Fast Fully Automatic Segmentation of the Human Placenta from Motion Corrupted MRI

Amir Alansary[1,*], Konstantinos Kamnitsas[1], Alice Davidson[2], Rostislav Khlebnikov[2], Martin Rajchl[1], Christina Malamateniou[2], Mary Rutherford[2], Joseph V. Hajnal[2], Ben Glocker[1], Daniel Rueckert[1], and Bernhard Kainz[1]

[1] Department of Computing, Imperial College London, UK
[2] King's College London, Division of Imaging Sciences, London, UK

Abstract. Recently, magnetic resonance imaging has revealed to be important for the evaluation of placenta's health during pregnancy. Quantitative assessment of the placenta requires a segmentation, which proves to be challenging because of the high variability of its position, orientation, shape and appearance. Moreover, image acquisition is corrupted by motion artifacts from both fetal and maternal movements. In this paper we propose a fully automatic segmentation framework of the placenta from structural T2-weighted scans of the whole uterus, as well as an extension in order to provide an intuitive pre-natal view into this vital organ. We adopt a 3D multi-scale convolutional neural network to automatically identify placental candidate pixels. The resulting classification is subsequently refined by a 3D dense conditional random field, so that a high resolution placental volume can be reconstructed from multiple overlapping stacks of slices. Our segmentation framework has been tested on 66 subjects at gestational ages 20–38 weeks achieving a Dice score of $71.95 \pm 19.79\%$ for healthy fetuses with a fixed scan sequence and $66.89 \pm 15.35\%$ for a cohort mixed with cases of intrauterine fetal growth restriction using varying scan parameters.

1 Introduction

The functions of the placenta affect the fetal birth weight, growth, prematurity, and neuro-development since it controls the transmission of nutrients from the maternal to the fetal circulatory system. Recent work [8] has shown that magnetic resonance imaging (MRI) can be used for the evaluation of the placenta during both normal and high-risk pregnancies. Particularly, quantitative measurements such as placental volume and surface attachment to the uterine wall, are required for identifying abnormalities. In addition, recording the structural appearance (*e.g.*, placental cotyledons and shape) is essential for clinical qualitative analysis. Moreover, the placenta is usually examined after birth, on a flat surface providing a standard representation for obstetricians. Flat cutting planes, as common in radiology, show only a small part of the placenta. A 3D visualization is considered useful in particular for cases that require preoperative planning or surgical navigation (e.g. treatment of twin-to-twin transfusion syndrome). Hence, fully automatic 3D segmentation, correction of motion artifacts,

and visualization is highly desirable for an efficient pre-natal examination of the placenta in the clinical practice.

Fast MRI acquisition techniques (single shot fast spin echo – ssFSE) allow acquiring single 2D images of the moving uterus and fetus fast enough so that motion does not affect the image quality. However, 3D data acquisition and subsequent automatic segmentation is challenging because maternal respiratory motion and fetal movements displace the overall anatomy, which causes motion artifacts between individual slices as shown in Fig. 1. Furthermore, a high variability of the placenta's position, orientation, thickness, shape and appearance inhibits conventional image analysis approaches to be successful.


(a) Axial   (b) Sagittal   (c) Coronal

Fig. 1. Three orthogonal 2D planes from a motion corrupted 3D stack of slices showing a delineated placenta. The native scan orientation (a) shows no motion artifacts, while (b) and (c) do.

Related work: To the best of our knowledge, fully automatic segmentation of the placenta from MRI has not been investigated before. Most previous work in fetal MRI was focused on brain segmentation [2] and very recently has been extended to localize other fetal organs [6]. These methods rely on engineering visual features for training a classifier such as random forests. Stevenson *et al.* [9] present a semi-automatic approach for measuring the placental volume from motion free 3D ultrasound with a random walker (RW) algorithm. Their method shows a good inter-observer reproducibility but requires extensive user interaction and several minutes per segmentation. Even though ultrasound is fast enough to acquire a motion free volume, the lack of structural information and weak tissue gradients make it only useful for volume measurements. Wang *et al.* [12] present an interactive method for the segmentation of the placenta from MR images, which requires user interaction to initialize the localization of the placenta. Their approach performs well on a small cohort of six subjects but shows a user-dependent variability in segmentation accuracy.

