Total Variation Regularization of Displacements in Parametric Image Registration

Valeriy Vishnevskiy, Tobias Gass, Gábor Székely, and Orcun Goksel

Computer Vision Laboratory, ETH Zurich, Switzerland
valery.vishnevskiy@inf.ethz.ch
{gass,szekely,ogoksel}@vision.ee.ethz.ch

Abstract. Spatial regularization is indispensable in image registration to avoid both physically implausible displacement fields and potential local minima in optimization methods. Typical ℓ_2-regularization is incapable of correctly recovering non-smooth displacement fields, such as at sliding organ boundaries during time-series of breathing motion. In this paper, Total Variation (TV) regularization is used to allow for accurate registration near such boundaries. We propose a novel formulation of TV-regularization for *parametric* displacement fields and introduce an efficient and general numerical solution scheme using the Alternating Directions Method of Multipliers (ADMM). Our method has been evaluated on two public datasets of 4D CT lung images as well as a dataset of 4D MR liver images, demonstrating accurate registrations both inside and outside moving organs. The target registration error of our method is 2.56 mm on average in the liver dataset, which indicates an improvement of over 24% in comparison to other published methods.

Keywords: Medical image registration, Total Variation, 4D CT, ADMM

1 Introduction

Image registration is an essential part of several applications in modern medical imaging, such as atlas-based segmentation, volumetric reconstruction from slices, tissue parameter estimation and motion tracking. All these applications rely on correctly estimating spatial anatomical correspondences between images. The existence of such correspondences and an underlying bijective mapping cannot be guaranteed in some registration contexts, such as inter-patient or 2D registration. However, in other contexts where the anatomy changes minimally, for example, in intra-patient registration of time-series 3D (4D) volumes, one-to-one correspondences do exist between images when the image margins that may move inside/outside the field-of-view are neglected. Such 4D motion, however, often involves sliding between different parts of anatomy, such as the liver or the lungs over their surrounding (cf. Fig. 1). The resulting non-smooth motion fields are difficult to represent and estimate using typical ℓ_2-regularized registration techniques.

A simple way to accommodate for discontinuities near sliding boundaries is to use binary masks for objects to be registered, such that deformations (or image

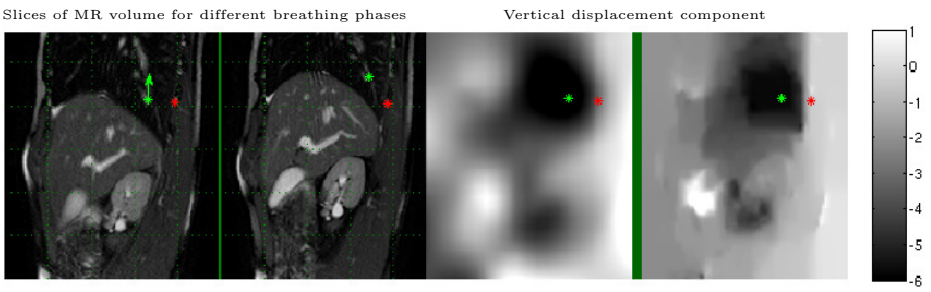

Fig. 1. Sliding motion example from a 4D MRI liver sequence during breathing. Left: Displacements for landmarks inside (green) and outside (red) the rib cage are visualized. Right: Inferior-superior component of displacement fields estimated by the smooth ℓ_2-regularized (left) and our TV-regularized (right) registration techniques. Note that the sliding organ boundaries are captured better by our TV-regularized method.

similarity) outside object's masks can be ignored. For motion estimation inside the lungs, masks were identified in [9] using automatic lung segmentation from 4D CT. To guarantee physically possible motions along the mask surface, in [4] two separate displacement fields were estimated simultaneously with additional constraints. Nevertheless, motion masking has several disadvantages: They require an initial segmentation stage that is cumbersome when done manually and error-prone when automatic. Furthermore, sliding can *only* be recovered at the *mask interface*, making registration efficacy dependent on the hard decision made at the prior masking stage. This problem was addressed by Kiriyanthan and Cattin in [10], where they have proposed to estimate a motion mask automatically *during* the registration process using the segmentation model of Chan et al [3].

