A Spectral Approach to Total Variation

Guy Gilboa

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

Abstract. The total variation (TV) functional is explored from a spectral perspective. We formulate a TV transform based on the second time derivative of the total variation flow, scaled by time. In the transformation domain disks yield impulse responses. This transformation can be viewed as a spectral domain, with somewhat similar intuition of classical Fourier analysis. A simple reconstruction formula from the TV spectral domain to the spatial domain is given. We can then design low-pass, high-pass and band-pass TV filters and obtain a TV spectrum of signals and images.

Keywords: Scale Space, total variation, image analysis.

1 Introduction

The total variation (TV) functional is today a fundamental regularizing tool in image processing. It is employed for denoising and deconvolution [30, 12, 26, 28, 27, 20], optical-flow [8], tomographic reconstruction [31], texture and image analysis [7, 4, 3, 35, 21] and more. Since its introduction in [30] in the context of image processing many studies have been devoted to its analysis and interpretation, e.g. [12, 26, 13, 14]. 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. [24, 33]). For more complex non-stationary signals, such as images and speech, harmonic analysis methods were developed in the form of wavelets [17, 25, 18], spectral graph theory [15] and diffusion maps [16]. We explore a way to provide spectral information for total variation analysis.

In [32] Steidl et al have shown the close relations, and equivalence in a 1D discrete setting, of the Haar wavelets to both TV regularization [30] and TV flow [1]. This was later developed for a 2D setting in [37]. The development of features in the scale space framework [38, 22, 29, 36] and the emergence of critical points were studied for example in [22, 9, 23, 34, 13, 21]. This work relies on the established theory of the TV flow proposed by Andreu et al in [1] and further developed in [2, 6, 32, 10, 5, 19] and the references therein.

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:

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

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 [26, 34], TV-flow [1, 2, 6], inverse-scale-space evolutions [11] 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.

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:

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

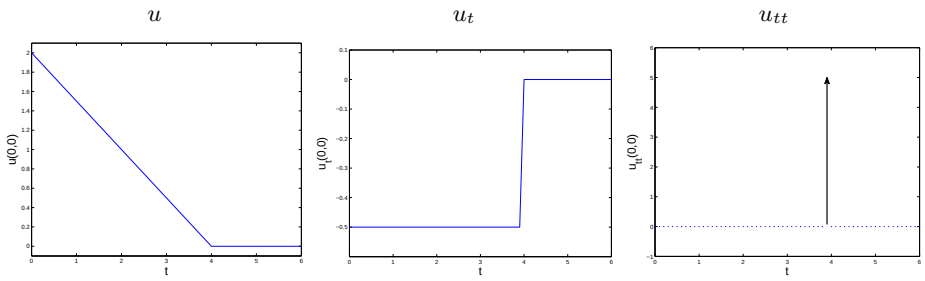

Fig. 1. Illustrating the evolution of a disk in $\mathcal{R}^2$. The value is within $|x| < r$, for example at $(x_1 = 0, x_2 = 0)$. The second derivative is an impulse at time t_d. [here we set $r = 4$, $h = 2$ and therefore $t_d = 4$].

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

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:

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

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 [19] 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. 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 $u|_0^\infty = \bar{f} - f$. □

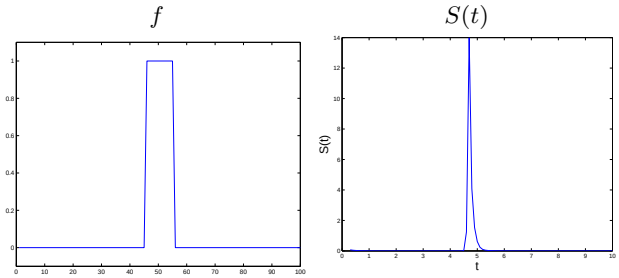

Fig. 2. A single one-dimensional disk and the corresponding numerical spectral response $S(t)$.

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

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

The spectral response roughly corresponds to the amplitude of the response in a Fourier domain (see Fig. 3). 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 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:

$$\phi_H(t) = \phi(t)H(t). \tag{6}$$

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

$$f_H(x) = \int_0^\infty \phi_H(t)dt + \bar{f}, \tag{7}$$

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 define the following ideal spectral filters:*

(i) Ideal low-pass filter:

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

(ii) Ideal high-pass filter:

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

(iii) Ideal band-pass filter:

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

(iv) Ideal band-stop filter:

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

2.3 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. 5, 6 for the spectral response and selected elements $\phi(t)$ of two image examples.

