Co-registration of high resolution MRI sub-volumes in non-human primates

Jérémy Lecoeur, Feng Wang, Li Min Chen, Rui Li, Malcolm J. Avison, Benoit M. Dawant

Co-registration of high resolution MRI sub-volumes in non-human primates

Jérémy Lecoeur[a], Feng Wang[b], Li Min Chen[b], Rui Li[a], Malcolm J. Avison[b], Benoit M. Dawant[a]

[a] Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA

[b] Department of Radiology and Radiological Science, Vanderbilt University Medical Center, Nashville, TN 37232, USA

ABSTRACT

Dynamic structural and functional remodeling of the Central Nervous System occurs throughout the lifespan of the organism from the molecular to the systems level. MRI offers several advantages to observe this phenomenon: it is non-invasive and non-destructive, the contrast can be tuned to interrogate different tissue properties and imaging resolution can range from cortical columns to whole brain networks in the same session. To measure these changes reliably, functional maps generated over time with high resolution fMRI need to be registered accurately. This article presents a new method for registering automatically thin cortical MR volumes that are aligned with the functional maps. These acquisitions focus on the primary somato-sensory cortex, a region in the anterior parietal part of the brain, responsible for fine touch and proprioception. Currently, these slabs are acquired in approximately the same orientation from acquisition to acquisition and then registered by hand. Because they only cover a small portion of the cortex, their direct automatic registration is difficult. To address this issue, we propose a method relying on an intermediate image, acquired with a surface coil that covers a larger portion of the head to which the slabs can be registered. Because images acquired with surface coils suffer from severe intensity attenuation artifact, we also propose a method to register these. The results from data sets obtained with three squirrel monkeys show a registration accuracy of thirty micrometers.

Keywords: Rigid registration, longitudinal studies, attenuation artifacts in MR, primates, functional studies.

1. INTRODUCTION

There is compelling evidence that plastic changes of the brain are a key etiologic factor in the emergence of a broad range of neurological, psychiatric and behavioral disorders. The development of successful therapies to remediate these pathological plastic changes requires, among other things, the availability of robust quantitative measures of plasticity that can link the spatio-temporal trajectories of cellular and molecular events in vitro to Central Nervous System (CNS) remodeling in vivo in animals, and ultimately in humans. The enhancement of existing and the development of new MRI methods for measuring CNS plasticity, together with tools that allow the MRI observations in animal models to be linked back to molecular and cellular events on the one hand, and to plastic changes in the human CNS on the other, will provide a methodological bridge across scales and species. Central to the success of such an enterprise is the ability to achieve stable, unbiased and ideally fully automated coregistration of structural images collected from the same subject longitudinally over time. The fidelity of this coregistration ultimately constrains the sensitivity with which MRI can detect changes in local structure and/or function associated with normal and pathological brain remodeling in vivo.

The goal of this study is to develop an accurate and reliable registration process for longitudinal acquisitions of MRI scans providing partial brain coverage at very high resolution in order to quantify functional and structural changes across time. The scans used in this study generally cover a thin slab (~ 8 mm) close to the surface of the brain. These slabs are centered on the somato-sensory cortex, a region of the cortex corresponding to the areas 1, 2 and 3 of Brodmann's nomenclature [1] and responsible for the sense of fine touch and proprioception of hand and face [2]. As these areas are quite small (between 2 and 3 mm^2) and our ultimate objective is to measure temporal changes in these images, a very accurate registration process is needed.

In current practice, the position of a slab is prescribed manually with the help of a user interface provided by the scanner manufacturer, using the lower resolution larger field of view images for guidance. This is a somewhat unreliable process and its accuracy largely depends on the skills of the operator. In general, however, the slabs acquired during different sessions, possibly separated by long periods of time, will not cover exactly the same brain volume nor are they

acquired at exactly the same angle. It is therefore necessary for slabs to be registered retrospectively. Currently, coregistration of overlapping but non-identical slabs acquired in different sessions is performed manually and this process introduces another source of inaccuracy. The goal of the present study was to develop and appraise the performance of an accurate, reliable, and automated process for coregistering MRI scans acquired longitudinally from the same subject in multiple imaging sessions. In particular, we sought to achieve high fidelity coregistration of very high resolution structural image sets providing partial brain coverage (referred to here as "slabs") acquired using surface coil transceiver coils whose image intensities are spatially inhomogeneous.

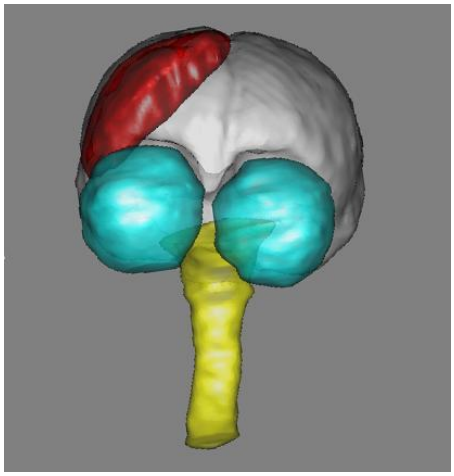

Figure 1 – 3D rendering of brain (white), eyes (cyan), brain stem and spine (yellow) and slab (red).

A large body of work dealing with the longitudinal registration of both human and animal head scans has been published over the years and a complete review of these methods (see for instance [3]) would be outside the scope of this article. However, these methods are generally applied to whole brain images or to slabs covering an entire cross-section of the head ([4-6]). In the present work we address a more general problem, that of coregistering very high resolution images of local brain areas collected from the same subject at different times. These high resolution image sets generally comprise sub-volumes (or "slabs") within a larger volume of brain imaged at lower spatial resolution. Figure 1 illustrates the typical position, orientation, and coverage of a slab. The unique nature of our data set and the accuracy requirements for our long term goals necessitated the development of a new registration procedure, which is described and evaluated in this article.

