
Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images¹

Francois Rousseau, PhD, Orit A. Glenn, MD, Bistra Iordanova, Claudia Rodriguez-Carranza, PhD, Daniel B. Vigneron, PhD
James A. Barkovich, MD, Colin Studholme, PhD

Rationale and Objectives. This paper describes a novel approach to forming high-resolution MR images of the human fetal brain. It addresses the key problem of fetal motion by proposing a registration-refined compounding of multiple sets of orthogonal fast two-dimensional MRI slices, which are currently acquired for clinical studies, into a single high-resolution MRI volume.

Materials and Methods. A robust multiresolution slice alignment is applied iteratively to the data to correct motion of the fetus that occurs between two-dimensional acquisitions. This is combined with an intensity correction step and a super-resolution reconstruction step, to form a single high isotropic resolution volume of the fetal brain.

Results. Experimental validation on synthetic image data with known motion types and underlying anatomy, together with retrospective application to sets of clinical acquisitions, are included.

Conclusion. Results indicate that this method promises a unique route to acquiring high-resolution MRI of the fetal brain in vivo allowing comparable quality to that of neonatal MRI. Such data provide a highly valuable window into the process of normal and abnormal brain development, which is directly applicable in a clinical setting.

Key Words. fetal MRI; registration; 3D reconstruction; high resolution

© AUR, 2006

Ultrasound is the imaging modality of choice for the screening evaluation of the developing fetal brain. However, once an abnormality is suspected by routine prenatal ultrasound, further evaluation with MRI is indicated (1,2). Indeed, fetal MRI can detect additional abnormalities not detectable by prenatal ultrasound, including gyral and sulcal abnormalities (2–6).

MRI allows for excellent tissue contrast and spatial resolution, uses no ionizing radiation, permits imaging in


multiple planes, has a large field of view, and allows for visualization of fetal anatomy in a manner not possible with ultrasound. The development of ultrafast two-dimensional (2D) acquisition sequences has led to significant improvements in the clinical utility of fetal MRI (7, 8). However, the slice acquisition time is still very critical and has to be as short as possible to reduce the impact of fetal and maternal motion on the exam and for the comfort of the mother, since fetal MRI is often performed without sedation. As a result, sets of thick 2D slices are generally acquired in clinical studies, with motion commonly occurring between slices. Overall, the resulting image data are limited in its geometric integrity between slices due to motion and in its through-plane spatial resolution. In addition, body surface coils used in the imaging processes do not provide a homogeneous sensitivity over the entire field of view and, combined with motion during the acquisition, can then produce changes over time in the

Acad Radiol 2006; 13:1072–1081

¹ Is now with the University Louis Pasteur, LSIT, Pole API, Bd S. Brant, Illkirch, France. Received January 31, 2006; accepted May 5, 2006. This work was supported by a Whitaker Foundation award (RG-01-0115), NIH grant R01-MH65392, and NIH Biomedical Research Partnership grant R01-EB0822. **Address correspondence to:** R.S. e-mail: rousseau@lsitit.u-strasbg.fr

© AUR, 2006

doi:10.1016/j.acra.2006.05.003



Low resolution MR image *i*

Figure 1. Coordinate systems used during the slice and image registration. Two rigid transformations are considered for each slice: T_{2k}^i between the slice k and the low-resolution image i and T_1^i between the low-resolution image i and the high-resolution volume.

distortion profile for the same anatomical region (see Fig. 3). This inhomogeneity is due to the fact that the gradients are linear over central region but then fall off along the length of the scanner. There are also regions at both ends where the gradient strength is the same as in the center.

This work is motivated by the observation that current clinical imaging protocols make use of multiple orthogonal 2D multiplanar acquisitions to study the fetal brain. However, because of the motion between the slices, current techniques do not allow direct quantitative measurements to be made on the 3D anatomy. The aim of this work is to develop and apply registration-based reconstruction methods to the problem of correcting motion and intensity distortions between the different 2D acquisitions and to then reconstruct a single higher-resolution MR image of the fetal brain.

Super Resolution Methods

The principle of super-resolution is to combine low-resolution images to produce an image that has a higher spatial resolution than the original images (9). This is a large research field encompassing many applications; however, the majority of the work has focused on using lower-resolution data acquired on a regular grid and often assuming simple translational motion between the lower-resolution sample grids, unlike our data, which are corrupted by full three-dimensional (3D) rigid motion on a slice-by-slice basis.

Overall, published methods can be categorized into two main divisions: frequency domain and spatial do-




Figure 2. Optimization example for slice registration exploiting interleaved acquisition. The slices are grouped by time (in this example, an image of eight slices is considered). At the first step, the optimization is performed for two groups of four slices each. At the next step, each group is divided into two groups. At the end, the optimization is performed for each slice independently.




