



HAL
open science

Discrete Multi Atlas Segmentation using Agreement Constraints

Stavros Alchatzidis, Aristeidis Sotiras, Nikos Paragios

► **To cite this version:**

Stavros Alchatzidis, Aristeidis Sotiras, Nikos Paragios. Discrete Multi Atlas Segmentation using Agreement Constraints. British Machine Vision Conference, British Machine Vision Association, Sep 2014, Nottingham, United Kingdom. hal-01061457

HAL Id: hal-01061457

<https://inria.hal.science/hal-01061457>

Submitted on 7 Sep 2014

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.

Discrete Multi Atlas Segmentation using Agreement Constraints

Stavros Alchatzidis¹³⁴
stavros.alchatzidis@ecp.fr

Aristeidis Sotiras²
aristeidis.sotiras@uphs.upenn.edu

Nikos Paragios¹³⁴
nikos.paragios@ecp.fr

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

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

³ Ecole des Ponts Paristech, Champs-sur-Marne, Île-de-France, France

⁴ Ecole Centrale de Paris, Châtenay-Malabry, Île-de-France, France

Abstract

In this paper we propose an approach for integrated volume segmentation and multi-atlas registration. The aim is to recover simultaneously all atlas deformations along with a segmentation mask that imposes consistency between their deformed segmentation maps. This is modeled through a pairwise graphical model where deformation variables are coupled with segmentation ones. Optimization via dual decomposition is used to recover the optimal solution of the method. Promising results on the IBSR dataset demonstrate the potential of our method where improvement is obtained when comparing to the post-registration label fusion.

1 Introduction

Segmentation is a well studied problem in medical imaging and computer vision. Despite the maturity of the field, segmenting multiple anatomical structures remains a challenging problem. Atlas-based segmentation [11] that addresses this problem through the registration of an atlas, and the associated annotation, to the query image has gained popularity in recent years due to the advance of registration methods. Nonetheless, this approach is limited with respect to the spectrum of variations that can be captured by the atlas.

Aiming to address this limitation, approaches that employ multiple atlases have been proposed. In this scenario, multiple registrations are performed resulting in multiple segmentation hypotheses. The final segmentation is produced by subsequently fusing the hypotheses either in a local [2, 6] or a global [1] fashion. Fusion strategies vary from simple majority voting to more sophisticated probabilistic frameworks [10, 13]. In general, the segmentation problem is solved in two discrete steps and registration is merely seen as a fixed preprocessing step. In this setting, the registration process does not profit from the segmentation fusion since it is not allowed to be updated towards refining correspondences and segmentation hypotheses.

In this paper, we aim to couple registration and segmentation problem through a unified formulation for multi-atlas segmentation. Registration terms seek optimal visual correspondences between atlases and target volumes while imposing smoothness. Segmentation terms seek voxel-wise consensus on the labeling of the target with respect to the deformed segmentation maps. Prior per voxel probabilities, produced by learning of local features, are taken into account in a seamless manner. In order to mathematically formulate these components, we adopt a pairwise Markov Random Field (MRF) graphical model where each atlas is associated with a deformation field, while the target image is associated with a segmentation map. Let us note that a similar approach was recently proposed by [4]. This approach is different from the proposed one in two regards: i) it employs an Expectation-Maximization framework for parameter optimization, and ii) it does not take into account discriminatively learned voxel likelihoods.

The remainder of this paper is organized as follows: Sec. 2 presents the theoretical model of our approach, while Sec. 3 discusses its discrete variant. Experimental results are part of Sec. 4, and discussion concludes the paper.

2 Problem Formulation

Without loss of generality, let us consider N annotated images which form the set $\mathbf{A} = \{A_0, \dots, A_{N-1}\}$. 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, $S_i(x) \in \{0, \dots, M-1\}$. As input, an image I is given to be segmented into anatomical regions. Henceforth, we are going to interchangeably refer to this image as either target or query image. The segmentation mask, S_I along with 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 are the outputs to our problem.

2.1 Method Outline

In the standard multi-atlas registration framework, the set of deformation fields is commonly estimated during an initial step by maximizing an intensity-based image similarity criterion, while penalizing non-smooth configurations. The proposed method simultaneously optimizes for the parameters of the final segmentation S_I as well as the set of deformation fields \mathbf{D} . The estimation of the deformation fields is also guided by the fusion result. Our basic premise is that *the combined strategy that allows the two problems to interact will lead to increased segmentation accuracy by introducing agreement between the segmentations proposed by each atlas*. Last, we also propose to take into account class specific appearance priors when optimizing for the final segmentation S_I . The idea behind incorporating prior segmentation probabilities is motivated by the fact that image registration is often trapped in local minima. This is particularly true in areas of high anatomical variability like brain cortex where appearance information is a more reliable cue and can robustly guide segmentation [14].

