

Polyakov action minimization for efficient color image processing

Guy Rosman ^{1*}, Xue-Cheng Tai ^{2,3**}, Lorina Dascal ¹, and Ron Kimmel ¹

¹ Dept. of Computer Science, Technion, Haifa, Israel,
{rosman,lorina,ron}@cs.technion.ac.il

² Dept. of Mathematics, Bergen University, Bergen, Norway, tai@math.uib.no

³ Division of Mathematical Sciences, School of Physical and Mathematical Sciences,
Nanyang Technological University, Singapore

Abstract. The Laplace-Beltrami operator is an extension of the Laplacian from flat domains to curved manifolds. It was proven to be useful for color image processing as it models a meaningful coupling between the color channels. This coupling is naturally expressed in the Beltrami framework in which a color image is regarded as a two dimensional manifold embedded in a hybrid, five-dimensional, spatial-chromatic (x, y, R, G, B) space.

The Beltrami filter defined by this framework minimizes the Polyakov action, adopted from high-energy physics, which measures the area of the image manifold. Minimization is usually obtained through a geometric heat equation defined by the Laplace-Beltrami operator. Though efficient simplifications such as the bilateral filter have been proposed for the single channel case, so far, the coupling between the color channel posed a non-trivial obstacle when designing fast Beltrami filters.

Here, we propose to use an augmented Lagrangian approach to design an efficient and accurate regularization framework for color image processing by minimizing the Polyakov action. We extend the augmented Lagrangian framework for total variation (TV) image denoising to the more general Polyakov action case for color images, and apply the proposed framework to denoise and deblur color images.

Key words: Laplace-Beltrami, diffusion, optimization, denoising, PDEs

1 Introduction

Variational, nonlinear diffusion filters have been extensively used in the last two decades for different image processing tasks. Numerical schemes implementing them are designed with an emphasis on accuracy, stability and computational efficiency. While a many of the works on image regularization involve greyscale

* This research was partly supported by ISF grant no. 1551/09 and by Rubin Scientific and Medical Fund.

** The research is also supported by MOE (Ministry of Education) Tier II project T207N2202 and IDM project NRF2007IDMIDM002-010.

images, only a small portion has made a coherent attempt at regularizing vector-valued signals.

These works, several of which are inspired by [39], describe regularization functionals which operate on vector-valued images. These include the works of Sapiro and Ringach [27], Blomgren and Chan [4], and Sochen et. al. [30], as well as more recent works such as ([34, 14]). The Beltrami framework [30] describes a regularizing functional, well suited for color image processing, which can be justified by the Lambertian model of color image formation. The framework considers the image as a 2-manifold embedded in a hybrid spatial-feature space. Regularization of the image in this framework is expressed as minimization of area surface. The Beltrami filter is strongly related to the bilateral filter (see [37], [28], [33], [29], [13], [3]), as well as to the nonlocal means filter, proposed in [1]. Minimization of the associated functional is usually done by evolving the image according to its Euler-Lagrange equation [30]. This evolution, using an explicit scheme, is limited in its time step, resulting in high computational complexity. Another possibility [2] is to perform a fixed-point iteration from the Euler-Lagrange equation. Recently, several approaches were suggested for improving the speed of computation of minimizers for the *Polyakov action* [22]. Those include an approximation of the Beltrami filter kernel [31], as well as employing vector extrapolation techniques [24], or operator splitting methods [11]. For the case of gray-scale images, the projection-based method [8] has been extended to the Polyakov function [5], with no suggestions made for vector-valued images.

In [32], the augmented Lagrangian method [15, 23] is used to perform TV regularization of images. In this paper we propose to use a similar constrained optimization for regularization of color images. Instead of discretizing the continuous optimality condition or the resulting Beltrami flow, we minimize the discretized Polyakov action itself. The resulting method is shown to be more efficient and accurate for image denoising and deblurring, compared to existing methods for Beltrami regularization in image processing. In Section 2 we review the Beltrami framework for color image regularization. In Section 3 we extend the coupled constrained optimization approach demonstrated in [32] to regularize color images by the Polyakov action. In Section 4 we display results of using our method for deblurring color images. Section 5 concludes the paper.

2 The Beltrami Framework

We now briefly review the Beltrami framework for non-linear diffusion in computer vision [18, 30, 38]. The basic notions used in this introduction are taken from Riemannian geometry, and we refer the reader to [12] for further reading.

