

Multi-atlas Segmentation as a Graph Labelling Problem: Application to Partially Annotated Atlas Data

Lisa M. Koch¹(✉), Martin Rajchl¹, Tong Tong¹, Jonathan Passerat-Palmbach¹, Paul Aljabar^{1,2}, and Daniel Rueckert¹

¹ Biomedical Image Analysis Group, Imperial College London, London, UK
l.koch@imperial.ac.uk

² Division of Imaging Sciences and Biomedical Engineering, King's College London, London, UK

Abstract. Manually annotating images for multi-atlas segmentation is an expensive and often limiting factor in reliable automated segmentation of large databases. Segmentation methods requiring only a proportion of each atlas image to be labelled could potentially reduce the workload on expert raters tasked with labelling images. However, exploiting such a database of partially labelled atlases is not possible with state-of-the-art multi-atlas segmentation methods. In this paper we revisit the problem of multi-atlas segmentation and formulate its solution in terms of graph-labelling. Our graphical approach uses a Markov Random Field (MRF) formulation of the problem and constructs a graph connecting atlases and the target image. This provides a unifying framework for label propagation. More importantly, the proposed method can be used for segmentation using only partially labelled atlases. We furthermore provide an extension to an existing continuous MRF optimisation method to solve the proposed problem formulation. We show that the proposed method, applied to hippocampal segmentation of 202 subjects from the ADNI database, remains robust and accurate even when the proportion of manually labelled slices in the atlases is reduced to 20 %.

1 Introduction

In recent years, major efforts have been undertaken towards building large medical image databases, e.g. in population studies. As the wealth of data increases, automated segmentation of images becomes crucial while manually annotating them becomes prohibitive. In particular, robust and accurate segmentation techniques relying on minimal manual input become increasingly desirable.

Atlas-based segmentation has proven to be a successful and robust tool for a number of applications. Many of these techniques [1, 2, 7] rely on label propagation from multiple suitable atlases after non-linear registration to a target image. The target segmentation can be formed by label fusion of the propagated labels, for example by applying a majority vote rule [1, 7] or another combination strategy such as a weighted average based on global or local similarity measures

between the target and atlas images [2]. Other combination strategies include STAPLE [20], where label fusion weights are estimated with an Expectation-Minimisation algorithm, or Joint Label Fusion [19], where correlations among atlases are taken into account. To account for high local anatomical variability between images, patch-based segmentation [5, 17] has been introduced, where the label fusion step employs a non-local weighted average of voxel labels in a small neighbourhood of the atlas images, with weights based on the similarities of patches centred on the compared voxels. Considerable improvements to results obtained with label propagation can be achieved by using the label propagation results as prior probabilities in subsequent refinement steps, combining them with regularisation terms and an intensity model in a Markov Random Field (MRF) formulation [11–13]. It has been shown that in general, segmentation accuracy decreases when fewer [7] or less suitable [1] atlases are used.

All of the above methods rely on the availability of fully annotated atlas data. However, segmentation methods requiring only a proportion of each atlas image to be labelled (while no knowledge is necessary about the remaining voxels) could reduce the workload of raters who manually label these atlases. [9] proposed an extension to STAPLE [20] which, to our knowledge, is the only existing multi-atlas segmentation method that uses partially annotated atlas data. However, partial annotations have been used in the context of interactive segmentation. In [4], regions of an image were manually labelled to enable automated segmentation of the remaining image. In particular, an MRF energy function [10] was formulated on a graph constructed on the regular image grid in which annotated voxels are connected to virtual terminal nodes with infinite weights. The MRF energy minimisation problem was efficiently solved by finding a min-cut/max-flow on the graph with graph-cuts [4]. An iterative graph-cuts approach [16] has been proposed to reduce user interaction compared to [4].

Some applications of graph-cuts employ MRF energy functions that have been formulated for labelling on graphs connecting more than one image. Recently, [6] applied graph-cuts for tumor segmentation based on PET and CT image pairs by minimising an MRF energy function which penalises segmentation differences between a PET and CT image of the same subject. [15] proposed a prostate segmentation algorithm in which multiple 2D slices were extracted from a 3D image at different angles. Exploiting axial symmetry, the 2D slices were simultaneously segmented using a max-flow algorithm which penalised segmentation differences between slices at similar angles. The max-flow algorithm they used is an extension of the recently proposed Continuous Max-Flow (CMF), which solves a continuous counterpart to the discrete min-cut/max-flow problem [21]. CMF can be computed using a reliable, highly parallelisable multiplier-based algorithm with guaranteed convergence, making it suitable for the optimisation of large labelling problems in parallel computing environments.

