IRWIN AND JOAN JACOBS
CENTER FOR COMMUNICATION AND INFORMATION TECHNOLOGIES

A Total Variation Spectral Framework for Scale and Texture Analysis

Guy Gilboa

CCIT Report #833
June 2013

DEPARTMENT OF ELECTRICAL ENGINEERING
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, HAIFA 32000, ISRAEL

CCIT Report #833 June 2013

A Total Variation Spectral Framework for Scale and Texture Analysis

Guy Gilboa

Department of Electrical Engineering,
Technion - Israel Institute of Technology,
Haifa 32000, Israel

June 30, 2013

Abstract

A new total variation (TV) spectral framework is presented. A TV transform is proposed which can be interpreted as a spectral domain, where elementary TV features, like disks, approach impulses. A reconstruction formula from the spectral to the spatial domain is given, allowing the design of new filters. The framework allows deeper understanding of scales in an L^1 sense and the ability to better analyze and process textures. An example of a texture processing application illustrates possible benefits of this new framework.

1 Introduction

The total variation (TV) functional is today a fundamental regularizing tool in image processing. It is employed for denoising and deconvolution [41, 15, 36, 38, 37, 27], optical-flow [10], tomographic reconstruction [42], texture and image analysis [8, 5, 3, 46, 28] and more. Since its introduction in [41] in the context of image processing many studies have been devoted to its analysis and interpretation, e.g. [15, 36, 16, 17]. We attempt in this paper to further enhance the intuition and applicability of this functional to feature extraction and image analysis by formulating a spectral framework, where one can decompose and reconstruct images using the basic TV elements of the image.

Spectral analysis has been used extensively in the analysis and processing of signals modelled as stationary random processes (see e.g. [33, 44]). For more complex non-stationary signals, such as images and speech, harmonic analysis methods were developed in the form of wavelets [21, 35, 23], spectral graph theory [18] and diffusion maps [19]. We explore a way to provide spectral information for total variation analysis.

In [43] Steidl et al have shown the close relations, and equivalence in a 1D discrete setting, of the Haar wavelets to both TV regularization [41] and TV flow [1]. This was later developed for a 2D setting in [48]. The development of features in the scale space framework [49, 31, 40, 47] and the emergence of critical points were studied for example in [31, 11, 32, 45, 16, 28]. This work relies on the established theory of the TV flow proposed by Andreu et al in [1] and further developed in [2, 7, 43, 12, 6, 25] and the references therein.

This manuscript is an extended version of the conference paper [26].

2 The TV Spectral Framework

The scale-space approach is a natural way to define scale:

$$u_t = -p,\ u|_{t=0} = f,\ p \in \partial_u J(u), \tag{1}$$

where $\partial_u J(u)$ denotes the subdifferential of some regularizing functional $J(u)$.

We are interested in the total variation functional:

$$J(u) = \int_\Omega |Du|, \tag{2}$$

where Du denotes the distributional gradient of u. It is therefore natural to examine the total variation scale-space, known as total-variation flow [1], formally written as:

$$\begin{array}{ll} \frac{\partial u}{\partial t} = \operatorname{div}\left(\frac{Du}{|Du|}\right), & \text{in } (0,\infty)\times\Omega \\ \frac{\partial u}{\partial n} = 0, & \text{on } (0,\infty)\times\partial\Omega \\ u(0,x) = f(x), & \text{in } x \in \Omega, \end{array} \tag{3}$$

where Ω is the image domain (a bounded set in $\mathcal{R}^N$ with Lipschitz continuous boundary $\partial\Omega$). We assume f has sufficient spatial regularity.

We now give our line of thought how the transform was derived. Similar results may probably be obtained using other, more formal, approaches.

In Fourier analysis, the sine and cosine functions (or exponents with imaginary arguments) are the basic functions of the transform. They form impulses in the Fourier domain. How can this be generalized to the total variation domain? We begin by examining some atom-like elements in the TV sense. It is well known that disks are elementary structures for the TV functional. For instance, they satisfy the eigenvalue problem in $\mathcal{R}^N$: $\partial_u J(u) = \lambda u$ (where $\lambda \in \mathcal{R}$), which implies their shape stays the same during the entire evolution (their height decreases until they disappear). Analytic solutions for disk regularizations and evolutions were obtained for

the TV regularization model [36, 45], TV-flow [1, 2, 7], inverse-scale-space evolutions [13] and more.

Let us recall the analytic solution of a simple case: evolution of a single disk in two dimensions. The indicator function of a disk of radius r in $\mathcal{R}^2$ is:

$$I(x) = \begin{cases} 1, & |x| < r \\ 0, & \text{otherwise} \end{cases}$$

For a disk of height h, $hI(x)$, we have that $\partial_u J(u) = \frac{2}{r} I(x)$ for all t until the disk disappears. We denote by $t_d = \frac{hr}{2}$ the disappearance time.

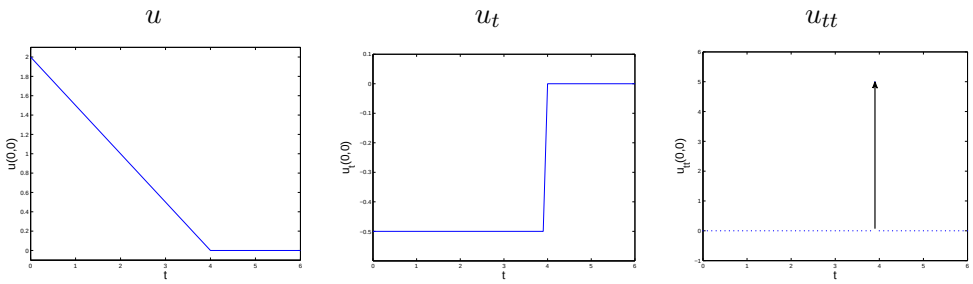

Figure 1: Illustrating the evolution of a disk in $\mathcal{R}^2$. The value is within $|x| < r$, for example at $(x_1 = 0, x_2 = 0)$. The second derivative is an impulse at time t_d. [here we set $r = 4$, $h = 2$ and therefore $t_d = 4$].

The solution of the TV flow for $u(t)$ is therefore

$$u(t) = \begin{cases} (h - \frac{2}{r}t)I(x), & 0 \le t < t_d \\ 0, & \text{otherwise} \end{cases}$$

The first and second derivatives in time are:

$$u_t(t) = \begin{cases} -\frac{2}{r}I(x), & 0 \le t < t_d \\ 0, & \text{otherwise} \end{cases}$$

$$u_{tt}(t) = \frac{2}{r}\delta(t - t_d)I(x),$$

where $\delta(t)$ denotes an impulse (Dirac delta) at $t = 0$. See Fig. 1 for an illustration.

We observe that u_{tt} yields an impulse of an elementary structure and is, therefore, a good candidate for a spectral representation. We would also like that the response will be invariant with respect to time. We normalize by multiplying it by the evolution time t. It will be seen later that this yields a straightforward reconstruction formula.

2.1 TV Transform

Let the TV transform be defined by

$$\phi(t) = u_{tt}t, \tag{4}$$

where $t \in (0, \infty)$ is the time parameter of the TV-flow, Equation (3), and u_{tt} is the second derivative in time of u in that flow.

