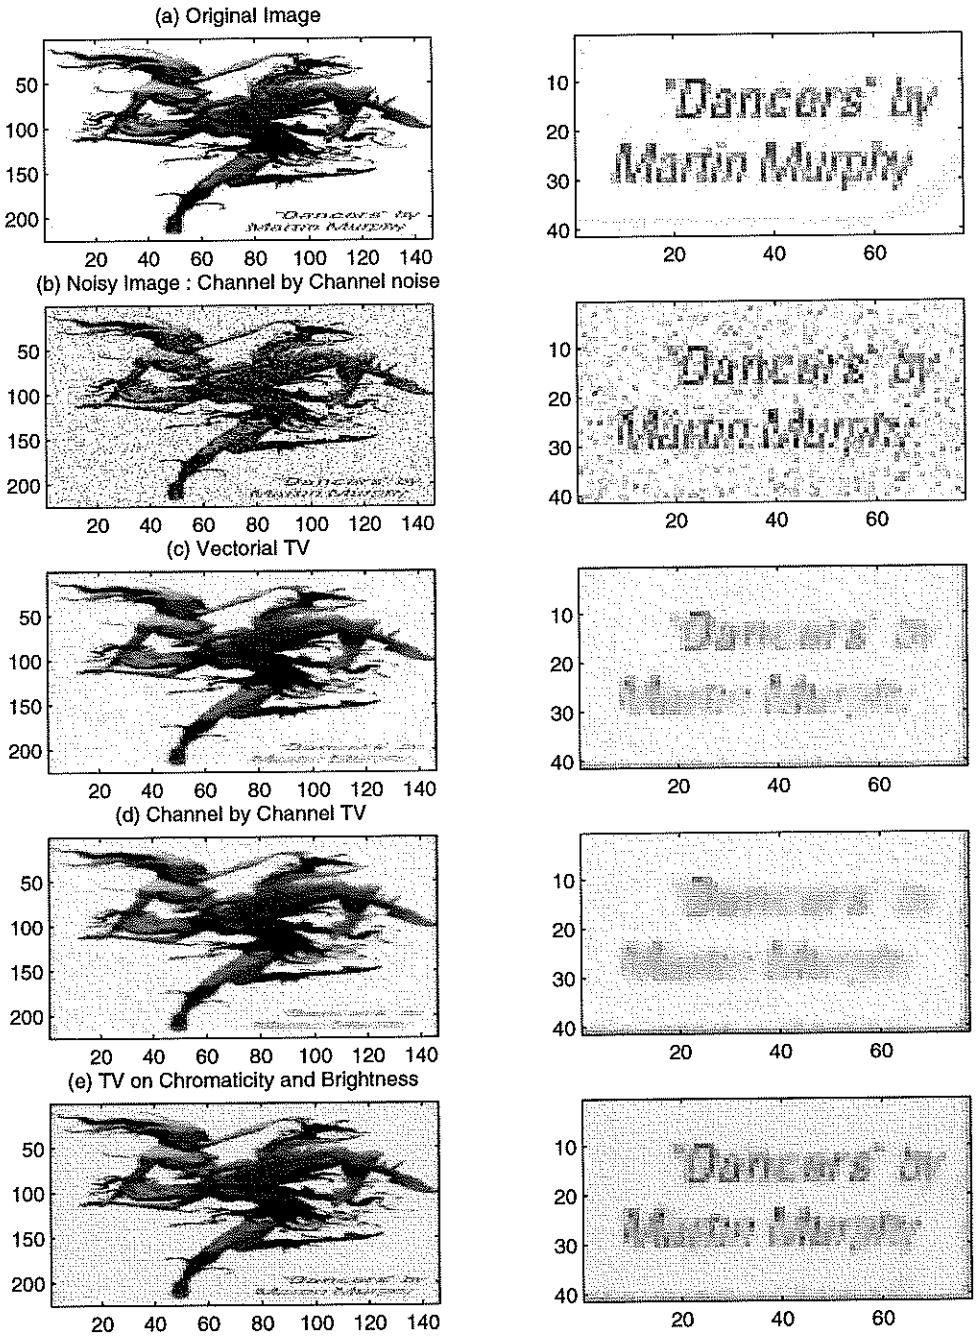

Figure 5: Comparison between the CB based TV and linear TV's (II): CB does not mix colors.