Another way to allow for sliding motion is to model motion trajectories in the time domain using intermediate images between breathing phases [2]. This approach necessitates simultaneous processing of several image pairs, which imposes severe time and memory limitations. Also, heuristic techniques exists that allow for sliding motion, such as the modified demons method [13], [11] or the anisotropic demons [17], in which the Gaussian smoothing step is replaced with anisotropic filtering. However, such methods do not have an explicit definition of regularization and hence lack a formal cost definition. Therefore, they do not allow for a coherent optimization scheme. Total Variation (TV), in contrast, can be defined as a penalty and hence the registration can be solved in a well-defined optimization framework as shown in this work.

Constraining spatial Total Variation penalizes spatial incoherence, but does not restrict the displacement field to be smooth. This property of the TV norm as a regularizer received much attention in the computer vision literature for correctly estimating optical flow between scenes with independent motion of overlapping objects [16]. A major difficulty with TV is the involvement of ℓ_1-norm, which is non-differentiable at zero leading to numerical instabilities for gradient-based methods. In [16], the use of its smooth seminorm approximations, i.e. $\|x\|_1 \approx \varphi_\varepsilon(x) = \sqrt{x^2 + \varepsilon}$, was proposed. However, for small ε, common

optimization algorithms easily become unstable. It is possible to employ discrete optimization methods, formulating the registration problem on Markov random fields [6]. This approach was applied for breathing motion estimation in [7], however commonly-used message passing and graph-cut based algorithms are inefficient for registering 3D volumes on fine grid resolutions. Fine displacement estimation requires a dense discretization of the displacement search space, which is extremely memory consuming for 3D data.

A recent approach for solving ℓ_1-norm problems is based on duality, which was applied to *non-parametric* medical image registration in [12] using the sum of square differences (SSD) metric. A census cost function was also used as image residual yielding improved results for 4D CT lung images with breathing motion [8]. However, such non-parametric pixel-level approaches easily lead to physically implausible motion fields, for which the authors proposed a combined median and Gaussian filtering of the displacement fields following each method iteration. Such heuristic regularization may lead to unstable or suboptimal solution schemes as mentioned above. Furthermore, non-parametric registration methods are highly susceptible to local-minima and not robust during optimization. As a remedy, a complex image metric based on local image statistics was proposed in [8], leading to time-ineffective implementations with still no theoretical guarantees on the validity of the solution. Due to these and other difficulties, most registration techniques use parametric deformation models such as B-splines, which offer physically plausible displacements in robust schemes with large displacement capture ranges. In this work, we present a *parametric* image registration approach formulated as a minimization problem with TV-regularization, to recover anatomical non-smooth (sliding) motion in a coherent optimization framework. To accommodate the ℓ_1-norm cost of TV, we introduce an efficient solution scheme using the *alternating directions method of multipliers* (ADMM).

1.1 Image Registration as Energy Minimization

Estimating the D-dimensional local transformation field $\mathbf{t} : \Omega \to \mathbb{R}^D$, that maps image $I_\mathrm{m} : \Omega \to \mathbb{R}$ to image $I_\mathrm{f} : \Omega \to \mathbb{R}$, is commonly formulated as the following optimization problem:

$$\min_{\mathbf{t}} \; E_\mathrm{D}(\mathbf{t};\, I_\mathrm{f}, I_\mathrm{m}) + \lambda E_\mathrm{R}(\mathbf{t}), \tag{1}$$

where Ω is the image domain, E_D is an image dissimilarity metric and E_R is a regularization term while λ controls the amount of regularization. Typically, E_D is a smooth metric such as the SSD, which allows for easy differentiation. For regularization, a popular choice is to component-wise penalize the ℓ_2-norm of displacement field derivatives, which leads to differentiable functionals [14]. Penalizing the ℓ_2-norm of the first derivative leads to the well-known *smoothness* regularization: $E_\mathrm{R}^\mathrm{smooth}(\mathbf{t}) = \sum_{d=1}^{D} \sum_{\mathbf{x}\in\Omega} \|\nabla t_d(\mathbf{x})\|_2^2 \delta\mathbf{x}$. Any gradient-based solver can then be used for minimizing (1) by computing the derivatives of the smooth metric E_D and smooth regularizer E_R with respect to the transformation parameters.

