



BEaST: Brain extraction using multiresolution nonlocal segmentation.

Simon Eskildsen, Pierrick Coupé, Kelvin Leung, Vladimir Fonov, Nicolas Guizard, Shafik Wassef, Lasse Ostergaard, Louis Collins

► To cite this version:

Simon Eskildsen, Pierrick Coupé, Kelvin Leung, Vladimir Fonov, Nicolas Guizard, et al.. BEaST: Brain extraction using multiresolution nonlocal segmentation.. MICCAI Workshop on Multi-Atlas Labeling and Statistical Fusion, Sep 2011, Toronto, Canada. pp.97-108, 2011. <hal-00614307>

HAL Id: hal-00614307

<https://hal.archives-ouvertes.fr/hal-00614307>

Submitted on 11 Aug 2011

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.

BEaST: Brain extraction using multiresolution nonlocal segmentation

Simon F. Eskildsen^{1,2}, Pierrick Coupé¹, Kelvin K. Leung³, Vladimir Fonov¹, Nicolas Guizard¹, Shafik N. Wassef¹, Lasse Riis Østergaard², D. Louis Collins¹, and the Alzheimer’s Disease Neuroimaging Initiative*

¹ McConnell Brain Imaging Centre, Montreal Neurological Institute, McGill University, 3801 University Street, Montreal, Canada

² Department of Health Science and Technology, Aalborg University, Fredrik Bajers Vej 7D, Aalborg, Denmark

³ Dementia Research Centre (DRC), UCL Institute of Neurology, Queens Square, London, UK

Abstract. Brain extraction is an important step in the analysis of brain images. Variability in brain morphology and intensity characteristics due to different imaging sequences makes the development of a general purpose brain extraction algorithm challenging. *Purpose:* To address this issue, we propose a new robust method (BEaST) for brain extraction. *Methods:* The method is based on nonlocal segmentation embedded in a multiresolution framework. A library of 50 priors are semi-automatically constructed from the NIHPD, ICBM, and ADNI databases. *Results:* A mean Dice coefficient of 0.9834 ± 0.0053 is obtained when performing leave-one-out cross validation. Validation using the online available Segmentation Validation Engine resulted in a top ranking position with a mean Dice coefficient of 0.9781 ± 0.0047 . *Conclusions:* The segmentation accuracy of the method is comparable to that of a recent label fusion approach, while being 40 times faster and requiring a much smaller library of priors.

Keywords: Brain extraction, skull stripping, patch-based segmentation, MRI, BET

1 Introduction

Brain extraction (or skull stripping) is an important step in many neuroimaging analyses, such as registration, tissue classification, and segmentation. While methods such as the estimation of intensity normalization fields and registration do not require perfect brain masks, other methods such as measuring cortical thickness rely on accurate brain extraction to work properly. In cases of incorrect brain extraction, subjects may be excluded from further processing, a potentially expensive consequence for most studies. The solution of manually correcting the brain masks is a labor intensive and time consuming task.

An accurate brain extraction method should separate all tissue external to the brain, such as skull, dura, and eyes, from the brain tissue without removing

* Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (www.loni.ucla.edu/ADNI). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://loni.ucla.edu/ADNI/Collaboration/ADNI_Authorship_list.pdf

any part of the brain. For instance, failure to remove the dura may lead to an overestimation of cortical thickness [1], while removing part of the brain would lead to an underestimation.

Several brain extraction methods have been proposed during the last decade [2–7], and studies evaluating these methods have found varying accuracy [8, 9]. While some methods are better at removing non-brain tissue, at the cost of removing brain tissue, others are better at including all brain tissue, at the cost of including non-brain tissue [8, 9] - classical example of sensitivity vs. specificity. A method that works reliably and robustly on a variety of different brain morphologies and acquisition sequences without requiring adjustment of parameters would greatly reduce the need for manual intervention and exclusion of subjects in neuroimaging studies.

