Bilateral Regularization in Reproducing Kernel Hilbert Spaces for Discontinuity Preserving Image Registration

Christoph Jud(✉), Nadia Möri, Benedikt Bitterli, and Philippe C. Cattin

Department of Biomedical Engineering, University of Basel, Switzerland
Gewerbestrasse. 14, 4123 Allschwil, Switzerland
christoph.jud@unibas.ch
http://dbe.unibas.ch

Abstract. The registration of abdominal images is an increasing field in research and forms the basis for studying the dynamic motion of organs. Particularly challenging therein are organs which slide along each other. They require discontinuous transform mappings at the sliding boundaries to be accurately aligned. In this paper, we present a novel approach for discontinuity preserving image registration. We base our method on a sparse kernel machine (SKM), where reproducing kernel Hilbert spaces serve as transformation models. We introduce a bilateral regularization term, where neighboring transform parameters are considered jointly. This regularizer enforces a bias to homogeneous regions in the transform mapping and simultaneously preserves discontinuous magnitude changes of parameter vectors pointing in equal directions. We prove a representer theorem for the overall cost function including this bilateral regularizer in order to guarantee a finite dimensional solution. In addition, we build direction-dependent basis functions within the SKM framework in order to elongate the transformations along the potential sliding organ boundaries. In the experiments, we evaluate the registration results of our method on a 4DCT dataset and show superior registration performance of our method over the tested methods.

1 Introduction

The image-based analysis of organ motion has attracted increasing attention in medical image analysis. Applications range from statistical motion modelling for motion compensation [4] over dynamic treatment planning [6] through to disease progression monitoring of the lung [1] to mention a few. Image registration of 4d datasets turned out to be the beating heart of a successful estimation of motion, where correspondence among temporal states is recovered by finding a meaningful spatial alignment of the image sequence. In contrast to classical registration approaches where smooth transformations are preferred, in the analysis of organ motion the spatial transform mappings are required to contain discontinuous transforms in between the organs along the sliding boundaries. Hence, if

L. Wang et al. (Eds.): MLMI 2016, LNCS 10019, pp. 10–17, 2016.
DOI: 10.1007/978-3-319-47157-0_2

these boundaries are not explicitly considered, the registration usually fails due to misalignments near organ boundaries.

In this paper, we present a new *bilateral* regularization which prefers transformations which are locally homogeneous while preserving discontinuous changes in magnitude of neighboring parameter vectors having equal and opposite directions. This has the effect that non-smooth transformations across organ boundaries stay admissible, with minor influence to the regularization penalty, and simultaneously regions within organs are transformed similarly. Hand in hand with this bilateral regularizer, we construct anisotropic basis functions which consider the structure tensor of the reference image. Thus, we align the transformations tangential to the potential organ boundaries which favours displacements along the boundaries and hinders them to evolve orthogonal to the boundaries.

Our method builds upon a recently presented sparse kernel machine (SKM) [3] for discontinuous registration. Our contribution is threefold: (1) we add bilateral regularization terms which *need not to be norms* and prove the corresponding representer theorem. (2) We integrate anisotropy into the spatially varying kernel of the SKM. Finally, (3) we provide a GPU accelerated implementation which we will make publicly available[1]. We evaluate our method on the 4DCT dataset of the POPI model [13].

The parametric approach to non-rigid image registration has been studied for more than two decades. After the introduction of the free-form deformations into registration [9], a lot of advanced registration approaches have been applied successfully to medical images. A comprehensive overview over non-rigid image registration methods, including non-parametric and discrete approaches, can be found in [11].

Discontinuity preserving attempts appeared increasingly in the past few years [5,8,10,14]. In [10], an anisotropic regularization has been introduced into a Demons like [12] framework, where the regularization has been splitted into tangential and normal directions of the image gradients. In [8], the Demons framework has been extended with an anisotropic *and* a bilateral filtering. Both approaches are non-parametric and specify the transformation properties indirectly through differential operators or smoothing filters and therefore are conceptually rather rigid. Parametric approaches for discontinuity preserving registration are less common [3,14]. In [14], discontinous basis functions are used jointly with a total-variation regularization. They have not explicitly considered organ boundaries and do not explicitly preserve discontinuous changes in the regularizer.

2 Background

In this section, we provide a brief overview over the SKM for image registration originally introduced in [3]. We especially pay attention to the representer theorem which allows for the discretization of the transformation model, having the

[1] https://github.com/ChristophJud/SKMImageRegistration.git.

guarantee of a finite dimensional minimizer. We will stick to the notation used in [3].