Having defined $\phi(t) \in L^1(\Omega)$, we now need the inverse transform, which reconstructs a signal from all $\phi(t)$ responses. The reconstruction formula is very simple and is defined as:

$$w(x) = \int_0^\infty \phi(t; x)dt + \bar{f(x)}, \tag{5}$$

where $\bar{f} = \frac{1}{\Omega}\int_\Omega f(x)dx$ is the mean value of the initial condition. Naturally, if we do not manipulate the spectral domain for filtering, we expect to reconstruct the image of the initial condition f, as stated in the following:

Theorem 1 *For $\phi(t)$ defined in (4), the reconstruction formula (5) recovers $f \in BV(\Omega) \cap L^\infty(\Omega)$, that is $w(x) = f(x)$.*

Proof. We examine the left-term on the right hand side of Eq. (5). Integration by parts yields

$$\int_0^\infty \phi(t)dt = \int_0^\infty u_{tt}tdt = u_t t|_0^\infty - u|_0^\infty.$$

We use the property of finite extinction time of the TV flow. A two-dimensional proof by energy methods is given in [2] Th. 5. A more recent proof for all dimensions using energy estimates and Sobolev inequalities is given in [25] Th. 2.4, 2.5. In essence, this property means that for some $t_1 \in (0, \infty)$ we have $u(t) \equiv const$, $\forall t > t_1$. Therefore also $u_t(t) \equiv 0$ in a similar time range and $u_t(t)t|_{t\to\infty} = 0$. The expression $u_t \in -\partial_u J(u)$ is finite for all $t \in [0, \infty)$ so that $u_t t|_{t=0} = 0$. We can therefore conclude that the left term $u_t t|_0^\infty = 0$. For Neumann boundary conditions the mean is unchanged, therefore $u|_{t\to\infty} = \bar{f}$. Using the initial condition we have $u|_0^\infty = \bar{f} - f$. □

One can generalize this theorem to functionals other than TV, where the final time extinction property may not be valid or is hard to prove. Also, from a computational viewpoint, one may not want to evolve the entire scale-space (until the extinction time) but still view the spectrum, perform filtering and reconstruct an image based on an evolution until some time T. In these cases we can use the following finite time reconstruction formula:

$$w_T(x) = \int_0^T \phi(t; x)dt + f_r(T; x), \tag{6}$$

where f_r is a residual part defined by

$$f_r(T;x) = u(T;x) - u_t(T;x)T. \tag{7}$$

We can prove a similar recovery of $w_T(x)$ to $f(x)$ (assuming $u_t(T;x)T \in L^\infty$) by using the same arguments as in Theorem 1, leaving the term $u_t(T;x)T$ as is.

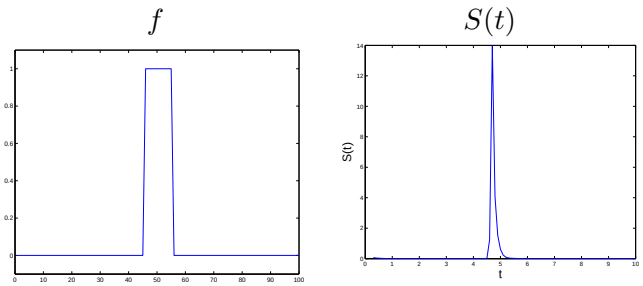

Figure 2: A single one-dimensional disk and the corresponding numerical spectral response $S(t)$.

Definition 1 (TV Spectral Response) *The TV spectral response for* $t \in (0,\infty)$ *is defined as:*

$$S(t) = \|\phi(t;x)\|_{L^1} = \int_\Omega |\phi(t;x)|dx. \tag{8}$$

The spectral response roughly corresponds to the amplitude of the response in a Fourier domain (see Fig. 4). If the response is high, a large "quantity" of the element $\phi(t)$ is contained in the image. If it is low, this element can be considered negligible. A response for one dimensional disk, as computed discretely, is depicted in Fig. 2. We will show in our experiments that, as can be expected, elements with high spectral response compose the main features of the image.

2.2 Nonlinear Eigenvalue Problem

We would like to obtain a connection between eigenfunctions and the spectral components $\phi(t)$.

The nonlinear eigenvalue problem with respect to the functional $J(u)$ is defined by:

$$p(u) = \alpha u,\ p(u) \in \partial J(u),\ \alpha \in \mathcal{R}. \tag{9}$$

We denote u which admits (9) an eigenfunction, where α is the corresponding eigenvalue.

In this manuscript we assume $J(u)$ is the total variation functional and $u \in BV$. The following theorem, however, can be generalized for any one-homogeneous functional.

We refer to a functional as positively homogeneous of degree one (or one-homogeneous for short) if

$$J(ku) = kJ(u),\ k \geq 0.$$

We recall the following property of one-homogeneous functionals:

$$J(u) = \langle p, u \rangle.$$

This can be shown by using the subgradient definition

$$\{p | J(v) - J(u) \geq \langle p, v - u \rangle, \forall v \in \mathcal{U}\},$$

assigning $v = ku$ and using $J(ku) = kJ(u)$ yields $J(u) \geq \langle p, u \rangle$. On the other hand, assigning $v = 0$ and using $J(0) = 0$ yields $J(u) \leq \langle p, u \rangle$.

Using the above, another useful property can be shown:

Lemma 1 *If $J(u)$ is one-homogeneous and $p \in \partial J(u)$ then $\forall k$, $0 < k < \infty$, $p \in \partial J(kf)$.*

Proof. From the one homogeneous property, $J(ku) = kJ(u)$, and the subgradient definition, $J(v) \geq J(u) + \langle p, v - u \rangle$, we have the bound $J(v) - J(ku) \geq (J(u) + \langle p, v - u \rangle) - kJ(u)$. Using $J(u) = \langle p, u \rangle$ we can conclude that $J(v) - J(ku) \geq \langle p, v - ku \rangle$ and therefore $p \in \partial J(kf)$. $\square$

Theorem 2 *For any $f \in BV$ which is a solution of the eigenvalue problem* (9)*, the following analytic solution holds:*

$$u(t; x) = \begin{cases} f(x)(1 - \alpha t), & 0 \leq t \leq \frac{1}{\alpha} \\ 0, & t > \frac{1}{\alpha} \end{cases} \tag{10}$$

$$\phi(t; x) = \delta(t - \frac{1}{\alpha}) f(x),$$

$$S(t) = \delta(t - \frac{1}{\alpha}) \| f(x) \|_{L^1}.$$

Proof. We take $\Delta t \to 0$ and use the partial derivative definition to rewrite $u_t = -p$, Eq. (3) top row, as

$$u(t + \Delta t; x) = u(t; x) - \Delta t p,\ p \in \partial J(u(t; x)).$$

For $t = 0$, $u(0; x) = f(x)$, and since $f(x)$ is an eigenfunction we have $p = \alpha f(x)$. Thus for $t = \Delta t$ the solution is $u(\Delta t; x) = f(x)(1 - \Delta t \alpha)$. Using Lemma 1 and assigning $k = (1 - \Delta t \alpha) > 0$ we have $p|_{t=\Delta t} = \alpha f(x)$, $p \in \partial J(u(\Delta t; x))$. By induction this follows until $t = \frac{1}{\alpha}$. We can thus write

$u_t(t;x) = -\alpha f(x)$, $0 \le t < \frac{1}{\alpha}$ and solve $u(t;x) = u(0;x) - t\alpha f(x)$ in this time range and u(t;x)=0 for $t \ge \frac{1}{\alpha}$. The expressions for $\phi(t)$ and $S(t)$ follow from their definitions (4) and (8), respectively. □

It is shown above that for the particular case of a signal composed of a single eigenfunction we get a single spectral component at some time t_1, where $\phi(t_1)$ is exactly this eigenfunction. For a signal composed of several eigenfunctions which are well separated spatially - we observe numerically that we get spectral peaks and $\phi(t)$ functions which correspond to the original eigenfunctions. The conditions for additivity of the eigenfunctions should be investigated more deeply. One may follow ideas of curve evolutions of the TV-flow given in [7, 2] to obtain precise formulations.

