



Local atlas selection for discrete multi-atlas segmentation

Stavros Alchatzidis, Aristeidis Sotiras, Nikos Paragios

► To cite this version:

Stavros Alchatzidis, Aristeidis Sotiras, Nikos Paragios. Local atlas selection for discrete multi-atlas segmentation. International Symposium on Biomedical Imaging, Apr 2015, New York, United States. 10.1109/isbi.2015.7163888 . hal-01145510

HAL Id: hal-01145510

<https://hal.science/hal-01145510>

Submitted on 27 Apr 2015

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.

LOCAL ATLAS SELECTION FOR DISCRETE MULTI-ATLAS SEGMENTATION

Stavros Alchatzidis¹³

Aristeidis Sotiras²

Nikos Paragios¹³

¹ Equipe GALEN, INRIA Saclay, Île-de-France, Orsay, France

² Section of Biomedical Image Analysis, University of Pennsylvania, Pennsylvania, USA

³ Ecole Centrale de Paris, Châtenay-Malabry, Île-de-France, France

ABSTRACT

Multi-atlas segmentation is commonly performed in two separate steps: i) multiple pairwise registrations, and ii) fusion of the deformed segmentation masks towards labeling objects of interest. In this paper we propose an approach for integrated volume segmentation through multi-atlas registration. To tackle this problem, we opt for a graphical model where registration and segmentation nodes are coupled. The aim is to recover simultaneously all atlas deformations along with selection masks quantifying the participation of each atlas per segmentation voxel. The above is modeled using a pairwise graphical model where deformation and segmentation variables are modeled explicitly. A sequential optimization relaxation is proposed for efficient inference. Promising performance is reported on the IBSR dataset when comparing to majority voting and local appearance-based weighted voting.

Index Terms— Multi-atlas, segmentation, medical imaging, discrete optimization, Markov Random Fields

1. INTRODUCTION

Multi-atlas segmentation methods have grown in popularity in recent years. Such methods use a database of n atlases (in this context *atlas* refers to the aggregation of an intensity and an annotated image) to segment a target image. Segmentation is addressed through deformable registration, where the transformation used to register each atlas to the target volume is used to deform the corresponding segmentation map, producing n -deformation maps and eventually n -segmentation maps. These maps are subsequently fused into one final segmentation by using either i) majority voting, ii) local per-voxel appearance based weighted voting [1], or iii) an EM approach [2] to estimate the confusion matrix for each atlas which is used to weigh votes.

The quality of the final segmentation depends greatly on the quality of the available hypotheses that is, in turn, influenced by registration errors. The correlated nature of the errors produced by different atlases is modeled in [3], while [4]

proposed a generative model for label fusion that allows for the estimation of a spatially varying membership function that models the target image as a combination of one or more atlases. Aiming to further reduce the influence of registration errors in the final segmentation, methods that couple registration and segmentation has been proposed [5, 6].

In this paper we propose a unified formulation for multi-atlas segmentation. Our work derives from a model with the same structure as [6]. The major differences between the two works are: i) [6] uses soft local atlas selection based on appearance, while our method explicitly models variables for selecting parts of the image in accordance with the global consensus. Atlas support masks are directly produced from optimizing the model. ii) Local atlas selection in [6] is based on a similarity criterion similar to the data cost function. A direct consequence of this is the need for a common similarity metric to be applicable on all possible atlas-target pairs. Our method achieves selection by comparing candidate segmentations. This enables the use of images of different modalities to be part of the dataset. Let's note here that MRF-based registration can incorporate all possible types of similarity functions, which in conjunction to our method, allows for choosing an optimal similarity function per atlas-target pair.

2. MATERIALS AND METHODS

We consider N images forming the set of available atlases $\mathbf{A} = \{A_0, \dots, A_{N-1}\}$. Areas of interest corresponding to each image are annotated in segmentation masks forming the set $\mathbf{S} = \{S_0, \dots, S_{N-1}\}$. Each voxel in the segmentation masks is assigned to a label corresponding to one of M areas of interest, $S_i(x) \in \{0, \dots, M-1\}$.