Let the reference and target image $I_R, I_T : \mathcal{X} \to \mathbb{R}$ map the d-dimensional input domain $\mathcal{X} \subset \mathbb{R}^d$ to scalar values. Given a spatial transform mapping $f : \mathcal{X} \to \mathbb{R}^d$ image registration can be formulated as a regularized functional minimization problem

$$\min_f \ \mathcal{D}[I_R, I_T, f] + \eta \mathcal{R}[f]. \tag{1}$$

The dissimilarity measure $\mathcal{D}$ integrates over a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ to indicate how good the target and the transformed reference image match

$$\mathcal{D}[I_R, I_T, f] := \int_{\mathcal{X}} \mathcal{L}(I_R(x + f(x)), I_T(x))dx. \tag{2}$$

The regularization term $\mathcal{R}$ enforces certain properties of the transformation. A reproducing kernel Hilbert space (RKHS) is defined as transformation model

$$\mathcal{H} := \left\{ f \middle| f(x) = \sum_{i=1}^{\infty} c_i k(x_i, x), \quad x_i \in \mathcal{X}, \quad c_i \in \mathbb{R}^d, \quad \|f\|_{\mathcal{H}} < \infty \right\}, \tag{3}$$

where $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a positive definite kernel. For a comprehensive introduction to kernel methods and RKHS we refer to [2]. The RKHS norm $\|f\|_{\mathcal{H}} = \sqrt{\langle f, f \rangle}$, with the inner product $\langle \cdot, \cdot \rangle$ of $\mathcal{H}$, measures the magnitude of f and can be defined as regularization term. Thus, functions within $\mathcal{H}$ which have smaller magnitude are preferred. For the SKM, a sparsity inducing l_1-type norm was introduced as regularization

$$\|f\|_1 := \sum_{i=1}^{N} \|c_i\|_1 + \|v_f\|_{\mathcal{H}}, \tag{4}$$

where $v_f \in \mathcal{H}$ is orthogonal to $\mathcal{H}_0$ and

$$\mathcal{H}_0 = \left\{ f_0 \in \mathcal{H} \middle| f_0(\cdot) = \sum_{i=1}^{N} c_i k(x_i, \cdot), \quad c_i \in \mathbb{R} \right\} \tag{5}$$

is a finite dimensional linear subspace of $\mathcal{H}$ with pair-wise distinct sampled points $\{x_i\}_{i=1}^N, x_i \in \mathcal{X}$.

Theorem 1. *Let the training data $\{(x_i, y_i) \in \mathcal{X} \times \mathbb{R} | i = 1, \ldots, N\}$, a loss function $\mathcal{L} : \mathcal{X} \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ and two functions $g : \mathbb{R} \to \mathbb{R}$ and $h : \mathbb{R} \to \mathbb{R}$ be given. If one of the two functions g or h is strictly increasing and the other one is nondecreasing, the minimization problem*

$$\min_{f \in \mathcal{H}} \ \sum_{i=1}^{N} \mathcal{L}\left(x_i, y_i, f(x_i)\right) + g(\|f\|_{\mathcal{H}}) + h(\|f\|_1) \tag{6}$$

has a minimizer taking the form

$$f(x) = \sum_{i=1}^{N} c_i k(x_i, x), \quad c_i \in \mathbb{R}^d. \tag{7}$$

In [3] it was proven that each f can be decomposed as

$$f(x) = \sum_{i=1}^{N} c_i k(x_i, x) + v_f \tag{8}$$

for *unique* c_i. Moreover, they showed that $\mathcal{L}$ is independent of v_f and v_f has to be zero for minimizing $g(\|f\|_{\mathcal{H}}) + h(\|f\|_1)$. Hence, each minimizer of (6) lies within $\mathcal{H}_0$. In the following, we will introduce an additional regularization term which will fulfill an extended representer theorem.

3 Method

Both of the above mentioned regularization terms consider the magnitude of the function f. Hence, functions which deviate from the zero transformation are penalized and the smoothness is implicitly given by the basis functions $k(x_i, \cdot)$. In addition, for the sliding organ boundaries, we have the following requirements on the regularization: (1) nearby transform parameters should be locally homogeneous (within organs), (2) they should be admissible if they are pointing in opposite direction (sliding organ boundaries) and (3) they should also be admissible if they have different magnitudes as long as they point in similar directions (moving tissue next to still structure).