Building on recent work on label fusion [10], multi-atlas propagation and segmentation (MAPS) [11] was adapted to brain extraction to address the problem of variability in anatomy and acquisition, producing more robust results [12]. In label fusion, multiple atlases are selected from a library of labeled images and, by means of nonrigid registration, merged in the target image to obtain a segmentation. The method is highly dependent on the accuracy of the nonrigid registrations. Registration errors may result in later segmentation errors, as all selected labels typically are weighted equally. Due to the time consuming multiple nonrigid registrations step in MAPS, the processing time per subject on an Intel[®] Xeon[®] CPU (3GHz) is 19 h.

A recent framework inspired by nonlocal means filtering [13] has been introduced to achieve the label fusion task. This method has demonstrated promising segmentation results without the need for nonrigid registrations [14]. Instead of performing the fusion of nonlinearly deformed template structures, this method achieves the labeling of each voxel individually by comparing its surrounding neighbourhood with patches in training subjects in which the labels of the central voxels are known. The patch-based segmentation method cannot be directly applied to brain extraction, because *i*) false positives are likely to occur as extracerebral tissue may resemble brain within the patch structure, and *ii*) the computational complexity is high and becomes a problem for large structures. To address these issues, we propose to apply a multiresolution approach to the patch based segmentation for performing brain extraction.

2 Methods

The proposed method, **Brain Extraction based on nonlocal Segmentation Technique (BEaST)**, is inspired by the patch-based segmentation published in [14]. As done in [14], we use sum of squared differences (SSD) as the metric for distance between patches. Thus, we perform intensity normalization and spatial normalization before constructing the library of priors. The gold standard of library priors is constructed using a semi-automatic method that involves manual correction of the brain masks, because manual brain segmentation from scratch is a time consuming process.

The following describes the normalization, construction of the library containing the priors, and fundamental patch-based segmentation method as well as our contribution of embedding the method in a multiresolution approach to improve segmentation accuracy and computation time. Our main contribution

is the development of a robust procedure to identify accurate brain masks with an extensive validation on multiple datasets acquired on different scanners.

2.1 Normalization

Image intensity normalization is performed by first applying the bias field correction algorithm N3 [16] followed by a linear intensity normalization, which scales the intensities to the range $[0;100]$ using 0.1%–99.9% of the voxels in the intensity histogram. Spatial normalization is done by 9-parameter linear registration to the publicly available ICBM152 average [17] (stereotaxic space= $193\times 229\times 193$ with isotropic 1 mm voxels).

2.2 Definition of brain mask

Our goal for the brain extraction is to include all grey matter (GM) and white matter (WM) as well as internal CSF and some external CSF. Thus ventricles and deep sulci should be included in the resulting segmentation while all dura, exterior blood vessels, and nerves should be excluded. This definition is similar to results produced by the popular Brain Extraction Toolkit (BET) [4], except we aim at consistently excluding the superior sagittal and transverse sinuses and the optic chiasms, which in most cases are included in BET segmentations. Furthermore, BET sometimes includes dura proximate to the brain, carotid arteries, ocular fat / muscle, and parts of the eyes, while cutting off gyri in cases of atrophied brains. Our aim is to eliminate such errors in a consistent manner.

2.3 Construction of library priors

The library of segmentation priors is built from data in the NIH Pediatric Database (NIHPD) [18] (age: 5–18y), the International Consortium for Brain Mapping (ICBM) database [19] (age: 18–43y), and the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database [20] (age: 55–91y). The NIHPD and ICBM databases consist of healthy subjects, while the ADNI database, in addition to cognitive normal (CN) subjects, contains scans of subjects with Alzheimer’s disease (AD) and mild cognitive impairment (MCI). This way the entire human life span is almost covered and subjects with atrophic morphology are included, which provides a representative library of priors for performing brain extraction.

We chose 10 random T1-weighted (T1w) magnetic resonance (MR) scans from each of the NIHPD and ICBM databases. From the ADNI database we chose 10 random T1w MR scans at the baseline timepoint from each class (CN, MCI, AD). In total, our library consists of 50 priors.