1.2 TV-Regularization

Anisotropic TV-regularization can be written as the ℓ_1-norm of the finite differences of displacement field components as follows:

$$E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{t}) = \sum_{d=1}^{D} \sum_{\mathbf{x}\in\Omega} \|\nabla t_d(\mathbf{x})\|_1 \; \delta\mathbf{x}, \tag{2}$$

where ∇ is a linear operator for finite differences, and $\delta\mathbf{x}$ is the volume of a voxel. This functional is convex but no longer smooth, i.e. it is non-differentiable with respect to $\mathbf{x}$ at $\|\nabla t(\mathbf{x})\|_1 = 0$. This leads to poor performance when using standard gradient-based optimization methods. This is addressed in our work by employing ADMM optimization [1, 5], described later below. First, we introduce our displacement parametrization approach which allows us to impose TV regularization on the displacement field simply and efficiently.

1.3 Parameterizing Displacement Fields

Parameterizing a displacement field greatly reduces the dimensionality of the registration optimization problem. This constrains the size of parameter search space, while discarding many physically meaningless transformations. This allows for robust numerical methods that are more resilient to suboptimal local minima.

We parametrize the displacement field $\mathbf{t}$ using trilinear interpolation (1st order B-splines) with displacements $\mathbf{k}$ on control points that are placed with N-pixel spacing, i.e. $\mathbf{t} = \mathbf{t}(\mathbf{k})$, where $\mathbf{k}$ and $\mathbf{t}$ are displacements stacked into column vectors. Unlike higher order B-splines, linear interpolation guarantees that the interpolated displacement values at grid knots are equal to the corresponding B-spline coefficients. We then impose regularization on the displacement control grid points $\mathbf{k}$ instead of the displacement field $\mathbf{t}$ itself. Such approximation is commonly used in smooth registration methods, where it was reported in [15] to be practically equivalent to imposing constraints directly on the displacement field itself.

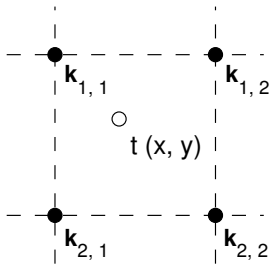

Fig. 2. Control grid points.

Parametric Upper Bound for TV. For linear interpolation in the 1D case, it is obvious that the TV of both the control grid and the underlying displacement field are exactly the same, i.e. $E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{t}(\mathbf{k}))=E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{k})$. For 2D and 3D cases this equality does not necessarily hold. However, an upper bound on such displacement field TV given the TV on the control grid can be computed which we show as given below.

For simplicity, we consider the TV of a single component of a 2D displacement field within a single patch shown in Fig. 2. The TV of the patch control points $\mathbf{k}$ and the field $\mathbf{t}$ are then, respectively:

$$E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{k}) = \left(|k(1,2) - k(1,1)| + |k(2,2) - k(2,1)|\right) N, \tag{3}$$

$$\begin{aligned} E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{t}) = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \Big(& \left|t\left(\tfrac{i+1}{N}, \tfrac{j}{N}\right) - t\left(\tfrac{i}{N}, \tfrac{j}{N}\right)\right| \\ & + \left|t\left(\tfrac{i}{N}, \tfrac{j+1}{N}\right) - t\left(\tfrac{i}{N}, \tfrac{j}{N}\right)\right| \Big). \end{aligned} \tag{4}$$

Using the fact that t is linearly interpolated from $\mathbf{k}$ and employing the triangle inequality, we can infer that:

$$E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{t}(\mathbf{k})) - E_{\mathrm{R}}^{\mathrm{TV}}(\mathbf{k}) \leq N \max(|\nabla \mathbf{k}|) \ . \tag{5}$$

In other words, the TV approximated on the control grid is an upper bound for TV on $\mathbf{t}$, and such approximation error decreases linearly as the parametrization grid is refined. Accordingly, we start with a coarse grid and refine it gradually by initializing the optimization at each step from upsampled lower-resolution grid.

