A Splitting Algorithm for Directional Regularization and Sparsification

Lars Lau Rakêt and Mads Nielsen
Department of Computer Science, University of Copenhagen
{larslau, madsn}@diku.dk

Abstract

We present a new split-type algorithm for the minimization of a p-harmonic energy with added data fidelity term. The half-quadratic splitting reduces the original problem to two straightforward problems, that can be minimized efficiently. The minimizers to the two sub-problems can typically be computed pointwise and are easily implemented on massively parallel processors. Furthermore the splitting method allows for the computation of solutions to a large number of more advanced directional regularization problems. In particular we are able to handle robust, non-convex data terms, and to define a 0-harmonic regularization energy where we sparsify directions by means of an L^0 norm.

1 Introduction

This paper is concerned with directional regularization of vector valued functional data, i.e. regularization where the angles between vectors are regularized, but magnitude remain unaffected. This type of regularization has many uses, in particular for data where magnitude and direction are corrupted in different ways. Examples are e.g. chromaticity denoising of images [8], inpainting [6], or regularization in optical flow estimation [1, 4].

We start out from the p-harmonic regularization of Vese and Osher [8], review the possibilities for adding a data term, and propose an effective splitting algorithm for minimizing the resulting energy functional. By modifying the sub-energies in the algorithm, we are able to define and solve several new types of directional regularization problems.

The p-harmonic energy is defined as

$$H_p(u) = \int \left\| \nabla \frac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} \right\|^p \mathrm{d}\boldsymbol{x} \tag{1}$$

where $u : \mathbb{R}^d \to \mathbb{R}^k$, $d, k \geq 2$ and $p \geq 1$. Vese and Osher derived the corresponding Euler-Lagrange equations and introduced a semi-implicit scheme for computing the steady state. The method is appealing in the sense that it is an unconstrained minimization problem, but still it is only the directions of the initial state that are regularized and data magnitudes stay intact. On the down side, this solution strategy can be quite slow, and does require some amount of hand tuning in order to obtain good results. To address these issues, we will add a data term to the p-harmonic energy, and derive a new algorithm with appealing computational properties for minimizing the resulting functional. The derived algorithm is generalized to handle robust data terms. Furthermore we extend the definition of p-harmonic to include $p = 0$, i.e. where we penalize the gradient using a sparsity imposing L^0 norm, and show how to generalize to anisotropic regularization.

The rest of the paper is organized as follows: In the next section we will review the problem of adding an L^2 data term to H_p. In Section 3 we derive the splitting algorithm for minimizing (1). Section 4 considers generalization of the derived algorithm, and finally we conclude and propose future directions in the last section.

2 Data terms for directional data

The energy (1) can either be used to compute a steady state of the directions in the data, or as showed in [8], one can merely diffuse the directions according to the energy as a denoising method. This will however typically require hand tuning of the stopping time, in order to obtain good results. To avoid this we propose to add a data term, which will give more control over the smoothness properties of the solutions. A standard choice is a quadratic term [4, 6], which, given data f, leads to the following energy

$$\lambda \int \|u(\boldsymbol{x}) - f(\boldsymbol{x})\|^2 \,\mathrm{d}\boldsymbol{x} + H_p(u),$$

where the parameter $\lambda > 0$ controls the trade-off between data fidelity and regularity. However (contrary to what was stated by Gai and Stevenson [4]) this data term will affect the magnitudes of the minimizer. To see this, we assume that $k = 2$ and use the polar representation $u = (r\cos\theta, r\sin\theta)$ and $f = (s\cos\vartheta, s\sin\vartheta)$, where the parameters are space dependent. Since the magnitude is only present in the data term, one can estimate r pointwise, and from the Euler-Lagrange equations one can calculate the solution

$$r = s\cos(\theta - \vartheta).$$

Different approaches can be taken in order to eliminate this problem. Here we take the simple approach of discarding information about magnitude completely from the data term, and minimize the energy

$$E(u) = \lambda \int \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - \tfrac{f(\boldsymbol{x})}{\|f(\boldsymbol{x})\|} \right\|^2 \mathrm{d}\boldsymbol{x} + H_p(u), \quad (2)$$

under the constraint that $\|u(\boldsymbol{x})\| = \|f(\boldsymbol{x})\|$, which can of course easily be included in the energy by means of an extra data term.

