Higher Degree Total Variation Regularization for 3-D Image Recovery

Yue Hu†, *Student Member, IEEE*, Greg Ongie†, and Mathews Jacob, *Senior Member, IEEE*

***Abstract*— We introduce a family of novel image regularization penalties for three-dimensional (3-D) signals called 3-D higher degree total variation (HDTV). These penalties further extend our previously introduced HDTV penalties, which generalize the popular total variation (TV) penalty to incorporate higher degree image derivatives. We make use of a fast alternating minimization algorithm for solving 3-D image recovery problems with HDTV regularization. Numerical experiments on compressed sensing recovery of 3-D magnetic resonance images show that HDTV and generalized HDTV improves the image quality significantly compared with TV. We also investigate the relationship between the recently introduced Hessian-Shatten norms and HDTV.**

I. INTRODUCTION

The total variation (TV) image regularization penalty is widely used in many image recovery problems, including denoising, compressed sensing, and deblurring [1]. The good performance of TV penalty may be attributed to its desirable properties such as convexity, invariance to rotations and translations, and ability to preserve image edges. However, one drawback of this scheme is that the reconstructed images can contain undesirable patchy or staircase artifacts, since TV regularization promotes sparse gradients.

We recently introduced a family of novel image regularization penalties termed as higher degree TV (HDTV) to overcome the above problem [2]. These penalties are defined as the mixed L_1-L_p norm ($p = 1$ or 2) of the nth degree directional image derivatives. The HDTV penalties inherit the above mentioned desirable properties of TV functional. Experiments on two-dimensional (2-D) images demonstrate that HDTV regularization minimizes the staircase and patchy artifacts brought in by TV, while still enhancing edge and ridge-like features in the image. Moreover, the algorithm is also observed to provide improved reconstructions, both visually and quantitatively.

The HDTV penalties were originally designed for 2-D image reconstruction problems and were defined solely in terms of 2-D directional derivatives. The direct extension of the scheme to 3-D was challenging due to the high computational complexity of our original iteratively reweighted majorize minimize (IRMM) algorithm [2] for solving HDTV regularized inverse problems.

Adapting an efficient algorithm we recently introduced for solving image recovery problems with 2-D HDTV regularization [6], in this work we extend the HDTV penalties to three-dimensional (3-D) signals. By design, 3-D HDTV penalties also inherit the desirable properties of TV and 2-D HDTV such as translation- and rotation-invariance, scale co-variance, and convexity. Our algorithm for solving 3-D HDTV regularized inverse problems is based on an alternating minimization scheme, which alternates between two efficiently solved subproblems given by a shrinkage and the inversion of a linear system.

We also study the relationship between 3-D HDTV and the recently proposed regularization penalties based on Hessian-Shatten norms [7]. Specifically, we show that the second degree HDTV penalty and the Hessian-Shatten $p = 1$ norm are equivalent as semi-norms. We demonstrate with numerical experiments that this equivalence carries over to the discrete setting, as well.

Finally, we demonstrate the utility of 3-D HDTV regularization in the context of practical inverse problems arising in medical imaging, including deblurring and denoising of 3D fluorescence microscope images, and compressed sensing MR image recovery of 3-D angiography datasets. We show that 3-D HDTV routinely outperforms TV in terms of the SNR of reconstructed images and its ability to preserve ridge-like details in the datasets.

II. HIGHER DEGREE TOTAL VARIATION

Image recovery with TV and HDTV regularization may be understood in a variational framework, where we consider the recovery of a continuously differentiable signal $f : \Omega \to \mathbb{C}$, $\Omega \subset \mathbb{R}^d$, from its noisy and degraded measurements $\mathbf{b}$. We model the measurements as $y = \mathcal{A}(f) + n$, where n is assumed to be Gaussian distributed white noise and $\mathcal{A}$ is a linear operator representing the degradation process. We may formulate the recovery of f as the following optimization problem

$$\min_f \|\mathcal{A}(f) - \mathbf{b}\|^2 + \lambda \mathcal{J}(f), \tag{1}$$

where $\|\mathcal{A}(f) - \mathbf{b}\|^2$ is the data fidelity term, $\mathcal{J}(f)$ is a regularization penalty, and the parameter λ balances the two terms, and is chosen so that the signal-to-error ratio is maximized.

