Robust non-rigid registration through agent-based action learning

Julian Krebs, Tommaso Mansi, Hervé Delingette, Li Zhang, Florin C Ghesu, Shun Miao, Andreas Maier, Nicholas Ayache, Rui Liao, Ali Kamen

► To cite this version:

Julian Krebs, Tommaso Mansi, Hervé Delingette, Li Zhang, Florin C Ghesu, et al.. Robust non-rigid registration through agent-based action learning. Medical Image Computing and Computer Assisted Interventions (MICCAI), Sep 2017, Quebec, Canada. pp.344-352, 10.1007/978-3-319-66182-7_40 . hal-01569447

HAL Id: hal-01569447

<https://hal.inria.fr/hal-01569447>

Submitted on 26 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust non-rigid registration through agent-based action learning

Julian Krebs¹²³, Tommaso Mansi¹, Hervé Delingette², Li Zhang¹, Florin C. Ghesu^{1,3}, Shun Miao¹, Andreas Maier³, Nicholas Ayache², Rui Liao¹, and Ali Kamen¹

¹ Siemens Healthineers, Medical Imaging Technologies, Princeton, New Jersey, USA

² Université Côte d’Azur, Inria, Asclepios Team, Sophia Antipolis, France
julian.krebs@inria.fr

³ Friedrich-Alexander-Universität, Pattern Recognition Lab, Erlangen, Germany

Abstract. Robust image registration in medical imaging is essential for comparison or fusion of images, acquired from various perspectives, modalities or at different times. Typically, an objective function needs to be minimized assuming specific a priori deformation models and pre-defined or learned similarity measures. However, these approaches have difficulties to cope with large deformations or a large variability in appearance. Using modern deep learning (DL) methods with automated feature design, these limitations could be resolved by learning the intrinsic mapping solely from experience. We investigate in this paper how DL could help organ-specific (ROI-specific) deformable registration, to solve motion compensation or atlas-based segmentation problems for instance in prostate diagnosis. An artificial agent is trained to solve the task of non-rigid registration by exploring the parametric space of a statistical deformation model built from training data. Since it is difficult to extract trustworthy ground-truth deformation fields, we present a training scheme with a large number of synthetically deformed image pairs requiring only a small number of real inter-subject pairs. Our approach was tested on inter-subject registration of prostate MR data and reached a median DICE score of .88 in 2-D and .76 in 3-D, therefore showing improved results compared to state-of-the-art registration algorithms.

1 Introduction

Registration of images with focus on the ROI is essential in fusion and atlas-based segmentation (e.g. [9]). Traditional algorithms try to compute the dense mapping between two images by minimizing an objective function with regard to some similarity criterion. However, besides challenges of solving the ill-posed and non-convex problem many approaches have difficulties in handling large deformations or large variability in appearance. Recently, promising results using deep representation learning have been presented for learning similarity metrics [8], predicting the optical flow [1] or the large deformation diffeomorphic metric mapping-momentum [10]. These approaches either only partially remove the

above-mentioned limitations as they stick to an energy minimization framework (cf. [8]) or rely on a large number of training samples derived from existing registration results (cf. [1,10]).

Inspired by the recent works in reinforcement learning [6,2], we propose a reformulation of the non-rigid registration problem following a similar methodology as in 3-D rigid registration of [4]: in order to optimize the parameters of a deformation model we apply an artificial agent – solely learned from experience – that does not require explicitly designed similarity measures, regularization and optimization strategy. Trained in a supervised way the agent explores the space of deformations by choosing from a set of actions that update the parameters. By iteratively selecting actions, the agent moves on a trajectory towards the final deformation parameters. To decide which action to take we present a deep dual-stream neural network for implicit image correspondence learning. This work generalizes [4] to non-rigid registration problems by using a larger number of actions with a low-dimensional parametric deformation model. Since ground-truth (GT) deformation fields are typically not available for deformable registration and training based on landmark-aligned images as in rigid registration (cf. [4]) is not applicable, we propose a novel GT generator combining synthetically deformed and real image pairs. The GT deformation parameters of the real training pairs were extracted by constraining existing registration algorithms with known correspondences in the ROI in order to get the best possible organ-focused results. Thus, the main contributions of this work are: (1) The creation and use of a low-dimensional parametric statistical deformation model for organ-focused deep learning-based non-rigid registration. (2) A ground truth generator which allows generating millions of synthetically deformed training samples requiring only a few (<1000) real deformation estimations. (3) A novel way of fuzzy action control.