3 Algorithm

Instead of directly minimizing (2) we propose to do a half-quadratic splitting of the energy in two coupled energies. These two sub-problems are then solved interleaved as the split energy is made converge to (2). In recent years, this technique has been used in a number of interesting applications, e.g. optical flow estimation [12, 10, 7], image denoising [9], and image processing [11]. We propose to split the energy as follows

$$\begin{aligned} E(u,v) = \lambda \int & \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - \tfrac{f(\boldsymbol{x})}{\|f(\boldsymbol{x})\|} \right\|^2 \mathrm{d}\boldsymbol{x} \\ + \beta \int & \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - v(\boldsymbol{x}) \right\|^2 \mathrm{d}\boldsymbol{x} \\ + \int & \|\nabla v(\boldsymbol{x})\|^p \, \mathrm{d}\boldsymbol{x} \end{aligned} \quad (3)$$

and obtain a minimizer of (2) by iteratively minimizing it in u and v as $\beta \to \infty$.

Consider first the minimization of (3) in u,

$$\begin{aligned} E_1(u) = \lambda \int & \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - \tfrac{f(\boldsymbol{x})}{\|f(\boldsymbol{x})\|} \right\|^2 \mathrm{d}\boldsymbol{x} \\ + \beta \int & \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - v(\boldsymbol{x}) \right\|^2 \mathrm{d}\boldsymbol{x} \end{aligned} \quad (4)$$

under the constraint that $\|u(\boldsymbol{x})\|$ should equal $\|f(\boldsymbol{x})\|$. Since no differential of u is involved the minimization can be done pointwise, and the solution can be shown to be

$$u(\boldsymbol{x}) = \|f(\boldsymbol{x})\| \, \pi \left(\lambda \frac{f(\boldsymbol{x})}{\|f(\boldsymbol{x})\|} + \beta v(\boldsymbol{x}) \right),$$

where π is the projection onto the k-dimensional unit ball.

The minimization problem in v gives the energy

$$\begin{aligned} E_2(v) = \beta \int & \left\| \tfrac{u(\boldsymbol{x})}{\|u(\boldsymbol{x})\|} - v(\boldsymbol{x}) \right\|^2 \mathrm{d}\boldsymbol{x} \\ + \int & \|\nabla v(\boldsymbol{x})\|^p \, \mathrm{d}\boldsymbol{x}, \end{aligned} \quad (5)$$

which can be solved by some of the many known methods for different values of p. The sequence of β's in the iterative minimization is chosen according to [11].

The case $p = 1$ is perhaps the most interesting as this corresponds to total variation (TV) regularization. In this case, depending on the interpretation of the vectorial gradient, E_2 can be minimized by the efficient dual methods of Bresson and Chan [2], or one can interpret the last term as the vectorial total variation of Goldluecke et al. [5], and use one of their methods. Here we will use the definition of Goldluecke et al., as this definition of total variation does not suffer from any channel smearing.

The algorithm has been implemented in CUDA C in order to take advantage of the thousand of cores on modern GPUs. The minimization procedure with $p = 1$ can be done in less than a second for a 640×480 color image on an NVIDIA Tesla C2050 GPU.

An example of chromaticity denoising using 1-harmonic regularization can be found in Figure 1.

4 Generalizations

The presented splitting algorithm can easily be modified to solve other directional problems. One can for example replace the regularization term in E_2 by a spatially weighted term or by an anisotropic term similar to the one of Werlberer et al. [10] to include e.g. magnitude information to guide the regularization directions. Here we will consider two generalizations in detail. First, the possibility of using more robust data terms, and secondly directional sparsification by means of 0-harmonic regularization.