The standard isotropic TV regularization penalty is defined as the L_1 norm of the gradient magnitude,

$$\mathrm{TV}(f) = \int_\Omega \|\nabla f(\mathbf{r})\| \, d\mathbf{r},$$

where $||\cdot||$ denotes the Euclidean norm. In [2] we showed that the 2-D TV penalty can be reinterpreted as the mixed L_1-L_2 or L_1-L_1 norm of image directional derivatives. This

†Both authors have contributed equally to the paper. Y. Hu is with the Department of Electrical and Computer Engineering, University of Rochester, NY, USA. G. Ongie is with the Department of Mathematics, University of Iowa, IA, USA. M. Jacob is with the Department of Electrical and Computer Engineering, University of Iowa, IA, USA.

This work is supported by grants NSF CCF-0844812, NSF CCF-1116067, NIH 1R21HL109710-01A1, and ACS RSG-11-267-01-CCE.

observation led us to propose two families of higher degree TV (HDTV) regularization penalties in 2-D, specified by

$$\text{I-HDTV}n(f) = \int_\Omega \left(\frac{1}{2\pi} \int_0^{2\pi} |f_{\theta,n}(\mathbf{r})|^2 d\theta \right)^{\frac{1}{2}} d\mathbf{r}, \quad (2)$$

$$\text{HDTV}n(f) = \int_\Omega \left(\frac{1}{2\pi} \int_0^{2\pi} |f_{\theta,n}(\mathbf{r})| d\theta \right) d\mathbf{r}, \quad (3)$$

where $f_{\theta,n}$ is the nth degree directional derivative of f in the direction $\mathbf{u}_\theta = [\cos(\theta), \sin(\theta)]$, defined as

$$f_{\theta,n}(\mathbf{r}) = \frac{\partial^n}{\partial \gamma^n} f(\mathbf{r} + \gamma \mathbf{u}_\theta) \Big|_{\gamma=0}.$$

The family of penalties defined by (2) and (3) were given the name *isotropic* and *anisotropic HDTV*, respectively. It is evident from (2) and (3) that the 2-D HDTV penalties preserve many of the desirable properties of the TV penalty, such as invariance under translations and rotations and scale co-variance. Furthermore, practical experiments in [2] demonstrate that HDTV regularization outperforms TV regularization in many image recovery tasks, in terms of both SNR and the visual quality of reconstructed images. Our experiments also indicate that the anisotropic case, which corresponds to the fully separable L_1-L_1 penalty, typically exhibits better performance in image recovery tasks over isotropic HDTV.

The HDTV penalties have a natural extension to 3-D signals by considering the L_p norm of all directional derivatives in 3-D. We define

$$\text{HDTV}[n,p](f) = \int_\Omega \left(\int_{\mathbb{S}^2} |f_{\mathbf{u},n}(\mathbf{r})|^p d\mathbf{u} \right)^{\frac{1}{p}} d\mathbf{r}, \quad (4)$$

where $\mathbb{S}^2 = \{\mathbf{u} \in \mathbb{R}^3 : ||\mathbf{u}|| = 1\}$ and $f_{\mathbf{u},n}$ is the nth degree directional derivative defined as

$$f_{\mathbf{u},n}(\mathbf{r}) = \frac{\partial^n}{\partial \gamma^n} f(\mathbf{r} + \gamma \mathbf{u}) \Big|_{\gamma=0}; \ \mathbf{u} \in \mathbb{S}^2.$$

Due to its importance in the sequel, the $p = 1$ penalty we will simply denote as HDTVn, i.e. HDTVn = HDTV$[n, 1]$.