Contribution: In this paper we propose for the first time a fully automatic segmentation framework for the placenta from motion corrupted fetal MRI. The proposed framework adopts convolutional neural networks (CNNs) as a strong classifier for image segmentation followed by a conditional random field (CRF) for refinement. Our approach scales well to real clinical applications. We propose how to use the resulting placental mask as initialization for slice-to-volume registration (SVR) techniques to compensate for motion artifacts. We also show how

the resulting reconstructed volume can be used to provide a novel standardized view into the placental structures by applying shape skeleton extraction and curved planar reformation for shape abstraction.

2   Method

The proposed approach combines a 3D multi-scale CNN architecture for segmentation with a 3D dense CRF for segmentation refinement. This approach can be extended to compensate for motion and to provide a clinically useful visualization. Figure 2 shows an overview of the proposed framework.
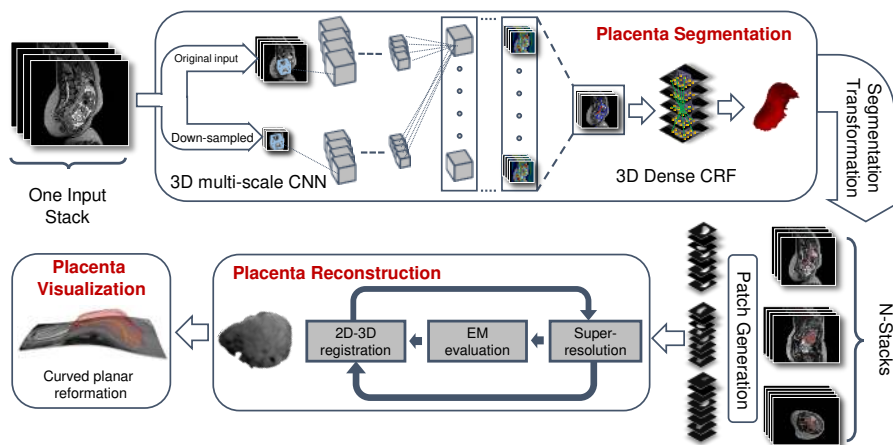

Fig. 2. The proposed framework for automatic placenta segmentation with extensions for motion correction and visualization.

Placenta segmentation: We adopt a 3D deep multi-scale CNN architecture [4] that is 11-layers deep and consists of two pathways to segment the placenta from the whole uterus. This multi-scale architecture has the advantage of capturing larger 3D contextual information, which is essential for detecting highly variable organs. Both pathways are complementary as the main pathway extracts local features, whereas the second one extracts larger contextual features. Multi-scale features are integrated efficiently by down-sampling the input image and processing the two pathways in parallel. In order to deal with the variations of the placenta's appearance, we apply data augmentation for training by flipping the image around the main 3D axes (maternal orientation).

Despite the fact that the multi-scale architecture can interpret contextual information, inference is subject to misclassification and errors. Hence, we apply a CRF to penalize inconsistencies of the segmentation by regularizing classification priors with the relational consistency of their neighbors. We use a 3D fully connected CRF model [7, 4] which applies a linear combination of Gaussian kernels to define the pairwise edge potentials. It is defined as $E(\mathbf{x}) =$

$\sum_{i \in N} U(x_i) + \sum_{i<j} V(x_i, x_j)$, where i and j are pixel indexes. The unary potential U is given by the probabilistic predictions of the CNN classification. Whereas the pairwise potential V is defined by

$$V(x_i, x_j) = \mu(x_i, x_j) \sum_{m=1}^{K} \left(\omega_1 e^{\left(-\frac{|p_i - p_j|^2}{2\theta_\alpha^2} - \frac{|I_i - I_j|^2}{2\theta_\beta^2} \right)} + \omega_2 e^{\left(\frac{|p_i - p_j|^2}{2\theta_\gamma^2} \right)} \right),$$

where I and p are intensity and position values. $\mu(x_i, x_j)$ is a simple label compatibility function given by the Potts model $[x_i \neq x_j]$. Here, ω_1 controls the importance of the appearance of nearby pixels to have similar labels. ω_2 controls the size of the smoothness kernel for removing isolated regions. θ_α, θ_β and θ_γ are used to adjust the degree of similarity and proximity. We have chosen the configuration parameters heuristically similar to [4]. Although this tissue classification approach is capable of segmenting the placenta robustly, the segmentation is still subject to inter-slice motion artifacts.