In the Beltrami framework, images are expressed as maps between two Riemannian manifolds. Denote such a map by $X : \Sigma \rightarrow M$, where Σ is a two-dimensional manifold, parameterized by global coordinates (σ^1, σ^2) , and M is the spatial-feature manifold, embedded in \mathbb{R}^{d+2} , where d is the number of image channels. For example, a gray-level image can be represented as a surface

embedded in \mathbb{R}^3 . The map X in this case is $X(\sigma^1, \sigma^2) = (\sigma^1, \sigma^2, I(\sigma^1, \sigma^2))$, where I is the image intensity. For color images, X is given by $X(\sigma^1, \sigma^2) = (\sigma^1, \sigma^2, I^1(\sigma^1, \sigma^2), I^2(\sigma^1, \sigma^2), I^3(\sigma^1, \sigma^2))$, where I^1, I^2, I^3 are the color components (for example, red, green, blue for the RGB color space).

Next, we choose a Riemannian metric on this surface. Its components are denoted by g_{ij} . The canonical choice of coordinates in image processing uses Cartesian coordinates $\sigma^1, \sigma^2 = x, y$. We denote the elements of the inverse of the metric by superscripts g^{ij} , and the determinant by $g = \det(g_{ij})$.

Once images are defined as Riemannian embeddings, we can look for a measure on this space of embedding maps. Denote by (Σ, g) the image manifold and its metric, and by (M, h) the space-feature manifold and its metric. The functional $S[X]$ characterizes the mapping $X : \Sigma \rightarrow M$, and is defined to be

$$S[X, g_{ij}, h_{ab}] = \int d^m \sigma \sqrt{g} \|dX\|_{g,h}^2, \quad (1)$$

where m is the dimension of Σ , g is the determinant of the image metric, and the range of indices is $i, j = 1, 2, \dots, \dim(\Sigma)$ and $a, b = 1, 2, \dots, \dim(M)$. The integrand $\|dX\|_{g,h}^2$ is given by $\|dX\|_{g,h}^2 = (\partial_{x_i} I^a) g^{ij} (\partial_{x_j} I^b) h_{ab}$. We use here Einstein's summation convention: identical indices that appear up and down are summed over. This functional, for $\dim(\Sigma) = 2$ and $h_{ab} = \delta_{ab}$, is known in string theory as the Polyakov action [22], and extends the action functional from classical mechanics to the relativistic case.

In the case of color images, where both the spatial and the color spaces are assumed to be Cartesian, the metric becomes

$$(g_{ij}) = \begin{pmatrix} 1 + \beta^2 \sum_{a=1}^3 (I_{x_1}^a)^2 & \beta^2 \sum_{a=1}^3 I_{x_1}^a I_{x_2}^a \\ \beta^2 \sum_{a=1}^3 I_{x_1}^a I_{x_2}^a & 1 + \beta^2 \sum_{a=1}^3 (I_{x_2}^a)^2 \end{pmatrix} = G,$$

where a subscript of I^a denotes a partial derivative and the parameter $\beta > 0$ determines the ratio between the spatial and color coordinates. The functional becomes

$$S(X) = \int \sqrt{g} d\sigma^1 d\sigma^2, \quad g = \det(G) = 1 + \beta^2 \sum_{a=1}^3 \|\nabla I^a\|^2 + \frac{\beta^4}{2} \sum_{a,b=1}^3 \|\nabla I^a \times \nabla I^b\|^2,$$

The role of the cross product term $\sum_{a,b=1}^3 \|\nabla I^a \times \nabla I^b\|^2$ in the regularization was explored in [18],[17]. It penalizes deviations from the Lambertian model of image formation [18], or specifically – misalignment of the gradient directions between color channels.

The functional S is usually minimized by time evolution of the image according to the Euler-Lagrange equations,

$$I_t^a = -\frac{1}{\sqrt{g}} h^{ab} \frac{\delta S}{\delta I^b} = \frac{1}{\sqrt{g}} \underbrace{\text{div}(D\nabla I^a)}_{\Delta_g I^a}, \quad (2)$$

where the matrix $D = \sqrt{g}G^{-1}$. See [30] for explicit derivation. The operator Δ_g generalizes the Laplacian to manifolds, and is called the Laplace-Beltrami operator. Evolution according to these equations result in the Beltrami scale-space.

