A New Fast Accurate Non-Linear Medical Image Registration Program Including Surface Preserving Regularization

Audrunas Gruslys, Julio Acosta-Cabronero, Peter J. Nestor, Guy B. Williams, and Richard E. Ansorge*., *Member, IEEE*

Abstract— Recently inexpensive graphical processing units (GPUs) have become established as a viable alternative to traditional CPUs for many medical image processing applications. GPUs offer the potential of very significant improvements in performance at low cost and with low power consumption. One way in which GPU programs differ from traditional CPU programs is that increasingly elaborate calculations per voxel may not impact of the overall processing time because memory access dominates execution time. This paper presents a new GPU based elastic image registration program named Ezys. The Ezys image registration algorithm belongs to the wide class of diffeomorphic demons but uses surface preserving image smoothing and regularization filters designed for a GPU that would be computationally expensive on a CPU. We describe the methods used in Ezys and present results from two important neuroscience applications. Firstly inter-subject registration for transfer of anatomical labels and secondly longitudinal intra-subject registration to quantify atrophy in individual subjects. Both experiments showed that Ezys registration compares favourably with other popular elastic image registration programs. We believe Ezys is a useful tool for neuroscience and other applications, and also demonstrates the value of developing of novel image processing filters specifically designed for GPUs.

Index Terms—**deformable registration, medical image analysis, demons algorithms, graphics processors, multi-core processing.**

Manuscript received ??? . This work was supported in part by EPSRC.
Asterisk indicate corresponding author
*A Gruslys is with the Cavendish Laboratory, J J Thomson Avenue, Cambridge CB3 0HE, UK (e-mail a.gruslys@cantab.net).
J. Acosta-Cabronero is with Department of Clinical Neurosciences; School of Clinical Medicine; University of Cambridge; Robinson Way, Cambridge CB2 0SZ, United Kingdom and The German Center for Neurodegenerative Diseases (DZNE), Leipziger Strasse 44, 39120 Magdeburg, Germany (e-mail jac@cantab.net).
P.J. Nestor is with The German Center for Neurodegenerative Diseases (DZNE), Leipziger Strasse 44, 39120 Magdeburg, Germany (e-mail peter.nestor@dzne.de).
G.B. Williams is with (e-mail gbw1000@wbic.cam.ac.uk).
R.E. Ansorge is with the Cavendish Laboratory, J J Thomson Avenue, Cambridge CB3 0HE, UK (e-mail rea1@cam.ac.uk).

.

I. Introduction

FREE-FORM image registration methods have many important applications in medical image processing. However such registrations are computationally demanding; free-form image registration algorithms have to optimize a large number, perhaps millions, of free parameters whilst also applying smoothness constraints to ensure physicality. In their 2009 publication [1] Klein et al. compared 14 popular free-form registration programs having typical processing time per registration of between many minutes and hours. More recently, graphical processing units (GPUs) have been applied to the problem and impressive speed ups have been demonstrated [2-4]. Importantly, this increase in computational speed not just convenient, it also enables the design of more complex image registration algorithms and processing pipelines.

The present paper introduces a new, GPU-enabled, free-form image registration program, Ezys, which incorporates surface preserving smoothing and regularization filters optimized for GPUs. Ezys is very fast, robust and performs well for important neuroscience applications including inter-subject registration for transfer of anatomical labels and intra-subject serial registration for detection of subtle longitudinal volume changes in dementia patients. We describe the Ezys algorithms in some detail, including novel image denoising and surface-based regularization methods.

As validation we present two studies. The first study repeats Klein et al.'s inter-subject anatomical label matching test where we find Ezys performs better than any of the 14 programs previously tested using the median target overlap measure. On this measure Ezys also outperforms a more recent version of the ANTS-SyN program [5] which itself was significantly better than the beta version used in [1] where it was found to be one of the best performing programs. This first study can be regarded as an application of voxel based morphometry (VBM) [6].

Our second validation study involves intra-subject registration to quantify progressive regional brain atrophy in individual subjects using the Jacobians of the derived deformation fields for tensor based morphometry (TBM). A total of 18 subjects were studied. In practice this method is very sensitive to the quality of the transformation fields. We repeated our Ezys TBM analysis using three other popular

image registration packages: FNIRT - part of the FSL 5.0 software toolkit [7], NiftyReg 1.3 - a GPU enabled program [3], and DARTEL - a diffeomorphic image registration algorithm [8]. DARTEL was one of the best performing methods evaluated by [1] and has been available as part of the statistical parametric mapping software package, SPM, since the SPM5 version. The SPM package is frequently employed for neuroscience image-analysis and a recent paper describes many applications [9].

Progressive brain atrophy is a well known feature in many neurodegenerative diseases. Detecting atrophy patterns in individual subjects is important for accurate diagnosis and measuring therapeutic outcomes. Typically sequential MRI scans are acquired at intervals of 6 months or more and many different techniques have been developed to quantitatively analyze such datasets. Such methods include voxel based morphometry (VBM) [6], the boundary shift integral (BSI) [10], structural image evaluation using normalization of atrophy (SIENA) [11] and tensor-based morphometry (TBM) [12-15]. Both BSI and SIENA use linear image registration to transform longitudinal images to a common coordinate system and then look for changes in tissue boundaries within the brain. In VBM, longitudinal images are also registered to a standard or custom cohort-specific space and a voxel level statistical analysis on the grey and/or white matter partitions is then used. In contrast TBM uses the deformation fields relating two structural images directly. The determinant of the Jacobian matrix of the transformation field is calculated at each voxel and is a direct measure of the local volume change. For brevity we will use the term Jacobian map to mean Jacobian determinant map in the rest of this paper. TBM has some advantages over other methods in that Jacobian maps give local volume changes between the two images at voxel resolution. Finally the full Jacobian matrix has potentially more information than just its determinant, for example the eigenvectors or Green's tensor can be used [16].

II. Materials and Methods

The details of the Ezys registration methods are described in sections A-F, the subjects and methods used in our VBM analysis in sections G-H and for the TBM analysis in sections I-K.

A. Overview Image Registration

Image registration requires matching a source image to a target image. This is achieved by finding a displacement map or transformation, *T*, relating each point in the target image to a corresponding point in the source image. Corresponding points are selected by optimizing a similarity measure (or objective cost function) such as the sum of squared intensity differences, normalized correlation [17] or normalized mutual information (NMI) [18] among others. The displacement map, *T*, can be used to estimate local volume differences between source and target brain structures in both longitudinal and cross-sectional study designs. In particular, the Jacobian determinant of *T* measures local volume changes. *T* is diffeomorphic if its Jacobian is positive everywhere. This is normally required for medical image registration making *T* invertible and avoiding fold over.

Both the source, I_s, and target image, I_t, can be regarded as functions that relate position with intensity, the intensity of an image voxel at coordinates $\underline{r}$ being $I(\underline{r})$. Medical images, including those acquired with MRI, are voxelated, hence exact intensity values are only defined at discrete coordinates. In the present context, a transformation, *T*, which deforms a source image to match the target, can be interpreted as a three-dimensional vector field, $T(\underline{r})$, that relates coordinates of any given point in a target image to the coordinates of a corresponding point in a source image. *T*, therefore, deforms a source image in the following way: $I_s^{def}(\underline{r}) = I_s(T(\underline{r}))$. Here, we will use the notation TI_s as shorthand for $I_s(T(\underline{r}))$. Note that the argument, $\underline{r}$, is defined over the domain of the target image. The aim of the image registration algorithm is to optimize *T* by maximizing (or minimizing if appropriate) the chosen similarity measure, subject to any smoothness constraints imposed to ensure physicality of the resulting transformation, to find the "best match" between the deformed source and the target images.

The final displacement map *T* will in general point to non-integer voxel positions in the source image, hence some method of choosing a suitable intensity value from the source image is required. Trilinear interpolation is commonly used for this purpose and gives a weighted average of the neighbouring source voxel values. As interpolation causes blurring of the transformed image it is important to concatenate displacements to avoid multiple interpolation steps during registration. Ezys implements nearest neighbour, trilinear and sinc interpolation methods; typically trilinear interpolation is used for transforming MR and Jacobian images and nearest neighbour for transforming segmentation maps.