Combining aspects of multi-atlas segmentation with recent developments in min-cut/max-flow methods [15, 21], and with the goal of successfully exploiting partially labelled images for segmentation, we propose a segmentation framework incorporating three novel contributions:

1. We revisit the labelling problem in existing multi-atlas segmentation methods and express the solution in terms of labelling a graph which connects the target and atlas images. We formulate an MRF energy function on this graph and show analytically that its solution is equivalent to existing multi-atlas segmentation methods [2, 7]. We then show how spatial regularisation and intensity models [11] can be readily incorporated in the proposed formulation.
2. We provide a generalisation of the Continuous Max-Flow algorithm which can efficiently minimise energy functions on graphs connecting multiple images. The proposed generalised CMF is applied to the graph labelling problem formulated here. However, as it is highly parallelisable, its scalability to large graphs could make it useful for solving large-scale MRF frameworks beyond the scope of this work.
3. We show that the proposed method can provide automated segmentations using partially annotated atlas data. To evaluate the method, we apply it to hippocampal segmentation in 202 Magnetic Resonance (MR) images of the ADNI database and investigate its performance under varying proportions of labelled voxels per atlas image.

2 Method

In the following sections, we revisit multi-atlas segmentation (Sect. 2.1) and show how it can be viewed as a labelling problem on a graph connecting atlases and the target image (Sect. 2.2). In Sect. 2.3 we discuss extensions to the proposed MRF energy function to integrate label propagation, regularisation, as well as intensity models into a single comprehensive framework. Furthermore, we show how the proposed methodology can be applied to segmentation problems where only partially annotated atlas data are available (Sect. 2.4). The optimisation of the proposed energy function is discussed in Sect. 2.5.

2.1 Multi-atlas Label Propagation

For multi-atlas segmentation [2, 7] using R images, all atlas images $j \in \{1, \dots, R\}$ are registered to the target image i . For convenience we assume $i = R + 1$. The label maps l_j associated with the atlas images j are then propagated to the target. Figure 1a shows an example atlas set with corresponding label maps, and an unlabelled target image. Each voxel $x \in \Omega$ in the target image i is labelled using some combination strategy, e.g. a weighted average of atlas labels $l_j(x)$:

$$l_i(x) = \arg \max_L \sum_{j=1}^R \beta_{ij}(x) \delta(l_j(x) = L) \quad (1)$$

Here $\delta(\cdot)$ is an indicator function. The weights $\beta_{ij}(x)$ can be uniform (which is equivalent to a majority vote rule) or based on global or local similarity measures between images i and j . Suitable measures were investigated in [2].

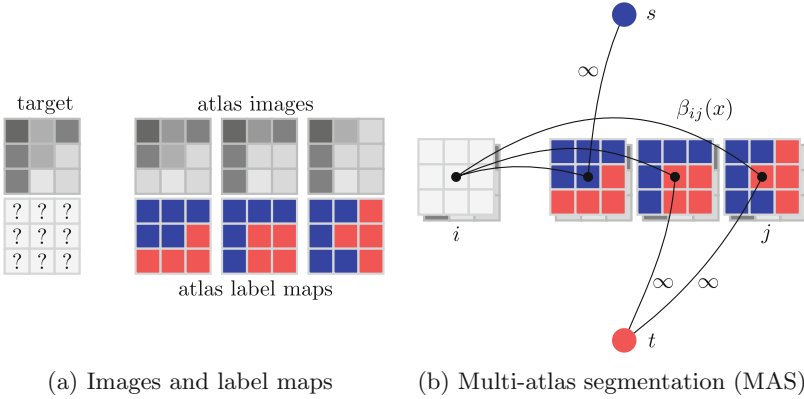


Fig. 1. (a) An example dataset with an unlabelled target image on the left, atlas images and corresponding manual annotations (blue and red depict different labels) on the right. (b) In MAS, each voxel x in target image i is labelled by label propagation from atlases $j \in \{1, \dots, R\}$ with fusion weights $\beta_{ij}(x)$. This can also be interpreted as a graph labelling problem, where atlas voxels are connected to the terminal nodes with infinitely weighted edges (Color figure online).

