A Discrete MRF Framework for Integrated Multi-Atlas Registration and Segmentation

Stavros Alchatzidis, Aristeidis Sotiras, Evangelia I. Zacharaki, Nikos Paragios

► To cite this version:

Stavros Alchatzidis, Aristeidis Sotiras, Evangelia I. Zacharaki, Nikos Paragios. A Discrete MRF Framework for Integrated Multi-Atlas Registration and Segmentation. International Journal of Computer Vision, 2017, 121, pp.169-181. 10.1007/s11263-016-0925-2 . hal-01359094

HAL Id: hal-01359094

<https://hal.science/hal-01359094>

Submitted on 2 Sep 2016

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.

A Discrete MRF Framework for Integrated Multi-Atlas Registration and Segmentation

Stavros Alchatzidis · Aristeidis Sotiras · Evangelia I. Zacharaki · Nikos Paragios

Received: date / Accepted: date

Abstract Multi-atlas segmentation has emerged in recent years as a simple yet powerful approach in medical image segmentation. It commonly comprises two steps: i) a series of pairwise registrations that establish correspondences between a query image and a number of atlases, and ii) the fusion of the available segmentation hypotheses towards labeling objects of interest. In this paper, we introduce a novel approach that solves simultaneously for the underlying segmentation labels and the multi-atlas registration. The proposed approach is formulated as a pairwise Markov Random Field, where registration and segmentation nodes are coupled towards simultaneously recovering all atlas deformations and labeling the query image. The coupling is achieved by promoting the consistency between selected deformed atlas segmentations and the estimated

query segmentation. Additional membership fields are estimated, determining the participation of each atlas in labeling each voxel. Inference is performed by using a sequential relaxation scheme. The proposed approach is validated on the IBSR dataset and is compared against standard post-registration label fusion strategies. Promising results demonstrate the potential of our method.

Keywords multi-atlas segmentation · medical imaging · Markov Random Fields · discrete optimization

1 Introduction

Segmentation, or the process of assigning voxels to distinct anatomical regions or tissue types, is a fundamental task in medical image analysis. The accurate delineation of anatomical structures is the cornerstone of quantitative analysis that aims, among other, to understand normal and diseased anatomical variability. The increasing availability and size of high resolution imaging data along with the widespread adoption of imaging across clinical and research practice further underline the importance of reliable, efficient and accurate image segmentation. While manual segmentation is the gold standard, the complexity of the task, as well as the high time requirement and the associated cost, make it prohibitive on a large scale. In such a setting, automatic image segmentation provides a valuable alternative.

As a consequence, important research efforts have been focused on developing automatic segmentation algorithms. Among the wealth of the developed techniques, segmentation via registration [26] stands as a unique example in medical image processing. Registration is used to map the grayscale image of an atlas to the query image, while the estimated deformation

S. Alchatzidis
Equipe GALEN, INRIA Saclay, Île-de-France, Orsay, France.
Center for Visual Computing, Department of Applied Mathematics,
Ecole Centrale de Paris, 92295 Châtenay-Malabry, France.
E-mail: stavros.alchatzidis@ecp.fr

A. Sotiras
Section of Biomedical Image Analysis, Department of Radiology, University of Pennsylvania, Pennsylvania, USA
E-mail: aristeidis.sotiras@uphs.upenn.edu

E. I. Zacharaki
Equipe GALEN, INRIA Saclay, Île-de-France, Orsay, France.
Center for Visual Computing, Department of Applied Mathematics,
Ecole Centrale de Paris, 92295 Châtenay-Malabry, France.
E-mail: evangelia.zacharaki@centralesupelec.fr

N. Paragios
Equipe GALEN, INRIA Saclay, Île-de-France, Orsay, France.
Center for Visual Computing, Department of Applied Mathematics,
Ecole Centrale de Paris, 92295 Châtenay-Malabry, France.
E-mail: nikos.paragios@ecp.fr

is subsequently used to propagate the available labels and provide an estimate of the segmentation. However, a single atlas is limited with respect to its ability to capture the variability of the population.

Extending registration-based segmentation by incorporating multiple atlases has thus emerged as a natural extension towards tackling the aforementioned limitation [15]. Multi-atlas segmentation, fueled by the maturity of the available registration techniques and advancements in computer hardware that partly alleviate its high computational cost, has gained significant popularity and found numerous applications in medical image analysis. For example, one may cite full brain segmentation [1, 3, 5, 12, 28], skull-stripping [10], hippocampal [9, 34], prostate [11, 18] and heart [16, 24] delineation.

Multi-atlas segmentation methods produce state of the art results in various settings by essentially applying a common pipeline. The pipeline comprises two basic steps. First, rigid or non-rigid registration is performed in order to align the atlas images to the query one and propagate their labels forming a set of candidate segmentations. Subsequently, the derived multiple segmentation hypotheses are fused to produce the final labeling of the query image.

One of the most popular fusion strategies is the majority voting one [26]. Since early approaches in multi-atlas segmentation, majority voting has been established as an intuitive, simple and robust baseline method. Many efforts have concentrated on developing more sophisticated fusion algorithms by incorporating additional information in terms of either local appearance [3], neighborhood information [13, 6], or by exploiting the statistical correlation of inter-atlas errors to calculate optimal weights [34].

Alternative fusion techniques adopt probabilistic models to derive per voxel weights. Modeling the segmentation candidates as noisy observations of the true segmentation and using Expectation-Maximization estimation of the per atlas confusion matrix has been proposed in [36, 27]. This method has been extended in many ways, for example including a smoothness prior [35], using appearance-based ranking [17], incorporating non-local intensity proposals [6] and performing iterative atlas selection [21]. In another related approach, a generative probabilistic model of label fusion was proposed in [30].

A common principle behind most multi-atlas segmentation approaches is that registering the atlas images to produce candidate solutions and segmenting the query image are treated separately, in two independent steps. However, registration could benefit from taking into account the underlying segmentation towards establishing more accurate correspondences. Thus, ap-

proaches that treat registration and segmentation through fusion as an inter-weaved process have recently appeared.

In [7] the authors segmented white and gray matter areas of the spinal cord by iteratively, affinely registering atlas masks to an estimated target segmentation, and fusing the registered masks and images to get a new estimate of the target segmentation. In a similar context, a probabilistic generative model based on the Demons registration framework was proposed by [12]. The model coupled registration by discouraging atlas to target transformations that disagree with atlas to precomputed mean template to target transformations. The model also imposed a distance transform based prior on the target mask, a membership field that corresponds voxels to atlases and a mixture of gaussians appearance model linking anatomical regions to target intensity probabilities. An approximate variational EM scheme was used to find the most probable model parameters. A similar work was presented in [33], where the Large Deformation Diffeomorphic Metric Mapping registration framework was employed.