Figure 3. Acquired fetal MR image data (top three rows) with a resolution of $1 \times 1 \times 3$ mm. It can be noted that strong motion artefact corrupt the coronal acquisition (second row). Lower row shows the resulting reconstructed high-resolution image (resolution $1 \times 1 \times 1$ mm).

main. Frequency domain methods are based on Fourier transform properties. Denote a continuous 2D image $f(x, y)$ and n shifted images $f_r(x + \delta x_r, y + \delta y_r)$. Let the set $\{f_r(i, j)\}$, with $i \in \{1 \dots M\}$, $j \in \{1 \dots N\}$ and $r \in \{1 \dots n\}$ be the discrete version of the set $\{f_r(x, y)\}$ using an ideal uniform sampling with periods T_x and T_y . The continuous Fourier transform (CFT) F_r of the continuous image and the discrete Fourier transform (DFT) Y_r of the shifted and discrete images are related via aliasing:

$$Y_r(m, n) = \frac{1}{T_x T_y} \sum_{i=-\infty}^{+\infty} \sum_{j=-\infty}^{+\infty} F_r\left(\frac{2\pi}{T_x}\left(\frac{m}{M} + i\right), \frac{2\pi}{T_y}\left(\frac{n}{N} + j\right)\right). \tag{1}$$

Assuming that $f(x, y)$ is band limited and using the shifting property of the Fourier transform, Equation 1 can be rewritten in matrix form as

$$Y = \Phi F \quad (2)$$

where Y is a vector containing the DFT coefficients $Y_r(m, n)$ and F is a vector containing the unknown CFT of $f(x, y)$. The super-resolution restoration requires finding the DFTs of the observed images, determining the motion matrix Φ , solving the system of Equations (2) for F , and applying the inverse DFT to compute the reconstructed image. The key advantages of frequency domain methods are their theoretical simplicity and the low computational complexity. Their disadvantages include the limitation to global translational motion, a spatially invariant degradation model, and limited ability to vary the point spread function over the image.

In spatial domain methods, the most common approach is to model the physical problem and to invert it. These approaches use a generic observation model of the form:

$$\mathbf{y}_r = DBW_r \mathbf{z} + \mathbf{n}_r \quad \text{for } 1 \leq r \leq n \quad (3)$$

where

- \mathbf{y}_r , $r \in \{1 \dots n\}$ denotes the low resolution images,
- \mathbf{z} is the high resolution image,
- \mathbf{n}_r represents observation noise,
- D is the subsampling matrix,
- B a blur matrix,
- W_r is the geometric transformation of r th low resolution image.

The three operators can be combined into a single matrix H_r : $H_r = DBW_r$. The matrix H_r thus incorporates motion compensation, degradation effects, and subsampling.

The problem is usually decomposed into two parts: motion estimation (estimating the matrix H_r for each \mathbf{y}_r) and super-resolution reconstruction (estimating \mathbf{z} using the observation model). Accurate motion estimation, which is a critical aspect of super-resolution methods, can be performed using block matching techniques or optical flow methods for instance. (Please see (10, 11) for further details about motion estimation.) The super-resolution reconstruction can be performed using Simulate and Correct methods (12, 13), Projection Onto Convex Sets methods (14), and Bayesian methods (15, 16). One can note that in (16), motion estimation and the high-resolution image are jointly estimated. These spatial domain methods offer important advantages in terms of flexibility but they can be computationally very expensive. Last, another spatial

domain approach is that of Interpolation-Restoration (17). In this case, a super-resolution image is reconstructed from low-resolution images using techniques for reconstruction from non-uniformly spaced samples, and restoration techniques can be applied to compensate for degradations. Such a general approach is suited to our problem of slice-by-slice motion corruption, and we will take this direction in our work.

Application to MRI

Building a high-resolution 3D MR image of the fetal brain is challenging because of the original 2D slice thickness, the intensity distortion, and the unpredictable fetal motion that can occur in any direction and can be as much as several centimeters in distance between slice acquisitions. As far as we know, this problem, which has been discussed before in the literature (18, 19), is still an open issue. Super-resolution has previously been investigated in MRI using a specialized protocol to acquire shifted images (20, 21), where the motion between slices is a known translation in the slice direction. Moore et al. (22) built a high-resolution dynamic heart model from isotropic coronal slices. The misalignment was corrected by registering a volume with sagittal and axial scout images. Image intensities were averaged together to generate a high-resolution volume. Lotjonen et al. (23) proposed an algorithm to correct translation-only movement artifacts simultaneously from short- and long-axis cardiac MR series. Locations of short-axis slices were optimized based on data from long-axis slices and vice versa. No reconstruction of a higher resolution volume of the heart was done.

