

Graphical Models and Deformable Diffeomorphic Population Registration Using Global and Local Metrics

Aristeidis Sotiras^{1,2}, Nikos Komodakis³, Ben Glocker^{1,4}, Jean-François Deux⁵, and Nikos Paragios^{1,2}

¹ Laboratoire des Mathématiques Appliquées aux Systèmes (MAS)
Ecole Centrale de Paris, France
aristeidis.sotiras@ecp.fr,

² Equipe GALEN, INRIA Saclay - Île de France, Orsay, France

³ Computer Science Department, University of Crete, Greece

⁴ Computer Aided Medical Procedures (CAMP), Technische Universität München
Munich, Germany

⁵ Service d'Imagerie Médicale, Centre Hospitalier Universitaire Henri Mondor, Créteil, France

Abstract. In this paper we propose a novel framework to unite a population to an optimal (unknown) pose through their mutual deformation. The registration criterion comprises three terms, the first imposes compactness on appearance of the registered population at the pixel level, the second tries to minimize the individual distances between all possible pairs of images, while the last is a regularization one imposing smoothness on the deformation fields. The problem is reformulated as a graphical model that consists of hidden (deformation fields) and observed variables (intensities). A novel deformation grid-based scheme is proposed that guarantees the diffeomorphism of the deformation and is computationally favorably compared to standard deformation methods. Towards addressing important deformations we propose a compositional approach where the deformations are recovered through the sub-optimal solutions of successive discrete MRFs by using efficient linear programming. Promising experimental results using real 2D data demonstrate the potentials of our approach.

1 Introduction

Population registration is defined as the identification of a homology between more than two images. Its importance is evident in problems like statistical modeling of variations and atlas construction. To solve the fore-mentioned problems, often a reference frame is chosen and all population members are mapped to this pose using pair-wise registration algorithms. The explicit selection of the reference image bias inherently the registration towards the chosen reference frame [1] and influence inherently its performance. Such a behavior is the opposite to the one expected towards appropriate representation of the population. Last but not least, these methods are not applicable when aiming statistical deformations population modeling using data coming from different modalities.

Methods that try to overcome the above-stated limitations can be subdivided into two classes. The first class of methods initially focuses on the appropriate selection of the reference [2] or constructs a reference template through the use of the population

statistics [3, 1]. Conventional registration methods based in pair-wise criteria are then considered towards the reference. The main limitation of these methods lies in the use of the template. The second class are template-free group-wise registration methods using either local pairwise relations or global population measurements. In [4], all possible pair-wise registrations were considered and a mean model was created by composing the deformations for each member into a mean deformation. In [5], the sum of univariate entropies along pixel stacks is introduced to address the problem of group-wise registration using an affine deformation model and is further extended by [6] to include FFD. Last but not least, in [7] local pair-wise relations were considered to deform mutually a population of images towards providing an atlas-based segmentation.

Template-driven methods introduce bias to the process through the selection of the reference, and treat individually examples of the population. On the other hand, template-free population-registration methods suffer from the lack of modularity with respect to the registration criterion and the deformation model, are sensitive to the initial conditions while being computationally inefficient. The scope of the objective function is limited to pairwise relations and computational approximations [5] are used in order to meet the high computational and memory demands.

In this paper, we propose a graphical model approach to population registration [Fig.1]. The latent variables of the model are n -deformations (Hermite-based polynomials) of the population examples and the optimal reference pose. The pose variables are connected with the observations and the corresponding deformation variables towards measuring the statistical compactness of the registration result at the pixel level. The registration variables are inter-connected and aim to decrease the cost of pair-wise comparisons between individual examples. Last, but not least the registration variables within an image are connected so as to impose smoothness. The resulting paradigm can easily encode different deformation interpolation methods, local similarity metrics and global statistical measurements while being computational efficient [when compared with the state of the art methods]. This graphical model is expressed in the form of a MRF. Towards validating the approach, we consider population registration of calf muscle MRI images.