2.3 Spectral Filtering

Let $H(t)$ be a filter defined in the TV spectral domain as a real valued function of t. The filtered response $\phi_H(t)$ in the spectral domain is defined by:

$$\phi_H(t) = \phi(t)H(t). \tag{11}$$

The filtered response in the spatial domain is then the corresponding reconstruction procedure

$$f_H(x) = \int_0^\infty \phi_H(t)dt + \bar{f}, \tag{12}$$

An ideal filter in Fourier analysis eliminates completely energy of undesired frequencies while perfectly retaining frequencies in the desired range. We can now define analogous ideal filters in the TV spectral sense:

Definition 2 (Ideal Spectral Filters) *Let $t_1, t_2 \in (0, \infty)$. We denote $\phi(\infty) \equiv u(\infty) = \bar{f}$, $\phi(0) = 0$. We can now define the following ideal spectral filters:*

(i) Ideal low-pass filter:

$$H_{LPF,t_1}(t) = \begin{cases} 0, & 0 \le t < t_1 \\ 1, & t_1 \le t \le \infty \end{cases}$$

(ii) Ideal high-pass filter:

$$H_{HPF,t_1}(t) = \begin{cases} 1, & 0 \le t < t_1 \\ 0, & t_1 \le t \le \infty \end{cases}$$

(iii) Ideal band-pass filter:

$$H_{BPF,t_1,t_2}(t) = \begin{cases} 0, & 0 \le t < t_1 \\ 1, & t_1 \le t < t_2 \\ 0, & t_2 \le t \le \infty \end{cases}$$

(iv) Ideal band-stop filter:

$$H_{BSF,t_1,t_2}(t) = \begin{cases} 1, & 0 \le t < t_1 \\ 0, & t_1 \le t < t_2 \\ 1, & t_2 \le t \le \infty \end{cases}$$

Naturally, there are close relations between the above filters. As in the case of Fourier filters, a LPF is the spectral complement of a HPF and vise versa. We can summarize it by the following statements:

Proposition 1 *For any* $0 < t_1, t_2 < \infty$ *the following identities hold:*

$$\begin{array}{lcl} H_{LPF,t_1}(t) & = & f - H_{HPF,t_1}(t); \\ H_{HPF,t_1}(t) & = & f - H_{LPF,t_1}(t); \\ H_{BPF,t_1,t_2}(t) & = & f - H_{BSF,t_1,t_2}(t); \\ H_{BSF,t_1,t_2}(t) & = & f - H_{BPF,t_1,t_2}(t); \end{array} \tag{13}$$

Proof. We rewrite f as an all-pass spectral filter, that is $H(t) \equiv 1$, $0 \le t \le \infty$. The identities then follow from the filters' definitions and the reconstruction formula, Eq. (5), (note to take into account $\phi(\infty) \equiv \bar{f}$). □

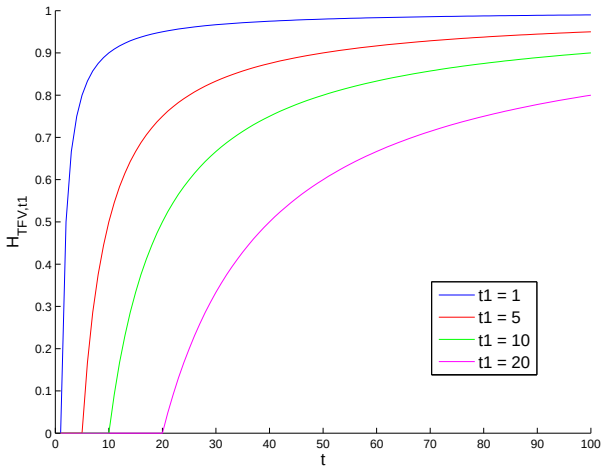

Figure 3: Illustration of Theorem 3. TV-Flow is equivalent to a non-ideal low-pass-filter in the spectral domain. In the plot four filters H_{TFV,t_1} are plotted for different values $t_1 = \{1, 5, 10, 20\}$.

2.4 TV Flow as a Low-Pass-Filter

We we'll now show that TV-Flow can be interpreted as a specific kind of low-pass-filter. This is formalized in the following theorem:

Theorem 3 *The solution $u(t_1)$ at time t_1 of the TV-flow process, Eq.* (3)*, is equivalent to spectral filtering with the following filter:*

$$H_{TVF,t_1}(t) = \begin{cases} 0, & 0 \le t < t_1 \\ \frac{t-t_1}{t}, & t_1 \le t < \infty \end{cases} \tag{14}$$

Proof. We apply the reconstruction formula, Eq. (12):

$$f_{H_{TVF,t_1}} = \int_0^\infty H_{TVF,t_1}\phi(t)dt + \bar{f},$$

which can be written as

$$f_{H_{TVF,t_1}} = \int_{t_1}^\infty t u_{tt} dt - t_1 \int_{t_1}^\infty u_{tt} dt + \bar{f}.$$

Using integration by parts, similar to Theorem 1, the left part of the right-hand-side is $-t_1 u_t(t_1) - u(\infty) + u(t_1)$. The right term is $t_1 u_t(t_1)$ (again we use the property of finite extinction time yielding $t_1 u_t(\infty) = 0$). Thus the terms $t_1 u_t(t_1)$ and $u(\infty) = \bar{f}$ cancel out and we get $f_{H_{TVF,t_1}} = u(t_1)$. □

It is therefore obvious that TV-flow (and ROF for that matter) is not an ideal low-pass-filter, as can be expected. It erodes the contrast of all the spectral components $\phi(t)$. The precise extent of this erosion is now quantified and given in Eq. (14). See Fig. 3 for plots of various TV-Flow responses for different t_1 values.