B. The Ezys Image Registration Algorithm

The Ezys image registration algorithm belongs to a wide class of demons algorithms first introduced by Thirion [19]; in particular, it belongs to the subclass of diffeomorphic demons. A Gaussian filter is used as the "fluid" regulariser, and our customized anisotropic diffusion filter (described in section F) is used as the diffusion regulariser. Similarity measures including NMI and normalized correlation ratio are implemented in Ezys, NMI was used as the similarity measure for all the results presented here. Several demons algorithms do not require explicit similarity measures; they only use forces. However such force fields, if conservative, can be regarded as the gradient of an objective cost function with respect to the free parameters defining the transformation.

After an initial affine transformation has been found our elastic registration algorithm runs through the following steps:

1. Reduce noise and enhance surfaces on both source and target images using the surface enhancement algorithm as discussed in section D below.
2. Detect surfaces on a target image as discussed in section E.

3. Assign an initial transformation, *T*, parameterized by a set of parameters p_i (indexed by *i*). These are estimated using geometric moments or by registering sub-sampled versions of source and target images.
4. Then iterate across steps 4.1-4.4 until convergence:
   4.1. Calculate the gradient, *G*, of the chosen similarity measure, *S*, with respect to the transformation parameters p_i, so that that $G_i = \partial S(TI_s, I_t)/\partial p_i$, where $S = S(TI_{s,} I_t)$. In our case the p_i are just the displacements at each target voxel, thus *G* is a vector field over the domain of the target voxels.
   4.2 Apply a 3D Gaussian filter to smooth *G*, reducing kernel width as registration progresses.
   4.3. Find an optimum step size λ_{opt} by performing the modified gradient descent steps 4.3.1 to 4.3.3:
      4.3.1. If λ_0 is the optimum λ value from the previous step, for trial values λ_{tr} in $\{\pm\lambda_0; \pm 4\lambda_0; \pm\lambda_0/4\}$, calculate transformations $T_{tr}(\underline{r}) = T(\underline{r} + \lambda_{tr} G(\underline{r}))$. Use of function composition rather than addition ensures that the T_{tr} are diffeomorphic provided T is diffeomorphic and λ_{tr} are not too large.
      4.3.2. For each T_{tr}, evaluate $S_{tr}(T_{tr} I_s, I_t)$ between TI_s and I_t. A parabolic fit to the S_{tr} against λ_{tr} is then used to find the new λ_{opt} value corresponding to an optimum value of S and set $T(\underline{r}) = T(\underline{r} + \lambda_{opt} G(\underline{r}))$.
      4.3.3. For one in five iterations a random non-optimal step is used instead of λ_{opt} to help escape from local stationery points.
   4.4. Given the new value of *T*, apply a regularization step to smooth the transformation using the surface preserving smoothing filter described in sections E and F below.
5. Done.

We found that using function composition rather than addition in step 4.3.1 contributes significantly to the overall performance of our algorithm, as it ensures that the resulting transformation is diffeomorphic. Most fluid-flow based image registration algorithms use a similar idea when exponentiating velocity fields, however our method is less computationally intensive and well suited to GPU based interpolation.

C. Symmetric Image Registration

Asymmetric image registration algorithms, where an interpolated source image is warped to match a stationary target, have recently been criticized for introducing bias [20, 21]. We implemented an option to average direct, $A \rightarrow B$, and indirect, $B \rightarrow A$, registrations, thus if the asymmetric registration $A \rightarrow B$ results in T_1, and the registration $B \rightarrow A$, in T_2, we compute $T_{A \rightarrow B}(\underline{r})$ as $(T_1(\underline{r}) + T_2^{-1}(\underline{r}))/2$, where T_2^{-1} is the inverse of the transformation T_2, [22]. This option increases symmetry and reduces potential bias caused by the asymmetric treatment of source and target images at the expense of doubling the number of registrations required.

D. Surface Preserving Noise Reduction

Prior to surface detection, a step to reduce image noise and to increase image contrast while preserving surfaces was implemented as follows:

For each given image voxel at $\underline{r}$ we assume it and its 26 nearest neighbours (forming a $3\times3\times3$ cube) can be split into two tissue classes by a cut-plane through the cube. The corrected image intensity at $\underline{r}$ is then computed as the median signal intensity of all voxels within the cube that belong to the same tissue class as $\underline{r}$. We choose the cut-pane by testing 39 candidate planes passing through the cube either parallel to or at 45° to one of the coordinate planes. For each candidate, the voxels are divided into two groups and a measure of separation is calculated as $\left|\bar{V}_{on} - \bar{V}_{off}\right|$, where $\bar{V}_{on}$ is the mean intensity of the voxels cut by the plane and $\bar{V}_{off}$ is the mean intensity of the voxels not cut by the plane. The 39 candidate planes are chosen such that each of these groups contains at least 9 voxels. The cut plane with the maximum measure of separation is selected and the voxel at $\underline{r}$ is assigned the median value of the $\bar{V}_{on}$ population if it is on the plane or otherwise the median value of the $\bar{V}_{off}$ population. An example of the filter is shown in Fig. 1 for one and two iterations. Two iterations were used for all the work presented here The algorithm both reduces image noise and increases tissue contrast. Note that although this filter is computationally expensive, it is well suited to parallel computation on GPUs.

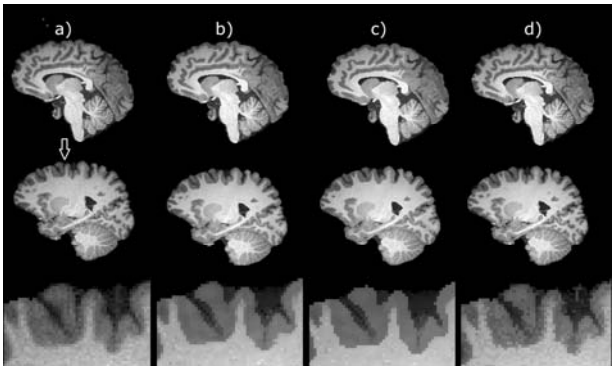

Fig.1. Example of Ezys surface enhancement algorithm. From left to right: a) original image, b) after one iteration, c) after two iterations and d) using a standard lower-upper-median (LUM) filter. The top two rows show sagittal slices from a T1 weighted structural image; the bottom row shows an enlarged view of the cortical region pointed by the arrow.

E. Surface Detection

Ezys performs an iterative surface-detection operation that assigns weights $W(\underline{r})$ to each voxel. The aim of this surface detection approach is to weight each voxel according to its

likelihood of belonging to a surface within the image volume. We consider a surface to be the boundary between image sub-volumes each of which is relatively homogeneous but with different mean intensities. The weights are used as a constraint on the Ezys regularization procedure described in section F. The method is as follows:

1. Initially set all weights equal to zero: $W(\underline{r}) = 0$ for all $\underline{r}$.
Repeat the following steps 2.1 - 2.5 for 5 iterations:
   2.1. For each voxel with coordinates $\underline{r} = (x, y, z)$ and for some selected distance *d* (typically d is 6 voxels), compute means and variances for voxels laying along one of the coordinate axes with distance less than *d* from $\underline{r}$. Along *x*, define a_{x+} and a_{x-} as the mean values at either side of r and v_{x+} and v_{x-} as the corresponding intensity variances; compute analogous means and variances along *y* and *z*.
   2.2. Compute surface sensitive signal to noise ratio along each axis, for *x* this is $s_x = (a_{x+} - a_{x-}) / \sqrt{v_{x+} + v_{x-}}$.
   2.3. For each Cartesian axis $j \in \{x, y, z\}$ calculate a distance n_j from the central voxel $\underline{r}$ to the nearest voxel along the *j* axis that has a weight greater or equal to the weight $W(\underline{r})$ of the central voxel. If there is no such voxel within a maximum distance *d*, then assign $n_j = d$.
   2.4. Compute coefficients $k_j = (1 - 1/(1 + 3n_j / d - 0.5))$ to determine the level of response along each axis and update all weights: $W(\underline{r}) = \sqrt{k_x s_x + k_y s_y + k_z s_z}$.
   2.5. If this is not the last iteration, apply an isotropic Gaussian filter to the surface estimates. The algorithm halves the width of the smoothing kernel in each iteration starting with 8 voxels.
3. After the last iteration, non-surface weights are set to zero. If current weights are *W*, calculate a smoothed version W_s of *W* by applying a 16-voxel wide isotropic Gaussian filter along each axis. If $W(\underline{r}) < W_s(\underline{r}) / 8$, set $W_{final}(\underline{r}) = 0$, otherwise set $W_{final}(\underline{r}) = W(\underline{r}) / W_s(\underline{r})$.

