Hindawi Publishing Corporation
BioMed Research International
Volume 2017, Article ID 3956363, 8 pages
http://dx.doi.org/10.1155/2017/3956363

Research Article

Histograms of Oriented 3D Gradients for Fully Automated Fetal Brain Localization and Robust Motion Correction in 3 T Magnetic Resonance Images

Ahmed Serag,[1] Gillian Macnaught,[2] Fiona C. Denison,[1] Rebecca M. Reynolds,[1] Scott I. Semple,[2] and James P. Boardman[1]

[1]*MRC Centre for Reproductive Health, University of Edinburgh, Edinburgh, UK*
[2]*Clinical Research Imaging Centre, University of Edinburgh, Edinburgh, UK*

Correspondence should be addressed to Ahmed Serag; a.f.serag@gmail.com

Received 22 August 2016; Revised 13 December 2016; Accepted 26 December 2016; Published 30 January 2017

Academic Editor: Yudong Cai

Fetal brain magnetic resonance imaging (MRI) is a rapidly emerging diagnostic imaging tool. However, automated fetal brain localization is one of the biggest obstacles in expediting and fully automating large-scale fetal MRI processing. We propose a method for automatic localization of fetal brain in 3 T MRI when the images are acquired as a stack of 2D slices that are misaligned due to fetal motion. First, the Histogram of Oriented Gradients (HOG) feature descriptor is extended from 2D to 3D images. Then, a sliding window is used to assign a score to all possible windows in an image, depending on the likelihood of it containing a brain, and the window with the highest score is selected. In our evaluation experiments using a leave-one-out cross-validation strategy, we achieved 96% of complete brain localization using a database of 104 MRI scans at gestational ages between 34 and 38 weeks. We carried out comparisons against template matching and random forest based regression methods and the proposed method showed superior performance. We also showed the application of the proposed method in the optimization of fetal motion correction and how it is essential for the reconstruction process. The method is robust and does not rely on any prior knowledge of fetal brain development.

1. Introduction

Recent successful application of magnetic resonance imaging (MRI) has provided us with an unprecedented opportunity to study, in intricate detail, the developing brain in the living fetus or neonate [1–7]. However, quantitative fetal brain analysis remains challenging due to the unique confrontations associated with fetal MRI such as motion artifacts, low signal, and spatial resolution as shown in Figure 1.

Advances in medical image processing techniques have facilitated the reconstruction of motion-corrected high-resolution 3D fetal MR images [8–12] from stacks of 2D intersecting images, which in turn have laid the foundation for modeling [13–15] and quantitative analysis [1, 15, 16] of the developing fetal brain. Such reconstruction methods rely on initial localization and cropping of the brain region from a standard wide field of view (FOV) MRI, to assist the slice-to-volume registration process [17] by excluding surrounding maternal tissues that can result in registration failure.

Fetal brain localization as an automated processing step is one of the biggest obstacles in expediting and fully automating large-scale fetal MRI processing and analysis. To date, the approaches proposed to address the issue of automating fetal MRI brain localization can be classified into two categories: template matching and machine learning approaches. In the earlier category, Anquez et al. [18] proposed an approach that starts by detecting the eyes using 3D template matching, followed by segmentation of the brain using a 2D graph-cut segmentation of the mid-sagittal slice rotated at several angles. The best matching segmentation is selected and used to initialize a 3D graph-cut segmentation. Taimouri et al. [19] proposed another template matching approach

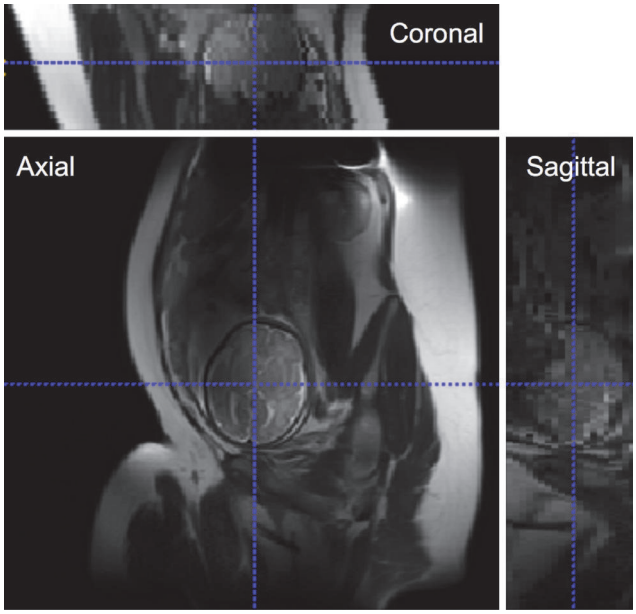

Figure 1: An example of a fetal MR scan at 35.5 weeks of gestational age (GA). The scan shows that the acquired image has an arbitrary fetal orientation, motion artifacts, and low spatial resolution.

that registers a template to each slice of the fetal brain in 3D.