2.5 Numerical Implementation

There are many ways to implement total variation flow. We are taking a variational approach. We would like to approximate $u_t = -p(u)$. In the explicit method we do the following sequence of iterations:

$$u(n+1) = u(n) - dtp(u(n)),$$

with a small time-step dt. Following a suggestion of [39] we do the implicit approximation which is unconditionally stable in dt:

$$u(n+1) = u(n) - dtp(u(n+1)).$$

We can write the above expression as $u(n+1) - u(n) + dtp(u(n+1)) = 0$, and see it coincides with the Euler-Lagrange of the following minimization:

$$E(u, u(n)) = J(u) + \frac{1}{2dt}\|u - u(n)\|_{l^2}^2,$$

where $u(n+1)$ is the minimizer u of $E(u, u(n))$ and $u(n)$ is fixed. We now have a standard ROF problem [41], which can be solved using various algorithms (e.g. [20, 14, 30]). In [39] it was suggested to use the Split-Bregman algorithm, which is probably the most efficient one. With current

implementation, given e.g. in [29], we get occasionally some random spikes in the spectrum $S(t)$. This is still under investigation. We currently use Chambolle's projection algorithm [14], which gives consistent results.

To approximate the second time derivative $u_{tt}(t)$ we store in memory 3 consecutive time steps of u and use the standard central scheme:

$$D^2u(n) = \frac{u(n-1) + u(n+1) - 2u(n)}{dt},$$

with $n = 1, 2, ..,$ $u(0) = f$, $D^2u(0) = 0$. The time t is discretized as $t(n) = ndt$. Therefore:

$$\phi(n) = D^2u(n)t(n) = n\left(u(n-1) + u(n+1) - 2u(n)\right), \tag{15}$$

and the spectrum is

$$S(n) = \sum_{i \in \mathcal{N}} |\phi_i(n)|, \tag{16}$$

where i is a pixel in the image and $\mathcal{N}$ is the image domain. Finally we discretize the residual part $f_r(t)$ using a forward time difference $Du(n) = (u(n+1) - u(n))/dt$, to have:

$$f_r(n) = u(n) - Du(n)n = (n+1)u(n) - nu(n+1). \tag{17}$$

The discrete reconstruction formula for iteration $n = N$ (equivalent to Eq. (6)) is:

$$w_N = \sum_{n=1}^{N} \phi(n) + f_r(N). \tag{18}$$

We can now show the discrete reconstruction property:

Theorem 4 *For any $N = 1, 2, ..$ the following reconstruction property holds: $w_N = f$, where f is the initial condition $u(0) = f$.*

Proof. We examine the expression $\sum_{n=1}^{N} \phi(n)$. For $N = 1$ we have

$$\sum_{n=1}^{1} \phi(n) = u(0) + u(2) - 2u(1).$$

For $N = 2$ we have

$$\sum_{n=1}^{2} \phi(n) = u(0) - 3u(2) + 2u(3).$$

We see that $u(1)$ was canceled out. It can be shown by induction that only $u(0)$ and the last two time instances of u are kept in the following way:

$$\sum_{n=1}^{N} \phi(n) = u(0) - (N+1)u(N) + Nu(N+1).$$

Taking into account that $f_r(N) = (N+1)u(N) - Nu(N+1)$ we get $w_N = u(0) = f$. $\square$

Note that the above theorem does not rely on any property related to BV functionals. Therefore, it can be valid for most scale-spaces. In particular, the numerical approximation of the TV-flow does not have to be precise to get a proper reconstruction (some properties, naturally, will be lost, and visually we'll see the spectrum smears).

3 Texture Processing

We present a texture processing application which can benefit from the TV spectral framework. The idea is that natural textures most often have many scales. One may want to remove, enhance or manipulate certain texture scales, but leave untouched other textural and structural components. This can be viewed as a sort of generalization of the structure-texture decomposition methods such as [46, 38, 5, 4, 34, 24].

The spectral framework allows a straightforward way to accomplish such tasks. The general approach can be describe as follows:

1. Compute the spectral components $\phi(t)$ and the spectrum $S(t)$.
2. Manually or automatically, analyze the spectrum and choose the desired spectral components $\phi(t)$, $t_1 \leq t \leq t_2$ to manipulate.
3. Change the targeted spectral components $\phi(t)$ to $\tilde{\phi}(t)$, $t_1 \leq t \leq t_2$ by attenuating, amplifying or doing any spatial filtering to them.
4. Reconstruct the image using the manipulated $\tilde{\phi}(t)$ for $t_1 \leq t \leq t_2$ and the original $\phi(t)$ otherwise.

In the simplest form, this is similar to some sort of band-pass filtering of particular texture-bands. However one can also filter spatially those bands, allowing more elaborated and new types of responses.

We'll show two examples of texture filtering. We use here color images, and therefore use a vector-valued total variation and a vectorial projection algorithm, following the work of [9].

See Figs. 7 and 8 for examples of such processes. Note that in Fig. 8 we first roughly segment the skin and produce a fuzzy mask. Then the attenuation of textures is only for the skin part (without processing the hair, eyes or mouth). Here we have selective spatial attenuation of the texture channels, which is more adapted to the image than band-pass filtering.

3.1 Feature Extraction

The spectral response $S(t)$ can be used to characterize an image. It informs us of the dominant scales and can be used when comparing images or as

features for a machine learning algorithms. See Figs. 6, 9 for the spectral response and selected elements $\phi(t)$ of two image examples.

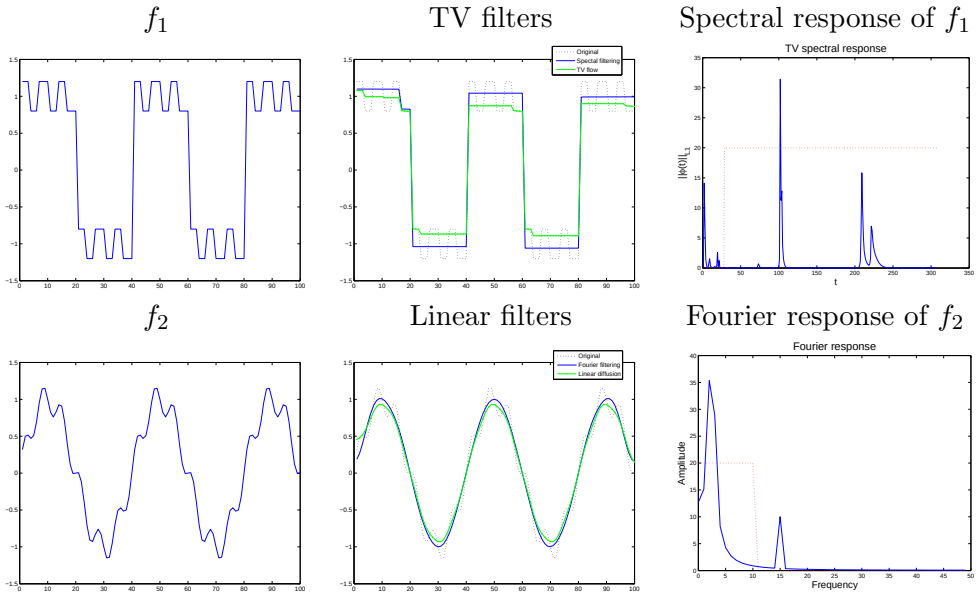

Figure 4: One dimensional example of ideal low pass filtering versus scale-space low pass filtering. Top row, processing f_1 (left), middle - response by spectral filtering (full blue line), and by TV flow (dotted green line). On the top right the spectral response is shown. On the bottom row an analogue linear case filters f_2 with Fourier ideal LPF (full blue line) versus linear diffusion (dotted green line).