The functional can also be generalized to the family of functionals

$$\int \sqrt{\beta_1 + \beta_2 \sum_{a=1}^3 \|\nabla I^a\|^2 + \beta_3 \sum_{a,b=1}^3 \|\nabla I^a \times \nabla I^b\|^2}, \quad (3)$$

for any positive $\beta_1, \beta_2, \beta_3$. While this approach cannot be explained by the minimal area interpretation, it makes sense in terms of color image restoration, and will be used in the results shown in Figure 4.

In the variational framework, the reconstructed image minimizes a cost-functional of the form

$$\Psi = \frac{\alpha}{2} \sum_{a=1}^3 \|KI^a - I_0^a\|^2 + S(X),$$

where K is a bounded linear operator. In the denoising case, K is the identity operator $Ku = u$, and in the deblurring case, $Ku = k * u$, where $k(x, y)$ is the blurring kernel. The parameter α controls the smoothness of the solution. This functional has been used for image denoising [30, 2] and blind deconvolution [16], and its relation to active contours explored in [6]. We introduce an approach for optimizing the functional Ψ using the augmented Lagrangian method.

3 An augmented Lagrangian approach for Beltrami regularization

In recent years, several attempts have been made of optimizing total variation functionals [26] using dual variables (we refer the reader to [9, 7, 8, 20, 35, 32, 36] and references therein, as well to more references found in the technical report [25]). These algorithms achieved great accuracy and efficiency, and are considered to be among state-of-the-art methods for TV restoration.

Specifically, in [32], total variation regularization is obtained by decoupling the optimization problem

$$\min_u |\nabla u| + \frac{\alpha}{2} \|Ku - f\|^2 \quad (4)$$

into a constrained optimization problem

$$\min_{u, \mathbf{q}} |\mathbf{q}| + \frac{\alpha}{2} \|Ku - f\|^2 \quad s.t. \quad \mathbf{q} = \nabla u, \quad (5)$$

where \mathbf{q} is a auxiliary field, parallel to the gradient of u . This constraint is then incorporated using an augmented Lagrangian penalty function of the form

$\rho_{\mu,r}(u, \mathbf{q}) = \boldsymbol{\mu}^T(\nabla u - \mathbf{q}) + \frac{r}{2}(\|\nabla u - \mathbf{q}\|^2)$. The penalty is used to enforce the constraint $\mathbf{q} = \nabla u$, without making the problem severely ill-conditioned.

We now describe a similar construction for the Polyakov action. Again, it is important to stress we are minimizing the functional itself, rather than discretizing the resulting minimizing PDE as in [30, 16, 31, 2, 24, 11].

We deal with the case of color images, for which the regularization offered by the Beltrami framework is more meaningful. Specifically, we replace the gradient norm penalty used in TV regularization by the action functional of Equation 1. This is done by replacing the first term in Equation 5 by the term

$$\int \sqrt{1 + \beta^2 \sum_{i \in \{R,G,B\}} \|\mathbf{q}_i\|^2 + \frac{\beta^4}{2} \sum_{i \in \{R,G,B\}} \sum_{j \neq i} \|\mathbf{q}_i \times \mathbf{q}_j\|^2}, \quad (6)$$

where β is the spatial-intensity aspect ratio, and $\{\mathbf{q}_i\}_{i \in \{R,G,B\}}$ denote components of the field \mathbf{q} , parallel to the gradient of each of the image channels. We then trivially extend the rest of the functional to the vectorial (per-pixel) case, obtaining the following functional

$$\mathcal{L}_{BEL}(u, \mathbf{q}, \boldsymbol{\mu}) =$$

$$\int \left\{ \sqrt{1 + \beta^2 \sum_{i \in \{R,G,B\}} \|\mathbf{q}_i\|^2 + \frac{\beta^4}{2} \sum_{i \in \{R,G,B\}} \sum_{j \neq i} \|\mathbf{q}_i \times \mathbf{q}_j\|^2} + \sum_{i \in \{R,G,B\}} \boldsymbol{\mu}_i^T(\mathbf{q}_i - \nabla u_i) + \frac{\alpha}{2} \|Ku - f\|^2 + \frac{r}{2} \sum_{i \in \{R,G,B\}} \|\mathbf{q}_i - \nabla u_i\|^2 \right\},$$

which corresponds to Beltrami regularization. The expressions optimizing u and $\boldsymbol{\mu}$ are replaced by their per-channel equivalents, $\{u_i\}$ and $\{\boldsymbol{\mu}_i\}$, for $i \in \{R, G, B\}$. The augmented Lagrangian algorithm for regularizing an image using the Polyakov action is given as Algorithm 1. At each inner iteration k ,

