

VARIATIONAL ENHANCEMENT AND DENOISING OF FLOW FIELD IMAGES

Pouya Dehghani Tafti*, Ricard Delgado-Gonzalo*, Aurélien F. Stalder[§] and Michael Unser*

*Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

[§]Department of Radiology, Xuanwu Hospital, Capital Medical University, Beijing, China

ABSTRACT

In this work we propose a variational reconstruction algorithm for enhancement and denoising of flow fields that is reminiscent of total-variation (TV) regularization used in image processing, but which also takes into account physical properties of flow such as curl and divergence. We point out the invariance properties of the scheme with respect to transformations of the coordinate system such as shifts, rotations, and changes of scale. To demonstrate the utility of the reconstruction method, we use it first to denoise a simulated phantom where the scheme is found to be superior to its quadratic (L_2) variant both in terms of SNR and in preservation of discontinuities. We then use the scheme to enhance the quality of pathline visualizations in an application to 4D (3D+time) flow-sensitive magnetic resonance imaging of blood flow in the aorta.

Index Terms— regularization, variational reconstruction, denoising, vector fields, curl, divergence, invariance, total variation, flow-sensitive MRI.

1. INTRODUCTION

With recent advances in medical imaging, it is now possible to measure three-dimensional vector flow in vivo using a variety of imaging modalities. Flow-sensitive magnetic resonance imaging (MRI), for instance, can be utilized to produce gated measurements of blood flow in the cardiovascular system, which permit one to form an analysis of blood circulation and vessel wall parameters [1, 2]. Acceleration is desirable in this context to shorten relatively long acquisition times; however, as a general rule, faster acquisition leads to lower quality, measured in terms of signal-to-noise ratio (SNR). In addition, physical properties of flow such as incompressibility are not automatically enforced in unprocessed acquisitions. Furthermore, in flow visualization based on particle tracing, some form of flow-field enhancement is often desirable to improve the visual quality and informativeness of renditions by emphasizing important features and suppressing noise.

The above observations point to a need for physically-motivated algorithms for flow-field correction, enhancement, and denoising, which combine the adaptability and physical soundness needed to address the noted issues in a satisfactory manner [3, 4]. In the present paper, we aim to propose one such algorithm that, together with its generalizations, belongs to the family of variational reconstruction methods. In addition to its theoretical appeal from a variety of standpoints (statistical interpretation, link with PDEs, etc.), variational reconstruction—often identified by the keyword *regularization*—can, under suitable conditions, lead to convex optimization problems that guarantee uniqueness and which can be solved using reliable techniques.

This work was supported by the Swiss National Science Foundation grant 200020-121763.

Specifically, we shall focus on reconstruction schemes which can be posed as the minimization of a cost functional

$$J(\mathbf{f}; \mathbf{y}) := d(\mathbf{f}; \mathbf{y}) + \sum_i \lambda_i \text{Reg}_i(\mathbf{f}) \quad (1)$$

where \mathbf{f} is a possible solution; \mathbf{y} is the vector of measurements (observations); d is a data fidelity criterion ($= 0$ if \mathbf{y} corresponds exactly to measurements made from \mathbf{f}); the Reg_i 's are regularization functionals that take large values if the solution is undesirable in a certain respect; and, finally, the λ_i 's are regularization parameters that endow the algorithm with some level of flexibility to adapt to different circumstances by adjusting the contribution of each of the component terms of J . From a different perspective, the same cost functional can be seen as the Lagrange relaxation of a constrained optimization problem with inequality constraints. In practice, however, the constraint bounds are often unknown; using the formulation given in (1) allows us to interpret the λ_i 's as tuning parameters of the algorithm, which can be adjusted to improve or optimize some measure of reconstruction quality.

Most of the variational algorithms proposed in the context of vector field reconstruction to date have used quadratic regularization [5–8]. In the present paper, we contrast quadratic schemes to those based on L_1 regularization that is reminiscent of total-variation (TV) regularization in 2D image processing [9, 10]. It is found that the latter display better preservation of discontinuities, while also improving upon quadratic regularization in SNR.

