

Directional Averages for Motion Segmentation in Discontinuity Preserving Image Registration

Christoph Jud^(✉), Robin Sandkühler, Nadia Möri, and Philippe C. Cattin

Department of Biomedical Engineering, University of Basel, Basel, Switzerland
christoph.jud@unibas.ch

Abstract. The registration of abdominal images is central in the analysis of motion patterns and physiological investigations of abdominal organs. Challenges which arise in this context are discontinuous changes in correspondence across sliding organ boundaries. Standard regularity criteria like smoothness, are not valid in such regions. In this paper, we introduce a novel regularity criterion which incorporates local motion segmentation in order to preserve discontinuous changes in the spatial mapping. Based on local directional statistics of the transformation parameters it is decided which part of a local neighborhood influences a parameter during registration. Thus, the mutual influence of neighboring parameters which are located on opposing sides of sliding organ boundaries is relaxed. The motion segmentation is performed within the regularizer as well as in the image similarity measure and is thus implicitly updated throughout the optimization. In the experiments on the 4DCT POPI dataset we achieve competitive registration performance compared to state-of-the-art methods.

Keywords: Image registration · Regularization · Motion segmentation

1 Introduction

Abdominal image registration is an active field of research with many applications such as the analysis of respiratory dynamics or the physiology of abdominal organs as for example the lung. The particular challenges which arise in this scenario are discontinuous changes in the correspondence that occur between organs sliding along each other. Standard smoothness regularity is thus inappropriate at the boundaries between the sliding organs.

In this paper, we present a parametric image registration method comprising a novel regularity criterion. The idea is that the spatial mapping is ought to be locally homogeneous while this criterion should be relaxed at sliding organ boundaries. With local homogeneity we mean that a displacement at a particular image location should be aligned to the average displacement in its neighborhood. To cope with sliding interfaces, we replace the average displacement by a *directional average* where only displacements which are aligned up to a certain degree are considered (see Fig. 1). The proposed *directional average* performs a

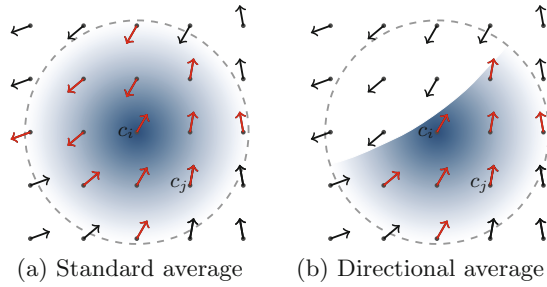


Fig. 1. The dashed circles mark the neighborhood of c_i while the radial shading visualizes the amount of influence a displacement has to the average. In the standard average (left), all displacements in the neighborhood, colored in red, are taken into account while in the directional average (right), only displacements which are approximately aligned to c_i are considered.

motion segmentation for each parameter separately. The homogeneity requirement for a certain parameter therefore only extends to an aligned neighborhood.

There are major trends which target the sliding organ problem. In [10, 14], image masks of parts in the image sliding along each other are used to register them independently. However, the expected sliding interfaces have to be known in advance. Various approaches [1, 3, 6, 7, 12] adapt their regularity criterion to local image features, as for example image gradients, in order to reduce the influence of the regularity across sliding organ boundaries. The assumption here is that sliding interfaces occur where high intensity changes in the images are present. This applies well to the interfaces between the lung and the thoracic cavity for example. However, it does not hold for generic sliding interfaces. In [5, 9], this problem is addressed by motion segmentation, which is performed exclusively on the spatial transformation. The generic recognized sliding interfaces are thus refined during the optimization. The motion segmentation and the image registration are separately formalized and intertwined by alternating optimization. To achieve a stable convergence remains, however, challenging. In [4, 13, 16], stationary sparse regularity is applied without explicitly considering sliding interfaces. The regularity only amounts to correct the image similarity term at sliding organ boundaries.