Algorithm 1 Augmented Lagrangian optimization of the Beltrami framework

- 1: $\boldsymbol{\mu}^0 \leftarrow 0$
- 2: **for** $k=0,1,\dots$ **do**
- 3: Update $\{u_i\}^k, \{\mathbf{q}_i\}^k$:

$$(\{u_i\}^k, \{\mathbf{q}_i\}^k) = \operatorname{argmin}_{\{u_i\}, \{\mathbf{q}_i\}} \mathcal{L}_{BEL}(\{u_i\}, \{\mathbf{q}_i\}, \{\boldsymbol{\mu}_i^k\}) \quad (7)$$

 according to Equation 8 and Subsection 3.1.

- 4: Update the Lagrange multipliers according to Equation 9
 - 5: **end for**
-

$\{u_i\}_{i \in \{R,G,B\}}$ is updated in the Fourier domain, as in [32],

$$u_i^k = \mathcal{F}^{-1} \left\{ \frac{\alpha \mathcal{F}\{K^*\} \mathcal{F}\{f_i\} - \mathcal{F}\{D_x^-\}((\mu_i^1)^k + r(p_i)^k) - \mathcal{F}\{D_y^-\}((\mu_i^2)^k + r(q_i)^k)}{\alpha \mathcal{F}\{K^*\} \mathcal{F}\{K\} - r \mathcal{F}\{\Delta\}} \right\}, \quad (8)$$

where D_x^-, D_y^-, Δ denote the backward derivative along the x and y directions, and the Laplacian operator, respectively, and $\mathcal{F}\{\cdot\}, \mathcal{F}^{-1}\{\cdot\}$ denote the Fourier transform and its inverse, respectively. We explicitly write $\mathbf{q}_i = (p_i, q_i), i \in \{R, G, B\}$, for the components of \mathbf{q} of each color channel, approximating its x and y derivatives, computed using backward differences.

We note that the optimization of u using the Fourier domain resembles, in a sense, the approach taken by [21]. Since, however, it is done with respect to the auxiliary field, iteratively, its effect is suited to the nonlinear nature of the Beltrami flow. An update rule for the auxiliary field \mathbf{q}_i of each channel is described in Subsection 3.1.

According to the augmented Lagrangian method, the Lagrange multipliers μ_i are updated so as to approximate the optimal Lagrange multipliers,

$$(\mu_i)^k = (\mu_i)^{k-1} + r ((\mathbf{q}_i)^k - (\nabla u_i)^k). \quad (9)$$

Finally, the coefficient r is updated between each outer iteration by multiplying r with a scalar $\gamma > 1$. We note r needs not be very large, thus avoiding ill-conditioning of the functional $\mathcal{L}_{BEL}(u, \mathbf{q}, \mu)$.

3.1 Updating the auxiliary field \mathbf{q}

For optimizing \mathbf{q} , a short inner-loop of a fixed-point solver with *iterative reweighted least squares* (IRLS) allows us to efficiently obtain a solution. In numerical experiments, optimization over \mathbf{q} takes less than half the CPU time of the algorithm. Furthermore, since this problem is solved per pixel, it can be easily parallelized, for example on a GPUs.

The update of $\mathbf{q}_i = (p_i, q_i), i \in \{R, G, B\}$, the components of \mathbf{q} at each pixel, is done by optimizing the function

$$\begin{aligned} & \sqrt{1 + \beta^2 \sum_i (p_i^2 + q_i^2) + \frac{\beta^4}{2} \sum_i \sum_{j \neq i} (p_i q_j - q_i p_j)^2} \\ & + \frac{r}{2} \sum_i \|\mathbf{q}_i - (\nabla u_i)\|^2 + \sum_{i \in \{R, G, B\}} (\mu_i^k)^T (\mathbf{q}_i - \nabla u_i), \end{aligned}$$

where $(\nabla u)_i = ((u_i)_x, (u_i)_y)^T$ denote the components of the various image channel gradients. Details of the update equations are given in the technical report [25].

4 Results

We now demonstrate the minimization of the Polyakov functional using the augmented Lagrangian method, for various applications. More examples are shown in the Technical report [25].