Another notable property of the family of algorithms considered here (inclusive of quadratic and L_1 variants) is that they exhibit invariance to coordinate transformations such as rotation and scaling, which goes hand-in-hand with their inherent connection to physical properties of flow such as irrotational vs solenoidal tendencies [11, 12].

In the remainder of the paper, we first describe the variational formulation of the reconstruction problem and propose our solution in the continuous setting (§2.1). This is followed by a brief description of one possible way of discretization and numerical resolution of the problem in §2.2. We then discuss applications of the resulting scheme to denoising and enhancement of simulated and real data in §3, before concluding with a summary of the main points and future research directions in §4.

2. VARIATIONAL RECONSTRUCTION OF FLOW FIELDS

2.1. Continuous formulation

The choice of the data fidelity criterion in (1) is in practice often dictated by statistical observations with regard to measurement noise as well as practical convenience. The most common choice, to which we shall adhere, is the mean squared error or, in finite sample sizes, simply the absolute squared error.

The regularization functionals in (1) typically take the form of integrals $\int_{\mathbb{R}^d} \Phi(Rf(\mathbf{x})) d\mathbf{x}$, where R is a differential-type operator and Φ is a convex function.

In many practical situations, a preferred choice of orientation and scale cannot be inferred from the data, at least within a wide range. It is therefore reasonable in such circumstances to require that the regularization functionals not have preferential behaviour with respect to orientation and scale. For the above functional to be invariant under rotation, shift, and scaling (up to a scalar factor), it is sufficient to take Φ to be a homogeneous function and choose a regularization operator R that commutes with the said transformations (up to a scalar factor where appropriate). For homogeneous Φ we can write $\text{Reg}(f)$ as

$$\text{Reg}(f) = \|Rf\|_p^p$$

where the L_p norm is in general computed for the magnitude of Rf (since Rf can be vector-valued).

As for the regularization operator R , one can directly verify that the appropriate commutation/invariance laws are satisfied by the curl and divergence, as well as their properly formulated combinations with their adjoints (among which one finds the vector Laplacian $\Delta = \mathbf{grad} \text{div} - \mathbf{curl} \text{curl}$). Note that coordinate rotations for vector fields follow a law that is different from scalar fields. This is because coordinates of a vector field are specified in the coordinate system of its domain; hence, when rotating the domain coordinates by some orthogonal matrix Ω , one must transform the vector by Ω^T to keep its direction fixed. The rotation formula for vector fields is therefore $f \mapsto \Omega^T f(\Omega \cdot)$.

Among cost functionals that fit the above description, here we shall focus on the following first-order one, which can be seen as the vector equivalent of TV regularization in some sense:

$$\begin{aligned} J^{(p)}(f; \mathbf{y}) &= \sum_m |f(\mathbf{m}) - \mathbf{y}[\mathbf{m}]|^2 + \lambda_c \|\mathbf{curl} f\|_p^p + \lambda_d \|\text{div} f\|_p^p \\ &= \sum_m |f(\mathbf{m}) - \mathbf{y}[\mathbf{m}]|^2 + \lambda_c \int_{\mathbb{R}^d} (\sqrt{|\mathbf{curl} f|^2})^p \\ &\quad + \lambda_d \int_{\mathbb{R}^d} (\sqrt{|\text{div} f|^2})^p. \end{aligned}$$

($|\cdot|$, when applied to a vector, denotes its magnitude). This cost functional allows us to independently penalize rotational and compressive/divergent behaviour of the reconstruction by means of the parameters λ_c and λ_d .

2.2. Discretization and implementation

Replacing partial derivatives by finite difference operators $\delta_i : f \mapsto f - f[\cdot - \mathbf{e}_i]$, where \mathbf{e}_i , $i = 1, \dots, d$, denotes the i th standard unit vector in \mathbb{R}^d , we arrive at the discrete formulation of the problem (in implementation, boundary conditions are incorporated in δ_i):

$$\begin{aligned} J_\delta^{(p)}(f; \mathbf{y}) &= \sum_m |f[\mathbf{m}] - \mathbf{y}[\mathbf{m}]|^2 + \lambda_c \sum_m (\sqrt{|\mathbf{curl}_\delta f[\mathbf{m}]|^2})^p \\ &\quad + \lambda_d \sum_m (\sqrt{|\text{div}_\delta f[\mathbf{m}]|^2})^p, \end{aligned} \quad (2)$$