Placenta segmentation recovery: To tackle these motion artifacts caused by fetal and maternal movements we combine our segmentation framework with flexible motion compensation algorithm based on patch-to-volume registration (PVR) [3]. This technique requires multiple orthogonal stacks of 2D slices to provide a better reconstruction quality. It is based on splitting the input data into overlapping square patches or superpixels [1]. The motion-free 3D image is then reconstructed from the extracted patches using iterative super-resolution and 2D/3D registration steps. The motion-corrupted and misaligned patches are excluded during the reconstruction using an EM-based outliers rejection model. We extend this process to allow propagation of the placental mask to the final reconstruction through evaluating an MR specific point spread function, registration-based transformation, and the learned confidence weights.

Placenta visualization: We present an extension of our placenta segmentation pipeline based on a novel application of shape abstraction using a flexible cutting plane. It is supported by a mean-curvature flow skeleton [10] generated from the triangulated polygonal mesh of the placenta segmentation and textured similar to curved planar reformation [5], see Figure 3. Although this part is not evaluated thoroughly, clinicians revealed that such a representation is potentially desirable since it compares well to a flattened placenta after birth.

3   Experimental Results

Data: We test our approach on two dissimilar datasets that are different in health status, gestational ages and acquired using different scanning parameters. All scans have been ethically approved. *Dataset I* contains 44 MR scans of healthy fetuses at gestational age between 20–25 weeks. The data has been acquired on a Philips Achieva 1.5T, the mother lying 20° tilt on the left side to avoid pressure on the inferior vena cava. ssFSE T2-weighted sequences are used to acquire stacks of images that are aligned to the main axes of the fetus. Usually three to six stacks are acquired for the whole womb and the placenta

Fig. 3. A native plane (a) cannot represent all structures of the placenta at once. Therefore, we use our segmentation method (b), correct the motion in this area using [3], project the placenta mask into the the resulting isotropically resolved volume (c), extract the mean curvature flow skeleton [10] (black lines in (d)), use the resulting points to support a curved surface plane (e) and visualize this plane with curved planar reformation [5] (f). The plane in (f) covers only relevant areas, hence gray value mapping can be adjusted automatically to emphasis placental structures.

with a voxel size of $1.25 \times 1.25 \times 2.50$mm. *Dataset II* contains 22 MR scans of healthy fetuses and fetuses with intrauterine fetal growth restriction (IUGR) at gestational age between 20–38 weeks. The data was acquired with a 1.5T Philips MRI system using ssFSE sequences and a voxel size of $0.8398 \times 0.8398 \times 4$mm. Ground truth labels for both datasets have been obtained manually slice-by-slice in 2D views from the original motion-corrupted stacks by a clinical expert.

Experiments: The proposed segmentation framework is evaluated using three main metrics: Dice similarity coefficient to measure the accuracy of the segmentation, absolute volume similarity to measure the volumetric error between the segmented and the ground truth volumes, and average Hausdorff distance as a distance error metric between the segmented and the ground truth surfaces.

We evaluate in a first experiment [exp-1] the automatic segmentation of the placenta on *Dataset I* using a 4-fold cross validation (11 test patients and 33 training patients per fold). The main aim of this experiment is to evaluate the performance of our segmentation framework on a healthy homogeneous dataset. The results for this experiment are $71.95 \pm 19.79\%$ Dice, $30.92 \pm 33.68\%$ absolute volume difference, and $4.94 \pm 6.93mm$ average Hausdorff distance.

In a second experiment [exp-2], we train the CNN using the whole 44 subject from *Dataset I* and test it on the 22 subjects from *Dataset II*. Where datasets *I* and *II* are significantly different using different scanners and scanning parameters. In addition, the gestational age range of the fetuses in *Dataset II* is wider, which has a big influence on the fetal body and placenta sizes. Hence we test the performance of our framework when it is used to test data from a different environment. The results of this experiment are $56.78 \pm 21.86\%$ Dice, $48.19 \pm 46.96\%$ absolute volume difference, and $8.41 \pm 7.1mm$ average Hausdorff distance.

