

## **BET: Brain Extraction Tool**

**Stephen M. Smith**

FMRIB technical Report TR00SMS2b

(A related paper has been accepted for publication in Human Brain Mapping)

FMRIB (Oxford Centre for Functional Magnetic Resonance Imaging of the Brain)

Department of Clinical Neurology, Oxford University, John Radcliffe Hospital,

Headington, Oxford OX3 9DU, UK

steve@fmrib.ox.ac.uk [www.fmrib.ox.ac.uk/~steve](http://www.fmrib.ox.ac.uk/~steve)

phone: +44 1865 222726 fax: +44 1865 222717

**Abstract** An automated method for segmenting MR head images into brain and non-brain has been developed. It is very robust and accurate and has been tested on thousands of data sets from a wide variety of scanners and taken with a wide variety of MR sequences. The method - BET (Brain Extraction Tool) - uses a deformable model which evolves to fit the brain's surface by the application of a set of locally adaptive model forces. The method is very fast and requires no pre-registration or other pre-processing before being applied. This report describes the new method and gives some example results and also the results of extensive quantitative testing against "gold-standard" hand segmentations and two other popular automated methods.

**Keywords:** Brain segmentation, cortical surface modelling

## **1 Introduction**

There are many applications related to brain imaging which either require, or benefit from, the ability to accurately segment brain from non-brain tissue. For example, in the registration of functional images to high resolution MR images, both FMRI and PET functional images often contain little non-brain tissue because of the nature of the imaging, whereas the high resolution MR image probably will contain a considerable amount - eyeballs, skin, fat, muscle, etc - and thus registration robustness is improved if these non-brain parts of the image can be automatically removed before registration. A second example application of brain/non-brain segmentation is as the first stage in cortical flattening procedures. A third example is in brain atrophy estimation in diseased subjects; after brain/non-brain segmentation, brain volume is measured at a single time point with respect to some normalising

volume such as skull or head size; alternatively, images from two or more time points are compared, to estimate how the brain has changed over time [13, 14]. Note that in this application, tissue-type segmentation is also used to help disambiguate brain tissue from other parts of the image such as CSF [16]. A fourth application is the removal of strong ghosting effects which can occur in functional MRI (eg with EPI - echo planar imaging). These artefacts can confound motion correction, global intensity normalisation and registration to a high resolution image. They can have an intensity as high as the “true” brain image, preventing the use of simple thresholding to eliminate the artefacts, whereas a geometric approach such as that presented here can remove the effects (though only from outside of the brain).

This report describes a complete method for achieving automated brain/non-brain segmentation. The method described here does not attempt to model the brain surface at the finest level, for example, following sulci and gyri, or separating cortex from cerebellum. This finer modelling would be a later stage, if required, after the brain/non-brain segmentation.

After a brief review of brain extraction, the brain extraction algorithm is described in detail, followed by a description of an addition which attempts to find the exterior surface of the skull. Finally, example qualitative results are presented, followed by the results of extensive quantitative evaluation against 45 “gold-standard” hand segmentations and comparisons, using this data, with two other popular automated methods.

## **2 Review**

To date, there have been three main methods proposed for achieving brain/non-brain segmentation; manual, thresholding-with-morphology and surface-model-based. In this review these will be briefly described and compared.

The problem of brain/non-brain segmentation is a subset of structural segmentation, which aims, for example, to segment the major brain structures such as cerebellum, cortex and ventricles. It is an image processing problem where a semi-global understanding of the image is required as well as a local understanding. This is often more difficult than situations where purely local or purely global solutions are appropriate. For an example of the difference between local and semi-global operations, take the finding of “corners” in images. In clean images with clean sharp corners, a good solution may be found by applying small locally-acting op-

erators to the image. However, in the presence of large amounts of noise, or if it is required to find less sharp corners, a larger-scale view must be taken - for example, two edges must be defined over a larger area, and their position of intersection found.

Manual brain/non-brain segmentation methods are, as a result of the complex information understanding involved, probably more *accurate* than fully automated methods are ever likely to achieve. This is the level in image processing where this is most true. At the lowest, most localised, level (for example, noise reduction or tissue-type segmentation), humans often cannot improve on the numerical accuracy and objectivity of a computational approach. The same also often holds at the highest, most global, level; for example, in image registration, humans cannot in general take in enough of the whole-image information to improve on the overall fit that a good registration program can achieve. However, with brain segmentation, the appropriate size of the image “neighbourhood” which is considered when outlining the brain surface is ideally suited to manual processing. For example, when following the external contours of gyri, differentiating between cerebellum and neighbouring veins, cutting out optic nerves, or taking into account unusual pathology, semi-global contextual information is crucial in helping the human optimally identify the correct brain surface.

Of course, there are serious enough problems with manual segmentation to prevent it from being a viable solution in most applications. The first is time cost - manual brain/non-brain segmentation typically takes between 15 minutes and 2 hours per 3D volume. The second is the requirement for sufficient training, and care during segmentation, that subjectivity is reduced to an acceptable level. For example, even a clinical researcher who has not been explicitly trained will be likely to make a mistake in the differentiation between lower cerebellum and neighbouring veins.

The second class of brain segmentation methods is thresholding-with-morphology, e.g. [6]. An initial segmentation into foreground/background is achieved using simple intensity thresholding. Lower and upper thresholds are determined which aim to separate the image into very bright parts (e.g. eyeballs and parts of the scalp), less bright parts (e.g. brain tissue), and the dark parts (including air and skull). Thus a binary image is produced. In the simplest cases, the brain can now be determined by finding the largest single contiguous non-background cluster. A binary brain mask then results; this can be applied to the original image. However, the brain cluster is almost always connected, often via fairly thin strands of bright voxels, to non-brain tissue such as the eyeballs or scalp. For example, this “bridge” can be caused either by the optic nerve, or simply at points around the brain where the dark skull gap is very narrow. Thus before the largest single cluster is used, it

must be disconnected from the non-brain bright tissue. This is normally achieved by morphological filtering; the bright regions in the binary image are eroded away until any links between brain and non-brain are eliminated, the largest single cluster is then chosen, and this is then dilated by the same extent as the erosion, hopefully resulting in an accurate brain mask.

Thresholding-with-morphology methods are mostly only semi-automated - the user is normally involved in helping choose the threshold(s) used in the initial segmentation. It is often necessary to try the full algorithm out with a variety of starting thresholds until a good output is achieved. A second problem is that it is very hard to produce a general algorithm for the morphology stage that will successfully separate brain from non-brain tissue; it has proved difficult to automatically cope with a range of MR sequences and resolutions. In general, results need some final hand editing. In part, this is due to the fact that it is hard to implement situation-specific logical constraints (e.g., prior knowledge about head images) with this approach.