The task of compounding multiple orthogonal sets of low-resolution fetal MRI slices into a single high-resolution dataset shares a number of common features with the published methods in the super-resolution literature but presents significant additional problems. Specifically, we have multiple orthogonal datasets that have nonisotropic voxel size, we have subject motion during slice acquisitions so that each 2D slice is misaligned with respect to others in the acquisition, and the resulting stack of low-resolution 2D slices does not have a consistent 3D geometry. This movement can also mean a loss of spatial information between slices, in regions of anatomy missed in the acquisition due to the motion. The fully 3D motion artifacts make a manual motion correction procedure unfeasible. Moreover, because the intensity inhomogeneity of the MRI acquisition, the movement of the fetus within

the field of view can induce large-scale changes in the illumination pattern between acquisitions.

In this article we report our work on this challenging problem and describe a fully automatic and accurate algorithm to correct slice misalignments, correct relative intensity distortions, and reconstruct a single high-resolution image from sets of clinically typical low-resolution fetal MRI data.

METHODS

Our overall aim is to reconstruct an image of the fetal brain that has increased and isotropic sampling resolution, increased overall spatial resolution (reduced point spread function), and consistent tissue intensity over the field of view. The resolution of the source data is typically 1×1 mm in plane with 3-mm-thick slices. Multiple sets consisting of between 20 and 40 slices each are acquired over a period of around 20 seconds for each set. The slices are commonly acquired in an interleaved pattern (usually alternating odd-even), and significant movement of 1 centimeter or more can occur during the acquisition of one set of slices. If the fetus moves quickly at the wrong time, single slices can be corrupted due to motion. This is relatively rare, but typically in the clinical protocol, these images are visually detected and the acquisition repeated, until at least two nonmotion degraded sets of axial, sagittal, and coronal slices of the fetal brain are obtained. It is therefore reasonable to assume (and has been confirmed visually) in this type of data that motion corrupted slices are rare within the data we have to reconstruct. However, the algorithm could be applied with slices removed where within slice motion occurs, by making use of the other orthogonal datasets to fill in the missing structure.

We posed the task of recovering the local relative alignment of each slice in three dimensions but can make use of the consistent geometry provided in the through-plane direction of a slice that is provided by the other orthogonal sets of slice data, to constrain the collective alignment of the multiple low-resolution images.

The proposed method is made up of three steps: motion correction, volume reconstruction, and contrast correction.

Notations.—Let us consider n low-resolution 3D images, denoted by I_{LR}^i , $i \in \{1 \dots n\}$. A so-called low-resolution 3D image is a stack of 2D thick slices. Such a notation is used in order to keep in mind that the low-resolution

source data are not geometrically consistent. The reconstructed high-resolution volume is denoted by I_{HR} .

Motion Correction

In order to correct fetal motion, we propose to use a registration-based method. Any registration method consists of three components (24): a similarity measure between the two images, a transformation relating the coordinate system of one image to that of another image, and an optimization procedure to improve the transformation between the two images with respect to the similarity measure.

Registration methods are usually classified as being either sparse (feature) or dense (voxel) based. In the case of fetal brain MRI registration, because of the 3D non-consistency of the scans, finding homologous 3D landmarks is a very challenging task and we have therefore approached the problem using a voxel-based method. Registration methods rely on the fact that the similarity measure is optimal when the two images are perfectly aligned. The choice of one measure relies on the assumptions relating intensities in the two images: an identity relationship, an affine relationship, a functional relationship, and a statistical relationship. Due to the motion and contrast artifacts corrupting fetal MRI, it is highly desirable to use a similarity measure that is not based on strong assumptions between the two images. Mutual information is a measure of the statistical dependence between two random variables (25, 26). The mutual information $MI(X, Y)$ is defined as follows:

$$MI(X, Y) = h(X) + h(Y) - h(X, Y) \quad (4)$$

where the marginal entropy $h(X)$ is $h(X) = -\int f_X(x) \times \log f_X(x) dx$ and where the joint entropy $h(X, Y)$ is $h(X, Y) = -\int f_{XY}(x, y) \log f_{XY}(x, y) dx dy$. To reduce the sensitivity to changes in overlap, we use the normalized mutual information (27) which is:

$$NMI(X, Y) = \frac{h(X) + h(Y)}{h(X, Y)}. \quad (5)$$

Maximization of $NMI(X, Y)$ seeks a transformation where joint entropy is minimized with respect to the marginal entropies.

Motion correction is performed by aligning globally the low-resolution images together and then by aligning every slice of the low-resolution images to the reconstructed high-resolution volume.

