NON-RIGID REGISTRATION GUIDED BY LANDMARKS AND LEARNING

Jutta Eckl[1]*, Volker Daum*[1]*, Joachim Hornegger*[1,2]*, Kilian M. Pohl*[3]

[1] Pattern Recognition Lab
[2] Erlangen Graduate School in Advanced Optical Technologies (SAOT)
Friedrich-Alexander-Universität Erlangen-Nürnberg
[3] Department of Radiology, University of Pennsylvania
E-Mail: jutta.eckl@cs.fau.de

ABSTRACT

Registration methods frequently rely on prior information in order to generate anatomical meaningful transformations between medical scans. In this paper, we propose a novel intensity based non-rigid registration framework, which is guided by landmarks and a regularizer based on Principle Component Analysis (PCA). Unlike existing methods in this domain, the computational complexity of our approach reduces with the number of landmarks. Furthermore, our PCA is invariant to translations. The additional regularizer is based on the outcome of this PCA. We register a skull CT scan to MR scans aquired by a MR/PET hybrid scanner. This aligned CT scan can then be used to gain an attenuation map for PET reconstruction. As a result we have a Dice coefficient for bone areas at 0.71 and a Dice coefficient for bone and soft issue areas at 0.97.

Index Terms— non-rigid registration, landmarks, regularizer based on PCA

1. INTRODUCTION

The registration of medical scans often relies on prior information to produce anatomical meaningful transformations [1]. Techniques, such as [2, 3] constrain the registration to a manifold. To capture anatomical meaningful deformations, the manifold is generally trained on a large set of scans. In this paper, we develop a non-rigid registration algorithm based on technology only requiring small data sets for training.

Our registration approach is guided by landmarks as well as PCA on deformation fields. Non-rigid landmark based methods, such as Fischer et al. [4], use the Lagrange multipliers to integrate the landmarks into the energy function capturing the registration problem. Hereby each Lagrange multiplier expands the dimensionality. This is inconsistent to the intuitive idea that landmarks should simplify the registration. We instead propose to fixate the deformation maps at those locations treating the correspondences as Dirichlet boundary conditions. Thus, the computational burden reduces with an increasing number of landmarks.

The deformation of the remaining image space is then defined by the minimum of an energy functional consisting of an image similarity term and two regularizers. The first regularizer enforces the smoothness of the deformation field, the second regularizer is based on the outcome of a PCA on the deformation maps. Compared to manifold learning techniques, PCA based non-rigid registration methods, such as [5], require less training samples as PCA is based on the assumption that the distribution of the deformation fields is Gaussian. We thus do not constrain the resulting deformation to the space defined by the PCA model, such as in [5], but instead only use the PCA model to guide the registration. We furthermore reduce the variation to be captured by PCA by defining a model that is invariant to translations.

We use our algorithm to register a skull CT scan to MR scans which were acquired by a MR/PET hybrid scanner. MR/PET hybrid scanners can help to detect tumors in an early stage and give information about the activity of the tumor, improve the treatment of patients in radiation therapy or assist medicines in differential diagnostics. Unlike for PET/CT hybrid scanners, where the attenuation map is constructed from the CT scan, the MR scan from a MR/PET hybrid scanner cannot be directly used for attenuation correction of PET due to the image inhomogeneity in MR scans. The skull CT scan is the average of various CT scans which we will call an atlas CT. The registration process of the atlas CT to the MR scan is the focus of our application.

2. METHODS

We now construct a non-rigid registration algorithm to map each voxel $\mathrm{x} \in \Omega_M$ of the image domain $\Omega_M \subset \mathbb{R}^d$ of the moving image M to a voxel $\mathrm{x} - \boldsymbol{u}(x) \in \Omega_F$ in the image domain $\Omega_F \subset \mathbb{R}^d$ of the fixed image F with the transformation $\boldsymbol{u}(\cdot)$. We compute the optimal mapping $\boldsymbol{u}^*$ by solving

$$\boldsymbol{u}^* := \mathrm{argmin}_{\boldsymbol{u}}\, \mathcal{D}(F, M, \boldsymbol{u}) + \alpha \cdot \mathcal{R}(\boldsymbol{u}) + \beta \cdot \mathcal{P}(\boldsymbol{u}) \quad (1)$$

where $\mathcal{D}(\cdot)$ is the image matching term, $\mathcal{R}(\cdot)$ the regularizer enforcing smoothness on $\boldsymbol{u}$, $\mathcal{P}(\cdot)$ is the regularizer based on our PCA and $\alpha, \beta \in \mathbb{R}$ are weighting factors.