A more sophisticated version of the above approach is given in [11]. Here a series of thresholding and morphology steps are applied, with each step carefully tuned to overcome specific problems, such as the thin strands joining brain to non-brain after thresholding. Whilst the results presented are impressive, this method is highly tuned to a narrow range of image sequence types. A second related example is presented in [12]. Here edge detection is used instead of thresholding, to separate different image regions. Next, morphology is used to process these regions, in order to separate the large region associated with the brain from non-brain regions. The resulting algorithm can therefore be more robust than some of the thresholding-with-morphology methods; this method (BSE) is used in the quantitative testing presented below. A third example is that implemented in AFNI [4, 15]. Here a Gaussian mixture model across the different image tissue types is fitted to the intensity histogram in order to estimate thresholds for the following slice-by-slice segmentation. This is followed by a surface-model-based surface smoothing, and finally with morphological “cleaning-up”. Again, this technique is used in the quantitative testing presented below. Yet another example is [1], where head/non-head segmentation is first performed, using thresholding and morphology. Next, anisotropic diffusion is applied, to reduce noise and “darken” some non-brain regions, followed by further thresholding and morphology, along with a heuristic method for identifying and removing the eyes. The final surface is modelled with a “snake” [8]. Further examples can be found in [2, 3, 10].

The third class of methods uses deformable surface models; for example, see [5, 9]. Here a surface model is defined - for example, a tessellated mesh of triangles. This model is then “fitted” to the brain surface in the image. Normally there are two

main constraints to the fitting - a part which enforces some form of smoothness on the surface (both to keep the surface well-conditioned and to match the physical smoothness of the actual brain surface) and a part which fits the model to the correct part of the image - in this case, the brain surface. The fitting is usually achieved by iteratively deforming the surface from its starting position until an optimal solution is found. This type of method has the advantages that it is relatively easy to impose physically-based constraints on the surface, and that the surface model achieves integration of information from a relatively large neighbourhood around any particular point of interest; this is therefore using semi-global processing, as described above. In general this kind of approach seems to be more robust, and easier to successfully automate, than the thresholding-with-morphology methods.

### **3 Method Detail**

#### **3.1 Overview of the Brain Extraction Method**

We start with a brief overview of the new method. Firstly, the intensity histogram is processed to find “robust” lower and upper intensity values for the image, and a rough brain/non-brain threshold. The centre-of-gravity of the head image is found, along with the rough size of the head in the image. Next a triangular tessellation of a sphere’s surface is initialised inside the brain, and allowed to slowly deform, one vertex at a time, following forces that keep the surface well-spaced and smooth, whilst attempting to move towards the brain’s edge. If a suitably clean solution is not arrived at then this process is re-run with a higher smoothness constraint. Finally, if required, the outer surface of the skull is estimated. A graphical overview is shown in Figure 1.

#### **3.2 Estimation of Basic Image and Brain Parameters**

The first processing that is carried out is the estimation of a few simple image parameters, to be used at various stages in subsequent analysis.

Firstly, the robust image intensity minimum and maximum are found. Here robust means the effective intensity extrema, calculated ignoring small numbers of voxels which have widely different values from the rest of the image. These are calculated by looking at the intensity histogram, and ignoring long low tails at each end. Thus, the intensity “minimum”, referred to as  $t_2$  is the intensity below which lies 2% of

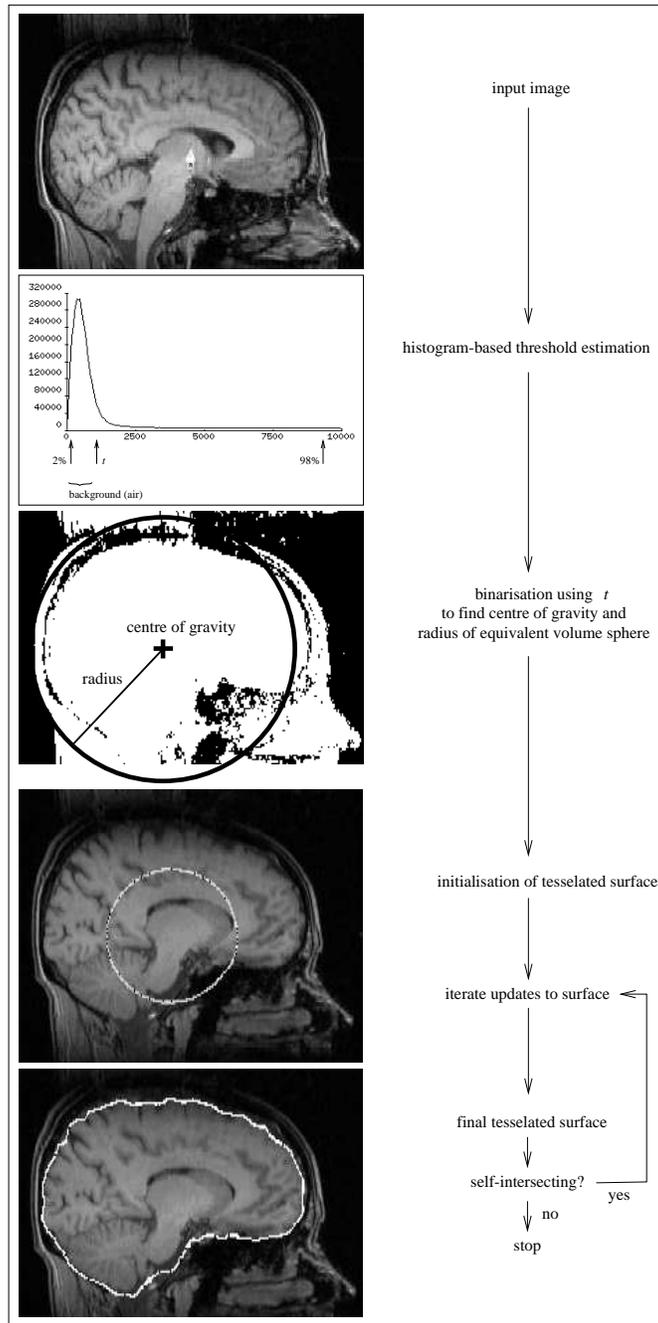


Figure 1: BET processing flowchart.

the cumulative histogram. Similarly,  $t_{98}$  is found. It is often important for the latter threshold to be calculated robustly, as it is quite common for brain images to contain a few high intensity “outlier” voxels; for example, the DC spike from image reconstruction, or arteries, which often appear much brighter than the rest of the image. Finally, a roughly chosen threshold is calculated which attempts to distinguish between brain matter and background. (Note that because bone appears dark in most MRI images, “background” is taken to include bone.) This  $t$  is simply set to lie 10% of the way between  $t_2$  and  $t_{98}$ .

The brain/background threshold  $t$  is used to roughly estimate the position of the centre of gravity (COG) of the brain/head in the image. For all voxels with intensity greater than  $t$ , their intensity (“mass”) is used in a standard weighted sum of positions. Intensity values are upper limited at  $t_{98}$ , so that extremely bright voxels do not skew the position of the COG.

Next the mean “radius” of the brain/head in the image is estimated. There is no distinction made here between estimating the radius of the brain and the head - this estimate is very rough, and simply used to get an idea of the size of the brain in the image; it is used for initialising the brain surface model. All voxels with intensity greater than  $t$  are counted, and a radius is found, taking into account voxel volume, assuming a spherical brain.

Finally, the median intensity of all points within a sphere of the estimated radius and centred on the estimated COG is found -  $t_m$ .

### 3.3 Surface Model and Initialisation

The brain surface is modelled by a surface tessellation using connected triangles. The initial model is a tessellated sphere, generated by starting with an icosahedron and iteratively subdividing each triangle into 4 smaller triangles, whilst adjusting each vertex’s distance from the centre to form as spherical a surface as possible. This is a common tessellation of the sphere. Each vertex has five or six neighbours, according to its position relative to the original icosahedron.