3.1 Bilateral Regularizer

Let us start with the following regularization on f

$$\mathcal{B}(f) := \sum_{i,j=1}^{N} \|c_i - c_j\|^2 k(x_i, x_j) + \|v_f\|_{\mathcal{H}}. \tag{9}$$

Dependent on the proximity of the parameters c_i, c_j indicated by $k(x_i, x_j)$ their radiometric difference is penalized. We extend the representer theorem as follows:

Theorem 2. *Let the training data* $\{(x_i, y_i) \in \mathcal{X} \times \mathbb{R} | i = 1, \ldots, N\}$ *and a loss function* $\mathcal{L} : \mathcal{X} \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ *be given. Furthermore, let the functions* $g_i : \mathbb{R} \to \mathbb{R}$ *and* $p_i : \mathcal{H}_0 \to \mathbb{R}, i = 1 \ldots l$ *be weakly lower semicontinuous and bounded from below. If at least one of the functions* $g_{i=j}$ *is strictly increasing and the other ones* $g_{i \neq j}$ *are nondecreasing, the minimization problem*

$$\min_{f \in \mathcal{H}} \sum_{i=1}^{N} \mathcal{L}\left(x_i, y_i, f(x_i)\right) + \sum_{i=1}^{l} g_i(p_i(f_0) + \|v_f\|_{\mathcal{H}}) \tag{10}$$

has a minimizer taking the form

$$f(x) = \sum_{i=1}^{N} c_i k(x_i, x), \quad c_i \in \mathbb{R}^d. \tag{11}$$

Proof. Since c_i are unique for a particular f, each summand $g_i(p_i(f_0) + \|v_f\|_{\mathcal{H}})$ is well-defined. As already proved in [3], $\mathcal{L}$ is independent of v_f. The following inequalities hold as g_i are nondecreasing

$$g_i(p_i(f_0) + \|v_f\|_{\mathcal{H}}) \geq g_i(p_i(f_0)) \quad \forall i. \tag{12}$$

For the functions g_i which are strictly increasing, we have that equality holds iff $v_f = 0$. As for the other functions the value does not increase when setting $v_f = 0$, we conclude that (10) can be minimized iff $v_f = 0$. □

For the bilateral regularizer $\mathcal{B}$ as an example, we set the corresponding p_i to $p_i(f_0) = \sum_{i,j=1}^{N} \|c_i - c_j\|^2 k(x_i, x_j)$. Note that p_i need *not* to be a norm.

With the regularizer $\mathcal{B}$, we meet the first of our requirements. For the specific case of discontinuity preserving registration we want to integrate the two remaining properties as well. Consider the regularizer

$$\mathcal{B}_p(f) := \sum_{i,j=1}^{N} \frac{s^2}{2} \log\left(1 + \frac{1}{s}\mathcal{P}(c_i, c_j)^2\right) k(x_i, x_j) + \|v_f\|_{\mathcal{H}}, \tag{13}$$

where $\mathcal{P} : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ measures the area of the parallelogram[2] which is spanned by c_i and c_j and a positive scaling s. With this regularizer, nearby parameters which are aligned in direction can discontinuously vary in magnitude and even in pointing in opposite direction without influencing the value of $\mathcal{B}_p$. Thus, $\mathcal{B}_p$ does meet all the above mentioned requirements. The proof of the representer theorem including $\mathcal{B}_p$ is now straight forward. The effect of applying $\mathcal{B}_p$ is shown in Fig. 1.

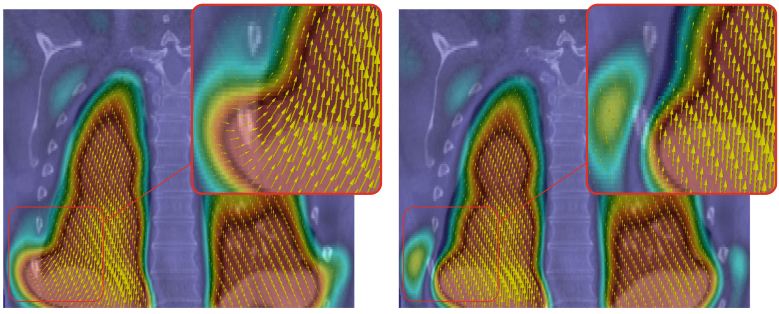

Fig. 1. Coronal slice of an example registration result applying $\mathcal{B}_p$ as regularizer once only slightly (left) and once very strongly (right). The over regularization (right) results in region-wise aligned transformations.

[2] Note that for $d = 3$, $\mathcal{P}$ is the magnitude of the cross-product $\|c_i \times c_j\|$.

3.2 Anisotropic Kernel

In addition to the introduced bilateral regularizers we further tailor the framework to sliding organ boundaries by introducing an anisotropic transformation model. We elongate the transformations along potential organ boundaries given a positive definite matrix-valued function Σ. Using Σ, we construct an *anisotropic* and *nonstationary* kernel as it has been proposed in [7]. Thus, we only need an appropriate Σ-function.