Sets $\mathbf{A}$ and $\mathbf{S}$ along with the query image I are given as input. Output for our problem are : i) The segmentation mask of the query image, S_I , ii) a set of deformation fields $\mathbf{D} = \{D_0, \dots, D_{N-1}\}$, where D_i denotes the deformation field mapping A_i to I iii) a set of atlas selection maps, $\mathbf{B} = \{B_0, \dots, B_{N-1}\}$, $B_i(x) \in [0, 1]$ quantifying how much an atlas affects the final segmentation over the image domain.

Research partially supported by the European Research Council Starting Grant Diocles (ERC-STG-259112).

2.1. Transformation model

In this work, the popular Free Form Deformations (FFDs) transformation model [7, 8] is used. Free Form Deformations parametrize the transformation $D(x)$ by a linear combination of K control points, $D(x) = x + \sum_{i=0}^{K-1} \omega_i(x) \phi_i$, where ϕ_i is the displacement of control point i and $\omega_i(x)$ is the weighting function corresponding to point i . We use N uniformly distributed grids of control points over the image domain (one corresponding to each atlas) and cubic B-splines as the weighting functions.

2.2. MRF Formulation: Graph Structure

We use Markov Random Field to formulate our 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 constructed graph should encode the registration, the segmentation, the coupling and the atlas selection components of the proposed model.

Registration 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 superimposed onto image A_i , there is a node $p_i \in \mathcal{V}_{D_i}$ that represents its displacement as a random variable. Since grids are isomorphic, p indexes a common to all grids control point position, while i indexes the grid. The edge system of each grid $\mathcal{E}_{D_i}$ is created by a regular connectivity scheme. Edges encode dependencies between random variables. The solution space around a control point is quantized and indexed by a the set of labels $\mathcal{L}_D$. This represents possible control point displacements $\mathbf{d}$. We refer to a potential displacement attributed to a node indexed by a label as $\mathbf{d}^{l^d}$.

Segmentation An additional grid $\mathcal{G}_S = (\mathcal{V}_S)$ is employed to represent segmentation. Each node $p \in \mathcal{V}_S$ corresponds to a random variable. 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 attributed to a segmentation node by l^s . In our experiments, every node $p \in \mathcal{V}_S$ corresponds to a voxel in the target image, so that s corresponds to a voxel position.

Coupling Integrating segmentation and multi-atlas registration is achieved by coupling segmentation and deformation grids. 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 in $p \in \mathcal{V}_D$ with nodes of $\mathcal{V}_S$ that correspond to voxels who belong to the image support of p .

Atlas Selection We encode atlas selection at deformation nodes. A consequence of using such a sparse field is that a candidate deformation, $\mathbf{d}^{l^d}$, corresponds to a segmentation mask that agrees (or disagrees) at a certain percentage with a

candidate segmentation mask in the support of a deformation node. To encode this, we consider a control point as active if the proposed segmentation over the control point's image support agrees with the estimated segmentation by at least a certain threshold. Otherwise, it is considered as disabled. If a control point is disabled it will not penalize inconsistent candidate segmentations. To model this, we consider the Cartesian product between the deformation label set $\mathcal{L}_D$ and a binary selection label set $\mathbf{e} = \{0, 1\}$ as the label set of deformation nodes. Thus, a label l^d indexes a pair $(\mathbf{d}^{l^d}, \mathbf{e}^{l^d})$. In the following section, only the relevant part of this pair will appear in the right hand side of the equations.

2.3. MRF Formulation: Energy terms

The above graph is associated with an energy of the form:

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

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

Registration For the matching term, we are interested in quantifying how well the assignment of a displacement label aligns atlas A_i to the target image:

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

$D_i^{l_{p_i}^d}$ is the transformation induced by the movement of the control point p in the i th deformation grid by the displacement $\mathbf{d}^{l_{p_i}^d}$. The support function $\hat{\omega}_{p_i}$ determines the contribution of the point x to the unary potential of the control point p .

Regarding the regularization term 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:

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

Segmentation We want to penalize candidate segmentations that disagree with a learned appearance model $p_x(l)$.

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

Coupling We want to encourage the agreement between the estimated segmentation and the warped segmentation mask.

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

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

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

```

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 worse 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  $|var(AG_p) - (var(M_i) + var(M_o))| > \epsilon$  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

```

Algorithm 1: Agreement percentages estimation algorithm.