The next section describes the data we have used in this study. Section III provides technical information on our registration procedure and on the validation technique we have used. Results are presented in Section IV. Our conclusions and directions for future work are discussed in Section V.

2. IMAGES AND ACQUISITION PROTOCOL

The MR images that were used were acquired from three male squirrel monkeys with age ranging from 4 to 5 years old, and body weight ranging from 800 g to 1100 g. The animals were scanned while anesthetized to reduce artifacts caused by intra-session motion. The induction anesthesia compounds were ketamine (10 – 30 mg/kg) and atropine (0.03 - 0.08 mg/kg) and anesthesia is then maintained by isoflurane (0.75 - 2%) with a NO2/O2 ratio of 7/3. Two data sets were acquired during each session: (1) an oblique scan, which covers a small portion of the cortex and is referred to as a slab, and (2) a head scan acquired that covers a larger part of the brain, which is referred to as a partial head scan. Acquisition details for these two data sets are detailed in the following two subsections.

2.1. Slabs

These images were acquired on a 9.4 Tesla Varian MRI using a 3 cm transmit/receive surface coil. They were oriented obliquely with respect to the scanner frame of reference. As discussed above, the coil is positioned manually and the acquisition parameters, which permit to select the slab, e.g., its orientation, are also set manually during each session using a software interface provided by the scanner manufacturer. A slab consists of 16 slices with 512 × 512 voxels each. The in-plane pixel resolution is 68.35 × 68.35 μm^2; the slice thickness is 0.5 mm. The complete slab

thickness is thus 8 mm and the field of view (FoV) is 36 × 36 mm^2. These images were acquired with a repetition time (TR) of 400 milliseconds and an echo time (TE) of 16 milliseconds with a flip angle of 25 degrees. Figure 2 shows a typical slab in three orthogonal directions.

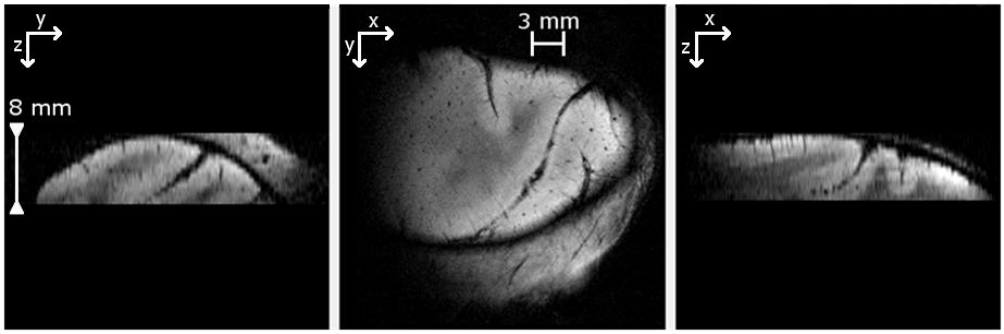

Figure 2 - Three orthogonal views of a slab.

2.2. Partial 3D Head Scans

These images were acquired with the same 3 cm surface coil used to acquire the slabs. The slab and the partial 3D head scans were acquired during the same session. Nor the animal nor the coil is displaced during acquisition. The partial 3D head scans are 128 × 128 × 128 voxels isotropic acquisitions with voxel dimension of 0.5 × 0.5 × 0.5 mm^3 (TR = 5 milliseconds, TE = 2.39 milliseconds with a flip angle of 15 degrees and FoV = 64 × 64 mm^2). As shown in Figure 3, these images are affected by severe signal attenuation artifacts as one moves away from the surface coil. Note that signal is present even far from the coil but the large intensity difference makes it difficult to adjust the contrast to permit the visualization of anatomical details in both high intensity and low intensity regions. Because only a portion of the head is visible in these images with one visualization parameter setting, they are referred to as partial 3D head scans.

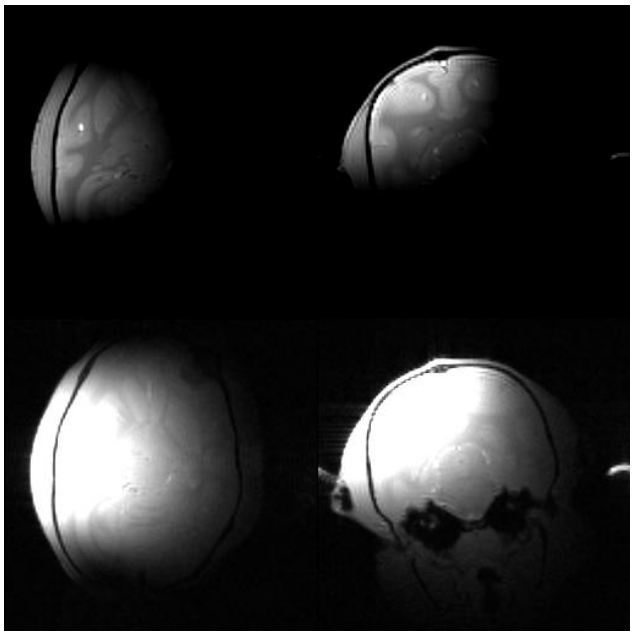

Figure 3 –Partial 3D head MRI scans. Because the signal drops rapidly from the coil, the contrast needs to be adjusted to permit visualization. Bottom row: two views with contrast adjusted in order to show details in regions of low intensity. Top row: Same images with contrast adjusted to show details in regions of high intensity. When the contrast is adjusted for high intensity regions, the rest of the head is not visible. When it is adjusted for low intensity regions, the area close to the coil is saturated.

