A DUAL ALGORITHM FOR L1-REGULARIZED RECONSTRUCTION OF VECTOR FIELDS

Emrah Bostan, Pouya Dehgani Tafti, and Michael Unser

Biomedical Imaging Group, EPFL, Switzerland

ABSTRACT

Recent advances in vector-field imaging have brought to the forefront the need for efficient denoising and reconstruction algorithms that take the physical properties of vector fields into account and can be applied to large volumes of data. With these requirements in mind, we propose a computationally efficient algorithm for variational denoising and reconstruction of vector fields. Our variational objective combines rotation- and scale-invariant regularization functionals that permit one to tune the algorithm to the physical characteristics of the underlying phenomenon. In addition, these regularization terms involve L_1 norms in the spirit of total-variation (TV) regularization, which, as in the scalar case, leads to better preservation of discontinuities and superior SNR performance compared to its quadratic alternative. Some experimental results are provided to illustrate and verify the proposed scheme.

***Index Terms*—** Vector fields, denoising, reconstruction from incomplete data, total-variation regularization, curl, divergence, splitting methods.

1. INTRODUCTION

Denoising and enhancement of vector fields arising in flow-sensitive medical imaging, optical-flow reconstruction, electromagnetic-field measurement, and a variety of other applications is an increasingly topical subject of research in image processing. This is in part due to advances in imaging and measurement technologies that have led to new modalities such as flow-sensitive magnetic resonance imaging (MRI) [1, 2]. Moreover, the rise in computational power and capacity over the past decade has made it possible to process large volumes of multidimensional data that result from measuring vector fields. A vector field in 3D consists of 3 components per voxel, which requires one to store $3N^3$ values for a volume of size $N \times N \times N$.

Since the measurements of a vector field are typically related in a fundamental way to some physical phenomenon/model, it is worthwhile to design reconstruction frameworks that take the physical properties of the model into account. Following this reasoning, Tafti and Unser [3] have recently introduced a regularized reconstruction method which is not only physically driven but is also based on the principle of invariance to coordinate transformations such as rotation and scaling. Their variational method takes the form of the optimization problem

$$\mathbf{f}_{\mathrm{opt}} = \underset{\mathbf{f}}{\operatorname{argmin}} \quad \mathfrak{D}(\mathbf{f};\mathbf{y}) + \sum_i \lambda_i \|\mathrm{R}_i \mathbf{f}\|_p^p, \tag{1}$$

where $\mathbf{f}_{\mathrm{opt}}$ is the reconstructed vector field, $\mathbf{y}$ is the vector of measurements, $\mathfrak{D}$ is the data fidelity functional, R_i is the i^{th} differential regularization operator, and $\|\cdot\|_p$ is the vector (or tensor) L_p norm (for vectors, it corresponds to the scalar L_p norm of the magnitude of the vector field) [3].

One of the examples of the regularization scheme in [3] involves the use of curl and divergence operators for regularization. In the context of quadratic regularization (*i.e.*, with $p = 2$), these regularization operators have been considered in the past by several authors: See for instance Suter and Chen [4], Dodu and Rabut [5], and Arigovindan *et al.* [6]. Consistent with the recent trend in biomedical image processing—which favors total-variation (TV) type and L_1 regularization—it is observed in Tafti and Unser [3] that the use of L_1 norms in (1) leads in practice to better performance then the quadratic (L_2) alternative in preserving discontinuities. This approach is further demonstrated in Tafti *et al.* [7], where it is used to enhance the pathline visualization of blood flow in the aorta.

Unlike in the case of quadratic regularization, where the solution to the optimization problem can be obtained using iterative linear methods, the numerical problem to solve in the L_1 setting is nonlinear and the need to identify efficient algorithms for minimizing (1) is far more pronounced. The algorithm proposed by Tafti and Unser [3] for the purpose of minimizing (1) with L_1 norms falls in the category of iteratively reweighted least square (IRLS) bound-optimization methods. While the said method is capable of producing results in reasonable time, we here show that we are able to gain in terms of objective minimization and reconstruction performance within same amount of time. To do so, we use the Legendre-Fenchel transformation [8] to reformulate the optimization problem in its dual form and adapt the forward-backward splitting-type algorithm [9] to our needs. The improved computational efficiency makes it more convenient to consider difficult problems such as reconstruction from incomplete measurements.