The effect of our surface estimation method compared to a simple smoothed gradient is that step 2.3 in our method produces surfaces which are 1 voxel thick

F. Surface Preserving Regularization

Free form registrations have to be constrained to prevent nonphysical results such as negative volume changes; this process is known as regularization. There are two classes of strategy commonly employed for regularization. In the first class, extra terms are added to the cost function to penalize for excessive local deformation [23], examples of such regularization terms include elastic energy and bending energy constraints. A disadvantage of these methods is the difficulty of selecting suitable weighting factors. The second class attempts to implicitly constrain the optimization to keep the transformation field sufficiently smooth. Methods in this category include B-spline based deformations [24] which impose smoothness by limiting the number of degrees of freedom. Alternatively demons algorithms [19] use diffusion or fluid regularisers to ensure some level of implicit smoothness. The regularization used by Ezys firstly applies a Gaussian smoothing filter to the gradient field (equivalent to a fluid regulariser) and secondly by modifying the displacement field by simulating diffusion constrained by the detected surfaces (a diffusion regulariser). Simulating diffusion to ensure smoothness is not novel, for example edge detection for 2D image processing using diffusion with a varying diffusion constant was discussed by Perona and Malik [25]. Adaptive regularization methods have also been discussed by many authors, including Hermosillo et al. [26] who include locally computed similarity measures in their survey of multimodality image matching, Stefanescu et al. [27] who use a local intensity gradient based regularizer and more recently by Simpson et al. [28] who use Bayesian inference methods for determining locally adaptive regularisation constraints. Nevertheless our particular choice of an anisotropic, spatially-dependent diffusion kernel, optimized for GPUs, is new. The registration performance of Ezys is strongly dependant on the choice of an appropriate kernel. As discussed in the previous section, weights $W(\underline{r})$ are found for each voxel; with positive values representing points more likely to be on a surface. These weights are incorporated into the diffusion kernel. If $W(\underline{r})$ is large at a point, less smoothing is applied at that point so as to avoid blurring of boundaries.

The Ezys diffusion kernel, $K(\delta\underline{r}, \underline{r})$, is defined with the following properties:

1. *K* is linear in $T(\underline{r})$ for performance reasons, ease of implementation, analysis and to avoid various non-linear effects, i.e. $T_{smoothed}(\underline{r}) = \sum_{\delta\underline{r}} K(\delta\underline{r}, \underline{r}) T(\underline{r} + \delta\underline{r})$;

2. Applying *K* does not cause net-expansion of the source image, thus the sum of weights must be 1, i.e. $\sum_{\delta\underline{r}} K(\delta\underline{r}, \underline{r}) = 1$;

3. *K* should be unbiased in the sense that it should not cause net-drift of the source image with respect to the target image after each application. Hence the first geometric moment of *K* should be 0: $\sum_{\delta\underline{r}} K(\delta\underline{r}, \underline{r}) \delta\underline{r} = 0$;

4. For numerical stability, neighbouring points, $\underline{r} + \delta\underline{r}$, should be weighted less than the point of interest: i.e. $K(\delta\underline{r}, \underline{r}) < K(0, \underline{r})$;

5. A large surface weight, $W(\underline{r} + \delta\underline{r})$, at the point $\delta\underline{r} + \underline{r}$ should cause $K(\delta\underline{r}, \underline{r})$ to be large: i.e. for any fixed $\underline{r}$ and $\delta\underline{r}$, $K(\delta\underline{r}, \underline{r})$ should always increase if $W(\underline{r} + \delta\underline{r})$ does;

6. Points will always see their weight diminish as a function of their distance from $\underline{r}$: i.e. $K(\delta\underline{r}_2, \underline{r}) < K(\delta\underline{r}_1, \underline{r})$ if $|\delta\underline{r}_2| > |\delta\underline{r}_1|$ and $W(\underline{r} + \delta\underline{r}_2) \le W(\underline{r} + \delta\underline{r}_1)$.

A function satisfying the above requirements was implemented to run on the GPU. For each point with coordinates $\underline{r}$ and any unit vector in voxel space $\underline{p}$ along a particular coordinate axis, we define the function

$f_{\underline{r},\underline{p}}(i) = W(\underline{r} + i\,\underline{p})$ where i is an integer. The algorithm then finds two local maxima for $f_{\underline{r},\underline{p}}(i)$: one negative $i = i_1 < 0$ and one positive $i = i_2 > 0$. Then we define:

$$M_{\underline{r},\underline{p}}(i) = 1/|i_1| \text{ for } i_1 \le i < 0,$$
$$M_{\underline{r},\underline{p}}(i) = 1/|i_2| \text{ for } 0 < i \le i_2,$$
$$M_{\underline{r},\underline{p}}(i) = 0 \text{ for } i < i_1 \text{ or } i > i_2 \text{ and}$$
$$M_{\underline{r},\underline{p}}(0) = \max(M_{\underline{r},\underline{p}}(-1), M_{\underline{r},\underline{p}}(1)).$$

At present we search for up to 5 steps in either direction, if no maximum is found the corresponding i_1 or i_2 value is set to 5. This function satisfies all the above conditions except (2); thus finally, we normalize it:
$K(\underline{p}_i, \underline{r}) = M_{\underline{r},\underline{p}}(i) / \sum_j M_{\underline{r},\underline{p}}(j)$, and $K(\delta\underline{r}, \underline{r}) = 0$ for directions $\delta\underline{r}$ that are not parallel to our selected orientation $\underline{p}$. This smoothing step is applied in turn for $\underline{p}$ along each coordinate axis.

G. VBM based Target Overlap Calculation

In [1] Klein et al. compare the performance of 14 registration programs using target overlap measures. Target overlap is a VBM based measure of registration accuracy. Given a set of images each having a voxel based labelled segmentation, pairs of images can be registered and the fraction of labelled voxels in the target image which match the propagated labels of the source image calculated. In fact two different definitions of target overlap are discussed in Klein's paper. Either calculate the fraction of matching labels over the whole image, $TO_1 = \frac{\sum_r \|S_r \cap T_r\|}{\sum_r \|T_r\|}$, where S_r is the r-th labelled region in the registered source image and T_r is the r-th labelled region in the target image, and $\|\ \|$ indicates a sum over voxels. Or calculate separate target overlap measures for each region r, as $TO_r = \frac{\sum_r \|S_r \cap T_r\|}{\sum_r \|T_r\|}$ and then define the target overlap as the average of the regional target overlaps: $TO_2 = \sum_r TO_r / N_{regions}$. The second measure is more sensitive to image registration quality as all regions have equal weight, and is the measure used in Figure 5 of Klein et al.'s paper.

We reproduced the calculations using Ezys for 80 normal subjects with T1 weighted MR scans and corresponding manually segmented label sets. The data comprised 40 subjects used to construct the LONI Probabilistic Brain Atlas (LPBA40), 18 subjects from the Internet Brain Segmentation Repository (IBSR18), 12 subjects from the Columbia University Medical Center (CUMC12) and 10 subjects scanned at the MGH/MIT/HMS Athinoula A. Martinos Center for Biomedical Imaging (MGH10). Full details of these datasets and the preprocessing performed including brain extraction are given in Klein et al.'s paper. We used their preprocessed image data downloaded from the associated web site (http://www.mindboggle.info/papers/). For all images Ezys was run with the VBM settings shown in Table I below.

The ANTS-SyN program [5] was one of the best performing programs evaluated by Klein et al. and we repeated the analysis using the more recent ANTs-1.9.x-win32.exe version from http://sourceforge.net/projects/advants/files/ANTS/. These registrations were processed using the same software tools that were used for processing the Ezys registrations.