The spherical tessellated surface is initially centered on the COG, with its radius set to half of the estimated brain/head radius, i.e. intentionally small. Allowing the surface to grow to the optimal estimate gives better results in general than setting the initial size to be equal to (or larger than) the estimated brain size (see Figure 7). An example final surface mesh can be seen in Figure 2.

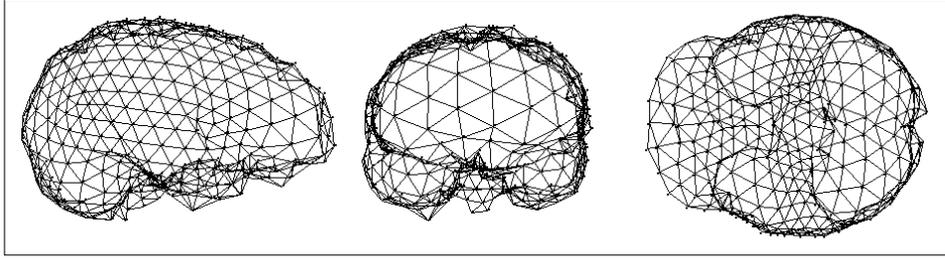


Figure 2: Three views of a typical surface mesh, shown for clarity with reduced mesh density.

The vertex positions are in real (floating point) space, i.e. not constrained to the voxel grid points. A major reason for this is that making incremental (small) adjustments to vertex positions would not be possible otherwise. Another obvious advantage is that the image does not need to be pre-processed to be made up of cubic voxels.

### 3.4 Main Iterated Loop

Each vertex in the tessellated surface is updated by estimating where best that vertex should move to, in order to improve the surface. In order to find an optimal solution, each individual movement is small, with many (typically 1000) iterations of each complete incremental surface update. In this context, “small movement” means small relative to the mean distance between neighbouring vertices. Thus for each vertex, a small update movement vector  $\mathbf{u}$  is calculated, using the following steps.

#### 3.4.1 Local Surface Normal

Firstly the local unit vector surface normal  $\hat{\mathbf{n}}$  is found. Each consecutive pair of [central vertex]-[neighbour A], [central vertex]-[neighbour B] vectors is taken and used to form the vector product (see Figure 3). The vector sum of these vectors is scaled to unit length to create  $\hat{\mathbf{n}}$ . By initially taking the sum of normal vectors *before* rescaling to unity, the sum is made relatively robust; the smaller a particular [central vertex]-[neighbour A]-[neighbour B] triangle is, the more poorly conditioned is the estimate of normal direction, but this normal will contribute less towards the sum of normals.

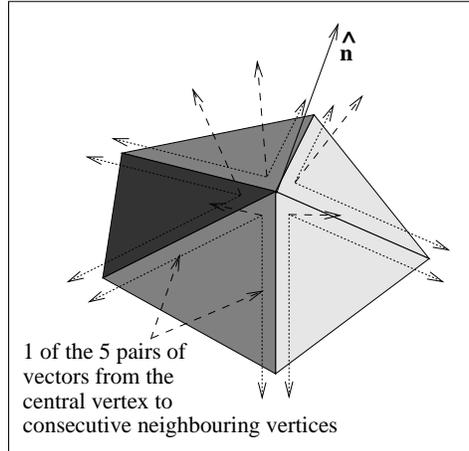


Figure 3: Creating local unit vector surface normal  $\hat{\mathbf{n}}$  from all neighbouring vertices.

### 3.4.2 Mean Position of Neighbours and Difference Vector

The next step is the calculation of the mean position of all vertices neighbouring the vertex in question. This is used to find a difference vector  $\mathbf{s}$ , the vector that takes the current vertex to the mean position of its neighbours. If this vector were minimised for all vertices (by positional updates), the surface would be forced to be smooth and all vertices would be equally spaced. Also, due to the fact that the surface is closed, the surface would gradually shrink.

Next,  $\mathbf{s}$  is decomposed into orthogonal components, normal and tangential to the local surface;

$$\mathbf{s}_n = (\mathbf{s} \cdot \hat{\mathbf{n}}) \hat{\mathbf{n}} \quad (1)$$

and

$$\mathbf{s}_t = \mathbf{s} - \mathbf{s}_n. \quad (2)$$

For the 2D case, see Figure 4 (the extension to 3D is conceptually trivial). It is these two orthogonal vectors which form the basis for the three components of the vertex's movement vector  $\mathbf{u}$ ; these components will be combined, with relative weightings, to create an update vector  $\mathbf{u}$  for every vertex in the surface. The three components of  $\mathbf{u}$  are now described.

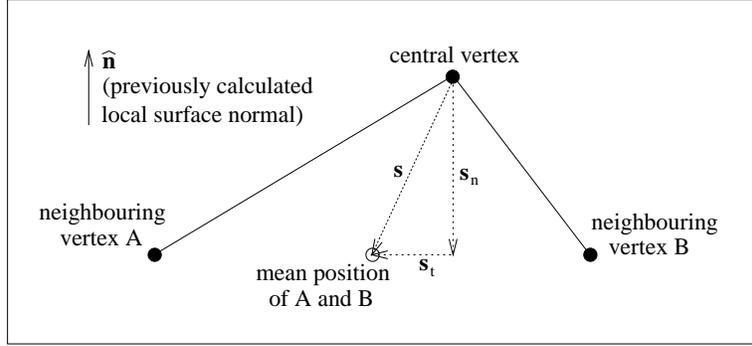


Figure 4: Decomposing the “perfect smoothness” vector  $s$  into components normal and tangential to the local surface.

### 3.4.3 Update Component 1: Within-Surface Vertex Spacing

The simplest component of the update movement vector  $\mathbf{u}$  is  $\mathbf{u}_1$ , the component that is tangential to the local surface. Its sole role is to keep all vertices in the surface equally spaced, moving them only *within* the surface. Thus  $\mathbf{u}_1$  is directly derived from  $s_t$ . In order to give simple stability to the update algorithm,  $\mathbf{u}_1$  is not set equal to  $s_t$ , but to  $s_t/2$ ; the current vertex is always tending towards the position of perfect within-surface spacing (as are all others).