Atlas Selection Smoothness A simple Potts model is used to impose smoothness over the selection field:

$$f_{p_i, q_i}^{AS}(l_{p_i}^d, l_{q_i}^d) = \text{Ind}(l_{p_i}^d, l_{q_i}^d).$$

Atlas Selection To enforce atlas selection we use an agreement percentage a_p to define an additional unary cost for deformation nodes. Note that the agreement percentage does not depend on the grid index i but only on the inter-grid control point index p :

$$g_{p_i}^{AS}(l_{p_i}^d) = (1 - a_p) \cdot (1 - \mathbf{e}^{l_{p_i}^d}),$$

where $\mathbf{e}^{l_{p_i}^d}$ is equal to 1 when p_i is enabled, and 0 otherwise. In addition, a coupling term is added:

$$f_{p_i, q_s}^{AS}(l_{p_i}^d, l_{q_s}^s) = -f_{p_i, q_s}^C(l_{p_i}^d, l_{q_s}^s) \cdot (1 - \mathbf{e}^{l_{p_i}^d})$$

so that if a node is disabled the sum of $f^C + f^{AS}$ will be zero. In other words, if a node p_i is disabled the cost to be paid will be $1 - a_p$ regardless of the level of disagreement of the deformation with its corresponding segmentation nodes. Note that a disabled deformation node will not affect the segmentation variables and conversely will not be affected by them.

Agreement estimation We seek a set of agreement thresholds that can quantify the level of local agreement. If the local proposal of an atlas does not agree by a certain threshold with the estimated segmentation, we consider the registration as having locally failed and thus, having nothing more to contribute. Thus, the extra constraint imposed on registration variables acts as noise and should be ignored, resulting in decoupling registration and segmentation variables.

Let us detail here the way the agreement percentages a_p are estimated. We use p_i to refer to control point p in deformation grid i like we did with random variables in the previous subsection. The algorithm used is provided in Alg.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 the threshold is defined accordingly (steps 11-15).

```

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 the model consisting of  $\mathcal{G}_{D_k}^C, \mathcal{G}_S$  and the edges of  $\mathcal{E}_C$  coupling the two.
5:   Update the deformation field  $D_k$  and selection image  $B_k$  from optimal labeling of  $\mathcal{G}_{D_k}$ .
6:   Recalculate  $f_{p_k, q_s}^C(0, l_{q_s}^s)$ 
7: end for

```

Algorithm 2: Sequential model optimization.

2.4. MRF optimization

We introduce an iterative procedure for optimization that estimates the deformation and selection fields for each atlas sequentially. 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), inferring deformation variables 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.

3. EXPERIMENTAL VALIDATION

In order to validate the proposed method, we used the dataset provided by the Internet Brain Segmentation Repository. The dataset includes 38 different regions annotated over all atlases with each atlas being annotated using 34 to 36 regions. The dataset has been pre-processed as detailed in [6] (linear registration and cropping). Per-voxel likelihoods have been computed in the same fashion as in [6]. Leave-one-out cross-validation was performed on a set of 18 atlases.

For all registrations, we use a multi-resolution scheme that employs a deformation grid of two different resolutions. The spacing was set to 7mm for the coarse resolution, and was halved for the fine one. A total number of 37 labels, 12 regularly sampled along each principal axis, were used. We iterate 4 times at each level, reducing with each iteration the radius of candidate displacements by a factor of 0.66. Normalized Cross Correlation (NCC) is the matching criterion that was used for all experiments.

We compare the proposed method against [6], as well as standard (decoupled) multi-atlas segmentation. For the latter, registration is performed using the same parameters as for the




Fig. 1: Sagittal views of the ground truth segmentation and the result produced by the proposed method, majority voting, local appearance-based voting and [6]. Ellipses highlight areas where the proposed method has recovered more details than the competing ones.




Fig. 2: Ground truth segmentation and pairs of warped atlas segmentation mask and selection field for the atlas with the least (5281 for the whole image) and most (9731 for the whole image) deselected control points. Note that the areas where segmentation errors occur (highlighted by red ellipses) have been correctly deselected.