2 Method

2.1 Training Artificial Agents

Image registration consists in finding a spatial transformation $\mathcal{T}_\theta$, parameterized by $\theta \in \mathbb{R}^d$ which best warps the moving image $\mathbf{M}$ as to match the fixed image $\mathbf{F}$. Traditionally, this is done by minimizing an objective function of the form: $\arg \min_\theta \mathcal{F}(\theta, \mathbf{M}, \mathbf{F}) = \mathcal{D}(\mathbf{F}, \mathbf{M} \circ \mathcal{T}_\theta) + \mathcal{R}(\mathcal{T}_\theta)$ with the image similarity metric $\mathcal{D}$ and a regularizer $\mathcal{R}$. In many cases, an iterative scheme is applied where at each iteration t the current parameter value θ_t is updated through gradient descent: $\theta_{t+1} = \theta_t + \lambda \nabla \mathcal{F}(\theta_t, \mathbf{M}_t, \mathbf{F})$ where $\mathbf{M}_t$ is the deformed moving image at time step t : $\mathbf{M} \circ \mathcal{T}_{\theta_t}$.

Inspired by [4], we propose an alternative approach to optimize θ based on an artificial agent which decides to perform a simple action a_t at each iteration t consisting in applying a fixed increment $\delta\theta_{a_t}$: $\theta_{t+1} = \theta_t + \delta\theta_{a_t}$. If θ is a d -dimensional vector of parameters, we define $2d$ possible actions $a \in \mathcal{A}$ such that $\delta\theta_{2i}[j] = \epsilon_i \delta_i^j$ and $\delta\theta_{2i+1}[j] = -\epsilon_i \delta_i^j$ with $i \in \{0..d-1\}$. In other words the application of an action a_t increases or decreases a specific parameter within θ_t

by a fixed amount where δ_i^j is an additional scaling factor per dimension that is set to 1 in our experiments but could be used e.g. to allow larger magnitudes first and smaller in later iterations for fine-tuning the registration.

The difficulty in this approach lies into selecting the action a_t as function of the current state s_t consisting of the fixed and current moving image: $s_t = (\mathbf{F}, \mathbf{M}_t)$. To this end, the framework models a Markov decision process (MDP), where the agent interacts with an environment getting feedbacks for each action. In reinforcement learning (RL) the best action is selected based on the maximization of the quality function $a_t = \arg \max_{a \in \mathcal{A}} Q^*(s_t, a)$. In the most general setting, this optimal action-value function is computed based on the reward function defined between two states $\mathcal{R}(s_1, a, s_2)$ which serves as the feed-back signal for the agent to quantify the improvement or worsening when applying a certain action. Thus, $Q^*(s_t, a)$ may take into account the immediate but also future rewards starting from state s_t , as to evaluate the performance of an action a .

Recently, in RL powerful deep neural networks have been presented that approximate the optimal Q^* [6]. Ghesu *et al.* [2] used deep reinforcement learning (DRL) for landmark detection in 2-D medical images. In the rigid registration approach by Liao *et al.* [4] the agent’s actions are defined as translation and rotation movements of the moving image in order to match the fixed image.

In this work, the quality function $\mathbf{y}_a(s_t) \approx Q^*(s_t, a)$ is learned in a supervised manner through a deep regression network. More precisely, we adopt a single-stage MDP for which $Q^*(s_t, a) = \mathcal{R}(s_t, a, s_{t+1})$, implying that only the immediate reward, i.e. the next best action, is accounted for. During training, a batch of random states, pairs of $\mathbf{F}$ and $\mathbf{M}$, is considered with known transformation $\mathcal{T}_{\theta_{GT}}$ (with $\mathbf{F} \approx \mathbf{M} \circ \mathcal{T}_{\theta_{GT}}$). The target quality is defined such that actions that bring the parameters closer to its ground truth value are rewarded:

$$Q^*(s_t, a) = \mathcal{R}(s_t, a, s_{t+1}) = \|\theta_{GT} - \theta_{s_t}\|_2 - \|\theta_{GT} - \theta_{s_{t+1}}^a\|_2 . \quad (1)$$

The training loss function consists of the sum of L_2 -norms between the explicitly computed Q -values (Eq. 1) for all actions $a \in \mathcal{A}$ and the network’s quality predictions $\mathbf{y}_a(s_t)$ per action. Having a training batch $\mathcal{B}$ with random states s_b the loss is defined as: $L = \sum_{s_b \in \mathcal{B}} \sum_{a \in \mathcal{A}} \|\mathbf{y}_a(s_b) - Q^*(s_b, a)\|^2$.