Scale-space, smoothing and denoising: In Figure 1, results are shown for smoothing an image using various values of α , which in a sense parallel samples along the Beltrami scale-space (defined by the flow). We used the same initial penalty parameter $r = 0.5$, for which the constraints were satisfied after very few iterations. Fixed-point iterations over \mathbf{q} were limited to 2 inner and 2 outer (IRLS) iterations for each cycle. The number of outer iterations, updating $\boldsymbol{\mu}$, in Figure 1 was 150, although fewer iterations suffice.

A comparison of the results of the augmented Lagrangian method and splitting schemes [11] shows that the augmented Lagrangian method converges faster, as can be seen in Figure 2. In this experiment, α was set for optimal results for both the augmented Lagrangian and the splitting methods. The PSNR plot also demonstrates the more accurate discretization of the proposed method. This can be easily seen in the preservation of edges in Figure 2. Experiments comparing our method to the explicit scheme showed a similar behavior.

Table 1 measures the CPU-time required for several images (shown in Figure 3) for our algorithm, compared to Beltrami filtering with operator splitting techniques. The time step used was the largest possible so as to avoid instabilities and inaccurate operator approximation.

Since the solution obtained by discretizing the functional and by discretizing the resulting Euler-Lagrange equation need not be the same, a different halting condition was used. After measuring the PSNR of each algorithm with respect to the original image, we measured the CPU time each algorithm took to gain 99% of the maximal rise in SNR. While this cannot be done in real applications, it does give us an objective measure of the time it takes to complete the convergence.

The speedups obtained are by at least of a factor of two compared to additive operator splitting (AOS) [19], which is one of the fastest methods for Beltrami regularization [11], even when the time step large enough to cause visible artifacts in the splitting results. The augmented Lagrangian method clearly gave still more accurate results in a shorter CPU time.

Deblurring: Deblurring results using the Beltrami framework are shown in Figure 4, with the blur kernel k a disc of radius 5 pixels. We compare our results to standard deblurring algorithms available in Matlab, as well as to BM3D deblurring [10], and to the FTVd algorithm [35]. Where the algorithms require a regularization parameter other than the noise level, it is empirically set to minimize the mean squared error. For Figure 4, we have chosen to use the functional shown in Equation 3. We set β_1 set to a very small positive constant, and set $\beta_2 = \beta^2, \beta_3 = \frac{\beta^4}{20}$, in order to slightly dampen the strength of the gradient coupling term.

The results clearly demonstrate the accurate deblurring obtained using the regularization offered by the Beltrami framework for natural color images, with slightly better PSNR compared to TV regularization. Beyond PSNR, careful examination of the images show the tendency of Beltrami regularization to avoid artifacts which do not fit the appearance of natural images, and discourage uneven coloring artifacts. This can be seen in Figures 4,5. The same discrepancy between PSNR reading and visual results in color image processing has already

been noted by Goldluecke and Cremers [14]. We iterate this word of caution, and refer the reader to the images themselves.

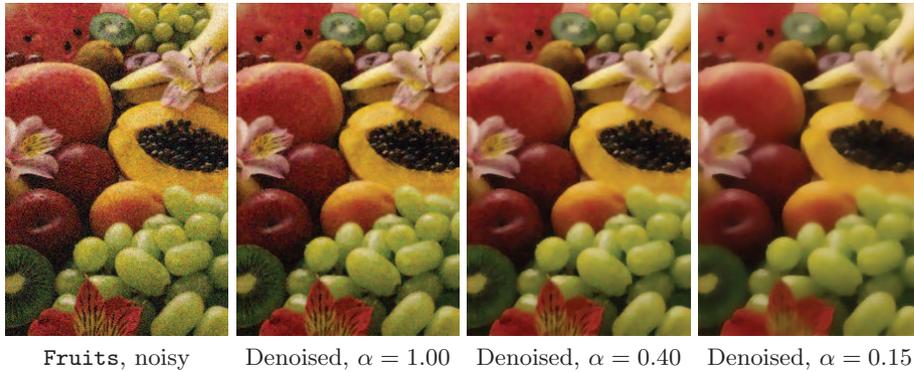


Fig. 1: Smoothing, under various α values, of the Fruits image, with added Gaussian noise with $\sigma = 20$ intensity levels per channel.

Image	CPU time, AL	CPU time, splitting
Astro	1.77s	3.5s
Fruits	2.97s	7.36s
Lion	21.23s	59.66s
Monarch	3.63s	7.71s

Table 1: Comparison of the CPU time required to complete 99% of the rise in SNR.