2.2 Continuous Energy Formulation

We formulate our problem as an energy minimization one. The proposed energy consists of three components: 1) 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) that enforces the smoothness of the deformation field; 2) a segmentation component comprising an appearance prior term (S_P), that measures the log-likelihood of segmentation with respect to the probabilities learned during a training phase. 3) a coupling term (C) that takes into account the labeling that is proposed by the atlases and encourages their agreement with the estimated segmentation S_I . The energy has the following form:

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

2.2.1 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 Given a smoothness inducing function ψ , the regularization term takes the following form:

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

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 [8, 12] 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 the weighting function corresponding to point i . In the current approach we will 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.

2.2.2 Segmentation

Segmentation Prior The quality of the segmentation hypotheses, that are provided by the estimated deformation of 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 these cases, one can exploit additional more robust cues to enhance segmentation estimation. Local appearance provides complementary information that can be incorporated into the segmentation modeling energy terms. Assuming a probability function on the candidate labeling of the form $\pi_x(l)$ where $l \in \{0, \dots, M-1\}$ we wish to penalize all segmentations that go against prior information:

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

Such probabilities can be learned using any modern classification method.

2.2.3 Coupling

In standard multi-atlas frameworks, fusion strategies follow multi-atlas registration in order to estimate the segmentation of the target image. These fusion strategies take into account the level of agreement between the hypotheses derived from the registered atlases to attribute a label. In this work, we close the circuit between multi-atlas registration and label fusion. We allow segmentation to influence registration. By imposing agreement between the segmentation estimate and the atlas hypotheses, improved segmentation accuracy can 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 deformations that lead to disagreement between the hypotheses and the estimated segmentation in a robust manner.

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

where $\kappa(a, b)$ is a penalizing disagreement function, for which we assume $\kappa(a, a) = 0$ and $\hat{\rho}$ is a similarity function allowing for atlases that match well the query image to influence more the segmentation variables. As widely used by label fusion schemes [2], this term acts as a sort of soft local atlas selection. In the context of the proposed method, it allows coupling information to be shared by atlases whose registration to the query image is locally reliable.

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:

$$E_{MRF}(\mathbf{I}) = \sum_{p \in \mathcal{V}} g_p(l_p) + \sum_{(p,q) \in \mathcal{E}} f_{pq}(l_p, l_q) \quad (7)$$

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 constraint that integrates the two problems. Let us see now 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 superimposed onto image A_i , there is a node $p \in \mathcal{V}_{D_i}$ that represents its displacement. The edge system of each grid \mathcal{E}_{D_i} is created by a regular connectivity scheme. Edges encode interactions between random variables. The solution space around a control point is quantized and indexed by a discrete set of variables \mathcal{L}_D . This set represents possible control point displacements. We refer to a potential control point displacement attributed to a deformation node by l^d .

Segmentation A set of nodes \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 .

Integrated Segmentation and Multi-Atlas Registration Integrating segmentation and multi-atlas registration is achieved by coupling segmentation and deformation variables. The set of edges \mathcal{E}_C connects nodes of \mathcal{V}_S with nodes of the set \mathcal{V}_D . Nodes coming from different graphs are connected if they have common image support.

3.2 MRF Energy

The continuous energy in Eq. 1 is mapped to a discrete MRF energy of the form in Eq. 7. In brief, we will map i) the matching term M (Eq. 2) to the unary potentials of deformation variables (Eq. 8), ii) the deformation smoothness penalty term R (Eq. 3) to pairwise potentials between deformation variables (Eq. 9), and iii) the coupling penalty C (Eq. 6) to pairwise potentials between deformation and segmentation variables (Eq. 11).

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 [5]. 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.

For the matching term, we are interested in quantifying how well the assignment of a displacement label $l_{p_i}^d \in \mathcal{L}_D$ to a node $p \in \mathcal{V}_{D_i}$ aligns atlas A_i to the target image. This is




Figure 1: Graph structure. Different colored edges / nodes correspond to different types energies. Note that nodes from different deformation grids are not directly connected. Only edges connecting a single deformation node from each grid to the segmentation variables are shown for clarity.

measured by the following unary potential:

$$g_{p_i}^M(I_{p_i}^d) = \int_{\Omega} \hat{\omega}_{p_i}(x) \rho(A_i \circ D_i^{I_{p_i}^d}, I(x)) dx. \quad (8)$$

$D_i^{I_{p_i}^d}$ is the transformation induced by the movement of the control point p in the i th deformation grid by the displacement $I_{p_i}^d$. 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, [5] 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:

$$f_{p_i q_i}^R(I_{p_i}^d, I_{q_i}^d) = \|\mathbf{d}^{I_{p_i}^d} - \mathbf{d}^{I_{q_i}^d}\|, \quad (9)$$

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

3.2.2 Segmentation

In order to assign a class label to every voxel of the target image, we take into account the learned appearance model for every class. The appearance model is encoded in the form of a probability distribution $\pi_x(l)$ 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:

$$g_{q_s}^{SP}(I_{q_s}^s) = \int_{\Omega} \hat{\omega}_{q_s}(x) (-\log(\pi_x(I_{q_s}^s))). \quad (10)$$

where $\hat{\omega}_{q_S}(x)$ denotes the support of the segmentation node. For our experiments we have corresponded each segmentation node to a voxel, but a coarser or softer assignment could be envisaged.