Low-resolution image registration.—The first step consists of approximately globally aligning the n low-resolution images. One low-resolution image is chosen arbitrarily as the reference image and is used to define a global coordinate system. The other low-resolution images are rigidly registered to the reference image (6 degrees of freedom [DOF]: three translations and three rotations). The 3D rigid transformation between an image I_{LR}^i and the reference image is denoted T_1^i . The high-resolution volume coordinate system is chosen to match with the global coordinate system defined previously. Figure 1 illustrates the two transformations for each slice and the different coordinate systems used during the slice and image registration.

Slice registration.—Once the low-resolution images are roughly globally aligned, slice motion artifacts are corrected. The low-resolution image acquisition is interleaved so that odd and then even slices are acquired in sequence. As the fetus can move during the exam and the observed structures are small, consecutive slices may present completely different anatomical structures. Therefore, we cannot register the slices using, for instance, a correlation assumption between consecutive slices. Slice motion is therefore corrected by rigidly registering a slice to the current reconstructed high-resolution volume. This high-resolution volume is built using the current estimate of slice positions.

We register the slices of low-resolution image I_{LR}^i with the volume reconstructed using the $n - 1$ other low-resolution images I_{LR}^j , $j \in \{1 \dots n, j \neq i\}$. A local coordinate system is defined for every slice and the transformation between a slice and the high-resolution volume is composed from the two rigid transformations:

$$\mathbf{y}_k^i = T_1^i T_{2_k}^i \mathbf{x}_k^i \quad (6)$$

where $T_{2_k}^i$ denotes the 3D rigid transformation (6 DOF) from the slice k of the low-resolution image i to low-resolution image i coordinates, \mathbf{x}_k^i is the slice pixel coordinates and \mathbf{y}_k^i the voxel coordinates in the compounded high-resolution volume (see Fig. 1).

Optimization.—For both image and slice registration, a gradient ascent method is used to maximize the normalized mutual information. Concerning slice registration, we exploit the interleaved acquisition to propose a hierarchical approach to provide a starting estimate to the optimization procedure. Slices of a low-resolution image are ordered by acquisition time and are separated into two groups. The optimization consists in finding the optimal

rigid transformations T_2 between the new groups of slices and the low-resolution image coordinate system. The similarity measure is computed between the current reconstructed volume and the group of slices. This is performed for the n low-resolution images. Once the convergence criterion is reached, each group is divided into two groups and the optimization procedure is then performed on these new groups of slices (Fig. 2).

For the final phase of alignment, we maximize independently the normalized mutual information between each slice and the current reconstructed high-resolution volume. This uses an iterative slice-by-slice scheme, looping over all slices until the sum of the slice-to-volume criteria (called the global convergence criterion) does not increase. For each 2D-3D registration step, a two-level multiresolution strategy is applied to avoid the pitfall of local optima.

Volume Reconstruction

The reconstruction problem can be classified as scattered interpolation (28). As a volume reconstruction is required repeatedly during the registration and contrast correction steps, the reconstruction algorithm must be computationally efficient. To deal with the computation time constraint, a local neighborhood approach has been adopted. Local methods are based on the assumption that each point influences the resulting volume only up to a certain finite distance.

For each pixel \mathbf{x} in the slices of the original low-resolution scans, the 3D coordinates \mathbf{y} of \mathbf{x} in the high-resolution volume are calculated taking into account the results of the image and slice registration. The intensity value of \mathbf{x} is injected into the reconstructed volume using a kernel f centered at \mathbf{y} . If the kernel is a narrow boxcar function, the reconstruction method approximates the pixel nearest neighbor (PNN) method (29). If the kernel is a pyramid function, the method is referred to pixel trilinear (PTL) interpolation (30). Our choice of kernel is motivated by need to match the underlying point spread function of the data. In the through slice direction this can be modeled as a Gaussian (31). The Gaussian kernel is perpendicular to the slice plane, and the standard deviation determining the profile is chosen to match the known acquisition profile. Within a slice plane of the MRI, the point spread function can often be modeled as a sinc, but for simplicity we have used a narrow Gaussian. Local methods such as these have the advantage of making it possible to update the compounded image incrementally within a limited spatial extent, thus requiring lower computational burden.

Using local reconstruction methods with a truncated kernel for speed, gaps can occur in the reconstructed volume when sudden significant motion occurs, and there is no data contributing to a given reconstructed point. Empty voxels are filled using the average of the surrounding voxels. This local approach has the advantage that it can be implemented very efficiently. In practice, using an oriented Gaussian kernel, we found that gaps do not occur using three or more orthogonal low-resolution images.