In the present paper, we first describe our new formulation (Section 2.1), and then apply it to the problems of vector-field denoising (Section 2.1) and reconstruction of vector fields from missing data (Section 2.2).

2. PROBLEM FORMULATION

In the continuous setting, we combine a quadratic data-fidelity term with L_1 div-$\mathbf{curl}$ regularization to arrive at the energy functional

$$\begin{aligned} J(\mathbf{f};\mathbf{y}) &= \sum_{\mathbf{m}} |\mathrm{A}\mathbf{f}(\mathbf{m}) - \mathbf{y}[\mathbf{m}]|^2 + \lambda_c \|\mathbf{curl}\ \mathbf{f}\|_1 \\ &\qquad + \lambda_d \|\mathrm{div}\ \mathbf{f}\|_1 \\ &= \sum_{\mathbf{m}} |\mathrm{A}\mathbf{f}(\mathbf{m}) - \mathbf{y}[\mathbf{m}]|^2 + \lambda_c \int_{\mathbb{R}^d} |\,\mathbf{curl}\ \mathbf{f}| \\ &\qquad + \lambda_d \int_{\mathbb{R}^d} |\,\mathrm{div}\ \mathbf{f}|, \end{aligned}$$

where A is a linear operator acting on $\mathbf{f}$ and $|\cdot|$ denotes the absolute value of a scalar, or the magnitude of a vector, as appropriate. For

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

discretization of the partial derivatives involved in the definition of curl and divergence, we use the finite-difference operators

$$\delta_i : \mathbf{f} \mapsto \mathbf{f} - \mathbf{f}[\cdot - \mathbf{e}_i],$$

where $\mathbf{e}_i$, $i = 1, \ldots, d$, denotes the i^{th} standard unit vector in $\mathbb{R}^d$. The adjoint of δ_i is the operator

$$\delta_i^* : \mathbf{f} \mapsto \mathbf{f} - \mathbf{f}[\cdot + \mathbf{e}_i].$$

In the implementation, δ_i also incorporates the desired boundary conditions and δ_i^* their adjoint conditions. In three dimensions (*i.e.*, with $d = 3$), where we have $\mathbf{f}[\mathbf{m}] = (f_1[\mathbf{m}], f_2[\mathbf{m}], f_3[\mathbf{m}])$, curl and divergence are discretized as

$$\begin{aligned}\mathbf{curl}_\delta\, \mathbf{f} &= (\delta_3 f_2 - \delta_2 f_3, \delta_1 f_3 - \delta_3 f_1, \delta_2 f_1 - \delta_1 f_2)\\ \operatorname{div}_\delta \mathbf{f} &= \delta_1 f_1 + \delta_2 f_2 + \delta_3 f_3.\end{aligned}$$

Their adjoints are found to be

$$\begin{aligned}\mathbf{curl}_\delta^*\, \mathbf{f} &= (-\delta_3^* f_2 + \delta_2^* f_3, \delta_3^* f_1 - \delta_1^* f_3, -\delta_2^* f_1 + \delta_1^* f_2)\\ \operatorname{div}_\delta^* \mathbf{f} &= (\delta_1^* f_1, \delta_2^* f_2, \delta_3^* f_3).\end{aligned}$$

