

Slic-Seg: Slice-by-Slice Segmentation Propagation of the Placenta in Fetal MRI Using One-Plane Scribbles and Online Learning

Guotai Wang¹, Maria A. Zuluaga¹, Rosalind Pratt^{1,2}, Michael Aertsen³,
Anna L. David², Jan Deprest⁴, Tom Vercauteren¹, and Sebastien Ourselin¹

¹Translational Imaging Group, CMIC, University College London, UK

²Institute for Women's Health, University College London, UK

³Department of Radiology, University Hospitals KU Leuven, Belgium

⁴Department of Obstetrics, University Hospitals KU Leuven, Belgium

Abstract. Segmentation of the placenta from fetal MRI is critical for planning of fetal surgical procedures. Unfortunately, it is made difficult by poor image quality due to sparse acquisition, inter-slice motion, and the widely varying position and orientation of the placenta between pregnant women. We propose a minimally interactive online learning-based method named Slic-Seg to obtain accurate placenta segmentations from MRI. An online random forest is first trained on data coming from scribbles provided by the user in one single selected start slice. This then forms the basis for a slice-by-slice framework that segments subsequent slices before incorporating them into the training set on the fly. The proposed method was compared with its offline counterpart that is with no retraining, and with two other widely used interactive methods. Experiments show that our method 1) has a high performance in the start slice even in cases where sparse scribbles provided by the user lead to poor results with the competitive approaches, 2) has a robust segmentation in subsequent slices, and 3) results in less variability between users.

1 Introduction

The placenta plays a critical role in the health of the fetus during pregnancy. Abnormalities in the placental vasculature such as occur in twin-to-twin trans-fusion syndrome (TTTS) [4], can result in unequal blood distribution and a poor outcome or death for one or both twins. Placenta accreta, which is caused by an abnormally adherent placenta invading the myometrium, increases the risk of heavy bleeding during delivery. Minimally-invasive fetoscopic surgery provides an effective treatment for such placental abnormalities, and surgical planning is critical to reduce treatment-related morbidity and mortality.

With advantages such as large field of view, lack of ionizing radiation and good soft tissue contrast, Magnetic Resonance Imaging (MRI) is widely used for general surgical planning, but high-quality MRI for a fetus is difficult to achieve, since the free movement of the fetus in the uterus can cause severe motion artifacts [7]. The Single Shot Fast Spin Echo (SSFSE) allows the motion




Fig. 1. Examples of fetal MRI. (a), (b) and (c) are from one patient while (d) is from another. Note the motion artifacts and different appearance in odd and even slices in (b) and (c). The position of the placenta is anterior in (a), but posterior in (d).

artifacts to be nearly absent in each slice, but inter-slice motions still corrupt the volumetric data. The slices are acquired in an interleaved spatial order, which leads to different appearance between odd and even slices as shown in Fig. 1. In addition, fetal MRI is usually sparsely acquired with a large inter-slice spacing for a good contrast-to-noise ratio. Although some novel reconstruction techniques [5] can get super-resolution volume data of fetal brain from sparsely acquired slices, they have yet to demonstrate their utility for placental imaging and require a dedicated non-standard acquisition protocol. These factors bring several challenges to the segmentation of the placenta from clinical MR data.

The low-quality volumetric data with high-quality slices motivates employing 2D segmentation methods with a slice-by-slice strategy. Automatic methods rarely work well with medical images due to ambiguous appearance cues. Prior-knowledge brought from different patients in the form of shape/appearance models or propagated atlases [6] may help make the segmentation more robust, but the position and orientation of the placenta within the uterus varies greatly between pregnancies (see Fig. 1(a) and Fig. 1(d)), making it hard to model such statistical prior-knowledge. In contrast, interactive segmentation has been widely used in practice, where scribbles given by user provide useful information for accurate segmentation. A convenient interactive method should make full use of scribbles to get accurate segmentation with only a few number of user interactions. Traditional methods such as snakes or generalized gradient vector flow (GGVF) [10] use only the spatial information of an initial contour provided by the user, others such as Graph Cuts [2] and Geodesic Framework [1,3] take advantage of low level features to estimate the probability that a pixel belongs to the foreground or background.

In this paper, we propose a learning-based semi-automatic approach named Slic-Seg for segmentation of the placenta in fetal MRI. It is different from traditional interactive segmentation methods in the following ways: 1) It aims to make full use of user inputs to improve the accuracy and reduce number of user interactions. 2) Online Random Forest (RF) is employed for effective learning based on mid-level features, allowing the training set to be expanded on the fly. As a result, the method can achieve a high performance with a minimal number of user inputs.