5 Conclusions

We presented an extension of the augmented Lagrangian method for color image processing with Beltrami regularization. Unlike existing techniques, the method discretizes the functional itself, rather than the resulting optimality conditions or minimizing flow. We present numerical examples demonstrating its efficiency and accuracy compared to existing techniques for variational regularization, and its effectiveness in image deblurring. In future work we intend to add a robust fidelity term [36], and explore other possible applications for our framework.

References

1. Bartomeu Coll Antoni Buades and Jean-Michel Morel, *A review of image denoising algorithms, with a new one*, SIAM Interdisciplinary Journal **4** (2005), 490–530.

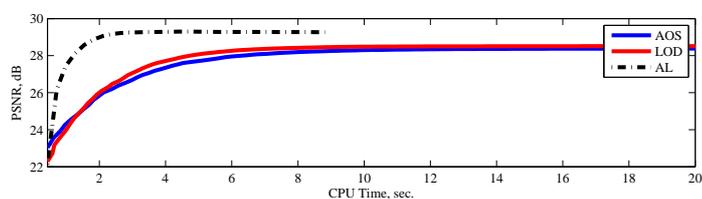
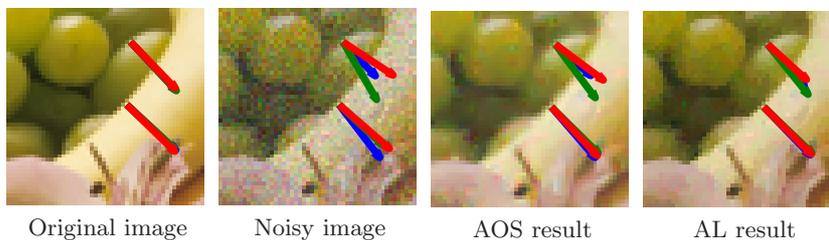


Fig. 2: A comparison of the results for the AOS scheme, and the augmented Lagrangian method, as well as the PSNR of splitting schemes and the AL method as a function of CPU time. The arrows in the images demonstrate gradient directions at each channel. The graph demonstrates a faster convergence of the augmented Lagrangian method, as well as a more accurate discretization.



Fig. 3: Images used to compare the computational cost of the augmented Lagrangian and splitting-based Beltrami regularization. Left to right: (a) Astro image. (b) Fruits image. (c) Lion image (d) Monarch image.

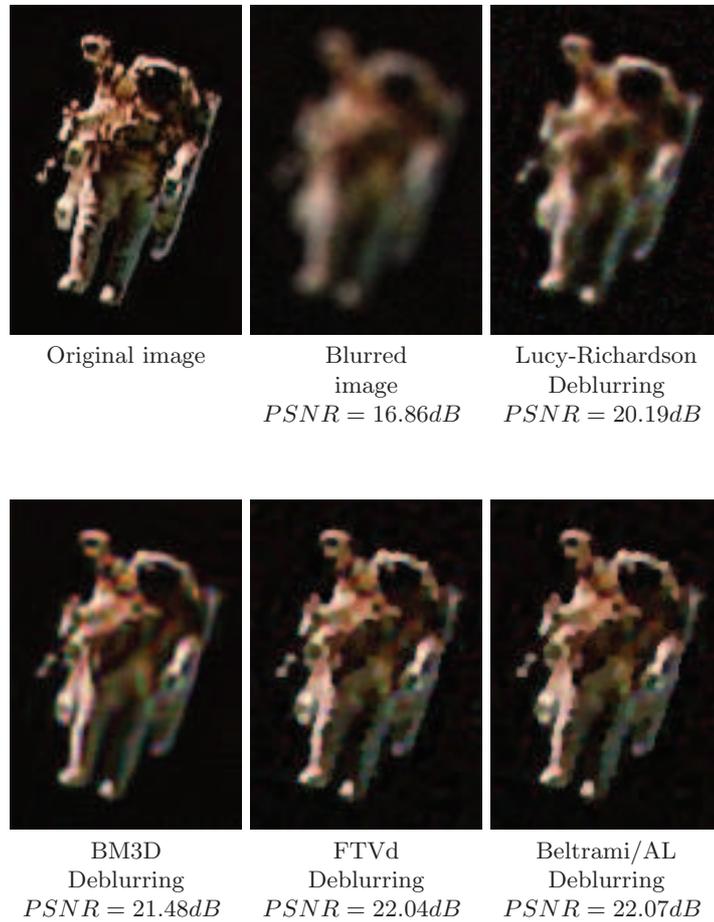


Fig. 4: Deblurring results with a disc blur filter of radius 5 and an Gaussian noise of $\sigma = 5$. Left to right, top to bottom: (a) The original image. (b) The blurred image. (c) Deblurring using the Lucy-Richardson algorithm (d) BM3D-based deblurring. (e) Deblurring using the FTVd method. (f) Deblurring using the Beltrami / augmented Lagrangian algorithm.