The NIHPD and ICBM databases contain T2w and PDw images in addition to T1w images. T1w images have high signal for the brain tissue, while T2w and PDw images have high signal for cerebro-spinal fluid (CSF) (see Fig. 1). We utilize this in defining the library priors. By adding intensities from the three different sequences we obtain an image with a very high signal for the intracranial cavity (ICC), which can be easily extracted using BET. From the ICC segmentation we use Fast Accurate Cortex Extraction (FACE) [21] to delineate the boundary between GM and CSF in the cerebrum. Cerebellum and brain stem

are added by non-linearly fitting masks in stereotaxic space. Finally, extensive and careful manual corrections were performed to get an optimal brain segmentation matching our definition. On average, such corrections took between 1-2 hours per brain.

Priors from the ADNI database are constructed using the semi-automatic segmentations used in MAPS [12], which are publicly available at the ADNI website. These segmentations are accurate definitions of the GM and WM of the brain, but all interior CSF is excluded (see Fig 2). Therefore, we deform a spherical mesh initialized around the brain to fit smoothly along the border of the segmentation. This way we obtain a similar definition of a brain segmentation as for the NIHPD and ICBM data. Finally, these segmentations were manually corrected in the same way as the NIHPD and ICBM data.

All library priors were flipped along the midsagittal plane to increase the size of the library utilizing the symmetric properties of the human brain.

2.4 Patch-based segmentation

We use the patch-based segmentation method described in [14]. In brief, a label is applied to a given voxel in the target image based on the similarity of its surrounding patch $P(x_i)$ to all the patches $P(x_{s,j})$ in the library. For each voxel x_i of the target image, the surrounding neighbourhood V_i is searched for similar patches in the N library images. A nonlocal means estimator $v(x_i)$ is used to estimate the label at x_i :

$$v(x_i) = \frac{\sum_{s=1}^N \sum_{j \in V_i} w(x_i, x_{s,j}) l(x_{s,j})}{\sum_{s=1}^N \sum_{j \in V_i} w(x_i, x_{s,j})}, \quad (1)$$

where $l(x_{s,j})$ is the label of voxel $x_{s,j}$ at location j in library image s . We used $l(x_{s,j}) \in \{0, 1\}$, where 0 is background and 1 is object (brain). The weight $w(x_i, x_{s,j})$ assigned to label $l(x_{s,j})$ depends on the similarity of $P(x_i)$ to $P(x_{s,j})$ and is computed as:

$$w(x_i, x_{s,j}) = e^{-\frac{\|P(x_i, P(x_{s,j}))\|_2^2}{h^2}}, \quad (2)$$

where $\|\cdot\|_2$ is the L2-norm, normalized by the number of elements and computed between each intensity of the elements of the patches $P(x_i)$ and $P(x_{s,j})$. The parameter h of the weighting function is locally adapted as in [14].

Computation time rises quickly if these calculations are to be made for all patches in all library images. To decrease computation time, the N closest images from the library are selected based on their similarity to the target image within a defined region of interest (ROI) (initialization mask, see eq. 3). The similarity is calculated as the SSD between the target and each library image. Furthermore, to reduce the number of patches to consider, preselection of patches is done as proposed in [14]. Finally, an initialization mask M is constructed as the union of all segmentation priors S_i minus the intersection of all S_i :

$$M = (S_1 \cup S_2 \cup \dots \cup S_N) \setminus (S_1 \cap S_2 \cap \dots \cap S_N). \quad (3)$$

The patch-based segmentation is performed within this ROI only under the assumption that the library is representative of all brain sizes after spatial normalization. This approach reduces the ROI by 50% compared with the union of all S_i and by 85% compared with the entire stereotaxic space (Fig. 3).