2 Global and Local Population Registration

Let us consider n images $\{I_1, \dots, I_n\}$, where each image is described by intensity values $I_i(\mathbf{x}_i)$ for different image domains $\Omega_i, \mathbf{x}_i \in \Omega_i$. The aim of the mutual population deformation is to determine a set of transformations $\mathbf{T} = \{T_i : \mathbf{x}_R = T_i(\mathbf{x}_i), i = \{1, \dots, n\}\}$ which maps mutually corresponding points from the n -image spaces to the same point of a reference frame Ω_R . In our case, we assume the reference pose to simply correspond to the geometry and not an image template.

Deformation Model: Let us consider a grid-based deformation model that can encode different interpolation methods in a way that the transformation is one-to-one and invertible. The deformation of an object is achieved by manipulating an underlying mesh of control points. We superimpose a deformation grid $G_i : [1, K] \times [1, L]$ onto each one of the images I_i and let us also consider a grid $G_0 : [1, K] \times [1, L]$ in the reference pose. The central idea of our approach is to deform the grids simultaneously

(with a given displacement vector $\mathbf{d}_{\mathbf{p}_i^k}$ for each control point k belonging to the grid G_i) such that meaningful correspondences between the population examples are obtained and their mapping to the reference pose creates a statistically compact variable. In this case, the transformation of an image pixel $\mathbf{x}_i = (x_i, y_i) \in \Omega_i$ can be written as $T_i(\mathbf{x}_i) = \mathbf{x}_i + D_i(\mathbf{x}_i)$ where $D_i(\mathbf{x}_i) = \sum_{\mathbf{p}_i^k \in G_i} \eta(|\mathbf{x}_i - \mathbf{p}_i^k|) \mathbf{d}_{\mathbf{p}_i^k}$ and $\eta(\cdot)$ is a weighting function that measures the contribution of the control point \mathbf{p}_i^k to the displacement field D_i .

Population-wise Global Comparisons: The first term of the objective criterion to be minimized is the global statistical compactness one. We consider the intensity values of the deforming images at corresponding coordinate locations as a distribution of a random variable $\pi(\mathbf{i}(\mathbf{x}))$, where $\mathbf{i}(\mathbf{x}) = \{I_1(T_1^{-1}(\mathbf{x})), \dots, I_n(T_n^{-1}(\mathbf{x}))\}$. In statistics, one can associate a random variable to a measure of compactness with respect to this density. Examples can refer to standard deviation, higher order moments, Shannon entropy, etc. It should be expected that as the images are aligned the compactness of the probability distribution should increase. We introduce the following global measurement towards population registration

$$E_g(T_1, \dots, T_n) = \iint_{\Omega_R} \gamma(\pi(\mathbf{i}(\mathbf{x}))) d\mathbf{x} \quad (1)$$

with γ being a monotonic function inversely proportional to the compactness of the intensity distribution at \mathbf{x} once all population examples have been mapped to the reference pose. Such an objective function introduces the inverse transformation, that is challenging from theoretical and practical point of view when referring to deformable deformation. An alternative criterion that can be considered is using the forward transformations and measure the similarity of the images on the intersection of the deformed images, or

$$E_g(\mathbf{T}) = \int \dots \int_{\Omega_i \cup \dots \cup \Omega_n} \phi(T_1(\mathbf{x}_1), \dots, T_n(\mathbf{x}_n)) \gamma(\lambda(\mathbf{x}_1, \dots, \mathbf{x}_n)) d\mathbf{x}_1 \dots d\mathbf{x}_n \quad (2)$$

where $\lambda(\mathbf{x}_1, \dots, \mathbf{x}_n) = \pi(I_1(T_1(\mathbf{x}_1)), \dots, I_n(T_n(\mathbf{x}_n)))$ and ϕ is a Dirac-driven function whose role is to define which pixels correspond to the same position at the reference pose defined as follows: $\prod_{(i,j) \in [1,n] \times [1,n]} \delta_\alpha(|\mathbf{x}_i - \mathbf{x}_j|)$.