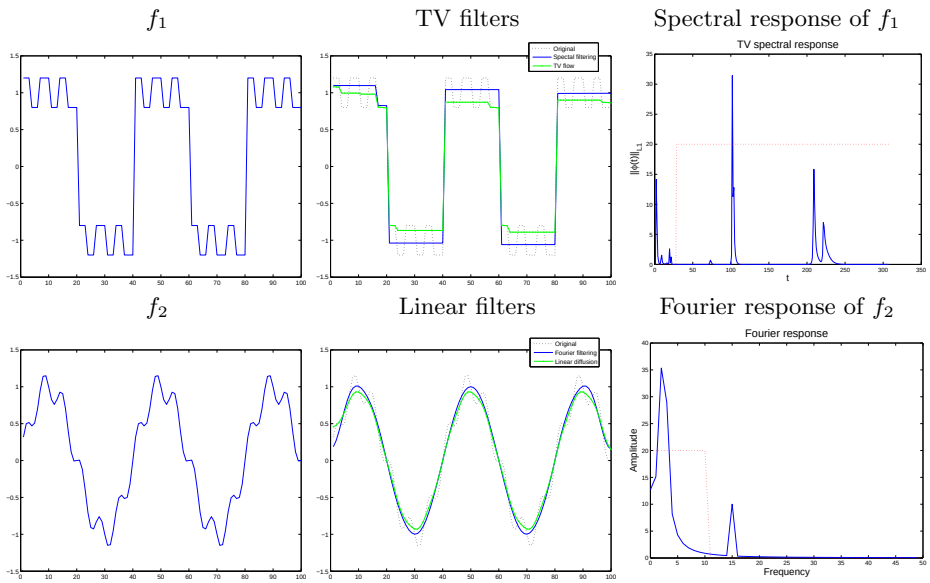

Fig. 3. One dimensional example of ideal low pass filtering versus scale-space low pass filtering. Top row, processing f_1 (left), middle - response by spectral filtering (full blue line), and by TV flow (dotted green line). On the top right the spectral response is shown. On the bottom row an analogue linear case filters f_2 with Fourier ideal LPF (full blue line) versus linear diffusion (dotted green line).

3 Examples

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

In Fig. 3 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. 3 top row: $f_1 = \text{sign}(\sin 2\pi\varphi_1) + 0.2\,\text{sign}(\sin 2\pi\varphi_2)$. The spectral 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

