

HIGHER DEGREE TOTAL VARIATION FOR 3-D IMAGE RECOVERY

Greg Ongie^{*†} Yue Hu^{*} Mathews Jacob[†]

^{*†} Department of Mathematics, University of Iowa, IA, USA.

^{*} Department of Electrical and Computer Engineering, University of Rochester, NY, USA.

[†] Department of Electrical and Computer Engineering, University of Iowa, IA, USA.

ABSTRACT

We extend the novel higher degree total variation (HDTV) image regularization penalties to 3-D signals. These penalties generalize the popular total variation (TV) penalty to incorporate higher degree image derivatives. We adapt a fast alternating minimization algorithm designed for solving 2-D image recovery problems with HDTV regularization to the 3-D setting. Numerical experiments on the compressed sensing recovery of 3-D magnetic resonance images show that HDTV and generalized HDTV improve the image quality significantly compared with TV. We also investigate the relationship between the recently introduced Hessian Schatten-norms and HDTV.

Index Terms— Higher degree total variation, 3-D image restoration, compressed sensing

1. 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 n th 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 provides improved reconstructions, both visually and quantitatively. Notably, it minimizes the staircase and patchy artifacts characteristic of TV, while still enhancing edge and ridge-like features in the image.

The HDTV penalties were originally designed for 2-D image recovery problems. The direct extension of the scheme to 3-D was challenging due to the high computational complexity of our original algorithm for solving HDTV regularized inverse problems. Adapting an efficient algorithm we recently introduced for solving image recovery problems with 2-D HDTV regularization [3], in this work we extend the HDTV penalties to three-dimensional (3-D) signals. The 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 present the modifications necessary to implement the algorithm in 3-D.

We also study the relationship between 3-D HDTV and the recently proposed regularization penalties based on Hessian Schatten-norms [4]. Specifically, we show that the second degree HDTV penalty and the Hessian Schatten-norm for $p = 1$ are equivalent as semi-norms. We demonstrate with numerical experiments that this equivalence carries over to the discrete setting, as well. The benefit of our scheme is that we may define HDTV penalties for arbitrarily high degree image derivatives, unlike the Hessian Schatten-norms which are strictly second degree penalties.

Finally, we demonstrate the utility of HDTV regularization in the context of compressed sensing MR image recovery of 3-D 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.

2. HIGHER DEGREE TOTAL VARIATION

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

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

This work is supported by grants NSF CCF-0844812, NSF CCF-1116067, NIH 1R21HL109710-01A1, ACS RSG-11-267-01-CCE, and ONR-N000141310202.

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,

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

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_p norm of image directional derivatives. This observation led us to propose the family of higher degree TV (HDTV) regularization penalties in 2-D, specified by

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

where $f_{\theta, n}$ is the n th degree directional derivative of f in the direction $\mathbf{u}_{\theta} = [\cos(\theta), \sin(\theta)]$, defined as $f_{\theta, n}(\mathbf{r}) = \frac{\partial^n}{\partial \gamma^n} f(\mathbf{r} + \gamma \mathbf{u}_{\theta})|_{\gamma=0}$. Our experiments also indicate that the anisotropic ($p = 1$) penalty typically exhibits better performance in image recovery tasks over other cases.

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

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

where $\mathbb{S}^2 = \{\mathbf{u} \in \mathbb{R}^3 : \|\mathbf{u}\| = 1\}$ and $f_{\mathbf{u}, n}$ is the n th degree directional derivative defined as $f_{\mathbf{u}, n}(\mathbf{r}) = \frac{\partial^n}{\partial \gamma^n} f(\mathbf{r} + \gamma \mathbf{u})|_{\gamma=0}$. Due to its importance in the sequel, the anisotropic ($p = 1$) penalty we will simply denote as HDTV_n , i.e. $\text{HDTV}_n = \text{HDTV}[n, 1]$.

By design the 3-D HDTV penalties are guaranteed to be rotation- and translation-invariant, and convex for $p \geq 1$. It is also clear they are also contrast and scale covariant, i.e. for all $\alpha \in \mathbb{C}$, $\text{HDTV}[n, p](\alpha \cdot f) = |\alpha| \text{HDTV}[n, p](f)$ and $\text{HDTV}[n, p](f_{\alpha}) = |\alpha|^{n-3} \text{HDTV}[n, p](f)$, where $f_{\alpha}(x) := f(\alpha \cdot x)$.