1.4 Numerical Scheme: ADMM

As was mentioned in the introduction, accurate and robust minimization of the parametrized TV-regularized energy (1) is difficult because the gradient required to *simultaneously* minimize both the data term E_{D} and the regularization term $E_{\mathrm{R}}^{\mathrm{TV}}$ does not always exist. To overcome this challenge, we utilize the alternate direction method of multipliers (ADMM), which allows us to iteratively minimize each of the two terms separately, while constraining their solutions to be similar.

To this end, we introduce a redundant variable $\mathbf{z}$ to (1), which leads to the following equivalent problem:

$$\min_{\mathbf{k},\mathbf{z}} \ \ E_{\mathrm{D}}(\mathbf{t}(\mathbf{k});\, I_{\mathrm{f}}, I_{\mathrm{m}}) + \lambda \|\mathbf{z}\|_1 \ , \qquad \text{s.t.} \quad \nabla \mathbf{k} = \mathbf{z}. \tag{6}$$

Writing the augmented scaled Lagrangian for the problem above and performing the dual descent method [1] results in the following iterative ADMM scheme, where the optimization of E_{D} and ℓ_1-norm terms are decoupled into two separate subproblems.

$$\mathbf{k}^{j+1} = \arg\min_{\mathbf{k}} \ \ E_{\mathrm{D}}(\mathbf{t}(\mathbf{k});\, I_{\mathrm{f}}, I_{\mathrm{m}}) + \frac{\rho}{2} \|\nabla \mathbf{k} - \mathbf{z}^j + \mathbf{u}^j\|_2^2, \tag{7a}$$

$$\mathbf{z}_d^{j+1} = \arg\min_{\mathbf{z}_d} \ \ \frac{\lambda}{\rho} \|\mathbf{z}_d\|_1 + \frac{1}{2} \|\nabla \mathbf{k}_d^{j+1} + \mathbf{u}_d^j - \mathbf{z}_d\|_2^2, \quad d = 1, \ldots, D, \tag{7b}$$

$$\mathbf{u}^{j+1} = \mathbf{u}^j + \mathbf{k}^{j+1} - \mathbf{z}^{j+1}. \tag{7c}$$

The $\mathbf{k}$-update step (7a) finds control grid displacements that optimize the image similarity metric. It is a smooth optimization problem that can be solved using

any gradient-based optimization technique by adding $\rho\nabla^\top(\nabla\mathbf{k}-\mathbf{z}^j+\mathbf{u}^j)$ to the gradient of the similarity metric. The penalty parameter ρ is usually set to 1 and can be updated heuristically to accelerate convergence [1]. $\mathbf{u}$ is a scaled dual variable. The $\mathbf{z}$-update step (7b) is called the proximity operator of the ℓ_1 norm and can be solved by element-wise shrinkage:

$$\mathbf{z}_d^{j+1} = \max\{\mathbf{u}^j + \mathbf{k}^{j+1} - \lambda/\rho, 0\} - \max\{-\mathbf{u}^j - \mathbf{k}^{j+1} - \lambda/\rho, 0\}. \qquad (8)$$

Implementation. Input image intensities are first scaled to the [0, 1] interval. Gaussian pyramids are used by starting the registration with both downsampled images and a coarse control grid. Registration at each consecutive image level is then initialized by interpolating from the previous level's control grid displacements. When the finest level of the image pyramid is reached, we start to subdivide the control grid, effectively decreasing the pixel span of each control patch and hence refining the estimated displacement resolution. We use SSD for CT images and normalized cross correlation for MRI as similarity metrics.

The method was implemented in `Matlab` with parallelized `mex`-functions for warping images and calculating E_D gradient. The $\mathbf{k}$-update step in (7a) is solved with the limited-memory BFGS method of the `minFunc`[1] package with the maximum number of iterations set to 10. Variables $\mathbf{x}$, $\mathbf{z}$, $\mathbf{u}$ were initialized with zeros. The algorithm was executed on a 6-core Intel Xeon 2.4 GHz processor.