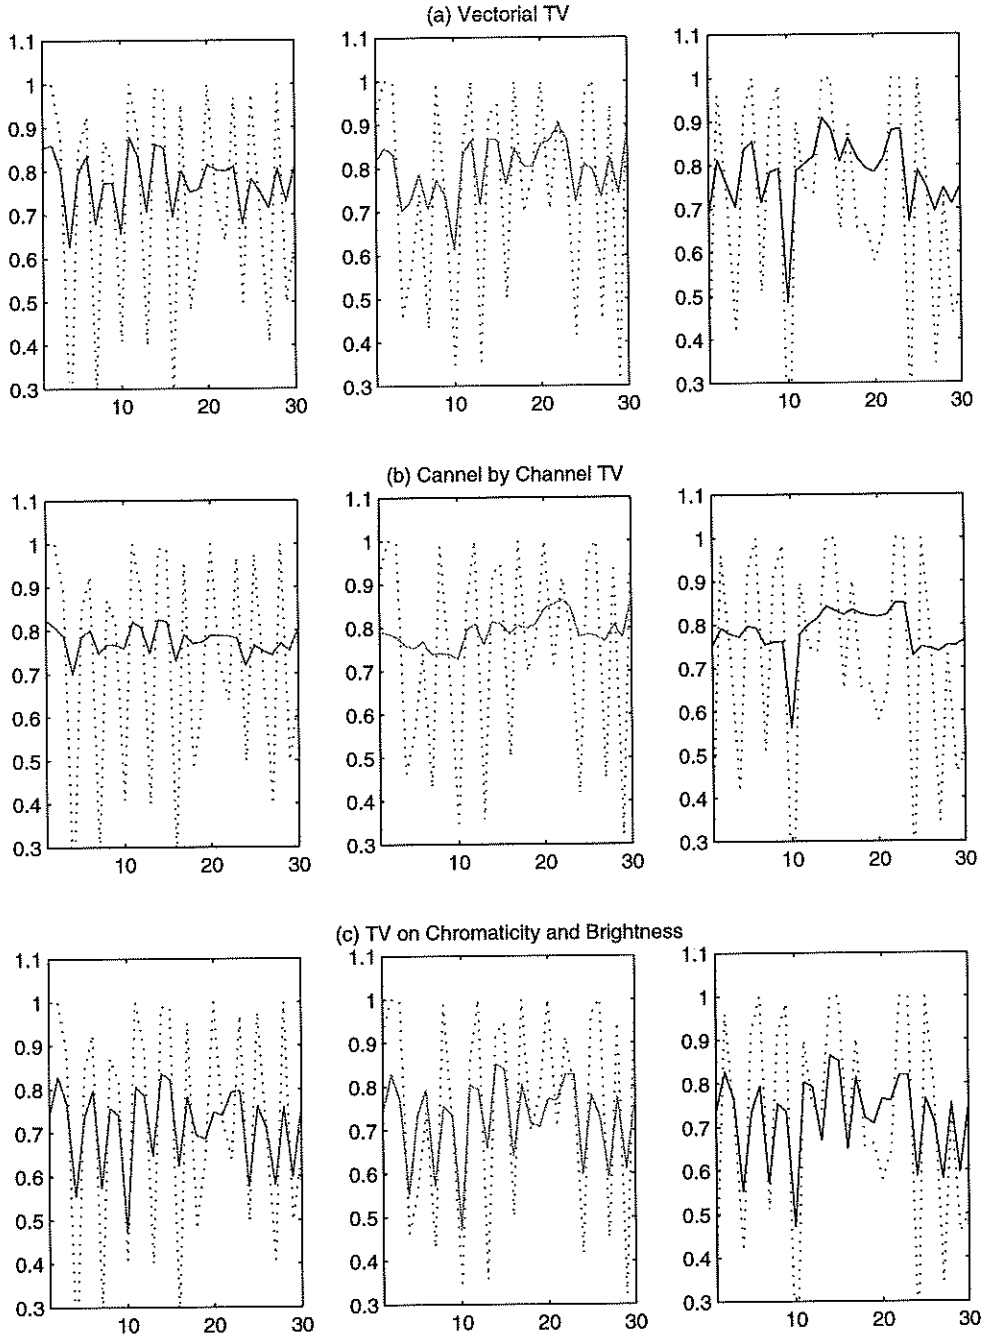

Figure 6: Comparison between the CB based TV and linear TV's (III): the 1-dimensional slice show. The identity of the pure gray color is best restored by the CB based TV.