3. METHODOLOGY

In this work, we assume that geometric distortion is minimal. With the high bandwidth acquisition sequences we use, this assumption is reasonable. Based on this assumption, all the transformations needed to register the various scans are rigid-body transformations, i.e., they only involve a translation vector $[t_x, t_y, t_z]$ and a rotation matrix $R(\varphi, \theta, \psi)$. The easiest approach to compute these values would be to register directly slabs acquired during different sessions to each other with a standard intensity-based algorithm. This approach was tried but, as shown in the results section, did not lead to satisfactory results.

The solution we have developed to address this issue involves two main steps: (1) register the slab to the partial 3D volume acquired during the same session and (2) register the partial 3D volumes acquired longitudinally. As will be discussed, several intermediate steps are also required to achieve the desired accuracy.

3.1. Slab to partial 3D head scan

Since the animal is anesthetized during acquisition, motion between the slab and the 3D scan can be assumed to be minimal. The transformation that registers these scans can thus be computed analytically using image coordinates and scanning information retrieved from the image headers (Euler angles used to acquire the scan and center of the volume). In principle, this approach should lead to an accurate registration but we have observed a small residual error, which we attribute to the hardware. To correct for this residual error, we use an in-house implementation of an intensity-based registration algorithm, which we call AMIR. This algorithm computes the transformation that maximizes the Normalized Mutual Information (NMI) [7] between the image volumes. Mutual Information (MI) and NMI have been used extensively as similarity measures for the registration of single and multi-modality images since the validation study conducted by Fitzpatrick et al. [8]. NMI has been documented as being more robust than MI for partially overlapping volumes [7]. NMI is defined as

$$\mathrm{NMI}(A;B) = \frac{H(A) + H(B)}{H(A,B)}$$

Where H(A) is the marginal entropy of image A, H(B) is the marginal entropy of image B and H(A,B) is the joint entropy of images A and B [9]. The marginal entropy of an image is computed from its intensity probability density function as:

$$H(A) = -\sum_{i \in A} p(i) \log[p(i)]$$

in which *i* is an intensity and *p(i)* is the probability of observing intensity *i* in the image. The intensity probability density function is estimated from the normalized intensity histogram. Similarly, the joint entropy of two images is computed as

$$H(A,B) = -\sum_{i \in A}\sum_{j \in B} p(i,j) \log[p(i,j)]$$

in which *p(i,j)* is the joint probability density function estimated from the normalized joint intensity histogram. Intensity histograms were built with 128 intensity bins.

3.2. Partial head scans co-registration

The easiest approach to register the inter-session partial 3D head scans is to apply an intensity based registration algorithm to these image volumes. However, as discussed previously these images are affected by severe intensity variation artifacts. These artifacts are challenging for intensity-based registration methods and, as shown in the results section, registration of the unprocessed volumes led to poor results. Methods have been proposed in the literature to correct for intensity variation (see for instance [10-14]). However, these methods were developed and are typically used for whole human head images acquired at 1.5 or 3T. One popular intensity correction algorithm is the N3 algorithm developed by Sled et al. [12]. This algorithm was applied to our images but without success. Rather than attempting to develop an intensity correction algorithm for the type of images used in the study, we opted to work with the gradient images. Others [15-17] have proposed to combine gradient and intensity images for intensity-based registration. Here, working with the gradient images alone led to accurate results for two reasons. First, sulci are clearly visible in the images and their edges provide strong registration features. Second, the rapid intensity drop-off visible in the images is

attenuated in the gradient image. Figure 4 shows illustrative examples of gradient magnitude images (the gradient images were computed with a standard 3D Sobel operator [18] as implemented in ITK [19]).

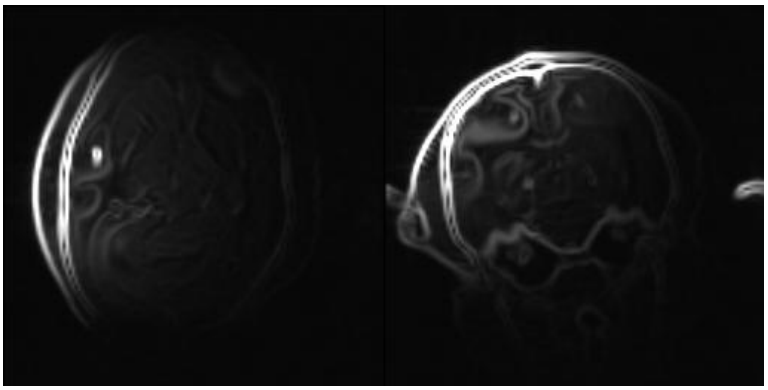

Figure 4 – Representative 3D partial head gradient magnitude images obtained with a 3D Sobel operator.

The bottom panel of Figure 5 shows the intensity profiles in the original and in the gradient magnitude images along the red line shown on the top panel. The procedure that was followed was thus to first compute the gradient magnitude of both images and then to register these with the same intensity-based registration algorithm that is used to register the slabs to their corresponding partial head volume. NMI was also used as the similarity measure. The intensity histograms were built with 128 bins.