4.1 Robust data terms

If data contains outliers or is severely degraded, an L^2 data term will often produce problematic results. In order to mitigate the effects of outliers, one can replace the quadratic norm by a robust data term. A popular choice is the L^1 norm, which combined with unconstrained TV regularization has some nice properties, e.g. contrast invariance [3].

Due to the projected terms $\frac{u}{\|u\|}$ it may however be very hard to find an explicit solution to (4) if the

Figure 1: 1-harmonic regularization of an image with noisy chromaticity. Left: Original image, middle: Noisy chromaticity (MSE 663.5), right: Result after regularization (MSE 32.7), $\lambda = 1.2$. Zoom-ins show original image and pure chromaticity, where color vector magnitude has been removed.

quadratic norm is replaced by a robust term ρ. However, as in the L^2 case the modified energy can be minimized pointwise, giving a minimization problem of the form

$$e_1(\boldsymbol{u}) = \lambda\, \rho\left(\frac{\boldsymbol{u}}{\|\boldsymbol{u}\|} - \frac{\boldsymbol{f}}{\|\boldsymbol{f}\|}\right) + \beta \left\|\frac{\boldsymbol{u}}{\|\boldsymbol{u}\|} - \boldsymbol{v}\right\|^2. \quad (6)$$

e_1 is a function of $\boldsymbol{u}$ only through its defining angles $\theta_i \in [0, 2\pi)$, $i = 1, \ldots, k-1$, and one can resolve to minimizing e_1 by a complete search in the space of angles $[0, 2\pi)^{k-1}$. For $k = 2$ the additional time consumption of the regularization procedure is less than a second, and for $k = 3$ a brute force search in angle space is still feasible with an additional computation time of approximately 10 minutes, in the setup from in the previous section.

We have considered two robust measures of data fit, namely an L^1 data term and a Cauchy/Lorentzian data term

$$\rho = \|\cdot\|, \qquad \rho = \log(1 + \|\cdot\|^2).$$

Figure 2 shows an image where the color data has been severely degraded (top row contains the image data, bottom row only the directional component). We see that the results produced by using the robust data terms better handle desaturation and have better vibrancy than the result produced by the L^2 term.

4.2 0-harmonic Regularization

1-harmonic regularization will evoke some sparsity in the regularized directions, due to the well known staircasing effect of TV regularization. But in some contexts we want to go even further than total variation can bring us. For this we can use an L^0 regularization of the gradient.

Consider the norm $\|x\|^p$, $x \in \mathbb{R}$. As $p \to 0$ $\|x\|^p \to 1$ for $x \neq 0$, and weighting the importance of continuity of the empty sum over continuity of the empty product, we can choose $\|0\|^0 = 0$, which defines the zero norm. Formally, the vectorial zero norm is defined as

$$\|u\|_{\ell_0} = |\{\boldsymbol{x} : \|u(\boldsymbol{x})\|_{\ell_1} \neq 0\}|.$$

Clearly, regularizing the gradient of an image with this norm will create a very sparse, cartoon like result with only few edges. Xu et al. [11] recently presented an efficient algorithm for regularization using the zero norm, i.e. for minimizing energies of the form

$$\beta \int \|v(\boldsymbol{x}) - g(\boldsymbol{x})\|^2 \,\mathrm{d}\boldsymbol{x} + \int \|\nabla v(\boldsymbol{x})\|_{\ell_0} \,\mathrm{d}\mu(\boldsymbol{x}),$$

where μ denotes the Hausdorff measure. By replacing E_2 with this type of energy we can regularize our observations such that magnitudes remain unaffected while the directions are sparsified. While an L^0 regularization of the gradient is not suitable for data where noise is significant, it has a number of other interesting uses. One example is data where one knows directions to be sparse, e.g. a fixed camera filming cars going by. Another possible applications is image compression, where sparse color directions may be very valuable, since the human visual system is quite tolerant to color errors, and because the prominent color features of the image is preserved very well using this form of regularization. This means that images with sparse color directions may be hard to distinguish from the original, but can be encoded much more efficiently. An example of 0-harmonic regularization can be found in Figure 3. Close inspection reveals that the curtain in the right side of the image seems somewhat more flat in the