To resemble a realistic transfer learning application we have designed a third experiment [exp-3] using both datasets. The network is evaluated with 2-fold cross validation, 10 test subjects from *Dataset II*, and 44+10 training subjects from *Dataset I* and *Dataset II*. This experiment yielded a Dice accuracy of $66.89 \pm 15.35\%$, an absolute volume difference of $33.05 \pm 30.71\%$, and an average Hausdorff distance of $5.8 \pm 4.24mm$. Detailed results are shown in Fig. 4. Training one fold takes approximately 40 hours and inference can be done within 2 minutes on an Nvidia Tesla K40.

(a) Dice   (b) AHD

Fig. 4. Evaluation of the proposed method using (a) average Dice coefficient and (b) average Hausdorff distance. [exp-1] refers to 4-fold evaluation on *Dataset I*, [exp-2] uses *Dataset I* for training and *Dataset II* for testing, and [exp-3] mixes both datasets for training and uses unseen examples from *Dataset II* for testing.

Evaluation of clinical parameters: We compare our work to known values from the clinical practice. [8] shows that the average placental volume increases from 252.4 cm^3 at 20 weeks to 1421.5 cm^3 at 37 weeks. Fig. 5 compares these values from the clinical literature [11] to our automatically measured volumes from our datasets. It shows that our approach achieves very similar volumetric results compared to both expert estimations and clinical literature. In addition, the slope parameters of our segmentation and ground truth are not significantly different with p-value 0.94. Pathological cases from *Dataset II* show differences to scans of healthy placentas in Fig. 5(c).

Motion compensation: Evaluating the quality of a motion compensated reconstruction is challenging due to the absence of the motion-free ground truth data. Assuming that the 2D in-plane patches from the original 3D stacks have no motion artifacts, reconstructed patches are evaluated using these motion-free 2D patches as ground truth. The average peak signal-to-noise ratio (PSNR) is calculated for all the patches of each subject. The baseline represents the quality of patches from non-native slice orientations, which is comparable to using a single motion corrupted stack for diagnostics with arbitrary cutting planes. The baseline has low PSNR values due to the motion artifacts between the input stacks. These values increase during the reconstruction iterations as a result of reducing the motion artifacts of the segmented placenta, see Fig. 6.

4 Discussion and Conclusion

We present a fully automatic segmentation framework for the human placenta from motion corrupted fetal MRI scans. We perform rigorous experiments on two different testing datasets in order to evaluate thoroughly the presented segmentation approach, which is based on a 3D deep multi-scale convolutional neural network combined with conditional random field segmentation refinement. Our experiments show that this framework can tackle motion artifacts by achieving segmentation accuracy of 71.95% for healthy fetuses. It is also capable of segmenting the placenta from dissimilar data by achieving segmentation accuracy of 66.89% for a cohort mixed with cases of intrauterine fetal growth restriction
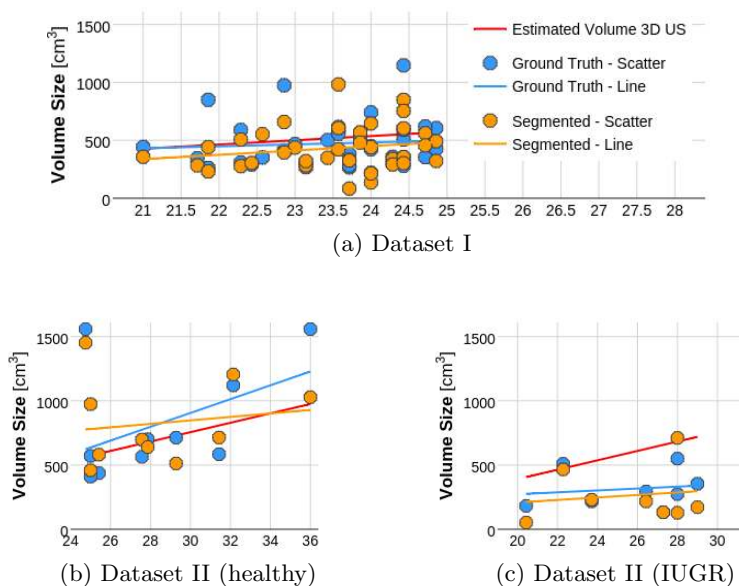
(a) Dataset I