Examples of machine learning approaches include the method proposed by Ison et al. [20], which is based on a two-phase random forest (RF) classifier to suppress the influence of maternal tissues and provide likely positions of tissue centroids inside the brain. An approximation of a high-order Markov Random Field (MRF) finds the optimal selection of landmarks, and these landmarks, along with a confidence weighted probability map, provide an estimate of the center of a region of interest around the brain. Keraudren et al. [21] described a different learning-based approach that decomposes the search space by performing a 2D detection process, based on Scale-Invariant-Feature-Transform (SIFT) features, before accumulating the votes in 3D. The approach relies on providing the gestational age as a prior knowledge to remove the scale component.

Recently, Kainz et al. [22] proposed a localization method where rotation invariant feature descriptors (based on Spherical Harmonics) for medical volumes of the uterus are calculated in order to learn the appearance of the fetal brain in the feature space. Alansary et al. [23] proposed a different localization method where superpixels are first extracted and a histogram of features for each superpixel is computed using bag of words based on dense SIFT descriptors. Then, a graph of superpixels is constructed and a RF classifier is trained to distinguish between brain and nonbrain superpixels.

Although most previously reported methods achieved accurate brain localization results, all methods were implemented and tested using 1.5 T MR images and the generalization of these approaches to 3 T MR images has not been examined. Fetal MRI performed at 3 T is emerging as a promising modality for the evaluation of fetal anatomy as it provides an increased signal-to-noise ratio (compared with 1.5 T). This increase in the signal-to-noise ratio can allow for decreased acquisition time, increased spatial resolution, or a combination of both, with the overall goal of obtaining more detailed imaging of fetal anatomy.

In this article, we focus on developing a novel method for fully automatic fetal brain localization from raw 3 T MR images, with no need for preprocessing (such as intensity inhomogeneity correction) or prior knowledge about the dataset under study (such as gestational age (GA) or imaging plane). We treat the problem as a machine learning problem where a model is trained using a sample of positive and negative examples. From each sample, discriminant features are extracted, for both positive (brain) and negative (nonbrain) examples, based on a 3D extension of the successful Histogram of Oriented Gradients (HOG) [24] feature descriptor. A sliding window classifier, based on Support Vector Regression (SVR), is used to assign a score to all possible windows in an image and the window with the highest score is selected.

The remainder of this article is organized as follows. Section 2 describes in detail the proposed method. In Section 3, the experimental results of applying the proposed method to 3 T MR fetal images are presented. We also compare the performance of our method against a template matching method and present its application in the optimization of fetal motion correction. Section 4 discusses our study and outlines areas for prospective work. This paper is an extension of previous preliminary work [25].

2. Methods

2.1. Data. The database for this study included a total of 104 MR images from 32 singleton fetuses of healthy women and women with diabetes, between 34.30 and 37.60 weeks of gestational age (mean and standard deviation of 35.92 ± 0.83). Ethical approval was obtained from the National Research Ethics Committee (South East Scotland Research Ethics Committee) and written informed consent was obtained.

2.2. MR Image Acquisition. Images were acquired on a Siemens Magneton Verio 3 T MRI clinical scanner (Siemens Healthcare GmbH, Erlangen, Germany). T2-weighted half-Fourier acquisition single-shot turbo spin-echo images were acquired of the fetal brain in sagittal, coronal, and transverse orientations where at least one stack is available in each anatomical direction (HASTE: TR/TE = 1800/86 ms, FOV = 400 × 400 mm, matrix = 192 (phase) × 256 (frequency), and voxel size = 1.5 × 1.5 × 6 mm).

2.3. The HOG Descriptor for 2D Images. In the 2D image domain, successful methods such as Harris-Corner detector [26] and the well-known SIFT descriptor [27] rely on aggregated gradients. However, the HOG descriptor [24] organizes gradients into histograms. As the first step, the gradient image is computed by convolving the input image with an appropriate filter mask. A grid of histograms is then constructed, where each histogram organizes the respective gradients into bins according to their orientation. To preserve locality, a

histogram is computed for each cell in an evenly spaced grid. Accordingly, each cell contains the same number of gradients (depending on the cell size) and gets assigned exactly one histogram. The cells themselves are then organized in rectangular blocks, which may overlap. The histogram values of all cells within one block are concatenated to form a vector. The vector of each block is then normalized and subsequently, the concatenation of all those block-vectors yields the final feature vector. Further details and an evaluation of the 2D HOG descriptor can be found in Dalal and Triggs [24].