Population-wise Local Comparisons: It may be the case that a distribution exhibits good compactness characteristics globally but certain members of the population can always be placed in the tale of the distribution. To avoid such cases, local pair-wise comparison between the members of the population are going to be considered. Let $\rho_{ij}(\cdot)$ be a similarity measurement used to compare the visual information for the images i and j . Then, if (without loss of generality) we consider for example pixel-based measurements, the pair of forward deformations T_i, T_j , should minimize the distance in the intersection of the deformed images:

$$E_l(T_i, T_j) = \iint_{\Omega_i \cup \Omega_j} \phi(|T_i(\mathbf{x}_i) - T_j(\mathbf{x}_j)|) \rho_{ij}(I_i(T_i(\mathbf{x}_i)), I_j(T_j(\mathbf{x}_j))) d\mathbf{x}_i d\mathbf{x}_j \quad (3)$$

In simple words, this quantity evaluates the pertinence of the correspondences between the two images using both definition domains Ω_i, Ω_j where only the pixels for which




Fig. 1. The node and the edge system of the constructed graph. With blue color the relationship between the grid nodes and the images is depicted (deformation model). The black edges represent the smoothness terms while the red ones encode the local dissimilarity measure. The global relationship between all the nodes at respective places in the grids is shown by the yellow edges. (For clarity a fraction of the edges is shown.)

correspondences between the two images have been found are considered. The criterion can be extended to deal with the case of n -images by simply considering all possible pairs of images.

Smoothness Constraints: Medical images capture properties of spatially continuous anatomical structures, therefore it is natural to assume that the deformation applied to them should be locally smooth. Opposite to the former cases, this constraint should be applied to each grid separately. This constraint can be defined on the grid as

$$E_s(T_1, \dots, T_n) = \sum_{i=1}^n \iint_{\Omega_i} \psi(\nabla_{T_i}(\mathbf{x}_i)) d\mathbf{x}_i \quad (4)$$

where ψ is a convex function imposing smoothness.

The optimal parameters of the deformation should be determined through the minimization of an objective function being composed of the above terms. Gradient descent method is the most common approach, but is unable to guarantee the recovery of the global minimum, is computationally inefficient, and far from being modular. Graphical models and the off-the-shelf discrete optimization methods being associated to them can address the above mentioned constraints.

3 Graphical Model towards Population Registration

In order to be able to use discrete optimization schemes the deformation space should be quantized. Let $\Theta = \{\mathbf{d}^1, \dots, \mathbf{d}^q\}$ be a quantized version of the deformation field, then a discrete set of labels $L = \{l^1, \dots, l^q\}$ can be corresponded to it. A label assignment $l_{\mathbf{p}}^{\xi}$, where $\xi \in \{1, \dots, q\}$, to a grid node \mathbf{p} is associated with displacing the node by the corresponding vector $\mathbf{d}_{\mathbf{p}}^{l_{\mathbf{p}}^{\xi}}$. If a label is assigned to every node we get a discrete labeling \mathbf{l} . The displacement field associated with a certain labeling \mathbf{l} becomes $D(\mathbf{x}) = \sum_{\mathbf{p} \in G} \eta(|\mathbf{x} - \mathbf{p}|) \mathbf{d}_{\mathbf{p}}^{l_{\mathbf{p}}}$. We have considered the Hermite splines. In this case $D = \sum_{l=0}^1 \sum_{m=0}^1 H_l(u) H_m(v) \mathbf{d}_{i+l, j+m}$. $i = \lfloor x/\delta_x \rfloor$, $j = \lfloor y/\delta_y \rfloor$, $u = x/\delta_x - \lfloor x/\delta_x \rfloor$ and $v = y/\delta_y - \lfloor y/\delta_y \rfloor$. H_l represents the l th basis function of the Hermite spline and $\delta_x = \frac{M}{K-1}$, $\delta_y = \frac{N}{L-1}$ denotes the control point spacing. Hermite

splines involve less computations than cubic- B splines while exhibiting the same desired properties.

