Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations

Grégoire Malandain,[a,*] Éric Bardinet,[a,b] Koen Nelissen,[c] and Wim Vanduffel[c,d]

[a] Epidaure team, INRIA, 06 902 Sophia-Antipolis cedex, France
[b] CNRS UPR 640-LENA, 75 651 Parix cedex 13, France
[c] Laboratorium voor Neuro-en Psychofysiologie, Katholieke Universiteit Leuven, Campus Gasthuisberg, Leuven B-3000, Belgium
[d] MGH/MIT/HMS Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA 02129, USA

In the past years, the development of 3-D medical imaging has enabled the 3-D imaging of in vivo tissues, from an anatomical (MR, CT) or even functional (fMRI, PET, SPECT) point of view. However, despite immense technological progress, the resolution of these images is still short of the level of anatomical or functional details that in vitro imaging (e.g., histology, autoradiography) permits.

The motivation of this work is to compare fMRI activations to activations observed in autoradiographic images from the same animals. We aim to fuse post-mortem autoradiographic data with a pre-mortem anatomical MR image.

We first reconstruct a 3-D volume from the 2-D autoradiographic sections, coherent both in geometry and intensity. Then, this volume is fused with the MR image. This way, we ensure that the reconstructed 3-D volume can be superimposed onto the MR image that represents the reference anatomy.

We demonstrate that this fusion can be achieved by using only simple global transformations (rigid and/or affine, 2-D and 3-D), while yielding very satisfactory results.
© 2004 Elsevier Inc. All rights reserved.

Keywords: Histology; Autoradiography; Mortem

Introduction

In the past years, the development of 3-D medical imaging devices has enabled anatomical (MR, CT), or even functional (fMRI, PET, SPECT) 3-D imaging of in vivo tissues. However, despite huge technological advances, the resolution of these images still does not match the level of anatomical or functional details revealed by in vitro imaging (e.g., histology, autoradiography) (Muthuswamy et al., 1996).

The motivation of this work is to compare activations detected by fMRI to activations obtained using autoradiographic (AR) imaging of the same subject. Data were acquired on awake behaving animals (rhesus monkeys).

The aim of the presented work is to fuse the autoradiographic data with the pre-mortem anatomical MR image to allow straightforward comparison of the detected activations in terms of location.

The challenge in fusing a set of 2-D autoradiographs—or more generally, a set of contiguous thin 2-D sections—with a 3-D MRI volume of the same individual, is to find for each 2-D section a correspondence in the 3-D MR volume. Fusion can be done following different strategies. Among these, the most intuitive one is twofold. First a 3-D volume is reconstructed by aligning the thin 2-D sections with respect to a chosen reference one. Then, this 3-D volume is co-registered to the MR one.

Reconstruction of a 3-D volume from a stack of 2-D images (histological slices or autoradiographs) has been largely investigated in the literature. It consists of registering every pair of consecutive images in the stack to recover their geometrically coherent 3-D alignment. It however differs from the image matching problem, as classically understood, in several ways:

- slices to be registered are not images of the *same* object, but rather of *similar* objects (i.e., two consecutive sections) that can exhibit differences in shape and/or in intensity;
- non-coherent distortions may occur from one slice to the next.

The most common method, manual registration (Deverell et al., 1993; Rydmark et al., 1992), despite its simplicity, suffers from several drawbacks: it is not reproducible, it is user-dependent, and it is very time consuming. Those inherent difficulties make it not suitable to process much data.

Alternatively, fiducial markers (obtained by sticking needles in the material before slicing) can be tracked over the whole stack to recover the original geometry (Ford-Holevinski et al., 1991; Goldszal et al., 1995, 1996; Humm et al., 1995). However, since it is usually done by least squares minimization, a bias may appear if needles are not orthogonal to the cutting planes. Moreover, tracking can sometimes be difficult, especially when needle holes collapse. Furthermore, the needles may destroy part of the tissues of interest in the sections.

* Corresponding author. Epidaure team, INRIA, 2004 route des lucioles, BP 93, 06 902 Sophia-Antipolis cedex, France. Fax: +33-4-92-38-76-69.
E-mail address: gregoire.malandain@sophia.inria.fr (G. Malandain).
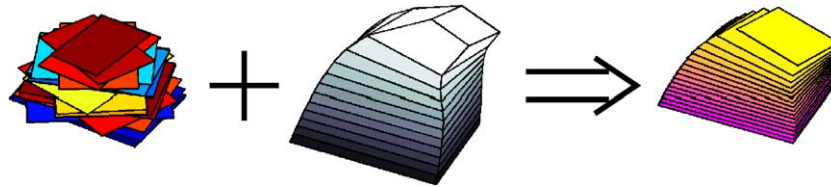Available online on ScienceDirect (www.sciencedirect.com.)

Fig. 1. The purpose of the described work is the fusion of a series of misaligned 2-D sections, that can also exhibit differences in intensity, (left) with a 3-D image (middle) to reconstruct a 3-D volume homogeneous in intensity with the geometry of the 3-D image (right).

Other images registration methods have also been investigated (Maintz and Viergever, 1998; van den Elsen et al., 1993). Those methods can be divided in two classes: the geometrical approaches which require the segmentation of some features (points, lines, or even objects of interest) and the iconic approaches based solely on image intensity.

Concerning the geometrical methods, registration can first be achieved with global descriptors, for example, centers of mass and principal axes. This proved to yield only limited precision (Schormann and Zilles, 1997), but may be used as initialization (Hess et al., 1998; Hibbard and Hawkins, 1988). More precise features, for example, contours (Cohen et al., 1998; Hibbard and Hawkins, 1988; Zhao et al., 1993), edges (Kay et al., 1996; Kim et al., 1995), or points (Rangarajan et al., 1997) can also be used.

Because of the difficulty to design a fully automatic and reliable segmentation method, or to manually extract features of interest, iconic (intensity based) methods have also been investigated. They are based on the minimization (or maximization) of a similarity measure of the images' intensities: among others, one finds cross-correlation (Hibbard and Hawkins, 1988; Ourselin et al., 2001) or mutual information (Kim et al., 1997).

Note that the work of Ourselin et al. (2001) can be considered as an hybrid ICP[1]-like approach (Besl and McKay, 1992) between geometrical and iconic methods. Blocks can be considered as geometrical features while co-registration of blocks is achieved by minimization (or maximization) of a similarity measure. We find it particularly suitable for the purpose of 3-D volume reconstruction.

So far, we have only addressed the problem of spatially aligning the 2-D slices to reconstruct a geometrically coherent 3-D volume. It should be emphasized that the problem of *intensities alignment* has rarely been addressed. Clearly, 2-D autoradiographs exhibit intensity inhomogeneities from one image to the next: sections of varying thickness present different amount of radio-isotopes, sensitivity may vary from film to film, etc. These can be compensated if appropriate standards (*micro-scales*) of known radioactivities are also imaged (and scanned) on each piece of film (Reisner et al., 1990). When they are not present, alternative approaches have to be proposed.

The previously cited literature is limited to the problem of the reconstruction of a 3-D volume from 2-D slices. However, it does not allow for an easy comparison with *more classical* 3-D modalities (e.g., MRI or CT). The fusion of such histological or autoradiographic data with other 3-D data has not been often addressed. In this case, the specific transformations due to the acquisition protocol (e.g., cutting, manipulation, chemical treatment, etc.) have to be compensated. To the best of our knowledge, histology has previously been co-registered with MRI (Schormann and Zilles, 1998; Schormann et al., 1995) and with PET data

(Mega et al., 1997). In both cases, photographs were acquired during the cutting process to serve as an intermediate modality. A 3-D linear transformation was used to map the photographic volume onto the MR volume, while 2-D highly non-linear transformations compensated residual in-plane misalignments between the histological sections and the photographs.

In this paper, we address the fusion of 2-D autoradiographic slices with a 3-D anatomical MR image, when no intermediate modality, such as photographs, is available (see Fig. 1). In this case, retrieving the correspondence for each 2-D section in the MR volume is an awkward task. Moreover, we will demonstrate that it is possible to achieve a satisfactory correspondence by using linear transformations only.

We present the data acquisition protocols in the following section. In Methods, we describe the fusion methodology, which includes a first reconstruction of a 3-D volume from the 2-D autoradiographs (the tools used for the different subtasks are described in appendices). The steps of the method are illustrated in Results. Finally, the proposed method is then discussed in Discussion.

Acquisition protocols and preprocessing

Acquisition protocols

MR image

Monkeys were trained for fMRI activations studies. During this period, several anatomical T1-weighted MR images were acquired, while animals were anesthetized to avoid motion artifacts (Vanduffel et al., 2001). Averaging those MRIs results in an anatomical MR image of size of $240 \times 256 \times 80$ and resolution of 1 mm^3.

Autoradiographs

Upon completion of the fMRI studies, monkeys were trained for the autoradiographic study, for which a double-label deoxyglucose technique (DG) was used to distinguish between two activation paradigms (please refer to (Geesaman et al., 1997; Vanduffel et al., 2000) for details). [^{3}H]deoxyglucose ([^{3}H]DG) was injected during the presentation of the first stimulus, then the second stimulus was presented and [^{14}C]deoxyglucose ([^{14}C]DG) was injected.

After a short delay, the monkey was sacrificed (via the injection of sodium pentobarbital) and perfused transcardially (first with a saline solution to wash out the blood and second with a fixative solution). The brain was then extracted from the skull. The hemisphere dedicated to our protocol was cut into two pieces and frozen as fast as possible (since the deoxyglucose diffuses rapidly at room temperature). Each block was put on the cutting plane to warrant the planarity of the bottom section after freezing.

---

[1] Iterative Closest Point.

Each block (posterior and anterior) was finally sectioned using a cryostat (slice thickness of 40 μm), resulting in 818 slices for the posterior block and 887 slices for the anterior one. Sections were mounted on coverslips and dried. All sections were first exposed against ^{3}H sensitive film and then against ^{14}C sensitive film. The autoradiographs were subsequently manually scanned.

The two images (^{3}H and ^{14}C) of a given section were co-registered (using the registration algorithm presented in Appendix A).

Acquisition analysis

The goal of this work is the fusion of the autoradiographic data (from now on, we consider only one label, e.g., ^{14}C, since the two labels of the same section have been co-registered) with the MR image, which serves as ground truth geometry. The above description of the acquisition may helps us understand the transformations that occurred and that have to be accounted for.

- Perfusion yields a global shrinkage of the brain.
- Brain extraction causes a global non-rigid deformation (mostly bending).
- CSF goes away, resulting in local deformations (e.g., ventricles collapse).
- In their position for freezing, the two blocks are submitted to biomechanical constraints (external pressure, gravity force, etc.) different from those exerted in vivo, this may yield non-rigid deformations.
- Freezing is also another cause of shrinkage.
- Sections cutting and positioning on films results in random in-plane transformations (e.g., cutting artifacts).

It should be pointed out that, while the first transformations are obviously three-dimensional and apply to the whole brain, the last item concerns 2-D slice-dependent transformations.

Preprocessing

Autoradiographs

Each brain section on the autoradiographic film was scanned into a 1276 × 1024 image, resulting in a pixel resolution of 40 μm^2 and a slice thickness of 40 μm (see Fig. 2). These images were then subsampled by a factor 4 which resulted in 818 images of size 320 × 256 with a pixel resolution of 160 μm^2.

For economical reasons, some of the digitized images contained more than one brain section. The use of very simple operators (thresholding to remove the background and mathematical morphology—erosion, selection of connected component, dilation—to separate overlapping sections) enabled us to extract the sections of interest in each image.

To automate this task, the scanning protocol was designed so that this particular section was centered in the image: this way, the component to be extracted is the one which contains the central point.

Some images (e.g., the most posterior sections, that are very small) do not follow this protocol: in this case, the sections of interest were manually selected.