Fig. 1. Construction of library priors using multiple modalities. A) Intensities from T1w, T2w, and PDw images are added. B) BET is used to produce an ICC mask. C) FACE is used to delineate the cortical boundary and produce a cerebrum mask. D) Cerebellum and brain stem are added by stereotaxic masks, and the mask is manually corrected.

2.5 Multiresolution framework

In order to obtain optimal performance for brain extraction, the patch size needs to be large compared to the patch sizes used for smaller structures, such as the hippocampus. A large patch size is needed to avoid inclusion of extra-cerebral tissue, such as dura. This is computationally impractical in the native resolution. Therefore, we suggest to embed the patch based segmentation in a multiresolution framework, which provides the opportunity to have large patch sizes while still being computationally practical.

In brief, the multiresolution framework enables propagation of segmentation across scale by using the resulting segmentation at the previous scale to initialize the segmentation at the current one.

The library images, labels, initialization mask, and target image at the stereotaxic resolution V^j are all resampled to a lower resolution V^{j-k} , and the patch-




Fig. 2. Adaptation of library priors using deformable surface. A) Semi-automatic GM/WM mask as used by MAPS [12]. B) Adapted mask generated by deforming a surface mesh to the boundary of the GM/WM mask and manually corrected.

based segmentation is performed. The nonlocal means estimator $v_V^{j-k}(x_i)$ at the V^{j-k} resolution is propagated to a higher resolution V^{j-k+1} by upsampling using trilinear interpolation. The estimator function $v_V^{j-k}(x_i)$ can be considered as the confidence level of which label to assign the voxel. Values close to 0 are likely background, while values close to 1 are likely object. We define a confidence level α to assign labels to the voxels at each scale. Voxels with $v_V^{j-k}(x_i) < \alpha$ are labelled background, and voxels with $v_V^{j-k}(x_i) > (1 - \alpha)$ are labelled object. Segmentation of these voxels is considered final, and they are excluded from further processing. Voxels with $v_V^{j-k}(x_i)$ in the range $[\alpha; 1 - \alpha]$ are propagated and processed at a higher resolution (V^{j-k+1}). This procedure is repeated until the resolution of the stereotaxic space V^j is reached. This way the initialization mask of each resolution step is limited to the voxels with uncertain segmentation at the previous step (Fig. 3). This greatly reduces the computational cost. At the stereotaxic resolution, segmentation is done by thresholding the estimator $v_V^j(x_i)$ at 0.5.

During experiments, we used three resolutions ($k = 2$) with isotropic voxel sizes respectively of 4 mm, 2 mm, and 1 mm (stereotaxic space resolution) (see Fig. 3). We empirically chose confidence level α and variable patch size and search area depending on the resolution (see table 1).

Voxel Size (mm)	Patch Size	Search Area	α
4×4×4	3×3×3	3×3×3	0.2
2×2×2	3×3×3	9×9×9	0.2
1×1×1	5×5×5	13×13×13	-

Table 1. Patch size, search area, and confidence level α chosen for the three resolutions.