Apart from the multi-atlas process our method can be seen in the same perspective as classical joint segmentation/registration methods. A first joint approach for voxel-based registration and segmentation is proposed in [37]. The joint class histogram is used to drive the Mutual Information metric for the registration step and a Gaussian mixture model is inferred to drive the segmentation step. Validation is carried out over rigid segmentations and 2D images. Application in medical images as well as a generative model to link the two processes can be found in [4] where intensity uniformity as well as spatial priors are introduced in the segmentation model. Use of MRFs can be seen again in the works of Xiaohua et al. first on MRI data [39] and then in a formulation for contrast enhanced breast MRI scans [38]. Recent works include lung [40] and prostate segmentation [22], as well as brain lesions identification in MRI data [32, 23].

We complemented previous approaches for integrated registration and label fusion segmentation in [2] by introducing a discrete formulation based on Markov Random Field theory. Latent variables include the displacements of the grid nodes of a B-Spline transformation model as well as voxel segmentation variables. Segmentation additionally takes into account class likelihoods produced by a discriminatively trained classifier. Constraints were imposed by taking into account how congruent the proposed segmentations are with respect to the proposals of the rest of the atlases as well as the classifier produced likelihoods. Registration and segmentation variables are coupled by using an appearance-based weighting similar to the one used in local fusion strate-

gies. As a consequence, as votes are weighted, atlases that do not match well locally will have a minor contribution to the inferred segmentation mask, resulting in a local "soft" atlas selection scheme.

In this work, we build upon the work presented in [2] by extending the formulation, providing a more detailed description of the method and reporting a more extensive validation setting. In this work, we estimate membership fields by introducing a local atlas selection scheme. This scheme, explicitly models variables for selecting parts of each atlas labeling by comparing them directly to the estimated underlying segmentation. As a consequence, membership field images are directly produced by optimizing the model. The fact that the atlas selection is achieved by comparing segmentations, facilitates the use of images of different modalities to be part of the dataset.

The remainder of this paper is organized as follows. In Section 2 we formulate the problem in the continuous domain, while in Section 3 we present the decomposition of the objective function toward a mapping into a discrete graph structure. Experimental validation is discussed in Section 4, while Section 5 concludes the paper.

2 Problem Formulation

We consider a dataset of N annotated images $\mathbf{A} = \{A_0, \dots, A_{N-1}\}$. Each image comes with a corresponding segmentation mask where the anatomical regions of interest have been annotated, forming the set $\mathbf{S} = \{S_0, \dots, S_{N-1}\}$. Each voxel in the segmentation image is assigned to a segmentation label corresponding to one of M anatomical classes, $S_i(x) \in \{0, \dots, M-1\}$. In this paper, we refer to an *atlas* as the aggregation of an intensity image A_i and its corresponding segmentation mask S_i .

Moreover, we consider that an image I is given as input to be segmented into anatomical regions. Henceforth, we are going to interchangeably refer to this image as either target or query image. The output of the proposed algorithms comprises: i) a set of membership field images $\mathbf{F} = \{F_0, \dots, F_{N-1}\}$, $F_i(x) \in \{0, 1\}$ denoting if an atlas influences a point x in the query image; ii) the segmentation mask S_I corresponding to the target image; and iii) a set of deformation fields $\mathbf{D} = \{D_0, \dots, D_{N-1}\}$, where D_i denotes the deformation field mapping of A_i to I .

2.1 Method Outline

The goal of the proposed method is to simultaneously solve for the parameters of the final segmentation S_I of the query image and the set of deformation fields $\mathbf{D}$. Hence, the agreement of the deformed segmentation masks with the underlying estimated segmentation may be taken into account during the estimation of the deformation fields, leading to more accurate correspondences and consequently, improved segmentation. Our basic premise is that *by allowing the two problems to interact, the quality of the respective solutions will rise due to the additional available information*.

Furthermore, we aim to improve the final segmentation S_I by taking into account class specific appearance priors. The motivation behind incorporating prior segmentation probabilities lies in that fact that image registration is often trapped in local minima when trying to match areas of high anatomical variability (*e.g.*, brain cortex). In such a setting, appearance information constitutes an alternative, more reliable cue that can robustly guide segmentation [31].

2.2 Continuous Energy Formulation

We formulate the problem as an energy minimization one. The proposed energy consists of three components: i) a registration component comprising a matching term (M), that quantifies the level of alignment between each atlas and the query image, and a regularization term (R_d) that enforces the smoothness of the deformation field; ii) a segmentation component comprising an appearance prior term (S_P), that measures the log-likelihood of the segmentation with respect to the probabilities (π) learned during a training phase. iii) a coupling term (C) that takes into account the labeling that is proposed by the atlases, over the domain indicated by $\mathbf{F}$, and encourages their agreement with the estimated segmentation S_I . Finally, a regularization term (R_f) is imposed on the membership fields allowing the atlases to influence the derived segmentation in a smooth spatially varying fashion.

The energy has the following form:

$$E(\mathbf{D}, \mathbf{A}, \mathbf{S}, S_I) = \underbrace{M(\mathbf{D}, \mathbf{A}, I) + R_d(\mathbf{D})}_{\text{Registration}} + \underbrace{S_P(S_I, \pi)}_{\text{Segmentation}} + \underbrace{C(\mathbf{D}, \mathbf{F}, I, \mathbf{S}, S_I) + R_f(\mathbf{F})}_{\text{Coupling}} \quad (1)$$

The first two energy terms correspond to the standard energy that is commonly minimized in multi-atlas segmentation frameworks, while the third term is common

in segmentation frameworks. The fourth term introduces the main novelty of this work, *i.e.* the coupling between the segmentation and the multi-atlas registration. Let us now detail each term of the previous energy.

2.2.1 Registration

The registration component aims to align all atlases to the query image. As such, it is a generalization of standard pairwise image registration.

Matching criterion Given any dissimilarity intensity-based criterion ρ , the matching term takes the following form:

$$M(\mathbf{D}, \mathbf{A}, I) = \sum_{i=0}^{N-1} \int_{\Omega} \rho(I, A_i \circ D_i(x)) dx. \quad (2)$$

This term is the summation of the independently evaluated dissimilarity criteria between all atlases and the query image I .

Deformation smoothness Image registration is an ill-posed problem according to Hadamard's definition. In order to account for this fact, regularization is necessary. Typical regularization settings penalize the non-smoothness of deformation fields. Thus, given a smoothness inducing function ψ , the regularization term takes the following form:

$$R_d(\mathbf{D}) = \sum_{i=0}^{N-1} \int_{\Omega} \psi(D_i(x)) dx. \quad (3)$$

In other words, this term evaluates the smoothness of all deformation fields mapping from an atlas to the target image and sums the independent evaluations.