2 Results and Discussion

We evaluated our *parametric TV* registration (pTV) method on the following three abdominal time-series datasets that involve breathing motion:

4D-CT POPI Dataset. The public POPI[2] dataset [18] consists of: 10 3D CT images of different phases of one breathing cycle; 41 corresponding landmarks for each of the images; a binary lung mask; and the deformation fields estimated by the POPI organizers using both the standard demons algorithm and the free-form deformation method (FFD). We resampled all images to an isotropic $2\times2\times2\,\mathrm{mm}^3$ pixel resolution, resulting in $235\times176\times141$ voxels/image. We then registered the first image (phase) to all other 9 images as in [18], resulting in 9

[1] `http://www.di.ens.fr/~mschmidt/Software/minFunc.html`
[2] `http://www.creatis.insa-lyon.fr/rio/popi-model`

Table 1. Mean TRE [mm] over 41 landmarks for the 9 POPI registrations.

	TRE for each image pair									**Average**
Method	#1	#2	#3	#4	#5	#6	#7	#8	#9	**TRE**
Demons	1.28	1.38	1.39	1.22	1.24	1.25	1.29	1.12	1.11	1.25
FFD	0.79	0.80	1.13	1.11	**1.10**	1.20	1.20	0.88	0.92	1.01
pTV	**0.72**	**0.71**	**1.12**	**1.01**	1.11	**1.03**	**1.06**	**0.84**	**0.81**	**0.93**

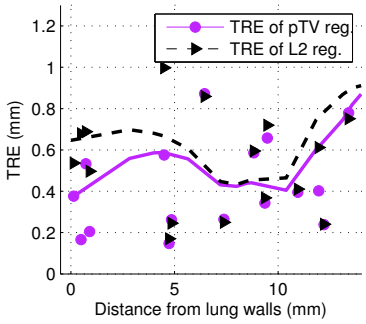

Fig. 3. Registration spatial accuracy with respect to target distance from the lung wall. Local regression curves are also presented for illustration purposes.

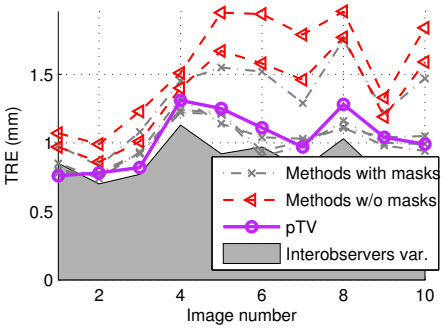

Fig. 4. Our proposed method pTV compared to the top 6 out of 20 results, published and reported on DIR website.

Table 2. Mean (snap-to-pixel) TRE on DIR data set, where cTV is the non-parametric TV method [8] with ℓ_1-norm census cost function and cTVmask is the same method with motion masks provided.

	TRE for each 4D sequence										
Method	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	**Mean**
cTVmask [8]	0.78	0.78	0.93	**1.24**	**1.22**	**0.94**	1.01	**1.11**	**0.98**	**0.94**	**0.99**
cTV [8]	0.79	0.80	1.02	1.23	1.27	1.09	1.87	3.01	1.11	1.17	1.37
pTV	**0.76**	**0.78**	**0.82**	1.31	1.25	1.11	**0.97**	1.28	1.04	0.99	1.03

individual registrations. The regularization parameter λ was set to 0.001. The average run-time of our registration method was 30 seconds per image pair.

Table 1 shows the target registration error (TRE) of our pTV algorithm in comparison to Demons and FFD results reported in [18]. It can be observed that our method achieves an improved registration accuracy overall. One source of such improvement is the replacement of ℓ_2-regularization of FFDs, which is unable to capture sliding motion near the lung wall. This is is demonstrated in Fig. 3 on one image pair for targets within 15 mm of the lung wall. Our method pTV is seen to have significantly lower TRE for targets closer to the lung wall.

4D-CT DIR Dataset. The public DIR dataset[3] contains 10 4D CT sequences from different individuals, each sequence with 10 images/breathing-phases with an average resolution of $1\times1\times2.5\,\text{mm}^3$ and sizes varying between $256\times256\times94$ and $512\times512\times136$ voxels (see Fig. 5). On this dataset, we empirically set λ to 0.001 and the finest control grid spacing to 3 voxels. This led to an average run-time of 3 minutes for pTV. To evaluate our method, from each sequence we used the images of extreme inhale and exhale breathing phases, which are annotated with 300 landmarks inside the lungs.