2.2 Reformulation as a Graph Labelling Problem

As an alternative perspective, we can construct a graph to model the relationship of shared information between the atlases and the target. According to the above labelling scenario, this graph connects each voxel x in the target image i to the corresponding voxels in the atlas images j with an edge weighted by $\beta_{ij}(x)$. Figure 1b visualises this configuration. To find a labelling on the graph, we can formulate a potential function V that penalises conflicting labels in voxels connected by a high weight $\beta_{ij}(x)$, e.g.

$$V(l_i(x), l_j(x)) = \beta_{ij}(x) \delta(l_j(x) \neq l_i(x)) \quad (2)$$

This assigns a high penalty when the target and atlas labels differ and the atlas is considered similar to the target i , as defined by the similarity measure $\beta_{ij}(x)$. In the case of a majority vote, the weights are uniform, e.g. $\beta_{ij}(x) = 1$. The cost for labelling an individual voxel x in image i can then be calculated as follows:

$$E_{\text{propagation}}(l_i(x)) = \sum_{j=1}^R V(l_i(x), l_j(x)) \quad (3)$$

$$= \sum_{j=1}^R \beta_{ij}(x) \delta(l_j(x) \neq l_i(x)) \quad (4)$$

$$= \sum_{j=1}^R \beta_{ij}(x) - \sum_{j=1}^R \beta_{ij}(x) \delta(l_j(x) = l_i(x)) \quad (5)$$

In this formulation the labels in the atlases are fixed. This is achieved by assigning infinite unary potentials to the atlas voxels (visualised as infinitely weighted edges to the terminal nodes in Fig. 1b for the binary case). Voxels in the target image are assumed to be conditionally independent since spatially neighbouring voxels in the target image are not connected in the graph (in contrast to the setting for regularisation in many vision problems [10]). The optimal label can therefore be found by minimising $E_{\text{propagation}}(l_i(x))$ independently for all voxels:

$$l_i(x) = \arg \min_L E_{\text{propagation}}(l_i(x) = L) \quad (6)$$

$$= \arg \min_L - \sum_{j=1}^R \beta_{ij}(x) \delta(l_j(x) = L) \quad (7)$$

$$= \arg \max_L \sum_{j=1}^R \beta_{ij}(x) \delta(l_j(x) = L), \quad (8)$$

which leads to the same result as the vote rule in Eq. 1. This demonstrates that multi-atlas segmentation can be expressed in terms of a graph optimisation problem. It is important to note that patch-based segmentation (PBS [5, 17]) can also be expressed in this framework. In this case we can use a slightly different graph structure as the label fusion step in PBS takes into account multiple voxels in a neighbourhood of x in each atlas instead of just one voxel at location x .

While this alternative problem formulation does not provide any immediate benefits over traditional multi-atlas segmentation in itself (because it is equivalent), it has two advantages: (1) it allows easy integration of additional components and therefore provides a unifying reformulation for existing multi-atlas segmentation methods, as shown in Sect. 2.3 and (2) the graphical approach extends to segmentation using partially annotated atlases (Sect. 2.4).

2.3 Unified Label Propagation Framework

As we are interpreting the whole dataset, comprising target and atlas images, as one graph satisfying Markov properties, we can assign unary potentials (data terms) to each voxel in all images, and pairwise potentials to each pair of voxels [10]. In the previous section, we proposed assigning pairwise potentials between target and atlas voxels for label propagation. In addition, we assigned infinite unary potentials to the atlas voxels since their labels are fixed. It is important to note that a data term can be specified for the target image as well, using prior probabilities, intensity models of the data, or a combination of both. Lastly, we can incorporate spatial regularisation with pairwise potentials between adjacent voxels within an image. The propagation, data, and regularisation terms can be combined to a comprehensive labelling energy function on the whole graph:

$$E(l) = E_{\text{data}}(l) + E_{\text{regularisation}}(l) + E_{\text{propagation}}(l) \quad (9)$$