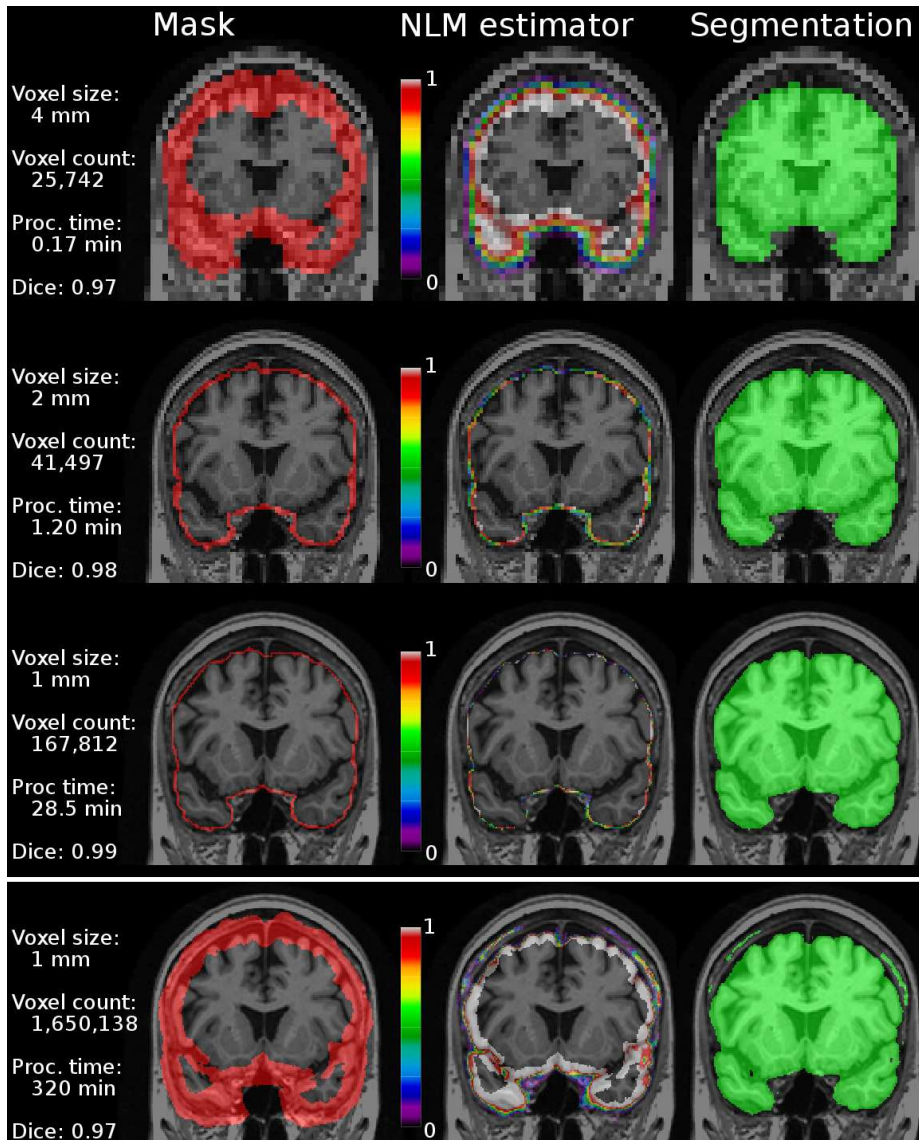


Fig. 3. The multiresolution segmentation process (row 1-3) compared to a single resolution approach (row 4). Column 1: Initialization mask. Column 2: Nonlocal means estimator map. Column 3: Segmentation by thresholding the NLM estimator and adding the intersection mask. Processing times are accumulated time from initialization. Notice the inclusion of dura in the single resolution approach.

3 Results

We measured the segmentation accuracy in a leave-one-out cross validation (LOOCV) fashion. Each of the 50 library images was processed with the re-

maintaining 49 images as priors, and the resulting segmentation was compared to manually corrected labels in the library. Fig. 4 shows the Dice coefficients for increasing number of priors selected from the library. As shown in [14], increasing the number of selected priors improves the segmentation accuracy. In our experiment, accuracy is high even when using only very few selected priors. Increasing the number of selected priors seems to make the segmentations more consistent. We chose $N = 20$ for further experiments. Fig. 6(a) shows the segmentation