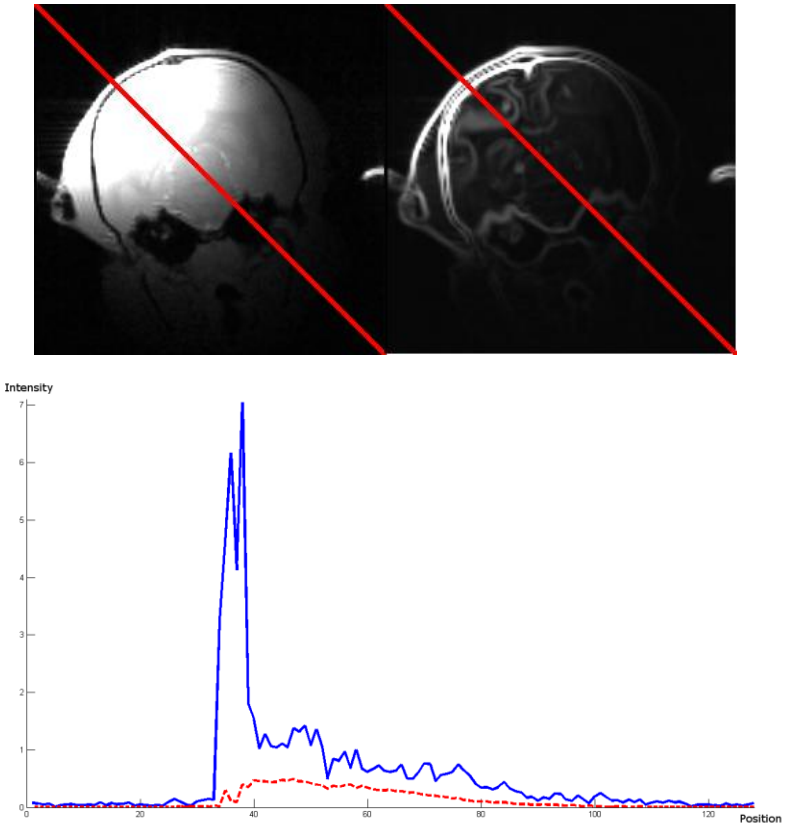

Figure 5: Top panel, original and gradient magnitude for one 3D head image. Bottom panel, intensity profiles along the line shown on the top panel images. Solid blue: intensity profile for the original partial 3D head scan. Dotted red: intensity profile for the gradient magnitude image.

3.3. Final refinement

Even though both the slab to partial 3D and the 3D to 3D steps produce results that are visually correct, the overall slab to slab registration obtained via the 3D slabs is not accurate enough. This is due to the difference in spatial resolution between the volumes. Indeed, the in-plane resolution of the partial head volumes is almost ten times lower than the resolution of the slabs. Slab-to-3D or 3D-to-3D registrations may appear accurate when evaluated with the 3D images but not when slabs are compared. Fortunately, the transformation obtained with the process described so far is close enough to the correct transformation to permit a final intensity-based registration of the slabs. Again, NMI was used as the similarity measure and the intensity histograms were built with 128 intensity bins.

3.4. Overall registration procedure

Figure 6 summarizes the various steps involved in the procedure we have developed. The first transformation is the concatenation of the transformations computed with the scanner information and with the intensity-based registration of a slab with its corresponding partial head volume. The second transformation is estimated by the intensity-based registration of the gradient magnitude images. T3 is obtained with scanner information and with the intensity-based registration. The last transformation (T4) is estimated by the intensity-based registration of the slabs.

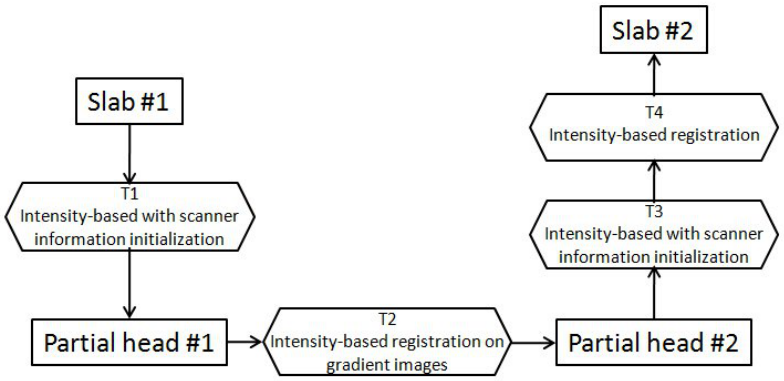

Figure 6: Overall process used in this study to register slabs acquired during two imaging sessions.

3.5. Evaluation of accuracy

Assessing the accuracy of automatic registration algorithms is a difficult task. When a ground truth transformation is known, the estimated transformation can be compared to it. In this study, no ground truth is available and a validation procedure based on visible landmarks has been used. The small dark circles visible in the image shown on the middle panel of figure 2 are trans-cortical vessels. Because these are easily visible and because they typically cover the entire image, they form a good set of reference points to evaluate our approach. The procedure that was used to create the set of landmarks for each monkey is as follows. First, vessel cross-sections were manually localized by the first author on all the slices of the second slab. Next, 50 of these cross-sections were randomly selected and circles were placed around these selected cross sections.