As mentioned in the introduction, many existing multi-atlas segmentation methods (e.g. [11, 13]) use an MRF formulation to improve label propagation results

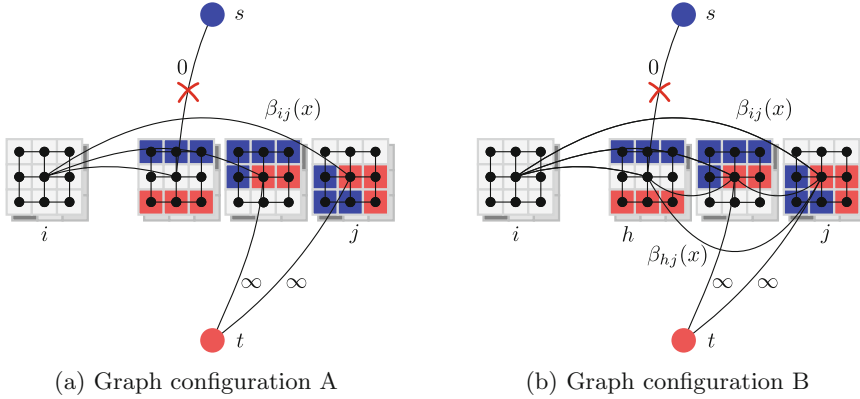


Fig. 2. Graph connections with partially annotated atlas data (blue and red depict different labels), based on the example dataset of Fig. 1a. Voxels with missing labels (white) are disconnected from terminal nodes. Spatial regularisation is enabled in all images. (a) Voxels at each location x in the target image are connected to voxels in atlases j . (b) Additionally, atlas voxels are connected to voxels in other atlases (Color figure online).

with the benefits of regularisation and intensity data models. However, these approaches use probabilistic label propagation results as prior probabilities (i.e. unary potentials) in a *subsequent* refinement step, therefore adding the MRF optimisation as a separate post-processing step. The above comprehensive formulation treats label propagation as part of the optimisation process, and therefore unifies all the components within a single framework. Using the above general formulation, it is possible to open-up a new field of applications, namely segmentation using partially annotated atlas data.

2.4 Segmentation Using Partially Annotated Atlas Data

In the graph setting proposed in the previous sections, an atlas label is characterised by an infinitely weighted connection to a terminal node. For a partially annotated atlas image, unlabelled voxels may simply be disconnected from the terminal nodes. In a multi-atlas segmentation problem as described by Fig. 1b, this translates to a segmentation based only on the available labels. In this case, however, there is no guarantee that each location in the target image can be segmented (depending on the extent of missing labels in the atlases at specific voxel locations), and performance is degraded for low proportions of labelled voxels. We therefore propose the following two graph configurations (Fig. 2): **(A)** Labels are propagated from atlases to the target, i.e. label propagation edges $\beta_{ij}(x)$ are used as in the multi-atlas segmentation case. Additionally, spatial regularisation is used in all images so that labels may be shared between similar regions with labelled and unlabelled voxels, and which can then be propagated to the target image. **(B)** Instead of only propagating labels between target and atlases,

atlas images are also interconnected. This serves to facilitate the propagation, especially when manual labels are very scarce.

2.5 Optimisation

It has been shown that MRF energy functions consisting of unary and pairwise terms can be minimised using min-cut/max-flow approaches if the pairwise terms are metric or semi-metric [3], yielding globally optimal results for binary labelling problems and approximately globally optimal results for multiple labels [3]. Recently, [21] proposed a continuous max-flow (CMF) algorithm in the 2D or 3D domain (i.e. a single image) which is highly parallelisable in contrast to discrete graph-based methods [21]. As the proposed energy function can be optimised for a large graph consisting of voxels in all images and their interactions, this approach was used and extended for graphs between multiple images.