H. Subjects and MRI Protocols for Longitudinal TBM

T1 weighted MRI scans of eighteen subjects were used for the TBM study: 5 elderly controls and 6 Alzheimer's disease (AD) subjects each with three time points at 6 monthly intervals and, 7 subjects with semantic dementia (SD) with two time points having intervals between 19 and 38 months. These subjects were selected as a representative subset of a larger group previously reported on in [29, 30]. The whole-head anatomical scans were acquired on a Siemens Trio 3T system (Siemens Medical Systems, Erlangen, Germany) with gradient coils capable of 45 mT/m and 200 T/m/s slew rate. A standard 12-channel phased-array total imaging matrix head-coil (Siemens Medical Systems, Erlangen, Germany) was used for radio-frequency reception. The structural scans consisted of 3D magnetization-prepared, rapid gradient-echo (MPRAGE) volumes acquired with the following imaging parameters: TR / TE / inversion time / flip angle = 2300 ms/2.86 ms/900 ms/ 9°, the image matrix was $192 \times 192 \times 144$ with isotropic 1.25 mm voxel dimensions. Whole head images were used for the registrations so as to minimize potential errors from inconsistent brain extraction between time points. Note, however, that for some of the post registration quantitative processing and preparation of figures SPM5 derived brain masks were used.

I. TBM Analysis using Ezys

For each of the 18 subjects all possible pairwise intra-subject registrations were performed, yielding 2 or 6 Jacobian maps in the case of 2 or 3 time points. For each subject all maps were transformed to a common frame and annual atrophy rates were then computed on a per voxel basis assuming a constant atrophy rate. To quantify regional changes, images were segmented using a manually improved Brodmann map [31]; note that the original version was obtained from the MRIcron software package [32]. In more detail the processing pipeline for the each subject and each available pair of time points was as follows:

1. The image at time-point A was registered to that at time-point B and vice versa, yielding transformations fields: $T_{A\to B}$ and $T_{B\to A}$. The Ezys settings used are shown in Table I below.
2. Next compute the transformation composition: $T_{A\to B\to A} = T_{B\to A}\,T_{A\to B} \equiv T_{B\to A}(T_{A\to B}(\underline{r}))$.

3. The Jacobian map $J_{A\to B\to A}$ corresponding to each $T_{A\to B\to A}$ was calculated. Ideally, $J_{A\to B\to A}(\underline{r})$ should be unity everywhere, however in practice it may not be. Making the unbiased assumption that $T_{A\to B}$ and $T_{B\to A}$ are both affected by the same multiplicative Jacobian error we obtain $J_{error}(\underline{r}) = \sqrt{J_{A\to B\to A}(\underline{r})}$.

4. Corrected Jacobian maps, $J^{c}_{A\to B}(\underline{r})$ and $J^{c}_{B\to A}(\underline{r})$, were then calculated by dividing by $J_{error}(\underline{r})$.

5. The volume change estimates in $J^{c}_{A\to B}$ are in the coordinate system of B; whereas whose of $J^{c}_{B\to A}$ are in A's coordinate system. Having reciprocal Jacobians in different coordinate systems complicates the analysis. To deal with this we transformed all Jacobians into a common frame (that of a particular elderly control) using transformations obtained by registering the original images to the common frame. This resulted in a set of Jacobians, either $J^{cf}_{A\to B}$ and $J^{cf}_{B\to A}$ in the case of two points or six such Jacobians in the case of three time points.

6. The set of the Jacobian maps for the subject ($J^{cf}_{A\to B}$ etc.) were then used to fit (voxel-by-voxel) for the annual atrophy rate assuming this was constant.

7. The atrophy rates from step 6 were averaged over each separate Brodmann region in grey matter using an SPM5 derived grey matter mask for the common frame and the Brodmann map discussed above.

8. Subsets of the Brodmann regions were combined to produce atrophy estimates of grey matter within the temporal and parietal lobes.

J. TBM Analysis using DARTEL

DARTEL registration requires prior segmentation of each brain into grey matter, white matter and (optionally) CSF tissue classes. This segmentation is part of the SPM software suite into which DARTEL is embedded [33]. In typical use the set of images being analyzed is first segmented, then individually rigidly registered to MNI152 space [34] and an average template image is constructed by iterative non-linear registration of all images to the average. Our DARTEL analysis had the following steps:

1. Firstly each image was segmented using the unified segmentation [9] in SPM8 with default settings, generating grey and white matter probability maps for each subject and time point in native space.

2. The segmented images were imported into DARTEL using the native 1.25 mm isotropic voxel spacing.

3. The segmented images were then used to create a common template using DARTEL and grey and while matter segmentations (CSF maps were not used). This produced transformation fields between each image's native space and the common template space. These fields can be used for both forward (native to template) and inverse (template to native) transformations

4. For each subject and pair of time points the ratio of the Jacobian maps was calculated to yield an estimate of the Jacobian map between the time points in template space. These maps could then be further transformed into the native space of the subject or the common frame used in step 5 of sub-section I.

5. Subsequent processing was the same as steps 6-8 in sub-section I.

K. TBM Analysis using NiftyReg 1.3 and FSL5

Processing with these programs followed the same steps as section I, but omitting the J_{error} correction. The programs were run with the finest practical warp resolutions in order to improve detail in Jacobian maps. For FNIRT we used `warpres = 7.5,7.5,7.5` and for NiftyReg we used `sx=sy=sz=2.0` and bending energy regularization `be=0.2`.

L. Processing

Programs were run on a PC running Windows 7 64bit with 8 GB of ram, i7-2600 (3.40 GHz) cpu and Nvidia GTX 580 CUDA capable graphics card with 512 cores and 1.5 GB of ram. FSL5 is not available for a native Windows OS and was run in on a virtual Ubuntu 12.04 hosted by Oracle VM Virtualbox, restricted to using 3 GB of host ram and only one processor.

.

III. RESULTS

A. Performance and Quality

Table I shows the parameter settings used for Ezys in the VBM and TBM studies and Table II gives the average times required for one pairwise registration. In the VBM study ANTS-SyN was about 18.5 times slower than Ezys on our PC and in fact used all the available CPU cores in parallel for some of its calculations.

TABLE I
EZYS PARAMETER SETTINGS

Parameter	VBM	TBM
MinGrid	8	4
MaxGrid	384	128
Steps	2	1
Reg Type	Direct	Both
Diffeomorphic	1	1
Diffusion	1	1
Elastic Termination	0.0001	0.0001
Cost Fun	NMI	NMI

Adjustable parameter settings for Ezys processing, Min/MaxGrid set the limits for the inverse widths of the Gaussian smoothing of displacement vectors, Steps controls the number resolution subdivisions used, Reg Type is direct for simple source to target registration or Both for an average of source to target and the inverse of target to source. The Diffeomorphic switch enables convolution of displacements for implicit prevention of negative Jacobian values and Diffusion enables the surface aware diffusion regularization filter. The Elastic Termination parameter sets minimum cost function change for convergence.

TABLE II
PROGRAM EXECUTION TIMES PER REGISTRATION IN SECONDS

Program	CUMC12	IBSR18	MGH10	LPBA40	WBIC TBM
Ezys	60.5	44.1	45.2	34.4	87
ANTS-SyN	1138	864	825	849	
FSL 5					1324
NiftyReg 1.3					111
DARTEL					605

Average execution times in seconds for one pairwise registration for programs listed in column 1. Columns 2-5 show the times for the four datasets used for the VBM comparison with Klein et al. and column 6 shows the times for the TBM atrophy analysis. Our DARTEL processing involved creation of a common template and subsequent registration to this template for 18 images, rather than 80 pairwise registrations. Thus in the table we show the total DARTEL processing time divided by 80 to be comparable with the other figures.

In Fig. 2 we show images for a typical inter-subject pairwise registration using direct registration VBM parameters for Ezys and ANTS-SyN with Klein et al.'s standard settings. While the overall match between the warped source and target images is good there are residual differences. Where the geometry of the sulci and gyri differ significantly, residual thin lines are left on the warped images, this may be due to the diffeomorphic registration constraint. The white arrows in the Fig. 2 point to examples. The difference images in columns 6 and 8 for Ezys and ANTS-SyN are similar with SyN being slightly smoother.