By design the 3-D HDTV penalties are guaranteed to be rotation- and translation-invariant, and convex. It is also clear they are also contrast and scale covariant, i.e. for all $\alpha \in \mathbb{C}$, HDTV$[n,p](\alpha \cdot f) = |\alpha|$HDTV$[n,p](f)$ and HDTV$[n,p](f_\alpha) = |\alpha|^{n-3}$HDTV$[n,p](f)$, where $f_\alpha(x) := f(\alpha \cdot x)$.

A. Relation to Hessian-Shatten Norms

Recently Lefkimmiatis *et al.* [7] introduced a family of second-degree regularization penalties known as *Hessian-Shatten norms*, defined as

$$\text{HS}p(f) = \int_\Omega ||\mathcal{H}f(\mathbf{r})||_{\mathcal{S}_p}, \quad \forall \ 1 \le p \le \infty, \quad (5)$$

where $\mathcal{H}f(\mathbf{r})$ is the Hessian matrix of f at $\mathbf{r}$, and $||\cdot||_{\mathcal{S}_p}$ is the Shatten p-norm, defined as $||X||_{\mathcal{S}_p} = ||\sigma(X)||_p$ where $\sigma(X)$ denotes the vector containing the singular values of the matrix X. There is a close relationship between the $p = 1$ case of (5) and the anisotropic second degree HDTV penalty, HDTV2, as demonstrated in the following proposition:

Proposition 1: The penalties HDTV2 and HS1 are equivalent as semi-norms over $\mathcal{C}^2(\Omega, \mathbb{R})$, $\Omega \subset \mathbb{R}^d$, in dimension $d = 2$ or 3, with bounds

$$(1 - \delta) \cdot \text{HS1}(f) \le C \cdot \text{HDTV2}(f) \le \text{HS1}(f),$$

where $\delta = 0.37$ for $d = 2$, $\delta = 0.43$ for $d = 3$, and C is a normalization constant independent of f.

We omit a full proof for brevity. The essential idea is to re-express the HDTV2 penalty as the integral of the function

$$\Phi(\mathbf{r}) = \int_{\mathbb{S}^{d-1}} \left| \mathbf{u}^T \mathbf{diag}[\boldsymbol{\lambda} f(\mathbf{r})] \, \mathbf{u} \right| d\mathbf{u},$$

where $\mathbf{diag}[\boldsymbol{\lambda}_f(\mathbf{r})]$ is the diagonal matrix whose entries are given by $\boldsymbol{\lambda}_f(\mathbf{r})$, the Hessian eigenvalues of f at $\mathbf{r}$. Then the result follows from bounding the quantity $\Phi(\mathbf{r})/||\boldsymbol{\lambda} f(\mathbf{r})||_{\ell_1}$.

In particular, one can show $C \cdot \text{HDTV2}(f) = \text{HS1}(f)$ if the Hessian matrices of f at all spatial locations are either positive or negative semi-definite, i.e. have all non-negative eigenvalues or all non-positive eigenvalues. In natural images only a fraction of the pixels or voxels will have Hessian eigenvalues with mixed sign, thus we expect the HS1 and HDTV2 penalties to be nearly proportional and to behave very similarly in applications. Our experiments in the results section are consistent with this observation.

III. Numerical Implementation

To solve the image recovery problem (1) with 3-D HDTV regularization we make use of a fast alternating minimization algorithm originally introduced for 2-D HDTV in [6]. The algorithm is an adaptation of a half-quadratic minimization method [8] used for solving image recovery problems with TV regularization [9], [10]. The algorithm alternates between two well-defined subproblems: a shrinkage step and the inversion of a linear system. The latter subproblem is much simpler to solve if the measurement operator $\mathcal{A}$ has a diagonal form in the Fourier domain, as is the case for many practical inverse problems we consider, such as denosing, deblurring, and compressed sensing MR images recovery. We also note that this algorithm is designed specifically for anisotropic ($p = 1$) HDTV penalties, hence we focus on those cases in this work.