(a) Original image (b) Degraded image (c) L^2 data term $\lambda = 0.5$ (d) L^1 data term $\lambda = 0.5$ (e) Cauchy data term $\lambda = 0.6$

Figure 2: Comparison of different data terms with 1-harmonic regularization of an image with noisy chromaticity (Gaussian noise with standard deviation of 0.2 and 40% impulse noise followed by renormalization of the directions). Parameters λ are chosen as the highest value for which no strong artifacts are noticeable in the reconstructed image. Mean squared errors: Degraded image 5322, L^2 data term 1285, L^1 data term 876, Cauchy data term 1145.

regularized image, but otherwise they are hard to distinguish. We note that the sparse color directions in Figure 3 (without magnitude information) can be encoded using only 11 % of the space of the original in PNG ISO/IEC 15948:2004 encoding with standard settings, and using a specialized encoding, this number may be further decreased.

5 Conclusions and Outlook

We have considered the problem of directional regularization, in particular in form of the p-harmonic regularization of Vese and Osher [8]. We have introduced new magnitude-preserving data fidelity terms that allows for better control of the regularization, and we have derived an efficient parallel algorithm for minimizing the resulting energy, and demonstrated a GPU implementation. In addition we have considered a novel generalizations of the proposed algorithm, and have outlined a general scheme of generalization. We have briefly touched the question of how to generalize to anisotropic regularization of angles, and have in detail considered how to use more advanced data terms, and how to extend p-harmonic regularization to the case $p = 0$ which gives an interesting method for sparsifying directions, which in addition to reconstructing sparse structures, may be used for compression purposes.

The proposed algorithm works for arbitrary domain and response dimensions. Here we have only considered the case of color images ($d = 2$, $k = 3$), but interesting applications may be found in a large number of other types of data. We have briefly mentioned optical flow ($d = k = 2$), and for three dimensional domains, many interesting applications are available in e.g. medical image analysis and video coding.

An interesting point of future research is how to couple angle and magnitude components for anisotropic regularization of different types of data. Moreover the use of 0-harmonic regularization for image compression is worth looking into, as preliminary results suggest that one can achieve very high compression of the chromaticity, at little or no cost in the perceived quality of the final image.

Acknowledgements

Photo in Figure 2 by Alfonsina Blyde[1].