Fig. 3 shows intra-subject deformations from the TBM study for a SD patient scanned at time points 1.4 years apart. Characteristic temporal lobe atrophy is clearly seen for all four programs however significantly better detail is seen in the Ezys and DARTEL images which are comparable.

The accuracy of the symmetric registration discussed in section II.C was assessed by calculating the voxel based displacement error, $\mathrm{DE}(\underline{r}) = \left| T_{A\to B}(T_{B\to A}(\underline{r})) - \underline{r} \right|$, for typical intra-subject registrations. For the case of symmetric registration we found a mean DE of 0.28 ± 0.22 voxels and for direct one way registration 0.61 ± 0.38 voxels.

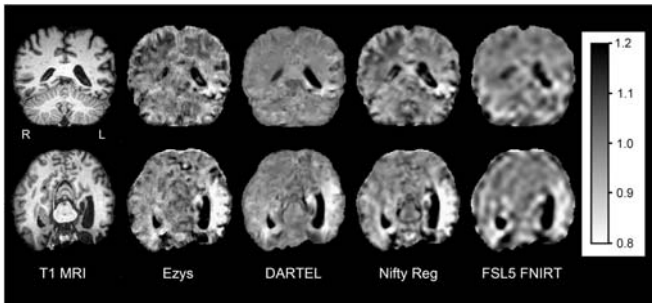

Fig. 3**.** Coronal (top row) and axial (bottom row) cross-sections of Jacobian maps inferred from four registration algorithms for a semantic dementia patient scanned twice, 1.4 years apart. The T1 MRI image for the latter time point is shown on the left hand column. Jacobian values represent voxel based volume changes between the two time points.

B. VBM Accuracy

We ran Ezys with the VMB settings in Table I for all image pairs in the four data sets used by Klein et al. in [1], the 2088 registrations took at total of 22 hours on our system. As a check we also ran the recent version 1.9.x of the ANTS-SyN diffeomorphic Symmetric-Normalization program [5] on the these datasets. Klein et al. had found SyN to be one of the best performing programs. Fig. 4 and Table III compare the

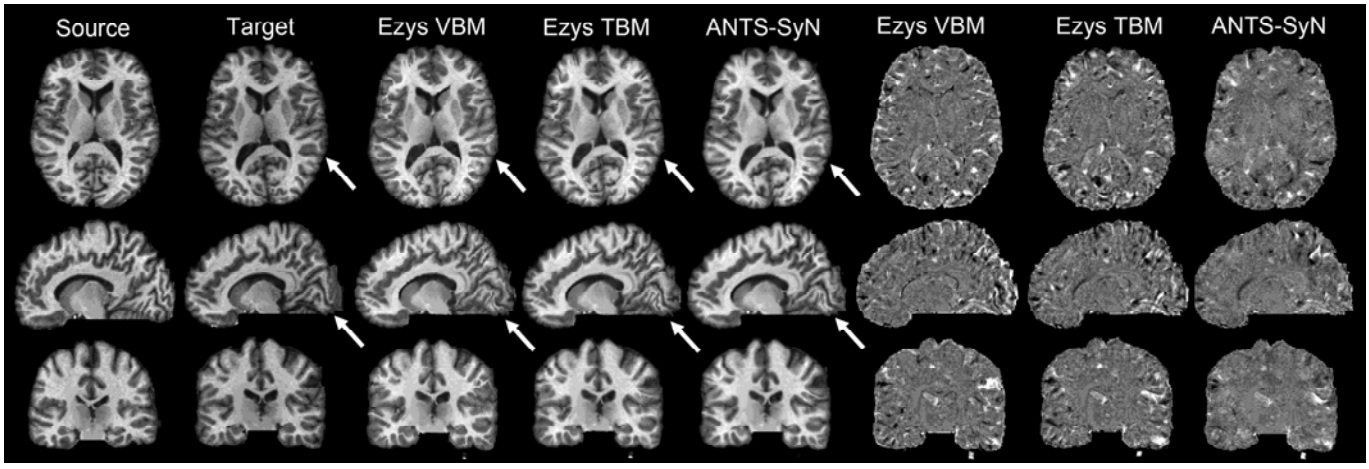

Fig. 2**.** Example of pairwise registration from VBM study, the rows show representative axial, sagittal and coronal slices for MGH subject 1 (source in column 1) registered to MGH subject 5 (target in column 2). Columns 3 to 5 show the registration results which should match column 2. Columns 3 and 5 show the Ezys - VBM and ANTS-SyN as used for this study. Column 4 shows the result for Ezys run in symmetric TBM mode for comparison. Columns 6 to 8 show the corresponding difference images obtained by subtracting the target image in column 2 from those in columns 3-5. The arrows point to features where the source and target differ.

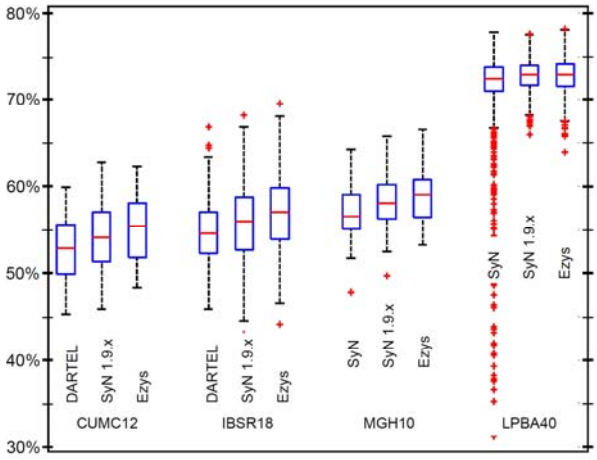

Fig. 4. Box and whisker plots comparing the distributions of mean target overlap (as percentages) for the sets of all possible pairwise registrations for each of the four datasets. In each case the best result from Figure 5 of Klein et al. (either ANTS-SyN or DARTEL on the left) is compared to the Ezys result (right) and our own reprocessing with a more recent version of ANTS-SyN (version 1.9.x centre).

TABLE III
MEDIAN TARGET OVERLAP VALUES IN PERCENT

Program	CUMC12	IBSR18	MGH10	LPBA40
Ezys	55.41±0.32	56.99±0.24	59.03±0.31	72.79±0.05
ANTS-SyN 1.9.x	54.07±0.32	55.92±0.24	58.05±0.32	72.79±0.04
ANTS-SyN (Klein)	51.46±0.29	52.98±0.24	56.52±0.31	72.48±0.13
DARTEL (Klein)	52.86±0.29	54.63±0.20	55.97±0.57	58.71±0.12

Median values (in percent) for data shown in Fig. 4. The values for Ezys and ANTS-SyN 1.9.x were calculated by us and the other values were reported by Klein et al. and obtained from their web page http://mindboggle.info/papers/ evaluation_NeuroImage2009.php. The errors are the standard errors on the disruption means.

performance of the best non-linear registration algorithms tested by Klein et al., (i.e. ANTS-SyN and DARTEL) the latest SyN implementation and Ezys for the four datasets.

Ezys performs significantly better than the algorithms tested by Klein et al. and matches or exceeds the performance of the newer version of ANTS-SyN which also performs significantly better than the version tested by Klein et al., the median overlaps are larger and the tails for small values are reduced. Ezys also required much less processing time. In Table IV we show the Ezys target overlaps for each of the individual labelled regions averaged over all 1560 pairwise registrations performed for the LPBA40 dataset. These are compared to the corresponding results for the ANTS-SyN program as tested by Klein et al. The differences are well correlated between left and right hemispheres suggesting that genuine anatomical structures are being co-registered. Similar results were observed for the other three datasets and more recent version of ANTS-SyN and Ezys have fewer differences. We also assessed the performance of NiftyReg - a recent GPU based program not tested by Klein et al. - on the MGH10 dataset and obtained a median TO_2 of 57.51% using, after some optimisation tests, the default $BE = 0.005$ and minimum spline spacing $sx = sy = sz = 3.5$ mm.