Contrast Correction

We use a contrast correction step to correct the local relative intensity distortion between the low-resolution images to allow accurate compounding. Intensity changes from one 2D slice to another are both global and local. To address this issue, one low-resolution image is used as a reference for tissue intensities and the other low-resolution images are corrected to it. We consider the relationship between contrasts in the low-resolution image to be corrected and the reference image as a spatially varying multiplicative field that is assumed to vary smoothly over the field of view. A direct approach to estimation is used where the two volumes are low pass filtered with a Gaussian kernel G . The relative scaling in intensity $\beta_i(\mathbf{x})$ at location \mathbf{x} of the low-resolution image $I_{LR}^i(\mathbf{x})$ is estimated for all slice pixels \mathbf{x} as follows:

$$\beta_i(\mathbf{x}) = a_i * \frac{G(IR(\mathbf{x}))}{G(I_{LR}^i(\mathbf{x}))} \text{ with } a_i = \frac{\sum_{\mathbf{x}} I_{LR}^i(\mathbf{x})}{\sum_{\mathbf{x}} I_{LR}^i(\mathbf{x}) \frac{G(IR(\mathbf{x}))}{G(I_{LR}^i(\mathbf{x}))}} \quad (7)$$

where IR is the intensity reference image, I_{LR}^i is the low-resolution image considered, and a_i is used to keep the mean value of the image I_{LR}^i unchanged. IR is an arbitrary low-resolution image projected in the reconstruction space. The corrected low-resolution image \hat{I}_{LR}^i is then: $\hat{I}_{LR}^i(x) = \beta_i(x) I_{LR}^i(x)$.

Overview of the Algorithm

The proposed method consists in the following steps:

1. Apply global rigid registration between all the low-resolution images.
2. Create initial slice registration estimate by hierarchical application of slice registration algorithm

3. For each resolution:
 - (a) Apply slice registration algorithm
 - (b) Test global criteria reached
4. Apply relative intensity distortion correction
5. Reconstruct the final high-resolution volume.

Where the slice registration algorithm is the following:

1. For each low-resolution image I_{LR}^i :
 - (a) Reconstruct a high-resolution volume using the low-resolution images $I_{LR}^j, j \in \{1 \dots n, j \neq i\}$,
 - (b) For each slice in I_{LR}^i , perform the 2D-3D registration.

Material and Evaluation Procedure

Simulated motions on real MR datasets.—Initially we want to explore the ability to recover slice to volume alignment for typical but known motion, on typical anatomical structures. In order to perform the most realistic simulation possible, we used as a starting point 3D MRI acquisitions of four premature neonates with high quality and no motion. The postconceptional age of the premature neonates at the scan acquisition was 28, 32.6, 34.3, and 42.6 weeks. T1-weighted SPGR images were acquired on a 1.5-T GE MRI scanner with a dedicated MRI compatible incubator and had a spatial resolution of $0.7 \times 0.7 \times 2.1$ mm. The 3D image was interpolated to isotropic voxels using B-spline interpolation. Figure 4 shows coronal views of the four premature neonates.

In order to evaluate the slice registration algorithm, anisotropic subsampled volumes with 6 DOF motion artifact were simulated from the isotropic neonatal image. Motion artifacts were defined by the maximum magnitude of the head motion and by the head motion type. The first simulated motion was a sudden motion. It simulated a short fast movement of the head during a short portion of the acquisition. The model used for the simulation incorporates the interleaved aspect of acquisition. The second type of motion was a temporally smooth motion simulated using B-splines. In both simulations, some information is lost due to missed portions of the underlying anatomy. The displacements were chosen from a uniform distribution with a varying range for $[-x, x]$ mm translations in each direction and between $[-x, x]$ degrees for each rotation, $x \in \{1, 5, 10\}$. Compared with the size of fetal brain, a head motion of ± 10 mm in translation is a significant fraction of the diameter of the brain.




Figure 4. Isotropic high-resolution premature neonatal MRI datasets used for the simulations showing large range of anatomy. The ages at exam are (a) 28 weeks, (b) 32.6 weeks, (c) 34.3 weeks, and (d) 42.6 weeks.

Figure 5 shows one set of isotropic premature neonate data and two examples of simulated low-resolution images using random motions (sudden motion and smooth motion with a maximum magnitude of 10 mm for the translation and 10 degrees for the rotations). These simulations from a known underlying anatomy allow us to assess both the accuracy and the robustness of the slice registration process. Considering n simulated low-resolution images already globally aligned, the hierarchical and the multiresolution strategies are applied for the slice registration. Evaluation of the algorithm has been done using 24 simulations.