Transformation model In this work, the popular Free Form Deformations (FFDs) transformation model [29, 20] is used. Free Form Deformations parametrize the transformation $D_j(x)$ by a linear combination of K control points:

$$D_j(x) = x + \sum_{i=0}^{K-1} \omega_i(x) \phi_i, \quad (4)$$

where ϕ_i is the displacement of control point i and $\omega_i(x)$ is an interpolation or weighting function that determines the influence of the control point i to the image point x . In the current approach, we use N uniformly distributed grids of control points superimposed over the image domain (one corresponding to each atlas) and cubic B-splines as the weighting functions.

2.2.2 Segmentation

Segmentation Prior The quality of the segmentation hypotheses that are provided by warping the given atlases is conditioned upon the quality of the registration. Image registration is often overwhelmed when trying to establish correspondences between highly variable anatomical regions leading to inaccurate results. In such cases, one can exploit additional cues to enhance segmentation estimation. Local appearance provides complementary information that can be incorporated into segmentation. Assuming a probability function on the candidate labeling of the form $\pi_x(l)$ where $l \in \{0, \dots, M-1\}$, we aim to penalize all segmentations that go against prior information:

$$S_P(S_I, \pi) = \int_{\Omega} -\log(\pi_x(S_I(x))) dx. \quad (5)$$

Such probabilities can be learned using any modern classification method.

2.2.3 Coupling

In this work, we close the circuit between multi-atlas registration and label fusion. We allow segmentation to influence registration by imposing agreement constraints between the estimation of the underlying segmentation and the atlas hypotheses. As a consequence, improved segmentation accuracy may be achieved through the refinement of the registration result.

The above is modeled by means of introducing a third term in the energy. This term penalizes in a controlled manner deformations that lead to disagreement between the hypotheses and the estimated segmentation. Specifically the penalties are not defined in advance but are updated based on the segmentation consistency and smoothness of membership fields at each iteration.

$$C(\mathbf{D}, \mathbf{F}, I, \mathbf{S}, S_I) = \sum_{i=0}^{N-1} \int_{\Omega} F_i(x) \kappa(S_i \circ D_i(x), S_I(x)) dx, \quad (6)$$

where $\kappa(a, b)$ is a penalizing disagreement function, for which we assume $\kappa(a, a) = 0$.

2.2.4 Membership field smoothness

We assume that only a spatially varying subset of the dataset is pertinent to deducing the correct segmentation mask at any point. Dealing with images, an important natural prior is spatial smoothness:

$$R_f(\mathbf{F}) = \sum_{i=0}^{N-1} \int_{\Omega} \psi(F_i(x)) dx, \quad (7)$$

where ψ is again a smoothness inducing function.

3 Markov Random Field Formulation

We use Markov Random Field (MRF) theory to formulate the above minimization problem in a discrete context. The problem is represented by a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ denotes the set of nodes that encode the latent variables, and $\mathcal{E}$ the set of edges that encode the interactions between the variables.

The graph is associated with an energy of the form:

$$E_{MRF}(\mathbf{l}) = \sum_{p \in \mathcal{V}} g_p(l_p) + \sum_{(p,q) \in \mathcal{E}} f_{pq}(l_p, l_q), \quad (8)$$

where random variables p take values from a discrete set of solutions $\mathcal{L}$, $g_p(l_p)$ measures the cost of assigning a value l_p to the variable p and $f_{pq}(l_p, l_q)$ is a pairwise function that determines the cost of assigning different values l_p and l_q to the variables p and q .

3.1 Graph Structure

The constructed graph should encode the multi-atlas registration, the segmentation and the constraints that integrate the two problems. Let us now detail how the graph is constructed to achieve this.

Multi-Atlas Registration Let us recall that the deformation model is parametrized by N deformation grids. This is encoded in the MRF graph $\mathcal{G}$ by a set of N isomorphic grid graphs $\mathcal{G}_D = \{\mathcal{G}_{D_0}, \dots, \mathcal{G}_{D_{N-1}}\}$. For every control point in the deformation grid that is superimposed onto image A_i , there is a node $p_i \in \mathcal{V}_{D_i}$ that represents its displacement. Since grids are isomorphic, p indexes a common control point position, while i indexes the grid. The edge system of each grid $\mathcal{E}_{D_i}$ is created by assuming a regular 6-connectivity scheme. The label set $\mathcal{L}_D$ for this set of variables is a quantized version of the displacement space. A label assignment $l_{p_i}^d \in \mathcal{L}_D$ (with $p_i \in \mathcal{V}_{D_i}$) is equivalent to displacing the control point p_i by displacement $\mathbf{d}_{p_i}$.

Segmentation An additional graph $\mathcal{G}_S = (\mathcal{V}_S)$ is employed to model segmentation. A node $p_s \in \mathcal{V}_S$ encodes a random variable and corresponds to a voxel in the target image whose position is indexed by the subscript s . We should also emphasize the fact that the nodes that form the segmentation graph are **not** connected to one another. The set of possible solutions $\mathcal{L}_S$ represents the set of anatomical regions augmented by the background label. We refer to a potential anatomical label in $\mathcal{L}_S$ by l^s .

Coupling Integrating segmentation and multi-atlas registration is achieved by coupling segmentation and deformation graphs. The set of edges $\mathcal{E}_C$ connects nodes of $\mathcal{V}_S$ with nodes of $\mathcal{V}_D$. In order to create the coupling edge system, we connect every node $p \in \mathcal{V}_D$ with nodes of $\mathcal{V}_S$ that correspond to voxels that are influenced by a displacement of p .

Local Atlas Selection We parametrize membership fields by taking into account the spatial support of the deformation nodes. Voxels within the support of a control point share the same membership state ($[0, 1]$). To model this, we augment the label set of the deformation nodes by considering the Cartesian product between the deformation label set $\mathcal{L}_D$ and a binary selection label set $\mathcal{L}_E = \{0, 1\}$. Thus, for a node p a label l^d indexes a pair $(\mathbf{d}_p, e_p^d)$, where $e_p^d \in \mathcal{L}_E$. A node p is *selected* when $e_p^d = 1$, otherwise it is *deselected*. If p is deselected, it will not penalize inconsistent candidate segmentations and it will not be influenced by them. In the following section, only the relevant part of the pair $(\mathbf{d}_p, e_p^d)$ appears in the right hand side of the equations.

3.2 MRF Energy

The continuous energy in Eq. 1 is mapped to a discrete MRF energy of the form in Eq. 8. In short, we map i) the matching term M (Eq. 2) to the unary potentials of the deformation variables (Eq. 9), ii) the deformation smoothness penalty term R_d (Eq. 3) to pairwise potentials between deformation variables (Eq. 10), and iii) the coupling penalty C (Eq. 6) to one pairwise potential between registration and segmentation variables and one unary potential over deformation variables (first and second part of the right hand side of Eq. 12, correspondingly), as well as iv) the membership field smoothness penalty term R_f to an additional pairwise potential between deformation nodes (Eq. 13).

3.2.1 Multi-Atlas Registration

Multi-atlas registration is performed by registering in a pairwise fashion all atlases to the target image. Formulating pairwise registration in a discrete setting has been shown in [14]. For completeness reasons, we briefly discuss here how the matching term M and the regularization term R of Eq. 1 are mapped to unary and pairwise potentials.