The contribution of this paper is twofold. First, we formulate a novel regularity criterion which is based on motion segmentation which is exclusively performed on the transform parameters and thus is implicitly integrated in the registration objective. Hence, it obviates alternated optimization schemes. Second, we modify the image similarity such that the dependency of gradients on opposing sides of sliding interfaces is relaxed. In the experiments on a 4DCT dataset, we compare our method to state-of-the-art registration methods and show competitive registration results.

2 Background

We recap the kernel-framework for image registration on which we base our method. It was elaborated in [3,4] and we borrow the notation used therein. Let a reference and target image $I_R, I_T : \mathcal{X} \rightarrow \mathbb{R}$ map the d -dimensional input domain $\mathcal{X} \subset \mathbb{R}^d$ to intensity values and a spatial mapping $f : \mathcal{X} \rightarrow \mathbb{R}^d$ transform the reference coordinate system. Image registration is performed by optimizing

$$\arg \min_u \int_{\mathcal{X}} \mathcal{L}(I_R(x + f(x)), I_T(x)) dx + \eta \mathcal{R}[f], \quad (1)$$

where \mathcal{L} is a loss-function and \mathcal{R} is a regularization term with η as trade-off parameter. As transformation model a reproducing kernel Hilbert space (RKHS) is defined

$$\mathcal{H} := \left\{ f \mid f(x) = \sum_{i=1}^{\infty} c_i k(x_i, x), \quad x_i \in \mathcal{X}, \quad c_i \in \mathbb{R}^d, \quad \|f\|_{\mathcal{H}} < \infty \right\}, \quad (2)$$

where $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a reproducing kernel and $\|\cdot\|_{\mathcal{H}}$ is the RKHS norm. For more details about kernel methods we refer to [2]. In [3], the existence of a finite dimensional solution to Eq. 1 with N pair-wise distinct sampled domain points x_i was shown applying a regularization term operating solely on the finite many transform parameters $c := \{c_i\}_{i=1}^N$

$$\mathcal{R}[f] := g(p(c)), \quad (3)$$

where $g : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly increasing function and $p : \mathbb{R}^{Nd} \rightarrow \mathbb{R}$ is weakly semi-continuous and bounded from below. As k , the combined Wendland kernel of [3] is used. Note, that in a homogeneous region the direction of a transform parameter is similar to an actual displacement if a radial basis function is used as kernel.

3 Method

Within organ regions, the transformation should be locally homogeneous while discontinuities across sliding interfaces are required. Thus, neighboring parameters should point in similar direction to fulfill the local homogeneity requirement. However, at a certain *misalignment* this similarity requirement should be relaxed. The assumption is that such misalignments appear for neighboring parameters which are located on opposing sides of sliding interfaces. As a criterion for the misalignment, we propose that the angle

$$\omega(c_i, c_j) = \cos^{-1} \left(\frac{c_i^T c_j}{\|c_i\|_{\epsilon} \|c_j\|_{\epsilon}} \right) \quad (4)$$

between the average direction of different sliding regions has to exceed a certain threshold θ , where the ϵ -norm¹ $\|\cdot\|_{\epsilon} = \sqrt{\|\cdot\|^2 + \epsilon}$. We specify a sigmoid function

¹ Although the ϵ -norm is not really a norm, we apply it because of numerical stability.

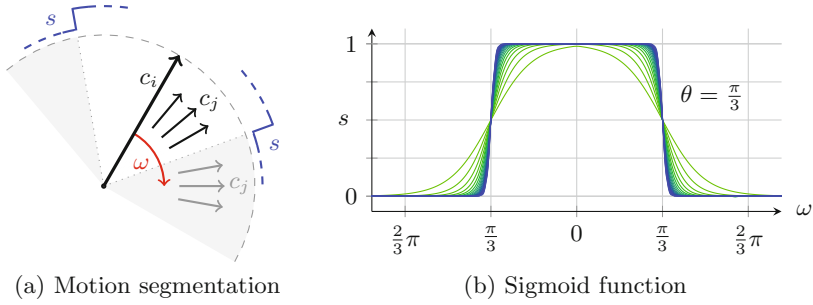


Fig. 2. The geometrical interpretation of s and ω (a). The angle between c_i and neighboring c_j is given by ω , while the sigmoid function s has the value 1 as long as the angle $\omega < \theta$. In (b), s is plotted with increasing slopes $\sigma = 1, 2, 4, 6, \dots, 20$.

which is one if the angle ω is small and zero if ω exceeds a threshold θ

$$s(\omega) = 1 - \frac{1}{2} (1 + \tanh(\sigma(\omega - \theta))), \tag{5}$$

where σ controls the slope of s (see Fig. 2).