2.1. Denoising of Vector Fields

When A is the identity operator, our problem simplifies to vector-field denoising. To present our formulation, we first introduce the Legendre-Fenchel transforms of our regularizers, which we then use to find an equivalent of (1). We introduce dual variables $\mathbf{u}$ (vector) and v (scalar)

$$\begin{aligned}\|\mathbf{curl}_\delta\ \mathbf{f}\|_1 &= \max_{\mathbf{u}\in\mathcal{B}_3} \langle \mathbf{curl}_\delta^*\, \mathbf{u}, \mathbf{f}\rangle\\ \|\operatorname{div}_\delta\ \mathbf{f}\|_1 &= \max_{v\in\mathcal{B}_1} \langle \operatorname{div}_\delta^*\, v, \mathbf{f}\rangle,\end{aligned}$$

where $\mathcal{B}_l$, $l = 1, 3$, are the unit balls of the scalar and vector dual (infinity) norms, respectively. They are defined by the identity $\mathcal{B}_l = \{u : \|u\|_\infty \leq 1\}$. We recall that the vector ℓ_∞ norm is defined as the standard ℓ_∞ norm of the magnitude of the vector field. The optimization problem then becomes

$$\min_{\mathbf{f}} \max_{\mathbf{u}\in\mathcal{B}_3} \max_{v\in\mathcal{B}_1} \Big(\sum_{\mathbf{m}} |\mathbf{f}[\mathbf{m}] - \mathbf{y}[\mathbf{m}]|^2 \tag{3}$$

$$+ \lambda_c \langle \mathbf{curl}_\delta^*\, \mathbf{u}, \mathbf{f}\rangle + \lambda_d \langle \operatorname{div}_\delta^*\, v, \mathbf{f}\rangle \Big). \tag{4}$$

At this point, following an argument similar to the one given by Beck and Teboulle [10], we exchange the order of minimization and maximization and obtain

$$\max_{\mathbf{u}\in\mathcal{B}_3} \max_{v\in\mathcal{B}_1} \min_{\mathbf{f}} \Big(\sum_{\mathbf{m}} |\mathbf{f}[\mathbf{m}] - \mathbf{y}[\mathbf{m}]|^2 \tag{5}$$

$$+ \lambda_c \langle \mathbf{curl}_\delta^*\, \mathbf{u}, \mathbf{f}\rangle + \lambda_d \langle \operatorname{div}_\delta^*\, v, \mathbf{f}\rangle \Big). \tag{6}$$

By the first-order optimality condition, the solution of the internal minimization problem is found to in terms of the dual variables $\mathbf{u}, v$ as

$$\mathbf{f} = \mathbf{y} - \frac{\lambda_c}{2}\, \mathbf{curl}_\delta^*\, \mathbf{u} - \frac{\lambda_d}{2}\, \operatorname{div}_\delta^*\, v. \tag{7}$$

After substituting (7) back into (6), we arrive at the optimization problem

$$\min_{\mathbf{u}\in\mathcal{B}_3} \min_{v\in\mathcal{B}_1} \{\Phi(\mathbf{u}, v) + \Psi(\mathbf{u}, v)\}, \tag{8}$$

where

$$\begin{aligned}\Phi(\mathbf{u}, v) = \quad & \sum_{\mathbf{m}} \left| \mathbf{y}[\mathbf{m}] - \frac{\lambda_c}{2}\, \mathbf{curl}_\delta^*\, \mathbf{u}[\mathbf{m}] \right|^2\\ & + \sum_{\mathbf{m}} \left| \mathbf{y}[\mathbf{m}] - \frac{\lambda_d}{2}\, \operatorname{div}_\delta^*\, v[\mathbf{m}] \right|^2\\ & + \frac{1}{2} \langle \lambda_c\, \mathbf{curl}_\delta^*\, \mathbf{u}, \lambda_d\, \operatorname{div}_\delta^*\, v \rangle\end{aligned}$$

and $\Psi(\mathbf{u}, v) = \xi_{\mathcal{B}_3}(\mathbf{u}) + \xi_{\mathcal{B}_1}(v)$ with $\xi_{\mathcal{B}_l}$, $l = 1, 3$, denoting the indicator function of the dual ball $\mathcal{B}_l$, defined as

$$\xi_{\mathcal{B}_l}(u) = \begin{cases} 0, & \text{if } u \in \mathcal{B}_l \\ \infty, & \text{otherwise.}\end{cases}$$