Analogous to discrete max-flow approaches, the energy function on the graph can be minimised by maximising a source flow p^s through the network, subject to flow conservation and capacity constraints on the edges. In the original CMF algorithm [21], spatial flows $\mathbf{p} = [p_x, p_y, p_z]^T$ exist between adjacent voxels in the image domain Ω (for regularisation) and source and sink flows $p^{s,t}$ between voxels and terminal nodes. The optimisation is performed with a variational approach by introducing a Lagrange multiplier $u(x)$ to incorporate the constraints [21]. It has been shown that the resulting $u(x)$ corresponds to the globally optimal labelling [21]. While CMF has been extended to multi-label segmentation problems [22], we restrict the scope of this paper to binary labelling problems.

In the following, we propose a generalisation of CMF from a single image to an arbitrary configuration of interconnected images to account for any user-defined choice of inter-image relationships $\beta_{ij}(x)$. Figure 3 shows the capacity constraints and introduces the notation for inter-image flows $r_{ij}(x)$ (for label propagation), spatial flows $\mathbf{p}_i(x)$ (for regularisation) and terminal flows $p_i^{s,t}(x)$ (for unary priors). The notation is similar to [15], where inter-image constraints were used in a different context. To satisfy flow conservation, the sum of all in-




Fig. 3. Notation for flow constraints $\beta_{ij}(x)$, $C_i^{s,t}(x)$, $\alpha_i(x)$ for propagation, data term and regularisation, and corresponding inter-image flows $r_{ij}(x)$, source and sink flows $p_i^{s,t}(x)$ and spatial flows $\mathbf{p}_i(x)$ at location x in image i .

and outgoing flows $\rho_i(x)$ at each node must be zero, i.e.

$$\rho_i(x) = \operatorname{div} \mathbf{p}_i(x) - p_i^s(x) + p_i^t(x) + \sum_{j=1, j \neq i}^n r_{ij}(x) = 0, \quad (10)$$

where $r_{ij}(x) = -r_{ji}(x)$ and n is the number of images in the graph. This leads to the Lagrangian function

$$L(u, p^s, p^t, \mathbf{p}, r) = \sum_{i=1}^n \left(\int_{\Omega} p_i^s dx + \langle u_i, \rho_i \rangle - \frac{c}{2} \|\rho_i\|^2 \right), \quad (11)$$

which can be solved iteratively by optimising each variable $u, p^s, p^t, \mathbf{p}, r$ separately [15, 21]. The novel component compared to [15, 21] is the use of inter-image flows $r_{ij}(x)$ between any pair of images i, j . We therefore show in particular the optimisation step at iteration k for $r_{ij}(x)$ while fixing all other variables:

$$r_{ij}^{k+1} = \arg \max_{|r_{ij}| \leq \beta_{ij}} L(u, p^s, p^t, \mathbf{p}, r) \quad (12)$$

$$= \arg \max_{|r_{ij}| \leq \beta_{ij}} u_i^k r_{ij}^k + u_j^k r_{ji}^k \quad (13)$$

$$\begin{aligned} & - \frac{c}{2} \left\| \left(\operatorname{div} \mathbf{p}_i - p_i^s + p_i^t \right)^k + \sum_{l=1, l \neq i, j}^n r_{il}^k + r_{ij}^k \right\|^2 \\ & - \frac{c}{2} \left\| \left(\operatorname{div} \mathbf{p}_j - p_j^s + p_j^t \right)^k + \sum_{l=1, l \neq j, i}^n r_{jl}^k + r_{ji}^k \right\|^2 \\ = \arg \max_{|r_{ij}| \leq \beta_{ij}} & - \frac{c}{2} \left\| \left(\operatorname{div} \mathbf{p}_i - p_i^s + p_i^t \right)^k + \sum_{l=1, l \neq i, j}^n r_{il}^k - \frac{u_i^k}{c} + r_{ij}^k \right\|^2 \quad (14) \end{aligned}$$

$$\begin{aligned} & - \frac{c}{2} \left\| \left(\operatorname{div} \mathbf{p}_j - p_j^s + p_j^t \right)^k + \sum_{l=1, l \neq j, i}^n r_{jl}^k - \frac{u_j^k}{c} - r_{ij}^k \right\|^2 \\ = \arg \max_{|r_{ij}| \leq \beta_{ij}} & - \frac{c}{2} \|J_i^k + r_{ij}^k\|^2 - \frac{c}{2} \|J_j^k - r_{ij}^k\|^2 \quad (15) \end{aligned}$$

where

$$J_i^k = \left(\operatorname{div} \mathbf{p}_i - p_i^s + p_i^t \right)^k + \sum_{l=1, l \neq i, j}^n r_{il}^k - \frac{u_i^k}{c} \quad (16)$$