2.1. Landmarks

To follow the intuitive idea mentioned in the introduction, our framework now efficiently embeds the landmarks into the registration framework by treating the landmarks as Dirichlet boundary conditions and determining $\boldsymbol{u}^*$ for the remaining image domain. In other words, let $\Omega_L \subset \Omega_M$ be the set of landmarks in the moving image and $\boldsymbol{u}_L(\mathrm{x})$ be the corresponding deformation mapping landmark $\mathrm{x} \in \Omega_L$ to its corresponding landmark in Ω_F. We then solve (1) for $\Omega_M \setminus \Omega_L$ using the constraint that $\boldsymbol{u}^*(\mathrm{x}) = \boldsymbol{u}_L(\mathrm{x})$ for $\mathrm{x} \in \Omega_L$. As we cut off the landmarks, the remaining image domain has no longer rectangular shape. As consequence, the respective matrices are asymmetric now. We address this problem in Section 2.3 where we specify the implementation of our registration framework.

2.2. Regularization based on PCA

The definition of $\mathcal{P}(\cdot)$ is based on our PCA of deformation fields which identifies the major modes of variation within a sampling set. Let $\mathcal{U}$ be a Hilbertspace. A functional PCA decomposition finds mutually orthogonal modes $\boldsymbol{v}_i$ of normal length that optimize

$$\begin{aligned} \boldsymbol{v}_i^*(\mathrm{x}) &:= \mathrm{argmax}_{\boldsymbol{v}_i}\, \mathcal{E}_{PCA}(\boldsymbol{v}_i) \\ &= \mathrm{argmax}_{\boldsymbol{v}_i} \sum_{j=1}^{m} \langle \boldsymbol{w}_j - \overline{\boldsymbol{w}}, \boldsymbol{v}_i \rangle^2 \end{aligned} \tag{2}$$

with $\overline{\boldsymbol{w}}(\mathrm{x}) := \frac{1}{m}\sum_{i=1}^{m} \boldsymbol{w}_i(\mathrm{x})$, $\boldsymbol{w}_i$ are the m sampling functions, i.e. the deformation maps of the training set, and $\langle x, y \rangle := x^T y$ is the inner product induced by the Hilbertspace $\mathcal{U}$. Accordingly, the values of $\mathcal{P}(\boldsymbol{u})$ relate to the residual of $\boldsymbol{u}$ with respect to the PCA space:

$$\mathcal{P}(\boldsymbol{u}) := \left|\left| \boldsymbol{u} - \left(\overline{\boldsymbol{w}} + \sum_{i=1}^{m} \boldsymbol{v}_i \langle \boldsymbol{v}_i^*, \boldsymbol{u} - \overline{\boldsymbol{w}} \rangle \right) \right|\right|_{\mathcal{U}} \tag{3}$$

where $||\cdot||_{\mathcal{U}}$ is the L^2-Norm.

In the remainder of this Section, we determine the major axis $\boldsymbol{v}_i$ fulfilling (2) by incorporating the constraint by a Langrangian multiplier ensuring the normal length.

$$\boldsymbol{v}_i^* = \mathrm{argmax}_{\boldsymbol{v}_i} \sum_{j=1}^{m} \langle \boldsymbol{w}_j - \overline{\boldsymbol{w}}, \boldsymbol{v}_i \rangle^2 - \lambda_i (\langle \boldsymbol{v}_i, \boldsymbol{v}_i \rangle - 1) \tag{4}$$

From this, we calculate the Gâteaux derivative $\mathrm{d}\mathcal{E}_{PCA}(\boldsymbol{v}_i; \boldsymbol{\eta})$ at $\boldsymbol{v}_i$ in the direction $\boldsymbol{\eta} \in \mathcal{U}$ to determine the extremal points what results in

$$\left\langle \boldsymbol{\eta}, \sum_{j=1}^{m} (\boldsymbol{w}_j - \overline{\boldsymbol{w}}) \langle \boldsymbol{w}_j - \overline{\boldsymbol{w}}, \boldsymbol{v}_i \rangle \right\rangle = \langle \boldsymbol{\eta}, \lambda_i \boldsymbol{v}_i(\mathrm{x}) \rangle \tag{5}$$

For details of the calculation of the Gâteaux derivative, see [6].