In (8), note that $\Phi \in \mathcal{C}^1$ whereas $\Psi \notin \mathcal{C}^1$. The dual problem in (8) is however solvable by means of the proximal gradient algorithm of Beck and Teboulle [10], which is a variant of the *forward-backward splitting* algorithm as explained by Combettes and Pesquet [9]. This approach involves the computation of the *proximity operator* associated with Ψ. For some $L_\mathbf{u}, L_v > 0$, at iteration k, this operator is given by

$$\begin{aligned}&\operatorname{prox}_\Psi(\mathbf{u}^{(k)}, v^{(k)}; L_\mathbf{u}; L_v) =\\ &\quad \operatorname*{argmin}_{(\mathbf{w},z)} \left(\Psi(\mathbf{w}, z) + \frac{1}{2} \left\| \mathbf{w} - \left(\mathbf{u}^{(k)} - \frac{1}{L_\mathbf{u}} \frac{\partial \Phi}{\partial \mathbf{u}} \right) \right\|_2^2 \right.\\ &\quad \left. + \frac{1}{2} \left\| z - \left(v^{(k)} - \frac{1}{L_v} \frac{\partial \Phi}{\partial v} \right) \right\|_2^2 \right).\end{aligned} \tag{9}$$

In our case, for $\Psi(\mathbf{u}, v) = \xi_{\mathcal{B}_3}(\mathbf{u}) + \xi_{\mathcal{B}_1}(v)$, the proximity operator can be applied element-wise. The parameters $L_\mathbf{u}, L_v$ in (9), which determine the step sizes, are computed by bounding the Lipschitz constants of the respective gradients. Note that we perform proximal projection in a separable fashion: first with respect to one of our dual variables assuming that the other one is fixed, and then with respect to the other dual variable assuming that the previous variable is fixed. Algorithmically, all this implies that

$$\begin{aligned}\mathbf{u}^{(k+1)} &= \operatorname{prox}_{\xi_{\mathcal{B}_3}}(\mathbf{u}^{(k)}; L_\mathbf{u})\\ &= \mathcal{P}_{\mathcal{B}_3} \left(\mathbf{u}^{(k)} - \frac{1}{L_\mathbf{u}} \frac{\partial \Phi}{\partial \mathbf{u}} \right)\\ v^{(k+1)} &= \operatorname{prox}_{\xi_{\mathcal{B}_1}}(v^{(k)}; L_v)\\ &= \mathcal{P}_{\mathcal{B}_1} \left(v^{(k)} - \frac{1}{L_v} \frac{\partial \Phi}{\partial v} \right),\end{aligned}$$

where $\mathcal{P}_{\mathcal{B}_l}$, $l = 1, 3$, the projection onto the ball $\mathcal{B}_l$, is implemented by "shrinking" those elements of the argument that have magnitude larger than 1. The gradients appearing above are computed by

$$\begin{aligned}\frac{\partial \Phi}{\partial \mathbf{u}} &= -\lambda_c\, \mathbf{curl}_\delta \left(\mathbf{y} - \frac{\lambda_c}{2}\, \mathbf{curl}_\delta^*\, \mathbf{u} \right) + \frac{\lambda_c \lambda_d}{2}\, \mathbf{curl}_\delta \operatorname{div}_\delta^*\, v\\ \frac{\partial \Phi}{\partial v} &= -\lambda_d \operatorname{div}_\delta \left(\mathbf{y} - \frac{\lambda_d}{2} \operatorname{div}_\delta^*\, v \right) + \frac{\lambda_c \lambda_d}{2} \operatorname{div}_\delta\, \mathbf{curl}_\delta^*\, \mathbf{u}.\end{aligned}$$

The final scheme for the denoising problem is given in Algorithm 1.