We assume, that high gradients in the image indicate potential sliding organ boundaries. Given a guidance image $I : \mathcal{X} \to \mathbb{R}$ (in our case the reference image) we derive the principal directions of intensity changes by performing eigen analysis on the structure tensor S, where $S(x) = U_x \Lambda_x \Lambda_x U_x^T$, U_x contains the eigenvectors column-wise and Λ_x is a diagonal matrix containing the eigenvalues λ_i of $S(x)$ on the diagonal. We use the largest eigenvalue to narrow the kernel in orthogonal direction to the sliding boundaries, while we stretch it in the remaining directions. We scale the eigenvalues

$$\tilde{\lambda_1} = \frac{1}{1+\lambda_1^\alpha}, \quad \tilde{\lambda_i} = 1 + \frac{\lambda_1^\alpha}{1+\lambda_1^\alpha}, \quad i = \{2 \dots d\}, \tag{14}$$

where α is a positive constant (e.g. $\alpha = 1.5$). The final covariances become $\Sigma(x) = U_x \tilde{\Lambda}_x \tilde{\Lambda}_x U_x^T$. As we assume, that the main motion induced by respiration is aligned with the z-axis, we scale λ_1 with $d^{-1/2}\mathcal{P}(u_1, e_z)^2$ where u_1 is the first eigenvector and e_z is the unit vector in z direction. Thus, the anisotropic distortion is damped, if the direction is not aligned to the respiratory motion.

4 Results

We evaluate our kernel machine with the $l1$-type regularizer (4) for a sparse transform mapping (*sKM*). Furthermore, denoted as *pKM*, we apply the parallelogram regularizer $\mathcal{B}_p$. We run *sKM* and *pKM* using the isotropic (iso) as well as the anisotropic (ani) transformation model. We used the combined *Wendland* kernel proposed in [3] and always apply the squared loss function.

As dataset, we used the 4DCT POPI model [13] containing a respiratory cycle of a patients thorax including ground truth landmarks. We manually optimized the parameters for image 7 and used the same configuration for all other time steps. The target image was image number 1. We performed gradient descent on three scales with an isotropic control point grid spacing of $\{16, 8, 4\}$ resulting in $\{25\,\text{k}, 189\,\text{k}, 1.5\,\text{m}\}$ parameters and uniformly sampled $\{133\,\text{k}, 258\,\text{k}, 619\,\text{k}\}$ image points. We compare our results to the free-form deformation method FFD [9] and an isotropic variant of the parametric total variation method pTV [14]. For the FFD, we took the transformations from the POPI homepage[3]. For pTV the authors of [14] have kindly provided their TRE[4] values. The expected target registration errors are listed in Table 1. On average (cf. Ø in Table 1), our

[3] http://www.creatis.insa-lyon.fr/rio/popi-model.

[4] Target registration error: Euclidean distance between ground truth landmarks and reference landmarks which have been warped by the resulting f.

Table 1. Expected TRE [mm] of the first 40 landmarks with respect to image 1.

Method	0	2	3	4	5	6	7	8	9	Ø	Ø*
No reg.	0.34	0.07	2.15	4.72	5.81	6.25	4.87	4.16	2.11	3.39	5.41
FFD [9]	0.87	0.56	1.01	1.05*	1.06*	1.02*	1.11*	0.76*	0.86	0.92	1.06
pTV [14]	0.64	0.73	1.08	0.99*	0.93*	0.94*	1.08*	0.82*	0.71	0.88	0.98
SKM [3]	0.65	0.50	1.14	0.97*	0.96*	0.90*	0.93*	0.73*	0.66*	0.83	0.94
sKM iso	0.67	0.55	1.00	0.94*	1.03*	0.92*	0.80*	0.69*	0.71	0.81	0.92
sKM ani	0.68	0.53	1.00	0.96*	0.93*	0.93*	0.89*	0.75*	0.71	0.82	0.93
pKM iso	0.69	0.54	0.96	0.94*	1.00*	0.92*	0.81*	0.69*	0.72	0.81	0.92
pKM ani	0.68	0.54	1.05	0.96*	0.93*	0.92*	0.88*	0.76*	0.71	0.83	0.92

The target registration errors have been fitted to the Maxwell-Boltzmann distribution where the derived mean values are listed in the table. In the last two columns, the average mean Ø and a significant average mean Ø* are provided.

methods perform similarly and outperform the tested methods FFD and pTV in terms of TRE. In addition, we highlight the cases 4 to 8, where all methods have a higher probability[5] than 0.95 to beat the non-registered TRE. Their average mean is listed as Ø*. These cases are in the exhalation phase and differ most from the reference image 1, which is in the inhalation state. In Fig. 2, we show qualitative improvements when applying the anisotropic methods.