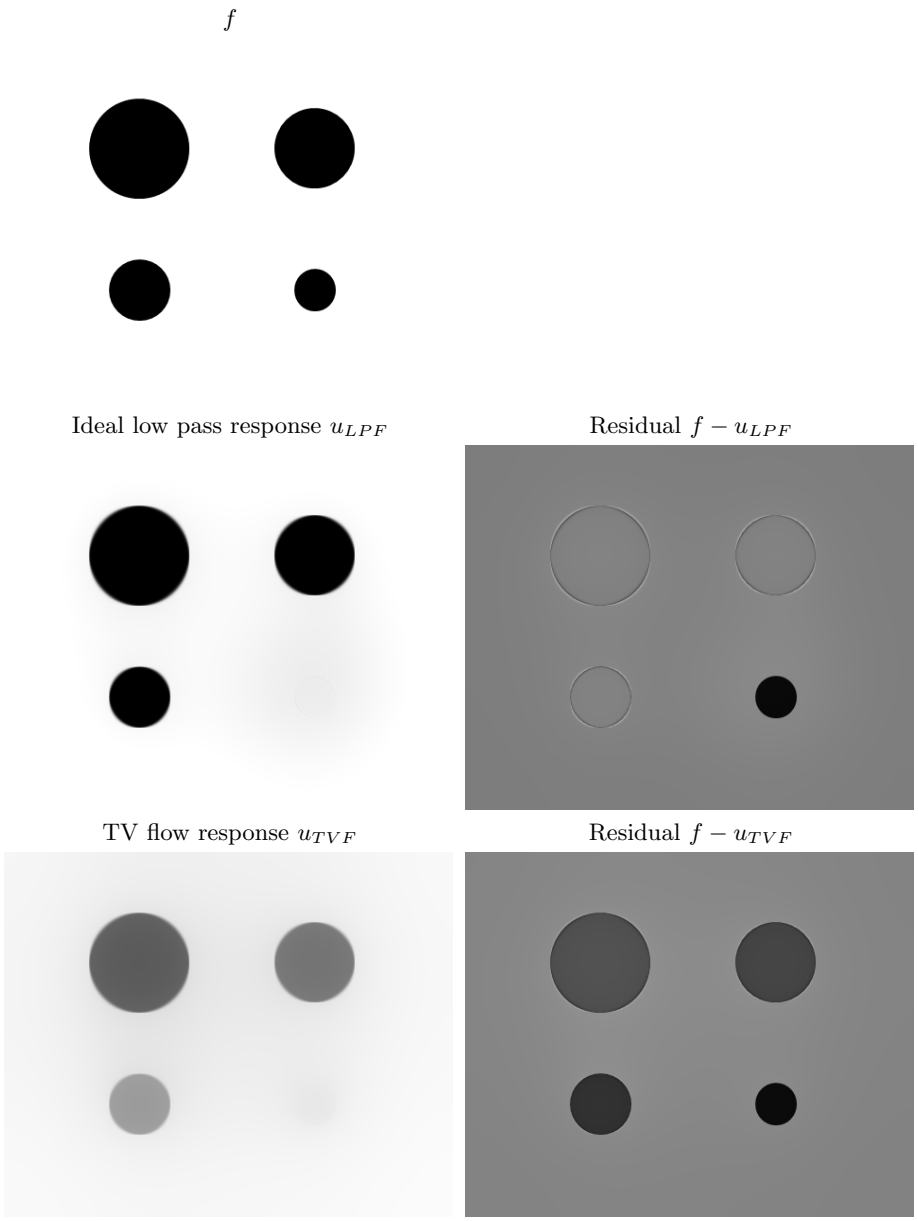

Fig. 4. Comparison between the ideal low-pass filter response and TV-flow. In both cases the response is shown for the minimal extent of filtering in which the smallest circle completely vanishes. One sees the considerable reduction of contrast of the larger circles in the TV-flow versus the sharp and stable results of the ideal TV LPF.