Due to the inconsistence of rigid registration, we want to make our model robust to translations of rigid alignment. Therefore d constant global rigid translation functions $\boldsymbol{t}_j$ with

$$\boldsymbol{t}_j(\mathrm{x}) = (0, \ldots, 0, \underbrace{1}_{\text{j-th entry}}, 0, \ldots, 0)^T \in \mathbb{R}^d \quad j = 1, \ldots, d \tag{6}$$

are introduced. Our translation invariant sampling functions $\widetilde{\boldsymbol{w}_i}$ are now defined as :

$$\widetilde{\boldsymbol{w}}_i(\mathrm{x}) = \boldsymbol{w}_i(\mathrm{x}) - \sum_{j=1}^{d} \boldsymbol{t}_j(\mathrm{x}) \langle \boldsymbol{t}_j, \boldsymbol{w}_i \rangle \quad i = 1, \ldots, n \tag{7}$$

These new sampling functions $\widetilde{\boldsymbol{w}}_i$ are orthogonal to $\boldsymbol{t}_j$. We recompute $\overline{\boldsymbol{w}}$ and $\boldsymbol{v}_i^*$ from (2) by replacing $\boldsymbol{w}_i$ with $\widetilde{\boldsymbol{w}}_i$. The modes $\boldsymbol{v}_k^*$ are orthogonal to $\boldsymbol{t}_j$ as each $\boldsymbol{v}_k^*$ is a linear combination of $\widetilde{\boldsymbol{w}}_i$. This means, that $\boldsymbol{t}_j$ are orthonormal to our PCA basis and we can augment this basis with the additional modes $\boldsymbol{v}_{m+j}^* = \boldsymbol{t}_j$. Hence, $\mathcal{P}(\boldsymbol{u})$ is translation invariant.

We also generate a translation invariant PCA model based on the derivatives of the deformations. The first derivative of a deformation is already translation invariant and the Laplacian $\Delta\boldsymbol{u}$ of a deformation field is invariant to global affine transformations. This Laplacian is the basis for the curvature PCA model. The mean $\overline{w}$ and the Eigenmodes $\boldsymbol{v}_i$ are computed on the Laplacian of sampling functions $\Delta\boldsymbol{w}$ by similar proceeding as described above with

$$\mathcal{P}_\Delta(\boldsymbol{u}) = \left|\left| \Delta\boldsymbol{u}(\mathrm{x}) - \left(\Delta\overline{\boldsymbol{w}} + \sum_{i=1}^{m} \boldsymbol{v}_i(\mathrm{x}) \langle \boldsymbol{v}_i, \Delta\boldsymbol{u} - \Delta\overline{\boldsymbol{w}} \rangle \right) \right|\right|_{\mathcal{U}}^2 \tag{8}$$

The registration in Section 3 is performed using $\mathcal{P}$ or $\mathcal{P}_\Delta$ as additional regularizer of (1).

2.3. Numerical Solution

Similar to [4], we determine the solution to (1) by solving

$$\nabla_{\boldsymbol{u}} \left(\mathcal{D}(F, M, \boldsymbol{u}) + \alpha \mathcal{R}(\boldsymbol{u}) + \beta \mathcal{P}(\boldsymbol{u}) \right) = 0 \tag{9}$$

where $\nabla_{\boldsymbol{u}}$ is the gradient with respect to $\boldsymbol{u}$. We solve this equation on a discrete image grid where $\boldsymbol{u}$ is a vector capturing the displacement of each voxel in Ω_M. Furthermore, we now assume without loss of generality that $\overline{\boldsymbol{w}}(\mathrm{x}) = 0$:

$$\nabla_{\boldsymbol{u}} \mathcal{P}(\boldsymbol{u}) = 2 \left(\boldsymbol{u} - \sum_{i=1}^{m} \boldsymbol{v}_i \langle \boldsymbol{v}_i, \boldsymbol{u} \rangle \right) - 2\widetilde{\boldsymbol{w}}. \tag{10}$$