The accuracy was assessed by computing a registration error measured on a set of four points P_i distributed

$$RMS = \sqrt{\frac{1}{n} \sum_{i=1}^n TRE_i^2},$$

where TRE is the target registration error defined as follows: $TRE_i = \|P_i - \hat{T}^{-1}(T^*(P_i))\|^2$. T^* denotes the known applied motion transformation and \hat{T} is the estimated geometric transformation. P_i forms a square of 100 mm by 100 mm in which the fetal brain is included. The error previously defined provides thus a maximum bound of the registration error for the region of interest.




Figure 5. (a) Original isotropic high-resolution premature neonatal brain MRI. (b) Simulated low-resolution volume using sudden motion and corresponding high-resolution reconstructed volume. (c) Simulated low-resolution volume using smooth motion and corresponding high-resolution reconstructed volume. The maximum magnitude of the simulated motion was 10 (see text for more details).

Fetal datasets.—A typical clinical study of a fetus at our institute follows the following procedure: a quick localizer sequence is obtained in less than 30 seconds during maternal free breathing in order to determine the location of the fetal head. No sedation or contrast agents were administered during the study, and this study was approved by the UCSF IRB. Conventional single-shot fast spin-echo (SSFSE) T2-weighted images are then acquired during normal maternal breathing (a single slice is acquired in less than 1 second). Sets of contiguous slices were acquired in the axial, sagittal, and coronal planes with respect to the fetal brain, based on the initial localizer. All images were acquired in an interleaved manner to reduce saturation of spins in adjacent slices due to fetal motion during the sequence. The sequence parameters were TR = 6000, TE effective = 90 ms, field of view




Figure 6. Simulation results for premature neonate MR images: starting RMS errors versus final RMS errors in millimeters using two different simulated motions (sudden motion and smooth motion). Each of the 24 simulations is represented by a point. For all cases, final RMS errors remained below 1 mm (please note the difference of scaling between starting RMS and final RMS).

(FOV) = 240 mm, in plane resolution about 1-mm, 3-mm slice thickness. The reconstruction method has been performed on a database of 17 subjects.

RESULTS

From the simulated motion experiments on premature neonatal data, we evaluated the RMS registration error for four points at the corners of a box within the brain tissue of size 100 mm \times 100 mm for each slice. These are presented in Figure 6. The simulated starting RMS error is up to 12 mm. Each point represents the result of one simulation and its coordinates are the starting RMS error and the final RMS error. For all cases, the final overall slice alignment error was significantly reduced by the alignment process. Overall, even with the presence of strong motion artifacts and significantly varying developmental age, RMS errors remained below 1 mm. In all the cases, visual inspection of the obtained results was also very satisfactory (see Fig. 5). These experiments show the robustness of the proposed method to significant fetal head motion and demonstrate that the stage of development of the underlying brain anatomy does not affect the registration accuracy.

We then applied the algorithm to clinically acquired fetal MRI datasets using the same optimization param-




Figure 7. High-resolution fetal MR image (second row) reconstructed from low-resolution images (first row, coronal acquisition). Third row: Surface rendering of semiautomatic brain surface segmentation performed on original low-resolution image and on high-resolution reconstructed image.

ters. Each dataset consists of three orthogonal low-resolution images (axial, coronal, sagittal). Figure 3 shows the high-resolution reconstructed volume compared to the three original low-resolution images. It clearly shows the quality improvement in terms of resolution and motion correction provided by the proposed compounding method, even with the presence of significant motion artifacts.

Figure 7 shows one low-resolution image and the high-resolution reconstructed image, and the corresponding segmented brain from the original coronal low resolution and the high resolution. The surface rendering of semi-interactive brain surface segmentation has been performed using the rview software (www.rview.colin-studholme.net). This figure clearly illustrates the importance of a geometrically consistent and intensity consistent image when interpreting cortical development. In the compounded image, early stages of the formation of the temporal and frontal gyri are seen. In addition, in the original data one can note that when the fetal brain is not entirely scanned in one of the sets, but compounding multiple image sets where the extremities were scanned, allows recovery of the whole brain.

Finally, Figure 8 shows the results of the proposed method at three different time points for the same subject (25, 30, and 34 weeks). Overall, the results in terms of




Figure 8. High-resolution reconstructed images at three different time points for the same subject (25, 30, and 34 weeks). This result illustrates that whatever the age of the subject, the quality of the reconstruction volumes is visually very satisfactory.

contrast and resolution are visually very satisfactory, whatever the age of the subject. Such quality of the reconstructed images allows us to study the early stages of the formation of the cortex.