where for the magnitudes of discrete curl and divergence we have,

$$\begin{aligned} |\mathbf{curl}_\delta f[\mathbf{m}]|^2 &= \sum_{1 \leq i < j \leq d} (\delta_i f_j[\mathbf{m}] - \delta_j f_i[\mathbf{m}])^2; \\ |\text{div}_\delta f[\mathbf{m}]|^2 &= \sum_{1 \leq i, j \leq d} \delta_i f_i[\mathbf{m}] \delta_j f_j[\mathbf{m}]. \end{aligned}$$

We shall consider the quadratic ($p = 2$) and L_1 ($p = 1$) cases. The former is comparable to what is given in Dodu and Rabut [7] and Arigovindan & al. [8], while the latter is more reminiscent of TV regularization. For $p = 2$, (2) describes a quadratic programme that can be solved using iterative linear solvers (the derivatives of a quadratic cost function are linear). For $p = 1$, on the other hand, we follow Figueiredo & al. [13], and form a sequence of quadratic upper bounds (majorizers) on the cost function, which can then be reduced sequentially in order to approach the global solution.

Specifically, using the inequality $\sqrt{a} \leq \sqrt{a'} + \frac{1}{2}(a - a')/\sqrt{a'}$ for positive a, a' and ignoring terms that depend only on f' and \mathbf{y} (which do not affect the solution), we define the upper bound at fixed f' as

$$\begin{aligned} Q_\delta(f, f'; \mathbf{y}) &:= \sum_m \sum_{1 \leq i \leq d} f_i[\mathbf{m}]^2 - \sum_m \sum_{1 \leq i \leq d} 2f_i[\mathbf{m}]y_i[\mathbf{m}] \\ &\quad + \lambda_c \sum_m c_m^{-1} |\mathbf{curl}_\delta f[\mathbf{m}]|^2 \\ &\quad + \lambda_d \sum_m d_m^{-1} |\text{div}_\delta f[\mathbf{m}]|^2, \end{aligned}$$

with

$$c_m := \sqrt{|\mathbf{curl}_\delta f'[\mathbf{m}]|^2}, \quad d_m := \sqrt{|\text{div}_\delta f'[\mathbf{m}]|^2}. \quad (3)$$

Next, consider the sequence

$$\tilde{f}_{(n)} := \arg \min_f Q_\delta(f, \tilde{f}_{(n-1)}; \mathbf{y}), \quad (4)$$

defined recursively with some initialization such as $\tilde{f}_{(0)} = \mathbf{0}$. For a given $\tilde{f}_{(n-1)}$, the above minimization is a quadratic re-weighted least square problem, which can be solved by a linear solver. Note that minimizing $Q(f, f'; \mathbf{y})$ is equivalent to minimizing $J_\delta^{(1)}(f; \mathbf{y})$. Furthermore, we have

$$\begin{aligned} Q(\tilde{f}_{(n)}, \tilde{f}_{(n)}; \mathbf{y}) &\leq Q_\delta(\tilde{f}_{(n)}, \tilde{f}_{(n-1)}; \mathbf{y}) \\ &< Q_\delta(\tilde{f}_{(n-1)}, \tilde{f}_{(n-1)}; \mathbf{y}) = Q(\tilde{f}_{(n-1)}, \tilde{f}_{(n-1)}; \mathbf{y}) \end{aligned}$$

which shows that, with increasing n , the $J_\delta^{(1)}(\tilde{f}_{(n)}; \mathbf{y})$'s form a decreasing sequence.

The final scheme for $p = 1$ is given in Algorithm 1. In practice the local minimization in the last step of the loop may be replaced by a fixed number of iterations of a linear solver such as conjugate gradient (CG), which will still reduce the global cost.