In testing, the agent iteratively selects the best action, updates the parameter θ_t and warps the moving image $\mathbf{M}_t$ as to converge to a final parameter set representing the best mapping from moving to fixed image (see Fig. 1b).

2.2 Statistical Deformation Model

One challenge of the proposed framework is to find a low dimensional representation of non-rigid transformations to minimize the number of possible actions (equal to $2d$), while keeping enough degrees of freedom to correctly match images. In this work, we base our registration method on statistical deformation models (SDM) defined from Free Form Deformations (FFD). Other parametrizations could work as well. Typically, the dense displacement field is defined as the

Fig. 1: (a) Training Data Generation: Synthetic deformations (blue arrows) and inter-subject GT deformations (black) are used for intra- (green) and inter-subject (red) image pairs for training. (b) Dual-stream network used for Q -value prediction $\mathbf{y}_a$ including complete single-stage Markov Decision Process for testing (blue background).

summation of tensor products of cubic B -splines on a rectangular grid. Rueckert *et al.* [7] proposed to further reduce the dimensionality by constructing an SDM through a principal component analysis (PCA) on the B -spline displacements.

We propose to use the modes of the PCA as the parameter vector θ describing the transformation $\mathcal{T}_\theta$ that the agent aims to optimize. The agent’s basic increment per action ϵ_i is normalized according to the mean value of each mode estimated in training. To have a stochastic exploration of the parameter space, predicted actions a_t are selected given fixed probabilities (see [4]).

Fuzzy Action Control. Since parameters θ are the amplitudes of principal components, the deviation of θ_{2m} and θ_{2m+1} from the mean μ_m should stay within k -times the standard deviation σ_m in testing. In order to keep θ inside this reasonable parametric space of the SDM, we propose fuzzy action controlling. Thus, actions that push parameter values of θ outside that space, are stochastically penalized – after being predicted by the network. Inspired by rejection sampling, if an action a moves parameter θ_m to a value f_m , then this move is accepted if a random number generated between $[0, 1]$ is less than the ratio $\mathcal{N}(f_m; \mu_m, \sigma_m) / \mathcal{N}(h; \mu_m, \sigma_m)$ where $h_m = \mu_m + k\sigma_m$, and $\mathcal{N}$ is the Gaussian distribution function. Therefore, if $|f_m - \mu_m| \leq k\sigma_m$, the ratio is greater than 1 and the action is accepted. If $|f_m - \mu_m| > k\sigma_m$ then the action is randomly accepted, but with a decreased likelihood as f_m moves far away from μ_m . This stochastic thresholding is performed for all actions at each iteration and rejection is translated into adding a large negative value to the quality function $\mathbf{y}_a$. The factor k controls the tightness of the parametric space and is empirically

chosen as 1.5. By introducing fuzzy action control, the MDP gets more robust since the agent’s access to the less known subspace of the SDM is restricted.

2.3 Training Data Generation

Since it is difficult to get trustworthy ground-truth (GT) deformation parameters θ_{GT} for training, we propose to generate two different kinds of training pairs: Inter- and intra-subject pairs where in both moving and fixed images are synthetically deformed. The latter pairs serve as a data augmentation method to improve the generalization of the neural network.

In order to produce the ground truth deformations of the available training images, one possibility would be to apply existing registration algorithms with optimally tuned parameters. However, this would imply that the trained artificial agent would only be as good as those already available algorithms. Instead, we make use of manually segmented regions of interest (ROI) available for both pairs of images. By constraining the registration algorithms to enforce the correspondence between the 2 ROIs (for instance by artificially outlining the ROIs in images as brighter voxels or using point correspondences in the ROI), the estimated registration improves significantly around the ROI. From the resulting deformations represented on an FFD grid, the d principal components are extracted. Finally, these modes are used to generate the synthetic training samples by warping the original training images based on randomly drawn deformation samples according to the SDM. Amplitudes of the modes are bounded to not exceed the variations experienced in the real image pairs, similar to [7].

Intra-subject training pairs can be all combinations of synthetically deformed images of the same subject. Since the ground-truth deformation parameters are exactly known, it is guaranteed that the agent learns correct deformations. In the case of inter-patient pairs a synthetic deformed image i_{mb} of one subject I_m is allowed to be paired with any synthetic deformed image i_{nc} of any other subject I_n with b, c denoting random synthetic deformations (see Fig. 1a). Thereby, the GT parameters θ_{GT} for image pair (i_{mb}, i_{nc}) are extracted via composition of the different known deformations such that $((i_{mb} \circ \mathcal{T}_{\theta}^{i_{mb}, I_m}) \circ \mathcal{T}_{\theta}^{I_m, I_n}) \circ \mathcal{T}_{\theta}^{I_n, i_{nc}}$. Note the first deformation would require the inverse of a known deformation that we approximate by its opposite parameters for reasons of computational efficiency. The additional error due to this approximation, computed on a few pairs, remained below 2% in terms of the DICE score.