As far as the matching term is concerned, we are interested in quantifying how well the assignment of a displacement label $l_{p_i}^d \in \mathcal{L}_D$ to a node $p_i \in \mathcal{V}_{D_i}$ aligns atlas A_i to the target image. This is measured by the following unary potential:

$$g_{p_i}^M(l_{p_i}^d) = \int_{\Omega} \hat{\omega}_{p_i}(x) \rho(A_i \circ D_{p_i}, I(x)) dx. \quad (9)$$

D_{p_i} is the transformation induced by the movement of the control point p in the i -th deformation grid by the displacement $l_{p_i}^d$. The weighting function $\hat{\omega}_{p_i}$ determines the contribution of the point x to the unary potential of the control point p . This function is similar to the ω weighting functions used in the FFD deformation model (Eq. 4).

Regarding the regularization term, [14] shows that it can be efficiently modeled by pairwise potentials. A discrete approximation of the gradient of the spatial transformation can be computed by taking the vector difference between the displacements of neighboring nodes that belong to the same deformation grid:

$$f_{p_i q_i}^R(l_{p_i}^d, l_{q_i}^d) = \|\mathbf{d}_{p_i} - \mathbf{d}_{q_i}\|, \quad (10)$$

where $\mathbf{d}_{p_i}$ is the displacement applied to node p in the i -th deformation grid, indexed by $l_{p_i}^d$.

3.2.2 Segmentation

In order to assign a class label to every voxel of the target image, we take into account learned appearance model for every class. The appearance model is encoded in the form of a probability distribution $\pi(l)$ and can be naturally incorporated in the MRF model by setting the unary potentials of the segmentation grid for every label to the negative log-probability of the respective class:

$$g_{q_s}^{SP}(l_{q_s}^s) = -\log(\pi(l_{q_s}^s)).$$

3.2.3 Integrated Segmentation and Multi-Atlas Registration

We want to encourage the agreement between the estimated segmentation and the warped segmentation masks. Thus, we penalize control point displacements of grid $\mathcal{G}_{D_i}$ that result in warping the segmentation mask of the corresponding atlas i in a fashion that does not agree with our final segmentation:

$$f_{p_i q_s}^C(l_{p_i}^d, l_{q_s}^s) = \hat{\omega}_{p_i}(s) \cdot \text{Ind}(S_i \circ D_{p_i}(s), l_{q_s}^s), \quad (11)$$

where p_i belongs to the grid $\mathcal{G}_{D_i}$, q_s belongs to $\mathcal{G}_S$ and

$$\text{Ind}(x, y) = \begin{cases} 0 & \text{when } x = y \\ 1 & \text{when } x \neq y. \end{cases}$$

Algorithm 1 Agreement percentages estimation algorithm.

```

1: Define  $a_{p_i} = \frac{\int_{\Omega} \hat{\omega}_{p_i}(x) (1 - \text{Ind}(\hat{S}(x), S_i(x))) dx}{\int_{\Omega} \hat{\omega}_{p_i}(x) dx}$ 
2: for each control point index  $p$  do
3:   Define  $AG_p = \cup_i(a_{p_i})$ 
4:   Sort and break  $AG_p$  into two sets:  $\hat{M}_i$  containing best half and  $\hat{M}_o$  containing the worst half.
5:   Define  $m = \frac{\sum \hat{M}_i}{2 * |\hat{M}_i|} + \frac{\sum \hat{M}_o}{2 * |\hat{M}_o|}$ 
6:   for N times do
7:     Initialize the inlier centroid  $\hat{c}_i \sim U(m, 1)$  and the outlier centroid  $\hat{c}_o \sim U(m, 0)$  where  $U(a, b)$  is the uniform distribution in  $[a, b]$ 
8:     Run 2-means. Get  $\tilde{M}_i, \tilde{M}_o, \tilde{s}$  as the clusters and score of the clustering.
9:     Keep the clusters with the minimum score in  $M_i, M_o$ .
10:  end for
11:  if  $|\text{var}(AG_p) - (\text{var}(M_i) + \text{var}(M_o))| > \epsilon$  then
12:    Set  $a_p = \frac{\min(M_i) + \max(M_o)}{2}$ 
13:  else
14:    Set  $a_p = 0$ 
15:  end if
16: end for

```

Local Atlas Selection Due to the sparse way we model selection variables, any candidate deformation $\mathbf{d}_p$ corresponds to a segmentation mask that agrees at a certain percentage to the segmentation variables within the support of p . In order to model local atlas selection, we need to determine a set of such percentages below which a candidate segmentation is unacceptably incongruent with the consensus segmentation, and thus should lead to a node being deselected. We call these percentages *agreement percentages* and index them with a_p . Note that the agreement percentage does not depend on the grid index i but only on the inter-grid control point index p . Given an agreement percentage a_p , atlas selection may be enforced by introducing an additional unary cost for deformation nodes. Thus the total coupling cost becomes

$$f_{p_i q_s}^{AS}(l_{p_i}^d, l_{q_s}^s) = e_{p_i}^d \cdot f_{p_i q_s}^C(l_{p_i}^d, l_{q_s}^s) + (1 - e_{p_i}^d) \cdot (1 - a_p), \quad (12)$$

where $e_{p_i}^d$ is equal to 1 when p_i is enabled, and 0 otherwise. As a consequence, nodes for which the (pairwise) segmentation cost f^C is very high when warping an atlas i , are disabled for this atlas. When a node p_i is disabled the cost to be paid is $1 - a_p$ regardless of the level of disagreement of the deformation with its corresponding segmentation nodes. Thus, a disabled deformation node will not affect segmentation variables and conversely will not be affected by them.

Agreement estimation Agreement percentages are estimated by locally comparing each segmentation mask

with an estimate of the consensus segmentation, followed by a clustering scheme to arrive at robust per control point estimates. This process is summarized in Algorithm 1. First, agreement percentages per control point are computed (step 1). Then a 2-means clustering algorithm is used to separate them into inliers and outliers (steps 3-10). Finally, a variance reduction criterion is used to determine the validity of the clustering into two sets and a_p is defined accordingly (steps 11-15).

Local Atlas Selection Smoothness Membership fields are encoded over deformation nodes. We want to enforce smoothness over the fields to achieve concise regions of influence for every atlas:

$$f_{p_i, q_i}^S(l_{p_i}^d, l_{q_i}^d) = \text{Ind}(e_{p_i}^d, e_{q_i}^d). \quad (13)$$

3.2.4 MRF energy parameterization

To conclude the outline of the discrete energy, we summarize the terms along with the parameters controlling each term's weight:

$$E_{MRF}(\mathbf{l}) = g_{p_i}^M(l_{p_i}^d) + \lambda f_{p_i, q_i}^R(l_{p_i}^d, l_{q_i}^d) + \sigma g_{q_s}^{SP}(l_{q_s}^s) + \alpha f_{p_i, q_s}^{AS}(l_{p_i}^d, l_{q_s}^s) + \beta f_{p_i, q_i}^S(l_{p_i}^d, l_{q_i}^d), \quad (14)$$

λ controls the deformation field smoothness, σ encodes the prior likelihood weight, α specifies how much registration influences segmentation and reversely how much segmentation affects registration, and β regulates the smoothness of membership fields.