Method	Dice	HD (mm)	SMSD (mm)	HE
Coupled	0.781 (0.019)	6.69 (1.19)	1.23 (0.69)	0.492 (0.020)
[6]	0.778 (0.019)	7.65 (1.16)	1.26 (0.65)	0.467 (0.019)
Decoupled Majority	0.765 (0.023)	6.27 (1.26)	1.28 (0.69)	0.6 (0.024)
Decoupled Local NCC	0.768 (0.023)	6.27 (1.24)	1.26 (0.69)	

Table 1: Quantitative results. The mean Dice overlap coefficient, HD, SMSD and Harmonic Energy (HE) are reported for all methods. Note that the proposed method achieves improved segmentation agreement, while producing smoother deformations.

proposed method, while fusion is performed by both majority and local appearance-based voting.

To qualitatively appraise the quality of the obtained results, we show in Fig.1 sagittal views of the ground truth segmentation as well as the results obtained with each of the methods we examined. We note that the proposed method has better retrieved details in the highlighted areas. In Fig.2 we show the ground truth segmentation and the corresponding views of segmentations proposals by the most and least deselected atlases along with their respective selection fields. We note that in the areas erroneous segmentation labels are proposed, the proposals have been correctly deselected.

To quantitatively evaluate the results, we estimate the mean Dice overlap coefficient, Hausdorff Distance (HD) and Symmetric Mean Surface Distance (SMSD) between the obtained segmentations and the ground truth one, over leave-one-out cross validation experiments (see Table 1). Similarly, we report the mean harmonic energy of the respectively estimated deformation fields. As far as the Dice overlap coefficient is concerned, the proposed method outperforms

both majority voting (paired t-test, $P = 0.000001$) and local appearance-based voting (paired t-test, $P = 0.00001$). Moreover, it performs better than [6] (paired t-test, $P = 0.011$) despite not exploiting appearance similarity information to weigh the coupling term. The proposed method also achieves smaller SMSD than all competing approaches. Note that the proposed method and [6] produce more accurate segmentations than standard decoupled multi-atlas segmentation variants, while producing smoother registration results (as measured by the Harmonic Energy).

4. DISCUSSION

In this paper, we proposed a novel method for multi-atlas segmentation that couples registration and segmentation variables. Experimental results demonstrate that segmentation and registration mutually benefit in this coupled setting, leading to more accurate segmentations and smoother registrations. To conclude, coupled approaches constitute a promising research direction.

5. REFERENCES

- [1] X. Artaechevarria, A. Munoz-Barrutia, and C. Ortiz-de Solorzano, "Combination Strategies in Multi-Atlas Image Segmentation: Application to Brain MR Data," *Medical Imaging, IEEE Transactions on*, vol. 28, no. 8, pp. 1266–1277, Aug. 2009.
- [2] Torsten Rohlfing, Daniel B Russakoff, and Calvin R Maurer, "Expectation Maximization Strategies for Multi-atlas Multi-label Segmentation," in *Information Processing in Medical Imaging*, vol. 2732 of *Lecture Notes in Computer Science*, pp. 210–221. Springer Berlin Heidelberg, 2003.
- [3] Hongzhi Wang, Jung Wook Suh, Sandhitsu R Das, John B Pluta, Caryne Craige, and Paul A Yushkevich, "Multi-atlas segmentation with joint label fusion," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 3, pp. 611–623, 2013.
- [4] M. R. Sabuncu, B. T. T. Yeo, K. Van Leemput, B. Fischl, and P. Golland, "A Generative Model for Image Segmentation Based on Label Fusion," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 10, pp. 1714–1729, Oct. 2010.
- [5] Juan Eugenio Iglesias, Mert Rory Sabuncu, and Koen Van Leemput, "A unified framework for cross-modality multi-atlas segmentation of brain MRI," *Medical Image Analysis*, vol. 17, no. 8, pp. 1181–1191, Dec. 2013.
- [6] Stavros Alchatzidis, Aristeidis Sotiras, and Nikos Paragios, "Discrete Multi Atlas Segmentation using Agreement Constraints," in *British Machine Vision Conference*, Sept. 2014.
- [7] Daniel Rueckert, Luke I Sonoda, Carmes Hayes, Derek L. G. Hill, Martin O. Leach, and David J. Hawkes, "Non-rigid registration using free-form deformations: application to breast mr images," *Medical Imaging, IEEE Transactions on*, vol. 18, no. 8, pp. 712–721, 1999.
- [8] Jan Kybic and Michael Unser, "Fast parametric elastic image registration," *Image Processing, IEEE Transactions on*, vol. 12, no. 11, pp. 1427–1442, 2003.