By applying this quantization of the deformation space one would like to reformulate the problem as a discrete multi-labeling problem. A common model for representing such problems are Graphical Models and MRFs. In the context of population registration, the graphical model will involve three terms, one singleton that measures the compactness and two pair-wise, one that account for smoothness at each deformation field and one that enforces pair-wise correspondences.

$$E_{GM}(G_0, T_1 \circ G_1, \dots, T_n \circ G_n) = \alpha \sum_{i=0}^n \sum_{\mathbf{p} \in G_0} V_{\mathbf{p}}(l_{\mathbf{p}}) + \beta_{intra} \sum_{i=0}^n \sum_{\mathbf{p} \in G_i} \sum_{\mathbf{q} \in (N(\mathbf{p}) \cap G_i)} V_{\mathbf{p}\mathbf{q}}(l_{\mathbf{p}}, l_{\mathbf{q}}) + \beta_{inter} \sum_{i=0}^n \sum_{\mathbf{p} \in G_i} \sum_{\mathbf{q} \in (N(\mathbf{p}) \setminus G_i)} V_{\mathbf{p}\mathbf{q}}(l_{\mathbf{p}}, l_{\mathbf{q}}) \quad (5)$$

where $V_{\mathbf{p}}(\cdot)$ are the unary potentials, $V_{\mathbf{p}\mathbf{q}}(\cdot, \cdot)$ are the pair-wise potentials and N represents the neighborhood system of the nodes [Fig.1]. α , β_{inter} and β_{intra} are weighting constants. The main challenge of discrete optimization methods is the quantization of the search space since it seeks for a compromise between computational complexity and the ability to capture a good minimum. This can be achieved through a compositional approach, where the final solution is obtained through successive optimization problems with respect to the deformation increment towards minimizing the objective function [8]. Thus, by keeping the set of the labels in a reasonable size it becomes possible to approximate the optimal solution in an efficient way.

3.1 Mapping of the Objective Function to the Graphical Model

Mapping global, local and smoothness costs to the graphical model consists of converting them to singleton and pair-wise terms. The most challenging case is the global cost due to the fact that in order to be properly determined it requires higher order cliques. The mapping of the other two terms is straightforward.

Singleton Term: The adoption of higher order cliques is possible within MRFs, however their use decreases significantly their computational efficiency. We consider an approximation of the global cost that consists of assuming that for a given node \mathbf{p} of a given deformation field/image i , the rest of the images do not move within the current iteration. This assumption is considered for all nodes, and for all deformation fields within a given iteration and therefore is not restrictive and quite common in minimizing graphical models through expansion moves. Then, the cost of a deformation will depend only on the label of this node, or,

$$V_{\mathbf{p}_i^k}^t(l_{\mathbf{p}_i^k}) \approx \int \dots \int_{\Omega_1 \cup \dots \cup \Omega_n} \eta_s^{-1}(\mathbf{x}_i, \mathbf{p}_i^k) \phi(T_1^{t-1}(\mathbf{x}_1), \dots, T_i^t(\mathbf{x}_i), \dots, T_n^{t-1}(\mathbf{x}_n)) \gamma(\lambda(I_1(T_1^{t-1}(\mathbf{x}_1)), \dots, I_i(T_i^t(\mathbf{x}_i)))) d\mathbf{x}_1 \dots d\mathbf{x}_n \quad (6)$$

where $\eta_s^{-1}(\mathbf{x}_i, \mathbf{p}_i^k) = \eta_s^{-1}(|\mathbf{x}_i - \mathbf{p}_i^k|) = \frac{\eta(|\mathbf{x}_i - \mathbf{p}_i^k|)}{\int_{\Omega_i} \eta(|\mathbf{y}_i - \mathbf{p}_i^k|) d\mathbf{y}_i}$. We have considered a congealing-like global cost that aims at minimizing the entropy of the pixel distributions upon registration. This term corresponds to the G_0 graphical model variables.