This leads to

$$r_{ij}^{k+1} = \begin{cases} -\beta_{ij}, & \frac{1}{2}(J_j^k - J_i^k) \leq -\beta_{ij}, \\ \frac{1}{2}(J_j^k - J_i^k), & |\frac{1}{2}(J_j^k - J_i^k)| \leq \beta_{ij}, \\ \beta_{ij}, & \text{otherwise.} \end{cases} \quad (17)$$

After convergence, a binary segmentation can be found by thresholding the resulting solution for u , e.g. at 0.5.

3 Application to Segmentation of Partially Labelled Hippocampus Data

Manually annotating medical images is very time consuming, placing a major burden on clinical experts tasked with labelling large datasets. The images are typically manually annotated slice-by-slice, therefore reducing the proportion of annotated slices while retaining robust and accurate segmentation is an important goal. With the following experiments we investigate the potential of the proposed method to provide accurate hippocampal segmentation using partially annotated atlas data. While demonstrating a scenario where a proportion of atlas *slices* are annotated, the framework proposed in this paper could be applied to different configurations of partial labels, for example to datasets where different regions are annotated in different images.

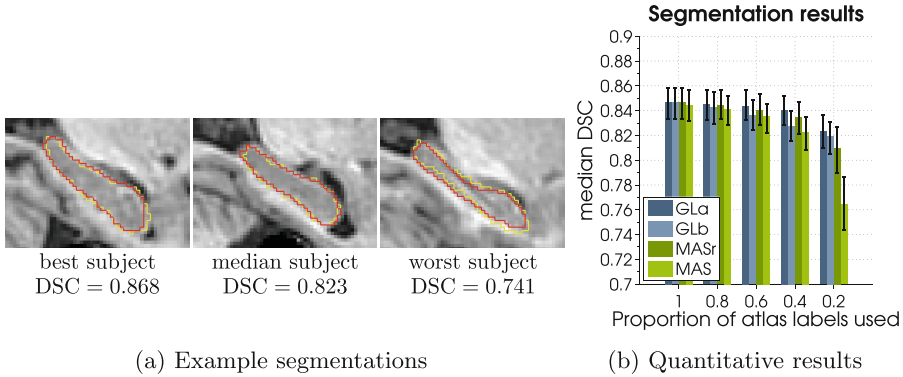
Data and Experiment Setup. The proposed method was applied to 202 images from the ADNI database [8] for which reference segmentations of the hippocampus were available through ADNI. All images were affinely aligned to the MNI152 template space and intensity-normalised [14]. A leave-one-out procedure was performed for evaluation where, for each target subject, the 10 most similar subjects were selected as atlases using normalised mutual information in a region of interest around the hippocampus. The selected atlases were non-rigidly registered to the target subject [18] with a control-point spacing of 5 mm.

Partially Annotated Atlases. Manual labels of a proportion q of slices in the atlas images were used for segmentation of the target image, while source and sink connections were removed in the remaining slices (as shown in Fig. 2). The selected slice positions were evenly distributed but varied in different atlases. For the spatial regularisation constraints (Fig. 3), we chose

$$\alpha_i(x) = c \exp\left(-\frac{\|\nabla I(x)\|^2}{2\sigma^2}\right), \quad (18)$$

which is a continuous equivalent of the regularisation term used in [11]. The parameters $c = 0.1$ and $\sigma = 50$ were heuristically tuned for the comparison methods introduced below. We report results for two graph configurations proposed in Sect. 2.4: Graph labelling using (a) connections between target and atlases (GLa) and (b) connections between all pairs of images (GLb). The similarity measure $\beta_{ij}(x)$ was chosen as a function of the local mean squared distance (LMSD) as recommended by [2] for hippocampus segmentation. The LMSD was evaluated in a cubic neighbourhood of radius 5 around each location [2].