Mini-batches are created online – during training – via random image pairing where intra- and inter-subject pairs are selected with the same probabilities. Through online random pairing the experience of new pairs is enforced since the number of possible image combinations can be extremely high (e.g. 10^{12}) depending on the number of synthetic deformations.

3 Experiments

We focused on organ-centered registration of MR prostate images in 2-D and 3-D with the use case of image fusion and atlas-based segmentation [9]. The

task is very challenging since texture and anatomical appearance can vary a lot. 25 volumes were selected from the MICCAI challenge *PROMISE12*⁴ and 16 from the *Prostate-3T* database⁵ including segmentations. Same images and the cases with rectal probes were excluded. Randomly 8 cases were chosen for testing (56 pairs), 33 for training. As preprocessing, translation-based registration for all pairs was carried out in 3-D using the *elastix*-framework[3] with standard parameters followed by cropping and down sampling the images (to 100x100/75x75x20 pixels in 2-D/3-D respectively). For the 2-D experiments, the middle slice of each volume has been taken. For the purpose of GT generation mutual information as similarity metric and a bending energy metric was used. The optimization function was further constrained by a Euclidean point correspondence metric. Therefore, equally distributed points were extracted from the given mask surfaces. *elastix* was used to retrieve the solution with the weights 1, 3 and 0.2 for the above-mentioned metrics and a B-spline spacing of 16x16(x8) voxels. As a surrogate measure of registration performance we used the DICE score and Hausdorff distance (HD) on the prostate region. The extracted GT resulted in median DICE coefficients of .96 in 2-D and .88 in 3-D. Given the B-spline displacements, the PCA was trained with $d = 15$ modes in 2-D, $d = 25$ in 3-D (leading to 30 respectively 50 actions with a reconstruction error $< 5\%$ (DICE score) as a compromise to keep the number of modes relatively small.

The network’s two independent processing streams contained 3 convolutional (with 32, 64, 64 filters and kernel size 3) and 2 max-pooling layers for feature extraction. The concatenated outputs of the two streams were processed in 3 fully-connected layers (with 128, 128, 64 knots) resulting in an output with size $2d$ (equals the number of actions). Batch normalization and ReLU units were used in all layers. The mini-batch size was 65/30 (2-D/3-D). For updating the network weights, we used the adaptive learning rate gradient-based method *RMSprop*. The learning rate was 0.001 with a decay factor of 0.8 every 10k mini-batch back-propagations. Training took about 12 hours/ 1 day for 2-D and 3-D respectively. All experiments were implemented in *Python* using the deep learning library *Theano* including *Lasagne*⁶. DL tasks ran on GPUs (*NVIDIA GeForce GTX TITAN X*). During testing 200 MDP iterations (incl. resampling of the moving image) took 10 seconds (GPU) in 2-D and 90 seconds in 3-D (GPU). The number of testing steps was set empirically since registration results only change marginally when increasing the number of steps. In empirical 2-D experiments with 1000 steps the agent’s convergence was observable.

For testing, the initial translation registration was done with *elastix* by registering each of the test images to an arbitrarily chosen template from the training base. Table 1 shows that our method reaches a median DICE coefficient of .88/.76 in 2-D/3-D and therefore shows similar performance as in [3] with the best reported median DICE of .76 on a different data set. However, on our challenging test data our method outperformed the *LCC-Demons*[5] algorithm with

⁴ <https://promise12.grand-challenge.org/>

⁵ <https://wiki.cancerimagingarchive.net/display/Public/Prostate-3T>

⁶ <https://lasagne.readthedocs.io/en/latest>

Table 1: Results of prostate MR registration on the 56 testing pairs. 2-D and 3-D results in comparison to *elastix* with B-spline spacing of 8 (e8) or 16 (e16) as proposed in [3] and the *LCC-Demons*[5] algorithm (dem). T are the initial scores after translation registration with *elastix*. 3-D* are results with perfect rigid alignment T*. **nfc** are our results with no fuzzy action control (HD in mm).