Nevertheless, the procedure still failed for a few images where sections were superimposed (e.g., Fig. 2(c)): here the sections were manually segmented.
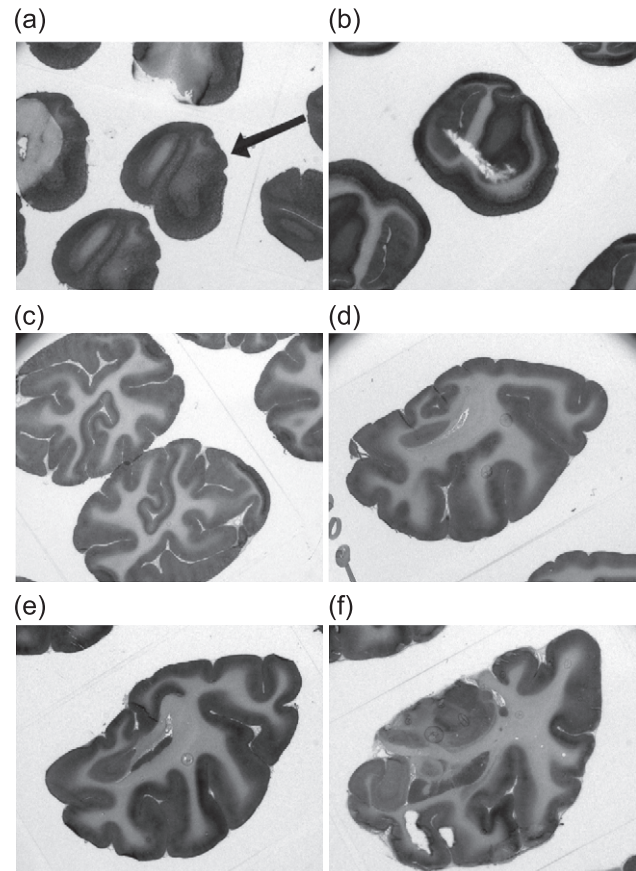

Fig. 2. Some of the original autoradiographic images (out of 818) of the posterior block of the monkey's brain.

These processed sections were put into a stack, resulting in a 3-D image of size 320 × 256 × 818 with a voxel size of 0.16 × 0.16 × 0.04 mm^3 (see top row of Fig. 6).

Anatomical MR image

The original MR volume had a size of 240 × 256 × 80 with a voxel size of 1 mm^3. We extracted in it the 66 × 89 × 48 sub-image which contained the right hemisphere of the brain.

We first performed an axis permutation on this sub-image so as to give it roughly the same geometry than the stack of autoradiographs. Finally, it was supersampled by a factor 4 using cubic splines (Thévenaz et al., 2000) to end up with an image of size 264 × 192 × 356 with a voxel size of 0.25 mm^3 (see Fig. 3).

Methods

We present in the following section the methodology of the fusion of the DG-MR volumes. We discuss in appendices the tools that were used to implement it since we consider that alternative choices for them can be done.

The proposed fusion of a 3-D MR image with the autoradiographs consists of

1. a first reconstruction of an autoradiographic volume, coherent with respect to both geometry and intensity, *without* the help of the MR image;
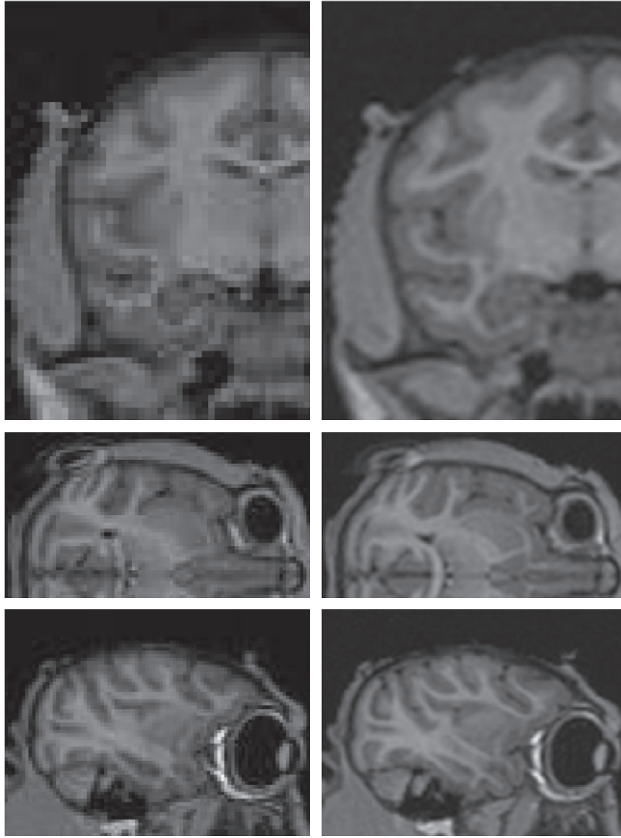
Fig. 3. Left column: a region of interest (ROI) from the original MR image centered on the right hemisphere. Right column: this ROI supersampled by a factor 4 with cubic splines.

2. a *fusion loop* that alternates between the reconstruction update of the autoradiographic volume with the help of the MR image and the 3-D registration of the autoradiographic volume to the MR image.

Notations

Let us consider two images I_1 and I_2. We will denote by $T_{I_2 \leftarrow I_1}$ the spatial transformation that maps the points of the frame attached to image I_1 onto points of the frame attached to image I_2. The image I_2 can also be viewed as a function that maps the

points of the frame attached to image I_2 into the set of intensities attached to image I_2. The composition of I_2 and $T_{I_2 \leftarrow I_1}$, denoted by $I_2 \circ T_{I_2 \leftarrow I_1}$, is then a function that maps the points of the frame attached to image I_1 into the set of intensities attached to image I_2, and represents the image I_2 in the geometry of image I_1 (in practice, after resampling).

The same kind of notations are used for the intensity transformation functions.

3-D reconstruction of an autoradiographic volume

The reconstruction of a 3-D autoradiographic volume from the 2-D autoradiographic sections S_i, $i = 1 \ldots N$ is twofold (see Fig. 4).

Geometry consistency

Following the scheme presented in Ourselin et al. (2001), we realign the 2D sections to build a geometrically coherent volume.

Each couple of consecutive sections, S_i and S_{i+1} is rigidly co-registered, with the method described below (see Appendix A), yielding a 2-D transformation $T_{i \leftarrow i+1}$ (or equivalently $T_{i+1 \leftarrow i}$). A reference section, of index ref$_1$, is then chosen, and the 2-D transformations $T_{i \leftarrow \text{ref}_1}$ are computed by composition of the $T_{i \leftarrow i+1}$ (if $i < \text{ref}_1$) or $T_{i+1 \leftarrow i}$ (if $i > \text{ref}_1$) transformations.

The resampled sections, $S_i \circ T_{i \leftarrow \text{ref}_1}$, can now be stacked to build a geometrically coherent 3-D volume. At this stage, visual inspection is necessary, to detect some eventual misalignments due to a wrong co-registration between two consecutive sections. These misalignments separate consistent sub-volumes in the 3-D reconstruction (see the second row of Fig. 6). Indeed, consider a misregistration for the couple of sections of indices e and $e+1$. The 2-D transformations with respect to the reference section of index ref$_1$ are computed by

$$T_{i \leftarrow \text{ref}_1} = T_{i \leftarrow i+1} \circ \cdots \circ T_{e-1 \leftarrow e} \circ T_{e \leftarrow e+1} \circ \cdots \circ T_{\text{ref}_1 - 1 \leftarrow \text{ref}_1}$$

for $i < e < \text{ref}_1$. The sections from index i to index e are then correctly co-aligned together (they form a consistent sub-volume) but are not correctly co-aligned with section of index $e+1$. By visually inspecting the reconstructed volume, these misregistrations are thus easily identifiable. Typically, this only concerns a small number of sections, less than 2% of the total. To correct them, we change the default registration parameters (e.g., the exploration neighborhood size r_x and r_y, or the fraction p of retained blocks, see Appendix A.1) and co-register again the misaligned couple of consecutive sections, until the transformation is satisfactory.
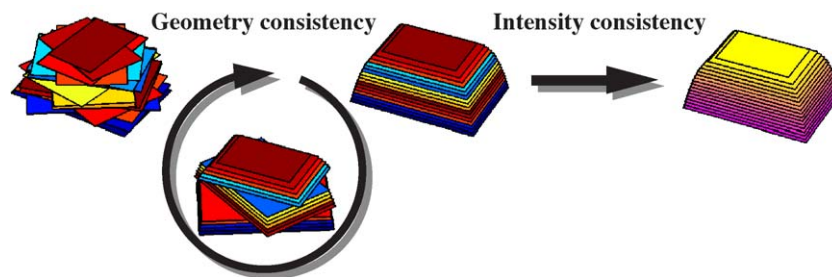

Fig. 4. The reconstruction of a 3-D volume from a series of 2-D sections consists in the co-registration of any two consecutive slices and in the correction of intensities inhomogeneities. The registration parameters of the couples of misregistered sections are changed until a satisfactory result is obtained.

Intensity consistency

Because the overall intensity may change from one autoradiographic image to the next, we compensate for these changes with a dedicated histogram matching algorithm applied to all pairs of consecutive sections (see Appendix B), yielding an affine intensity transformation $f_{i \leftarrow i+1}$ (or equivalently $f_{i+1 \leftarrow i}$) between two consecutive sections. A reference section, of index ref_2, is chosen, and the intensity transformations $f_{\text{ref}_2 \leftarrow i}$ are computed by composition of the $f_{i \leftarrow i+1}$ or $f_{i+1 \leftarrow i}$ intensity transformations.

To summarize, the 3-D reconstruction of an autoradiographic volume, coherent with respect to both geometry and intensity, is achieved by the superposition of the resampled sections $f_{\text{ref}_2 \leftarrow i} \circ S_i \circ T_{i \leftarrow \text{ref}_1}$.

Initial 3-D registration MRI/autoradiographic volume

The reconstructed autoradiographic volume AV is then coregistered with the MR volume (see Appendix A) using a 3-D rigid transformation. This provides an initial solution to our fusion problem.

Fusion loops

This loop aims to fuse the 2-D autoradiographic sections with the 3-D MR image. Let $AV^{(0)}$ be the first 3-D reconstruction of the autoradiographic volume and $T_{\text{MR} \leftarrow AV^{(0)}}$ the initial 3-D transformation (as obtained above). Each iteration k of the loop is divided into four steps.

1. The first step is twofold;
(a) Resampling of the MR volume into the geometry of $AV^{(k)}$ yielding the volume $\widetilde{MR}^{(k)} = \text{MR} \circ T_{\text{MR} \leftarrow AV(k)}$. This allows to extract MR slices, $\widetilde{MR}_i^{(k)}$, $i = 1 \ldots N$, that correspond to autoradiographic slices.
(b) 2-D registrations of each autoradiographic slice $AV_i^{(k)}$ against the corresponding resampled MR slice, resulting in N 2-D independent transformations $T_{AV_i^{(k)} \leftarrow \widetilde{MR}_i^{(k)}}$ (see Appendix A).
2. Filtering of the 2-D independent transformations $T_{AV_i^{(k)} \leftarrow \widetilde{MR}_i^{(k)}}$, yielding the 2-D correlated transformations $\hat{T}_{AV_i^{(k)} \leftarrow \widetilde{MR}_i^{(k)}}$ (see Appendix C).
3. Building of a new autoradiographic $AV^{(k+1)}$ volume by superposing the slices $AV_i^{(k)} \circ \hat{T}_{AV_i^{(k)} \leftarrow \widetilde{MR}_i^{(k)}}$.
4. 3-D registration of the autoradiographic volume $AV^{(k+1)}$ against the 3-D MR image, yielding a 3-D transformation $T_{\text{MR} \leftarrow AV^{(k+1)}}$ (see Appendix A).

This loop, that alternates between the reconstruction update of the autoradiographic volume and the 3-D registration of the autoradiographic volume against the MR image, is summarized in Fig. 5.

We only use global parametric transformations in this loop, for example, rigid transformations, similarities, or affine transformations. Table 2 reports the selected transformations for each loop.