3.2.3 Integrated Segmentation and Multi-Atlas Registration

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

$$f_{p_i q_S}^C(l_{p_i}^d, l_{q_S}^s) = \int_{\Omega} \hat{\omega}_{q_S}(x) \hat{\omega}_{p_i}(x) \hat{\rho}(A_i \circ D_i^{p_i}, I(x)) \text{Ind}(S_i \circ D_i^{p_i}(x), l_{q_S}^s) dx \quad (11)$$

where p_i belongs to the grid \mathcal{G}_{D_i} and q_S belongs to \mathcal{G}_S . $\text{Ind}(x, y) = 1$ except from $\text{Ind}(x, x) = 0$. Note that this potential is non zero only for pairs of segmentation/deformation nodes that share some common support in the image domain, *i.e.*, when $\exists x \mid \hat{\omega}_{q_S}(x) \neq 0, \hat{\omega}_{p_i}(x) \neq 0$.

3.3 MRF Optimization through Dual Decomposition

DD-MRF [7] 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 solution and update their costs. 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 To optimize graph \mathcal{G} , DD-MRF requires as input as set of subproblems $\mathbf{SD} = \{SD_i, \dots, SD_n\}$, such that: $\mathcal{G} = \bigcup_i SD_i$. We decompose the problem into a series of subproblems which can be exactly optimized through dynamic programming. The grid subgraph G_{D_i} will be decomposed into chain subproblems. C_{D_i} will represent the result of this decomposition, where every element of the set corresponds to a chain subproblem.

The agreement term is incorporated into tree subproblems. There is one such subproblem per segmentation node. It consists of the segmentation node, the deformation nodes to which it is connected, as well as the edges connecting them. Tree subproblems form the set T_S . Thus, our subproblem decomposition will be: $\mathbf{SD} = \{T_S, CD_0, \dots, CD_N\}$.

4 Validation

We have validated our method using data available on the Internet Brain Segmentation Repository (IBSR). We specifically use the skull stripped version of the dataset provided in [9]. 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 have made experiments with fully annotated images (35 annotation types on average per image). Atlases have been pre-registered in all possible




Figure 2: Left: Overlap results (DICE coefficient). Right: Overlap gain per registration comparing to the majority voting of the decoupled model for 'Labeling' 'Local' and 'Majority' and their corresponding results for 'Inter' and 'Atlas'. Refer to the Results Discussion in Chapter 4 for information on the types of results presented. 'coupled' refers to the method presented while 'decoupled' refers to the result of independent pairwise registration.

combinations using highly regularized deformable registration to find the ones closer to another. A weighted sum of harmonic energy and global data cost have been used to rank them. Experiments have been conducted using the 6 and 10 closest atlases to each target image.

Registration parameters The same iterative procedure as [5] 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 ($15 \times 18 \times 14$ and $29 \times 35 \times 27$). We use the deformation fields produced by the coarse level to initialize the finer one. For deformation nodes, candidate displacements span regularly 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. $(1 - NCC)^2$ was used as the similarity function ($\hat{\rho}$).

Segmentation Likelihoods Training has been conducted by leaving out 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 iv) normalized voxel positions and distance from the center voxel. The total size of the feature vector was 258. We use the Random Forest framework [3] 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.




Figure 3: Visual results. 1st row : Ground truth. 2nd row: Segmentation node labeling. 3rd row: Locally weighted fusion using deformations produced by the decoupled model. 4th row: Majority voting fusion using deformations produced by the decoupled model.

Optimization convergence Concerning DD-MRF convergence, we consider further optimization iterations useless 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.

Decoupled Model As a baseline to compare to we consider N pairwise registrations with identical parameters as the ones already presented, optimized in the same manner as our method. A way to think of these registrations is as if there were no segmentation variables and we had N separate (decoupled) deformation grids.

Results Discussion Fig. 2 presents overlap results of certain types. *Local* corresponds to deformed atlases fused using local weighted voting where the same similarity criterion as the one used in registration has been employed. *Majority* corresponds to deformed atlases fused using majority voting. *Labeling* corresponds to the labeling of the segmentation nodes as given by our method. *Class_argmax* corresponds to the most probable label per voxel as given by the Random Forest output. *Atlas* corresponds to the mean overlap of each deformed atlas compared to the ground truth. *Inter* corresponds to the mean overlap of every possible comparison between deformed atlases. Please note here that the reported standard deviations of the absolute overlap results (Left of Fig. 2) are overwhelmed by the intra-variability of the dataset. Overlap gain results (Right of Fig. 2), computed over the results of the decoupled model, provide more insight on the statistical properties of the method.