Some modifications to the algorithm must be made in the 3-D setting. Specifically, the method by which we discretize the angular integral in (4) and discretize the directional derivative operators is different in 3-D. We present these details below.

A. Sampling the unit sphere in 3-D

We approximate the inner integral in (4) with a Riemann sum by uniformly sampling points in $\mathbb{S}^2$. In the 2-D case, this can be achieved by parameterizing $\mathbf{u} \in \mathbb{S}^1$ as $\mathbf{u}_\theta = [\cos(\theta), \sin(\theta)]$, then discretizing the parameter θ as $\theta_i = i\frac{2\pi}{K}$, for $i = 1, ..., K$, where K is the specified number of

sample points. However, in the 3-D case, if we sample the usual parameterization of $\mathbb{S}^2$, i.e.

$$\mathbf{u}_{\theta,\phi} = [\cos(\theta)\sin(\phi), \sin(\theta)\sin(\phi), \cos(\phi)]$$

for $\theta \in [0, 2\pi], \phi \in [0, \pi]$, by uniformly discretizing θ and ϕ, the samples we obtain are heavily clustered at the poles of the sphere, providing a poor approximation of the integral. Instead, we make use of the ISOI software package [11], [12] based on the HEALPix spherical sampling method [13] to deterministically generate uniformly spaced samples of $\mathbb{S}^2$. We find $K \approx 64$ samples are sufficient to approximate the integral in (4). Note that these sample points may be computed in advance and stored in memory to reduce the computational overhead.

B. Steerability of Directional Derivatives

The direct evaluation of (4) by discretizing $\mathbb{S}^2$ is computationally expensive. The computational complexity can be considerably reduced by exploiting the *rotation steerability* of nth degree directional derivatives. Namely, the first degree directional derivatives $f_{\mathbf{u},1}$ have the equivalent expression

$$f_{\mathbf{u},1}(\mathbf{r}) = \mathbf{u}^T \nabla f(\mathbf{r}); \quad \forall\ \mathbf{u} \in \mathbb{S}^2.$$

Similarly, higher degree directional derivatives $f_{\mathbf{u},n}(\mathbf{r})$ can be expressed as the separable vector product

$$f_{\mathbf{u},n}(\mathbf{r}) = \mathbf{s}(\mathbf{u})^T \nabla_n f(\mathbf{r}); \quad \forall\ \mathbf{u} \in \mathbb{S}^2,$$

where, $\mathbf{s}(\mathbf{u})$ is the vector of polynomials in the components of $\mathbf{u}$ and $\nabla_n f(\mathbf{r})$ is the vector of all nth degree partial derivatives of f. For example, in the second degree case ($n = 2$), we may choose

$$\mathbf{s}(\mathbf{u}) = \left[u_x^2, u_y^2, u_z^2, 2u_x u_y, 2u_y u_z, 2u_x u_z\right]^T,$$
$$\nabla_2 f(\mathbf{r}) = \left[f_{xx}(\mathbf{r}), f_{yy}(\mathbf{r}), f_{zz}(\mathbf{r}), f_{xy}(\mathbf{r}), f_{yz}(\mathbf{r}), f_{xz}(\mathbf{r})\right]^T.$$

This shows that we may compute the nth degree directional derivatives for all $\mathbf{u} \in \mathbb{S}^2$ at all voxels with only a small number of filtering operations (six, in the above example).