Algorithm 1: Algorithm for L_1 regularization

```

input:  $\mathbf{y}$ ;
 $\tilde{f}_{(0)} \leftarrow \mathbf{0}$ ;
repeat
   $n \leftarrow n + 1$ ;
  for all data coordinates  $\mathbf{m}$  do
     $c_m \leftarrow \sqrt{|\mathbf{curl}_\delta f'[\mathbf{m}]|^2}$ ;  $d_m \leftarrow \sqrt{|\text{div}_\delta f'[\mathbf{m}]|^2}$ ;
  end
   $\tilde{f}_{(n)} \leftarrow \arg \min_f Q_\delta(f, \tilde{f}_{(n-1)}; \mathbf{y})$ ;
until stopping criteria are met;
return  $\tilde{f}_{(n)}$ .

```

3. EXPERIMENTS

3.1. Simulation

As our first experiment we considered the problem of denoising velocities sampled from a phantom modelling laminar flow with a hyperbolic profile inside a tube passing through a torus in which a uniform flow circulates (see sample visualizations of input and output in Fig. 1). We repeated the experiment with different levels of independent additive white Gaussian noise, and compared L_1 and L_2 regularization (corresponding to $p = 1$ and 2 in (2)) for each instance. For L_1 denoising, a CG solver with 100 iterations was used to approximately solve the quadratic problem in each iteration of Algorithm 1, and the number of external iterations was fixed at 20. For fair comparison, in the case of L_2 regularization we allowed a maximum of 100×20 CG iterations to solve the linear problem, although in practice the CG solver always converged before reaching this limit.

Results are summarized in Table 1. L_1 regularization is seen to be superior to L_2 regularization in all cases in terms of SNR, and visual inspection also showed it to be better at preserving discontinuities, as is the case for its scalar counterpart (TV).

Table 1: Comparison of denoising algorithms for phantom data

dataset	SNR [dB]		
input	0.00	10.00	20.00
L_1 regularized	9.01	14.53	21.22
L_2 regularized	8.77	13.79	20.85

Although in the results reported in Table 1 we used an oracle to optimize the parameters λ_c, λ_d (done separately for L_1 and L_2 regularization), comparison with the Monte Carlo SURE method of Ramani & al. [15], which does not use the ground truth, showed that under the Gaussian noise hypothesis, optimal parameters can be accurately estimated from the noisy data alone at twice the computational cost of using an oracle.

3.2. 4D flow-sensitive MRI

In our second experiment, we used flow-sensitive phase-contrast MRI recordings of blood flow in the thoracic aorta of a healthy subject. The dataset was acquired on a 3T MR system (Magnetom TRIO Tim, Siemens Medical Solutions, Erlangen, Germany) using a 12-channel body coil. All measurements were made using an RF-spoiled gradient echo sequence with prospective ECG gating and respiratory gating [16]. The sequence made use of parallel imaging based on the GRAPPA technique [17] with 24 auto-calibration lines and undersampling factor $R = 3$. The sequence included correction for Maxwell terms and the data were further corrected for background phase error using a second-order model [18].

To this data we applied the proposed L_1 algorithm for flow enhancement. In this case the Gaussian noise assumption was not valid in this case, and algorithm parameters predicted by Monte Carlo SURE were overly conservative. We therefore set the parameters manually so as to correct for non-zero divergence in the measured data and thereby enhance visualization based on particle tracing, by strongly penalizing non-zero divergence (i.e. by setting $\lambda_d \gg \lambda_c$); this was justified on the basis of the near-incompressibility (hence almost zero divergence) of blood. Pathline visualizations of the flow, before and after correction using L_1 regularization, can be seen in Fig. 2. Pathlines (also

called particle traces) represent the trajectories of virtual particles in a velocity field which, in our case, were emitted from a plane in the ascending aorta and integrated over two cardiac cycles in the blood flow velocity field. Here we used a commercial 3D visualization software package (EnSight, CEI, NC, USA) to produce the visualizations.

As is seen in Fig. 2, pathline visualizations of the uncorrected MR dataset produced many pathlines flowing outside of the lumen volume. Consequently, there were few pathlines going all the way through the aorta. In visualizations created from the corrected dataset, on the other hand, fewer pathlines were flowing out of the arterial volume and there were more trajectories going down the aorta all the way to the descending aorta. Because of these two aspects, the apparent quality of pathline visualizations made from the corrected dataset was substantially improved compared to those created from the original uncorrected dataset. Due to the limited resolution of the data and of the proximity between the thoracic trunk and the aorta, in both visualizations one can see cases of pathlines that cross vessel walls.