Algorithm 1: Vector-Field Denoising

input: $\mathbf{y}, \lambda_c \in \mathbb{R}_{>0}, \lambda_d \in \mathbb{R}_{>0}$
set: $k \leftarrow 0; \mathbf{u}^{(1)}, \mathbf{u}^{(0)} \leftarrow \mathbf{0}; v^{(1)}, v^{(0)} \leftarrow 0; t^0 \leftarrow 1$
$L_{\mathbf{u}} \leftarrow 12\lambda_c^2; L_v \leftarrow 12\lambda_d^2$
repeat
  $k \leftarrow k+1$
  $\mathbf{u}^{(k)} \leftarrow \mathrm{prox}_{1/L_{\mathbf{u}}}(\mathbf{u}^{(k)})$
  $v^{(k)} \leftarrow \mathrm{prox}_{1/L_v}(v^{(k)})$
  $t^{(k)} \leftarrow \frac{1+\sqrt{1+(4t^{(k-1)})^2}}{2}$
  $\mathbf{u}^{(k+1)} \leftarrow \mathbf{u}^{(k)} + \left(\frac{t^{(k-1)}-1}{t^{(k)}}\right)\left(\mathbf{u}^{(k)} - \mathbf{u}^{(k-1)}\right)$
  $v^{(k+1)} \leftarrow v^{(k)} + \left(\frac{t^{(k-1)}-1}{t^{(k)}}\right)\left(v^{(k)} - v^{(k-1)}\right)$
until stopping criterion is met
return $\mathbf{f}_{\mathrm{opt}} \leftarrow \mathbf{y} - \frac{\lambda_c}{2}\,\mathbf{curl}_\delta^*\,\mathbf{u}^{(k+1)} - \frac{\lambda_d}{2}\,\mathrm{div}_\delta^*\,v^{(k+1)}$

2.2. Reconstruction from Missing Measurements

We now consider the problem of reconstructing the vector field from incomplete measurements. In this case, the operator A is modeled by a self-adjoint diagonal operator. It is obtained by replacing a subset of rows of an identity operator with zero vectors, implying that the spectral norm of $\mathrm{A}^*\mathrm{A}$ is 1. The difficulty now is that the proximity operator does not admit a closed-form solution anymore. In effect, the proximity operator at each iteration now corresponds to a denoising problem in its own right, which we can solve approximately. We denote this approximate proximity operator by $\mathrm{prox}^\epsilon(\cdot, \lambda_c, \lambda_d)$, and use Algorithm 1 to implement it. We summarize our approach in Algorithm 2.

Algorithm 2: Reconstruction from Missing Data

input: $\mathbf{y}, \lambda_c \in \mathbb{R}_{>0}, \lambda_d \in \mathbb{R}_{>0}$
set: $k \leftarrow 0; \mathbf{f}^{(1)}, \mathbf{f}^{(0)} \leftarrow 0; t^0 \leftarrow 1$
repeat
  $k \leftarrow k+1$
  $\mathbf{b}^{(k)} \leftarrow \mathbf{f}^{(k)} - \mathrm{A}^*\left(\mathrm{A}\mathbf{f}^{(k)} - \mathbf{y}\right)$
  $\mathbf{f}^{(k)} \leftarrow \mathrm{prox}^\epsilon(\mathbf{b}, \lambda_c, \lambda_d)$
  $t^{(k)} \leftarrow \frac{1+\sqrt{1+(4t^{(k-1)})^2}}{2}$
  $\mathbf{f}^{(k+1)} \leftarrow \mathbf{f}^{(k)} + \left(\frac{t^{(k-1)}-1}{t^{(k)}}\right)\left(\mathbf{f}^{(k)} - \mathbf{f}^{(k-1)}\right)$
until stopping criterion is met
return $\mathbf{f}_{\mathrm{opt}} \leftarrow \mathbf{f}^{(k+1)}$

3. EXPERIMENTS AND RESULTS

In our vector-field denoising experiments, we consider a simulated 3D phantom model consisting of a fully developed laminar flow in a tube that is encircled by constant flow inside a torus (see Figure 1). The measurements are degraded by additive white Gaussian noise so that the input SNR is 0 dB. To assess the efficieny of the proposed denoising scheme, we compare its performance with that of the IRLS method of [3] with 300 conjugate-gradient (CG) iterations. Both methods are allowed to run for approximately 120s CPU time and results are reported in Table 1. As is clear from Table 1, the new dual scheme results in a significant SNR improvement and a lower objective value. We also illustrate visually the outputs in Figure 1 using ParaView 3.12.0.