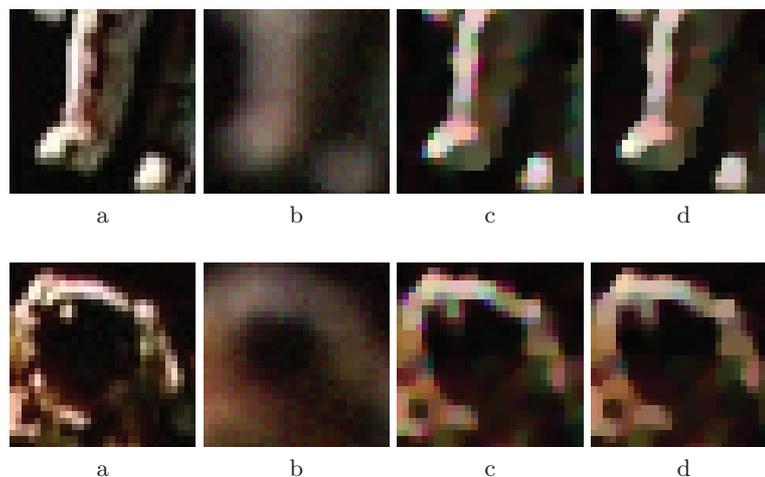


Fig. 5: Each row represents two regions zoomed-in from Figure 4. (a) The original image. (b) The corrupted image. (c) TV results. (d) Beltrami / AL restoration results.

2. Leah Bar, Alexndar Brook, Nir Sochen, and Nahum Kiryati, *Deblurring of color images corrupted by impulsive noise*, IEEE Trans. Image Process. **16** (2007), no. 4, 1101–1111.
3. Danny Barash, *A fundamental relationship between bilateral filtering, adaptive smoothing and the nonlinear diffusion equation*, IEEE Trans. Pattern Anal. Mach. Intell. **24** (2002), no. 6, 844–847.
4. Peter Blomgren and Tony F. Chan, *Color TV: Total variation methods for restoration of vector valued images*, IEEE Trans. Image Processing **7** (1996), 304–309.
5. Xavier Bresson and Tony Chan, *Fast dual minimization of the vectorial total variation norm and applications to color image processing*, CAM-Report 07-25, UCLA, 2007.
6. Xavier Bresson, Pierre Vandergheynst, and Jean-Philippe Thiran, *Multiscale active contours*, Int. J. of Comp. Vision **70** (2006), no. 3, 197–211.
7. Janylle Laurice Carter, *Dual methods for total variation-based image restoration*, CAM-Report 02-13, UCLA, April 2002.
8. Antonin Chambolle, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vis. **20** (2004), no. 1-2, 89–97.
9. Tony F. Chan, Gene H. Golub, and Pep Mulet, *A nonlinear primal-dual method for total variation-based image restoration*, SIAM J. Sci. Comput. **20** (1999), 1964–1977.
10. Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian, *Image restoration by sparse 3D transform-domain collaborative filtering*, Proc. SPIE (Jaakko T. Astola, Karen O. Egiazarian, and Edward R. Dougherty, eds.), vol. 6812, 2008.
11. Lorina Dascal, Guy Rosman, Xue-Cheng Tai, and Ron Kimmel, *On semi-implicit splitting schemes for the Beltrami color flow*, SSVM (Berlin, Heidelberg), Springer-Verlag, 2009, pp. 259–270.