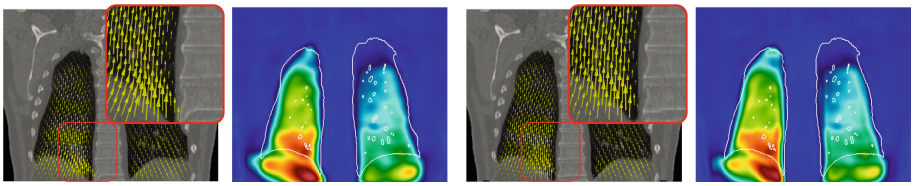

Fig. 2. Registration results (CT slice and corresponding transformation magnitudes) of pKM iso (left) and pKM ani (right). One can see, that with the anisotropic version the lower vertebra is not distorted and the ground truth outline (white) of the lung is better represented in the transformations.

5 Conclusion

We presented a novel method to approach discontinuity preserving image registration. With the introduction of bilateral regularization in a reproducing kernel Hilbert space we opened new possibilities to enforce certain neighborhood

[5] We defined the probability of a method H_a to beat the baseline H_0 as $P(X < Y)$, where the *independent* random variables X and Y are distributed according to the Maxwell-Boltzmann distribution of the respective method H_a and H_0.

properties of the transform mapping. By design, the parallelogram regularizer $\mathcal{B}_p$ allows discontinuous changes in neighboring transform parameters as long as they are aligned in direction while it prefers locally homogeneous parameters otherwise. We additionally integrated an anisotropic transformation model into the SKM framework to align the transformations with potential sliding organ boundaries. Our method outperforms the tested approaches on the POPI dataset which was shown in the experiments. However, a thorough evaluation on further datasets is needed and planned in near future.

References

1. Gorbunova, V., Lo, P., Ashraf, H., Dirksen, A., Nielsen, M., de Bruijne, M.: Weight preserving image registration for monitoring disease progression in lung CT. In: Axel, L., Fichtinger, G., Metaxas, D., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 863–870. Springer, Heidelberg (2008)
2. Hofmann, T., Schölkopf, B., Smola, A.J.: Kernel methods in machine learning. Ann. Stat. **36**, 1171–1220 (2008)
3. Jud, C., Möri, N., Cattin, P.C.: Sparse kernel machines for discontinuous registration. In: 7th International Workshop on Biomedical Image Registration (2016)
4. Jud, C., Preiswerk, F., Cattin, P.C.: Respiratory motion compensation with topology independent surrogates. In: Workshop on Imaging and Computer Assistance in Radiation Therapy (2015)
5. Kiriyanthan, S., Fundana, K., Majeed, T., Cattin, P.C.: A primal-dual approach for discontinuity preserving image registration through motion segmentation. Int. J. Comput. Math. Methods Med. (2016, in press)
6. Möri, N., Jud, C., Salomir, R., Cattin, P.C.: Leveraging respiratory organ motion for non-invasive tumor treatment devices: a feasibility study. Phys. Med. Biol. **61**(11), 4247–4267 (2016)
7. Paciorek, C.J., Schervish, M.J.: Spatial modelling using a new class of nonstationary covariance functions. Environmetrics **17**(5), 483–506 (2006)
8. Papież, 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)
9. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Imaging **18**(8), 712–721 (1999)
10. Schmidt-Richberg, A.: Sliding motion in image registration. Registration Methods for Pulmonary Image Analysis, pp. 65–78. Springer, Wiesbaden (2014)
11. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: a survey. IEEE Trans. Med. Imaging **32**(7), 1153–1190 (2013)
12. Thirion, J.-P.: Image matching as a diffusion process: an analogy with Maxwell's demons. Med. Image Anal. **2**(3), 243–260 (1998)
13. 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)
14. Vishnevskiy, V., Gass, T., Székely, G., Goksel, O.: Total variation regularization of displacements in parametric image registration. In: Yoshida, H., Näppi, J.J., Saini, S. (eds.) ABDI 2014. LNCS, vol. 8676, pp. 211–220. Springer, Heidelberg (2014)

http://www.springer.com/978-3-319-47156-3

Machine Learning in Medical Imaging
7th International Workshop, MLMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 17, 2016, Proceedings
Wang, L.; Adeli, E.; Wang, Q.; Shi, Y.; Suk, H.-I. (Eds.)
2016, XIV, 324 p. 127 illus., Softcover
ISBN: 978-3-319-47156-3