Figures from the fully annotated 6 atlas case in Fig. 2 strongly support that the enforced agreement constraints enhance pairwise registration: overlap is clearly improved (0.70 ± 0.02 from 0.68 ± 0.02 with a gain of 0.023 ± 0.004) while the resulting deformation fields are much smoother (harmonic energy of 0.58 ± 0.02 for our method comparing to 0.63 ± 0.03 for the decoupled registrations with a gain of -0.04 ± 0.008). At this point let us recall, that segmentation through registration draws its power from the common spatial structure of the query and atlas images. Combining segmentations from different structures of many images solely based on good intensity based matchings, while beneficial for the mean overlap, reduces the reliability of the resulting segmentation. Our method has a substantial impact on the agreement between atlases (0.715 ± 0.01 mean Dice of all possible registrations between atlases for our method, comparing to 0.66 ± 0.015 for the decoupled case with a gain of 0.053 ± 0.004) producing concordant pairwise registrations, while outperforming the label fusion methods in segmentation accuracy. Using the 10 closest atlases increases overlap as expected while retaining the trends established by the 6-atlas case.

Visual results from the 6 atlas case in Fig. 3 verify further our method’s properties. It is easy to see that the agreement constraints create, while no explicit smoothness term is present in the energy function, much smoother results by disallowing aggressive displacement fields that might correspond well visually but are not in accordance with the atlas group consensus.

5 Discussion and future work

In this paper, we presented a method that integrates registration and segmentation fusion. The proposed approach allows registration parameters to be updated based on the segmentation estimation. As a consequence, the quality of the obtained matching is refined, resulting in increased segmentation accuracy. The experimental results demonstrate the potential of the proposed approach.

References

- [1] P. Aljabar, R.A. Heckemann, A. Hammers, J.V. Hajnal, and D. Rueckert. Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy. *NeuroImage*, 46(3):726–738, 2009.
- [2] 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*, 28(8):1266–1277, August 2009.
- [3] Leo Breiman. Random Forests. *Mach. Learn.*, 45, 2001.
- [4] 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*, 17(8):1181–1191, December 2013. ISSN 13618415. doi: 10.1016/j.media.2013.08.001. URL <http://dx.doi.org/10.1016/j.media.2013.08.001>.
- [5] Ben Glocker, Aristeidis Sotiras, Nikos Komodakis, and Nikos Paragios. Deformable Medical Image Registration: Setting the State of the Art with Discrete Methods*. *Annual Review of Biomedical Engineering*, 13(1):219–244, 2011.
- [6] I. Isgum, M. Staring, A. Rutten, M. Prokop, M. A. Viergever, and B. van Ginneken. 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, July 2009. URL <http://dx.doi.org/10.1109/tmi.2008.2011480>.
- [7] Nikos Komodakis, Nikos Paragios, and Georgios Tziritas. MRF energy minimization and beyond via dual decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33:531–552, 2011.
- [8] Jan Kybic and Michael Unser. Fast parametric elastic image registration. *Image Processing, IEEE Transactions on*, 12(11):1427–1442, 2003.
- [9] Torsten Rohlfing. Image similarity and tissue overlaps as surrogates for image registration accuracy: widely used but unreliable. *IEEE transactions on medical imaging*, 31, 2012.
- [10] Torsten Rohlfing, DanielB Russakoff, and CalvinR Maurer. Expectation Maximization Strategies for Multi-atlas Multi-label Segmentation. In *Information Processing in Medical Imaging*, volume 2732 of *Lecture Notes in Computer Science*, pages 210–221. Springer Berlin Heidelberg, 2003. URL http://dx.doi.org/10.1007/978-3-540-45087-0_18.
- [11] Torsten Rohlfing, Robert Brandt, Randolph Menzel, and Calvin R Maurer Jr. Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *NeuroImage*, 21(4):1428–1442, 2004.
- [12] Daniel Rueckert, Luke I Sonoda, Carmes Hayes, Derek L. G. Hill, Martin O. Leach, and David J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *Medical Imaging, IEEE Transactions on*, 18(8):712–721, 1999.

- [13] 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*, 29(10):1714–1729, October 2010. ISSN 0278-0062. doi: 10.1109/tmi.2010.2050897. URL <http://dx.doi.org/10.1109/tmi.2010.2050897>.
- [14] M. Sdika. Combining atlas based segmentation and intensity classification with nearest neighbor transform and accuracy weighted vote. *Medical Image Analysis*, 14(2):219 – 226, 2010.