C. Discretization of the derivative operators

The standard approach to approximate partial derivatives is by using finite difference operators. For example, the derivative of a 2-D signal along the x dimension is approximated as $\partial_x f[k_1, k_2] = f[k_1 + 1, k_2] - f[k_1, k_2]$. This approximation can be viewed as the convolution of f by $\Delta_1[k] = \varphi(k + \frac{1}{2})$, where $\varphi(x) = \partial\beta_1(x)/\partial x$ and $\beta_1(x)$ is the first degree B-spline [14]. However, this approximation does not possess rotation steerability, i.e. the directional derivative can not be expressed as the linear combination of the finite differences along x and y directions.

To obtain discrete operators that are approximately rotation steerable we approximate the nth order partial derivatives, $\partial^\alpha := \partial_x^{\alpha_1}\partial_y^{\alpha_2}\partial_z^{\alpha_3}$ for all mutli-indicies $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ with $|\alpha| = n$, as the convolution of the signal with the tensor product of derivatives of one-dimensional B-spline functions. That is, $\partial^\alpha f[\mathbf{k}] = (D^\alpha * f)[\mathbf{k}]$ where

$$D^\alpha[\mathbf{k}] = \beta_n^{(\alpha_1)}(k_1 + \delta) \otimes \beta_n^{(\alpha_2)}(k_2 + \delta) \otimes \beta_n^{(\alpha_3)}(k_3 + \delta), \quad (6)$$

for all $\mathbf{k} = (k_1, k_2, k_3) \in \mathbb{N}^3$, and where $\beta_n^{(m)}(x)$ denotes the mth order derivative of a nth degree B-spline. In order to obtain filters with small spacial support, we choose the $\delta = 1/2$ if n odd and $\delta = 0$ if n even; this results in filters supported in a $(n + 1)^3$ volume.

While the tensor product of B-spline functions are not strictly rotation steerable, B-splines approximate Gaussian functions as their degree increases, and the tensor product of Gaussians is exactly steerable. Thus, the approximation of derivatives we define above is approximately rotation steerable; see Fig. 1.

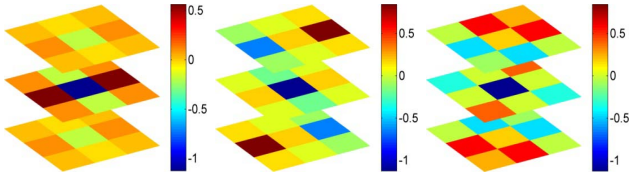

(a) $\theta = \phi = 0°$ (b) $\theta = 0°, \phi = 45°$ (c) $\theta = \phi = 45°$

Fig. 1. Second degree discrete directional derivative operators given by the B-spline approximation (6) at different orientations, specified by azimuthal angle θ and polar angle ϕ. Note that the operators are approximately rotation steerable.

IV. EXPERIMENTS

To demonstrate the utility of 3-D HDTV penalties, we investigate their use in the compressed sensing recovery of 3-D MR angiography datasets. Our experiments show that HDTV regularization routinely outperforms TV in terms of the SNR of reconstructed images and its ability to preserve ridge-like details in the 3-D datasets. Additionally, to verify the conclusions of Proposition 1 carry over to the discrete setting, we conduct experiments comparing the performance of the HDTV2 and HS1 penalties in denoising and deblurring 2-D images.

In each experiment, we optimize the regularization parameters to obtain the optimized signal-to-noise ratio (SNR) to ensure fair comparisons between different schemes. The SNR of the reconstruction is computed as:

$$\text{SNR} = -10\log_{10}\left(\frac{\|f_{\text{orig}} - \hat{f}\|_F^2}{\|f_{\text{orig}}\|_F^2}\right),$$

where $\hat{f}$ is the reconstructed image, f_{orig} is the original image, and $\|\cdot\|_F$ is the Frobenius norm.