3.3 MRF Optimization through Dual Decomposition

DD-MRF [19] has been introduced as a framework for MRF optimization, offering global optimality guarantees. Its flexibility in terms of possible energy types, its ability to report the quality of the final solution as well as its optimality guarantees are the merits we considered in opting for its use. DD-MRF works by receiving as input a decomposition of the initial graph (primal problem) into subgraphs (dual problems). It initializes the costs of the dual problems using the costs of the primal problem. It then proceeds by iteratively finding a global optimum for each subproblem, compare the subproblem solutions and update their costs.

In short, at each iteration: i) dual subproblems are solved in an optimal fashion (usually by dynamic programming); ii) the dual energy, defined as the sum of the energies of the optimal dual solutions, is computed; iii) a solution to the primal is inferred by the multitude of, possibly conflicting, dual solutions; iv) the difference between the energy of the primal and the sum

of the duals is computed (referred to as primal-dual gap); v) the primal dual gap is used to update subproblem costs. This way agreement is induced between the next subproblem solutions leading iteratively to a coherent, globally optimal solution. The way dual costs are updated guarantees that the euclidean distance of the current solution to the set of the globally optimum solutions will decrease monotonically.

Subproblem decomposition In order to optimize a graph $\mathcal{G}$, DD-MRF requires an input set of subproblems $\mathbf{SD} = \{SD_i, \dots, SD_n\}$, such that $\bigcup_i SD_i = \mathcal{G}$. Thus, we decompose the problem into a series of subproblems that can be exactly optimized through dynamic programming. Deformation grid subgraphs G_{D_i} are decomposed into sets C_{D_i} where every element of C_{D_i} corresponds to a chain subproblem.

The coupling term is divided into tree subproblems. There is one such subproblem per deformation node consisting of the deformation node p , the segmentation nodes within its support to which it is connected, as well as the edges connecting them. We call such a tree subproblem T_p and $T_{D_i} = \{T_p | p \in G_{D_i}\}$ denotes the tree subproblems that include all the nodes of a grid subgraph G_{D_i} . Thus, the subproblem decomposition for the problem at hand is given by the set $\mathbf{SD} = \{T_{D_0}, \dots, T_{D_N}, C_{D_0}, \dots, C_{D_N}\}$.

Local atlas selection model optimization However, a direct one-shot optimization of the above decomposition is not viable for a problem of this size and type. The inclusion of the switching potential that is encoded in the coupling cost makes optimization very hard. This has a crippling effect on both the time required for convergence and the quality of the resulting solution. In order to facilitate the solution of the problem, we further decompose the problem by breaking $\mathbf{SD}$ into N parts which we will optimize sequentially using DD-MRF, while updating the remaining parts with the solutions of the parts that have been already solved.

This iterative procedure estimates the deformation and selection fields for each atlas sequentially. For a given atlas i we project the coupling potentials of every other atlas by considering them stationary (zero vector displacement label for every deformation node). We then proceed by solving the single atlas decomposition over the updated segmentation potentials.

The algorithm consists of the steps shown in Alg. 2. The process starts by computing the potential functions for all candidate solutions. It then iterates over the atlases, projecting the updated constraints to the segmentation grid (step 3), optimizing the corresponding part of the decomposition to infer deformation variables

Fig. 1 Left: Closeness ranking of atlases to target IDs for leave-one-out cross-validation. Lighter boxes signify closeness of an atlas to a target. Right: Percentage of deselected nodes per atlas with respect to target ID as returned by our method when using all 17 atlases. The reported percentages are much more significant than what they appear, considering that 30%-40% of control points have a support that is essentially background. See Fig.5 for a visual representation. Elements in the diagonals should be ignored as naturally no atlas was used to segment itself.

and the best current segmentation labeling (step 4), updating deformation fields and selection images (step 5) and the corresponding constraints (step 6). The final segmentation is the one produced by the last iteration.

Algorithm 2 Sequential model optimization.

- 1: Define the potentials as described in previous sections.
 - 2: **for** each atlas k **do**
 - 3: Define $g_{q_S}^{Proj}(l_{q_S}^s) = \sum_{i \in N \setminus k} f_{p_i, q_S}^C(0, l_{q_S}^s)$ over $\mathcal{G}_S$
 - 4: Optimize using DD-MRF the subproblem set $\{T_{D_k}, C_{D_k}\}$
 - 5: Update the deformation field D_k and membership field image F_k from optimal labeling of $\mathcal{G}_{D_k}$.
 - 6: Recalculate $f_{p_k, q_S}^C(0, l_{q_S}^s)$
 - 7: **end for**
-

4 Validation

We have validated our method using leave-one-out cross-correlation over publicly data available on the Internet Brain Segmentation Repository (IBSR). We specifically use the skull stripped version of the dataset provided in [25]. The dataset consists of 18 T1-weighted MR Images with 1.5mm slice thickness. Images and masks have been linearly registered and cropped to $145 \times 158 \times 123$ from their initial resolution of $256 \times 256 \times 128$. We report mean results over labels that were annotated in more than half of the images of the dataset. Results are reported over 14 symmetric (left, right hemisphere) annotations (Thalamus, Caudate, Putamen, Pallidum,

Hippocampus, Amygdala, Accumbens, Lateral Ventricle, Inferior Lateral Ventricle, Cerebral White Matter, Cerebral Cortex, Cerebellum White Matter, Cerebellum Cortex, vessel) and 3 non-symmetric (3rd Ventricle, 4th Ventricle, Brain Stem).

Given a target image, all other images have been registered using affine registration and compared using Normalized Cross Correlation (NCC) to create a closeness ranking (see Fig. 1 (left)). Experiments using less than the full dataset ($N < 17$ atlases), refer to this ranking to select the N closest atlases to the target image. This ranking is also used by the sequential optimization detailed in Sec. 3.3 to determine the sequence of the atlases, from the closest to the furthest.

Registration parameters The same iterative procedure as [14] has been used in order to best cope with the computational efficiency vs accuracy trade-off. Dense deformation fields are produced by control point displacements using Cubic B-Spline functions. For all registrations we use two levels of deformation control points with spacing 7mm and 3.5mm. We use the deformation fields produced by the coarse level to initialize the finer one. For deformation nodes, candidate displacements are uniformly sampled over each axis, 12 per axis for a total of 37. For each control point level we iterate 5 times, reducing with each iteration the area of candidate displacements by a factor of 0.66. Normalized Cross Correlation (NCC) is the dissimilarity function (ρ) used for the matching criterion in all experiments.

