

LINEAR IMAGE REGISTRATION THROUGH MRF OPTIMIZATION

Ben Glocker^{1,2}, Darko Zikic¹, Nikos Komodakis³, Nikos Paragios^{2,4}, Nassir Navab¹

¹ Computer Aided Medical Procedures (CAMP), Technische Universität München, Germany

² Laboratoire MAS, Ecole Centrale Paris, Chatenay-Malabry, France

³ Computer Science Department, University of Crete, Greece

⁴ Equipe GALEN, INRIA Saclay - Ile-de-France, Orsay, France

ABSTRACT

We propose a Markov Random Field formulation for the linear image registration problem. Transformation parameters are represented by nodes in a fully connected graph where the edges model pairwise dependencies. Parameter estimation is then solved through iterative discrete labeling and discrete optimization while a label space refinement strategy is employed to achieve sub-millimeter accuracy. Our framework can encode any similarity measure, allows for automatic reduction of the degrees of freedom by simple changes on the MRF topology, and is robust to initialization. Promising results on real data and random studies demonstrate the potential of our approach.

Index Terms— Linear Image Registration, Markov Random Fields, Discrete Optimization

1. INTRODUCTION

In the last years, the solution of computer vision problems by Markov Random Fields (MRFs) [1] and discrete optimization has become increasingly popular for medical imaging applications such as segmentation [2] or non-linear registration [3]. Recent advances [4, 5] in discrete optimization methods make this approach very attractive.

However, linear image registration [6, 7], which is one of the most important and basic techniques in medical image processing has not yet been solved by discrete optimization.

In this paper we address this issue and present a linear registration method which formulates the problem as an MRF model and solves it by discrete optimization.

2. METHOD

The task of image registration is to recover a transformation T which aligns two images $I, J : \Omega \mapsto \mathbb{R}$. A common approach for modeling this problem is by energy minimization as

$$\hat{T} = \arg \min_T \xi(I, J \circ T) \quad , \quad (1)$$

The authors gratefully acknowledge the partial support of this work by Siemens Corporate Research (SCR), Princeton, NJ, USA.

where \hat{T} is the optimal transformation and ξ is an arbitrary intensity-based similarity measure. We consider T to be a linear transformation, more precisely we will focus on affine, anisotropic similarity, and rigid transformations which are often used in medical settings. There is a wide choice of similarity measures which can be employed, ranging from the sum of squared differences (SSD) to more sophisticated statistical measures such as mutual information (MI) [8, 9]. Formulating Eq. (1) as an MRF problem has the advantage that the similarity measure is easily interchangeable, since no derivatives of the measure are required. Furthermore, discrete optimization methods have generally a large capture range. Combined with intelligent discretization strategies such methods have potential for robust image registration.

In Sec. 2.1, we present the parameterization of the transformation model. Sections 2.2 and 2.3 describe our MRF model and discretization strategy and in Sec. 2.4 further features of the framework are given.

2.1. Parameterization

Affine transformations in homogeneous coordinates are linear transformations which can be written in the form

$$\begin{bmatrix} \hat{A} \\ \mathbf{0} & 1 \end{bmatrix} \quad , \quad (2)$$

where \hat{A} is a 2×3 or 3×4 matrix for 2D and 3D respectively, resulting in 6 degrees of freedom (DOF) in 2D and 12 DOF in 3D. We employ a parameterization in which the affine transformation is decomposed as

$$A = M_t R_\phi R_\theta^{-1} D_s R_\theta \quad . \quad (3)$$

Here, M_t is a translation, R_ϕ is a rotation, and $R_\theta^{-1} D_s R_\theta$ represents the shearing transformation. For the shearing, R_θ is a rotation and D_s is a diagonal matrix, representing anisotropic scaling. We parameterize the single matrices of Eq. (3) by respective parameters, compare also [10]. The 3D rotation matrices are parameterized by Euler angles. The resulting pa-




Fig. 1: MRF topology for 2D linear registration in case of rigid (left) and affine transformations (right).