A. Compressed Sensing MRI Recovery With 3-D HDTV

In these experiments we consider the compressed sensing recovery of a single coil 3-D MR angiography dataset from noisy and undersampled measurements. We experiment on a 512×512×76 MR dataset obtained from [15], which we retroactively undersample using a variable density random Fourier encoding with acceleration factor of 5. To these samples we also added 5 dB Gaussian noise with standard deviation 0.53. Shown in Fig. 2 are the maximum intensity projections (MIP) of the reconstructions obtained using various schemes. We note that 3-D HDTV preserves more edge-

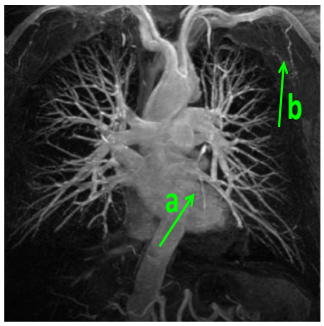

(a) actual

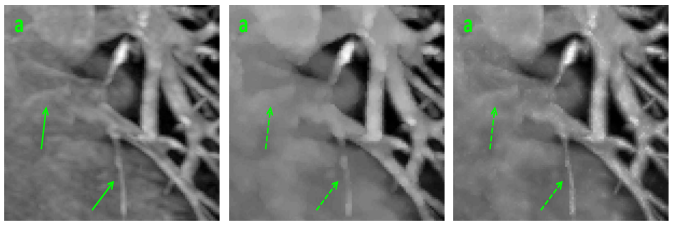

(b) actual (c) TV (d) HDTV2

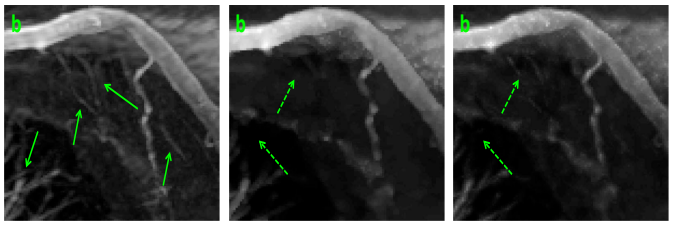

(e) actual (f) TV (g) HDTV2

Fig. 2. Compressed sensing recovery of 3-D MR angiography dataset from noisy and undersampled Fourier data (acceleration of 5 with 5 dB additive Gaussian noise). (a) shows the MIP of the original dataset, (b)-(d) and (e)-(g) show zoomed versions two regions of the original dataset and its reconstructions obtained using TV and HDTV2. We observe that 3D-HDTV methods preserve more line-like features compared with 3D-TV (indicated by green arrows).

and ridge-like details compared with standard 3D-TV. In Table 1 we provide quantitative comparisons between HDTV and TV at different accelerations. Note that both the HDTV2 and HDTV3 penalties routinely outperform TV in SNR.

	MRA (A=5)	**MRA**(A=1.5)	**Cardiac**
TV	13.87	14.53	18.37
HDTV2	**14.23**	**15.11**	**18.56**
HDTV3	14.01	14.70	18.50

Table 1: SNR (in dB) comparison of TV and HDTV for compressed sensing MR image recovery of 3-D datasets. 'MRA' is the dataset shown in Fig. (2), where A denotes the acceleration factor, and 'Cardiac' is a cardiac MR dataset.

B. Comparison of HS1 and HDTV2 Penalties

In Table 2 we present quantitative comparisons of the performance of the HS1 and HDTV2 regularization penalties on 2-D test images in the context of denoising, and deblurring. These experiments confirm that the discrete versions of these penalties perform similarly in image reconstruction tasks, as predicted by Proposition 1. Fig. 3 shows that the reconstructions are visually similar, as well.

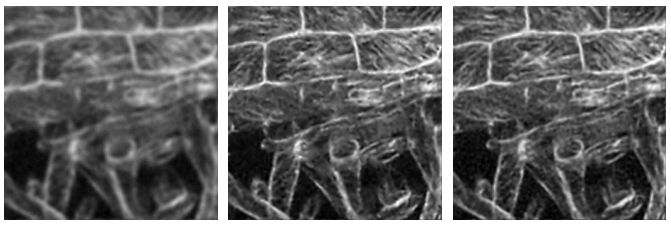

(a) Blurred image (b) HDTV2, 17.21dB (c) HS1, 17.13 dB

Fig. 3. Deblurring of a microscopy cell image. (a) is the blurred image, (b) and (c) show the deblurred images using HDTV2 and HS1 regularization, respectively. Both methods provide similar results, both visually and quantitatively, as predicted by Proposition 1.