3.1 Directional Average Regularizer

We define the directional average regularizer using Eq. 3 as

$$\mathcal{R}_{\text{da}}(f) := \sum_{i=1}^N \|c_i - \mu(c_i)\|_{\epsilon}. \tag{6}$$

The directional average μ for a certain parameter c_i becomes

$$\mu(c_i) = \frac{1}{Z(c_i)} \sum_{j=1, j \neq i}^N c_j k(x_i, x_j) s(\omega_{ij}), \quad Z(c_i) = \epsilon + \sum_{j=1, j \neq i}^N k(x_i, x_j) s(\omega_{ij}). \tag{7}$$

For convenience, we write ω_{ij} for $\omega(c_i, c_j)$. Actually, $\mu(c_i)$ is a weighted average over the remaining parameters c_j , where the kernel k serves as weighting function. As we use compact kernels, only *neighboring* parameters c_j are considered. The sigmoid function s additionally decides if a certain c_j contributes to the average dependent on the angle between c_i and c_j .

3.2 Directional Average Similarity Metric

With the directional average regularizer of Eq. 6 we specify what kind of parameter configurations should be preserved or penalized respectively. The problem is, that it is only an additive term in the overall registration objective in Eq. 1 and the similarity metric does not take the directional average into account.

This is critical because image similarity gradients across sliding interfaces influence a parameter update in the optimization. This can be verified when looking at the derivative of

$$\frac{\partial \mathcal{L}}{\partial c_i} = \frac{\partial \mathcal{L}}{\partial I_R(x + f(x))} \nabla I_R(x + f(x)) k(x_i, x) \quad (8)$$

where to the right side we differentiate \mathcal{L} with respect to its first argument. The gradient for a parameter c_i is influenced by all neighboring points x regardless of their alignment to c_i . We unravel now the strict distinction between similarity and regularization term and propose to modify the similarity metric as follows

$$\vec{\mathcal{D}}[I_R, I_T, f] := \int_{\mathcal{X}} \mathcal{L}(I_R(x + f(x)), I_T(x)) \mathcal{A}(f(x), c) dx. \quad (9)$$

We simplify notation with $f_x = f(x)$ and specify the factor \mathcal{A} as

$$\mathcal{A}(f_x, c) = 1 - \frac{1}{2} \left\| \frac{f_x}{\|f_x\|_\epsilon} - \frac{\mu_A(f_x, c)}{\|\mu_A(f_x, c)\|_\epsilon} \right\|, \quad (10)$$

where the average function μ_A is defined as

$$\mu_A(f_x, c) = \frac{\sum_j c_j k(x_j, x) s(\omega(f_x, c_j))}{\epsilon + \sum_j k(x_j, x) s(\omega(f_x, c_j))}. \quad (11)$$

The term \mathcal{A} and the regularizer \mathcal{R}_{da} mainly differ in two ways. First, \mathcal{A} compares a full displacement f_x with surrounding parameter vectors c_j ; there is no explicit exclusion of a parameter c_i in the directional average. Second, we avoid to compare apples and pears by considering only normalized versions of f_x and μ_A . Thus, \mathcal{A} evaluates to one if f_x is aligned to neighboring parameter vectors. In this case, the full image similarity is considered. Otherwise, \mathcal{A} becomes zero and the contribution of the image similarity is discarded. Since k smoothes across sliding organ boundaries the direction of an actual displacement f_x does not necessarily correspond to those of its surrounding transform parameters. In these particular regions, \mathcal{A} masks out the contribution of such a point x to the similarity metric.