4. CONCLUSION

In the present paper we considered the related problems of denoising and correction/enhancement of vector field data, and proposed variational algorithms for this purpose based on the notion of invariance to coordinate transformations such as rotation and scaling. Among variational algorithms that satisfy these invariances, we focused on a first-order variational scheme with L_1 regularization which is evocative of TV regularization for scalar fields, but is different in that it additionally takes account of vector qualities without scalar parallels, such as rotation and divergence.

We described an implementation of the variational scheme using quadratic upper bounds and re-weighted least squares, and compared it against its L_2 variant (comparable to smoothing spline schemes in image processing). In experiments with simulated phantoms, the proposed algorithm was found to be superior both in terms of SNR and in preserving discontinuities, over a range of input SNRs. We then discussed an application of the scheme to real data obtained by means of flow-sensitive MRI, for the purpose of improving the characteristics and visual representation of blood flow by correcting for non-zero divergence.

Possible directions for future research include higher order algorithms in the same family not covered in the present paper, faster implementations, and alternative means of predicting regularization parameters and assessing results, as well as applications in resolution enhancement, deconvolution, and other areas where reconstruction of vector fields might be of interest, such as image registration and estimation of optical flow.

5. REFERENCES

- [1] A. F. Stalder, M. F. Russe, A. Frydrychowicz, J. Bock, J. Hennig, and M. Markl, "Quantitative 2D and 3D phase contrast MRI: Optimized analysis of blood flow and vessel wall parameters," *Magn. Reson. Med.*, vol. 60, no. 5, pp. 1218–1231, 2008.
- [2] A. Frydrychowicz, R. Arnold, D. Hirtler, C. Schlensak, A. F. Stalder, J. Hennig, M. Langer, and M. Markl, "Multidirectional flow analysis by cardiovascular magnetic resonance in aneurysm development following repair of aortic coarctation," *J. Cardiovasc. Magn. Reson.*, vol. 10, no. 1, 2008.

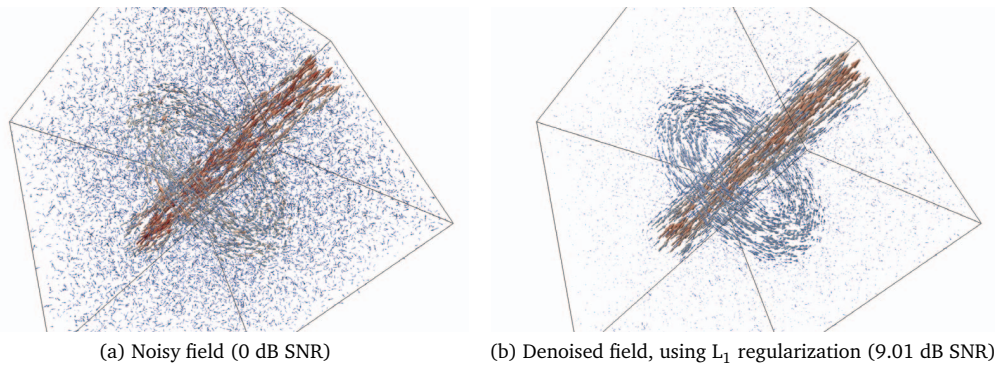


Fig. 1: Phantom before and after denoising; visualized using ParaView 3.8.0 [14].