4 Examples

Examples demonstrating the qualitative properties of this transform are shown below.

In Fig. 4 a 1D example is shown and compared with classical low-pass-filtering in the Fourier domain. In the classical linear setting (bottom row) we have: $f_2 = \sin 2\pi\varphi_1 + 0.2 \sin 2\pi\varphi_2$, (in this specific example $\varphi_1 = 0.025$, $\varphi_2 = 0.15$). We compare two linear low-pass filters (LPF) - an ideal LPF and linear diffusion. The ideal LPF (shown on bottom, right, dotted line) keeps all low frequencies and sets to zero all frequencies above the threshold. The diffusion processes attenuates more softly the frequencies near the threshold (as it is not an ideal LPF). We observe that the ideal LPF retains the low frequency with better contrast.

A signal with similar properties, adapted for the TV case, is shown in Fig. 4 top row: $f_1 = \text{sign}(\sin 2\pi\varphi_1) + 0.2\,\text{sign}(\sin 2\pi\varphi_2)$. The spectral

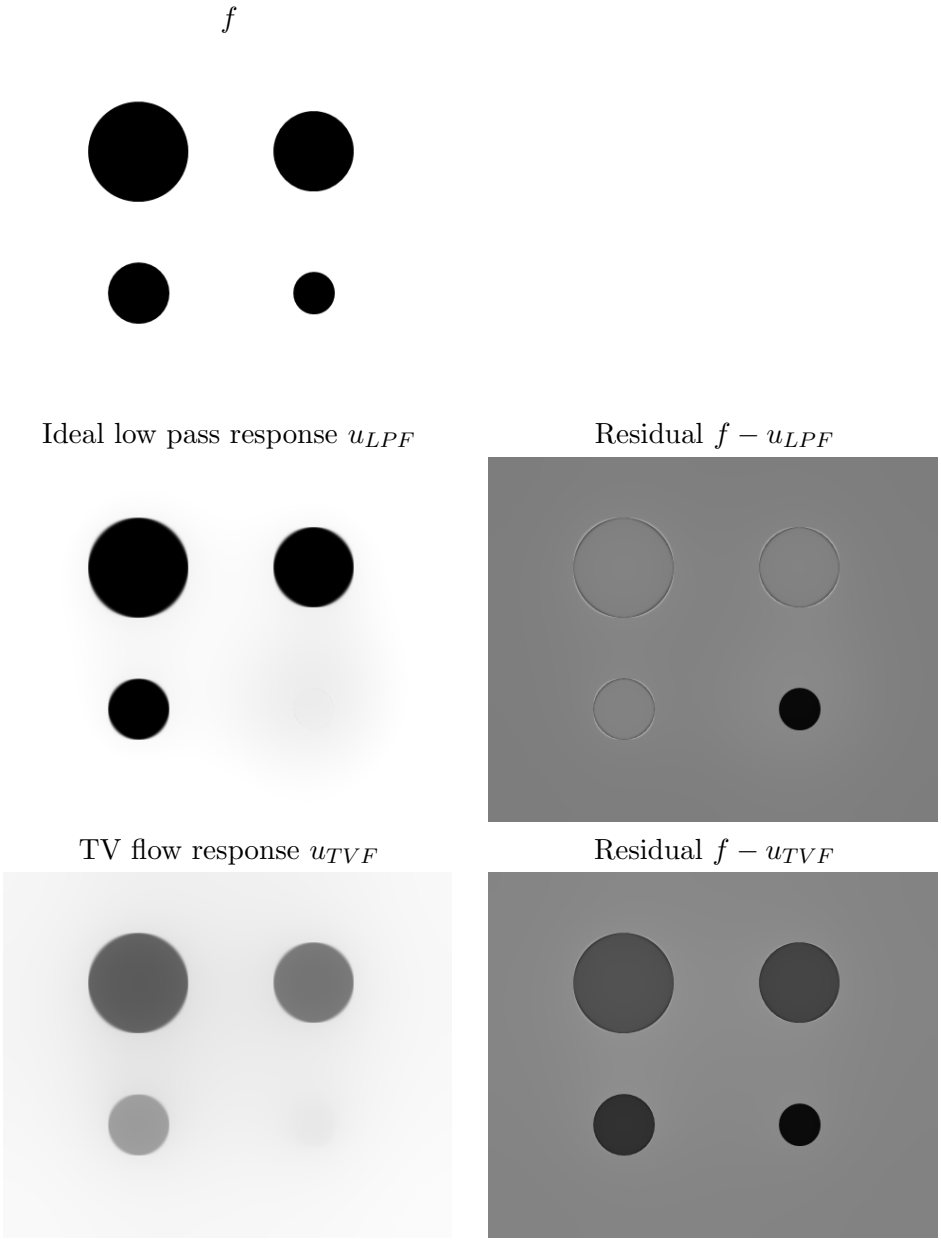

Figure 5: Comparison between the ideal low-pass filter response and TV-flow. In both cases the response is shown for the minimal extent of filtering in which the smallest circle completely vanishes. One sees the considerable reduction of contrast of the larger circles in the TV-flow versus the sharp and stable results of the ideal TV LPF.