Fully Annotated Atlases. We compared the proposed method to locally weighted multi-atlas segmentation (MAS [2]) with different numbers of atlases. The propagated atlas labels were fused with a weighted average as in Eq. 1. The label fusion was implemented using the proposed framework (since we have shown in Eq. 8 that this formulation is equivalent to MAS). We also compared



(a) Example segmentations

(b) Quantitative results

Fig. 4. (a) Best, median, and worst segmentation results (red) obtained with GLa using 20 % of available atlas slices, compared to the reference segmentation (yellow) (b) Median DSC and interquartile range for the proposed (blue) and comparison (green) methods using different proportions of atlas labels. The proposed methods use a proportion of slices, whereas the comparison methods use a proportion of images for segmentation (Color figure online).

against a variation of MAS using regularisation in the target image (MASr). Regularisation parameters (for MASr) and $\beta_{ij}(x)$ were chosen as in the proposed method to guarantee fair comparisons.

Results. Experiments were run for $q = \{1, 0.8, 0.6, 0.4, 0.2\}$, where $q = 1$ means that the 10 selected atlases were fully labelled. To measure the effect of partial annotations compared to full annotations, for proportion q , exactly the same number of labelled voxels were available to both the proposed methods and the comparison methods: either as a proportion of slices in each image (labelled voxels were distributed amongst all atlases), or as a proportion of the complete atlas set. This means that for the comparison methods MAS and MASr, fewer (but fully labelled) atlas images were used. Segmentation accuracy is reported with the median Dice Similarity Coefficient (DSC) and interquartile range. As the results were not normally distributed, statistical significance is measured using the two-sided Wilcoxon signed-rank test and is reported for $p < 10^{-4}$.

As expected, segmentation accuracy decreased when less atlases were used [7]. This was observed most strongly in MAS, while spatial regularisation (MASr) increased segmentation accuracy considerably, particularly when fewer atlases were used. The proposed methods GLa and GLb are identical to MASr for $q = 1$ where all atlas voxels are labelled. As the proportion of labelled voxels decreased, GLa was significantly more accurate than MASr for $q = 0.6$, $q = 0.4$ and $q = 0.2$. In GLb, where atlas images were densely interconnected, improvements over MAS and MASr could only be found for $q = 0.2$. The most distinct differences between the methods could be observed when only 20 % of the atlas voxels were labelled. Median DSC values are 0.823(0.03) and 0.819(0.03) for GLa and GLb compared to 0.764(0.04) and 0.810(0.04) for MAS and MASr. Figure 4a shows

example segmentations for GLa, for which the runtime per subject on a single CPU was 5 min when computation was limited to the hippocampal region.

4 Discussion and Conclusion

The experiments show that with the proposed method and graph configuration A (Fig. 2a), segmentation accuracy remained stable when reducing the proportion of labelled atlas voxels down to 40%. The segmentation results were still relatively accurate at 20%, so the required user input could be reduced considerably. In particular, the proposed method significantly outperformed multi-atlas segmentation using fewer, but fully labelled atlases. This suggests that it could be worthwhile to allocate resources for partially annotating more images instead of fully annotating few images. More generally, we demonstrated that partially annotated images, a resource that is not usually exploited, can be utilised for multi-atlas segmentation.

In order to provide a framework for multi-atlas segmentation using partially annotated atlases, we make two key methodological contributions: We take a new perspective on the well-studied problem of multi-atlas segmentation and formulate it as a graph-labelling problem on a graph between target and atlas images. The proposed framework unifies a number of existing atlas-based segmentation methods (e.g. [2, 7, 11]) and can incorporate label propagation, regularisation, and intensity models in a single and comprehensive optimisation problem. In the scope of this paper, we have limited ourselves to label propagation and regularisation. In future work, these components will be further developed, and suitable intensity models will be investigated.

Furthermore, we have proposed a generalisation of the Continuous Max-Flow [21] algorithm for parallelisable optimisation of energy functions on graphs between interconnected images of any configuration. In this work, graphs were constructed between target and atlas images to establish an analogy to multi-atlas segmentation scenarios. However, the proposed method could be used for segmenting large-scale, partially labelled datasets.