4 Results

We evaluate our new Directional Averages Motion Segmentation (DAMS) method on a publicly available abdominal 4DCT dataset. We register on three scale levels using an average stochastic gradient descent optimizer [8] where we upsample the transform parameters using nearest neighbor interpolation. A prominent sliding organ boundary in the dataset is located between the lung and the thoracic cavity well expressible by the structure tensor of the image. Therefore, in the first scale level, we apply the anisotropic kernel introduced in [3], where the basis functions are stretched with respect to the structure tensor of

the reference image. In the remaining levels, we use the combined Wendland kernel [3]. We apply the robust Cauchy loss $\mathcal{L}(x, x') := \frac{\beta^2}{2} \log \left(1 + (x - x')^2 / \beta^2 \right)$ where $\beta = 1$ controls the scale. In the directional average computation we have set the angle $\theta = \frac{\pi}{3} = 60^\circ$. We will make our implementation and configuration files publicly available².

4.1 POPI Model

We test our method on the 4DCT POPI dataset [15] containing 10 3D-states of a respiratory cycle of a thorax. We empirically set the registration parameters for image number 7 and fixed them for all other time steps. The initial step sizes for the optimizer have been adjusted to each case separately. Image number 1 was the target image. We compare our method to the FFD [11], the pTV [16], the SKM [4] and the bKM [3] method. For the FFD method, we took the target registration errors (TRE) from the POPI homepage³, for the other methods the authors of [16] resp. [3, 4] have kindly provided their TRE values. We calculated the expected TRE [mm] of the 40 first provided ground truth landmarks which are listed in Table 1. Our method performs on par with the tested methods in terms of TRE.

Table 1. Expected TRE and std dev [mm]. Last column: average mean \emptyset .

	0		2		3		4		5		6		7		8		9		\emptyset	
No reg	0.48	2.4	0.49	2.6	2.16	6.6	4.33	10	5.75	12	6.01	14	5.03	12	3.68	6.2	2.07	4.5	3.35	14
FFD	0.79	1.5	0.81	2.2	1.14	2.8	1.11	2.4	1.11	3.2	1.20	3.2	1.20	3.0	0.88	2.3	0.92	2.0	1.02	3.2
SKM	0.66	1.5	0.65	1.7	1.17	2.9	1.07	2.3	1.13	3.1	1.00	2.6	1.05	3.1	0.75	1.7	0.83	2.1	0.92	2.1
bKM	0.66	0.3	0.64	0.4	1.11	0.7	1.01	0.5	1.08	0.6	0.95	0.4	1.02	0.6	0.75	0.4	0.81	0.4	0.89	0.5
pTV	0.66	0.3	0.67	0.4	1.07	0.5	0.98	0.4	1.08	0.5	1.06	0.5	1.07	0.6	0.76	0.4	0.82	0.4	0.91	0.5
DAMS	0.65	0.3	0.60	0.4	1.12	0.7	1.06	0.6	1.10	0.6	0.95	0.5	1.03	0.5	0.75	0.4	0.83	0.4	0.90	0.5




Fig. 3. The transformation magnitude of a coronal slice through case 7 is visualized (all three scale levels) where we have overlaid the reference image to emphasize the spine. A clean outline of the lung at the lower vertebrae is visible.

² <https://github.com/ChristophJud/DAMSRegistration>.

³ <http://www.creatis.insa-lyon.fr/rio/popi-model>.




Fig. 4. Left: transform parameters $\{c_i^0\}_{i=1}^{N_0}$ after the first scale level as yellow arrows (amplified by 10). Right: sub-sampled f^3 after the third scale level as yellow arrows. Background: a coronal slice of the transformed target image.

In Fig. 3, we show a sample coronal slice through case 7 where we highlight the transformation magnitude. The transformation becomes finer in higher scale levels and greatly outlines the lung with a clear cutting transition at the lower vertebrae. In Fig. 4, the parameters of the first scale level and final displacements of the third level are plotted. One can clearly identify abrupt directional changes of neighboring parameters at the lower left rib and the lower vertebra.