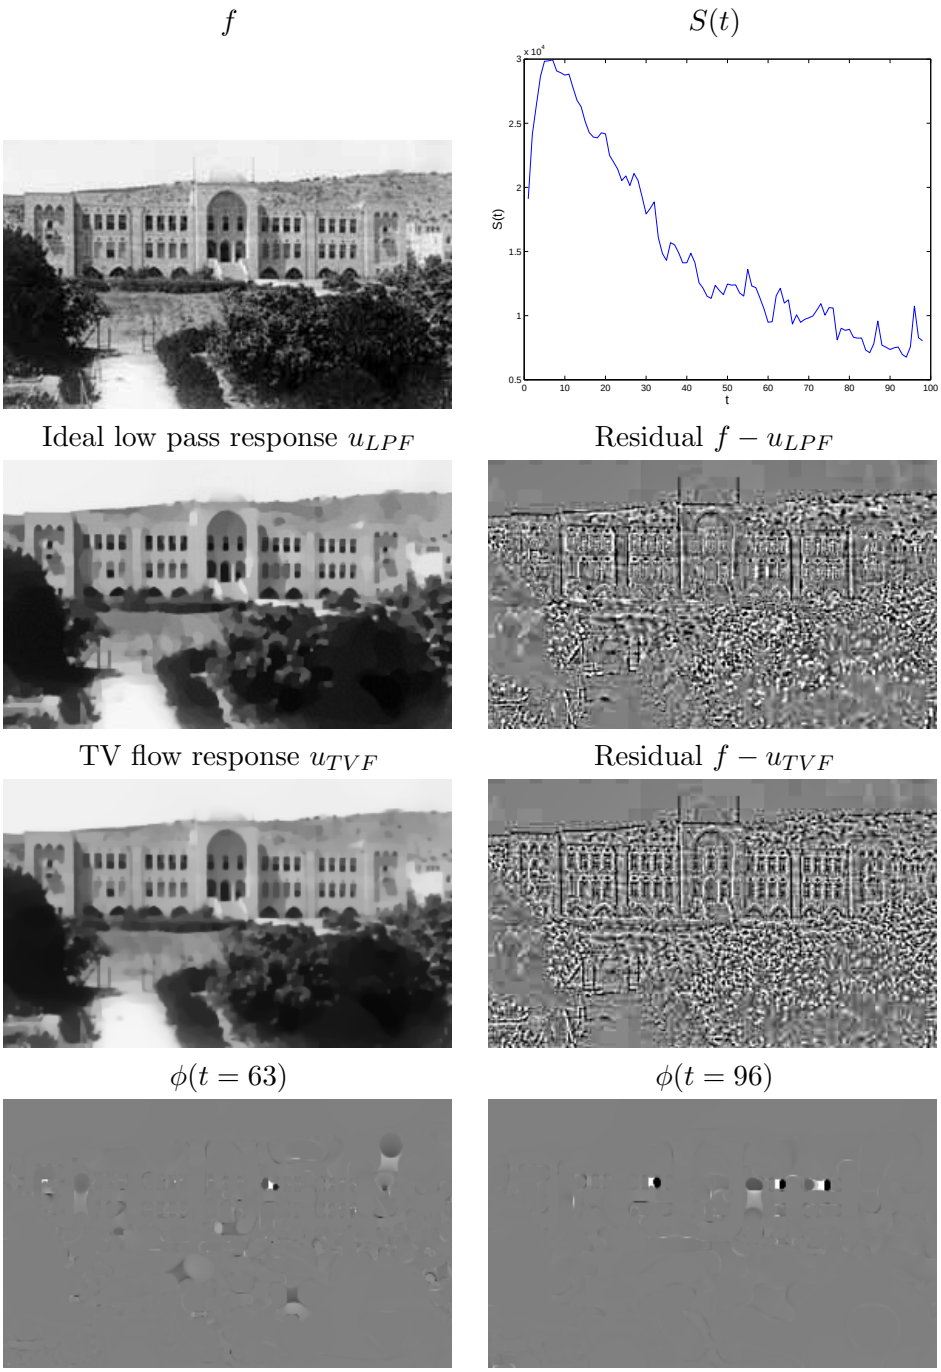

Figure 6: Old Technion image. Results of ideal low pass filtering. This is compared to TV-flow with equivalent filtering in the L^2 sense (the norm of the residual, $\|f - u\|_{L^2}$, is the same). In addition, two examples of $\phi(t)$ are shown for different t values.

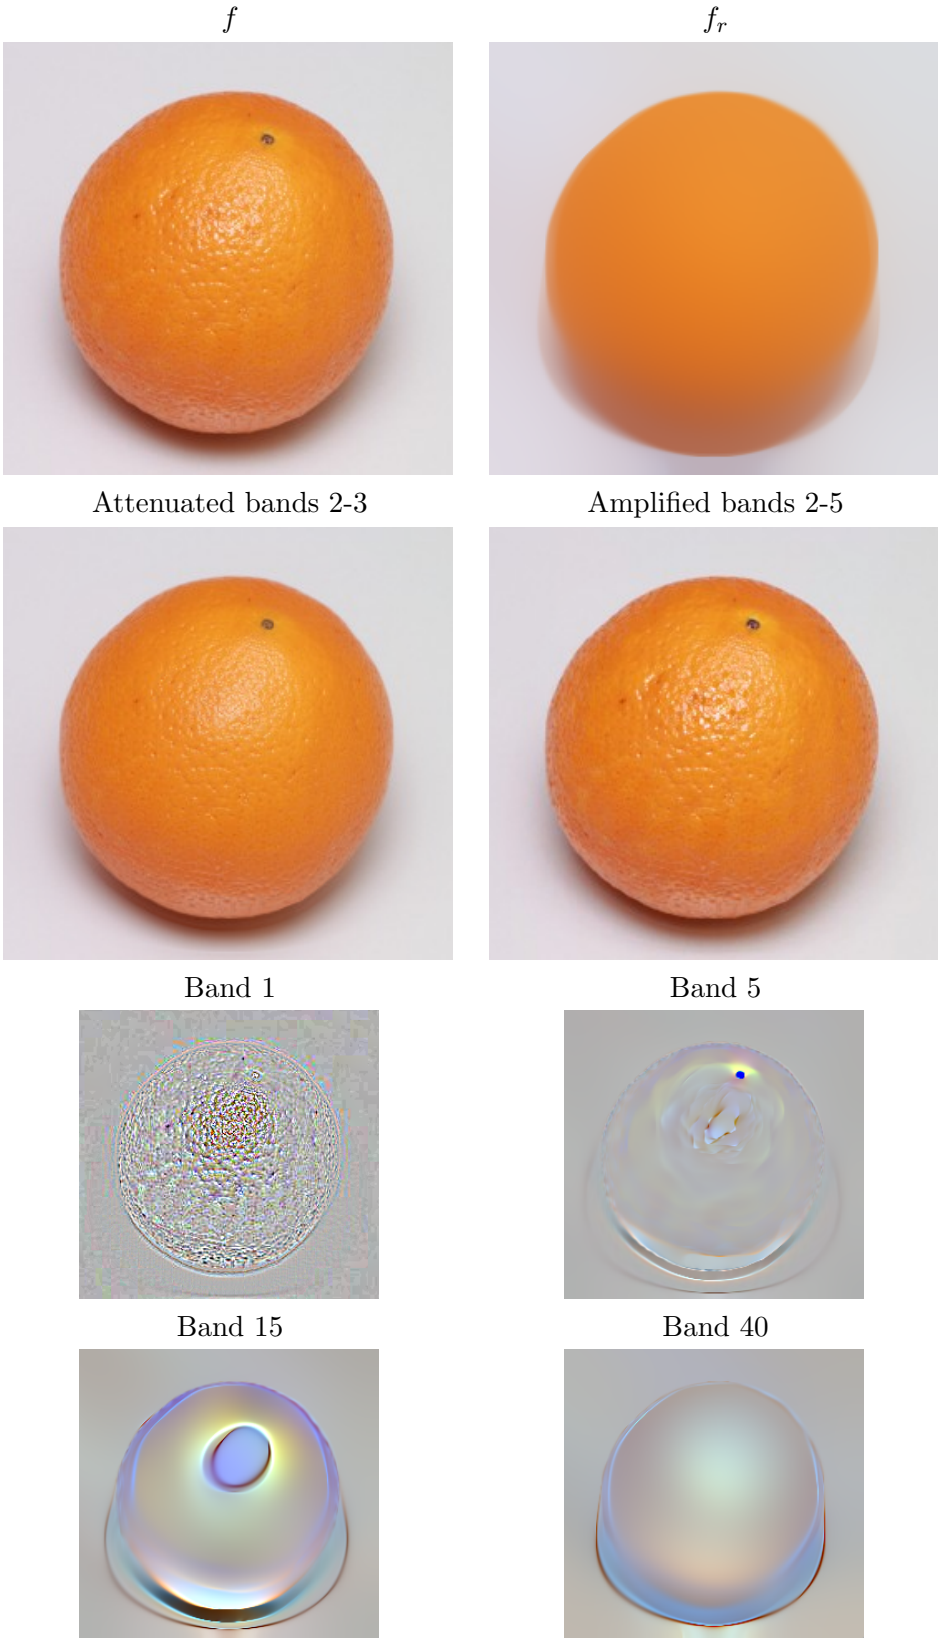

Figure 7: Texture processing. Attenuating or amplifying certain texture layers of an orange. Image taken from [22].