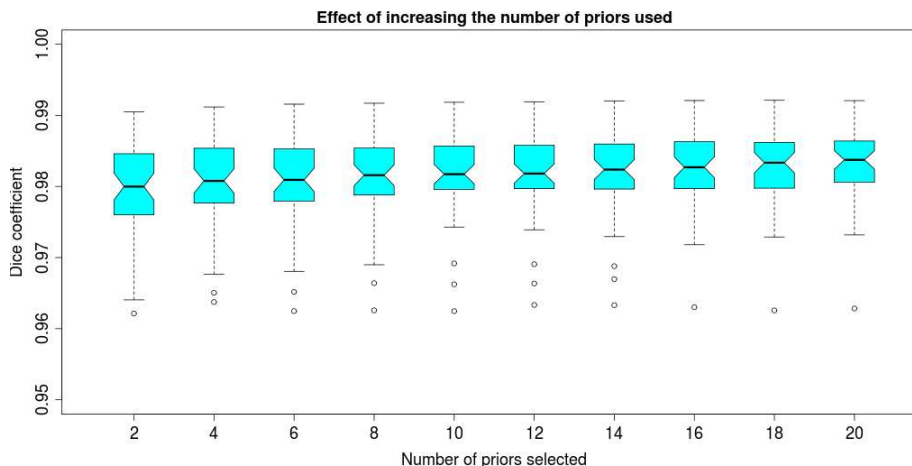


Fig. 4. Dice coefficient of segmentations using an increasing number of priors from the library. Experiment performed by leave-one-out.

accuracy within the different groups used in the experiment. The accuracy on ICBM data is significantly higher than the accuracy on the other groups tested. This may be due to the fact that ICBM data are more homogeneous than the other groups, which lead to higher redundancy and better matches of patches during the segmentation process.

A comparison to BET was performed by running BET with the optimal parameters reported in [12]. Fig. 5 shows typical examples of brains masks obtained by BET and BEaST on the five different groups tested here (NIHPD, ICBM, ADNI-CN, ADNI-MCI, ADNI-AD). On NIHPD and ICBM data, BET behaves quite well with only minor segmentation errors, such as inclusion of the transverse sinus and part of the eye sockets. On ADNI data, more serious errors are found using BET. These include inclusion of dura and cartilage while gyri are often cut off in atrophic brains. BEaST generally provides a more consistent and robust segmentation without serious errors.

We measured the segmentation output at each resolution by thresholding the nonlocal means estimator at 0.5. Fig. 6(b) shows the resulting Dice coefficient compared to BET. As shown, the accuracy increases along with scale, and at 2 mm voxel sizes (requiring about 1.2 minutes) BEaST has higher accuracy than BET. The difference in Dice coefficients may seem small. However, when measuring Dice coefficients in the context of whole brain segmentations, small




Fig. 5. Typical results using BET and BEaST on the five test groups. Column 1: Manual segmentation. Column 2: BET segmentation. Column 3: BEaST segmentation. Blue voxels are overlapping voxels in the segmentation compared to the gold standard. Green voxels are false positives and red voxels are false negatives. Generally, BEaST consistently removes dura, sinuses and optic chiasms, while this is not the case for BET.

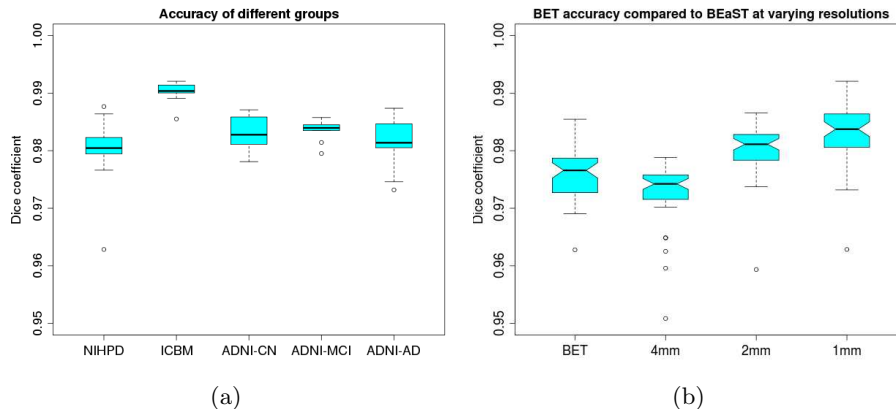


Fig. 6. a) Accuracy of segmentation within groups. b) Segmentation accuracy measured at varying voxel sizes compared to accuracy of BET.