Fig. 2 Top: Symmetric mean surface distance with respect to the ID of the target image. Bottom: Dice coefficient with respect to the ID of the target image. All 17 available atlases have been used to segment a target image. *Majority* refers to majority voting after pairwise registration, *Local* refers to the appearance-based locally weighted voting, *Soft Selection* refers to [2], while *Hard Selection* corresponds to the proposed method.

Segmentation Likelihoods Training has been conducted by removing the target image from the training set. The following configuration has been used to learn prior per voxel likelihoods. For each training set image, we sampled up to 150 samples belonging to each label and then sampled in a spatially uniform manner another 15000 samples avoiding duplicates. Three types of features were used: i) median, entropy, standard deviation, kurtosis and skewness sample statistics on a 3D patch with sides of 5, 7 and 9 voxels; ii) Gabor features using 6 per axis orientations and 3 scales; iii) HOG3D features on a $11 \times 11 \times 11$ patch broken up to 8 (2 per dimension) subpatches using 4 orientations; and iv) normalized voxel positions and distance from the center voxel. The total size of the feature vector was 258. We used the Random Forest framework [8] to discriminatively learn local per voxel probabilities for our target image.

We used 200 trees of maximum depth 20 and opted to stop splitting at 20 samples per node.

Optimization convergence Concerning DD-MRF convergence, we stop iteration when the mean change rate of the primal dual gap computed over the last 10 iterations reaches the 10% of the mean change rate of the primal dual gap computed over all iterations. This criterion leads to a uniform optimization quality over all possible target images, number of atlases and models.

Methods We compare our method against: i) pairwise registration fused using majority voting referred in the figures as *Majority*, ii) pairwise registration fused using appearance-based locally weighted fusion. We use an exponential weighting function over a local patch: $e^{-\rho NCC_s(x)}$ where ρ is a slope parameter, $NCC_s(x)$ the NCC comparison of an atlas and a target patch around

Method	Dice	SMSD (mm)
Hard S.	0.789 (0.018)	1.019 (0.114)
Soft S.	0.786 (0.019)	1.059 (0.105)
Major.	0.775 (0.018)	1.072 (0.138)
Local	0.779 (0.017)	1.052 (0.122)

Table 1 Mean Dice and SMSD using 17 atlases for all four methods.

point x and s the diameter of the patch. We refer to this method in the figures as *Local*, iii) [2], which is dubbed as *Soft Selection* because it couples registration and segmentation through the use of appearance-based local soft selection. The proposed method is termed as *Hard Selection* in the figures.

Parameter estimation Parameters were estimated by maximizing the Dice coefficient when segmenting 4 hand-picked target images by using 4 different datasets that consist of 7 atlases. First, the parameter λ was found for the pairwise registrations ($\lambda = 0.01$). For the appearance-based locally weighted fusion the slope parameter was found ($\rho = 3$) and the patch size was set ($s = 9$). Keeping λ the same, the parameters σ and α were set for [2] ($\sigma = 1.5, \alpha = 0.3$). Finally σ and α and β were set for the proposed method ($\sigma = 2.5, \alpha = 0.1, \beta = 0.025$). The ϵ parameter introduced in Alg. 1 was set to $\frac{3}{N^2}$ where N the number of atlases in the dataset.

4.1 Results

In this section, we validate the proposed method through qualitative and quantitative comparisons with competing methods. Moreover, we opt to shed light upon the mechanics of our method and in particular, how local atlas selection affects the segmentation framework. When we refer to a specific atlas id x out of the dataset we will use the notation A^x . In terms of quantitative results the complementarity of overlap and surface distance measures has been recently made evident in [25]. Overlap is measured by the Dice coefficient that is defined as $Dice(A, B) = \frac{2*|A \cap B|}{|A| + |B|}$ over annotations A and B . Surface distance is measured by the Symmetric Mean Surface Distance (SMSD) defined as follows: If $\hat{A}$ is the set of surface voxels of volume A and

$$d(A, B) = \frac{\sum_{i \in A} \min_{j \in B} e(i, j)}{|A|}$$

with $e(i, j)$ being the euclidean distance, then

$$SMSD(A, B) = \frac{d(\hat{A}, \hat{B}) + d(\hat{B}, \hat{A})}{2}.$$

In Fig. 2 we present comparison results of segmentations produced using all 17 available atlases. The proposed method is consistently better than competing methods although [2] seems to perform in some cases (e.g A^2) particularly well. As far as Dice coefficient is concerned, the results are statistically significant. The proposed method outperforms both majority voting (paired t-test, $P < 10^{-5}$) and local appearance-based voting (paired t-test, $P < 10^{-5}$). Moreover, it performs better than [2] (paired t-test, $P < 10^{-2}$) despite not exploiting appearance similarity information to weigh the coupling term. Regarding the SMSD metric, the proposed method outperforms majority voting (paired t-test, $P < 10^{-3}$), local appearance-based voting (paired t-test, $P < 10^{-2}$) and [2] (paired t-test, $P < 10^{-2}$). In Table 1 mean Dice overlap, SMSD are presented for the methods considered.

To qualitatively appraise the quality of the obtained results, we show in Fig. 4 sagittal views of the ground truth segmentation as well as the results obtained with all examined methods. We note that the proposed method has retrieved more details in the highlighted areas. In Fig. 5, we show the ground truth segmentation and the corresponding views of segmentation proposals by the most and least deselected atlases along with their respective membership fields. The two atlases appear to propose very different candidate segmentations. The segmentation that is proposed by the most deselected atlas, differs in many locations when compared to the ground truth. The proposed method correctly estimates most areas where erroneous proposals have been made and correctly disables the membership field. Moreover, correctly estimating the membership field assures that atlases with bad matching will not act as noise in the coupled registration/segmentation scheme.

A close look at Fig. 1 (right) adds additional insight into the atlas selection mechanism. The figure shows us how the algorithm prunes atlases in a concise manner. For example, we see how A^9 is treated in average as an outlier, while A^0 is consistently part of the set of atlases that drives the method. Additionally, we can see a clear clustering pattern. $A^0 - A^4$ and $A^{12} - A^{17}$ form a group in which atlases seem to segment one another, while the rest of the dataset seems to have a much more uniform node selection pattern.

Lastly, in Fig. 3 we address the issue of scaling with respect to the number of employed atlases. As far as the Dice coefficient is concerned (Fig. 3 (left)), [2] outperforms the proposed method when few atlases are used. This is due to the fact that the definition of the thresholds in Alg.1 is a noisy procedure.

As the number of used atlases increases, the mean quality of the dataset decreases. As consequence, our

Fig. 3 Left: Mean Dice coefficient with respect to the number of atlases used. Right: SMSD with respect to the number of atlases used. *Majority* refers to majority voting after pairwise registration, *Local* corresponds to appearance-based locally weighted voting, *Soft Selection* refers to [2], while *Hard Selection* corresponds to the proposed method.