C. *TBM Analysis of Atrophy*

Fig. 2 shows Jacobian maps for one pair of images from an SD subject computed using Ezys, DARTEL, FSL5-FNRIT and NiftyReg. As it would be expected in SD, temporal lobe atrophy and ventricular expansion are the main features. Fig. 5 shows tissue atrophy rate maps for two control, two AD and two SD subjects. For each subject these maps were computed using Jacobians from all possible pairwise registrations and assuming a constant atrophy rate as detailed in section II.I. To aid visual comparison the maps are shown warped to a common reference image space, (that of an elderly control subject). Serial changes - predominantly in the white matter - were detected in dementia patients with clearly recognizable patterns of progressive atrophy - widespread temporo-parietal regions in AD, and more focal temporal lobe atrophy in SD. Elderly controls, in contrast, showed less widespread and overall lower rates of progressive grey matter loss.

Scatter plots of average grey matter atrophy rates in temporal and parietal lobes are shown in Fig. 6 for all 18 subjects using the four registration programs: Ezys, DARTEL, NiftyReg and FSL5 FNIRT. Ezys yielded the best separation between control, AD and SD groups. Note that with Ezys, the six subjects with the most progressive parietal lobe atrophy were all the AD cases, and the six subjects with the highest rate of temporal lobe atrophy were all SD. It was also noted that the SD patient with the most "stable" temporal lobe had the most severely lesioned hippocampi (data not shown). Note that for all four approaches we used images from all time points symmetrically hence avoiding possible biases caused by preferential treatment of baseline images. Note also that we chose to present volume change rates for these two grey matter areas because they are known to be differentially affected in AD and SD; this was confirmed by the results.

IV. DISCUSSION

The purpose of this paper is to introduce our new GPU based image registration program Ezys. We have shown that for a number of neuroscience applications Ezys performs as well as the best existing programs and is faster.

The use of surface information, particularly corresponding positions of sulci and gyri is not new, but may depend on manual identification of landmarks [35] or prior segmentation [36, 37], additionally several software tools are available for segmentation including Freesurfer [38], BrainSuite [39] BrainVoyager [40] and BrainVisa [41]. In contrast the Ezys surface-preserving regularization methods are fully integrated into the registration procedure and use algorithms for local surface preservation which are particularly suitable for GPU implementation.

TABLE IV
REGION OVERLAPS FOR LPBA40

LPBA40 Region	Ezys		SyN		Difference	
	L	R	L	R	L	R
superior frontal gyrus	82.7	82.2	81.5	80.7	1.3	1.5
middle frontal gyrus	80.8	80.3	79.2	78.4	1.7	2.0
inferior frontal gyrus	75.2	75.2	73.6	72.9	1.7	2.3
precentral gyrus	73.1	74.2	74.2	73.3	**-1.1**	0.8
middle orbitofrontal gyrus	71.0	70.5	68.6	68.7	2.4	1.9
lateral orbitofrontal gyrus	62.9	60.2	60.7	59.1	2.2	1.1
gyrus rectus	68.8	70.3	68.9	71.2	**-0.1**	**-0.9**
postcentral gyrus	64.5	67.4	68.1	68.4	**-3.6**	**-0.9**
superior parietal gyrus	72.9	73.1	73.4	72.5	**-0.5**	0.6
supramarginal gyrus	66.3	66.4	67.6	66.5	**-1.3**	**-0.1**
angular gyrus	65.8	67.6	64.2	65.9	1.6	1.7
precuneus	69.1	69.4	67.4	68.5	1.7	0.9
superior occipital gyrus	61.6	59.1	60.6	58.8	1.0	0.3
middle occipital gyrus	70.4	69.9	68.9	68.3	1.5	1.6
inferior occipital gyrus	69.6	70.3	68.1	68.0	1.5	2.3
cuneus	67.4	66.7	66.9	66.1	0.5	0.5
superior temporal gyrus	77.3	78.3	77.0	77.8	0.3	0.5
middle temporal gyrus	67.4	70.7	67.8	70.6	**-0.4**	0.1
inferior temporal gyrus	69.4	70.9	67.2	68.4	2.1	2.5
parahippocampal gyrus	72.6	71.5	72.0	70.8	0.6	0.7
lingual gyrus	73.1	74.8	72.7	74.6	0.5	0.2
fusiform gyrus	71.7	72.5	70.9	70.8	0.9	1.6
insular cortex	84.3	82.5	79.1	77.6	5.2	5.0
cingulate gyrus	73.7	73.2	72.0	71.1	1.7	2.1
caudate	81.3	80.2	74.9	74.5	6.4	5.7
putamen	82.0	80.9	77.8	77.6	4.2	3.4
hippocampus	77.8	77.8	74.8	75.2	3.1	2.6
cerebellum	86.4		87.6		**-1.3**	
brainstem	81.3		79.9		1.4	

Average region overlaps, in percent, for 1560 pairwise registrations on LPBA40 dataset. The Ezys values for the left and right hemispheres are in columns 2 and 3, the corresponding results for SyN from Klein et al.'s analysis are in columns 4 and 5. The differences between Ezys and SyN are in columns 6 and 7, negative values (where SyN has outperformed Ezys) are shown in bold. The last two rows have undivided structures.

Using one displacement vector per voxel rather than, say B-spline interpolation, helps detect small displacements near surfaces while the surface preserving regularization ensures smooth transformations within individual anatomical regions without blurring boundaries. This effectively allows us to have more degrees of freedom where they are needed typically near boundaries between white and grey matter and CSF, while constraining displacements elsewhere. Ezys' high quality Jacobian map shown in Fig. 3 is a direct result of this feature. In a control experiment with pseudo ground truths, we showed an improvement over established registration methods such as FSL-FNIRT, NiftyReg, ANTS-SyN and DARTEL (see Fig. 4 and Table III) at a fraction of the time (Table II). We also demonstrated that Ezys can quantify subtle anatomical changes in longitudinal MRI datasets of single dementia patients. Invertible Jacobian determinant values at each voxel (for all available pairwise intra-subject registrations) were used to generate TBM maps that measured annual rates of brain tissue volume change (see Fig. 5). In patients, extensive patterns of atrophy were observed for all four algorithms. All maps resulted in consistent - with prior knowledge [42, 43] - patterns of longitudinal change in single patients and relatively stable behaviours for control subjects. The quantitative scatter plots shown in Fig. 6 revealed total separation between patients and controls for Ezys and DARTEL, whereas group differences for other methods were less striking. The distribution of AD and SD atrophy rates also differed considerably: as it was expected, SD patients presented with fast rates of temporal lobe atrophy but parietal lobes remained relatively stable; whereas AD subjects showed a simultaneous increase in temporal and parietal lobe shrinkage rates. Analogous plots for the other two algorithms in Fig. 6 showed comparable qualitative behaviours, but the distributions were more diffused and group results largely overlapped.

Ezys supports both direct source to target registration and a form of symmetric registration as discussed in section II.C. We found direct registration gave the best results in our inter-subject VBM study whereas symmetric registration was better for the intra-subject TBM study.

The internal data representation used by Ezys stores displacements of each voxel directly. Using such a simple internal representation allows the calculation displacement field compositions efficiently, this operation is more complex using B-spline based displacement fields. Ezys updates displacement fields at each iteration using displacement composition rather than addition, thus ensuring an invertible and diffeomorphic displacement field. In addition, smoothness is maintained by firstly applying a Gaussian filter to the gradients along each axis, which are approximated as three moving averages, and secondly by smoothing the displacement fields themselves. The smoothing of displacement fields is done by simulating diffusion constrained by surfaces, which is also implemented by using co-axial moving averages. These smoothing steps can very efficiently implemented on the GPU. In our implementation most computationally expensive tasks such as similarity evaluations, gradients, Gaussian filtering, regularizations and other image processing routines are implemented on the GPU. while other less performance critical parts of the algorithm are written in C++ and run on host CPU. Ezys has both graphical and command line user interfaces and can process long script files. Thus Ezys can easily be integrated with complex workflows. A GPU is required to run Ezys, but these are relatively inexpensive and indeed Ezys can run satisfactorily on suitably equipped laptops. The present version of Ezys requires a CUDA capable GPU of compute capability 2.0 or higher with a minimum of 1.5 GB of on board ram for typical applications. The results presented here used a GeForce GTX 580 and we have also obtained good results with a more recent GeForce GTX 760.