Figure 8: Texture processing. Removing texture layers from a face. Image taken from [22].

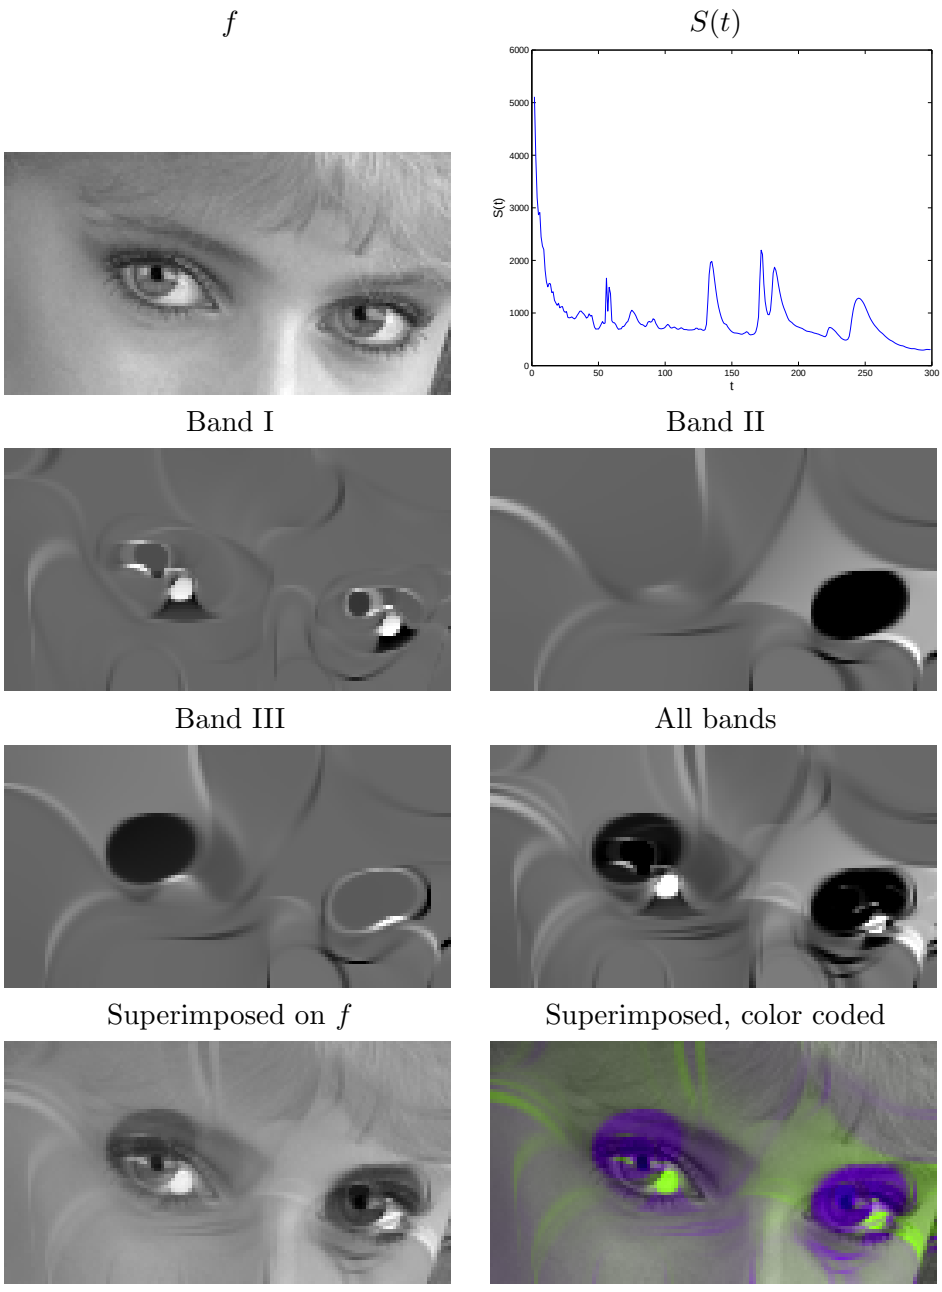

Figure 9: Feature extraction example. Salient features ate depicted as spectral peaks (top right). The first three spectral peaks are shown as Bands I-III. These bands are reconstructed together at the third row, right. This reconstruction is then superimposed on the image to show the localization of the bands. Bottom right - a color coded visualization of the image with the selected bands.

response $S(t)$ shows three active bands ($t < 30$ high oscillations, $t \approx 100$ low oscillations, and $200 < t < 250$ low amplitude step). TV flow is compared to ideal TV LPF, as defined above with filter threshold $t_1 = 30$. The filter response is illustrated in a dotted line at the top right. Note that in the TV spectral setting high frequencies are on the left side (small t values) as oppose to Fourier domain.

One can observe the very sharp transitions of the ideal LPF using the spectral filtering. Note that filtering with ideal LPF may result in too sharp transitions which can produce some reconstruction artifacts. This can be the case both in the linear and TV settings.

In Fig. 5 four circles of different sizes are processed. The ideal LPF is compared to TV-flow. In both cases the extent of filtering is such that the smallest circle completely vanishes. One can observe that the ideal LPF retains almost perfectly the larger three circles, whereas TV-flow erodes their contrast considerably.

In Fig. 6 an image of a building with landscape is examined. The ideal LPF response is shown along with a standard TV-flow filtering. In both cases the L^2 norm of the residual $f - u$ is the same. The ideal LPF exhibits sharper features. In addition two spectral elements $\phi(t)$ are shown. One can observe that the spatial response for any $\phi(t)$ is highly localized with very particular structures that emerge. The responses for the building windows (seen as black and white structure on the bottom row) highly resemble 2D Haar wavelets, which can be related to the analysis of [43, 48]. Other structures can be related to the explicit solutions of structures which retain their characteristic function, as analyzed in [7].

In Fig. 9 a possible direction for image analysis is shown. The first most salient peaks in the spectrum are examined (around times 60, 130, 170). We band-pass filter them, as the response is not fully concentrated near a singular time point. The composed three bands are shown on the third row, right. They are superimposed back on the original image. It is shown that they contain meaningful and well localized features with semantic meaning (in this case the eyes). Therefore they may serve as good candidates for image features in higher-level vision algorithms (e.g. face detection).

5 Conclusion

In this study a TV transform and a corresponding reconstruction formula were presented. This transform yields large response to all image structure which disappear at highly concentrated time intervals during the TV flow evolution. We can regard these structure as the "atoms" of the image, with respect to the total variation functional and gain a spectral understanding in the TV sense.

We have shown numerically that these structures are well localized spa-

tially and often represent significant image features with semantic meaning. Thus they can serve for image analysis and as input features to higher-level vision processing.

Extensions of this framework and relations to other TV-based formulations should be further investigated. For example, it may be the case that inverse-scale-space [13] can be interpreted as TV spectral low-pass filtering. Also other scale-spaces and regularization procedure, not based on the TV-functional, may be generalized using a similar approach.