Results

Reconstruction of an autoradiographic volume

A first autoradiographic volume is reconstructed by registering each couple of consecutive sections using the 2-D block matching algorithm (typical values of parameters are given in Table 1). The
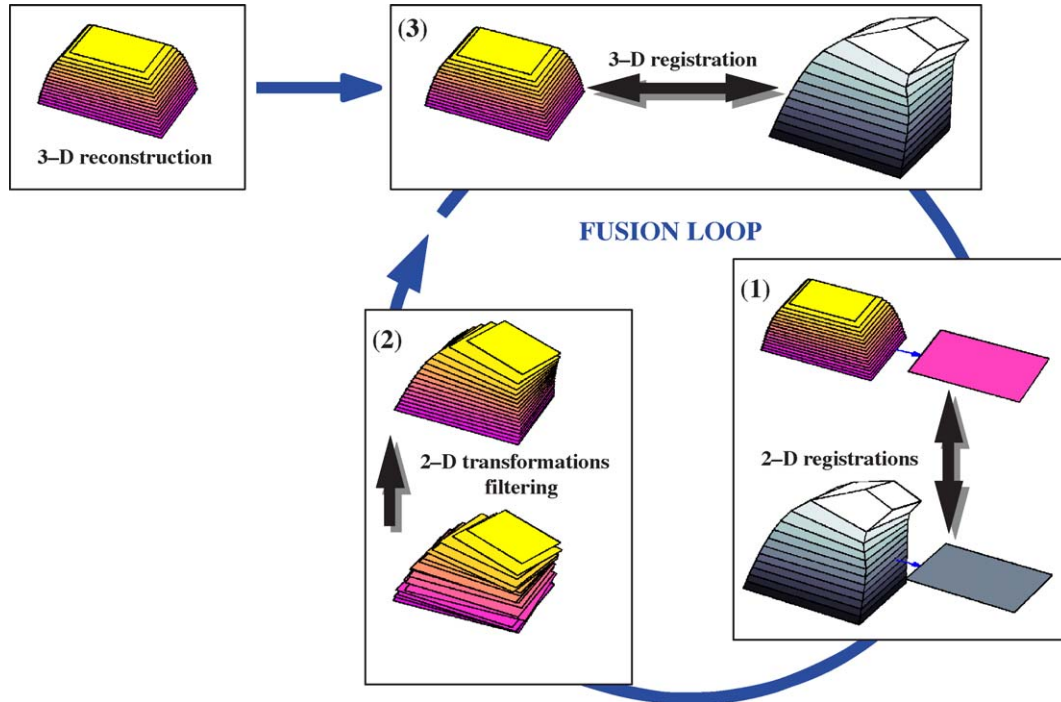

Fig. 5. Fusion loop (see text): (1) 2-D independent registrations of each AR section against the corresponding slice of the resampled MR volume; (2) 2-D transformations filtering and reconstruction of a 3-D AR volume; (3) 3-D registration of the AR volume against the MR volume.

Table 1
Typical values of the parameters for both the 2-D and the 3-D registration algorithms

	d_x	d_y	d_z	s_x	s_y	s_z	r_x	r_y	r_z	$p(\%)$	l
2-D	4	4	–	3	3	–	3	3	–	80	6
3-D	4	4	4	3	3	3	3	3	3	100	3

composition of the computed transformations allows the registration of all the sections against a reference one (taken at the middle of the stack).

Visual inspection of the reconstructed volume enabled the detection of possible remaining registration errors that were subsequently corrected by changing registration parameters. This control step was repeated until the result was satisfactory, that is, until the reconstructed volume seemed geometrically consistent. The first row of Fig. 6 shows the stack of autoradiographs to be compared to the obtained reconstruction (third row of same figure).

Intensity consistency was also obtained via histogram matching (see Appendix B) for each couple of consecutive sections. Note that this computation is independent from the above geometrical registration. Among the available criteria (sum of squared differences (SSD), correlation, maximum of likelihood (ML)), the SSD and the ML yield satisfactory results. We computed then several intensity-corrected reconstruction with different kernel's smoothing factor ($\sigma \in \{3, 5, 8, 12\}$). The best result (based on visual impression) was selected for the further steps. The last row of Fig. 6 displays the final reconstruction.

Initial registration with MRI

After reconstruction of the autoradiographic volume, the latter can be registered (here rigidly) against the MR sub-volume of interest (i.e., the corresponding hemisphere). The top row of Fig. 7 shows cross-sections of the reconstructed AR volume while the bottom row shows the corresponding cross-sections of the resampled MR volume after registration. Both volumes are roughly similar, but large differences can be seen especially in posterior area.

Fusion loops

The fusion loops consist in iterating four successive steps:

1. 2-D independent registrations of each AR section against the corresponding slice of the resampled MR volume,
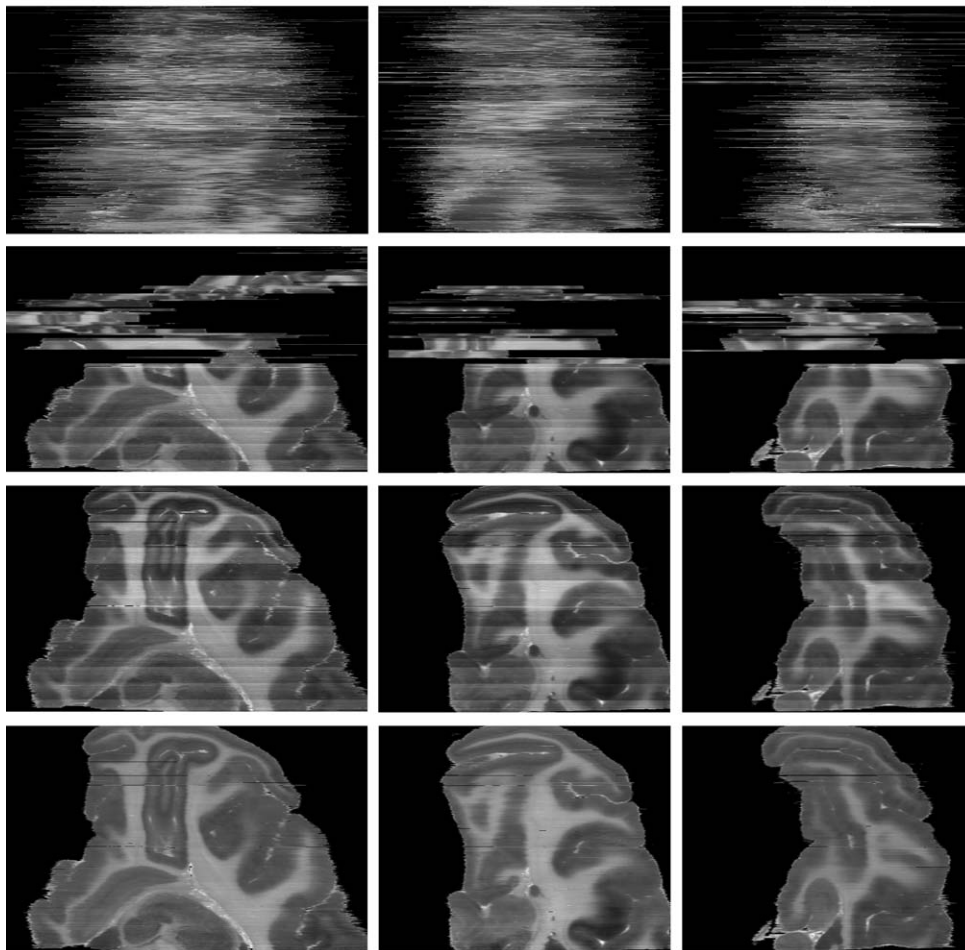

Fig. 6. From left to right: one axial and two sagittal slices of the reconstructed autoradiographic volume of the posterior block. First row: before any correction. Second row: reconstruction with the default registration parameters; failures are easily identifiable. Third row: after geometrical alignment of the slices, and before intensity correction. Fourth row: after intensity correction.
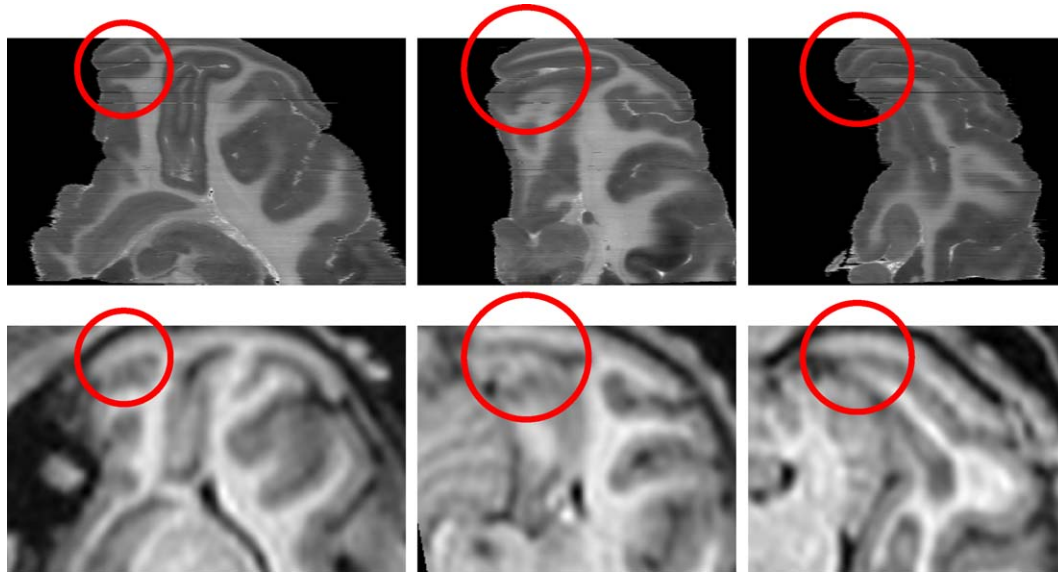
Fig. 7. Top row: one axial and two sagittal slices of the reconstructed autoradiographic volume after geometric alignment and intensity correction (id. last row of Fig. 6). Bottom row: corresponding slices of the resampled MR volume after 3-D rigid registration. The circles highlight registration errors, which are particularly visible in the posterior area.

2. 2-D transformations filtering,
3. reconstruction of a new 3-D AR volume,
4. and 3-D registration of the newly reconstructed AR volume against the MR volume.

The different transformation classes that were considered through the fusion loops can be found in Table 2. We stopped after the third loop since we did not notice any evolution with respect to the second one.

Intermediate results for loop # 1 are given in Fig. 8. Fig. 9 shows the final alignment of the histological slices with respect to the MR volume.

Fig. 10 highlights the advantages of the proposed fusion methodology: top, middle, and bottom rows show cross-sections of the original MR sub-volume of interest (the right hemisphere), the first reconstructed and resampled AR volume after rigid registration, and the final reconstructed and resampled AR volume after affine registration, respectively. Contours extracted from the presented MR cross-sections are superimposed on the AR cross-sections in Fig. 11.

Improvements along the fusion process can be seen in Fig. 12 which presents the differences between any two successive autoradiographic volumes (registered against the MR image). Displacements are large after loop #1 and become smaller between loop #2 and loop #3. The computation of the mutual information criterium

(which is not the measure that is optimized), through the fusion process, offers also a mean to assess the benefits of the proposed method (see Fig. 13).

The residual differences that can be seen in these figures suggest that further improvements could be reached by gradually increasing the number of degrees of freedom of the transformations.

We have also fused the anterior part of the same monkey's brain (see Fig. 14) to the MR volume. As we did observe only minor changes between loops #1 and #2, we did not go into the third loop. Indeed, the so-called banana problem (see Discussion) was less acute in this case.

Visual-based inspection

The results were checked using two different methods. First, the two volumes in the same geometry (here the one of the MR) can be seen in two synchronized 3-D viewers: both of them always display the same cross-sections (axial, sagittal and coronal) as well as a cursor at the same position (see Fig. 14(A) and Fig. 14(B)). Both volumes can alternatively be visually superimposed, possibly using different color maps and transparency (see Fig. 14(C)) (Delingette et al., 2001).

Both tools allow the navigation in the fused 3-D volumes which permits to check the accuracy of the fusion, for example, by following sulci. At the level of the visual cortex, the correspondence is accurate, as well as for most of the brain. Some subtle mismatch errors can be seen at the medial side of the caudate nucleus (due to collapsed ventricles).