### 3.4.4 Update Component 2: Surface Smoothness Control

The remaining two components of  $\mathbf{u}$  act parallel to the local surface normal. The first,  $\mathbf{u}_2$ , is derived directly from  $s_n$ , and acts to move the current vertex into line with its neighbours, thus increasing the smoothness of the surface. A simple rule here would be to take a constant fraction of  $s_n$ , in a manner equivalent to that of the previous component  $\mathbf{u}_1$ :

$$\mathbf{u}_2 = f_2 s_n, \quad (3)$$

where  $f_2$  is the fractional update constant. Most other methods of surface modelling have taken this approach. However, a great improvement can be made by using a nonlinear function of  $s_n$ . The primary aim is to smooth high curvature in the surface model more heavily than low curvature. The reason for this is that whilst high curvature is undesirable in the brain surface model, forcing surface smoothing to an extent which gives stable and good results (in removing high curvature) weights too heavily against successful following of the low curvature parts

of the surface. In other words, in order to keep the surface model sufficiently smooth for the overall algorithm to proceed stably, the surface is forced to be over-smooth, causing the underestimation of curvature at certain parts, i.e., the “cutting of corners”. It has been found that this problem is not overcome by allowing  $f_2$  to vary during the series of iterations (this a natural improvement on a constant update fraction). Instead, a nonlinear function is used, starting by finding the local radius of curvature,  $r$ :

$$r = \frac{l^2}{2|s_n|}, \quad (4)$$

where  $l$  is the mean distance from vertex to neighbouring vertex across the whole surface - see Figure 5. Now, a sigmoid function of  $r$  is applied, to find the update

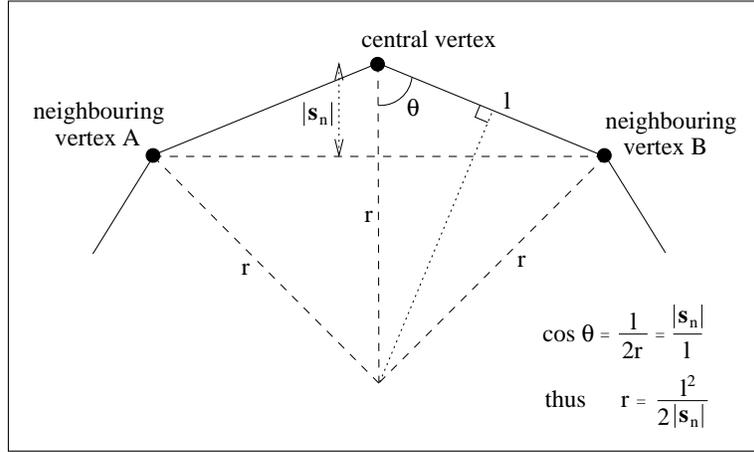


Figure 5: The relationship between local curvature  $r$ , vertex spacing  $l$  and the perpendicular component of the difference vector,  $|s_n|$ .

fraction:

$$f_2 = (1 + \tanh(F * (1/r - E)))/2, \quad (5)$$

where  $E$  and  $F$  control the scale and offset of the sigmoid. These are derived from a minimum and maximum radius of curvature; below the minimum  $r$ , heavy smoothing takes place (i.e., the surface deformation remains stable and highly curved features are smoothed), whilst above the maximum  $r$ , little surface smoothing is carried out (i.e., “long slow” curves are not over-smoothed). The empirically optimised values for  $r_{min}$  and  $r_{max}$  are suited for typical geometries found in the human brain. Consideration of the  $\tanh$  function suggests:

$$E = (1/r_{min} + 1/r_{max})/2, \quad (6)$$

and

$$F = 6/(1/r_{min} - 1/r_{max}). \quad (7)$$

For example, see Figure 6. The resulting smoothness term gives much better results

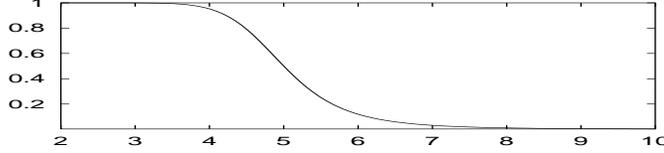


Figure 6: Smoothness update fraction vs local radius of curvature, given  $r_{max} = 10\text{mm}$ ,  $r_{min} = 3.33\text{mm}$

than a constant update fraction, both in ability to accurately model brain surface and in developmental stability during the many iterations.

### 3.4.5 Update Component 3: Brain Surface Selection Term

The final update component,  $\mathbf{u}_3$ , is also parallel to  $\mathbf{s}_n$ , and is the term which actually interacts with the image, attempting to force the surface model to fit to the real brain surface. This term was originally inspired by the intensity term in [5]:

$$\hat{\mathbf{n}} \prod_{d=1}^{30} \max(0, \tanh(I(\mathbf{x} - d\mathbf{n}) - I_{thresh})), \quad (8)$$

where the limits on  $d$  control a search amongst all image points  $\mathbf{x} - d\mathbf{n}$  along the surface normal pointing inwards from the current vertex at  $\mathbf{x}$ , and taking the product requires all intensities to be above a preset threshold. Thus whilst the surface lies within the brain, the resulting force is outwards. As soon as the surface moves outside of the brain (e.g., into CSF or bone), one or more elements inside the product become zero and the product becomes zero. One limitation of this equation is that it can only push outwards - thus the resulting surface is forced to be convex. A second limitation is the use of a single global intensity threshold  $I_{thresh}$ ; ideally, this should be optimally varied over the image.

Thus, instead of the above equation, a much simpler core equation is used, embodying the same idea, but this is then extended to give greater robustness in a wider range of imaging sequences. Firstly, along a line pointing inwards from the current vertex, minimum and maximum intensities are found:

$$I_{min} = \text{MAX}(t_2, \text{MIN}(t_m, I(0), I(1), \dots, I(d_1))), \quad (9)$$

$$I_{max} = MIN(t_m, MAX(t, I(0), I(1), \dots, I(d_2))), \quad (10)$$

where  $d_1$  determines how far into the brain the minimum intensity is searched for, and  $d_2$  determines how far into the brain the maximum intensity is searched for. Typically,  $d_1 = 20mm$  and  $d_2 = d_1/2$  (this ratio is empirically optimised, and reflects the relatively larger spatial reliability of the search for maximum intensity compared with the minimum).  $t_m$ ,  $t_2$  and  $t$  are used to limit the effect of very dark or very bright voxels. Note that the image positions where intensities are measured are in general between voxels as we are working in real (floating point) space - thus intensity interpolation needs to be used, interpolating between original voxel intensities. It was found that nearest neighbour interpolation gave better results than trilinear or higher order interpolations, presumably because it is more important to have access to the original (un-interpolated, and therefore “unblurred”) intensities than that the values reflect optimal estimates of intensities at the correct point in space.

Now,  $I_{max}$  is used to create  $t_l$ , a locally appropriate intensity threshold which distinguishes between brain and background:

$$t_l = (I_{max} - t_2) * b_t + t_2. \quad (11)$$

It lies a preset fraction of the way between the global robust low intensity threshold  $t_2$  and the local maximum intensity  $I_{max}$ , according to fractional constant  $b_t$ . This preset constant is the main parameter which BET can take as input. The default value of 0.5 has been found to give excellent results for most input images - with certain image intensity distributions it can be varied (in the range 0 to 1) to give optimal results. The necessity for this is rare, and for an MR sequence which requires changing  $b_t$ , one value normally works for all other images taken with the same sequence. (The only other input parameter, and one which needs changing from the default even less often than  $b_t$ , for example, if there is very strong vertical intensity inhomogeneity, causes  $b_t$  to vary linearly with  $Z$  in the image, causing “tighter” brain estimation at the top of the brain, and “looser” estimation at the bottom, or vice versa.) The update “fraction” is then given by:

$$f_3 = \frac{2(I_{min} - t_l)}{I_{max} - t_2}, \quad (12)$$

with the factor of 2 causing a range in values of  $f_3$  of roughly -1 to 1. If  $I_{min}$  is lower than local threshold  $t_l$ ,  $f_3$  is negative, causing the surface to move inwards at the current point. If it is higher, then the surface move outwards.