(b) Dataset II (healthy)

(c) Dataset II (IUGR)

Fig. 5. A graph comparing automatic segmentations (our approach), the ground truth (expert) and the linear estimations [11] of the placental volumes versus their gestational ages. (a) shows that our results from the first experiment [exp-1] using *Dataset I* are very close to both the expert and theoretical estimations. However, the third experiment [exp-3] using healthy subjects from *Dataset II* shows less consistency of the segmented volumes (b) due to the large dissimilarity test data. Moreover, (c) shows more inconsistency by testing on fetuses with IUGR from [exp-3]. (fetuses with unknown gestational age were excluded)


Fig. 6. A comparisons between the 3D reconstructed placenta using superpixel-based PVR reconstructions [3] with initial superpixel size $20x20$ pixels and a baseline using directly 2D slices from the input stacks for multi-planar examination. (stacks with high motion artifacts were excluded)

from different scanners. Moreover, we extend our framework scope to real clinical applications by compensating motion artifacts using slice to volume registration techniques, as well as providing a novel standardized view into the placental structures using skeleton extraction and curved planar reformation. In future work we will investigate the potential use of the standardized placenta views for image-based classification and automatic detection of abnormalities.

References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Susstrunk, S.: SLIC superpixels compared to state-of-the-art superpixel methods. Pattern Analysis and Machine Intelligence, IEEE Transactions on 34(11), 2274–2282 (2012)
2. Alansary, A., Lee, M., Keraudren, K., Kainz, B., Malamateniou, C., Rutherford, M., Hajnal, J.V., Glocker, B., Rueckert, D.: Automatic Brain Localization in Fetal MRI Using Superpixel Graphs. In: ICML, pp. 13–22. LNCS 9487 (2015)
3. Kainz, B., Alansary, A., Malamateniou, C., Keraudren, K., Rutherford, M., Hajnal, J.V., Rueckert, D.: Flexible reconstruction and correction of unpredictable motion from stacks of 2D images. In: MICCAI, pp. 555–562. Springer (2015)
4. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient Multi-Scale 3D CNN with Fully Connected CRF for Accurate Brain Lesion Segmentation. arXiv preprint arXiv: 1603.05959 (2016)
5. Kanitsar, A., Fleischmann, D., Wegenkittl, R., Felkel, P., Groller, E.: CPR – curved planar reformation. In: IEEE VIS 2002. pp. 37–44 (Nov 2002)
6. Keraudren, K., Kainz, B., Oktay, O., Kyriakopoulou, V., Rutherford, M., Hajnal, J.V., Rueckert, D.: Automated Localization of Fetal Organs in MRI Using Random Forests with Steerable Features. In: MICCAI, pp. 620–627. LNCS 9351 (2015)
7. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected CRFs with gaussian edge potentials. Adv. Neural Inf. Process. Syst (2012)
8. Routledge, E., Malamateniou, C., Rutherford, M.: Can MR imaging of the placenta demonstrate a distinct placental phenotype in normal and abnormal pregnancies? Tech. rep., King's College London, University of London (April 2015)
9. Stevenson, G.N., Collins, S.L., Ding, J., Impey, L., Noble, J.A.: 3-D Ultrasound Segmentation of the Placenta Using the Random Walker Algorithm: Reliability and Agreement. Ultrasound in medicine & biology 41(12), 3182–3193 (2015)
10. Tagliasacchi, A., Alhashim, I., Olson, M., Zhang, H.: Mean Curvature Skeletons. Computer Graphics Forum 31(5), 1735–1744 (2012)
11. Thame, M., Osmond, C., Bennett, F., Wilks, R., Forrester, T.: Fetal growth is directly related to maternal anthropometry and placental volume. European journal of clinical nutrition 58(6), 894–900 (2004)
12. Wang, G., Zuluaga, M.A., Pratt, R., Aertsen, M., David, A.L., Deprest, J., Vercauteren, T., Ourselin, S.: Slic-Seg: slice-by-slice segmentation propagation of the placenta in fetal MRI using one-plane scribbles and online learning. In: MICCAI, pp. 29–37. Springer (2015)

---

[1] Medical Imaging Interaction Toolkit (MITK) http://mitk.org/