5 Conclusion

We presented a novel regularity criterion which is targeted to discontinuity preserving image registration. The main contribution is the motion segmentation which is performed exclusively on the spatial transformation and which we integrated into the registration objective. It is based on a directional average of the transform parameters and can be directly optimized using gradient descent. In the experiments with a 4DCT dataset we have achieved competitive registration performance. It would be interesting to investigate further possibilities to adjust the similarity metric based on local motion segmentation.

References

1. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: MRF-based deformable registration and ventilation estimation of lung CT. *IEEE Trans. Med. Imaging* **32**(7), 1239–1248 (2013)
2. Hofmann, T., Schölkopf, B., Smola, A.J.: Kernel methods in machine learning. *Ann. Stat.* **36**, 1171–1220 (2008)
3. Jud, C., Möri, N., Bitterli, B., Cattin, P.C.: Bilateral regularization in reproducing kernel hilbert spaces for discontinuity preserving image registration. In: Wang, L., Adeli, E., Wang, Q., Shi, Y., Suk, H.-I. (eds.) *MLMI 2016*. LNCS, vol. 10019, pp. 10–17. Springer, Cham (2016). doi:[10.1007/978-3-319-47157-0_2](https://doi.org/10.1007/978-3-319-47157-0_2)

4. Jud, C., Möri, N., Cattin, P.C.: Sparse kernel machines for discontinuous registration and nonstationary regularization. In: Proceedings of the International Workshop on Biomedical Image Registration, pp. 9–16 (2016)
5. Kiriyanthan, S., Fundana, K., Majeed, T., Cattin, P.C.: Discontinuity preserving image registration through motion segmentation: a primal-dual approach. *Comput. Math. Methods Med.* **2016** (2016). Article ID 9504949
6. Pace, D.F., Aylward, S.R., Niethammer, M.: A locally adaptive regularization based on anisotropic diffusion for deformable image registration of sliding organs. *IEEE Trans. Med. Imaging* **32**(11), 2114–2126 (2013)
7. Papi ez, B.W., Heinrich, M.P., Fehrenbach, J., Risser, L., Schnabel, J.A.: An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration. *Med. Image Anal.* **18**(8), 1299–1311 (2014)
8. Polyak, B.T., Juditsky, A.B.: Acceleration of stochastic approximation by averaging. *SIAM J. Control Optim.* **30**(4), 838–855 (1992)
9. Preston, J.S., Joshi, S., Whitaker, R.: Deformation estimation with automatic sliding boundary computation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9902, pp. 72–80. Springer, Cham (2016). doi:[10.1007/978-3-319-46726-9_9](https://doi.org/10.1007/978-3-319-46726-9_9)
10. 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. *Med. Image Anal.* **17**(2), 182–193 (2013)
11. 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 Trans. Med. Imaging* **18**(8), 712–721 (1999)
12. Schmidt-Richberg, A., Werner, R., Handels, H., Ehrhardt, J.: Estimation of slipping organ motion by registration with direction-dependent regularization. *Med. Image Anal.* **16**(1), 150–159 (2012)
13. Shi, W., Jantsch, M., Aljabar, P., Pizarro, L., Bai, W., Wang, H., ORegan, D., Zhuang, X., Rueckert, D.: Temporal sparse free-form deformations. *Med. Image Anal.* **17**(7), 779–789 (2013)
14. von Siebenthal, M., Székely, G., Gamper, U., Boesiger, P., Lomax, A.J., Cattin, P.C.: 4D MR imaging of respiratory organ motion and its variability. *Phys. Med. Biol.* **52**(6), 1547 (2007)
15. Vandemeulebroucke, J., Sarrut, D., Clarysse, P.: The POPI-model, a point-validated pixel-based breathing thorax model. In: International Conference on the Use of Computers in Radiation Therapy, vol. 2, pp. 195–199 (2007)
16. Vishnevskiy, V., Gass, T., Szekely, G., Tanner, C., Goksel, O.: Isotropic total variation regularization of displacements in parametric image registration. *IEEE Trans. Med. Imaging* **36**, 385–395 (2016)