2.4. Extending HOG to 3D Images. Due to the nature of the analyzed data, that is, 3D volumetric images, 3D feature descriptors are ideal. It also has been demonstrated that 3D descriptors are much more descriptive than their corresponding 2D, as richer information is encoded into the histograms [28].

Here, our extension of the HOG approach to 3D medical images consists of two steps. First, we need to extract gradients from the images. Second, we need to organize these three-dimensional gradients into bins using appropriate histograms computed over uniformly spaced grid-blocks. This step is straightforward, as we simply extend the grid and histogram dimension, each by one. We then can convert each gradient into spherical coordinates (1) and bin it according to its orientation (azimuth $\theta \in [0, 2\pi)$ and zenith $\phi \in [0, \pi]$).

$$\begin{pmatrix} r \\ \theta \\ \phi \end{pmatrix} = \begin{pmatrix} \sqrt{x^2 + y^2 + z^2} \\ \tan^{-1}\left(\frac{y}{x}\right) \\ \cos^{-1}\left(\frac{z}{r}\right) \end{pmatrix}. \quad (1)$$

The first step, however, does not generalize as easily. To compute the image-gradient, several approaches might be considered. According to Dalal and Triggs [24], the convolution of the image with a 1D $[-1, 0, 1]$ filter mask is most suitable. This approximates the partial first-order derivative according to

$$\nabla f(x, y) = \left(\frac{\delta f}{\delta x}, \frac{\delta f}{\delta y}\right)^T \quad (2)$$

$$\approx \begin{pmatrix} f(x+1, y) - f(x-1, y) \\ f(x, y+1) - f(x, y-1) \end{pmatrix}, \quad (3)$$

$$\nabla f(x, y, z) = \left(\frac{\delta f}{\delta x}, \frac{\delta f}{\delta y}, \frac{\delta f}{\delta z}\right)^T \quad (4)$$

$$\approx \begin{pmatrix} f(x+1, y, z) - f(x-1, y, z) \\ f(x, y+1, z) - f(x, y-1, z) \\ f(x, y, z+1) - f(x, y, z-1) \end{pmatrix}. \quad (5)$$

Based on the description above, Algorithm 1 summarizes the extraction of the 3DHOG feature vector from 3D fetal MR images.

It is worth mentioning that Histograms of Oriented 3D Gradients have been previously reported in the literature, for example, in the areas of video sequence analysis [29, 30] or 3D object retrieval [31]. Those 3D extended HOGs are different from ours as their data representation was 2D + time [29, 30] or 3D mesh models [31]. Therefore, the previously proposed extensions of HOG cannot be directly (or simply) applied to our data as we deal with 3D volumetric medical images.

```
Set the block size s for image I;
Set the number of histogram bins B;
Set f_3dhog to represent the 3DHOG feature vector;
Compute the image gradient G;
foreach block do
    foreach gradient g in G do
        Transform g to spherical coordinates;
        Insert g into corresponding histogram bins;
    end
    Normalize block-vector b;
    Append normalized b to f_3dhog;
end
```

Algorithm 1: Histograms of Oriented 3D Gradients (3DHOG).

2.5. Brain Localization Using Sliding Window. A sliding window is used to move over all possible windows in an image, and the window with the highest score is selected. Given that one of our aims is to provide a method that does not require any prior knowledge, we use a sliding window of a fixed size that is large enough to hold the largest brain in our database. To assign a score to each window, we use Support Vector Regression (SVR). For our experiments, the detector has the following default settings: $[-1, 0, 1]$ gradient filter with no smoothing; 9 orientation bins in 0–180 degrees for azimuth and zenith; $3 \times 3 \times 3$ voxel blocks with an overlap of half the block size; $L2$-norm block normalization; $40 \times 40 \times 5$ step size for sliding window; SVR ($C = 1, \epsilon = 0.1$) with a first-order polynomial kernel. In the Results, we perform an evaluation to measure the performance with respect to the variation of these parameters and to find the best values for our learning task.

2.6. Comparisons against Other Methods. We implemented a basic method of template matching, with the template representing an average fetal brain, to compare our proposed method against. We will call the template $T(x_t, y_t, z_t)$, where (x_t, y_t, z_t) represent the coordinates of each voxel in the template. We will call the input or search image $S(x, y, z)$, where (x, y, z) represent the coordinates of each voxel in the search image. We then simply move the center (or the origin) of the template $T(x_t, y_t, z_t)$ over each (x, y, z) point in the search image and calculate the sum of products between the coefficients in $S(x, y, z)$ and $T(x_t, y_t, z_t)$ over the whole area spanned by the template. As all possible positions of the template with respect to the search image are considered, the position with the best score is the best position. Here, the best score is defined as the lowest Sum of Absolute Differences (SAD):

$$\mathrm{SAD}(x, y, z) = \sum_{i=0}^{N}\sum_{j=0}^{M}\sum_{k=0}^{L}\mathrm{Diff}(x+i, y+j, z+k, i, j, k), \quad (6)$$

where N, M, and L denote the sizes of each dimension of the template image. We also experimented replacing the support vector regressor with a random forest regressor [32] using the following parameter settings: number of trees in the forest = 10, minimum number of samples required to be at a leaf node = 1, and unlimited number of leaf nodes.