To have a point of comparison in terms of denoising performance, we also implemented the quadratic equivalent of the proposed algorithm (with the L_1 norms replaced by L_2 norms). The solution to this quadratic (least-squares) scheme are found by inverting a linear system using the CG method. Regularization parameters are optimized using an oracle to achieve the smallest-possible MSE for both regularizations. The results are reported in Table 2. As expected, L_1 regularization proves to be superior in terms of SNR for the particular flow model that we considered.

Table 1. Comparison of denoising algorithms for same CPU time.

Algorithm	Objective value	SNR improvement [dB]
Proposed method	1.95×10^4	12.87
IRLS	2.24×10^4	9.19

Finally, we reconstruct a 2D wind flow profile provided by MATLAB, which involves missing data on a uniform grid. Note that, in the 2D setting, we have to modify the formulation of the problem and redefine the curl and divergence operators, their adjoints, and the bounds on the Lipschitz constants. Again, the regularization parameters, λ_c and λ_d, were tuned using an oracle so as to get the smallest-possible MSE. For the $\mathrm{prox}^\epsilon(\cdot, \lambda_c, \lambda_d)$ step in Algorithm 2, 20 denoising iterations are used. We illustrate the results of the experiment in Figure 2.

All of the experiments are carried out on a Mac Pro with a 4×2.66 GHz processor and the algorithms were implemented in MATLAB.

Table 2. Comparison of denoising performances.

Regularization	input SNR [dB]	SNR improvement [dB]
L_1	0	12.87
L_2	0	8.64
L_1	10	10.09
L_2	10	3.79

4. CONCLUSION

We considered denoising and reconstruction of flow fields from missing data by means of regularized variational reconstruction. The regularizers were chosen so as to make the class of algorithms invariant to transformations of the coordinate system such as rotation and scaling. For efficient reconstruction, we reformulated the denoising problem in its equivalent dual form, to which we were then able to apply the proximal gradient algorithm of Beck and Teboulle. In phantom experiments, we found that our dual-based scheme is more efficient in terms of objective minimization and reconstruction performance than the IRLS approach for the denoising problem. We then showed that the reconstruction of a flow field from incomplete measurements can be implemented by solving successive denoising problems.

Some directions for future investigation include flow reconstruction from incomplete pulsed-wave ultrasound Doppler data where the sampling is not uniform, temporal regularization, optical flow estimation, and resolution enhancement in 4D flow-sensitive MRI.

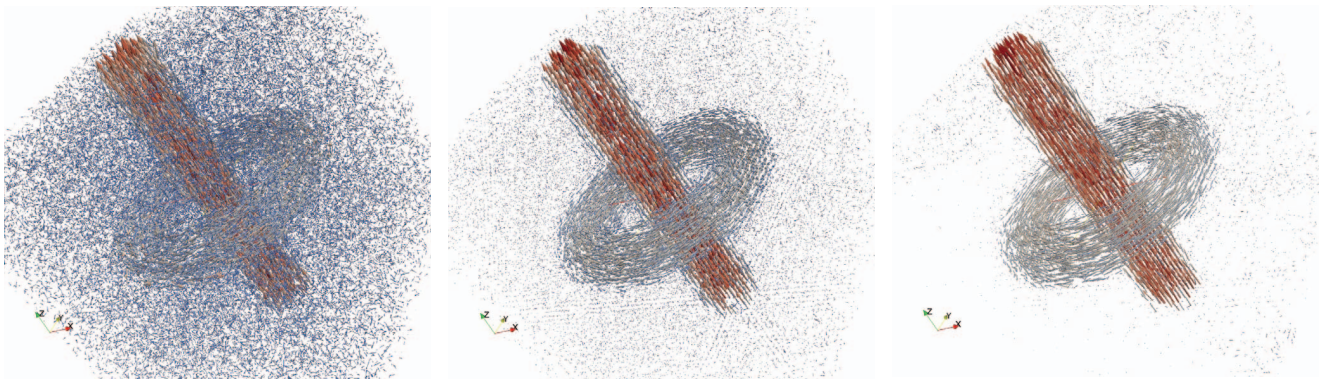

(a) (b) (c)