	2-D				3-D					3-D*				
	T	e16	e8	our	T	e16	e8	dem	our	T*	e16	e8	dem	nfc our
DICE Mean	.76	.74	.77	.87	.62	.63	.64	.67	.75	.74	.72	.67	.79	.79 .80
DICE Med.	.78	.79	.81	.88	.61	.71	.70	.67	.76	.75	.77	.76	.80	.79 .81
DICE StD.	.10	.15	.13	.05	.11	.22	.20	.11	.06	.08	.17	.23	.07	.05 .04
HD Mean	11.6	15.2	14.5	7.7	16.1	21.2	25.3	15.9	11.8	9.2	13.4	14.5	10.4	8.9 8.0
HD Med.	11.7	13.2	13.0	7.2	15.2	18.0	21.7	15.8	11.2	9.0	11.6	13.5	10.8	8.8 7.9
HD StD.	4.3	6.8	6.7	2.5	3.9	10.7	10.9	3.9	2.9	2.3	6.8	6.4	2.5	2.2 1.9

manually tuned parameters and *elastix*, using similar parameters as proposed for prostate registration [3] using B-spline spacing of 8 and 16 pixels. We found that better rigid registration can significantly improve the algorithm’s performance as shown in the experiments with perfect rigid alignment according to the segmentation (3-D*). Extreme results are shown visually in Fig. 2.

Regarding the results of *elastix* and *LCC-Demons*, a rising DICE score was observed while HD increased due to local spikes introduced in the masks (visible in Fig. 2b) as we focused on the DICE scores during optimization for fair comparisons. In the 3-D* setting, DICE scores and HDs improved when applying fuzzy action control compared to not applying any constraints (see Table 1).

4 Conclusion

In this work, we presented a generic learning-based framework using an artificial agent for approaching organ-focused non-rigid registration tasks appearing in image fusion and atlas-based segmentation. The proposed method overcomes limitations of traditional algorithms by learning optimal features for decision-making. Therefore, segmentation or handcrafted features are not required for the registration during testing. Additionally, we propose a novel ground-truth generator to learn from synthetically deformed and inter-subject image pairs.

In conclusion, we evaluated our approach on inter-subject registration of prostate MR images showing first promising results in 2-D and 3-D. In future work, the deformation parametrization needs to be further evaluated. Rigid registration as in [4] could be included in the network or applied as preprocessing to improve results as shown in the experiments. Besides, the extension to multi-modal registration is desirable.

Disclaimer This feature is based on research and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.

(a) 2-D: Moving, Fixed, *elastix-e8* (.84), *elastix-e16* (.70), ours (.94).

(b) 3-D: Moving, Fixed, *elastix-e8* (.49), *elastix-e16* (.59), *LCC-Demons* (.67), ours (.79).

Fig. 2: 2-D and 3-D registration results of extreme cases with segmentation masks overlays (fixed: green, moving: orange) and DICE scores in parenthesis.

References

1. Dosovitskiy, A., Fischer, P., Ilg, E., Häusser, P., Hazirbas, C., Golkov, V., v.d. Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: IEEE International Conference on Computer Vision (ICCV) (2015)
2. Ghesu, F.C., Georgescu, B., Mansi, T., Neumann, D., Hornegger, J., Comaniciu, D.: An artificial agent for anatomical landmark detection in medical images. In: Medical Image Computing and Computer Assisted Intervention (MICCAI) (2016)
3. Klein, S., Staring, M., et al.: Elastix: a toolbox for intensity-based medical image registration. *IEEE transactions on medical imaging* 29(1), 196–205 (2010)
4. Liao, R., Miao, S., de Tournemire, P., Grbic, S., Kamen, A., Mansi, T., Comaniciu, D.: An artificial agent for robust image registration. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17) (2017)
5. Lorenzi, M., Ayache, N., Frisoni, G.B., et al.: LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage* 81, 470–483 (2013)
6. Mnih, V., Kavukcuoglu, K., Silver, D., others.: Human-level control through deep reinforcement learning. *Nature* 518(7540), 529–533 (2015)
7. Rueckert, D., Frangi, A.F., Schnabel, J.A.: Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. *IEEE transactions on medical imaging* 22(8), 1014–1025 (2003)
8. Simonovsky, M., Gutiérrez-Becker, B., Mateus, D., Navab, N., Komodakis, N.: A deep metric for multimodal registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 10–18. Springer (2016)
9. Tian, Z., Liu, L., Fei, B.: A fully automatic multi-atlas based segmentation method for prostate MR images. In: SPIE Medical Imaging. pp. 941340–941340 (2015)
10. Yang, X., Kwitt, R., Niethammer, M.: Fast predictive image registration. In: International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis. pp. 48–57. Springer (2016)