2.7. Validation. Leave-one-out and 10-fold cross-validation procedure were performed for the 104 MR images. In case of leave-one-out cross-validation, each image in turn was left out as a test sample and the remaining 103 images were used as the training dataset. In 10-fold cross-validation, the original dataset was randomly partitioned into 10 equal sized subsamples. Of the 10 subsamples, a single subsample was retained as the test data, and the remaining 9 subsamples were used as training data. The cross-validation process was then repeated 10 times (the folds), with each of the 10 subsamples used exactly once as the validation data. A correct detection or complete brain localization was defined as the enclosure of all brain voxels within the detected box.

We also measured the distance d between the center of the ground truth bounding box and the center of the detected bounding box, and median absolute error, MdAE, was obtained using the following formula:

$$\mathrm{MdAE} = \underset{i=1,\ldots,n}{\mathrm{median}} |d_i|, \quad (7)$$

where n is the total number of images used for validation.

3. Results

Using leave-one-out cross-validation, the proposed method provided 96% success rate in brain localization and MdAE of 6.95 mm. When the 10-fold cross-validation was performed, the method provided comparable results with success rate of 95% and MdAE of 7.73 mm.

3.1. Influence of Parameters. Figure 2 summarizes the influence of the various HOG parameters on overall detection performance. Detector performance is sensitive to the way in which gradients are computed, but the simplest scheme provided better accuracy. As shown in Figure 2, increasing the number of orientation bins did not greatly improve performance up to 9 bins, but beyond this, performance started to degrade. Gradient strengths vary over a wide range due to intensity inhomogeneity; therefore effective local contrast normalization between overlapping blocks is essential for good performance. From our experiments, it turns out that $L1$-norm is better than $L2$-norm. When tested for different block sizes, the block size of $3 \times 3 \times 3$ voxels yielded the best performance. Testing for different sliding window step sizes, the step size of $40 \times 40 \times 5$ had the best performance. Experiments also showed that SVR with a second- and third-order polynomial kernels outperformed linear SVR.

3.2. Comparisons against Other Methods. To date, no method has been proposed for automatic localization of the fetal brain using 3 T MR images, making it difficult to compare our proposed method against a benchmark method. However, we compared our results against a basic method of template matching, with the template representing an average fetal brain. The median localization error for the implemented template-based matching approach was MdAE = 12.5 mm. This shows a degradation of brain localization accuracy compared to proposed method using Histograms of Oriented 3D Gradients. When we experimented replacing the support vector regressor with a random forest regressor, we obtained results with MdAE of 7.9 mm.

3.3. Application to Robust Motion Correction. Due to the nature of the acquisition, that is, acquiring stacks of 2D slices in real-time MRI, in order to reduce the scan time while avoiding slice cross-talk artifacts, slices are quite often misaligned due to fetal motion and form an inconsistent 3D volume (see Figure 1). Registration-based approaches for reconstructing motion-corrected high-resolution fetal brain volumes require a cropped box around the brain to exclude irrelevant tissues, otherwise the registration, and the reconstruction, will likely fail. Figure 3 shows an example of three different subjects that were successfully reconstructed using the Baby ToolKit [17]. It is clear that brain cropping is essential for the reconstruction process, and the registration fails when the whole FOV is used.

4. Discussion

In this article, Histogram of Oriented Gradients (HOG) is proposed to automatically localize the fetal brain in 3 T MR images. The main contribution of the article is the extension of HOG to 3D for fetal brain MRI localization. We chose to use HOG features partly because of its demonstrated superiority to other widely used features like SIFT/SURF in many applications [24, 33, 34] and partly because the use of HOG features in our task is rational as head boundaries are HOG rich relative to surrounding maternal tissues (see Figure 4).

We also chose to use Support Vector Regression (SVR) on HOG features, instead of Support Vector Machine (SVM), which was proposed by Dalal and Triggs [24] in the original HOG feature descriptor. Typically, SVR is used with a sliding window approach to assign a score to all possible windows in an image and the window with the highest score is then selected. Therefore, by using SVR over SVM, we avoid an inevitable problem that arises in overlapping sliding window approach, which is the occurrence of multiple detection windows in the same neighborhood area [35].