2.1. Relation to Hessian Schatten-Norms

Recently Lefkimmiatis *et al.* [4] introduced a family of second degree regularization penalties known as *Hessian Schatten-norms*, defined as

$$\text{HSp}(f) = \int_{\Omega} \|\mathcal{H}f(\mathbf{r})\|_{\mathcal{S}_p}, \quad \forall 1 \leq p \leq \infty, \quad (3)$$

where $\mathcal{H}f(\mathbf{r})$ is the Hessian matrix of f at \mathbf{r} , and $\|\cdot\|_{\mathcal{S}_p}$ is the Schatten 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 (3) and the anisotropic second degree HDTV penalty, HDTV_2 , as demonstrated in the following proposition:

Proposition 2.1 *The penalties HDTV_2 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*

$$(1 - \delta) \cdot \text{HS1}(f) \leq C \cdot \text{HDTV}_2(f) \leq \text{HS1}(f),$$

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 use $f_{\mathbf{u}, 2}(\mathbf{r}) = \mathbf{u}^T \mathcal{H}f(\mathbf{r}) \mathbf{u}$ and a change of variables to re-express the HDTV_2 penalty as the integral of the function

$$\Phi(\mathbf{r}) = \int_{\mathbb{S}^{d-1}} |\mathbf{u}^T \text{diag}[\lambda_f(\mathbf{r})] \mathbf{u}| d\mathbf{u},$$

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

In particular, one can show $C \cdot \text{HDTV}_2(f) = \text{HS1}(f)$ if the Hessian matrices of f at all spatial locations are either positive or negative semi-definite. In natural images only a fraction of the pixels or voxels will have Hessian eigenvalues with mixed sign, thus we expect the HS1 and HDTV_2 penalties to be nearly proportional and to behave very similarly in applications. Our experiments in the results section are consistent with this observation.

3. 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 [3]. 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 image 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 (2) and discretize the directional derivative operators is different in 3-D. We present these details below.

3.1. Sampling the unit sphere in 3-D

We approximate the inner integral in (2) 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 $i = 1, \dots, K$, where K is the specified number of

sample points. In the 3-D case, we could rewrite the integral using the usual parameterization of \mathbb{S}^2 , i.e.

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

for $\theta \in [0, 2\pi]$, $\phi \in [0, \pi]$, with volume element $d\mathbf{u} = \sin(\phi)d\phi d\theta$. However, if we uniformly discretize θ and ϕ , the samples we obtain are heavily clustered at the poles of the sphere, providing a poor approximation of the integral when the number of samples is small. Instead, we make use of the ISOI software package [5] to deterministically generate uniformly spaced samples of \mathbb{S}^2 . In practice we find $K \approx 64$ samples provide a sufficient approximation. Note that these sample points may be computed in advance and stored in memory to reduce the computational overhead.

3.2. Steerability of Directional Derivatives

The direct evaluation of (2) by discretizing \mathbb{S}^2 is computationally expensive. The computational complexity can be considerably reduced by exploiting the *rotation steerability* of n th 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 vector of polynomials in the components of \mathbf{u} and $\nabla_n f(\mathbf{r})$ is the vector of all n th degree partial derivatives of f . For example, in the second degree case ($n = 2$), we may choose $\mathbf{s}(\mathbf{u}) = [u_x^2, u_y^2, u_z^2, 2u_x u_y, 2u_y u_z, 2u_x u_z]^T$ and $\nabla_2 f(\mathbf{r}) = [f_{xx}(\mathbf{r}), f_{yy}(\mathbf{r}), f_{zz}(\mathbf{r}), f_{xy}(\mathbf{r}), f_{yz}(\mathbf{r}), f_{xz}(\mathbf{r})]^T$. This shows that we may compute the n th 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) corresponding to each n th degree partial derivative.

3.3. Discretization of the derivative operators

To obtain discrete operators that are approximately rotation steerable we approximate the n th order partial derivatives, $\partial^\alpha := \partial_{x_1}^{\alpha_1} \partial_{x_2}^{\alpha_2} \partial_{x_3}^{\alpha_3}$ for all multi-indices $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ 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

$$D^\alpha[\mathbf{k}] = \beta_n^{(\alpha_1)}(k_1 + \delta) \otimes \beta_n^{(\alpha_2)}(k_2 + \delta) \otimes \beta_n^{(\alpha_3)}(k_3 + \delta), \quad (4)$$

for all $\mathbf{k} = (k_1, k_2, k_3) \in \mathbb{N}^3$, and where $\beta_n^{(m)}(x)$ denotes the m th order derivative of a n th 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.