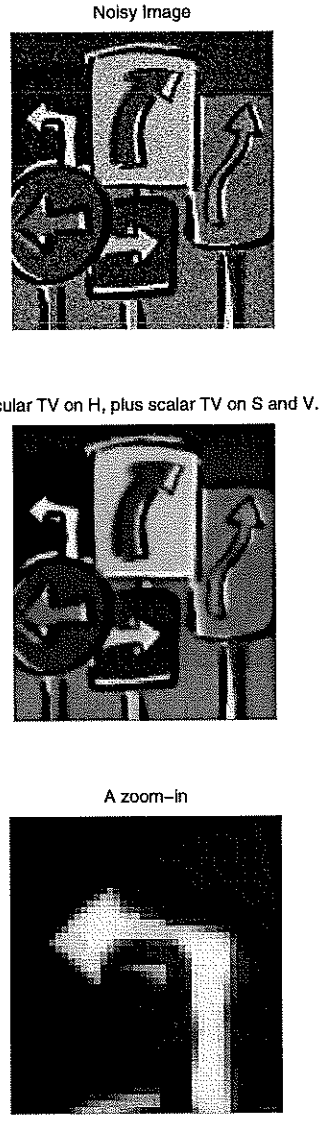

Figure 7: Circular TV on H, and scalar TV's on S and V. The unsatisfactory performance shows the inappropriateness of treating the H and S components separately.

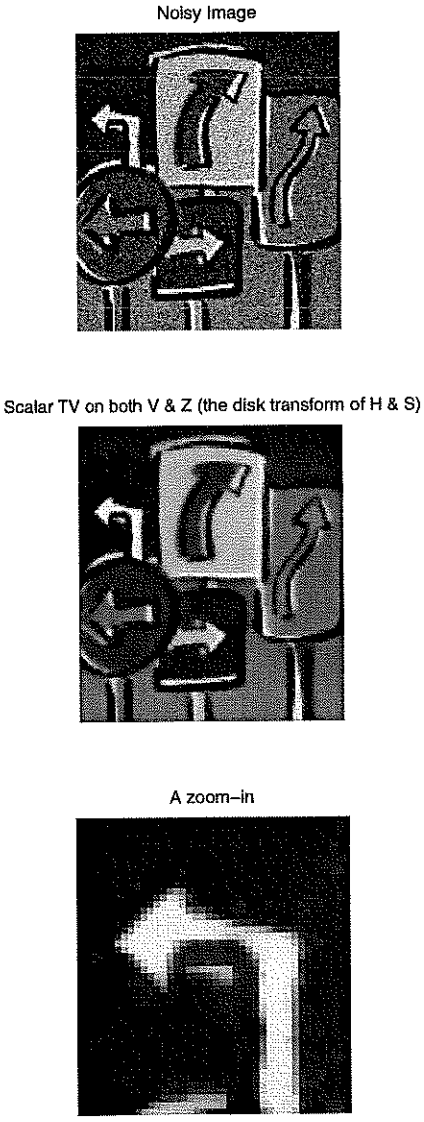

Figure 8: Scalar TV's on both the real function V and the complex function Z, i.e., the disk transform of H and S.

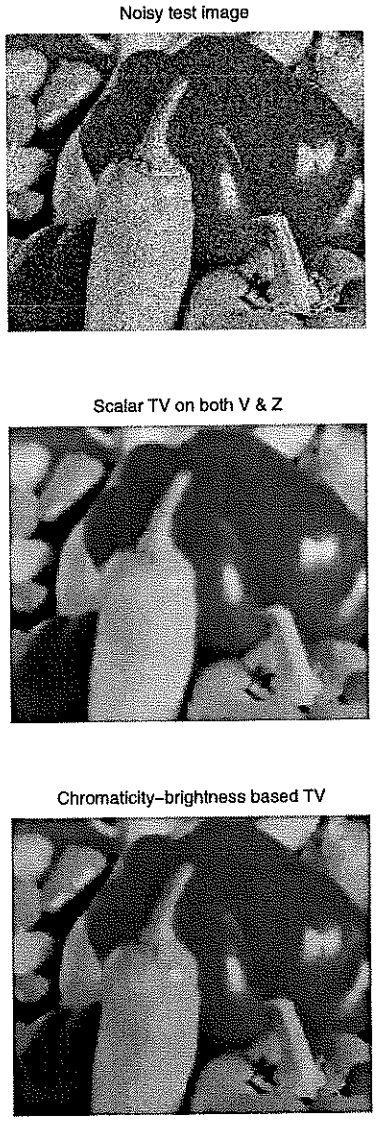

Figure 9: TV's based on the CB and HSV (under the disk transform) color models have the similar performance.