We examined the influence of different HOG parameters on overall detection performance in terms of median absolute error of a bounding box detected around the fetal brain and the ground truth (defined manually). In our results, we achieved 96% for complete brain localization, and the

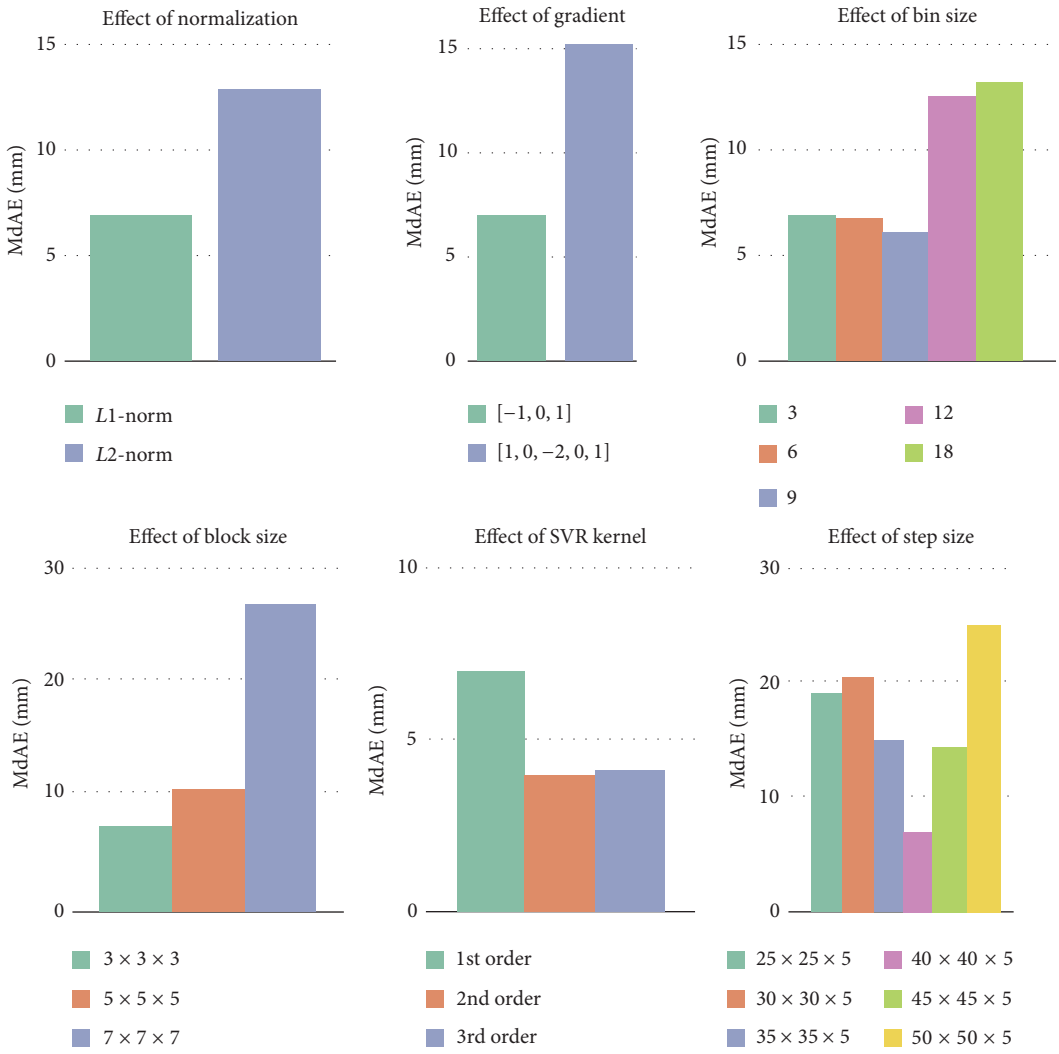

Figure 2: The effect of the various HOG parameters on overall localization performance.

automated detection process takes less than 5 seconds on a normal computer. Our method is more general than Anquez et al. [18] as it does not rely on localizing the eyes. It also does not require any prior knowledge (such as gestational age) as in Keraudren et al. [21].

We carried out comparisons against template matching and random forest based regression methods and the proposed method showed superior performance. We also showed the application of the proposed method in the optimization of fetal motion correction and how it is essential for the motion correction process.

Noteworthy is the fact that the original HOG feature descriptor is not rotation invariant, given that we relied on the large variation of the training dataset to handle different orientations. Using a database of 104 MRI scans aged between 34 and 38 weeks of GA, the rotation variance is represented within the training samples and hence the model will learn to classify it. In addition, the training samples were acquired in different planes (axial, coronal, and sagittal), which is another strength of the proposed method.