parameter vectors for the 2D and 3D case are

$$\mathbf{p} = [\phi, t_x, t_y, s_x, s_y, \theta] \quad (4)$$

$$\mathbf{p} = [\phi_x, \phi_y, \phi_z, t_x, t_y, t_z, s_x, s_y, s_z, \theta_x, \theta_y, \theta_z] \quad (5)$$

An advantage of this parameterization is that it can be used also for less general transformations. Rigid transformations are obtained, if $D = R_\theta = \text{Id}$, similarity transformations are encoded for $R_\theta = \text{Id}$.¹ Similarly, pure translation, rotation, scaling, or shearing transformation, or any combination thereof can be achieved.

2.2. Markov Random Fields for Linear Registration

The key idea of our approach is to reformulate parameter estimation for linear registration as a discrete labeling problem. To this end, we construct a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consisting of a set of nodes \mathcal{V} and a set of edges \mathcal{E} . Each node $v_i \in \mathcal{V}$ corresponds to one parameter p_i of the linear transformation, while an edge $(v_i, v_j) \in \mathcal{E}$ between two nodes represents their dependency. For each node, we introduce a set of discrete labels \mathcal{L}_i which is associated with a quantized version of the parameter space Φ_i . This way, image registration becomes a labeling task where one seeks to assign an optimal label $x_i \in \mathcal{L}_i$ to each of the nodes such that the resulting transformation aligns the two images.

A common approach for modeling discrete labeling tasks in terms of energy minimization is the usage of first-order MRFs. The general form of the MRF energy is defined by two sums over potential functions

$$E_{\text{mrf}} = \sum_{v_i \in \mathcal{V}} \psi_i(x_i) + \sum_{(v_i, v_j) \in \mathcal{E}} \psi_{ij}(x_i, x_j) \quad (6)$$

where ψ_i are the *unary potentials* determining the energy of assigning a label x_i to a node v_i , and ψ_{ij} are the *pairwise potentials* determining the energy of assigning a pair of labels (x_i, x_j) to the connected nodes v_i and v_j .

Note that the unary potentials are modeling the energy for a single parameter independently of all others. Since the linear transformation parameters are highly dependent, unary potentials are not appropriate for their modeling. In contrast,

¹Please note that the similarity transformation includes anisotropic scaling which is important for inter-subject registration.

the pairwise terms are able to model the energy resulting from a simultaneous change of two parameters. Hence, we use the pairwise potentials to model the energy corresponding to a simultaneous variation of the parameters p_i and p_j as

$$\psi_{ij}(x_i, x_j) = \xi(I, J \circ T_{ij}) \quad (7)$$

where the transformation T_{ij} encodes the change in the parameters p_i and p_j .

In order to force the global consistency of the transformation parameters, we connect all nodes between each other such that we obtain a fully connected graph $\mathcal{G}^* = (\mathcal{V}, \mathcal{E}^*)$. Exemplary configurations for different types of linear transformations are illustrated in Fig. 1.

Finally, we define the MRF energy for the linear registration problem as posed in Eq. (1) as

$$E_{\text{linreg}} = \sum_{(v_i, v_j) \in \mathcal{E}^*} \psi_{ij}(x_i, x_j) \quad (8)$$

$$= \sum_{(v_i, v_j) \in \mathcal{E}^*} \xi(I, J \circ T_{ij}) \quad (9)$$

In the literature, many optimization algorithms exist for efficiently solving discrete labeling problems in form of an MRF. We use a recently proposed method called FastPD [5] which has successfully been used for non-linear registration [3]. Due to the limited space, we refer the reader to the given references for more details about the optimization algorithm.

In the following, we describe the important part of parameter space discretization.

2.3. Parameter Space Discretization

We consider the following discretization strategy for obtaining the quantized versions Φ_i of the parameter space. For each parameter, we define a minimum and maximum value which determines the parameter range and perform a uniform sampling of the range. Here, the number of labels plays a crucial role. A sparse sampling might result in inaccurate registration while dense sampling increases complexity. To overcome this problem, we employ a successive label space refinement. In each iteration k , we rescale the parameter range by a factor α^k and compute a new MRF labeling. The successive label space refinement allows us to keep the number of labels quite small and we can start with a large parameter range, while being able to achieve sub-millimeter registration accuracy.