Discussion

We first address some methodological issues: first, the sole problem of reconstructing a volume from a series of 2-D autoradiographs and second, the fusion of the 2-D autoradiographs with

Table 2
Fusion loop, choice of the parameters: transformation classes and 2-D transformation filtering

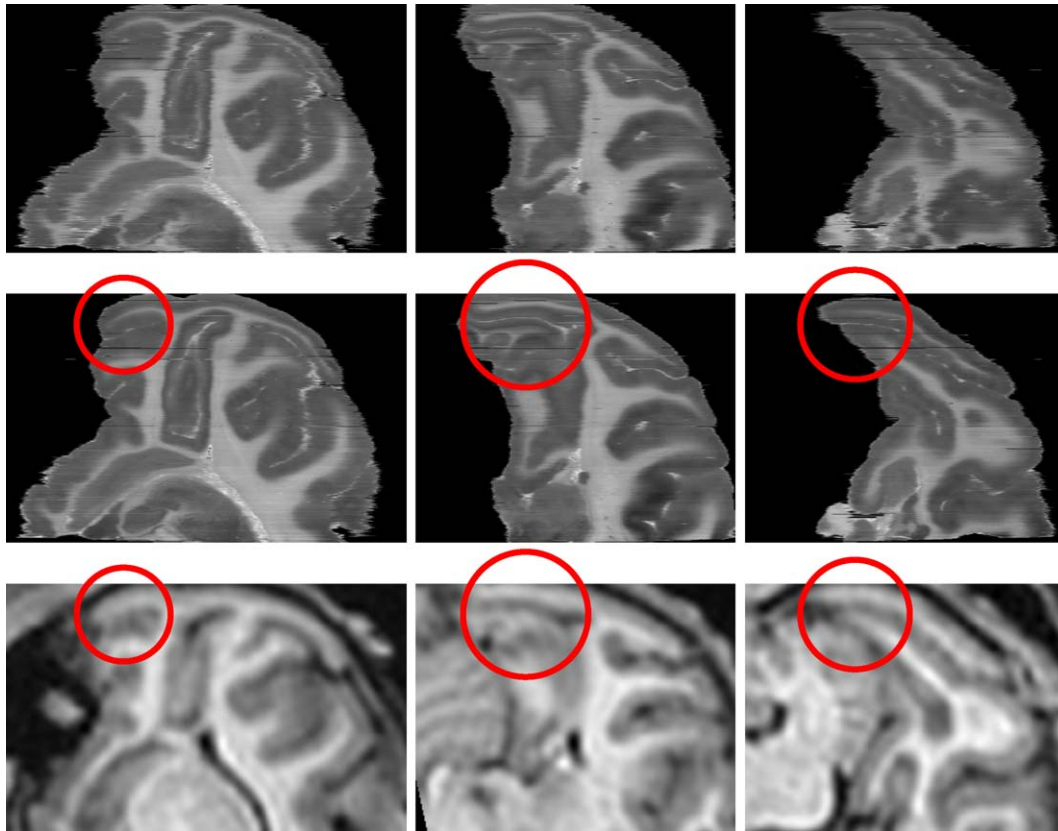|	2-D registrations $T_{AV_i^{(k)} \leftarrow \widetilde{MR}_i^{(k)}}$	2-D transformation filtering $\sigma_{(k)}$	3-D registration $T_{\mathrm{MR} \leftarrow \mathrm{AV}^{(k+1)}}$
Loop # 1	rigid	10	affine
Loop # 2	rigid	8	affine
Loop # 3	affine	10	affine

Fig. 8. Fusion loop # 1. From left to right: one axial and two sagittal slices. First row: autoradiographic volume after the independent 2-D slice registration. Second row: autoradiographic volume after smoothing of the 2-D transformation. Third row: MR image resampled after affine registration against the above-reconstructed AR volume. The effect of this first loop can be appreciate by comparing the middle row to the first row of Fig. 7.

the MR volume. Then, we discuss the conditions of applicability of our method.

Reconstruction

This step consists of two parts, a geometrical reconstruction and an intensity compensation.

Concerning the geometrical reconstruction, a visual inspection (and possibly manual correction) is necessary in case of mismatch between two consecutive sections. It should be pointed out that *mismatch* does not mean *poor accuracy*. Indeed, it has been proven that the block matching algorithm gives accurate results (Ourselin, 2002; Ourselin et al., 2001) given initial conditions in the convergence basin. If not, for example, in case of failure, there is a huge discrepancy between the obtained transformation and the expected one so the failure can be easily visually recognized. This is all the more noticeable when reconstructing a full set as two mismatched consecutive sections will separate two geometrically consistent 3-D sub-volumes. We corrected these errors by changing the registration parameters. This modified the shape of the convergence basin, now including the initial condition. Another solution could have been to change the initial transformation parameters.

Visual inspection of the intensity-based reconstruction is facilitated by similar considerations. Indeed, a single error (a mismatch between two consecutive sections) results in a clear separation between two intensity-consistent 3-D sub-volumes.

One might question the necessity of the intensity-based reconstruction. When no additional assumption can be made about the autoradiographs' acquisition protocol, it appears to be required. Indeed, the 3-D reconstructed AR volume has to be registered against the MR volume. To do so, 3-D sub-images (blocks) of the AR volume will be compared to 3-D blocks of the MR volume. The correlation coefficient assumes an affine relationship between intensities of these two blocks and thus needs each of both blocks to be consistent in intensity. Obviously, it is not possible to find the right correspondence for a block that would not be consistent in intensity. Since our implementation of block matching is robust (it can reject outliers), it may still compute the expected 3-D transformation if there are not *too many* corrupted blocks, for example, if there are only very few intensity *jumps* in the AR volume. Without such additional assumption, intensity-based reconstruction is thus necessary. However, it also appears that the visual examination of the 3-D reconstructed AR volume is more comfortable with intensity compensation than without.

Fusion

Fusion of a set of 2-D autoradiographs, or more generally of a set of contiguous thin 2-D sections, with an MR volume of the same individual, can be done following a twofold approach. First, the thin 2-D sections are aligned with respect to a chosen reference section, which yields a reconstructed 3-D volume; then this 3-D volume is co-registered to the MR volume.
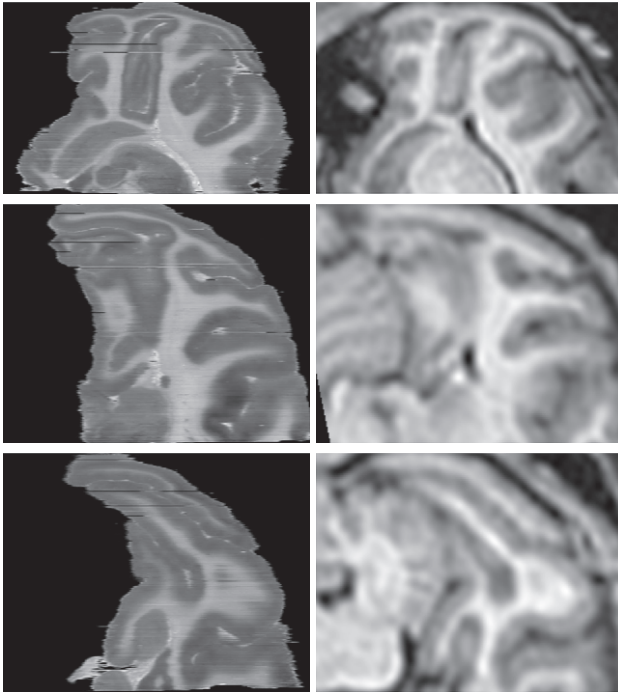
Fig. 9. Fusion loop # 3. From top to bottom: one axial and two sagittal slices. Left column: autoradiographic volume after smoothing of the 2-D transformation. Right column: MR image resampled after affine registration against the above reconstructed AR volume.

Obviously, as a result of such a co-registration, one expects to find a point-to-point correspondence between the MR volume and the stack of 2-D sections. This implicitly assumes that the alignment of the 2-D sections respects the real underlying anatomy. The above twofold approach is clearly not design to encounter such an expectation, which is generally not achieved in practice (see Fig. 7).

This difficulty in the reconstruction of a 3-D volume from 2-D sections (either histological or autoradiographic) comes from the so-called banana problem (illustrated in Fig. 15): namely a 3-D curved object cannot be reconstructed from cross-sections without any additional information.

The acquisition protocol can be designed to include additional information in the form of fiducials. To do so, needles can be sticked in the material before slicing, which provides correspondences to drive the geometry-based reconstruction (Ford-Holevinski et al., 1991; Goldszal et al., 1995, 1996; Humm et al., 1995). However, reconstruction can be awkward, particularly if needles are not orthogonal to the cutting plane or if needle holes collapse during the histological treatment. Another way is to take photographs of the unstained surface of the material during the slicing process, including a reference system fixed on the cryomicrotome (Mega et al., 1997; Schormann and Zilles, 1998; Schormann et al., 1995). Alignment of the photographs using the reference system will then provide a photographic 3-D volume with the real geometry of the object under study, for example, the brain after extraction from the skull but before slicing and staining. Thus, further reconstruction of the 2-D histological sections into a geometrically consistent 3-D volume can be made by independent registrations of the 2-D sections with the corresponding previously aligned photographs.

When such a priori additional information about the real geometry of the object under study is available, the banana problem does not occur. Nevertheless, for practical reasons, such information is not always available. In that case, the only information about the geometry of the brain (its real anatomy) before slicing is given by the anatomical MR volume.

A direct and strict twofold approach, namely the reconstruction of a 3-D volume by alignment of the 2-D sections followed by the registration of this volume with the anatomical MR volume could have provided a satisfactory result. This implicitly assumes that either the computed 2-D transformations will compensate the random in-plane transformations due to the acquisition procedure (this is not realistic for curved objects), or that further 3-D elastic registration may compensate for the residual distortions. However, existing elastic transformations implemented in such algorithms are obviously not adapted. First, they consider in a similar manner the three directions of space. In our particular problem, one direction (the one orthogonal to the cutting plane) clearly plays a different role. Moreover, from a methodological point of view, they do not take the acquisition procedure into account (see Acquisition analysis). Namely, transformations that have occurred during the acquisition are of different types, that is, 3-D and applied to the whole brain, or 2-D and independent from section to section.

The proposed fusion methodology mimics the acquisition procedure by considering both a stack of 2-D transformations (that correspond to the displacements of the AR sections) and a 3-D transformation corresponding to the registration against the MR volume.

More precisely, after an initial reconstruction of the 2-D sections, we alternate between the correction of this reconstructed AR volume (by recomputing the 2-D transformations) and a 3-D registration with the MR volume. By doing so, we allow the 2-D sections to slide on each other, and we expect to estimate more optimally the random in-plane transformations due to the acquisition procedure, together with the relative position of the reconstructed volume with respect to the MR volume. The choice of the transformation search spaces during these fusion loops is made so as to preserve the integrity (e.g., constant slice thickness, parallelism of sections) of the AR sections. That is, the optimization starts with strongly constrained transformations (i.e., rigid), which are then slightly relaxed (ending up with affine transformations both in 2-D and 3-D). In other words, we do not allow the AR sections to be strongly deformed before their corresponding MR slices have been localized for a given transformation search space.

It should be pointed out that the obtained result is visually correct, although it only involves a stack of 2-D and 3-D linear transformations.

Fig. 16 illustrates the banana problem with our data. It shows (in blue) the vertical central line of the first 3-D reconstruction of the AR volume (last row of Fig. 6) that could be considered as the symmetry axis of the cylindrical reconstructed object, and the deformation of this line (in red) after the fusion: a curvature appears.