Figure 7 shows one slice on which selected cross-sections have been encircled. Finally, the same cross sections were localized by the first author in the volume obtained after registering slab 1 to slab 2 and circles were placed around the selected cross sections. For each monkey, this resulted in pairs of images in which 50 homologous cross-sections were labeled. Two raters were then asked to localize the center of the selected cross-sections. These raters were provided with a printout of each slice of each volume on which cross-sections were selected and they picked the center of these cross-sections using a visualization tool developed in house. To evaluate intra-rater variability, these two raters were asked to repeat the experiment at a 4 week interval. The target registration error (TRE [20]) is defined as the distance between the centers of two homologous cross-sections. The intra-rater variability is defined as the distance between the centers of the same cross section selected by a rater the first and the second time. The inter-rater variability is defined as the difference between the center of the same cross-section selected by rater 1 and rater 2.

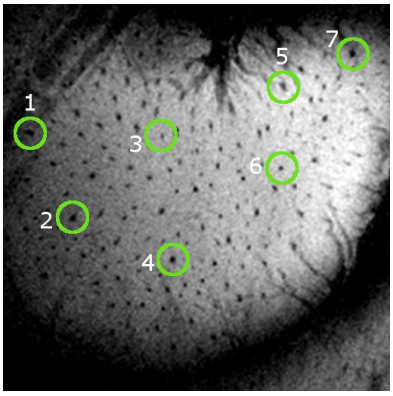

Figure 7 – Example of printout for the selection of vessels center

4. RESULTS

4.1. Direct slab-to-slab registration

Figure 8 shows results that have been obtained when attempting to register directly one slab to the other for two different monkeys. Each row in this figure shows one slice in the first slab on the left panel and the corresponding slice in the second slab on the right panel. The middle panel shows the same slice in the volume obtained after registering the first slab to the second one. As can be seen, differences in the position of the slab during acquisition lead to substantial difference in the images. In particular, some sulci can be seen in one image and different sulci in the other image. Because sulci are strong features in these images, the lack of correspondence between these makes the registration difficult.

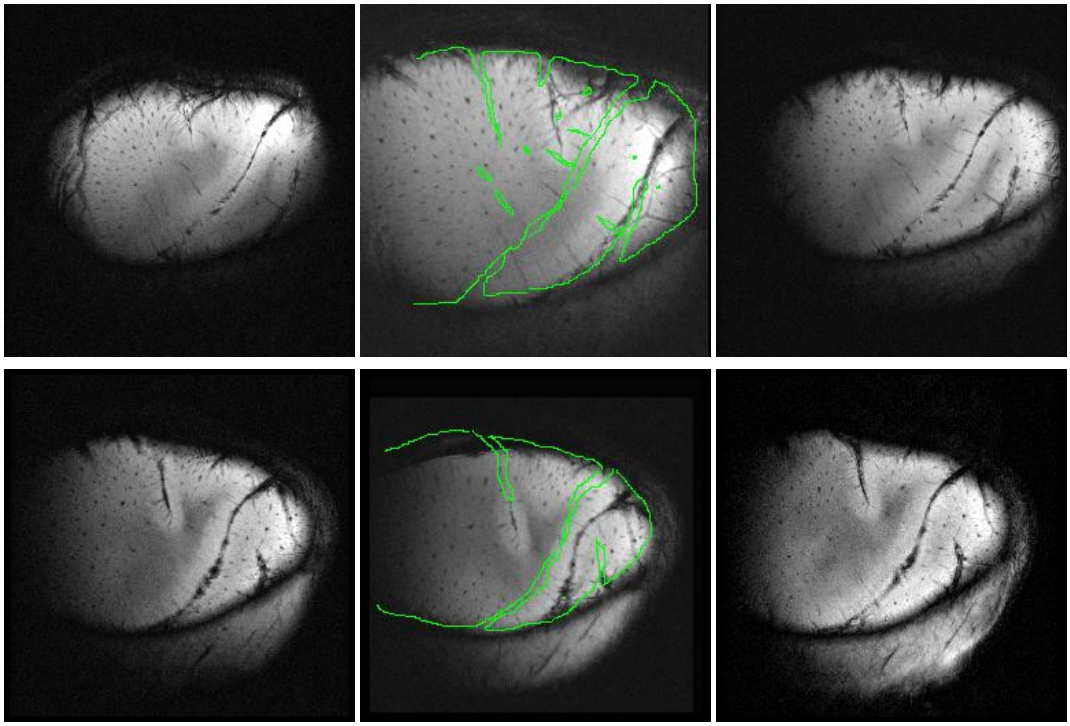

Figure 8 – Direct slab-to-slab registration is not efficient. Two examples of different monkeys' slab registered with contours of the other slab projected. Left panels: slice of first slab. Right panels: corresponding slice of second slab. Middle panels: first slab registered to second slab. Lines in the middle panel are contours of the second slab. Top: TRE is about 75 voxels. Bottom: TRE is about 35 voxels.

4.2. Slab-to-partial-3D registration

A representative example of slab-to-partial-3D registration results obtained only with information derived from the image headers is shown in the left panels of Figure 9. Although the global orientation and position of the slab is correct, it can be observed that the sulci and the edge of the cortex are misplaced. However, because the transformation computed with the header information is close to the correct transformation, an additional intensity-based registration permits the accurate registration of the slab with its corresponding partial head volume. Results obtained after this step are shown on the right panels of Figure 9.