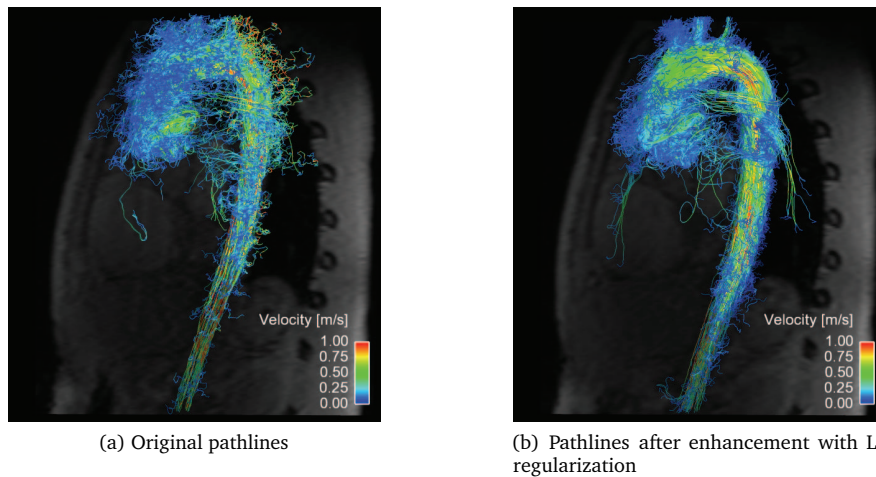


Fig. 2: Flow-sensitive MRI recordings of blood flow in the aorta before and after denoising.

- [3] M. H. Buonocore, "Algorithms for improving calculated streamlines in 3-D phase contrast angiography," *Magn. Reson. Med.*, vol. 31, no. 1, pp. 22–30, 1994.
- [4] S. M. Song, S. Napel, G. H. Glover, and N. J. Pelc, "Noise reduction in three-dimensional phase-contrast MR velocity measurements," *J. Magn. Reson. Imaging*, vol. 3, no. 4, pp. 587–596, 1993.
- [5] D. Suter and F. Chen, "Left ventricular motion reconstruction based on elastic vector splines," *IEEE Trans. Med. Imaging*, vol. 19, no. 4, pp. 295–305, 2000.
- [6] L. Amodè and M. N. Benbourhim, "A vector spline approximation," *J. Approx. Theory*, vol. 67, pp. 51–79, 1991.
- [7] F. Dodu and C. Rabut, "Irrational or divergence-free interpolation," *Numer. Math.*, vol. 98, pp. 477–498, 2004.
- [8] M. Arigovindan, M. Sühling, C. Jansen, P. Hunziker, and M. Unser, "Full motion and flow field recovery from echo Doppler data," *IEEE Trans. Med. Imaging*, vol. 26, no. 1, pp. 31–45, 2007.
- [9] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [10] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, 1992.
- [11] M. Arigovindan, "Variational reconstruction of vector and scalar images from non-uniform samples," Ph.D. dissertation, EPFL, Biomedical Imaging Group, 2005.
- [12] P. D. Tafti and M. Unser, "Fractional Brownian vector fields," *Multiscale Model. Simul.*, vol. 8, no. 5, pp. 1645–1670, 2010.
- [13] M. A. T. Figueiredo, J. B. Dias, J. P. Oliveira, and R. D. Nowak, "On total variation denoising: A new majorization-minimization algorithm and an experimental comparison with wavelet denoising," in *Proc. 2006 IEEE Int. Conf. Image Process. (ICIP 2006)*, 2006, pp. 2633–2636.
- [14] J. Ahrens, B. Geveci, and C. Law, "Paraview: An end-user tool for large-data visualization," in *Visualization Handbook*, C. D. Hansen and C. R. Johnson, Eds. Burlington: Butterworth-Heinemann, 2005, pp. 717–732.
- [15] S. Ramani, T. Blu, and M. Unser, "Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Trans. Image Process.*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [16] M. Markl, A. Harloff, T. A. Bley, M. Zaitsev, B. Jung, E. Weigang, M. Langer, J. Hennig, and A. Frydrychowicz, "Time-resolved 3D MR velocity mapping at 3T: Improved navigator-gated assessment of vascular anatomy and blood flow," *J. Magn. Reson. Imaging*, vol. 25, no. 4, pp. 824–831, 2007.
- [17] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. Med.*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [18] M. A. Bernstein, X. J. Zhou, J. A. Polzin, K. F. King, A. Ganin, N. J. Pelc, and G. H. Glover, "Concomitant gradient terms in phase contrast MR: Analysis and correction," *Magn. Reson. Med.*, vol. 39, no. 2, pp. 300–308, 1998.