Fig. 17 compares the result of the fusion using our approach and a twofold approach. It shows (middle row) a direct registration using affine transforms of the reconstructed AR volume with the MR volume, compared to the final result using fusion loops (bottom row). This demonstrates that a direct 3-D affine transformation is obviously not sufficient to model the deformation between the initial AR reconstructed volume and the MR volume. On the middle row of Fig. 17, the anterior part of the AR volume is well registered,
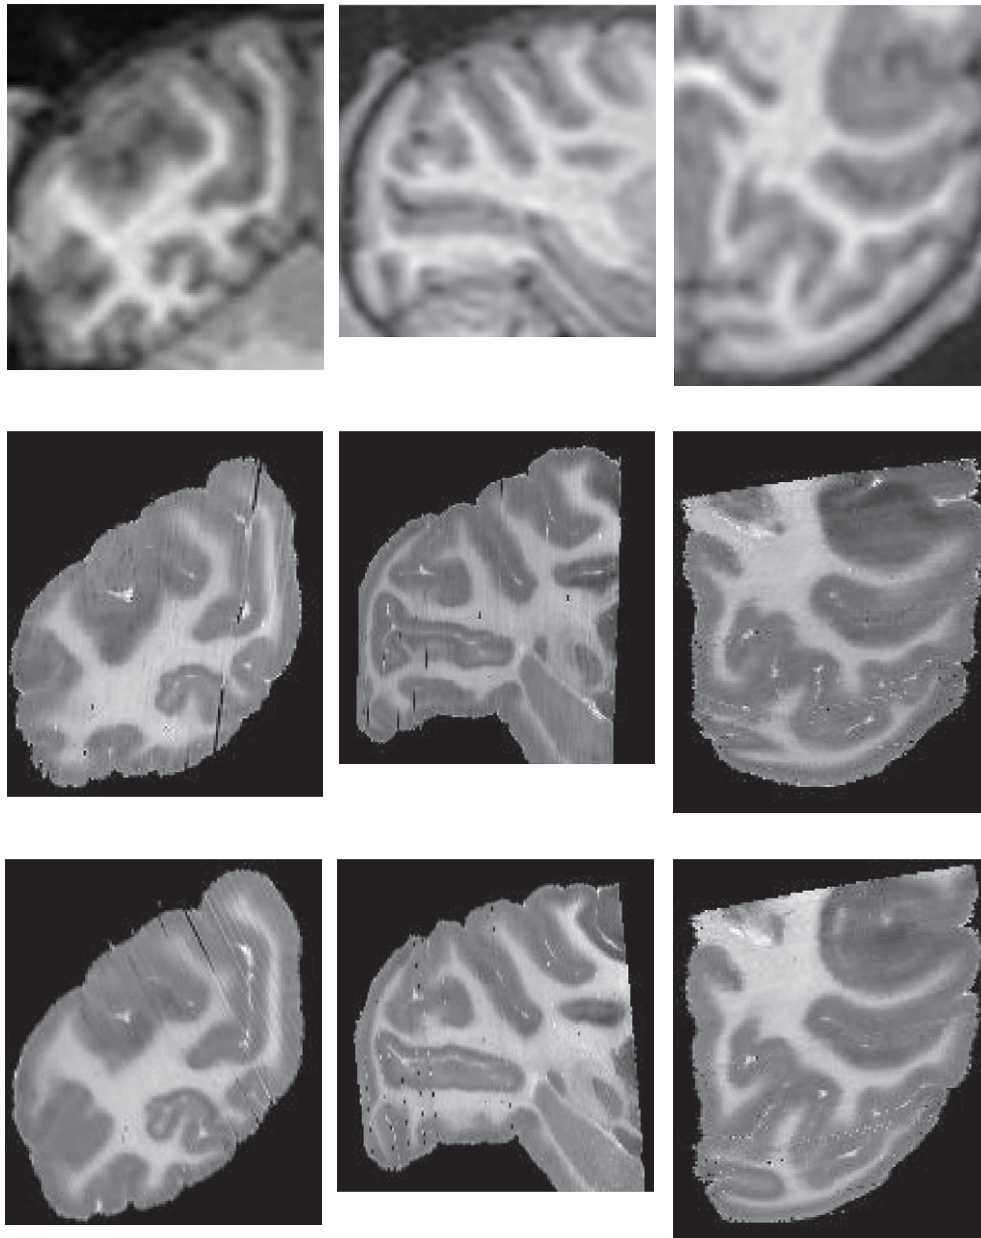
Fig. 10. From left to right: coronal, sagittal, and axial slices in the acquisition geometry of the MR image. Top row: the MR image (resampled with cubic splines, see Fig. 3). Middle row: the three same slices of the first reconstructed autoradiographic volume (before any fusion loop) rigidly registered against the MR image. Bottom row: the three same slices of the last reconstructed autoradiographic volume (after all fusion loops) registered against the MR image.

while a large discrepancy appears in the posterior part. This suggests that a further 3-D elastic registration (usually initialized by a rigid or affine transformation) would result in inhomogeneous deformations of the AR slices along the antero-posterior axis. Thus, we suspect that the latter procedure would alter the integrity of the AR slices and would not yield the most exact point-to-point correspondence in the MR volume for each 2-D AR section.

Conditions of use

If additional fiducial markers cannot, or have not, be included in the acquisition protocol (see the above discussion about fusion),

we demonstrated that it is still possible to fuse a stack of 2-D sections with a 3-D image. However, there is some limitations to our approach with respect to the above-mentioned one.

- To obtain the first reconstruction (section 3.1), the registration of two consecutive sections must make sense. It means that there should be enough information redundancy (similar structures should appear, etc) in these two sections so that a registration algorithm could co-align them (if there are no common structures, co-registration cannot succeed). Dealing with thin sections sounds obviously better than thick ones, however, this does not intrinsically put an absolute constraint (say 40 or 20 μm)
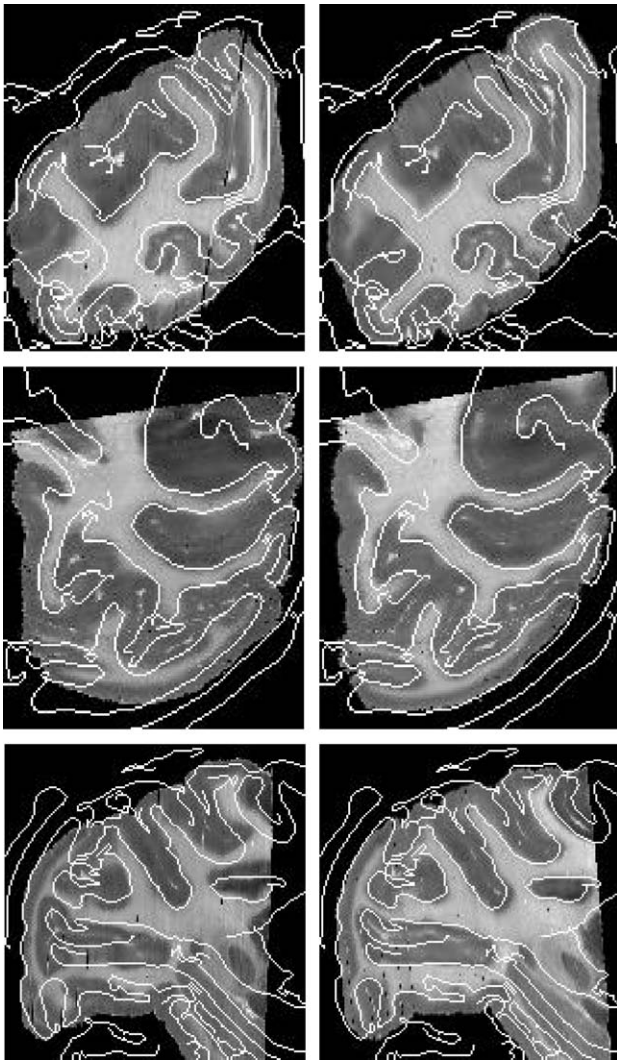
Fig. 11. The contours extracted from the MR image (first row of Fig. 10) superimposed on the corresponding slices of the middle (left column) and last (right column) rows from the same figure.

on the slice thickness since it highly depends on the imaged data. The condition for a successful reconstruction can be formulated as "enough 3-D structures can be seen in a number of consecutive sections."

- To co-register a 3-D reconstruction with a 3-D image, or a 2-D section with the corresponding resampled slice in the 3-D image, the same condition holds, that is, registration must make sense. It means that we must have two 3-D volumes (the total thickness of the sectioned material should be of the same order than the in-plane sections' dimensions), and that there is enough overlap and information redundancy between the two volumes.

- In this work, we only use linear (rigid or affine) transformations. It was sufficient for the proposed application, since we manipulate frozen sections, so that in-plane distortions are unlikely to occur, or are of limited amplitude (e.g., near the collapsed ventricles)

In other situations (e.g., histological stained sections), such 2-D distortions will certainly occur, and the assumption of global linear (affine or rigid) transformation will not hold anymore. The introduction of more complex transformations when co-regis-

tering two consecutive sections (e.g., 2-D piecewise affine transformations (Pitiot et al., 2003)) enables to correctly reconstruct a 3-D volume from a stack of stained sections (in our fusion scheme, this only concerns the first step, see 3-D reconstruction of an autoradiographic volume). However, to compensate for the distortions in the fusion loops is still under study.

Conclusion

This paper describes a methodology that allows the fusion of 2-D sections (autoradiographs) with a 3-D volume (MR). The iteration of a so-called fusion loop, which alternates between the correction of the reconstructed AR volume (by recomputing the 2-D transformations) and a 3-D registration, delivers very satisfactory results as confirmed by visual inspection.

Since such 2-D section (autoradiographs or histological stained sections) can be considered as imaging the *ground truth*, to fuse them with in vivo 3-D imaging modalities provides a way to assess the signal (or results emanating from the processing of the signal) of these modalities.

The proposed method address the case where additional fiducial markers (or intermediary images, e.g., photographs) cannot, or have not, be included in the acquisition protocol (see the above discussion about fusion).

From a methodological point of view, it should be pointed out that the fusion only involves a stack of 2-D affine transformations and a 3-D affine transformation (no complex 3-D or 2-D elastic transformations are required): we argue that this is appropriate since it faithfully models the acquisition procedure.

Yet, we still observe a few mismatched areas in our results. Following our strategy, we could perform an additional 2-D registration step between AR and corresponding MR sections using transformations with more degrees of freedom. One way would be to use transformations used by Mega et al. (1997) and Schormann and Zilles (1998), and another way would be to design a transformation that models the specific 2-D piecewise distortions (Pitiot et al., 2003).

Finally, thanks to the acquisition procedure (the brain and the AR sections are almost always manipulated when frozen), geometrical distortions within an AR section are somehow minimized. It will be challenging to process histological stained data with the same method since, in the later case, strong distortions will certainly occur because of the staining procedure.