References

[1] F. Andreu, C. Ballester, V. Caselles, and J. M. Mazn. Minimizing total variation flow. *Differential and Integral Equations*, 14(3):321–360, 2001.

[2] F. Andreu, V. Caselles, JI Dıaz, and JM Mazón. Some qualitative properties for the total variation flow. *Journal of Functional Analysis*, 188(2):516–547, 2002.

[3] J.F. Aujol and A. Chambolle. Dual norms and image decomposition models. *IJCV*, 63(1):85–104, June 2005.

[4] J.F. Aujol and G. Gilboa. Constrained and SNR-based solutions for TV-Hilbert space image denoising. *JMIV*, 26:217–237, 2006.

[5] J.F. Aujol, G. Gilboa, T. Chan, and S. Osher. Structure-texture image decomposition – modeling, algorithms, and parameter selection. *International Journal of Computer Vision*, 67(1):111–136, 2006.

[6] S. Bartels, R.H. Nochetto, J. Abner, and A.J. Salgado. Discrete total variation flows without regularization. *arXiv preprint arXiv:1212.1137*, 2012.

[7] G. Bellettini, V. Caselles, and M. Novaga. The total variation flow in R^N. *Journal of Differential Equations*, 184(2):475–525, 2002.

[8] B. Berkels, M. Burger, M. Droske, O. Nemitz, and M. Rumpf. Cartoon extraction based on anisotropic image classification. In *Vision, Modeling, and Visualization Proceedings*, pages 293–300, 2006.

[9] X. Bresson and T.F. Chan. Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Problems and Imaging*, 2(4):455–484, 2008.

[10] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV 2004*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36, 2004.

[11] T. Brox and J. Weickert. A tv flow based local scale estimate and its application to texture discrimination. *Journal of Visual Communication and Image Representation*, 17(5):1053–1073, 2006.

[12] M. Burger, K. Frick, S. Osher, and O. Scherzer. Inverse total variation flow. *Multiscale Modeling & Simulation*, 6(2):366–395, 2007.

[13] M. Burger, G. Gilboa, S. Osher, and J. Xu. Nonlinear inverse scale space methods. *Comm. in Math. Sci.*, 4(1):179–212, 2006.

[14] A. Chambolle. An algorithm for total variation minimization and applications. *JMIV*, 20:89–97, 2004.

[15] A. Chambolle and P.L. Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76(3):167–188, 1997.

[16] T.F. Chan and S. Esedoglu. Aspects of total variation regularized l 1 function approximation. *SIAM Journal on Applied Mathematics*, 65(5):1817–1837, 2005.

[17] T.F. Chan and J. Shen. A good image model eases restoration - on the contribution of Rudin-Osher-Fatemi's BV image model, 2002. IMA preprints 1829.

[18] F.R.K. Chung. *Spectral graph theory*, volume 92. Amer Mathematical Society, 1997.

[19] R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.

[20] J. Darbon and M. Sigelle. Image restoration with discrete constrained total variation part i: Fast and exact optimization. *Journal of Mathematical Imaging and Vision*, 2006.

[21] I. Daubechies et al. *Ten lectures on wavelets*, volume 61. SIAM, 1992.

[22] Public domain photos website. http://publicphoto.org/. 2013.

[23] D.L. Donoho. De-noising by soft-thresholding. *Information Theory, IEEE Transactions on*, 41(3):613–627, 1995.

[24] Vincent Duval, Jean-François Aujol, and Luminita Vese. Projected gradient based color image decomposition. In *Scale Space and Variational Methods in Computer Vision*, pages 295–306. Springer, 2009.

[25] Y. Giga and R.V. Kohn. Scale-invariant extinction time estimates for some singular diffusion equations. *Hokkaido University Preprint Series in Mathematics*, (963), 2010.

[26] G. Gilboa. A spectral approach to total variation. In *A. Kuijper et al. (Eds.): SSVM 2013*, volume 7893 of *Lecture Notes in Computer Science*, pages 36–47. Springer, 2013.

[27] G. Gilboa, N. Sochen, and Y.Y. Zeevi. Estimation of optimal PDE-based denoising in the SNR sense. *IEEE Trans. on Image Processing*, 15(8):2269–2280, 2006.

[28] G. Gilboa, N. Sochen, and Y.Y. Zeevi. Variational denoising of partly-textured images by spatially varying constraints. *IEEE Trans. on Image Processing*, 15(8):2280–2289, 2006.

[29] T. Goldstein. Split bregman page and code. http://tag7.web.rice.edu/split_bregman.html. 2013.

[30] T. Goldstein and S. Osher. The split bregman method for l1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.

[31] J.J. Koenderink. The structure of images. *Biol. Cybern.*, 50:363–370, 1984.

[32] B. Luo, J.F. Aujol, and Y. Gousseau. Local scale measure from the topographic map and application to remote sensing images. *Multiscale Modeling & Simulation*, 8(1):1–29, 2009.

[33] S.L. Marple Jr and W.M. Carey. Digital spectral analysis with applications. *The Journal of the Acoustical Society of America*, 86:2043, 1989.

[34] P. Maurel, J.-F. Aujol, and G. Peyré. Locally parallel texture modeling. *SIAM Journal on Imaging Sciences*, 4(1):413–447, 2011.

[35] Y. Meyer. Wavelets-algorithms and applications. *Wavelets-Algorithms and applications Society for Industrial and Applied Mathematics Translation., 142 p.*, 1, 1993.

[36] Y. Meyer. Oscillating patterns in image processing and in some nonlinear evolution equations, March 2001. The 15th Dean Jacquelines B. Lewis Memorial Lectures.

[37] M. Nikolova. A variational approach to remove outliers and impulse noise. *JMIV*, 20(1-2):99–120, 2004.

[38] S. Osher, A. Sole, and L. Vese. Image decomposition and restoration using total variation minimization and the H^{-1} norm. *SIAM Multiscale Modeling and Simulation*, 1(3):349–370, 2003.

[39] Stanley Osher. Private communication. 2013.

[40] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *PAMI*, 12(7):629–639, 1990.

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

[42] E.Y. Sidky and X. Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in medicine and biology*, 53(17):4777, 2008.

[43] G. Steidl, J. Weickert, T. Brox, P. Mrzek, and M. Welk. On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDEs. *SIAM Journal on Numerical Analysis*, 42(2):686–713, 2004.

[44] P. Stoica and R.L. Moses. *Introduction to spectral analysis*, volume 89. Prentice Hall Upper Saddle River, NJ, 1997.

[45] D. Strong and T.F. Chan. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems*, 19(6):S165–S187, 2003.

[46] L. Vese and S. Osher. Modeling textures with total variation minimization and oscillating patterns in image processing. *Journal of Scientific Computing*, 19:553–572, 2003.

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

[48] M. Welk, G. Steidl, and J. Weickert. Locally analytic schemes: A link between diffusion filtering and wavelet shrinkage. *Applied and Computational Harmonic Analysis*, 24(2):195–224, 2008.

[49] A.P. Witkin. Scale space filtering. In *Proc. Int. Joint Conf. On Artificial Intelligence*, pages 1019–1023, 1983.