Finally, the full update term is  $0.05f_3 l$ ; the update fraction is multiplied by a relative weighting constant, 0.05, and the mean inter-vertex distance,  $l$ . The weighting

constant sets the balance between the smoothness term and the intensity-based term - it is found empirically, but because all terms in BET are invariant to changes in image voxel size, image contrast, mesh density, etc., this constant is not a “worked-once” heuristic - it is always appropriate.

### 3.4.6 Final Update Equation

Thus the total update equation, for each vertex, is

$$\mathbf{u} = 0.5\mathbf{s}_t + f_2\mathbf{s}_n + 0.05f_3 l \hat{\mathbf{s}}_n. \quad (13)$$

### 3.4.7 Second Pass - Increased Smoothing

One obvious constraint on the brain surface model is that it should not self-intersect. Although it would be straightforward to force this constraint by adding an appropriate term to the update equation, in practice this check is extremely computationally expensive as it involves comparing the position of each vertex with that of every other at every iteration. As it stands, the algorithm already described rarely (around 5% of images) results in self-intersection. A more feasible alternative is to run the standard algorithm and then perform a check for self-intersection. If the surface is found to self-intersect, the algorithm is re-run, with much higher smoothness constraint (applied to concave parts of the surface only - it is not necessary for the convex parts) for the first 75% of the iterations; the smoothness weighting then linearly drops down to the original level over the remaining iterations. This results in preventing self-intersection in almost all cases.

It has been suggested that there might be some value in re-running BET on its own output; whilst areas incorrectly “left in” after a first run might get removed on subsequent runs, it is our experience that this is not in general successful, presumably because the overall algorithm is not designed for this application.

## 3.5 Exterior Skull Surface Estimation

A few applications require the estimation of the position of the skull in the image. A major example is in the measurement of brain atrophy [13]. Before brain change can be measured, two images of the brain, taken several months apart, have to be registered. Clearly this registration cannot allow rescaling, otherwise the overall

atrophy will be underestimated. However, because of possible changes in scanner geometry over time, it is necessary to hold the scale constant somehow. This can be achieved by using the exterior skull surface, which is assumed to be relatively constant in size, as a scaling constraint in the registration.

In most MR images, the skull appears very dark. In T1-weighted images the internal surface of the skull is largely indistinguishable from the cerebro-spinal-fluid (CSF), which is also dark. Thus the exterior surface is found. This also can be difficult to identify, even for trained clinical experts, but the algorithm is largely successful in its aim.

For each voxel lying on the brain surface found by BET, a line perpendicular to the local surface, pointing outward, is searched for the exterior surface of the skull, according to the following algorithm:

- Search outwards from the brain surface, a distance of 30mm, recording the maximum intensity and its position, and the minimum intensity.
- If the maximum intensity is not higher than  $t$ , assume that the skull is not measurable at this point, as there is no bright signal (scalp) on the far side of the skull. In this case do not proceed with the search at this point in the image. This would normally be due to signal loss at an image extreme, for example, at the top of the head.
- Find the point at greatest distance from brain surface  $d$  which has low intensity, according to maximisation of the term  $d/30 - I(d)/(t_{98} - t_2)$ . The first part weights in favour of increased distance, the second part weights in favour of low intensity. The search is continued only out to the previously found position of maximum intensity. The resulting point should be close to the exterior surface of the skull.
- Finally, search outwards from the previous point until the first maximum in intensity gradient is found. This is the estimated position of the exterior surface of the skull. This final stage gives a more well-defined position for the surface - it does not depend on the weightings in the maximised term in the previous section, i.e., is more objective. For example, if the skull/scalp boundary is at all blurred, the final position will be less affected than the previous stage.

This method is quite successful, even when fairly dark muscle lies between the skull and the brighter skin and fat. It is also mainly successful in ignoring the marrow within the bone, which sometimes is quite bright.

## 4 Results

### 4.1 Example Results

Figure 7 shows an example of surface model development as the main loop iterates, with a T1-weighted image as input, finishing with the estimation shown in Figures 8 and 9. Figures 10-12 show example results on T2-weighted, proton density and EPI (echo planar imaging, widely used for FMRI) images. Figure 13 shows an example estimate of the exterior skull surface.

Figure 14 shows the result of running BET on an EPI FMRI image which is heavily affected by ghosting. Clearly BET has worked well, both in removing the (outside-brain) ghosting, and also in allowing registration (using FLIRT [7]) to succeed.

### 4.2 Quantitative Testing Against “Gold-Standard” and Other Methods

An extensive quantitative and objective test of BET has been carried out. We used 45 MR images, taken from 15 different scanners (mostly 1.5T and some 3T, from 6 different manufacturers), using a wide range of slice thicknesses (between 0.8 and 6mm) and a variety of sequences (35 T1-weighted, 6 T2-weighted and 4 proton density). Hand segmentation of these images into brain/non-brain<sup>1</sup> was carried out. Thus a simple binary mask was generated from each input image. Some slices from an example hand segmentation are shown in the second column of Figure 15. Then, BET and two other popular automated methods (“AFNI” and “BSE”) were tested against the hand segmentations.<sup>2</sup>

The AFNI method [4, 15], though claiming to be fully automated, gave very poor results on most images (way off the scale on Figure 16), due to the failure of the initial histogram-based choice of thresholds. Much better results were obtained by setting the initial thresholds using the simpler but more robust method described

---

<sup>1</sup>We defined cerebellum and internal CSF as “brain” - i.e. matching the definitions used by the methods tested. Structures/tissues such as sagittal sinus, optic nerves, external CSF and dura are all ideally eliminated by all the methods (as can be confirmed by their results on “ideal” input images), and also by the hand segmentation.

<sup>2</sup>The versions of these algorithms were: BET version 1.1 from FSL version 1.3; BSE version 2.09; AFNI version 2.29 (with modifications described in the text). All are easily accessible on the internet; to the best of our knowledge, these are the only freely available widely used standalone brain/non-brain segmentation algorithms.