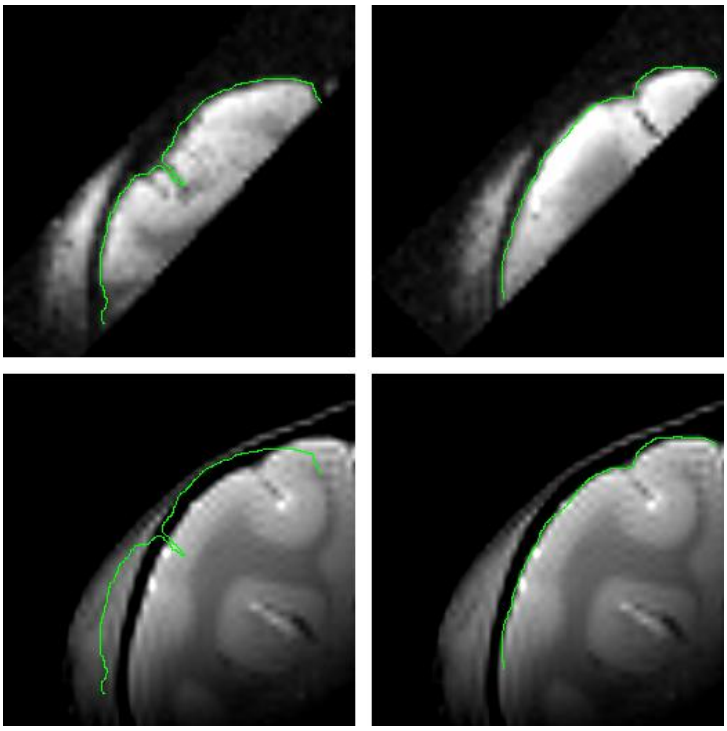

Figure 9: Registration of a slab with its corresponding partial head image. Left column with scanner-derived information only; Right column after an additional intensity-based step. Top row shows the slab image, the bottom row the partial head image. Contours have been drawn on the top images and copied on the bottom ones.

4.3. Partial-3D scans coregistration

The NMI can be viewed as a measure of the spread of the joint intensity histograms. The top panels of figure 10 shows the joint intensity histogram of the original images (left panel) and of the gradient magnitude images (right panel) when the images are registered. These figures show that the joint histogram of the original image is more diffuse than that of the gradient magnitude image. This, in turn, suggests that the NMI will be less sensitive to misalignment in the former images than the latter. The bottom panel confirms this. It shows the value of the NMI as a function of misalignment for the original images (dotted lines) and the gradient magnitude images (solid lines). For illustration purposes, a simple rotation around one axis is used to produce the misalignment. It is clear from this figure that the NMI has a peak that is much more pronounced at registration for the gradient magnitude images than for the original images. This facilitates the convergence to a local maximum, and explains why we have obtained better results with the gradient images than with the original images.

4.4. Final refinement

Figure 11 shows the effect of the last slab-to-slab intensity-based registration algorithm. The panel on the right is one slice in the second slab. The panel on the left is the corresponding slice in the volume obtained by registering slab 1 to

slab 2 after step 3 (see Figure 6). The panel in the middle is the corresponding slice in the volume obtained after the final step. After step 3, a registration error on the order of 3 voxels (~200 micrometers) can be observed. After step 4, the registration seems accurate, at least visually.

4.5. Quantitative evaluation of the complete process

Table 1 reports the target registration error as well as the inter- and intra-rater variability for the 3 data sets included in this study. The average registration error is 30.35 μm (0.444 voxels). Mean intra- and inter-rater variability are 28.43 μm and 28.71 μm respectively.

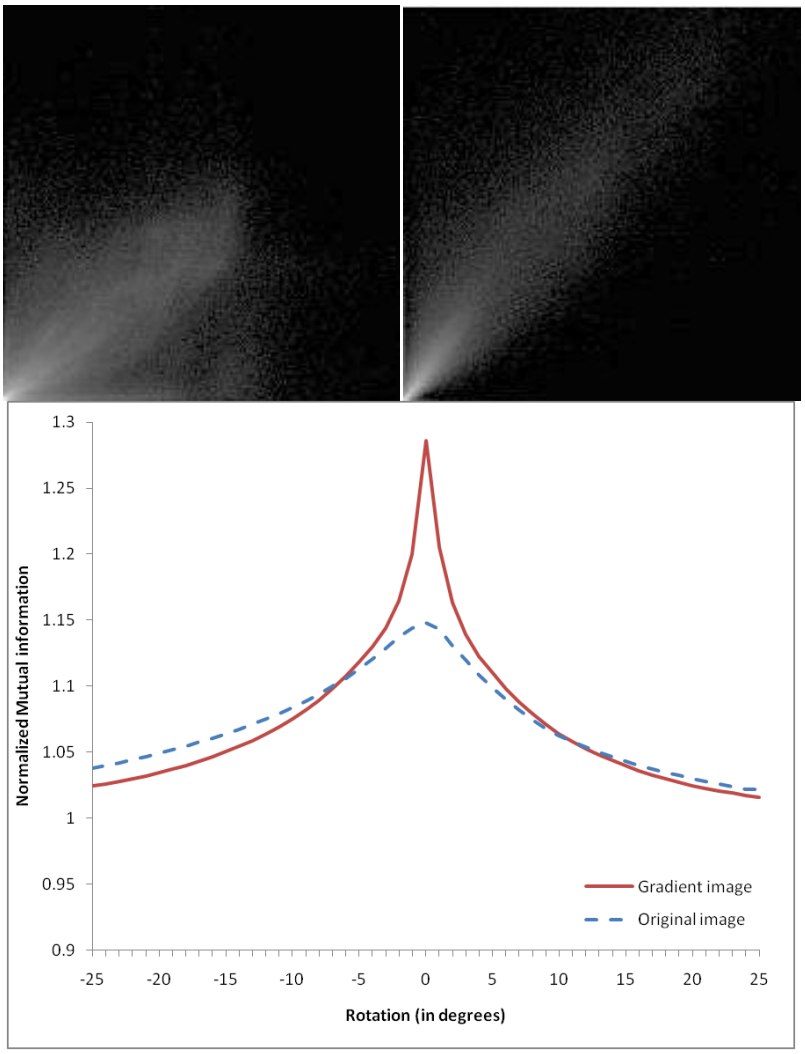

Figure 10 – Top left: Normalized joint histogram of partial 3D images after registration computed with gradient images – Top right: Normalized joint histogram of the gradient images. Bottom: NMI of gradient images and original images versus rotation. The joint histogram of the gradient image spreads less than the other one, showing a better correspondence of the voxels intensities. The NMI function of the gradient image is sharper and allows a better registration.