By simplifying $\mathcal{R}(\boldsymbol{u})$ to an inner product $\boldsymbol{u}^T \boldsymbol{A} \boldsymbol{u}$, the above equation simplifies to

$$\nabla_{\boldsymbol{u}} \mathcal{D}(F, M, \boldsymbol{u}) + \alpha \boldsymbol{A}\boldsymbol{u} + \beta \nabla_{\boldsymbol{u}} \mathcal{P}(\boldsymbol{u}) = 0 \tag{11}$$

In practice, we will solve this problem with the semi-implicit gradient descent which updates $\boldsymbol{u}$ according to solving the linear equation system

$$(\boldsymbol{I} + s{\cdot}\alpha \cdot \boldsymbol{A}) \cdot \boldsymbol{u}^{(i+1)} = \boldsymbol{u}^{(i)} s \cdot \nabla_{\boldsymbol{u}} \left(\mathcal{D}(F, M, \boldsymbol{u}^{(i)}) + \beta \mathcal{P}(\boldsymbol{u}^{(i)}) \right) \quad (12)$$

with $\boldsymbol{u}^{(i)}$ being the mapping at the i-th iteration, $\boldsymbol{I}$ is the identity matrix and s is the step size. For more details see [6].

We determine the initial mapping via $\boldsymbol{A} \cdot \boldsymbol{u}^{(0)} = y$, where y is a zero vector except for the entries corresponding to landmarks. There it has the values $\boldsymbol{u}_L$. Thus $\boldsymbol{u}^{(0)}$ correctly aligns the landmarks between the scans and fulfills the smoothness constraints but ignores the image information. Alternatively, we fix $\boldsymbol{u}^{(i+1)}$ at the landmarks and update $\boldsymbol{u}^{(i+1)}$ at all other voxels in the image domain. We do so by simply setting the rows associated with landmarks in $\boldsymbol{A}$ and $\nabla_{\boldsymbol{u}} \left(\mathcal{D}(F, M, \boldsymbol{u}^{(i)}) + \mathcal{P}(\boldsymbol{u}^{(i)}) \right)$ to zero. Now, $(\boldsymbol{I} + s \cdot \alpha \cdot \boldsymbol{A})$ is asymmetric. We could either resymmetrize the matrices or apply methods not requiring symmetry like the stabilized bi-conjugate gradient method. The second kind of solution is employed in our experiments. The approach thus can guarantee proper alignment in the vicinity of the landmarks but is a little bit slower than updating the entire deformation field.

We outlined our numerical solution for the semi-implicit gradient descent (12) as it is easier and more clearly to illustrate the integration of Dirichlet boundary conditions into this scheme. Due to former analysis of optimization schemes, we implemented for our experiments the Newton scheme after Broyden, Fletcher, Goldfarb and Shanno [7]. The iteration step via Newton holds

$$\boldsymbol{u}^{(i+1)} = \boldsymbol{u}^{(i)} s \left(\boldsymbol{H}_{\mathcal{D}(F,M,\boldsymbol{u}(i))} + \alpha \boldsymbol{A} + \beta \boldsymbol{H}_{\mathcal{P}} \right)^{-1} \left(\nabla_{\boldsymbol{u}} \mathcal{D}(F, M, \boldsymbol{u}(i) + \alpha \boldsymbol{A} \boldsymbol{u}^{(i)} + \beta \boldsymbol{H}_{\mathcal{P}} \boldsymbol{u}^{(i)} \right) \quad (13)$$

where $\boldsymbol{H}_{\mathcal{D}(F,M,\boldsymbol{u}(i))}$ is the Hessian of $\mathcal{D}(F, M, \boldsymbol{u}(i))$ and $\boldsymbol{H}_{\mathcal{P}}$ is the Hessian of $\mathcal{P}$. To avoid solving the large and densely filled Hessian $\boldsymbol{H}$, we employ an approximation based on Sherman-Morrison-Woodbury [8]. We make use of mutual information [9] as distance measure $\mathcal{D}$ in conjunction with the curvature regularizer $\mathcal{R}$ [10].