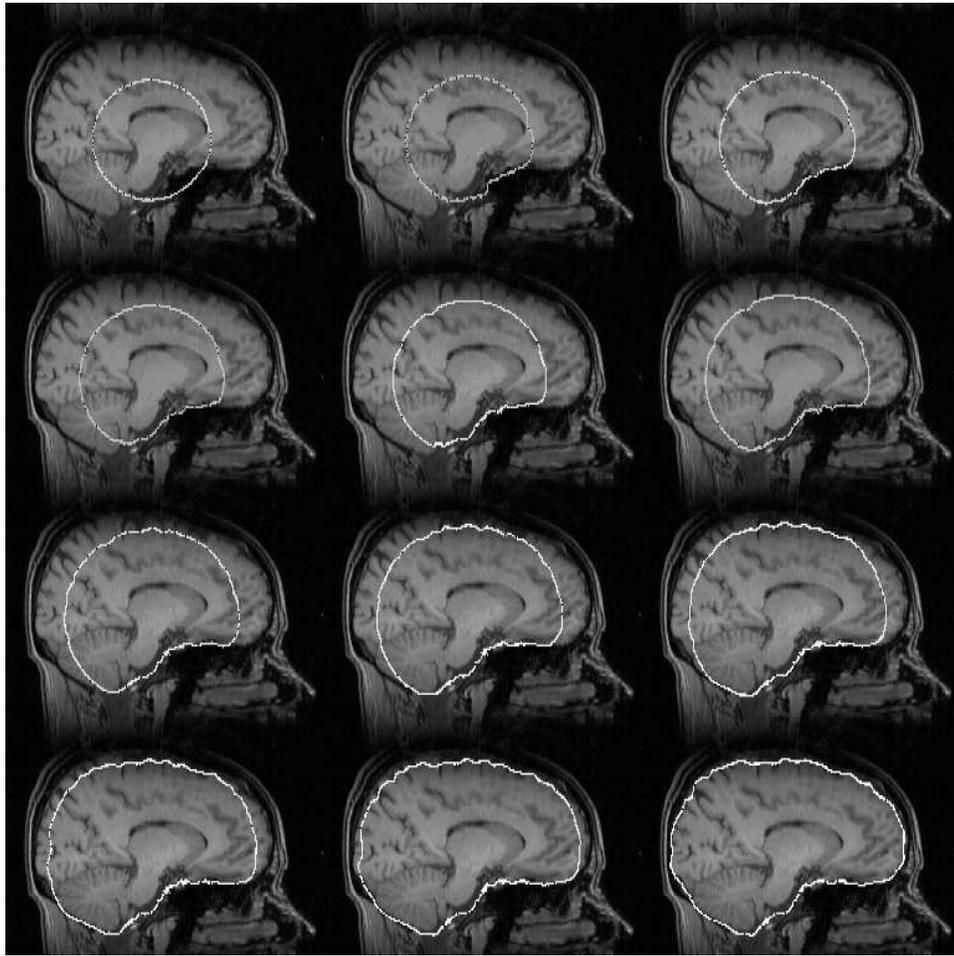


Figure 7: Example of surface model development as the main loop iterates. The dark points within the model outline are vertices.



Figure 8: Example brain surface generated by BET.

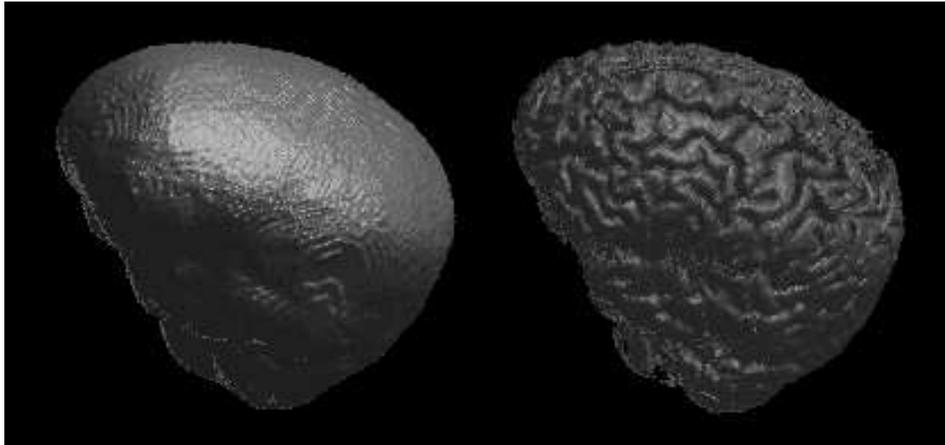


Figure 9: Example brain surface model (left) and resulting brain surface (right) generated by BET.

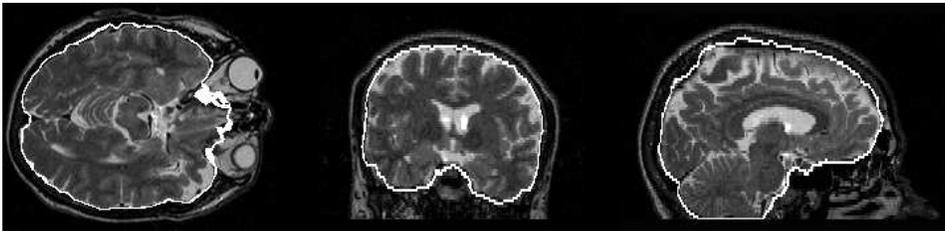


Figure 10: Example brain surface from a T2-weighted image.

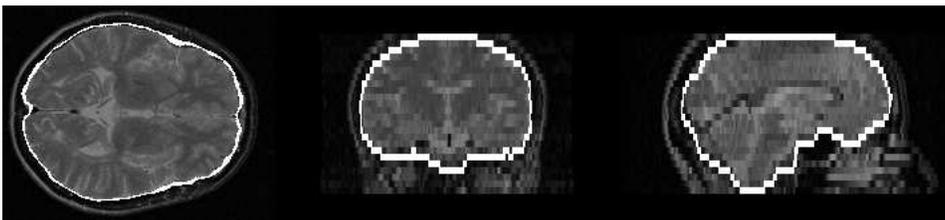


Figure 11: Example brain surface from a proton density image.

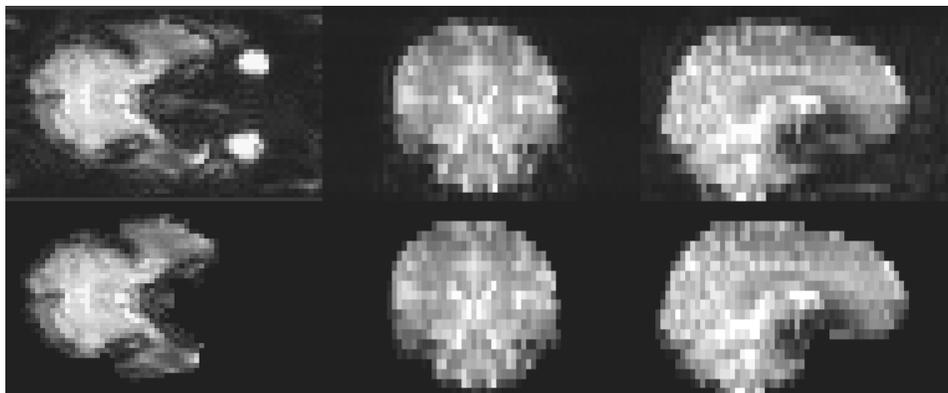


Figure 12: Example segmentation of an EPI image.

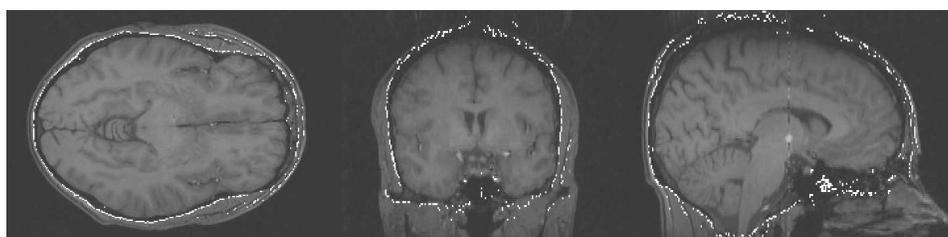


Figure 13: Example exterior skull surface generated by BET.