Acknowledgments

Appendix A. 2-D and 3-D Registration

In principle, any registration method can be used to register two images when the latter are two 2-D autoradiographs or the 3-D autoradiographic reconstructed volume and the 3-D MR volume,
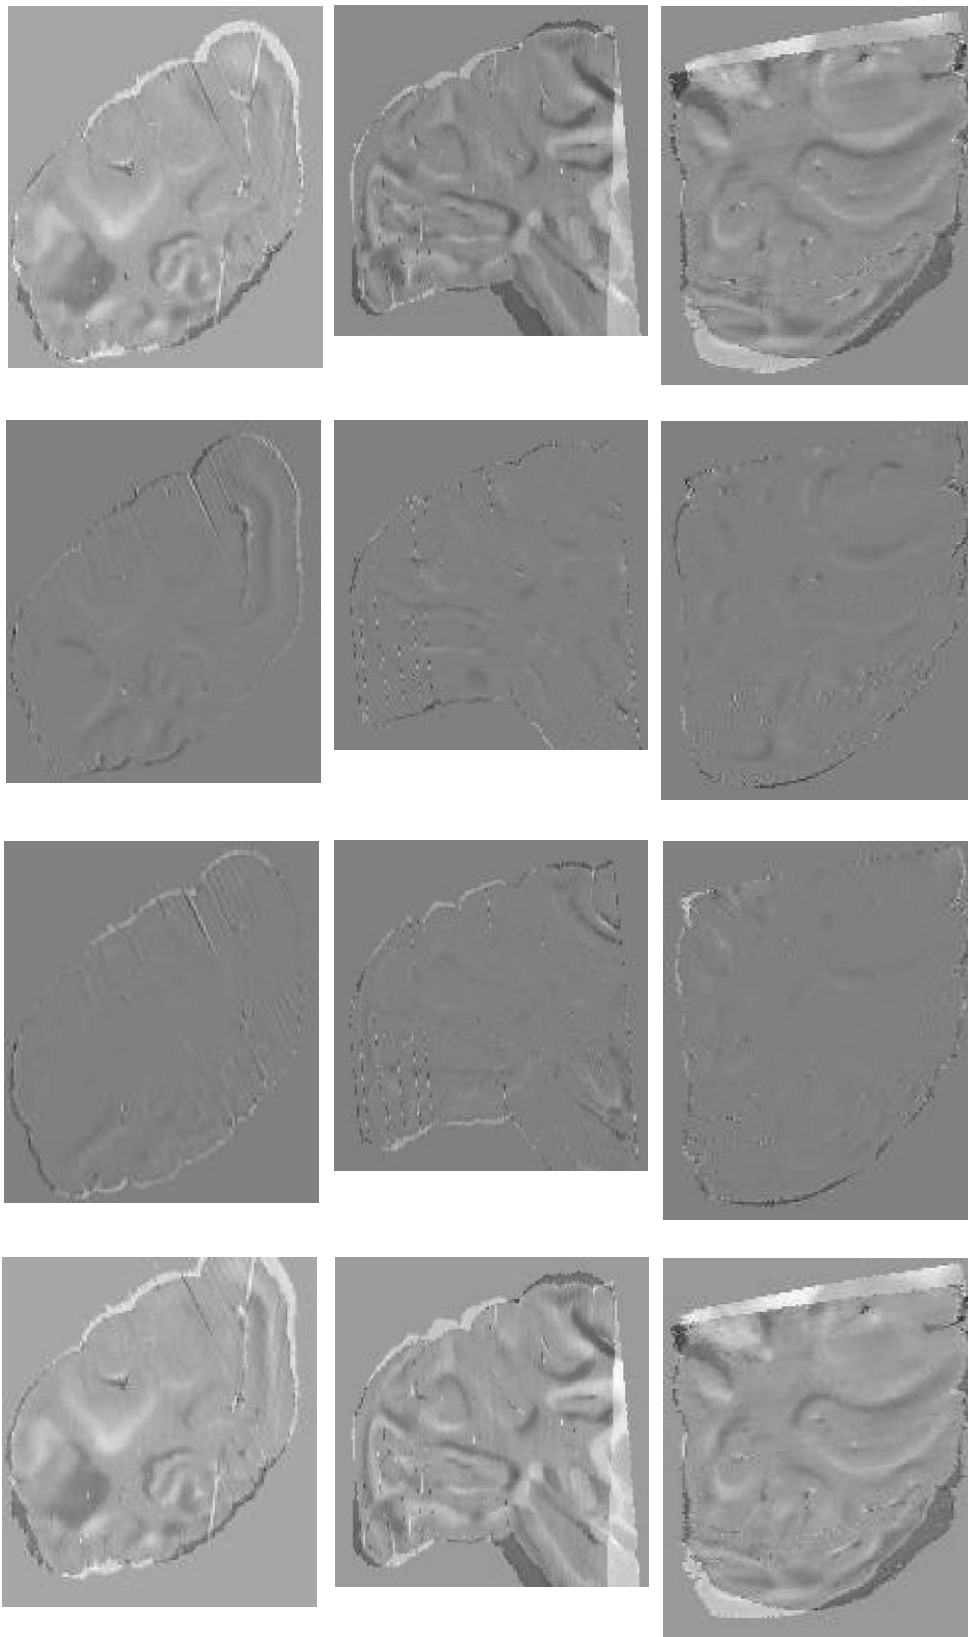
Fig. 12. Difference images between the successive autoradiographic volumes (AV) registered against the MR image. First row: AV after loop #1 minus first reconstructed AV. Second row: AV after loop #2 minus AV after loop #1. Third row: AV after loop #3 minus AV after loop #2. Fourth row: AV after loop #3 minus AV after loop #0.
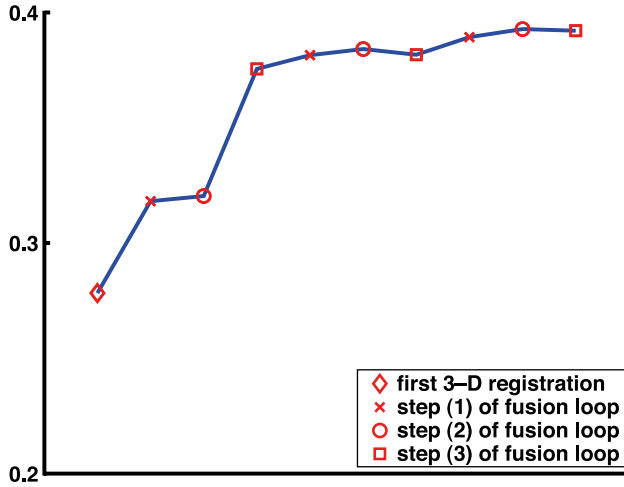
Fig. 13. Measure of the mutual information through the fusion procedure: it provides a mean to assess the benefits of the proposed method.

(Maintz and Viergever, 1998; van den Elsen et al., 1993). The method we choose, namely block matching, is particularly well suited for 3-D reconstruction as has been described and discussed in Ourselin et al. (2001), but can be used as well for 3-D registration (Ourselin et al., 2000).

This method is very similar to the ICP[2] algorithm (Besl and McKay, 1992) which consists in extracting feature points in the two images (say the reference and the floating images) to be registered and in iterating the following steps until convergence:

1. to pair each feature point of the floating image with the closest feature point in the reference image,
2. to compute the transformation that will best superimpose the paired points, and,
3. to apply this transformation to the feature points of the floating image.

Indeed, after applying the transformation, the pairings may have changed, thus a better transformation may be found by iterating these three steps.

In the block matching algorithm, we do not extract feature points but consider sub-images (i.e., blocks) in the floating image that will be paired to the most similar sub-image in the reference image. The computed transformation is the one that will best superimpose the centers of the paired sub-images. More details are given below.

A.1. Block matching

This algorithm takes as input a reference image I and a floating image J, and aims to estimate a transformation T such as $J \circ T$ can be superimposed on I. It is done through an iterative scheme: at each stage, correspondences are computed thanks to a block matching algorithm, and a transformation δT is estimated from these correspondences to allow an update of the searched transformation T.

As described, the transformation T is computed iteratively. Let us denote its estimation at step k by T_k (with $T_0 =$ Identity).

A block B (x, y) in image I is defined as the sub-image of I with upper left corner (x, y) and dimensions $d_x \times d_y$. For each block, we can compute its average value $\bar{B}(x, y)$ and its standard deviation $\sigma(x, y)$. We define the set B$_{a,b}$ of blocks by

$$\mathcal{B}_{s_x,s_y} = \{B(x,y) \subset I \text{ such that } x = as_x,$$

$$y = bs_y \text{ with } a, b \in \mathbb{N}\}$$

As it is, $\mathcal{B}_{1,1}$ contains all the possible blocks of size $d_x \times d_y$ included in I, while $\mathcal{B}_{s_x,s_y}$ with $s_x > 1$ or $s_y > 1$ contains only a subset of them.

While large values of s_x and s_y reduce the number of considered blocks in I, and thus decrease the computational cost, important features of I may be missed. We then choose a not too large s_x and s_y, but remove from the set $\mathcal{B}_{s_x,s_y}$ the blocks with the lowest standard deviation value.[3] We thus only keep a fraction p of blocks of $\mathcal{B}_{s_x,s_y}$, to compose the final set of blocks $\mathcal{D}_{s_x,s_y}$ of I. We have then $\mathcal{D}_{s_x,s_y} \subset \mathcal{B}_{s_x,s_y}$ and card $(\mathcal{D}_{s_x,s_y}) = p$card $(\mathcal{B}_{s_x,s_y})$.

Once the blocks to be considered in I are defined, we compute, for each of them their best corresponding block in $J \circ T_k$ (i.e., J resampled with T_k). To achieve this, we look for the block $B'(x', y')$ of upper left corner (x', y') and size $d_x \times d_y$ in $J \circ T_k$ that optimizes a given similarity measure S

$$B'(x',y') = \arg \max_{B'(u,v) \subset J \circ T_k} S(B(x,y), B'(u,v))$$

For practical reasons, the search for $B'(x', y')$ is limited in $J \circ T_k$ to a region of interest (exploration neighborhood), denoted by N(x, y), centered in (x, y), and defined by N$(x, y) = [x - r_x, x + r_x] \times [y - r_y, y + r_y]$.

$$B'(x',y') = \arg \max_{(u,v) \in N(x,y)} S(B(x,y), B'(u,v))$$

The choice of the similarity measure should depend on the expected relationship between the block intensities. Considering that a block may contain up to two (and rarely three) different tissues, an affine relationship seems reasonable. Hence, we choose the correlation coefficient as similarity measure (Roche et al., 2000):

$$S(B(x,y), B'(u,v))$$
$$= \frac{1}{d_x d_y} \frac{1}{\sigma(x,y)\sigma'(u,v)} \sum_{a=0\ldots d_x-1} \sum_{b=0\ldots d_y-1} [I(x+a, y+b)$$
$$- \bar{B}(x,y)][(J \circ T_k)(u+a, v+b) - \bar{B}'(u,v)]$$

where $\bar{B}'(u, v)$ and $\sigma'(u, v)$ are the average value and the standard deviation of the intensities of block $B'(u, v)$.

From each pair of corresponding blocks, $B(x, y)$ and $B'(x', y')$, we extract a pair of corresponding points, $C(x, y)$ and $C'(x', y')$: the centers of blocks $B(x, y)$ and $B'(x', y')$, respectively.

The optimal transformation δT_k between the images I and $J \circ T_k$ is the one that minimizes the residuals $\| C(x,y) - \delta T(C'(x', y')) \|^2$ in

---

[2] Iterative Closest Point.

[3] Blocks with low standard deviation are likely to contain only one type of tissue, yielding not reliable correspondences, while the ones with the highest standard deviation are more likely to contain different types of tissue, and yield more reliable correspondences.
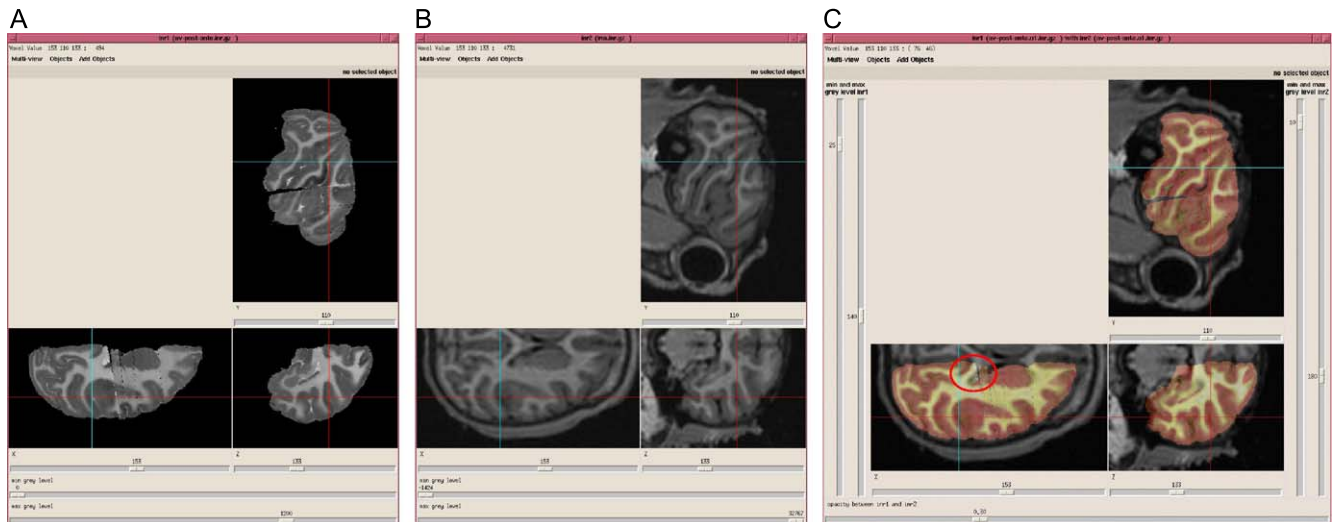
Fig. 14. Fused images can be presented in two different viewers, but with synchronized cursors, that is, at the same position in both viewers, (A and B), or in a single viewer with different color tables (C). Some subtle residual distortions, indicated by the circle and due to collapsed ventricles, can be seen.

a least squares sense. To reject outliers, a robust estimation with a weighted Least Trimmed Squares (LTS) was selected (Rousseeuw and Leroy, 1987). *M* estimators could have been chosen as well (Zhang, 1997).

After this estimation, the value of the transformation *T* is updated: $T_{k+1} = \delta T_k \circ T_k$. This iterative procedure stops as soon as no significant change occurs in the transformation evaluation ($\delta T_k < \varepsilon$).