5. Conclusion

The main motivation behind this work was to automate the fetal brain localization step which is one of the obstacles preventing rapid deployment and full automation of large-scale fetal MRI postprocessing. Due to the nature of the acquisition, that is, acquiring stacks of 2D slices in real-time MRI, in order to reduce the scan time while avoiding slice cross-talk artifacts, slices are quite often misaligned due to fetal motion and form an inconsistent 3D volume. Registration-based approaches for reconstructing motion-corrected high-resolution fetal brain volumes require a cropped box around the brain to exclude irrelevant tissues, otherwise the registration, and the reconstruction, will likely fail. Future work will include applying and evaluating the brain localization framework to a larger GA range (i.e., second trimester), and different MRI modalities of the fetal brain, such as T1-weighted, Echo Planar Imaging (EPI), and Diffusion Tensor Imaging (DTI). This will allow for more accurate assessment and analysis of the developing fetal

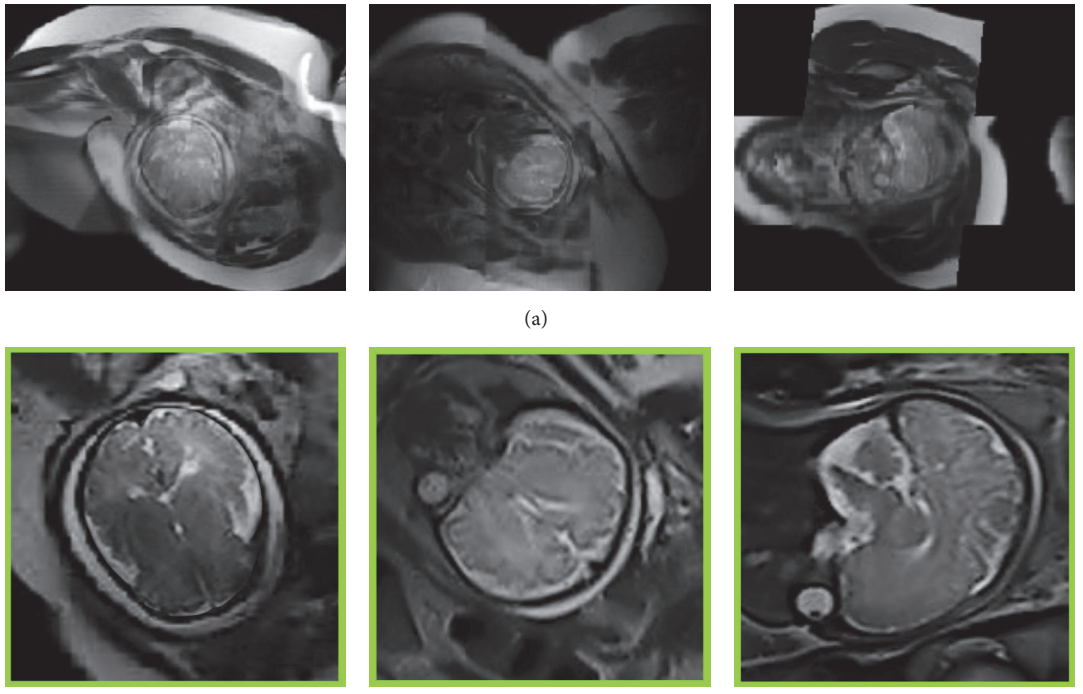

(a)

(b)

Figure 3: Three MRI scans reconstructed using raw MRI without brain localization and cropping (a). Inside the green colored frames (b), the same images were motion-corrected after automatic brain localization and cropping using the proposed method.

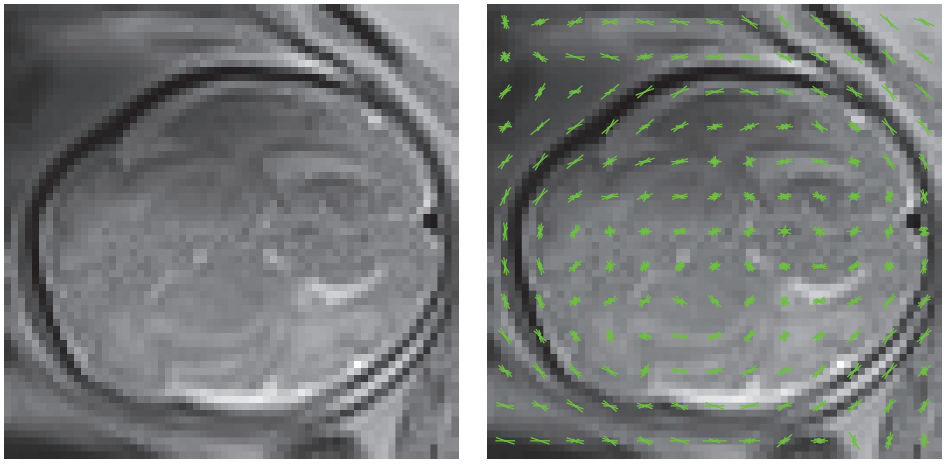

(a) (b)

Figure 4: An example of an axial fetal MR image used from the training database (a) and the resulting computed HOG descriptor overlaid on the example image (b).

brain across several stages of development using different imaging modalities. The method is available to the research community at http://brainsquare.org.