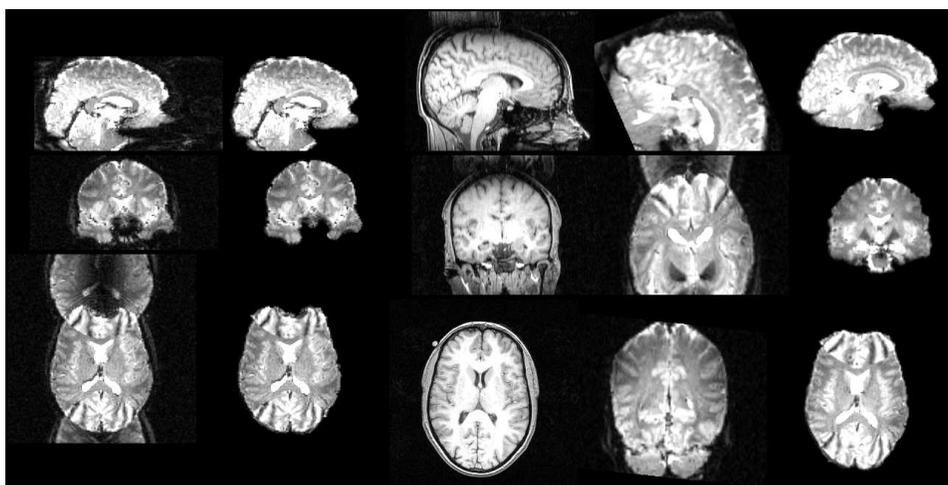


Figure 14: Left to right: the original FMRI image; BET output from the FMRI image; T1-weighted structural image; (failed) registration without using BET; successful registration if BET is used.

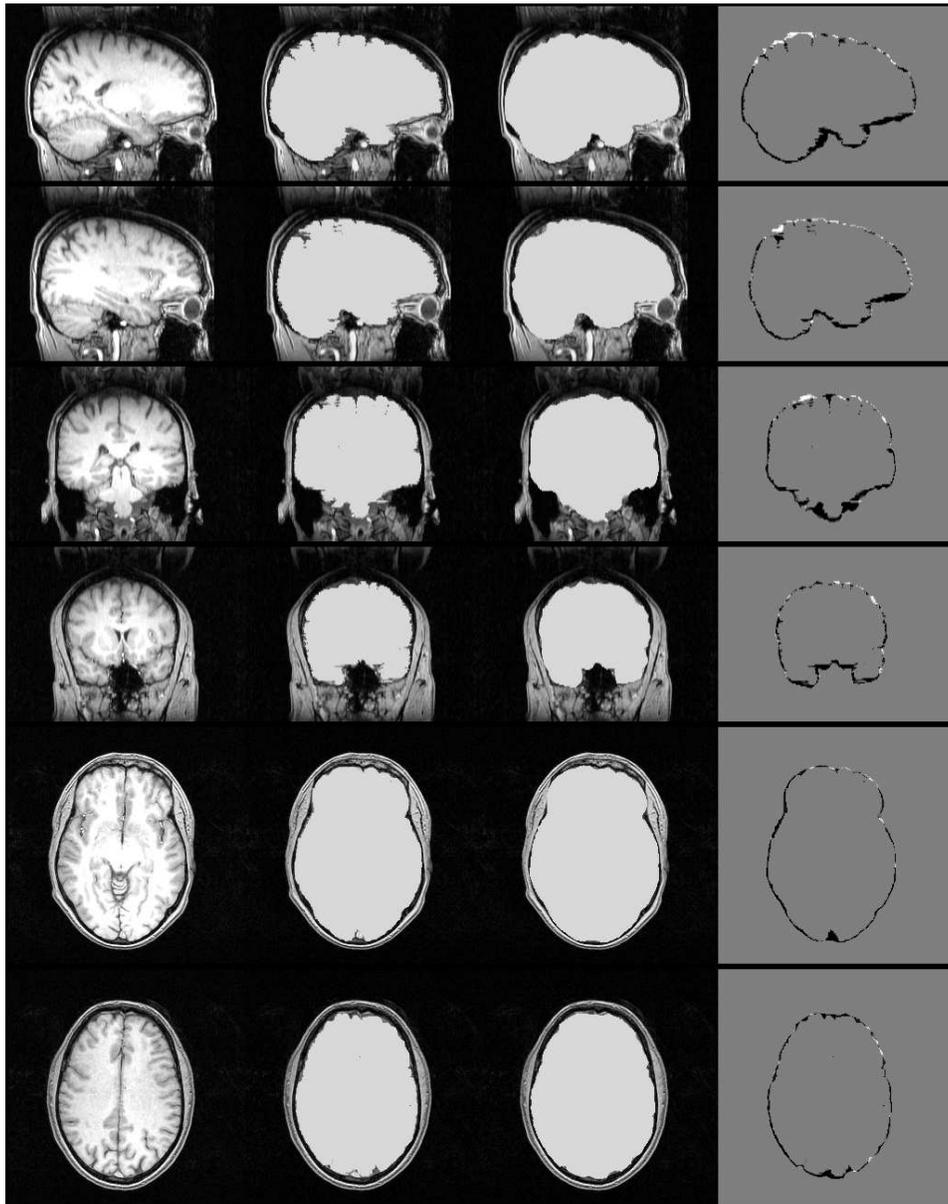


Figure 15: Left to right: Example original whole-head MR image; hand segmentation; fully automatic BET masking; hand mask minus BET mask.

in Section 3.2. The upper threshold was set to  $t_{98}$  and the lower threshold to 40% between  $t_2$  and  $t_{98}$ . The refined method is referred to below as “AFNI\*”.

The results of the three methods were evaluated using a simple % error formulation of  $0.5 * 100 * \text{volume}(\text{total nonintersection}) / \text{volume}(\text{hand mask})$ . The main results for the fully automated methods are shown on the left in Figure 16; the mean % error over the 45 images is shown for each method. The mean error is more meaningful than any robust measure (e.g. median) because outliers are considered relevant - the methods needs to be robust as well as accurate to be useful (though note that using median values instead gives the same relative results). The short bars show the results for the 35 T1-weighted input images only. BET gives significantly better results than the other two methods. Some slices of a typical BET segmentation<sup>3</sup> are displayed in the third column of Figure 15. The fourth column shows the hand mask minus the BET mask; in general BET is slightly overestimating the boundary (by approximately one voxel, except in the more complicated inferior regions), and of course smoothing across fine sulci.

It was also considered of interest to investigate the same test if initial controlling parameters were “hand-optimised” (i.e making the methods only “nearly-fully-optimised”). In order to carry this out in a reasonably objective manner, the primary controlling parameter for each method was varied over a wide range and the best result (comparing output with hand segmentation) was recorded. Fortunately, each method has one controlling parameter which has much greater effect on output than others, so the choice of which parameter to vary was simple.<sup>4</sup> The range over which each method’s principal controlling parameter was varied was chosen by hand such that the extremes were just having some useful effect in a few images. Each method was then run with the controlling parameter at 9 different levels within the range specified. The results are shown on the right in Figure 16; the methods all improve, to varying degrees. BET is still the best method, just beating AFNI\*. However, the most important message from these results is that although BET is the most accurate and robust method in both tests, it is also the most successfully “fully automated”, in that its results when run fully-automated are nearly as good as those when it is “hand-optimised”.

Note that all the above comments on the quantitative results also hold when only

---

<sup>3</sup>The chosen image gave an error close to BET’s mean error.