[3] `http://www.dir-lab.com`

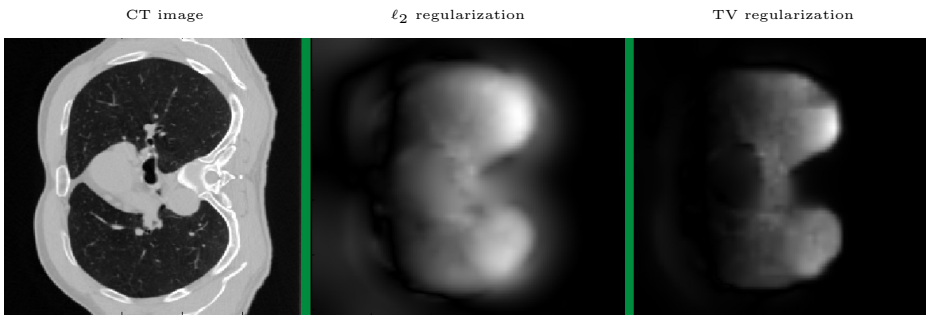

Fig. 5. An axial slice of a CT image from DIR dataset and the corresponding inferior-superior component of deformation fields for different regularization techniques. Note how ℓ_2 regularization oversmoothes deformations near lung borders at the sliding interface.

DIR is an ongoing benchmark, where TRE is computed by the organizers from submitted participant registration results (using a snap-to-voxel fashion for consistency with annotations). Those results are published online, allowing for a relative assessment of the performance of our pTV algorithm. Fig. 4 shows the average TRE for each sequence, comparing pTV to the top 6 published methods (out of 20 in total at the time of evaluation) from DIR. Note that most DIR submissions are customized particularly for lung (pulmonary CT) registration, and several of them use lung masks to recover motion *only* inside the lungs. This is depicted in Fig. 4 by representing methods using masks with different markers. Despite not using any customization particular for the lungs, our general image registration method pTV yields comparable results to the other techniques. Furthermore, pTV outperforms all other methods that do not utilize a lung mask, despite the relatively small margin for consistent improvement over the reported inter-observer variability. A slice of the estimated displacement field is seen in Fig. 5 to better capture sliding motion.

We also compare pTV with the non-parametric TV-regularized registration method of [8] that uses a *census* cost function. Two results with and without using lung masks are reported in [8], namely *cTV* and *cTVmask*, with which we compare our pTV results in Table 2. A 24% reduction in TRE can be observed when comparing pTV to the cTV version without mask.

4D-MRI liver sequences. Since the above lung datasets have landmarks only inside the lungs, we also evaluated our method on the dataset from [17]. Therein, landmarks are provided on both sides of the sliding liver interface, allowing for improved assessment of sliding-motion registration algorithms. 32 and 20 landmarks were provided inside and outside the liver, respectively, in the extreme breathing phases of 8 4D-MR sequences, where two sequences each were recorded from 4 volunteers. The images have $1.37\times1.37\times4\,\mathrm{mm}^3$ resolution and an average size of $164\times189\times23$ voxels, for which pTV required an average run-time of 50 s. Table 3 shows substantial improvement of pTV (19% inside and 35% outside the liver) in comparison to the results reported in [17].

Table 3. Average mean TRE [mm] for the liver dataset. For e_i^j being the TRE for landmark i in sequence j, each cell reports: $\mathbf{mean_j\{mean_i\, e_i^j\}}$ $(\mathbf{mean_j\{max_i\, e_i^j\}})$. The proposed method pTV (with $\lambda = 0.0005$) is compared to regular and anisotropic demons registrations.