Competing Interests

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interests.

Authors' Contributions

Ahmed Serag designed and performed the experiments, analyzed output data, and wrote the manuscript; James P. Boardman, Fiona C. Denison, and Rebecca M. Reynolds prepared the grant application. Gillian Macnaught and Scott I. Semple acquired imaging data. All authors approved the final submitted version and agree to be accountable for its content.

Acknowledgments

This work was supported by the Theirworld (http://www.theirworld.org), NHS Research Scotland, and NHS Lothian Research and Development. This work was undertaken in the MRC Centre for Reproductive Health which is funded by the MRC Centre Grant MR/N022556/1. The authors are grateful to the families who consented to take part in the study and to the midwives for recruiting.

References

[1] J. A. Scott, P. A. Habas, K. Kim et al., "Growth trajectories of the human fetal brain tissues estimated from 3D reconstructed in utero MRI," *International Journal of Developmental Neuroscience*, vol. 29, no. 5, pp. 529–536, 2011.

[2] P. S. Hüppi, S. Warfield, R. Kikinis et al., "Quantitative magnetic resonance imaging of brain development in premature and mature newborns," *Annals of Neurology*, vol. 43, no. 2, pp. 224–235, 1998.

[3] A. Serag, I. S. Gousias, A. Makropoulos et al., "Unsupervised learning of shape complexity: application to brain development," in *Spatio-Temporal Image Analysis for Longitudinal and Time-Series Image Data*, Lecture Notes in Computer Science, Springer, 2012.

[4] P. Batchelor, A. Castellano Smith, D. Hill, D. Hawkes, T. Cox, and A. Dean, "Measures of folding applied to the development of the human fetal brain," *IEEE Transactions on Medical Imaging*, vol. 21, no. 8, pp. 953–965, 2002.

[5] L. R. Ment, D. Hirtz, and P. S. Hüppi, "Imaging biomarkers of outcome in the developing preterm brain," *The Lancet Neurology*, vol. 8, no. 11, pp. 1042–1055, 2009.

[6] A. Serag, M. Blesa, E. J. Moore et al., "Accurate Learning with Few Atlases (ALFA): an algorithm for MRI neonatal brain extraction and comparison with 11 publicly available methods," *Scientific Reports*, vol. 6, Article ID 23470, 2016.

[7] N. N. Andescavage, A. du Plessis, R. McCarter et al., "Complex trajectories of brain development in the healthy human fetus," *Cerebral Cortex*, 2016.

[8] F. Rousseau, O. A. Glenn, B. Iordanova et al., "Registration-based approach for reconstruction of high resolution in utero MR brain images," *Academic Radiology*, vol. 13, no. 9, pp. 1072–1081, 2006.

[9] A. Gholipour, J. A. Estroff, and S. K. Warfield, "Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI," *IEEE Transactions on Medical Imaging*, vol. 29, no. 10, pp. 1739–1758, 2010.

[10] K. Kim, P. A. Habas, F. Rousseau, O. A. Glenn, A. J. Barkovich, and C. Studholme, "Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation," *IEEE Transactions on Medical Imaging*, vol. 29, no. 1, pp. 146–158, 2010.

[11] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel, "Reconstruction of fetal brain MRI with intensity matching and complete outlier removal," *Medical Image Analysis*, vol. 16, no. 8, pp. 1550–1564, 2012.

[12] B. Kainz, M. Steinberger, W. Wein et al., "Fast volume reconstruction from motion corrupted stacks of 2D slices," *IEEE Transactions on Medical Imaging*, vol. 34, no. 9, pp. 1901–1913, 2015.

[13] E. Dittrich, T. Riklin Raviv, G. Kasprian et al., "A spatio-temporal latent atlas for semi-supervised learning of fetal brain segmentations and morphological age estimation," *Medical Image Analysis*, vol. 18, no. 1, pp. 9–21, 2014.

[14] R. Wright, D. Vatansever, V. Kyriakopoulou et al., "Age dependent fetal MR segmentation using manual and automated approaches," *MICCAI Workshop on Perinatal and Paediatric Imaging*, pp. 97–104, 2012.

[15] A. Serag, V. Kyriakopoulou, M. A. Rutherford et al., "A multi-channel 4D probabilistic atlas of the developing brain: application to fetuses and neonates," *Annals of the BMVA*, vol. 2012, no. 3, pp. 1–14, 2012.