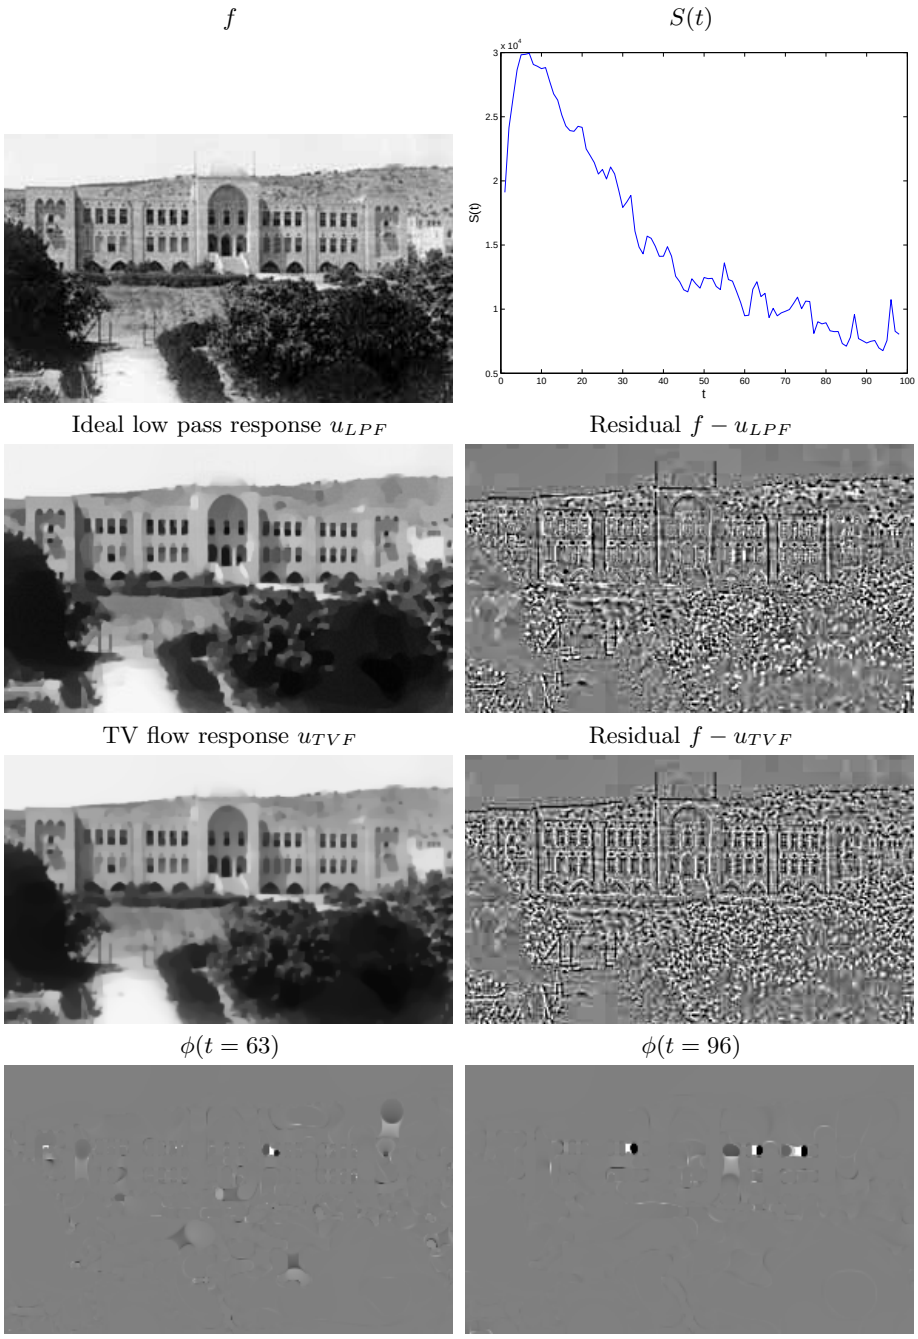

Fig. 5. Old Technion image. Results of ideal low pass filtering. This is compared to TV-flow with equivalent filtering in the L^2 sense (the norm of the residual, $\|f - u\|_{L^2}$, is the same). In addition, two examples of $\phi(t)$ are shown for different t values.