Fig. 1. Second degree discrete directional derivative operators given by the B-spline approximation (4) at different orientations. Note that the operators are approximately rotation steerable.

4. 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. 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.

4.1. Compressed Sensing MRI Recovery With 3-D HDTV

We consider the recovery of a single coil 3-D MR angiography dataset from noisy and undersampled measurements. We experiment on a $512 \times 512 \times 76$ MR dataset obtained from [6], 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- and ridge-like details compared with standard 3D-TV. Additionally, in Table 1 we provide quantitative comparisons between HDTV and TV with more MR datasets. Note that both the HDTV2 and HDTV3 penalties routinely outperform TV in SNR.

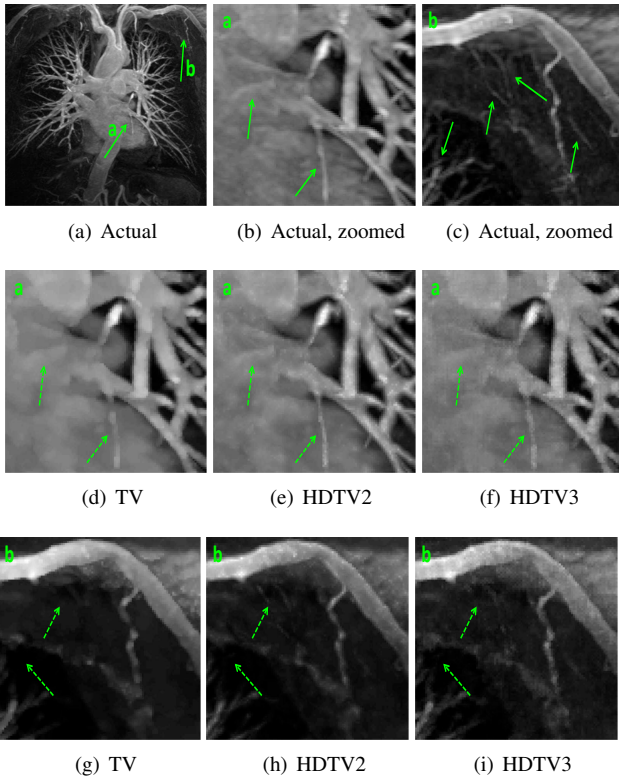


Fig. 2. Compressed sensing recovery of a 3-D MR angiography dataset. (a)-(c) show the MIP of the original dataset, (d)-(i) show two zoomed regions of the original dataset and their reconstructions obtained using TV and HDTV.




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.

4.2. Comparison of HS1 and HDTV Penalties

In Table 1 we compare the performance of HDTV and HS1 penalties in denoising and deblurring 2-D images. Note that all higher degree schemes outperform TV. We also observe the SNR values are approximately the same in the case of HDTV2 and HS1. Fig. 3 shows that the reconstructions are visually similar, as well. These experiments confirm that the discrete versions of these two penalties perform similarly in image recovery tasks, as predicted by Proposition 1.

	3-D Compressed Sensing			2-D Denoising		2-D Deblurring	
	mra5	mra1.5	cardiac	brain	lena	cell1	cell2
TV	13.87	14.53	18.37	27.60	27.35	15.66	16.67
HDTV2	14.23	15.11	18.56	28.05	27.65	16.19	17.21
HDTV3	14.01	14.70	18.50	28.30	27.45	16.17	17.20
HS1	—	—	—	28.08	27.51	16.17	17.13

Table 1: SNRs (in dB) of recovered test images under different regularization.

5. CONCLUSIONS

We extend the HDTV image regularization penalties to 3-D signals. The extension to 3-D is mainly enabled by the speedup offered by a recently introduced fast alternating minimization algorithm for 2-D HDTV, which we adapt here to the 3-D case. Our experiments demonstrate the improvement in image quality offered by the proposed scheme over TV for 3-D compressed sensing MR image reconstruction. We also show that second degree HDTV penalty and the $p = 1$ Hessian Schatten-norm closely approximate each other and perform similarly in image recovery settings.

6. 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 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.
- [4] S Lefkimmiatis, J P 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.
- [5] 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.
- [6] “<http://physionet.org/physiobank/database/images/>,”.