[16] A. Gholipour, A. Akhondi-Asl, J. A. Estroff, and S. K. Warfield, "Multi-atlas multi-shape segmentation of fetal brain MRI for volumetric and morphometric analysis of ventriculomegaly," *NeuroImage*, vol. 60, no. 3, pp. 1819–1831, 2012.

[17] F. Rousseau, E. Oubel, J. Pontabry et al., "BTK: an open-source toolkit for fetal brain MR image processing," *Computer Methods and Programs in Biomedicine*, vol. 109, no. 1, pp. 65–73, 2013.

[18] J. Anquez, E. D. Angelini, and I. Bloch, "Automatic segmentation of head structures on fetal MRI," in *Proceedings of the IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI '09)*, pp. 109–112, IEEE, Boston, Mass, USA, July 2009.

[19] V. Taimouri, A. Gholipour, C. Velasco-Annis, J. A. Estroff, and S. K. Warfield, "A template-to-slice block matching approach for automatic localization of brain in fetal MRI," in *Proceedings of the 12th IEEE International Symposium on Biomedical Imaging (ISBI '15)*, Brooklyn, NY, USA, April 2015.

[20] M. Ison, R. Donner, E. Weigl, and G. Langs, "Fully automated brain extraction and orientation in raw fetal MRI," in *Proceedings of the MICCAI Workshop—PAPI*, pp. 17–24, October 2012.

[21] K. Keraudren, V. Kyriakopoulou, M. Rutherford, J. V. Hajnal, and D. Rueckert, "Localisation of the brain in fetal MRI using bundled SIFT features," *Medical Image Computing and Computer-Assisted Intervention*, vol. 16, part 1, pp. 582–589, 2013.

[22] B. Kainz, K. Keraudren, V. Kyriakopoulou, M. Rutherford, J. V. Hajnal, and D. Rueckert, "Fast fully automatic brain detection in fetal MRI using dense rotation invariant image descriptors," in *Proceedings of the IEEE 11th International Symposium on Biomedical Imaging (ISBI '14)*, pp. 1230–1233, Beijing, China, May 2014.

[23] A. Alansary, M. Lee, K. Keraudren et al., "Automatic brain localization in fetal MRI using superpixel graphs," in *Proceedings of the 1st International Workshop on Machine Learning Meets Medical Imaging ICML Workshop (MLMMI '15)*, vol. 9487, Lille, France, July 2015.

[24] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, vol. 1, June 2005.

[25] A. Serag, J. P. Boardman, R. M. Reynolds, F. C. Denison, G. Macnaught, and S. I. Semple, "Automatic fetal brain localization in 3T MR images using Histograms of Oriented 3D Gradients," in *Proceedings of the 24th Signal Processing and Communication Application Conference (SIU '16)*, pp. 2273–2276, Zonguldak, Turkey, May 2016.

[26] C. Harris and M. Stephens, "A combined corner and edge detector," in *Proceedings of the Alvey Vision Conference*, C. J. Taylor, Ed., vol. 15, pp. 23.1–23.6, September 1988.

[27] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[28] P. Scovanner, S. Ali, and M. Shah, "A 3-dimensional sift descriptor and its application to action recognition," in *Proceedings of the 15th ACM International Conference on Multimedia (MM '07)*, pp. 357–360, Augsburg, Germany, September 2007.

[29] A. Kläser, M. Marszałek, and C. Schmid, "A spatio-temporal descriptor based on 3D-gradients," in *Proceedings of the 19th British Machine Vision Conference (BMVC '08)*, BMVA Press, September 2008.

[30] N. Buch, J. Orwell, and S. A. Velastin, "3D extended histogram of oriented gradients (3DHOG) for classification of road users in urban scenes," in *Proceedings of the 20th British Machine Vision Conference (BMVC '09)*, September 2009.

[31] M. Scherer, M. Walter, and T. Schreck, "Histograms of oriented gradients for 3D object retrieval," in *Proceedings of the International Conferences in Central Europe on Computer Graphics, Visualization and Computer Vision*, 2010.

[32] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[33] C. Yi, X. Yang, and Y. Tian, "Feature representations for scene text character recognition: a comparative study," in *Proceedings of the 12th International Conference on Document Analysis and Recognition (ICDAR '13)*, pp. 907–911, IEEE, Washington, DC, USA, August 2013.

[34] A. Kembhavi, D. Harwood, and L. S. Davis, "Vehicle detection using partial least squares," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 6, pp. 1250–1265, 2011.

[35] C. Premebida and U. Nunes, "Fusing LIDAR, camera and semantic information: a context-based approach for pedestrian detection," *International Journal of Robotics Research*, vol. 32, no. 3, pp. 371–384, 2013.