12. Manfredo Perdigao do Carmo, *Riemannian geometry*, Birkhäuser Verlag, Boston, MA, 1992.
13. Michael Elad, *On the bilateral filter and ways to improve it*, IEEE Trans. Image Process. **11** (2002), no. 10, 1141–1151.
14. Bastian Goldluecke and Daniel Cremers, *An approach to vectorial total variation based on geometric measure theory*, Computer Vision and Pattern Recognition, 2010.
15. Magnus R. Hesteness, *Multipliers and gradient methods*, Journal of Optimization Theory and Applications **4** (1969), 303–320.
16. Ran Kaftory, Nir Sochen, and Yehoshua Y. Zeevi, *Variational blind deconvolution of multi-channel images*, IJIST **15** (2005), no. 1, 56–63.
17. Ron Kimmel, *Numerical geometry of images: Theory, algorithms, and applications*, Springer Verlag, 2003.
18. Ron Kimmel, Ravi Malladi, and Nir Sochen, *Images as embedding maps and minimal surfaces: Movies, color, texture, and volumetric medical images*, Int. J. of Comp. Vision **39** (2000), no. 2, 111–129.
19. T. Lü, Pekka Neittaanmäki, and Xue-Cheng Tai, *A parallel splitting up method and its application to Navier-Stokes equations*, Applied Mathematics Letters **4** (1991), no. 2, 25–29.
20. Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin, *An iterative regularization method for total variation-based image restoration*, Simul **4** (2005), 460–489.
21. Sylvain Paris and Frédo Durand, *A fast approximation of the bilateral filter using a signal processing approach*, Int. J. of Comp. Vision **81** (2009), no. 1, 24–52.
22. A. M. Polyakov, *Quantum geometry of bosonic strings*, Physics Letters **103 B** (1981), no. 3, 207–210.
23. Michael J.D. Powell, *Optimization*, ch. A method for nonlinear constraints in minimization problems, pp. 283–298, Academic Press, 1969.
24. Guy Rosman, Lorina Dascal, Ron Kimmel, and Avram Sidi, *Efficient beltrami image filtering via vector extrapolation methods*, SIAM J. Imag. Sci. (2008), no. 3, 858–878.
25. Guy Rosman, Xue-Cheng Tai, Ron Kimmel, and Lorina Dascal, *Polyakov action minimization for efficient color image processing*, Technical Report CIS-2010-04, Technion, 2010.
26. Leonid I. Rudin, Stanley Osher, and Emad Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D Letters **60** (1992), 259–268.
27. Guillermo Sapiro and Dario L. Ringach, *Anisotropic diffusion of multivalued images with applications to color filtering*, IEEE Trans. Image Process. **5** (1996), no. 11, 1582–1586.
28. Stephen M. Smith and J.M. Brady, *SUSAN – A new approach to low level image processing*, Int. J. of Comp. Vision **23** (1997), 45–78.
29. Nir Sochen, Ron Kimmel, and Alfred M. Bruckstein, *Diffusions and confusions in signal and image processing*, J. of Math. in Imag. and Vis. **14** (2001), no. 3, 195–209.
30. Nir Sochen, Ron Kimmel, and Ravi Maladi, *A general framework for low level vision*, IEEE Trans. Image Process. **7** (1998), no. 3, 310–318.
31. Alon Spira, Ron Kimmel, and Nir A. Sochen, *A short-time Beltrami kernel for smoothing images and manifolds.*, IEEE Trans. Image Process. **16** (2007), no. 6, 1628–1636.
32. Xue-Cheng Tai and Chunlin Wu, *Augmented Lagrangian method, dual methods and split Bregman iteration for ROF model*, SSVM, 2009, pp. 502–513.

33. Carlo Tomasi and Roberto Manduchi, *Bilateral filtering for gray and color images*, Int. Conf. on Comp. Vision (1998), 836–846.
34. David Tschumperle and Rachid Deriche, *Vector-valued image regularization with pdes: A common framework for different applications*, IEEE Transactions on Pattern Analysis and Machine Intelligence **27** (2005), 506–517.
35. Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang, *A new alternating minimization algorithm for total variation image reconstruction*, SIAM J. Imag. Sci. **1** (2008), no. 3, 248–272.
36. Chunlin Wu, Juyong Zhang, and Xue-Cheng Tai, *Augmented Lagrangian method for total variation restoration with non-quadratic fidelity*, CAM Report 09-82, UCLA, December 2009.
37. Leonid P. Yaroslavsky, *Digital picture processing*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1985.
38. Anthony J. Yezzi, *Modified curvature motion for image smoothing and enhancement*, IEEE Trans. Image Process. **7** (1998), no. 3, 345–352.
39. Silvano Di Zenzo, *A note on the gradient of a multi-image*, Computer Vision, Graphics, and Image Processing **33** (1986), no. 1, 116–125.