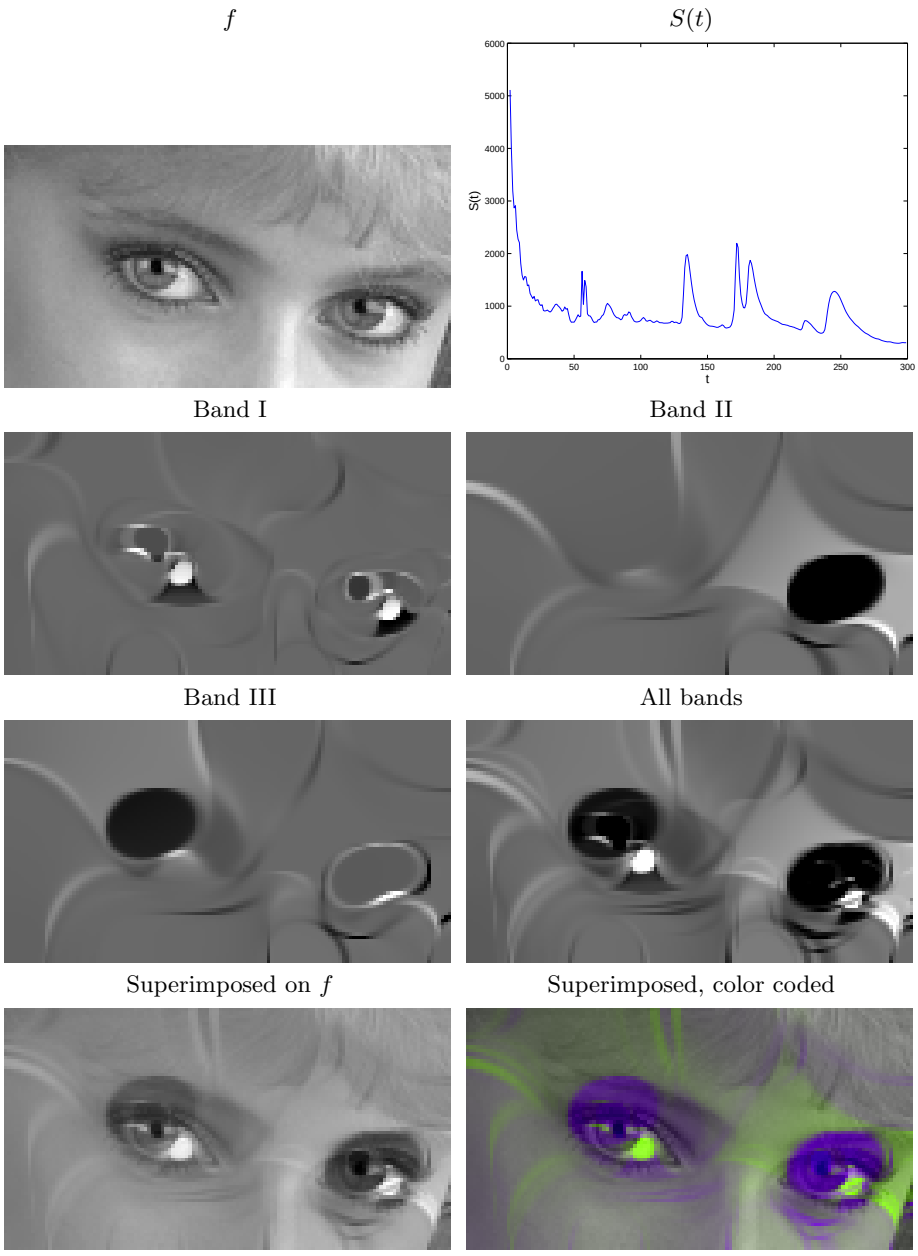

Fig. 6. Feature extraction example. Salient features ate depicted as spectral peaks (top right). The first three spectral peaks are shown as Bands I-III. These bands are reconstructed together at the third row, right. This reconstruction is then superimposed on the image to show the localization of the bands. Bottom right - a color coded visualization of the image with the selected bands.

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. 4 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. 5 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 [32, 37]. Other structures can be related to the explicit solutions of structures which retain their characteristic function, as analyzed in [6].

In Fig. 6 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).

4 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 spatially 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 [11] 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, 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.
5. S. Bartels, R.H. Nochetto, J. Abner, and A.J. Salgado. Discrete total variation flows without regularization. *arXiv preprint arXiv:1212.1137*, 2012.
6. G. Bellettini, V. Caselles, and M. Novaga. The total variation flow in R^N. *Journal of Differential Equations*, 184(2):475–525, 2002.
7. 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.
8. 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.
9. 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.
10. M. Burger, K. Frick, S. Osher, and O. Scherzer. Inverse total variation flow. *Multiscale Modeling & Simulation*, 6(2):366–395, 2007.
11. M. Burger, G. Gilboa, S. Osher, and J. Xu. Nonlinear inverse scale space methods. *Comm. in Math. Sci.*, 4(1):179–212, 2006.
12. A. Chambolle and P.L. Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76(3):167–188, 1997.
13. 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.
14. 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.
15. F.R.K. Chung. *Spectral graph theory*, volume 92. Amer Mathematical Society, 1997.
16. R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
17. I. Daubechies et al. *Ten lectures on wavelets*, volume 61. SIAM, 1992.
18. D.L. Donoho. De-noising by soft-thresholding. *Information Theory, IEEE Transactions on*, 41(3):613–627, 1995.
19. Y. Giga and R.V. Kohn. Scale-invariant extinction time estimates for some singular diffusion equations. *Hokkaido University Preprint Series in Mathematics*, (963), 2010.
20. 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.

21. 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.
22. J.J. Koenderink. The structure of images. *Biol. Cybern.*, 50:363–370, 1984.
23. 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.
24. S.L. Marple Jr and W.M. Carey. Digital spectral analysis with applications. *The Journal of the Acoustical Society of America*, 86:2043, 1989.
25. Y. Meyer. Wavelets-algorithms and applications. *Wavelets-Algorithms and applications Society for Industrial and Applied Mathematics Translation., 142 p.*, 1, 1993.
26. Y. Meyer. Oscillating patterns in image processing and in some nonlinear evolution equations, March 2001. The 15th Dean Jacquelines B. Lewis Memorial Lectures.
27. M. Nikolova. A variational approach to remove outliers and impulse noise. *JMIV*, 20(1-2):99–120, 2004.
28. 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.
29. P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *PAMI*, 12(7):629–639, 1990.
30. L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
31. 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.
32. 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.
33. P. Stoica and R.L. Moses. *Introduction to spectral analysis*, volume 89. Prentice Hall Upper Saddle River, NJ, 1997.
34. D. Strong and T.F. Chan. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems*, 19(6):S165–S187, 2003.
35. 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.
36. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner-Verlag, Stuttgart, Germany, 1998.
37. Martin Welk, Gabriele Steidl, and Joachim Weickert. Locally analytic schemes: A link between diffusion filtering and wavelet shrinkage. *Applied and Computational Harmonic Analysis*, 24(2):195–224, 2008.
38. A.P. Witkin. Scale space filtering. In *Proc. Int. Joint Conf. On Artificial Intelligence*, pages 1019–1023, 1983.