Method	Inside the liver	Outside the liver	Overall
Demons	3.57 (12.87)	4.25 (14.38)	3.83 (15.77)
Aniso. Demons [17]	3.00 (11.86)	4.05 (13.15)	3.40 (14.03)
pTV	**2.42** (11.0)	**2.64** (7.52)	**2.56** (9.78)

3 Conclusions

In this paper we have presented a parametric approach for image registration with total variation regularization. We have evaluated our method on datasets of different imaging modalities and anatomies. Our method was shown to accurately estimate anatomical displacements near breathing-induced sliding boundaries. For thoracic 4D CT images, the registration accuracy of the proposed algorithm was shown to be comparable to the best methods from the DIR dataset, which rely on lung segmentation masks, and to outperform the best *mask-free* methods published to date. For 4D breathing motion estimation, our method was also shown to outperform both the non-parametric TV-regularized registration method of [8] and the anisotropic smoothing method of [17]; with an overall reduction of the TRE by 24%. Furthermore, our proposed approach allows for using any registration method based on energy minimization for solving the $\mathbf{k}$ update step in the ADMM solver. This can facilitate adapting the algorithm to specific image modalities or existing workflows in clinical practice, which is an interesting direction for further research.

References

1. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends in Machine Learning 3(1), 1–122 (2011)
2. Castillo, E., Castillo, R., Martinez, J., Shenoy, M., Guerrero, T.: Four-dimensional deformable image registration using trajectory modeling. Phys Med Biol 55(1), 305–27 (2010)
3. Chan, T.F., Esedoglu, S., Nikolova, M.: Algorithms for finding global minimizers of image segmentation and denoising models. SIAM Journal on Applied Mathematics 66(5), 1632–1648 (2006)
4. Delmon, V., Rit, S., Pinho, R., Sarrut, D.: Registration of sliding objects using direction dependent B-splines decomposition. Phys Med Biol 58(5), 1303–14 (2013)
5. Figueiredo, M.A., Bioucas-Dias, J.M.: Algorithms for imaging inverse problems under sparsity regularization. In: IEEE Int Workshop on Cognitive Information Processing (CIP). pp. 1–6 (2012)
6. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. Medical Image Analysis 12(6), 731–741 (2008)

7. Heinrich, H., Jenkinson, M., Brady, M., Schnabel, J.A.: MRF-based deformable registration and ventilation estimation of lung CT. IEEE T Med Imag 32(7), 1239–1248 (2013)
8. Hermann, S., Werner, R.: TV-L1-based 3D medical image registration with the census cost function. In: Image and Video Technology, pp. 149–61 (2014)
9. Hu, S., Hoffman, E.A., Reinhardt, J.M.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. IEEE T Med Imag 20(6), 490–8 (2001)
10. Kiriyanthan, S., Fundana, K., Cattin, P.C.: Discontinuity preserving registration of abdominal MR images with apparent sliding organ motion. In: MICCAI Workshop on Abdominal Imaging. Computational and Clinical Applications, pp. 231–239. Springer (2012)
11. Papież, B.W., Heinrich, M.P., Risser, L., Schnabel, J.A.: Complex lung motion estimation via adaptive bilateral filtering of the deformation field. In: MICCAI 2013, pp. 25–32. Springer (2013)
12. Pock, T., Urschler, M., Zach, C., Beichel, R., Bischof, H.: A duality based algorithm for TV-L1-optical-flow image registration. In: MICCAI, pp. 511–8 (2007)
13. Risser, L., Vialard, F.X., Baluwala, H.Y., Schnabel, J.A.: Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions. Medical Image Analysis 17(2), 182–193 (2013)
14. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast MR images. IEEE T Med Imag 18(8), 712–21 (1999)
15. Schwarz, L.A.: Non-rigid registration using free-form deformations. Ph.D. thesis, Technische Universität München, Germany (2007)
16. Sun, D., Roth, S., Black, M.J.: Secrets of optical flow estimation and their principles. In: IEEE Conf Comp Vis and Pattern Recog (CVPR). pp. 2432–9 (2010)
17. Tanner, C., Samei, G., Székely, G.: Investigating anisotropic diffusion for the registration of abdominal MR images. In: IEEE Int Symp Biomedical Imaging (ISBI). pp. 484–7 (2013)
18. Vandemeulebroucke, J., Sarrut, D., Clarysse, P., et al.: The POPI-model, a point-validated pixel-based breathing thorax model. In: Int Conf Computers in Radiation Therapy (ICCR). pp. 195–9 (2007)