Ezys has potential applications in other areas of medical imaging, for example inter-modality imaging such as MR-CT or MR-PET. The accuracy and speed of would also be helpful in applications such as treatment planning for radiotherapy [44].

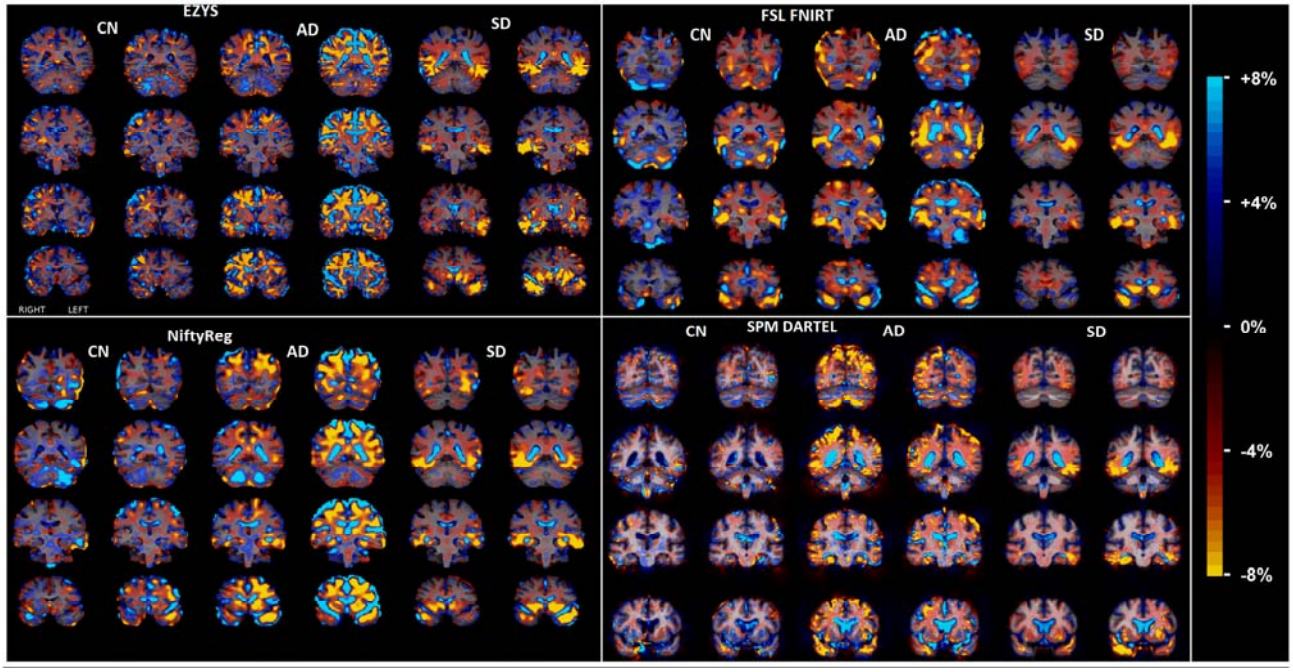

Fig. 5 Atrophy rates in percent per year for 2 controls (CN), 2 Alzheimer's disease (AD) and 2 semantic dementia (SD) subjects shown in a common reference frame, (that of a control subject). Results for Ezys, FSL-FNIRT and NiftyReg are shown on a common-frame, however in case of DARTEL they are shown in a frame of the template image produced by co-registration.

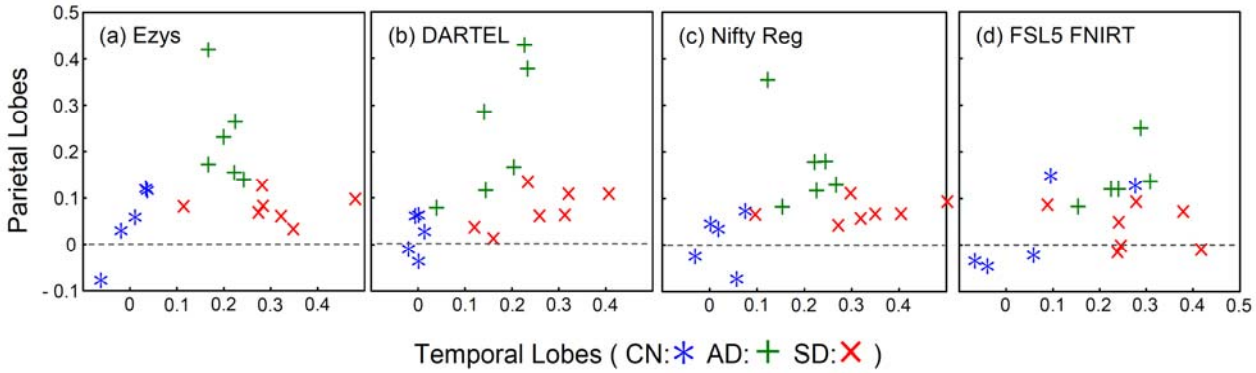

Fig. 6: Scatter plots of temporal versus parietal lobe grey matter atrophy rates in percent per annum for 18 elderly subjects, comprising 5 normal controls (CN), 6 Alzheimer's (AD) and 7 semantic dementia (SD) patients. Atrophy rates were derived using tensor based morphometry. Results from Ezys and three other programs are compared.

V. Conclusions

Ezys is a fast and accurate image registration program suitable for both inter-subject and intra-subject studies. We have shown both accurate anatomical label propagation between subjects and detection of disease-specific, longitudinal brain volume changes in single patients.

The image regularization and smoothing algorithms presented here are somewhat heuristic and could perhaps be further improved, however they do demonstrate that there is scope for developing novel algorithms targeted at SIMD architectures having a very large number of processing cores such as GPUs.

Ezys is freely available, the program version used for this work together with supporting material can be downloaded from: http://www.bss.phy.cam.ac.uk/~rea1/ezys.

References

[1] A. Klein, J. Andersson, B. A. Ardekani, J. Ashburner, B. Avants, M.-C. Chiang, G. E. Christensen, D. L. Collins, J. Gee, P. Hellier, J. H. Song, M. Jenkinson, C. Lepage, D. Rueckert, P. Thompson, T. Vercauteren, R. P. Woods, J. J. Mann, and R. V. Parsey, "Evaluation of 14 nonlinear deformation algorithms applied to

human brain MRI registration," *Neuroimage*, vol. 46, pp. 786-802, Jul 1 2009.

[2] R. E. Ansorge, S. J. Sawiak, and G. B. Williams, "Exceptionally Fast non-linear 3D Image Registration using GPUs," in *2009 Ieee Nuclear Science Symposium Conference Record, Vols 1-5*, B. Yu, Ed., ed, 2009, pp. 3207-3213.

[3] M. Modat, G. R. Ridgway, Z. A. Taylor, M. Lehmann, J. Barnes, D. J. Hawkes, N. C. Fox, and S. Ourselin, "Fast free-form deformation using graphics processing units," *Computer Methods and Programs in Biomedicine*, vol. 98, pp. 278-284, Jun 2010.

[4] K. O. Noe, K. Tanderup, J. C. Lindegaard, C. Grau, and T. S. Sorensen, "GPU Accelerated Viscous-fluid Deformable Registration for Radiotherapy," *Medicine Meets Virtual Reality 16*, vol. 132, pp. 327-332, 2008 2008.

[5] B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," *Medical Image Analysis*, vol. 12, pp. 26-41, Feb 2008.

[6] J. Ashburner and K. J. Friston, "Voxel-based morphometry - The methods," *Neuroimage*, vol. 11, pp. 805-821, Jun 2000.

[7] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith, "FSL," *Neuroimage*, vol. 62, pp. 782-790, Aug 15 2012.

[8] J. Ashburner, "A fast diffeomorphic image registration algorithm," *Neuroimage*, vol. 38, pp. 95-113, Oct 2007.