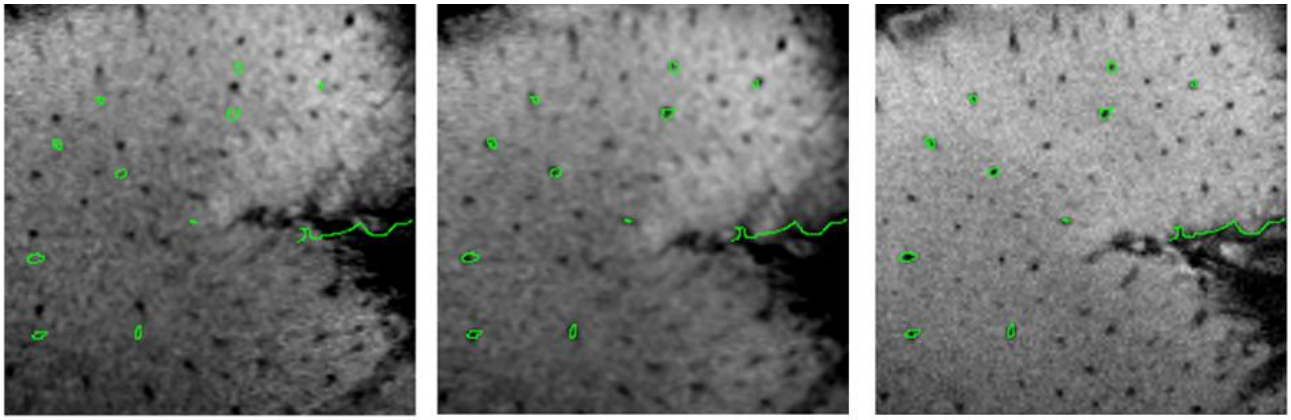

Figure 11 –Effect of the final AMIR step. Left: Slab 1 registered to slab 2 with contours of slab 2's vessel projected onto it before step 3. Middle: Slab 1 registered to slab 2 after step3. Left: Slab 2 with his vessels highlighted. Typical diameter of a vessel: 100 μm.

Table 1 – Registration error and inter- and intra-rater variability. All measures are expressed in μm. For each case, mean, standard deviation (SD), maximum and median (Med) are given.

	Registration error				Localization error of rater 1				Localization error of rater 2				Inter-rater localization error			
	Mean	SD	Max	Med	Mean	SD	Max	Med	Mean	SD	Max	Med	Mean	SD	Max	Med
Monkey 1	29.94	14.49	59.06	27.89	27.41	14.49	62.48	25.84	28.98	15.04	60.16	25.70	27.69	14.77	72.32	25.16
Monkey 2	30.90	18.25	67.88	28.23	28.92	15.38	69.45	25.22	27.55	12.17	53.84	24.75	29.46	15.04	66.92	29.26
Monkey 3	30.28	17.29	57.08	27.07	28.57	11.55	50.31	27.14	29.19	14.63	54.76	26.25	29.05	15.38	73.42	27.68
Overall	30.35	16.75	67.88	27.55	28.29	13.83	69.45	26.18	28.57	13.95	60.16	25.70	28.71	15.04	73.42	26.80

5. DISCUSSION

To the best of our knowledge, this is the first automatic registration method proposed for the longitudinal registration of small cortical slabs, especially at the spatial resolution at which we work. An analysis of the problem has shown that a simple slab-to-slab approach does not work. This resulted in a modification of the acquisition protocol for registration purposes.

A partial 3D volume could be acquired with the same surface coil used for the acquisition of the slab. However, severe intensity attenuation artifacts in these images proved to be challenging for their direct registration and intensity correction algorithm developed for human images acquired at lower field strength did not produce convincing results. This was solved by registering gradient images. Scanner imperfection and differences in spatial resolution between the slabs and the 3D images required several small adjustments to lead to an overall registration accuracy of 30.35 μm (0.444 voxels).

Other studies of registration of the somato-sensory cortex can be found in the literature but none is really comparable to ours, mostly because they rely on 3D full head MRI scans and also because they are conducted on different species, either human or rat. Schweizer et al.[5] used the FSL registration tool [21] with rigid body transformation and sinc interpolation with good results in human population but on our slabs, this led to unsatisfactory results. Others, like Sydekum et al.[6] preferred using directly the Functional Echo Planar Images (EPI) to compute the registration while we think that the registration of functional EPI should be performed using the transformation computed for anatomical images acquired during the same session as the EPI. Indeed, the anatomical images have a higher resolution; hence, structures and textures are sharper and allow a better registration. Moreover, the anatomical images are not subject to distortions.