	Denoising		Deblurring	
	brain	lena	cell1	cell2
TV	27.60	27.35	15.66	16.67
HDTV2	28.05	**27.65**	**16.19**	**17.21**
HS1	**28.08**	27.51	16.17	17.13

Table 2: SNR (in dB) comparison of deblurring and denoising of 2-D images using TV, HDTV2, and HS1 regularization. Note the SNR values are approximately the same in the case of HDTV2 and HS1, and both consistently outperform TV.

V. CONCLUSIONS

We extend our novel higher degree total variation (HDTV) image regularization penalty to 3-D signals. The penalty is essentially the L_1 norm of all directional derivatives at each voxel. We adapt a fast alternating minimization algorithm for 3-D HDTV that is considerably faster than our previous implementation for HDTV; the extension of the proposed scheme to 3-D is mainly enabled by this speedup. Our experiments demonstrate the improvement in image quality offered by the proposed scheme over TV for compressed sensing MR image reconstruction.

ACKNOWLEDGMENTS

The authors would like to thank Satish Ramani for helpful suggestions regarding the optimization algorithm.

REFERENCES

[1] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, Jan 1992.

[2] Y. Hu and M. Jacob, "Higher degree total variation (HDTV) regularization for image recovery," *IEEE Transactions on Image Processing*, vol. 21, no. 5, pp. 2559–2571, May 2012.

[3] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM J. Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.

[4] C. Li, W. Yin, H. Jiang, and Y. Zhang, "An efficient augmented lagrangian method with applications to total variation minimization," Rice University, CAAM Technical Report 12-14, 2012.

[5] C. Wu and X.-C. Tai, "Augmented Lagragian method, dual methods and split Bregman iteration for ROF, vectorial TV, and higher order models," *SIAM J. Imaging Sciences*, vol. 3, no. 3, pp. 300–339, 2010.

[6] Y. Hu, S. Ramani, and M. Jacob, "A fast majorize minimize algorithm for higher degree total variation regularization," in *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*. IEEE, 2013, pp. 326–329.

[7] S. Lefkimmiatis, J. Ward, and M. Unser, "Hessian Schatten-norm regularization for linear inverse problems," *Image Processing, IEEE Transactions on*, vol. 22, no. 5, pp. 1873–1888, 2013.

[8] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *Image Processing, IEEE Transactions on*, vol. 4, no. 7, pp. 932–946, 1995.

[9] M. Nikolova and M. K. Ng, "Analysis of half-quadratic minimization methods for signal and image recovery," *SIAM Journal on Scientific computing*, vol. 27, no. 3, pp. 937–966, 2005.

[10] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.

[11] J. C. Mitchell, "Sampling rotation groups by successive orthogonal images," *SIAM Journal on Scientific Computing*, vol. 30, no. 1, pp. 525–547, 2008.

[12] A. Yershova, S. Jain, S. M. LaValle, and J. C. Mitchell, "Generating uniform incremental grids on SO(3) using the Hopf fibration," *The International journal of robotics research*, vol. 29, no. 7, pp. 801–812, 2010.

[13] K. M. Gorski, E. Hivon, A. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, "Healpix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere," *The Astrophysical Journal*, vol. 622, no. 2, p. 759, 2005.

[14] M. Unser and T. Blu, "Wavelet theory demystified," *Signal Processing, IEEE Transactions on*, vol. 51, no. 2, pp. 470–483, 2003.

[15] "http://physionet.org/physiobank/database/images/."