<sup>4</sup>With BET, the parameter varied affects the setting of a local brain/non-brain threshold;  $b_t$  in equation 11 was varied from 0.1-0.9. With AFNI\*, the setting of the lower intensity threshold was varied; instead of using 40% between robust intensity limits as described above, a range of 10-70% was used. With BSE, the “edge detection sigma”, which controls the initial edge detection, was varied from 0.5-1.5.

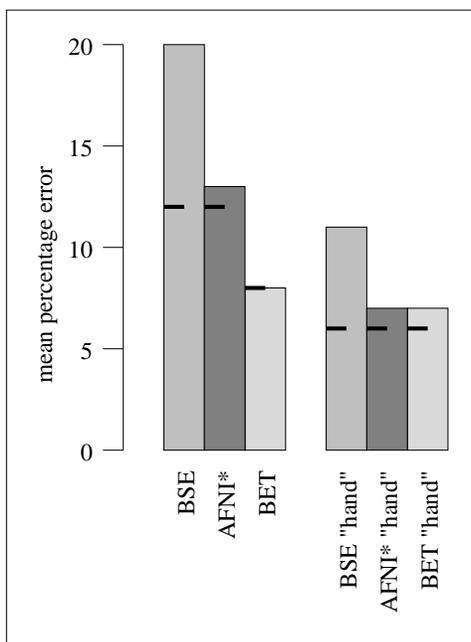


Figure 16: Mean % error over 45 MR images for three brain extraction methods, compared with hand segmentation; on the left are the results of testing the fully-automated methods, on the right are the “hand-optimised” results. The short bars show the results over only the 35 T1-weighted images

the 35 T1-weighted images are considered.

In theory it might be possible to “hand-tune” a method once for a given MR pulse sequence, and the resulting parameters then work well for all images of all subjects acquired using this sequence. If this were the case then possibly the results of AFNI\* and BET could be viewed as similarly successful (assuming that our improvements to AFNI were implemented). This is, however, not the case, as there was found to be significant variation in optimal controlling parameters for AFNI\* (within sequence type).

Finally, note that results from a brain extraction algorithm may improve if the image is pre-processed in certain ways, such as with an intensity inhomogeneity reduction algorithm. However, it is our experience that the best intensity inhomogeneity reduction methods require brain extraction to have already been carried out.

### **4.3 Results of using BET in Higher-Level Systems**

BET is used as the first stage in the measurement of longitudinal (two-time-point) atrophy in the SIENA system (Structural Image Evaluation, using Normalisation, of Atrophy) described in [13]. Using various different methods, including a 40 subject (120 image) dataset, a 500 subject dataset and a 16 subject (256 image) dataset, the estimation of percentage brain volume change error (which, amongst other things, is dependent on the robustness and accuracy of the brain extraction) was found to be around 0.2%. In all of these studies it was possible to use BET in an automated manner. As mentioned above, where the input parameter needed changing, for example, when proton density images were analysed, the input parameter was set once for all subjects, i.e. did not require tuning for each new subject.

BET is also used in the cross-sectional (single-time-point) atrophy estimation system SIENAX. Test-retest results of SIENAX show an error of between 0.5 and 1% (of brain volume, depending on image quality) - this has clear implications about the test-retest accuracy of BET. Again, this system has been run on hundreds of datasets successfully.

## 5 Conclusion

An automated method for segmenting MR head images into brain and non-brain has been developed. It is very robust and accurate and has been tested on thousands of data sets from a wide variety of scanners and taken with a wide variety of MR sequences. BET takes about 5-20 seconds to run on a modern workstation and is freely available (as a standalone program which can be run from the UNIX command line or from a simple TCL/TK GUI) as part of FSL, from [www.fmrib.ox.ac.uk/fsl](http://www.fmrib.ox.ac.uk/fsl)

## 6 Acknowledgements

The author acknowledges support to the FMRIB Centre from the UK Medical Research Council, and is very grateful to Prof. De Stefano, University of Siena, Italy, (for organising the sabbatical during which much of this research was carried out), and to Tomoyuki Mitsumori (for his painstaking work in creating the hand segmentations used in the quantitative evaluations presented here).

## References

- [1] M.S. Atkins and B.T. Mackiewich. Fully automatic segmentation of the brain in MRI. *IEEE Trans. on Medical Imaging*, 17(1):98–107, 1998.
- [2] M. Bomans, K.-H. Höhne, U. Tiede, and M. Riemer. 3-D segmentation of MR images of the head for 3-D display. *IEEE Trans. on Medical Imaging*, 9(2):177–183, 1990.
- [3] M.E. Brummer, R.M. Mersereau, R.L. Eisner, and R.R. Lewine. Automatic detection of brain contours in MRI data sets. *IEEE Trans. on Medical Imaging*, 12(2):153–166, 1993.
- [4] R. Cox. AFNI software package. [afni.nimh.nih.gov](http://afni.nimh.nih.gov).
- [5] A.M. Dale, B. Fischl, and M.I. Sereno. Cortical surface-based analysis I: Segmentation and surface reconstruction. *NeuroImage*, 9:179–194, 1999.
- [6] K.H. Höhne and W.A. Hanson. Interactive 3D segmentation of MRI and CT volumes using morphological operations. *Journal of Computer Assisted Tomography*, 16(2):185–294, 1992.

- [7] M. Jenkinson and S.M. Smith. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156, June 2001.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. In *Proc. 1st Int. Conf. on Computer Vision*, pages 259–268, 1987.
- [9] A. Kelemen, G. Székely, and G. Gerig. Elastic model-based segmentation of 3-D neuroradiological data sets. *IEEE Trans. on Medical Imaging*, 18(10):828–839, 1999.
- [10] F. Kruggel and D.Y. von Cramen. Alignment of magnetic-resonance brain datasets with the stereotactical coordinate system. *Medical Image Analysis*, 3(2):175–185, 1999.
- [11] L. Lemieux, G. Hagemann, K. Krakow, and F. G. Woermann. Fast, accurate, and reproducible automatic segmentation of the brain in T1-weighted volume MRI data. *Magnetic Resonance in Medicine*, 42(1):127–35, Jul 1999.
- [12] S. Sandor and R. Leahy. Surface-based labeling of cortical anatomy using a deformable database. *IEEE Trans. on Medical Imaging*, 16(1):41–54, 1997.
- [13] S.M. Smith, N. De Stefano, M. Jenkinson, and P.M. Matthews. Normalised accurate measurement of longitudinal brain change. *Journal of Computer Assisted Tomography*, 25(3):466–475, May/June 2001.
- [14] S.M. Smith, Y. Zhang, M. Jenkinson, J. Chen, P.M. Matthews, A. Federico, and N. De Stefano. Accurate, robust and automated longitudinal and cross-sectional brain change analysis. *NeuroImage*, 2002. In publication.
- [15] B.D. Ward. Intracranial segmentation. Technical report, Medical College of Wisconsin, 1999. [afni.nimh.nih.gov/afni/docpdf/3dIntracranial.pdf](http://afni.nimh.nih.gov/afni/docpdf/3dIntracranial.pdf).
- [16] Y. Zhang, M. Brady, and S. Smith. Segmentation of brain MR images through a hidden Markov random field model and the expectation maximization algorithm. *IEEE Trans. on Medical Imaging*, 20(1):45–57, 2001.