Fig. 2. Results obtained for the muscle image data set (mean and variance image). From left to right, the initial images, the result of the group-wise registration, the result of two pair-wise registrations.

Pair-Wise Terms: Two different cases have to be discerned, one that accounts for pair-wise registrations between all image pairs and one that imposes smoothness on the deformation fields. The adaptation of the local registration costs involves connections between the nodes $\mathbf{p}_i^k, \mathbf{q}_j^k$, that are placed in respective places k in grids that belong to two different images i and j . The inter pair-wise potential are defined as

$$V_{\mathbf{p}_i^k \mathbf{q}_j^k}(l_{\mathbf{p}_i^k}, l_{\mathbf{q}_j^k}) \approx \int_{\Omega_i \cup \Omega_j} \eta_p^{-1}(\mathbf{x}_i, \mathbf{p}_i^k, \mathbf{x}_j, \mathbf{q}_j^k) \phi(|T_i(\mathbf{x}_i) - T_j(\mathbf{x}_j)|) \cdot \rho(I_i(T_i(\mathbf{x}_i)), I_j(T_j(\mathbf{x}_j))) d\mathbf{x}_i d\mathbf{x}_j \quad (7)$$

where η^{-1} are inverse projection functions that depend on the distances between the pixel and the different deformation grids and are defined as: $\eta_p^{-1}(\mathbf{x}_i, \mathbf{p}_i^k, \mathbf{x}_j, \mathbf{q}_j^k) = \frac{\eta(|\mathbf{x}_i - \mathbf{p}_i^k|) \eta(|\mathbf{x}_j - \mathbf{q}_j^k|)}{\int_{\Omega_i \cup \Omega_j} \delta(T_i(\mathbf{y}_i), T_j(\mathbf{y}_j)) \eta(|\mathbf{y}_i - \mathbf{p}_i^k|) \eta(|\mathbf{y}_j - \mathbf{q}_j^k|) d\mathbf{y}_i d\mathbf{y}_j}$. The image metric used in the context of pair-wise image comparisons of our approach was the sum of absolute differences.

Last, but not least imposing smoothness on the deformation fields can be done by defining a distance function computing the magnitude of vector differences [8]

$$V_{\mathbf{p}_i \mathbf{q}_i}(l_{\mathbf{p}_i}^\xi, l_{\mathbf{q}_i}^\nu) = |\mathbf{d}^{\mathbf{p}_i} - \mathbf{d}^{\mathbf{q}_i}|.$$

To minimize the successive MRFs, that is to assign a label l to all the nodes \mathbf{p} of the constructed graph, an efficient linear programming method is used [9]. The last constraint to be addressed refers to the diffeomorphic property of the proposed population registration framework. This can easily be introduced by imposing hard constraints to the allowed deformations [10]. Following [11], the bound for the maximum displacement towards guaranteeing diffeomorphic deformations, in the case of the cubic Hermite spline, is proven to be 0.25 times the grid spacing.

4 Experimental Validation

To evaluate the performance of the method, the population registration of 2D MR human skeletal muscle calf images has been considered. The images were acquired with




Fig. 3. Comparison between group-wise and pair-wise registration. The Dice coefficients obtained through pair-wise registration with respect to all plausible individual template choices are compared with the population registration result.

a 1.5T Siemens scanner, with parameters $TR=711$, $TE=11$. Each volume consists of 90 slices of 4mm thickness with voxel spacing of $0.7812 \times 0.7812 \times 4$ mm. From the original volumes slices that correspond to respective positions were selected to for the data set. Segmentations for the data were provided by an expert. The parameters of the method, $\alpha, \beta_{inter}, \beta_{intra}$ were set to 10, 1 and 0.1, respectively. We have used a multi-scale implementation with 3 levels, an initial grid resolution of 8×8 , and a final one of 32×32 . A number of $2 \times 4 + 1$ labels were used per iteration cycle, sampled along the principal horizontal and vertical directions.