DISCUSSION

The ability to study the developing fetal brain in high-resolution promises to provide a vital source of clinical information that could contribute directly to a number of challenging clinical questions. It will permit the use of many quantitative morphometric analysis methods, originally developed to study the adult (32) and neonatal (33,34) brain, to be applied to examine the process of in utero brain development. Critically, high-resolution imaging is the key to seeing the process of cortical folding (in utero (35,19) and post mortem (36)), while consistent contrast allows us to study patterns of myelination. Such data are valuable both to specific clinical questions and, more fundamentally, to an improved understanding of the process of human brain development.

In this paper we have described a new methodology to reconstruct in utero high-resolution 3D MRI scans of the fetal brain by proposing to compound multiple sets of orthogonal 2D MRI slices. The method makes use of a novel combination of 2D to 3D registration, relative inhomogeneity correction, and high-resolution reconstruction from sparse data.

This is achieved by first globally registering the low-resolution images and then applying an iterative slice alignment scheme that seeks to refine the 3D positioning of each slice to the current combined high-resolution volume. This step is embedded in a spatiotemporal acquisition-ordered hierarchical scheme to recover the fetal brain trajectory, which uses multiple image resolution refinement at each step. This is driven by normalized mutual information to provide robustness to contrast variation induced by motion of the fetal brain with respect to the imaging coil in the magnet. As a final step, a relative intensity correction is applied between the low-resolution images to remove the differences in relative signal strength across the different acquisitions in each region of the fetal brain.

In the proposed approach, an arbitrary low-resolution image is used as a reference image for the registration step and for the intensity correction method. Since this reference is only used for the estimation of the initial global rigid transformation, the choice of image does not have a significant impact on the final registration result. The choice of the reference image for the intensity correction step is also not critical, since the purpose of this part of the algorithm is simply to correct the relative contrast variations between the acquisitions for compounding. An absolute intensity correction can then be applied after reconstruction. One assumption we have made is that fetal motion within the time period of one slice is negligible. Since this assumption is reasonable in most of the cases, it may happen that one slice is obscured by motion artifact during the time of imaging. In such case, the corrupted slice can be manually removed from the stack of slices. Further work concerning this aspect would be to remove automatically corrupted slices using intensity information between consecutive slices.

We have used a kernel-based local approach for the reconstruction step of the algorithm. This approach can handle data with anisotropic resolution and is computationally efficient. The overall algorithm thus takes less than 20 minutes to register three low-resolution images and to reconstruct a high-resolution volume, running on a 3-GHz Pentium 4 CPU. Although the proposed method allows us to successfully recover the 3D geometry of the fetal brain, there is still a slight loss of contrast in the reconstructed images. The reconstruction issue can be seen as an approximation problem with sparse data and anisotropic points. A possible way to improve the final quality of the reconstructed volume could be to develop a dedicated approximation method fast enough to handle

these specific data. Moreover, further work is needed to evaluate the influence of the different sources of errors on the reconstructed volume quality: registration errors, intensity contrast variations, and the reconstruction method. An elegant manner to reconstruct high-resolution fetal MR volume would be to include all these sources of errors in an unified framework.

Visual inspection of the obtained results on clinical data is very encouraging. The developed algorithm automatically reconstructs a 3D high-resolution geometrically consistent volume, and it has proved to be robust to large artifacts. This represents an important step toward 3D quantitative analysis of the fetal brain. Further work is needed to better understand the capabilities and limitations of the approach on a range of fetal motion profiles and image quality. Moreover, in the future, we would like to investigate the use of the proposed method to study other fetal organs and also for other modalities such as functional MRI or diffusion tensor imaging when images are motion corrupted.