Fig. 4 Sagittal views of the ground truth segmentation and the result produced by the proposed method, majority voting, local appearance-based voting and [2]. For the segmentation shown all 17 atlases have been used. Ellipses highlight areas where the proposed method has recovered more details than the competing ones.

Fig. 5 Ground truth segmentation and pairs of warped atlas segmentation mask and membership field for the atlas with the least (6% of the total number of control points) and most (11% of the total number of control points) deselected control points. For the segmentation shown all 17 atlases have been used. Note that the areas where segmentation errors occur (highlighted by red ellipses) have been correctly deselected.

method starts to outperform [2] because of its ability to extract good quality matchings from the newly added atlases, without allowing them to affect already established matchings. An analogous pattern can be seen in Fig. 3 (right) for the case of SMSD.

5 Discussion

In this paper, we presented a method that integrates registration and segmentation fusion in a pairwise MRF framework. The proposed approach allows registration parameters to be updated based on the segmentation estimation while membership fields estimation assures that bad matchings will not deteriorate segmentation

quality. The experimental results demonstrate the potential of the proposed approach.

6 Acknowledgements

This research was partially supported by European Research Council Grant Diocles (ERC-STG-259112).

References

1. Akhondi-Asl, A., Warfield, S.K.: Simultaneous Truth and Performance Level Estimation Through Fusion of Probabilistic Segmentations. *Transactions on Medical Imaging* 32, 1840+ (Oct 2013)

2. Alchatzidis, S., Sotiras, A., Paragios, N.: Discrete Multi Atlas Segmentation using Agreement Constraints. In: British Machine Vision Conference (Sep 2014), <http://hal.inria.fr/hal-01061457>
3. Artaechevarria, X., Munoz-Barrutia, A., Ortizde Solorzano, C.: Combination Strategies in Multi-Atlas Image Segmentation: Application to Brain MR Data. *Medical Imaging, IEEE Transactions on* 28(8), 1266–1277 (Aug 2009), <http://dx.doi.org/10.1109/tmi.2009.2014372>
4. Ashburner, J., Friston, K.J.: Unified segmentation. *NeuroImage* 26(3), 839–851 (Jul 2005), <http://dx.doi.org/10.1016/j.neuroimage.2005.02.018>
5. Asman, A.J., Landman, B.A.: Robust Statistical Label Fusion Through Consensus Level, Labeler Accuracy, and Truth Estimation (COLLATE). *Medical Imaging, IEEE Transactions on* 30(10), 1779–1794 (Oct 2011), <http://dx.doi.org/10.1109/tmi.2011.2147795>
6. Asman, A.J., Landman, B.A.: Non-local statistical label fusion for multi-atlas segmentation. *Medical Image Analysis* 17(2), 194–208 (Feb 2013), <http://dx.doi.org/10.1016/j.media.2012.10.002>
7. Asman, A., Smith, S., Reich, D., Landman, B.: Robust GM/WM Segmentation of the Spinal Cord with Iterative Non-local Statistical Fusion. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *Medical Image Computing and Computer-Assisted Intervention MICCAI 2013, Lecture Notes in Computer Science*, vol. 8149, pp. 759–767. Springer Berlin Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-40811-3_95
8. Breiman, L.: Random Forests. *Mach. Learn.* 45(1), 5–32 (Oct 2001), <http://dx.doi.org/10.1023/a:1010933404324>
9. Coupé, P., Manjón, J.V., Fonov, V., Pruessner, J., Robles, M., Collins, D.L.: Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage* 54(2), 940–954 (Jan 2011), <http://dx.doi.org/10.1016/j.neuroimage.2010.09.018>
10. Doshi, J., Erus, G., Ou, Y., Gaonkar, B., Davatzikos, C.: Multi-atlas skull-stripping. *Academic radiology* 20(12), 1566–1576 (2013)
11. Dowling, J., Frapp, J., Chandra, S., Pluim, J., Lambert, J., Parker, J., Denham, J., Greer, P., Salvado, O.: Fast Automatic Multi-atlas Segmentation of the Prostate from 3D MR Images. In: Madabhushi, A., Dowling, J., Huisman, H., Barratt, D. (eds.) *Prostate Cancer Imaging. Image Analysis and Image-Guided Interventions, Lecture Notes in Computer Science*, vol. 6963, pp. 10–21. Springer Berlin Heidelberg (2011), http://dx.doi.org/10.1007/978-3-642-23944-1_2
12. Eugenio Iglesias, J., Rory Sabuncu, M., Van Leemput, K.: A unified framework for cross-modality multi-atlas segmentation of brain MRI. *Medical Image Analysis* 17(8), 1181–1191 (Dec 2013), <http://dx.doi.org/10.1016/j.media.2013.08.001>
13. Fonov, V., Coupé, P., Eskildsen, S., Manjon, J., Collins, L.: Multi-atlas labeling with population-specific template and non-local patch-based label fusion. In: *MICCAI 2012 Workshop on Multi-Atlas Labeling*, pp. 63–66 (Oct 2012), <https://hal.inria.fr/hal-00739285/>
14. Glocker, B., Sotiras, A., Komodakis, N., Paragios, N.: Deformable Medical Image Registration: Setting the State of the Art with Discrete Methods*. *Annual Review of Biomedical Engineering* 13(1), 219–244 (2011)
15. Heckemann, R.A., Hajnal, J.V., Aljabar, P., Rueckert, D., Hammers, A.: Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33(1), 115–126 (Oct 2006), <http://dx.doi.org/10.1016/j.neuroimage.2006.05.061>
16. Isgum, I., Staring, M., Rutten, A., Prokop, M., Viergever, M.A., van Ginneken, B.: Multi-Atlas-Based Segmentation With Local Decision Fusion & Application to Cardiac and Aortic Segmentation in CT Scans. *Medical Imaging, IEEE Transactions on* 28(7), 1000–1010 (Jul 2009), <http://dx.doi.org/10.1109/tmi.2008.2011480>
17. Jorge Cardoso, M., Leung, K., Modat, M., Keihaninejad, S., Cash, D., Barnes, J., Fox, N.C., Ourselin, S.: STEPS: Similarity and Truth Estimation for Propagated Segmentations and its application to hippocampal segmentation and brain parcellation. *Medical Image Analysis* 17(6), 671–684 (Aug 2013), <http://dx.doi.org/10.1016/j.media.2013.02.006>
18. Klein, S., van der Heide, U.A., Lips, I.M., van Vulpen, M., Staring, M., Pluim, J.P.W.: Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information. *Medical Physics* 35(4), 1407–1417 (Apr 2008), <http://dx.doi.org/10.1118/1.2842076>
19. Komodakis, N., Paragios, N., Tziritas, G.: MRF energy minimization and beyond via dual decomposition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33(3), 531–552 (Mar 2011), <http://dx.doi.org/10.1109/tpami.2010.108>
20. Kybic, J., Unser, M.: Fast parametric elastic image registration. *Image Processing, IEEE Transactions on* 12(11), 1427–1442 (2003)
21. Langerak, T.R., van der Heide, U.A., Kotte, A.N.T.J., Viergever, M.A., van Vulpen, M., Pluim, J.P.W.: Label Fusion in Atlas-Based Segmentation Using a Selective and Iterative Method for Performance Level Estimation (SIMPLE). *Medical Imaging, IEEE Transactions on* 29(12), 2000–2008 (Dec 2010), <http://dx.doi.org/10.1109/tmi.2010.2057442>
22. Ou, Y., Shen, D., Feldman, M., Tomaszewski, J., Davatzikos, C.: Non-rigid registration between histological and MR images of the prostate: A joint segmentation and registration framework. In: *Computer Vision and Pattern Recognition Workshops, 2009. CVPR Workshops 2009. IEEE Computer Society Conference on*, pp. 125–132. IEEE (Jun 2009), <http://dx.doi.org/10.1109/cvpr.2009.5204347>
23. Parisot, S., Duffau, H., Chemouny, S., Paragios, N.: Joint Tumor Segmentation and Dense Deformable Registration of Brain MR Images. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *Medical Image Computing and Computer-Assisted Intervention MICCAI 2012, Lecture Notes in Computer Science*, vol. 7511, pp. 651–658. Springer Berlin Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-33418-4_80
24. van Rikxoort, E.M., Isgum, I., Arzhaeva, Y., Staring, M., Klein, S., Viergever, M.A., Pluim, J.P.W., van Ginneken, B.: Adaptive local multi-atlas segmentation: Application to the heart and the caudate nucleus. *Medical Image Analysis* 14(1), 39–49 (Feb 2010), <http://dx.doi.org/10.1016/j.media.2009.10.001>
25. Rohlfing, T.: Image similarity and tissue overlaps as surrogates for image registration accuracy: widely used but unreliable. *IEEE transactions on medical imaging* 31(2), 153–163 (Feb 2012), <http://dx.doi.org/10.1109/tmi.2011.2163944>
26. Rohlfing, T., Brandt, R., Menzel, R., Russakoff, D., Maurer, C.: Quo Vadis, Atlas-Based Segmentation? In: Suri, J., Wilson, D., Laxminarayan,