3. EXPERIMENTS

We perform two different types of experiments. First we consider images of a PET/CT scanner and compare different kind of registration methods. Then we apply out body scans of a MR/PET hybrid scanner and align an atlas CT to the MR scan.

3.1. Influence of landmarks on a PET/CT registration

To illustrate the influence of landmarks, we visually compare the outcome of different registration approaches to align a CT to the corresponding PET scan (Fig. 1). There were 7 landmarks specified: the superior and inferior renal capsules of the kidneys (4 landmarks), the lower tip of the liver (1 landmark) and the diaphragm (2 landmarks). Although the non-rigid registration shown in picture (c) is not state-of-the-art, we just want to illustrate the influence of landmarks as only non-rigid registration with landmarks (d) does match the CT scan with the PET scan as expected.

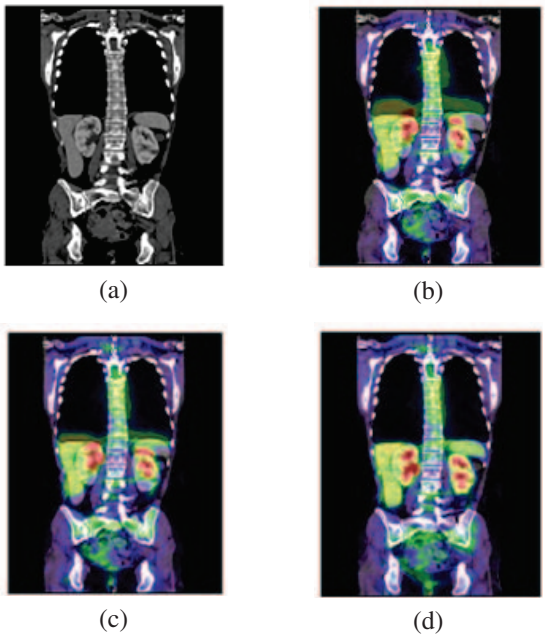


Fig. 1. (a) CT, (b) PET/CT rigid, (c) PET/CT non-rigid, (d) PET/CT non-rigid with landmarks

3.2. Atlas registration in a MR/PET application scenario

As the MR scan cannot be directly used for attenuation correction, an atlas CT is registered to a MR scan enabling attenuation correction of PET scans acquired in MR/PET hybrid scanners. The quality of this CT-MR registration is the focus of our experiment. We have 34 skull CT scans of 34 different patients and all in all 25 T1 weighted and 17 T2 weighted MR scans of those patients. For efficiency reasons we downsample the data to an isotropic resolution of 1.95mm and a volume size of $128 \times 128 \times 71$ to keep computational time down. Non-anatomical structures, such as the table, are masked and ignored. To determine the quality of the registered atlas CT, the CT scan of each patient is used as groundtruth employing the following four measures: the Root Mean Square Error in Hounsfield unit (HU) (RMSE), the Mean Absolute Error in HU (MAE), the Dice coefficient of thresholded bone areas (thresholded at 600HU) (BDICE) and the Dice coefficient of thresholded bone and soft issue areas (thresholded at -200HU) (STDICE).

As the employed Newton scheme [7] is a Quasi-Newton scheme, the inverse of the Hessian has to be calculated. This is either done exactly or with the approximation [8] as mentioned in Section 2.3. We compare the accuracy of our algorithm with respect to 5 different implementations : without PCA regularization (standard), with PCA regularization $\mathcal{P}(\cdot)$ on the deformation map (PCA exact) as well as PCA regularization on the curvature $\mathcal{P}_{\Delta}(\cdot)$ (curv PCA exact), and finally both PCA implementations with approximated Hessian (PCA approx and curv PCA approx). The algorithms with additional PCA regularization were employed with several number of eigenmodes m to test their impact. In Table 1 the results are listed for $m = 10$. Our results were fairly insensitive for $m > 10$. The average E and standard deviation σ is summarized in Fig. 2.