For the computation of Eq. (7), we need to compute the potential transformations $T_{ij} = T(\hat{\mathbf{p}}_{ij})$, where $\hat{\mathbf{p}}_{ij}$ is the variation of the current parameter vector \mathbf{p} by $\Delta \mathbf{p}_{ij}$ at parameters p_i and p_j . For all parameters \hat{p}_k , except for scaling, the variation is achieved by addition, i.e. $\hat{p}_k = p_k + \Delta p_k$, while for the scaling parameters \hat{p}_s multiplication is used $\hat{p}_s = p_s \Delta p_s$.




Fig. 2: Error distribution of the 3D affine random study test. Of the 200 trials, 3 were classified as failed with an error $>10\text{mm}$. The histogram of the successful trials is presented above and the distribution has a mean of 2.02mm ($\pm 0.84\text{mm}$), and median of 1.88mm .

2.4. Further Method Features


The described registration method is performed with a standard multi-scale strategy. The method also supports random subsampling strategies, which are often used to improve the runtime, cf. e.g. [11] and references therein. Furthermore, it is easy to implement a stratified approach, which is often used in real applications [12] and consists in sequentially performing registrations of increasing complexity. Due to the flexible parameterization as described in Sec.2.1, automatic reduction of the degrees of freedom is achieved by simple changes on the MRF topology. Since our method performs these modifications automatically if the search range for one parameter is set to zero, there is no need for explicit implementation of the single registration methods.

3. EVALUATION


We test our method by random studies in 2D and 3D on mono-modal brain (Sec. 3.2, 3.1) as well as by a series of multi-modal registrations (Sec. 3.3). All experiments are performed on brain data from the the Retrospective Image Registration Evaluation Project (RIRE) database. We employ a CT image with a resolution of $512 \times 512 \times 29$ and a physical voxel size of $0.65 \times 0.65 \times 4\text{mm}$ and 3 MR images (T1,T2,PD) with a resolution of $256 \times 256 \times 26$ and a voxel size of $1.25 \times 1.25 \times 4\text{mm}$. The number of labels per label space L_i is set to 9 and we found a fixed scaling factor of $\alpha = 2/3$ to be efficient.

3.1. 3D Random Study

The 3D random study is performed on a CT volume. Random affine transformations are applied to the source image, and the warped image is used as target. The transformation parameters are uniformly drawn from certain ranges. Translations result from $\pm 10\text{mm}$, rotations lie between $\pm 20^\circ$, scalings range from $[0.8, 1.2]$, and the shearing angle parameters θ are drawn from $\pm 5^\circ$. We performed 200 experiments with the same parameters, which were set manually. We use a subsampling strategy with 5% of randomly chosen voxels and SSD as similarity measure. The registration error is computed as the average distance (AD) of the corner points of a centered 200mm cube, warped by the ground truth and the estimated



(a) Error distribution for all classes.



(b) Boxplot description of error for single test classes.

Fig. 3: Error distribution of the 2D affine random study test. (a) Error over all classes. Of the 7000 trials, 19 from class 7 had an error $>1\text{mm}$ and were classified as failed. The histogram of the successful trials is presented above and the distribution has a mean of 0.05mm ($\pm 0.03\text{mm}$), and median of 0.04mm . (b) The error distribution for single classes as boxplot. The mean is presented by the blue cross, the median by the green line.

transformation. The results are summarized in Fig. 2. The average runtime is about 3 minutes.