A.2. Multi-scale implementation

Large transformations can only be captured by using both large regions of search (large values of r_x and r_y), and large blocks (to avoid local minima in the similarity measure), and thus at a high computational cost, while smaller transformations can be captured with smaller blocks and smaller regions of search.

Consequently, to allow for the capture of both small and huge displacements, we have implemented the block matching algorithm within a multi-scale strategy.

Each image is represented by a pyramid of *l* levels: from one level to the next, the image is subsampled by a factor 2 along every dimension. The registration is done at each level, while the initial transformation comes from the previous level.
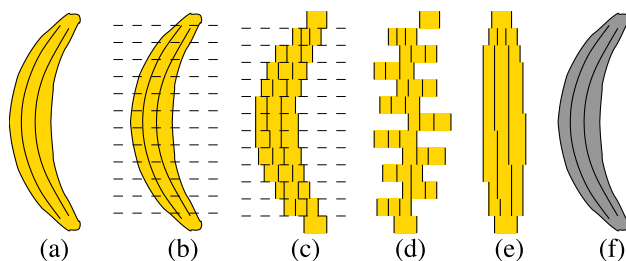

Fig. 15. The banana problem: the 3-D reconstruction of a 3-D curved object is not easy. (a) Take a 3-D curved object (e.g., a banana); (b) cut it into slices; (c) digitize the slices; (d) mix the digitized slices; (e) the 3-D reconstruction results in a cylindrical banana. (f) Using a shape prior (e.g., MRI) may help to reconstruct the curved banana.

Appendix B. Intensity compensation

This section addresses the problem of intensities alignment, that is, the correction of global intensity inhomogeneities from one 2-D image to the next. However,

- the histograms of two successive sections are not identical because the sections are intrinsically different and because of the noise, and
- there does not exist a one-to-one mapping between the intensities of two successive sections (the brighter one will use a larger interval of intensities).

To address these issues, we propose

- to use a continuous estimation of the intensity distribution (so that a one-to-one mapping can be easily built), and
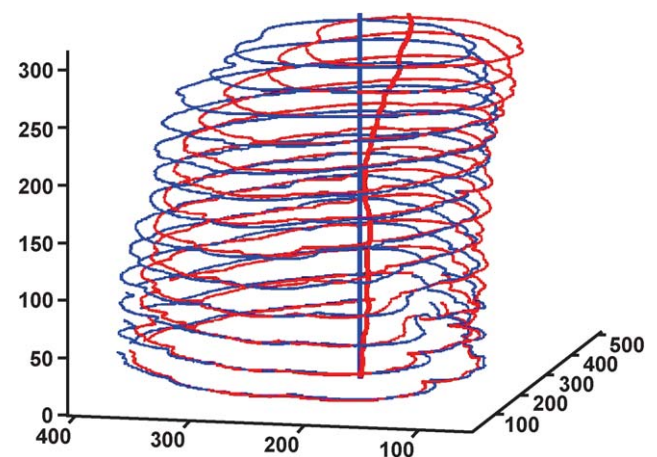

Fig. 16. The banana problem: illustration with the posterior block. In blue, some outer 2-D contours and the central line of the image after the first 3-D reconstruction (last row of Fig. 6). In red, the same after applying the fusion loops' transformations (left column of Fig. 9).

- to make the intensity distributions of two successive sections similar instead of strictly equal.

Some details on the proposed approach are given below, and more details can be found in Malandain and Bardinet (2003).

Consider two discrete histograms or equivalently two discrete probability density functions (PDF) $p(x_i)$ and $q(y_j)$, with $x_i, y_j \in \mathbb{Z}$, the aim of this section is to estimate an intensity transformation f such that the distribution p is *similar* to $q \circ f$. This problem is known as histogram matching.

Classically, this can be solved via histogram equalization. This consists in transforming one histogram into a flat histogram (Castleman, 1996) where all intensities have the same probability. Thus, matching two histograms can be achieved by (implicitly) using this flat histogram as intermediary (Hildebolt et al., 1996). However, this approach is not robust at all, as it acutely depends on the extremal values of the intensity histogram.
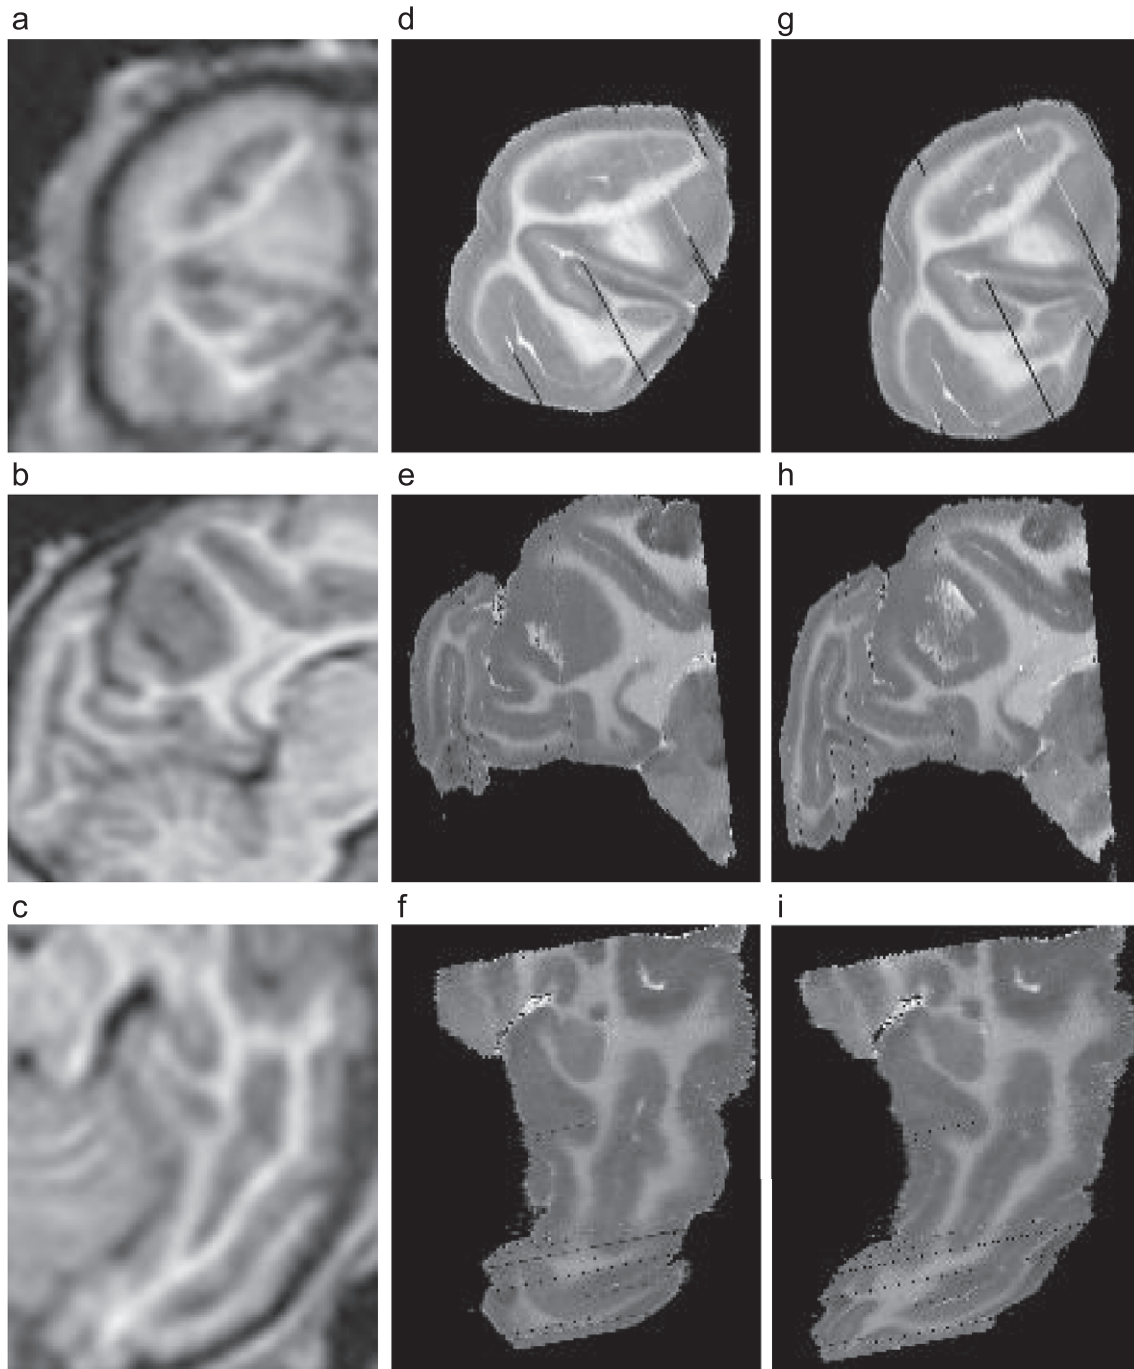

Fig. 17. From left to right: coronal, sagittal, and axial slices in the acquisition geometry of the MR image. Top row: the MR image (resampled with cubic splines, see Fig. 3). Middle row: the three same slices of the first reconstructed autoradiographic volume registered (affine transformation) against the MR image. Bottom row: the three same slices of the last reconstructed autoradiographic volume (after all fusion loops) registered against the MR image.

Moreover, we chose to estimate global parametric (e.g., affine or polynomial) intensity functions, with few degrees of freedom, since they are sufficient for our purpose. The above described approach does not allow to impose such a constraint: this also makes it not suitable.

Thus, our problem is as follows: given an intensity transformation function class $\mathcal{F}$, we are looking for the *best* function $\hat{f} \in \mathcal{F}$ such that $y = f(x)$:

$$\hat{f} = \arg \min_{f \in \mathcal{F}} S(p(x_i), q(y_j), f)$$

where S is some similarity measure of histograms.

Because of the discrete quantization of intensities (binning, etc), it is not possible to directly compare the discrete pdf $p(x_i)$ to $q \circ f(x_i)$. Instead, we will first estimate the continuous pdf corresponding to q, transform it with function f, and build from this transformed continuous pdf the discrete one corresponding to the x_i bins.

There exists many ways to estimate a continuous pdf from discrete samples or from an histogram. We chose the most common one, that is, the so-called Akaike–Parzen–Rosenblatt windowing technique (Akaike, 1954; Parzen, 1962; Rosenblatt, 1956).

The continuous pdf $Q_s(x)$ estimated from $q(y_j)$ is written

$$Q_s(y) = \sum_j q(y_j) K_s(y - y_j)$$

where $s > 0$ is a smoothing factor, K is a non-negative absolutely integrable function (called the kernel).

We chose a Gaussian kernel function. The smoothing factor is the standard deviation σ.

We then built the discrete pdf $q \circ f(x_i)$ as follows: for each bin x_i, we integrate $Q_s(y)$ in the interval $[x_i^{\min}, x_i^{\max}]$ that defines the bin x_i.

$$(q \circ f)_s(x_i) = \sum_j q(y_j) \int_{f(x_i^{\min}) - y_j}^{f(x_i^{\max}) - y_j} K_s(u) du$$

Because of the *smoothing* effect of the kernel, we do not compare $(q \circ f)_s(x_i)$ directly against $p(x_i)$ but against $p_s(x_i)$ with

$$p_s(x_j) = \sum_i p(x_i) \int_{x_j^{\min} - x_i}^{x_j^{\max} - x_i} K_s(x) dx$$

This formulation may however introduce some asymmetry in the criterion to be minimized. To overcome this issue, we modify it as follows:

$$\hat{f} = \arg \min_{f \in \mathcal{F}} (S(p_s(x_i), (q \circ f)_s(x_i)) + S((p \circ f^{-1})_s(y_j), q_s(y_j)))$$

Our practical choice for this application is to use affine function f. Indeed, the histogram of the 2-D image exhibits two peaks for, respectively, the grey and the white matter. To match two histograms, an affine function is thus sufficient. Minimization is done with a classical Powell-Brent procedure. We implemented several similarity measures (sum of squared differences (SSD), correlation, maximum of likelihood (ML)) which all yielded visually satisfactory results.

Appendix C. Transformations filtering

Here, we describe how to apply a low-pass filter to global parametric transformations (rigid or affine). This is used to add coherency between the independent 2-D transformations computed in the fusion loop.