[1] http://www.flickr.com/photos/alfon18/2476621123/

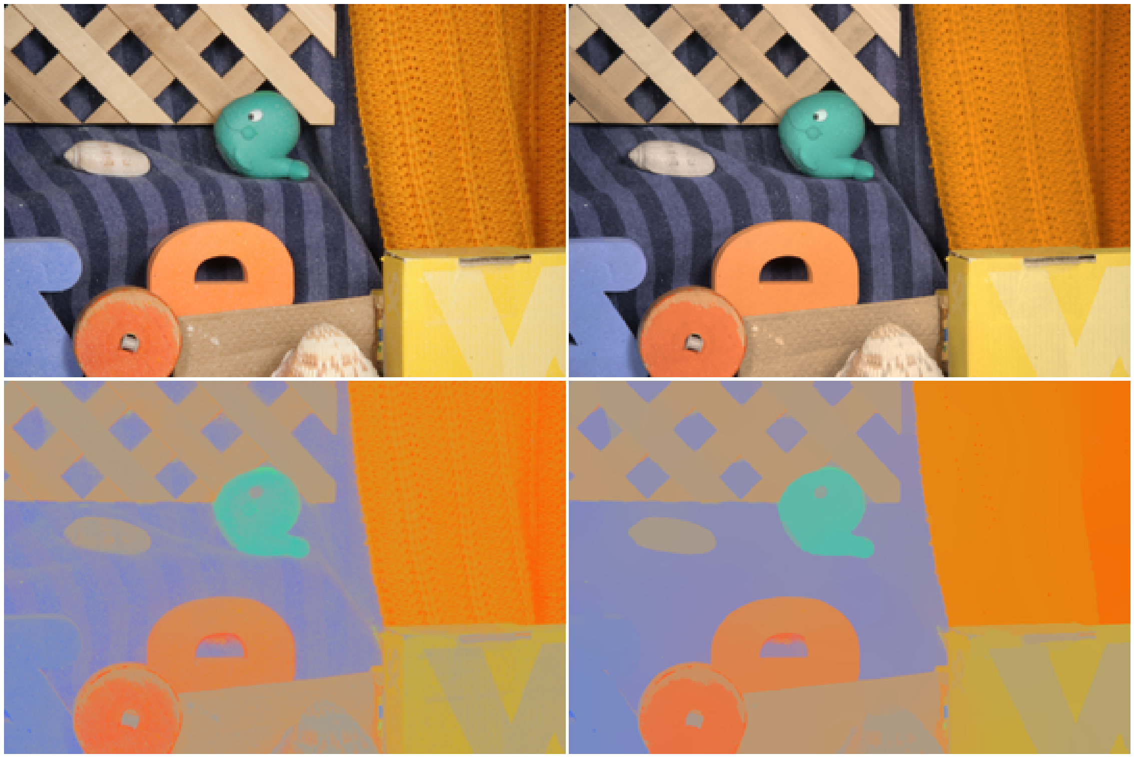

(a) Original image

(b) With sparse RGB directions

Figure 3: 0-harmonic regularization (MSE 16.6), original image and chromaticity, $\lambda = 5$

References

[1] Y. Adato, T. Zickler, and O. Ben-Shahar. A polar representation of motion and implications for optical flow. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1145–1152, 2011.

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

[3] J. Darbon and M. Sigelle. Image restoration with discrete constrained total variation part i: Fast and exact optimization. *Journal of Mathematical Imaging and Vision*, 26:261–276, 2006.

[4] J. Gai and R. L. Stevenson. Optical flow estimation with p-harmonic regularization. In *Proceedings of the International Conference on Image Processing, ICIP 2010, September 26-29, Hong Kong, China*, pages 1969–1972. IEEE, 2010.

[5] B. Goldluecke, E. Strekalovskiy, and D. Cremers. The natural total variation which arises from geometric measure theory. *SIAM Journal on Imaging Sciences*, 2012. to appear.

[6] J. Hahn, X.-C. Tai, S. Borok, and A. M. Bruckstein. Orientation-matching minimization for image denoising and inpainting. *International Journal of Computer Vision*, 92(3):308–324, 2011.

[7] L. L. Rakêt, L. Roholm, M. Nielsen, and F. Lauze. TV-L^1 optical flow for vector valued images. In Y. Boykov, F. Kahl, V. Lempitsky, and F. Schmidt, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 6819 of *Lecture Notes in Computer Science*, pages 329–343. Springer, 2011.

[8] L. A. Vese and S. J. Osher. Numerical methods for p-harmonic flows and applications to image processing. *SIAM Journal on Numerical Analysis*, 40(6):2085–2104, 2002.

[9] Y. Wang, J. Yang, W. Yin, and Y. Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*, 1(3):248–272, Aug. 2008.

[10] M. Werlberger, W. Trobin, T. Pock, A. Wedel, D. Cremers, and H. Bischof. Anisotropic huber-l1 optical flow. In *Proceedings of the British Machine Vision Conference (BMVC)*, London, UK, September 2009. to appear.

[11] L. Xu, C. Lu, Y. Xu, and J. Jia. Image smoothing via L_0 gradient minimization. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 30(4), 2011.

[12] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime TV-L^1 optical flow. In F. Hamprecht, C. Schnörr, and B. Jähne, editors, *Pattern Recognition*, volume 4713 of *Lecture Notes in Computer Science*, pages 214–223. Springer, 2007.