As mentioned earlier, assessing the accuracy of automatic registration is usually done by comparing the computed transformation to a ground truth, when the transformation is already known. As none was available in our data set, we

used the trans-cortical vessels as landmarks in our images to compute this accuracy. Arguably, the manual localization of landmarks is not the ultimate gold standard. However, in this study, the vessels were picked by the same rater on each corresponding volume for the validation purpose. In our case, the TRE is about 0.4 voxels. When comparing the inter-rater and intra-rater variability to the TRE, it seemed fair to assume that the main component of the TRE is due to the localization error.

CONCLUSIONS

This study proved that fully automatic registration of longitudinal high resolutions slabs is possible. So far, the method has been applied retrospectively with success. The next step will be to use the same approach prospectively to guide the acquisition process. In this approach the position of the slab will be determined manually during the first acquisition. For subsequent acquisitions, the monkey will be placed in the scanner and the surface coil placed approximately at the same location. A partial 3D acquisition will be performed and the method we have developed will be used to determine automatically the acquisition parameters for the slab.

The ultimate objective of the work is to assess brain plasticity. To that end, fMRI images acquired in the same orientation as the slabs and covering the same space will need to be registered to each other. Geometric distortions present in EPI will need to be taken into account in the process.

ACKNOWLEDGMENTS

This project is supported, in parts, by NIH Grant R01DA028303 from the National Institute of Biomedical Imaging and Bioengineering. The content is solely the responsibility of the authors and does not necessarily represent the official views of these institutions.

REFERENCES

[1] Brodmann, K., [Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues], Johann Ambrosius Barth, Leipzig, (1909).

[2] Penfield, W. and Rasmussen, T., [Secondary Sensory and Motor Representation. The Cerebral Cortex of Man: A Clinical Study of Localization of Function], Macmillan, New-York, (1950).

[3] Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A., "Mutual-information-based registration of medical images: A survey" IEEE Transactions on Medical Imaging 22, 986-1004 (2003).

[4] Baker, J. T., Patel, G. H., Corbetta, M., and Snyder, L. H., "Distribution of activity across the monkey cerebral cortical surface, thalamus and midbrain during rapid, visually guided saccades," Cerebral Cortex 16, 447-459 (2006).

[5] Schweizer, R., Voit, D., and Frahm, J., "Finger representations in human primary somatosensory cortex as revealed by high-resolution functional MRI of tactile stimulation," NeuroImage 42(1), 28-35 (2008).

[6] Sydekum, E., Baltes, C., Ghosh, A., Mueggler, T., Schwab, M., and Rudin, M., "Functional reorganization in rat somatosensory cortex assessed by fMRI: Elastic image registration based on structural landmarks in fMRI images and application to spinal cord injured rats," NeuroImage 44(4), 1345-1354 (2009).

[7] Studholme, C., Hill, D., and Hawkes, D., "An overlap invariant entropy measure of 3D medical image alignment," Pattern Recognition 32, 71-86 (1999).

[8] Fitzpatrick, J. M. and West, J., "A blinded evaluation and comparison of image registration methods," [Empirical Evaluation Techniques in Computer Vision], K. Bowyer and P. Phillips, 12-27 (1998).

[9] Shannon, C. E., "A mathematical theory of communication," Bell System Technical Journal 27, 379-423 & 623-656 (1948).

[10] Dawant, B. M., Zijdenbos, A., and Margolin, R., "Correction of intensity variations in MR images for computer-aided tissues classification," IEEE Transactions on Medical Imaging 12(4), 770-781 (1993).

[11] Ashburner, J. and Friston, K. J., "MRI sensitivity correction and tissue classification," Neuroimage 7(4), S107 (1998).

[12] Sled, J.-G., Zijdenbos, A.-P., and Evans, A.-C., "A nonparametric method for automatic correction of intensity nonuniformity in MRI data," IEEE Transactions on Medical Imaging 17, 87-97 (1998).

[13] Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P., "Automated model-based tissue classification of MR images of the brain," IEEE Transactions on Medical Imaging 18(10), 897-908 (1999).

[14] Shattuck, D., Sandor-Leahy, S. R., Schaper, K. A., Rottenberg, D. A., and Leahy, R. M., "Magnetic resonance image tissue classification using a partial volume model," NeuroImage 13(5), 856-876 (2001).
[15] Pluim, J. P., Maintz, J. B. A., and Viergever, M. A., "Image registration by maximization of combined mutual information and gradient information," IEEE Transactions on Medical Imaging 19, 809-814 (2000).
[16] Shams, R., Kennedy, R. A., Sadeghi, P., and Hartley, R., "Gradient intensity-based registration of multi-modal images of the brain," International Conference on Computer Vision, 1-8 (2007).
[17] Liu, J., Tian, J., and Dai, Y., "Multi-modal medical image registration based on adaptive combination of intensity and gradient field mutual information," EMBS Annual International Conference, 1429-1432 (2006).
[18] Sobel, I., "An isotropic 3x3x3 volume gradient operator," tech. rep., Hewlett-Packard Laboratory, (1995).
[19] Ibanez, L., Schroeder, W., Ng, L., and Cates, J., [The ITK Software Guide], Kitware Inc., http://www.itk.org/ItkSoftwareGuide.pdf, second ed., (2005).
[20] Maurer, C. R. J., Fitzpatrick, J. M., W., M. Y., Galloway, R. L. J., Maciunas, R. J., and Allen, G. S., "Registration of Head Volume Images Using Implantable Fiducial markers," IEEE Transactions on Medical Imaging 16, 447-462 (1997).
[21] Jenkinson, M. and Smith, S., "A global optimization method for robust affine registration of brain images," Medical Image Analysis 5(2), 143-156 (2001).