Given a set of transformations $T(i)$ which belongs to a transformation class $\mathcal{T}$, we want to estimate a filtered transformation $\hat{T}(i) \in \mathcal{T}$ such that:

$$\hat{T}(i) = \arg \min_{T \in \mathcal{T}} \sum_j g(i - j) \text{dist}(T(j), T)^2$$

where g is some low pass filter (e.g., a Gaussian function), and dist represents a distance for T.

Such a computation is not straightforward (Pennec and Ayache, 1998) but can be managed with the Frechet expectation. Since this expectation can be approximated by the standard expectation near the origin (i.e., the identity), we use this property to compute the Frechet expectation with an iterative procedure that then needs a first estimate of $\hat{T}(i)$, for example, $\hat{T}_{(0)}(i) = T(i)$:

$$\hat{T}_{(k+1)}(i) = \hat{T}_{(k)}(i) \circ \left[\sum_j g(i - j) \left(\hat{T}_{(k)}^{(-1)}(i) \circ T(j) \right) \right]$$

It stops when $\sum_j g(i-j) (\hat{T}_{(k)}^{(-1)}(i) \circ T(j))$ is close enough to the identity. In the above formula, we must ensure that the weighted sum belongs to the transformation class $\mathcal{T}$, so we have to detail how to compute a weighted sum of transformations $\sum_j w_j T(j)$ *within* a given class of transformation $\mathcal{T}$.

- Such a computation is straightforward for affine transformations that can be represented by matrices $T(j) = [t_{u,v}(j)]$ in homogeneous coordinates since the weighted sum can be achieved on the matrices' elements.

$$\hat{T}(i) = [\hat{t}_{u,v}(i)] \quad with \quad \hat{t}_{u,v}(i) = \sum_j w_j t_{u,v}(j)$$

- For rigid transformations, we must choose another parameterization since a element-by-element sum of two matrices which both represent rigid transformations is generally not a rigid transformation. The rotation and translation vectors perfectly fit for that purpose. With $T(j) = (\mathbf{v}(j), \mathbf{t}(j))$, we then have

$$\hat{T}(i) = (\hat{\mathbf{v}}(i), \hat{\mathbf{t}}(i)) \quad with \quad \hat{\mathbf{v}}(i) = \sum_j w_j \mathbf{v}(j)$$

$$and \quad \hat{\mathbf{t}}(i) = \sum_j w_j \mathbf{t}(j)$$

References

Akaike, H., 1954. An approximation to the density function. Ann. Inst. Stat. Math. 6, 127–132.

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14, 239–256 (February).

Castleman, K.R., 1996. Digital image processing. Chap. Point Operations, International ed. Prentice Hall, New Jersey, pp. 83–97.

Cohen, F.S., Yang, Z., Huang, Z., Nissanov, J., 1998. Automatic matching of homologous histological sections. IEEE Trans. Biomed. Eng. 45 (5), 642–649.

Delingette, H., Bardinet, E., Rey, D., Lemarechal, J.-D., Montagnat, J., Ourselin, S., Roche, A., Dormont, D., Yelnik, J., Ayache, N., 2001 (Oct.). YAV++: a Software platform for medical image processing and visualization. Workshop on Interactive Medical Image Visualization and Analysis satellite symposia of MICCAI, IMIVA'01.

Deverell, M.H., Salisbury, J.R., Cookson, M.J., Holman, J.G., Dykes, E., Whimster, F., 1993. Three-dimensional reconstruction: methods of improving image registration and interpretation. Anal. Cell. Pathol. 5, 253–263.

Ford-Holevinski, T.S., Castle, M.R., Herman, J.P., Watson, S.J., 1991. Microcomputer-based three-dimensional reconstruction of in situ hybridization autoradiographs. J. Chem. Neuroanat. 4 (5), 373–385.

Geesaman, B.J., Born, R.T., Andersen, R.A., Tootell, R.B., 1997. Maps of complex motion selectivity in the superior temporal cortex of the alert macaque monkey: a double-label 2-deoxyglucose study. Cereb. Cortex 7 (8), 749–757.

Goldszal, A.F., Tretiak, O.J., Hand, P.J., Bhasin, S., McEachron, D.L., 1995. Three-dimensional reconstruction of activated columns from 2-[^{14}C]deoxy-D-glucose data. NeuroImage 2 (1), 9–20.

Goldszal, A.F., Tretiak, O.J., Liu, D.D., Hand, P.J., 1996. Multimodality multidimensional image analysis of cortical and subcortical plasticity in the rat brain. Ann. Biomed. Eng. 24 (3), 430–439.

Hess, A., Lohmann, K., Gundelfinger, E.D., Scheich, H., 1998. A new method for reliable and efficient reconstruction of 3-dimensional images from autoradiographs of brain sections. J. Neurosci. Methods 84 (1–2), 77–86.

Hibbard, L.S., Hawkins, R.A., 1988. Objective image alignment for three-dimensional reconstruction of digital autoradiograms. J. Neurosci. Methods 26 (1), 55–74.

Hildebolt, C.F., Walkup, R.K., Conover, G.L., Yokoyama-Crothers, N., Bartlett, T.Q., Vannier, M.W., Shrout, M.K., Camp, J.J., 1996. Histogram-matching and histogram-flattening contrast correction methods: a comparison. Dentomaxillofacial Radiol. 25 (1), 42–47.

Humm, J.L., Macklis, R.M., Lu, X.Q., Yang, Y., Bump, K., Beresford, B., Chin, L.M., 1995. The spatial accuracy of cellular dose estimates obtained from 3D reconstructed serial tissue autoradiographs. Phys. Med. Biol. 40 (1), 163–180.

Kay, P.A., Robb, R.A., Bostwick, D.G., Camp, J.J., 1996. Robust 3-D reconstruction and analysis of microstructures from serial histologic sections, with emphasis on microvessels in prostate cancer. In: Höhne, K.H., Kikinis, R. (Eds.), Visualisation in Biomedical Computing. Lect. Notes Comput. Sci., vol. 1131. Springer, Hamburg, Germany, pp. 129–134.

Kim, B., Frey, K.A., Mukhopadhyay, S., Ross, B.D., Meyer, C.R., 1995. Co-registration of MRI and autoradiography of rat brain in three-dimensions following automatic reconstruction of 2D data set. In: Ayache, N. (Ed.), Computer Vision, Virtual Reality and Robotics in Medicine. Lect. Notes Comput. Sci., vol. 905. Springer, Nice, France, pp. 262–266.

Kim, B., Boes, J.L., Frey, K.A., Meyer, C.R., 1997. Mutual information for automated unwarping of rat brain autoradiographs. NeuroImage 5 (1), 31–40.

Maintz, J.B.A., Viergever, M.A., 1998. A survey of medical image registration. Med. Image Anal. 2 (1), 1–36.

Malandain, G., Bardinet, E., 2003. Intensity compensation within series of images. In: Ellis, R.E., Peters, T.M. (Eds.), Medical Image Computing and Computer-Assisted Intervention (MICCAI 2003)LNCS, vol. 2879. Springer Verlag, Montreal, Canada, pp. 41–49.

Mega, M.S., Chen, S.S., Thompson, P.M., Woods, R.P., Karaca, T.J., Tiwari, A., Vinters, H.V., Small, G.W., Toga, A.W., 1997. Mapping histology to metabolism: coregistration of stained whole-brain sections to premortem PET in Alzheimer's disease. NeuroImage 5 (2), 147–153.

Muthuswamy, M.S., Roberson, P.L., Haken, R.K.T., Buchsbaum, D.J., 1996. A quantitative study of radionuclide characteristics for radioimmunotherapy from 3D reconstructions using serial autoradiography. Int. J. Radiat. Oncol. Biol. Phys. 35 (1), 165–172.

Ourselin, S. 2002. (Jan.). Recalage d'images médicales par appariement de régions-Application à la construction d'atlas histologiques 3D. Thèse de sciences, Université de Nice Sophia-Antipolis.

Ourselin, S., Roche, A., Prima, S., Ayache, N., 2000. Block matching: a general framework to improve robustness of rigid registration of medical images. In: DiGioia, A.M., Delp, S. (Eds.), Third International Conference on Medical Robotics, Imaging And Computer Assisted Surgery (MICCAI 2000)Lect. Notes Comput. Sci., vol. 1935. Springer, Pittsburgh, PA, USA, pp. 557–566.

Ourselin, S., Roche, A., Subsol, G., Pennec, X., Ayache, N., 2001. Reconstructing a 3D structure from serial histological sections. Image Vis. Comput. 19 (1–2), 25–31.

Parzen, E., 1962. On the estimation of a probability density function and the mode. Ann. Math. Stat. 33, 1065–1076.

Pennec, X., Ayache, N., 1998. Uniform distribution, distance and expectation problems for geometric features processing. J. Math. Imaging Vis. 9 (1), 49–67.

Pitiot, A., Malandain, G., Bardinet, E., Thompson, P., 2003. Piecewise affine registration of biological images. In: Gee, J.C., Maintz, J.B.A., Vannier, M.W. (Eds.), Second International Workshop on Biomedical Image Registration WBIR'03. Lect. Notes Comput. Sci., vol. 2717. Springer-Verlag, Philadelphia, PA USA, pp. 91–101. Also research report INRIA RR-4866.

Rangarajan, A., Chui, H., Mjolsness, E., Pappu, S., Davachi, L., Goldman-Rakic, P., Duncan, J., 1997. A robust point-matching algorithm for autoradiograph alignment. Med. Image Anal. 1 (4), 379–398.

Reisner, A.H., Bucholtz, C.A., Bell, G.A., Tsui, K., Rosenfeld, D., Herman, G.T., 1990. Two- and three-dimensional image reconstructions from stained and autoradiographed histological sections. Comput. Appl. Biosci. 6 (3), 253–261.

Roche, A., Malandain, G., Ayache, N., 2000. Unifying maximum likelihood approaches in medical image registration. Int. J. Imaging Syst. Technol. Spec. Issue 3D Imaging 11 (1), 71–80.

Rosenblatt, M., 1956. Remark on some nonparametric estimates of a density function. Ann. Math. Stat. 27, 832–837.

Rousseeuw, P.J., Leroy, A.M., 1987. Robust Regression and Outlier Detection. John Wiley & Sons, New York.

Rydmark, M., Jansson, T., Berthold, C.H., Gustavsson, T., 1992. Computer assisted realignment of light micrograph images from consecutive sections series of cat cerebral cortex. J. Microsc. 165, 29–47.

Schormann, T., Zilles, K., 1997. Limitation of the principal axes theory. IEEE Trans. Med. Imag. 16 (6), 942–947.

Schormann, T., Zilles, K., 1998. Three-dimensional linear and nonlinear transformations: an integration of light microscopical and MRI data. Hum. Brain Mapp. 6 (5–6), 339–347.

Schormann, T., Dabringhaus, A., Zilles, K., 1995. Statistics of deformations in histology and application to improved alignment with MRI. IEEE Trans. Med. Imag. 14 (1), 25–35.

Thévenaz, P., Blu, T., Unser, M., 2000. Interpolation revisited. IEEE Trans. Med. Imag. 19 (7), 739–758.

van den Elsen, P.A., Pol, E.J.D., Viergever, M.A., 1993. Medical image matching—A review with classification. IEEE Eng. Med. Biol. 12 (4), 26–39.

Vanduffel, W., Tootell, R.B., Orban, G.A., 2000. Attention-dependent suppression of metabolic activity in the early stages of the macaque visual system. Cereb. Cortex 10 (2), 109–126.

Vanduffel, W., Fize, D., Mandeville, J.B Nelissen, K., Van Hecke, P., Rosen, B.R., Tootell, R.B., Orban, G.A., 2001. Visual motion processing investigated using contrast agent-enhanced fMRI in awake behaving monkeys. Neuron 32 (4), 565–577.

Zhang, Z., 1997. Parameter estimation techniques: a tutorial with application to conic fitting. Image Vis. Comput. J. 15 (1), 59–76.

Zhao, W., Young, T.Y., Ginsberg, M.D., 1993. Registration and three- dimensional reconstruction of autoradiographic images by the disparity analysis method. IEEE Trans. Med. Imag. 12 (4), 782–791.