changes in the coefficient corresponds to large changes in volume. In our case a change of 0.01 in Dice coefficient corresponds to about 30-40 cm^3 depending on brain size and the false positives - false negatives ratio. This volume is relatively large when compared to the size of the structures, which are usually measured in neuro-imaging studies. E.g. the size of the human hippocampus is about 3.5 cm^3

We validated BEaST using an independent test set. Validation was performed using the online Segmentation Validation Engine (SVE) of brain segmentation methods [9]. The web service allows the comparison of results with 40 hand-corrected brain masks. The website contains an archive of all uploaded results, which enables segmentation methods to be objectively benchmarked and compared. Images from the independent test set were normalized the same way as the library images. Validation using the test set resulted in a mean Dice coefficient of 0.9781 ± 0.0047 (see <http://sve.loni.ucla.edu/archive/study/?id=244>). At the time of writing, this result was the best of all the methods published on the website. MAPS had a second place with a Dice coefficient of 0.9767 ± 0.0021 . When compared with BEaST, these results are statistically significant (Student's group t-test, $p=0.05$). Furthermore, MAPS requires 19 hours of computation time, compared to 30 minutes for BEaST on similar hardware.

3.1 Computation time

In our experiments with 20 images selected from the library, the total processing time on an Intel[®] i7[®] processor at 3 GHz was less than 30 min per subject. With 10 images, the processing time was less than 20 min per subject. By contrast, without the multiresolution step and using the initialization mask, the processing time was around 320 min. Removing the initialization mask increased the processing time to 42 h.

The average processing time of BET was about 2 minutes. Obtaining the segmentation of BEaST at 2 mm voxel sizes takes about 1.25 minutes and the

corresponding Dice coefficients are higher than BET. This suggests that a fast low resolution result is in many cases sufficient for the subsequent analysis.

4 Discussion

The leave-one-out cross validation showed that the segmentation accuracy is consistently high (average Dice coefficient: 0.9834 ± 0.0053) and that selecting more than 10 priors does not noticeably increase the accuracy. Our results are similar to recent ones from a label fusion approach [12], where the library is almost 14 times larger and the processing time, about 40 times longer. The short processing time in BEaST (<30 min) results from only needing linear registrations and the advantage of using the ROI in the multiresolution strategy. The current implementation runs as a single thread. However, the nonlocal means calculations can easily be parallelized and implemented to exploit GPU processing [22], which will decrease processing time significantly.

Compared to the widely used brain extraction method, BET, BEaST more than halves the segmentation error increasing the average Dice coefficient from 0.9646 to 0.9834. In terms of speed, BET is faster than BEaST, if the segmentations are performed at the highest resolution. However, stopping the processing at 2 mm voxel sizes results in computation times similar to BET, while still obtaining higher segmentation accuracy.

The results showed a higher accuracy on ICBM data. This may be caused by the fact that the anatomical variability within this group is smaller than the other groups studied. This suggests that the accuracy may be improved by extending the number of priors for the groups with higher anatomical variability. Though the results show that only a relatively small library is needed, the library still needs to be representative for the patch-based segmentation to work optimal.

Using the online segmentation validation engine [9] we obtained a truly objective measure of the performance of BEaST. A mean Dice coefficient of 0.9781 is significantly better than the best score by MAPS (0.9767). Compared to MAPS, a relatively small number of library priors are needed in the patch-based segmentation. This makes it feasible to distribute the method as downloadable software. We intend to make BEaST available online if permission to redistribute the data in the library can be obtained.

As in label fusion, the nonlocal segmentation approach enables the segmentation of different targets simultaneously. For example, the intracranial cavity may be obtained by generating priors using appropriate methods, such as the multi-modal approach used as intermediate step to obtain the brain segmentation priors. Also, separation of cerebellum and brain stem from the cerebrum may be achieved with high accuracy if the structural priors are available.