References

1. Aljabar, P., Heckemann, R.A., Hammers, A., Hajnal, J.V., Rueckert, D.: Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *NeuroImage* **46**(3), 726–38 (2009)
2. Artaechevarria, X., Munoz-Barrutia, A., Ortiz-de Solórzano, C.: Combination strategies in multi-atlas image segmentation: application to brain MR data. *IEEE Trans. Med. Imaging* **28**(8), 1266–1277 (2009)
3. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Trans. PAMI* **23**(11), 1222–1239 (2001)
4. Boykov, Y., Jolly, M.-P.: Interactive organ segmentation using graph cuts. In: Delp, S.L., DiGoia, A.M., Jaramaz, B. (eds.) *MICCAI 2000*. LNCS, vol. 1935, pp. 276–286. Springer, Heidelberg (2000)
5. 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 (2011)

6. Han, D., Bayouth, J., Song, Q., Taurani, A., Sonka, M., Buatti, J., Wu, X.: Globally optimal tumor segmentation in PET-CT Images: a graph-based co-segmentation method. In: Székely, G., Hahn, H.K. (eds.) IPMI 2011. LNCS, vol. 6801, pp. 245–256. Springer, Heidelberg (2011)
7. 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 (2006)
8. Jack, C.R., Bernstein, M., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P., Whitwell, J., Ward, C., Dale, A., Felmlee, J., Gunter, J., Hill, D., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., DeCarli, C., Krueger, G., Ward, H., Metzger, G., Scott, K., Mallozzi, R., Blezek, D., Levy, J., Debbins, J., Fleisher, A., Albert, M., Green, R., Bartzokis, G., Glover, G., Mugler, J., Weiner, M.: The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *Magn. Reson. Imaging* **27**(4), 685–691 (2008)
9. Landman, B.A., Asman, A., Scoggins, A., Bogovic, J., Xing, F., Prince, J.: Robust statistical fusion of image labels. *IEEE Trans. Med. Imaging* **31**(2), 512–522 (2012)
10. Li, S.: Markov random field models in computer vision. In: *ECCV*, pp. 361–370 (1994)
11. van der Lijn, F., den Heijer, T., Breteler, M., Niessen, W.J.: Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts. *NeuroImage* **43**(4), 708–720 (2008)
12. Lötjönen, J.M., Wolz, R., Koikkalainen, J.R., Thurfjell, L., Waldemar, G., Soininen, H., Rueckert, D.: Fast and robust multi-atlas segmentation of brain magnetic resonance images. *NeuroImage* **49**(3), 2352–2365 (2010)
13. Makropoulos, A., Gousias, I.S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J.V., Edwards, D., Counsell, S.J., Rueckert, D.: Automatic whole brain MRI segmentation of the developing neonatal brain. *IEEE Trans. Med. Imaging* **33**(9), 1818–1831 (2014)
14. Nyúl, L.G., Udupa, J.K.: On standardizing the MR image intensity scale. *Magn. Reson. Med.* **42**(6), 1072–1081 (1999)
15. Qiu, W., Yuan, J., Ukwatta, E., Sun, Y., Rajchl, M., Fenster, A.: Prostate segmentation: an efficient convex optimization approach with axial symmetry using 3D TRUS and MR images. *IEEE Trans. Med. Imaging* **33**(4), 947–960 (2014)
16. Rother, C., Kolmogorov, V., Blake, A.: Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* **23**(3), 309–314 (2004)
17. Rousseau, F.: A supervised patch-based approach for human brain labeling. *IEEE Trans. Med. Imaging* **30**(10), 1852–1862 (2011)
18. 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. *IEEE Trans. Med. Imaging* **18**(8), 712–721 (1999)
19. Wang, H., Suh, J., Das, S., Pluta, J., Craige, C., Yushkevich, P.: Multi-atlas segmentation with joint label fusion. *IEEE Trans. PAMI* **35**(3), 611–623 (2012)
20. 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 Trans. Med. Imaging* **23**(7), 903–921 (2004)
21. Yuan, J., Bae, E., Tai, X.: A study on continuous max-flow and min-cut approaches. In: *CVPR*, pp. 2217–2224 (2010)
22. Yuan, J., Bae, E., Tai, X.-C., Boykov, Y.: A continuous max-flow approach to potts model. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *ECCV 2010, Part VI*. LNCS, vol. 6316, pp. 379–392. Springer, Heidelberg (2010)