[9] J. Ashburner, "SPM: A history," *Neuroimage*, vol. 62, pp. 791-800, Aug 2012.

[10] P. A. Freeborough and N. C. Fox, "The boundary shift integral: An accurate and robust measure of cerebral volume changes from registered repeat MRI," *Ieee Transactions on Medical Imaging*, vol. 16, pp. 623-629, Oct 1997.

[11] S. M. Smith, Y. 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*, vol. 17, pp. 479-489, Sep 2002.

[12] R. S. J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. A. Zeki, J., and W. D. Penny, *Human Brain Function*, 2nd ed.: Academic Press, 2003.

[13] C. Gaser, H. P. Volz, S. Kiebel, S. Riehemann, and H. Sauer, "Detecting structural changes in whole brain based on nonlinear deformations - Application to schizophrenia research," *Neuroimage*, vol. 10, pp. 107-113, Aug 1999.

[14] X. Hua, A. D. Leow, N. Parikshak, S. Lee, M. C. Chiang, A. W. Toga, C. R. Jack, M. W. Weiner, P. M. Thompson, and I. Alzheimer's Dis Neuroimaging, "Tensor-based morphometry as a neuroimaging biomarker for Alzheimer's disease: An MRI study of 676 AD, MCI, and normal subjects," *Neuroimage*, vol. 43, pp. 458-469, Nov 2008.

[15] C. Studholme, V. Cardenas, R. Blumenfeld, N. Schuff, H. J. Rosen, B. Miller, and M. Weiner, "Deformation tensor morphometry of sematic dementia with quantitive validation," *Neuroimage*, vol. 21, pp. 1387-1398, Apr 2004.

[16] N. Lepore, C. Brun, Y. Y. Chou, M. C. Chiang, R. A. Dutton, K. M. Hayashi, E. Luders, O. L. Lopez, H. J. Aizenstein, A. W. Toga, J. T. Becker, and P. M. Thompson, "Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors," *Ieee Transactions on Medical Imaging*, vol. 27, pp. 129-141, Jan 2008.

[17] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, "Automatic 3D Intersubject Registration of MR Volumetric Data in Standardized Talairach Space," *Journal of Computer Assisted Tomography*, vol. 18, pp. 192-205, Mar-Apr 1994.

[18] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recognition*, vol. 32, pp. 71-86, Jan 1999.

[19] J. P. Thirion, "Non-rigid matching using demons," in *1996 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Proceedings*, ed, 1996, pp. 245-251.

[20] N. C. Fox, G. R. Ridgway, and J. M. Schott, "Algorithms, atrophy and Alzheimer's disease: Cautionary tales for clinical trials," *Neuroimage*, vol. 57, pp. 15-18, Jul 1 2011.

[21] W. K. Thompson, D. Holland, and I. Alzheimer's Dis Neuroimaging, "Bias in tensor based morphometry Stat-ROI measures may result in unrealistic power estimates," *Neuroimage*, vol. 57, pp. 1-4, Jul 1 2011.

[22] A. Gruslys, S. Sawiak, and R. Ansorge, "3000 Non-Rigid Medical Image Registrations Overnight on a Single PC," in *2011 IEEE Nuclear Science Symposium and Medical Imaging Conference*, ed, 2011, pp. 3073-3080.

[23] P. Hellier and C. Barillot, "Multimodal non-rigid warping for correction of distortions in functional MRI," in *Medical Image Computing and Computer-Assisted Intervention - Miccai 2000*. vol. 1935, S. Delp, A. M. DiGioia, and B. Jaramaz, Eds., ed, 2000, pp. 512-520.

[24] V. Koivunen and J. M. Vezien, "Machine vision tools for CAGD," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 10, pp. 165-182, Mar 1996.

[25] P. Perona and J. Malik, "Scale-Space And Edge-Detection Using Anisotropic Diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 629-639, Jul 1990.

[26] G. Hermosillo, C. Chefd'Hotel, and O. Faugeras, "Variational methods for multimodal image matching," *International Journal of Computer Vision*, vol. 50, pp. 329-343, Dec 2002.

[27] R. Stefanescu, X. Pennec, and N. Ayache, "Grid powered nonlinear image registration with locally adaptive regularization," *Medical Image Analysis*, vol. 8, pp. 325-342, Sep 2004.

[28] I. J. Simpson, M. W. Woolrich, M. J. Cardoso, D. M. Cash, M. Modat, J. A. Schnabel, and S. Ourselin, "A Bayesian Approach for Spatially Adaptive Regularisation in Non-rigid Registration," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, ed: Springer, 2013, pp. 10-18.

[29] J. Acosta-Cabronero, K. Patterson, T. D. Fryer, J. R. Hodges, G. Pengas, G. B. Williams, and P. J. Nestor, "Atrophy, hypometabolism and white matter abnormalities in semantic dementia tell a coherent story," *Brain*, vol. 134, pp. 2025-2035, Jul 2011.

[30] J. Acosta-Cabronero, G. B. Williams, G. Pengas, and P. J. Nestor, "Absolute diffusivities define the landscape of white matter degeneration in Alzheimer's disease," *Brain*, vol. 133, pp. 529-539, Feb 2010.

[31] H. M. Duvernoy, "Vascularization of the cerebral cortex," *Revue Neurologique*, vol. 155, pp. 684-687, 1999 1999.

[32] C. Rorden. *MRIcro Software*. Available: http://www.mccauslandcenter.sc.edu/mricro/

[33] J. Ashburner and K. J. Friston, "Unified segmentation," *Neuroimage*, vol. 26, pp. 839-851, Jul 2005.

[34] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster, "A Probabilistic Atlas Of The Human Brain - Theory And Rationale For Its Development," *Neuroimage*, vol. 2, pp. 89-101, Jun 1995.

[35] A. A. Joshi, D. W. Shattuck, P. M. Thompson, and R. M. Leahy, "Surface-constrained volumetric brain registration using harmonic mappings," *Medical Imaging, IEEE Transactions on*, vol. 26, pp. 1657-1669, 2007.

[36] G. Postelnicu, L. Zollei, and B. Fischl, "Combined volumetric and surface registration," *Medical Imaging, IEEE Transactions on*, vol. 28, pp. 508-522, 2009.

[37] D. Shen and C. Davatzikos, "HAMMER: hierarchical attribute matching mechanism for elastic registration," *Medical Imaging, IEEE Transactions on*, vol. 21, pp. 1421-1439, 2002.

[38] B. Fischl, "FreeSurfer," *Neuroimage*, vol. 62, pp. 774-781, Aug 15 2012.

[39] D. W. Shattuck and R. M. Leahy, "BrainSuite: An automated cortical surface identification tool," *Medical Image Analysis*, vol. 6, pp. 129-142, Jun 2002.

[40] R. Goebel, F. Esposito, and E. Formisano, "Analysis of Functional Image Analysis Contest (FIAC) data with BrainVoyager QX: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis," *Human Brain Mapping*, vol. 27, pp. 392-401, May 2006.

[41] D. Rivière, D. Geffroy, I. Denghien, N. Souedet, and Y. Cointepas, "BrainVISA: an extensible software environment for sharing multimodal neuroimaging data and processing tools," *Neuroimage*, vol. 47, p. S163, 2009.

[42] P. Bright, H. E. Moss, E. A. Stamatakis, and L. K. Tyler, "Longitudinal studies of semantic dementia: The relationship between structural and functional changes over time," *Neuropsychologia*, vol. 46, pp. 2177-2188, 2008 2008.

[43] G. Chetelat, B. Landeau, F. Eustache, F. Mezenge, F. Viader, V. de la Sayette, B. Desgranges, and J. C. Baron, "Using voxel-based morphometry to map the structural changes associated with rapid conversion in MCI: A longitudinal MRI study," *Neuroimage*, vol. 27, pp. 934-946, Oct 1 2005.

[44] D. Ciardo, M. Peroni, M. Riboldi, D. Alterio, G. Baroni, and R. Orecchia, "The Role of Regularization in Deformable Image Registration for Head and Neck Adaptive Radiotherapy," *Technology in Cancer Research & Treatment*, vol. 12, pp. 323-331, Aug 2013.