REFERENCES

- Levine D. Fetal magnetic resonance imaging. *J Maternal-Fetal Neonat Med* 2004; 15:85–94.
- Coakley FV, Glenn OA, Qayyum A, et al. Fetal MR imaging: A developing modality for the developing patient. *AJR Am J Roentgenol* 2004; 182:243–252.
- Sonigo PC, Rypens FF, Carteret M, Delezoide A. MR imaging of fetal cerebral anomalies. *Pediatr Radiol* 1998; 28:212–222.
- Garel C, Chantrel E, Brisse H, et al. Fetal cerebral cortex: Normal gestational landmarks identified using prenatal MR imaging. *AJNR Am J Neuroradiol* 2001; 22:184–189.
- Girard N, Raybaud C, Gambarelli D, Figarella-Branger D. Fetal brain MR imaging. *MRI Clin North Am* 2001; 9:19–56.
- Glenn OA. Fetal central nervous system MR imaging. *Neuroimag Clin North Am* 2006; 16:1–17.
- Huppi PS, Inder TE. Magnetic resonance techniques in the evaluation of the perinatal brain: Recent advances and future directions. *Semin Neonatol* 2001; 6:195–210.
- Prayer D, Brugger PC, Prayer L. Fetal MRI: Techniques and protocols. *Pediatr Radiol* 2004; 34:685–693.
- Borman S. Topics in multiframe superresolution restoration. PhD thesis, University of Notre Dame, Indiana, 2004.
- Mitiche A, Boutheimy P. Computation and analysis of image motion: A synopsis of current problems and methods. *Int J Comput Vision* 1996; 19:29–55.
- Stiller C, Konrad J. Estimating motion in image sequences. *IEEE Signal Processing Magazine* 1999; 16:70–91.
- Peleg S, Keren D, Schweitzer L. Improving image resolution using sub-pixel motion. *Pattern Recognition Letters* 1987; 5:223–226.
- Irani M, Peleg S. Improving resolution by image registration. *Graphical Models Image Processing* 1991; 53:231–239.
- Patti AJ, Altunbasak Y. Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time. *IEEE Trans Image Processing* 1997; 6:1064–1076.
- Schultz RR, Stevenson RL. Extraction of high resolution frames from video sequences. *IEEE Trans Image Processing* 1996; 5:996–1011.
- Hardie RC, Barnad KJ, Armstrong EE. Joint map registration and high-resolution image estimation using a sequence of undersampled images. *IEEE Trans Image Processing* 1997; 6:1621–1633.
- Tekalp AM, Özkan MK, Sezan MI. High-resolution image reconstruction from lower-resolution image sequences and space-varying image restoration. *IEEE ICASSP* 1992; 169–172.
- Schierlitz L, Dumanli H, Robinson JN, et al. Three-dimensional magnetic resonance imaging of fetal brains. *Lancet* 2001; 357:1177–1178.
- Cachia A, Mangin JF, Rivière D, et al. A primal sketch of the cortex mean curvature: A morphogenesis based approach to study the variability of the folding patterns. *IEEE Trans Medical Imaging* 2003; 22:754–765.
- Peeters RR, Kornprobst P, Nikolova M, et al. The use of super-resolution techniques to reduce slice thickness in functional MRI. *Int J Imaging Syst Technol* 2004; 14:131–138.
- Greenspan H, Oz G, Kiryati N, Peled S. MRI inter-slice reconstruction using super-resolution. *Magn Reson Imaging* 2002; 20:437–446.
- Moore J, Drangova M, Wierzbicki M, et al. A high resolution dynamic heart model based on averaged MRI data. *Proc Med Image Comput Assist Intervent* 2003; 2878:549–555.
- Lotjonen J, Pollari M, Kivisto S, Lauerma K. Correction of movement artifacts from 4-D cardiac short- and long-axis MR data. *Proc Med Image Comput Assist Intervent* 2004; 3217:405–412.
- Maintz JB, Viergever MA. A survey of medical image registration. *Med Image Analysis* 1998; 2:1–36.
- Collignon A, Maes F, Delaere D, et al. Automated multimodality image registration using information theory. *Proc Inform Process Med Imaging* 1995; 287–298.
- Viola P, Wells W. Alignment by maximisation of mutual information. *Proc Int Conf Computer Vision* 1995; 15–23.
- Studholme C, Hill DLG, Hawkes DJ. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recogn* 1999; 32:71–86.
- Mitas L, Mitasova H. Spatial interpolation. *Geographical Information Systems: Principles and Applications* 1999; 481–492.
- Rohling R, Gee A, Berman L. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Med Image Analysis* 1999; 3:339–359.
- Gobbi D. Brain deformation correction using interactive 3D ultrasound imaging. PhD thesis, University of Western Ontario, 2003.
- Ohbuchi R, Chen D, Fuchs H. Incremental volume reconstruction and rendering for 3D ultrasound imaging. *SPIE Visual Biomed Comput* 1992; 1808:312–323.
- Studholme C, Cardenas V, Blumenfeld R, et al. A deformation tensor morphometry study of semantic dementia with quantitative validation. *NeuroImage* 2004; 21:1387–1398.
- Studholme C, Rodriguez-Carranza C, Cardenas V, et al. A deformation morphometry study of the influences on the pattern of brain tissue development in premature neonates. *Proceedings of Human Brain Mapping*, 2005.
- Rodriguez-Carranza C, Rousseau F, Iordanova B, et al. An iso-surface folding analysis method applied to premature neonatal brain development. *Proceedings of SPIE Medical Imaging, Image Processing* 2006; 6144, San Diego, CA.
- Inder TE, Huppi PS, Warfield S, et al. Periventricular white matter injury in the premature infant is followed by reduced cerebral cortical gray matter volume at term. *Ann Neurol* 1999; 46:755–760.
- Batchelor PG, Castellano Smith AD, Hill DLG, et al. Measures of folding applied to the development of the human fetal brain. *IEEE Trans Med Imaging* 2002; 21:953–965.