- S. (eds.) Handbook of Biomedical Image Analysis, pp. 435–486. Topics in Biomedical Engineering International Book Series, Springer US (2005), http://dx.doi.org/10.1007/0-306-48608-3_11
27. Rohlfing, T., Russakoff, D., Maurer, C.: Expectation Maximization Strategies for Multi-atlas Multi-label Segmentation. In: Taylor, C., Noble (eds.) Information Processing in Medical Imaging, Lecture Notes in Computer Science, vol. 2732, pp. 210–221. Springer Berlin Heidelberg (2003), http://dx.doi.org/10.1007/978-3-540-45087-0_18
28. Rousseau, F., Habas, P.A., Studholme, C.: A Supervised Patch-Based Approach for Human Brain Labeling. Medical Imaging, IEEE Transactions on 30(10), 1852–1862 (Oct 2011), <http://dx.doi.org/10.1109/tmi.2011.2156806>
29. 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. In: IEEE Transactions on Medical Imaging. pp. 712–721 (1999)
30. Sabuncu, M.R., Yeo, B.T.T., Van Leemput, K., Fischl, B., Golland, P.: A Generative Model for Image Segmentation Based on Label Fusion. Medical Imaging, IEEE Transactions on 29(10), 1714–1729 (Oct 2010), <http://dx.doi.org/10.1109/tmi.2010.2050897>
31. Sdika, M.: Combining atlas based segmentation and intensity classification with nearest neighbor transform and accuracy weighted vote. Medical Image Analysis 14(2), 219–226 (2010)
32. Seghier, M.L., Ramackhansingh, A., Crinion, J., Leff, A.P., Price, C.J.: Lesion identification using unified segmentation-normalisation models and fuzzy clustering. NeuroImage 41(4), 1253–1266 (Jul 2008), <http://dx.doi.org/10.1016/j.neuroimage.2008.03.028>
33. Tang, X., Oishi, K., Faria, A.V., Hillis, A.E., Albert, M.S., Mori, S., Miller, M.I.: Bayesian parameter estimation and segmentation in the multi-atlas random orbit model (2013)
34. Wang, H., Suh, J.W., Das, S.R., Pluta, J.B., Craige, C., Yushkevich, P.A.: Multi-Atlas Segmentation with Joint Label Fusion. Pattern Analysis and Machine Intelligence, IEEE Transactions on 35(3), 611–623 (Mar 2013), <http://dx.doi.org/10.1109/tpami.2012.143>
35. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. IEEE transactions on medical imaging 23(7), 903–921 (Jul 2004), <http://dx.doi.org/10.1109/tmi.2004.828354>
36. Warfield, S., Zou, K., Wells, W.: Validation of Image Segmentation and Expert Quality with an Expectation-Maximization Algorithm. In: Dohi, T., Kikinis, R. (eds.) Medical Image Computing and Computer-Assisted Intervention MICCAI 2002, Lecture Notes in Computer Science, vol. 2488, pp. 298–306. Springer Berlin Heidelberg (2002), http://dx.doi.org/10.1007/3-540-45786-0_37
37. Wyatt, P.P., Noble: MAP MRF joint segmentation and registration of medical images. Medical Image Analysis 7(4), 539–552 (Dec 2003), [http://dx.doi.org/10.1016/s1361-8415\(03\)00067-7](http://dx.doi.org/10.1016/s1361-8415(03)00067-7)
38. Xiaohua, C., Brady, M., Lo, J.C., Moore, N.: Simultaneous Segmentation and Registration of Contrast-Enhanced Breast MRI. In: Christensen, G., Sonka, M. (eds.) Information Processing in Medical Imaging, Lecture Notes in Computer Science, vol. 3565, pp. 126–137. Springer Berlin Heidelberg (2005), http://dx.doi.org/10.1007/11505730_11
39. Xiaohua, C., Brady, M., Rueckert, D.: Simultaneous Segmentation and Registration for Medical Image. In: Barillot, C., Haynor, D., Hellier, P. (eds.) Medical Image Computing and Computer-Assisted Intervention MICCAI 2004, Lecture Notes in Computer Science, vol. 3216, pp. 663–670. Springer Berlin Heidelberg (2004), http://dx.doi.org/10.1007/978-3-540-30135-6_81
40. Xue, Z., Wong, K., Wong, S.T.C.: Joint registration and segmentation of serial lung CT images for image-guided lung cancer diagnosis and therapy. Computerized Medical Imaging and Graphics 34(1), 55–60 (Jan 2010), <http://dx.doi.org/10.1016/j.compmedimag.2009.05.007>