Fig. 1. Denoising of a 3D vector field after approximately 120s CPU time: (a) noisy vector field, SNR= 0 dB; (b) denoised vector field with IRLS, SNR = 9.19 dB; (c) denoised vector field with the proposed scheme, SNR = 12.87 dB.

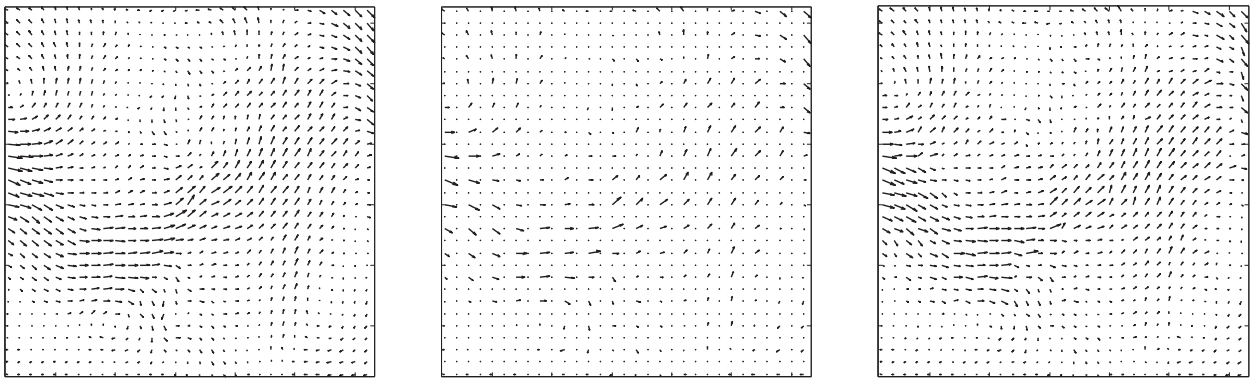

(a) (b) (c)

Fig. 2. Reconstruction of a 2D vector field from incomplete measurements: (a) original vector field; (b) incomplete measurements; (c) reconstructed vector field, ISNR = 8.49 dB.

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," *Magnetic Resonance in Medicine*, 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," *Journal of Cardiovascular Magnetic Resonance*, vol. 10, no. 1, pp. 30, 2008.

[3] P. D. Tafti and M. Unser, "On regularized reconstruction of vector fields," *IEEE Transactions on Image Processing*, vol. 20, no. 11, pp. 3163–3178, 2011.

[4] S. David and C. Fang, "Left ventricular motion reconstruction based on elastic vector splines," *IEEE Transactions on Medical Imaging*, vol. 19, no. 4, pp. 295–305, 2000.

[5] F. Dodu and C. Rabut, "Irrotational or divergence-free interpolation," *Numerische Mathematik*, vol. 98, no. 3, pp. 477–498, 2004.

[6] M. Arigovindan, M. Sühling, C. Jansen, P. Hunziker, and M. Unser, "Full motion and flow field recovery from echo Doppler data," *IEEE Transactions on Medical Imaging*, vol. 26, no. 1, pp. 31–45, 2007.

[7] P. D. Tafti, R. Delgado-Gonzalo, A. F. Stalder, and M. Unser, "Variational enhancement and denoising of flow field images," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2011, pp. 1061–1064.

[8] A. Chambolle, "An algorithm for total variation minimization and applications," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1–2, pp. 89–97, 2004.

[9] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 185–212. Springer-Verlag, 2011.

[10] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Transactions on Image Processing*, vol. 18, no. 11, pp. 2419–2434, 2009.