Results for the T1 weighted images								
	RMSE		MAE		BDICE		STDICE	
	E	σ	E	σ	E	σ	E	σ
standard	177	5.63	76	4.36	0.67	0.33	0.97	0.07
PCA approx	169	5.85	74	4.39	0.69	0.34	0.97	0.075
PCA exact	168	5.76	74	4.32	0.69	0.33	0.97	0.075
curv PCA approx	169	5.85	75	4.38	0.69	0.34	0.97	0.074
curv PCA exact	170	5.86	75	4.39	0.69	0.34	0.97	0.074

Results for the T2 weighted images								
	RMSE		MAE		BDICE		STDICE	
	E	σ	E	σ	E	σ	E	σ
standard	189	7.72	79	5.14	0.67	0.33	0.97	0.11
PCA approx	168	4.58	73	3.46	0.71	0.28	0.97	0.086
PCA exact	168	4.64	73	3.47	0.71	0.28	0.97	0.086
curv PCA approx	168	4.57	73	3.45	0.71	0.28	0.97	0.086
curv PCA exact	168	4.59	73	3.46	0.71	0.28	0.97	0.087

Fig. 2. Table of Results

All in all, PCA regularized implementations compared to the standard implementation clearly enhance the results. For example, the Dice coefficient for bone areas for the PCA approximate for T1 weighted images is 0,69, i.e. that the first 10 eigenmodes cover 69% of the bone areas. The corresponding Dice coefficient for bone and soft issue areas is 0.97, so the content of the remaining eigenmodes is irrelevant. As we achieve the coverage for bone and soft issue areas of 97% even with the approximate algorithm, we keep the computational work low by increasing simultaneously the robustness of the registration.

4. CONCLUSION

In this paper we have proposed a novel method how to integrate prior information to a non-rigid registration. The algorithm is guided by landmarks and the outcome of a PCA on deformation maps. Unlike previous approaches, we reduced the computational complexity with our method treating the landmarks as Dirichlet boundary conditions. In addition, we developed PCA model based regularizers that are invariant to global translations. By this means, we demonstrated on the one hand the great impact of non-rigid registration with landmarks and on the other hand we illustrated that our additional PCA regularizers help to increase the robustness of an atlas registration.

5. REFERENCES

[1] Glocker B., Komodakis N., Navab N., Tziritas G., and Paragios N., "Dense Registration with Deformation Priors," *Information processing in medical imaging, Vol. 21, pp. 540-551*, 2009.

[2] Wolz R., Aljabar P., Hajnal J., Hammers A., Rueckert D., and ADNI, "Leap: Learning embeddings for atlas propagation," *Neuroimage, Vol. 49, pp. 1316-1325*, 2010.

[3] Gerber S., Tasdizen T., Fletscher P. T., Joshi S., and Whitaker R., "Manifold modeling for brain population analysis," *Medical Image Analysis, Vol. 14, pp. 642-653*, 2010.

[4] Fischer B. and Modersitzki J., "Combining landmark and intensity driven registrations," *Applied Mathematics and Mechanics, Vol. 3, No.1, pp.32-35*, 2003.

[5] Wang Y. Staib L.H., "Physical Based Non-Rigid Registration Incorporating Statistical Shape Information," *Medical Image Analysis, Vol. 1, pp. 35-51*, 2000.

[6] Daum V., "Model-Constrained Non-Rigid Registration in Medicine," *PhD thesis, University of Erlangen-Nuremberg*, 2011.

[7] Nocedal J. and Wright S.J., "Numerical optimization," *Springer Series in Operation Research*, 1999.

[8] Petersen K.B. and Pedersen M.S., *The Matrix cookbook*, Technical University of Denmark, 2008.

[9] Maes F., Collignon A., Vandermeulen D., Marchal G., and Suetens P., "Multimodality image registration by maximization of mutual information," *IEEE Transactions on Medical Imaging, Vol. 16, No. 2, pp. 187-198*, 1997.

[10] Fischer B. and Modersitzki J., "Curvature based image registration," *Journal of Mathematical Imaging and Vision, Vol. 18, No.1, pp. 81-85*, 2003.