In conclusion, we have proposed a new brain extraction method, BEaST, based on nonlocal segmentation embedded in a multiresolution framework. The accuracy of the method is higher than BET and similar to that of a recent label fusion method, while being much faster and requiring a smaller library of priors.

Acknowledgements

The authors would like to thank Professor Nick Fox, Dementia Research Centre, Institute of Neurology, London, for contributing with the ADNI semi-automatic

brain segmentations. KKL acknowledges support from the MRC, ARUK and the NIHR.

References

1. van der Kouwe et al.: Brain morphometry with multiecho MPRAGE. *NeuroImage*, 40(2), 559-569 (2008)
2. Lemieux et al.: Fast, accurate, and reproducible automatic segmentation of the brain in T1-weighted volume MRI data. *Magnet. Reson. Med.* 42(1), 127-135 (1999).
3. Shattuck et al.: Magnetic resonance image tissue classification using a partial volume model. *NeuroImage* 13(5), 856-876 (2001).
4. Smith: Fast Robust Automated Brain Extraction. *Hum. Brain Mapp.* 17(3), 143-155 (2002)
5. Ségonne et al.: A hybrid approach to the skull stripping problem in MRI. *NeuroImage* 22(3), 1060-1075 (2004)
6. Zhuang et al.: Skull-stripping magnetic resonance brain images using a model-based level set. *NeuroImage* 32(1), 79-92 (2006)
7. Sadananthan et al.: Skull stripping using graph cuts. *NeuroImage* 49(1), 225-239 (2010).
8. Fennema-Notestine et al.: Quantitative evaluation of automated skull-stripping methods applied to contemporary and legacy images: Effects of diagnosis, bias correction, and slice location. *Hum. Brain Mapp.* 27(2), 99-113 (2006)
9. Shattuck et al.: Online resource for validation of brain segmentation methods. *NeuroImage* 45(2), 431-439 (2009)
10. Collins and Pruessner: Towards accurate, automatic segmentation of the hippocampus and amygdala from MRI by augmenting ANIMAL with a template library and label fusion, *NeuroImage* 52(4), 1355-1366 (2010)
11. Leung et al.: Automated cross-sectional and longitudinal hippocampal volume measurement in mild cognitive impairment and Alzheimer's disease, *NeuroImage* 51(4), 1345-1359 (2010)
12. Leung et al.: Brain MAPS: An automated, accurate and robust brain extraction technique using a template library. *NeuroImage* (in press)
13. Coupé et al.: An Optimized Blockwise Nonlocal Means Denoising Filter for 3-D Magnetic Resonance Images, *IEEE T. Med. Imaging* 27(4), 425-441 (2008)
14. Coupé et al.: Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage* 54(2), 940-954 (2011)
15. Dice: Measures of the Amount of Ecologic Association Between Species. *Ecology* 26, 297302 (1945)
16. Sled et al.: A nonparametric method for automatic correction of intensity nonuniformity in mri data. *IEEE T. Med. Imaging* 17(1), 87-97 (1998)
17. Fonov et al., the Brain Development Cooperative Group: Unbiased average age-appropriate atlases for pediatric studies, *NeuroImage* 54(1), 313-327 (2011)
18. Evans: The NIH MRI study of normal brain development, *NeuroImage* 30(1), 184-202 (2006)
19. Mazziotta et al.: A Probabilistic Atlas of the Human Brain: Theory and Rationale for Its Development : The International Consortium for Brain Mapping (ICBM), *NeuroImage* 2(2), 89-101 (1995)
20. Mueller et al.: The Alzheimer's Disease Neuroimaging Initiative, *Neuroimaging Clin N Am* 15(4), 869-877 (2005)
21. Eskildsen and Østergaard: Active Surface Approach for Extraction of the Human Cerebral Cortex from MRI, *LNCS* 4191, 823-830 (2006)
22. de Fontes et al.: Real time ultrasound image denoising. *J Real-Time Image Proc* 6(1), 1522 (2011)