3.2. 2D Random Study

The 2D random study is performed on a slice from the CT image in same way as the 3D study. We additionally bin the random transformations into 7 classes of increasing complexity in order to study the capture range of the method. The classes sample the parameter space equidistantly, for translation ($[0, 35]\text{mm}$), rotation ($[0, 42]^\circ$), scaling (1 ± 0.025 to 1 ± 0.175), and shearing angle ($[0, 7]^\circ$). For example, class 7 contains translations from $[30, 35]\text{mm}$, rotations from $[36, 42]^\circ$, scalings from 1 ± 0.175 , and shearing rotation of $[6, 7]^\circ$. The error is evaluated on the corners of a 200mm centered square. We perform 1000 tests for every single class, and evaluate the registration error for the single classes as well as the complete range in Fig. 3. The average runtime is about 20 seconds.

In addition, we test the robustness w.r.t. to initialization by testing the registration for large rotations. For 1000 trials we generate targets by small translation ($\pm 10\text{mm}$) and scaling (1 ± 0.1), and a large rotation ($[90, 180]^\circ$). In 99.3% of the cases the method successfully recovers the transformation with sufficient precision for further steps (error below 20mm).




Fig. 4: Checkerboard visualization of the multi-modal registration. The errors from left to right are 1.36, 4.20, and 2.27mm.

3.3. 3D Multi-Modal Registration

Our last experiment is a multi-modal registration scenario. For all registrations, a ground truth transformation is available obtained through a fiducial landmark registration performed by the providers of the database. A CT image is registered to three MR images acquired with different protocols (T1, T2, and PD). For the tests, we use a subsampling strategy with 20% of randomly chosen voxels. We also use the stratified approach (Sec. 2.4), where translation correction is performed prior to rigid registration. Using the ground truth, we determine the error by measuring the AD for a centered cube of $334 \times 334 \times 112$ mm. The AD for the initial alignment before registration is 46.18, 56.13, and 67.61mm for T1, T2, and PD respectively. After registration the error is 1.36, 4.20, and 2.27mm. For all registrations, we use the entropy correlation coefficient (ECC) proposed in [8]. Visual results are illustrated in Fig. 4. Here, the average runtime is about 5 minutes.

4. CONCLUSION

In this paper, we have derived an MRF formulation for the linear image registration problem. The transformation model is decomposed into the individual parameters which are represented by nodes in a fully connected graph. Parameter estimation is then solved through iterative discrete labeling and discrete optimization while using a label space refinement strategy and a multi-scale approach. Our framework can encode any similarity measure without need for substantial modifications and allows for stratified registration since any subclass of affine transformation can be modeled by simple changes on the MRF topology. We evaluated our method on real data and extensive random studies. We also showed that our method is robust w.r.t. initial alignment.

In future work, performance comparison with other methods will be performed. Since the main computation is based on image warping and evaluation of the similarity measure, we believe that our method is especially suitable for GPU optimization.

5. REFERENCES

- [1] S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," in *PAMI*, 1984.
- [2] Y.Y. Boykov and M.-P. Jolly, "Interactive graph cuts for optimal boundary & region segmentation of objects in n-d images," in *ICCV*, 2001.
- [3] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, "Dense image registration through mrfs and efficient linear programming," *Med. Im. Anal.*, 2008.
- [4] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," in *ICCV*, 1999.
- [5] Nikos Komodakis, Georgios Tziritas, and Nikos Paragios, "Fast, approximately optimal solutions for single and dynamic mrfs," in *CVPR*, 2007.
- [6] J.V. Hajnal, D.J. Hawkes, and D.L.G. Hill, *Medical Image Registration*, CRC Press, 2001.
- [7] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, 2003.
- [8] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *TMI*, 1997.
- [9] Paul Viola and William M. Wells, "Alignment by maximization of mutual information," *IJCV*, 1997.
- [10] R. Hartley and A. Zisserman, "Multiple view geometry in computer vision," 2003.
- [11] JPW Pluim, JBA Maintz, and MA Viergever, "Mutual-information-based registration of medical images: a survey," *TMI*, 2003.
- [12] Luis Ibanez, Will Schroeder, Lydia Ng, and Josh Cates, *The ITK Software Guide*, Insight Software Consortium, 2005.