The qualitative results of the group-wise registration of the muscle data are presented in [Fig.2]. Comparing visually the mean and the variance image of the population before and after the group-wise registration the success of the registration process can be assessed qualitatively. The mean image is far more sharp than the one before the registration process, while the variance image emphasizes the decrease of the intensity differences along the registered data.

To further appraise the performance of the proposed method, it was compared to a state of the art pair-wise registration method [8]. Similar parameters and deformation grids were used for both methods with the difference that for the group-wise registration scheme, Hermite weighting functions were used instead of cubic B -splines. The performance of the pairwise registration was exhaustively evaluated as all possible images were used as targets. The distributions of the Dice values for each image target are reported in [Fig.3], where a boxplot is given for every image target. The results for the pair-wise registration are given from column 1 to 18, while the last column corresponds to the results obtained by the proposed group-wise registration framework. By simple observation, it can be concluded that the group-wise registration outperforms the pair-wise method for the majority of the cases.

The results depicted in the graph suggest that considering the population as a whole and registering subjects jointly brings the population into better alignment than matching each subject to a target image. This is implied by the decrease of the dispersion of

the Dice values that is observed in the group-wise case. The results presented in the figure [Fig. 3] point out the intrinsic drawbacks of the pair-wise registration process whose performance is greatly influenced by the choice of the target image.

A Matlab implementation of our approach takes approximately 30 min on an Apple Mac with 4GB memory and 2.5GHZ Processor, for a population registration of 20 examples (256×256) and a final resolution grid of 32×32 per image. However, since our graph is similar to the one in [8] and the same optimization technique is used, a C++ implementation should decrease the running time to a couple of minutes.

5 Discussion

In this paper we have proposed a novel approach to unbiased diffeomorphic deformable population registration using graphical models and discrete optimization. Our approach is gradient free, modular in terms of the image and smoothness components and can encode global population criteria and pair-wise comparisons.

The extension of the method to deal with 3D data is natural and straightforward future direction. Furthermore, the use of higher-order MRFs towards proper approximation of the global costs will improve the performance of the method in terms of ability to capture the global optimum. Last, but not least the ability to construct an unbiased statistical anatomical atlas using the proposed concept could be a useful tool in a number of applications in medical imaging.

References

1. Bhatia, K., Hajnal, J., Puri, B., Edwards, A., Rueckert, D.: Consistent groupwise non-rigid registration for atlas construction. In: ISBI. (2004)
2. Park, H., Bland, P., Hero, A., Meyer, C.: Least biased target selection in probabilistic atlas construction. In: MICCAI. (2005)
3. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. Neuroimage (2004)
4. Seghers, D., D'Agostino, E., Maes, F., Vandermeulen, D., Suetens, P.: Construction of a brain template from mr images using state-of-the-art registration and segmentation techniques. In: MICCAI. (2004)
5. Zollei, L., Learned-Miller, E., Grimson, E., Wells, W.: Efficient population registration of 3d data. In: ICCV. (2005)
6. Balci, S., Golland, P., Shenton, M., Wells, W.: Free-form b-spline deformation model for groupwise registration. In: MICCAI. (2007)
7. Sotiras, A., Komodakis, N., Langs, G., Paragios, N.: Atlas-based deformable mutual population segmentation. In: ISBI. (2009)
8. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through mrfs and efficient linear programming. MIA (2008)
9. Komodakis, N., Tziritas, G., Paragios, N.: Performance vs computational efficiency for optimizing single and dynamic mrfs: Setting the state of the art with primal-dual strategies. CVIU (2008)
10. Rueckert, D., Aljabar, P., Heckemann, R., Hajnal, J., Hammers, A.: Diffeomorphic registration using b-splines. In: MICCAI. (2006)
11. Choi, Y., Lee, S.: Injectivity conditions of 2d and 3d uniform cubic b-spline functions. Graphical Models (2000)