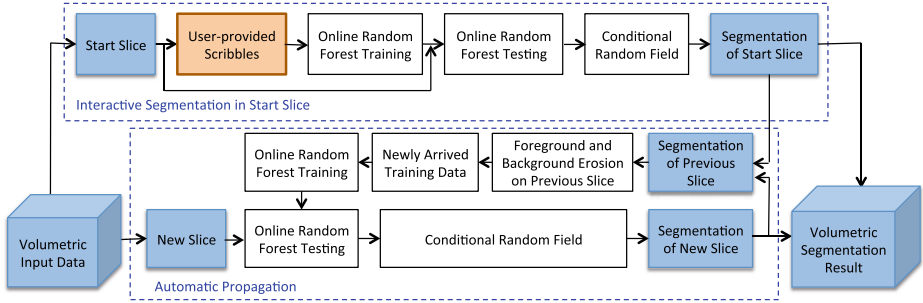


Fig. 2. The workflow of our Slic-Seg framework. User interaction is only required in the start slice. Other slices are segmented sequentially and automatically.

2 Methods

The workflow of our proposed Slic-Seg is shown in Fig. 2. A user selects a start slice and draws a few scribbles in that slice to indicate foreground and background. Online RF efficiently learns from these inputs and predicts the probability that an unlabeled pixel belongs to foreground or background. That probability is incorporated into a Conditional Random Field (CRF) to get the segmentation result based on which new training data is automatically obtained and added to the training set of RF predictor on the fly. To get the segmentation result from a volumetric placenta data, other slices are segmented sequentially and automatically without more user interactions.

Preprocess and Feature Extraction. Odd and even slices are rigidly aligned together to correct the motion artifacts, and histogram matching is implemented to address the different contrast between slices. For each pixel, features are extracted from a 9×9 pixel region of interest (ROI) centered on it. In each ROI, we extract gray level features including mean and standard deviation of intensity, texture features acquired by gray level co-occurrence matrix (GLCM) and wavelet coefficient features based on Haar wavelet.

Online Random Forests Training. A Random Forest [9] is a collection of binary decision trees composed of split nodes and leaf nodes. The training set of each tree is randomly resampled from the entire labeled training set (label 1 for the placenta and label 0 for background). At a split node, a binary test is executed to minimize the uncertainty of the class label in the subsets based on Information Gain. The test functions are of the form $f(\mathbf{x}) > \theta$, where \mathbf{x} is the feature vector of one sample, $f(\cdot)$ is a linear function, and θ is a threshold. At a leaf node, labels of all the training samples that have been propagated to that node are averaged, and the average label is interpreted as the posterior probability of a sample belonging to the placenta, given that the sample has fallen into that leaf node.

The training data in our application is obtained in one of two ways according on segmentation stage. For the start slice, training data comes from the scribbles provided by the user. During the propagation, after one slice is segmented, skeletonization of the placenta was implemented by morphological operators, and the background is eroded by a kernel with a certain radius (e.g., 10 pixels). New training data is obtained from the morphological operation results in that slice and added to existing training set of RF on the fly. To deal with online training, we use the online bagging [8] method to model the sequential arrival of training data as a Poisson distribution $\text{Pois}(\lambda)$ where λ is set to a constant number. Each tree is updated on each new training sample k times in a row where k is a random number generated by $\text{Pois}(\lambda)$.

Online Random Forests Testing. During the testing, each pixel sample \mathbf{x} is propagated through all trees. For the n th tree, a posterior probability $p_n(\mathbf{x})$ is obtained from the leaf that the test sample falls into. The final posterior is achieved as the average across all the N trees.

$$p(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N p_n(\mathbf{x}) \quad (1)$$

Inference Using Conditional Random Field. In the prediction of RF, the posterior probability for each pixel is obtained independently and it is sensitive to noise. To reduce the effect of noise and obtain the final label set for all the pixels in a slice, a CRF is used for a global spatial regularization. The label set of a slice is determined by minimizing the following energy function:

$$E(\mathbf{c}) = -\alpha \sum_i \Psi(\mathbf{c}_i | \mathbf{x}_i, I) - \sum_{i,j} \Phi(\mathbf{c}_i, \mathbf{c}_j | I) \quad (2)$$

where the unary potential $\Psi(\mathbf{c}_i | \mathbf{x}_i, I)$ is computed as $\log p(\mathbf{c}_i | \mathbf{x}_i, I)$ for assigning a class label \mathbf{c}_i to the i th pixel in a slice I , and p comes from the output of RF. The pairwise potential $\Phi(\mathbf{c}_i, \mathbf{c}_j | I)$ is defined as a contrast sensitive Potts model $\phi(\mathbf{c}_i, \mathbf{c}_j, \mathbf{g}_{ij})$ [2] where \mathbf{g}_{ij} measures the difference in intensity between the neighboring pixels and can be computed very efficiently. α is a coefficient to adjust the weight between unary potential and pairwise potential. The energy minimization is solved by a max flow algorithm [2]. A CRF is used in every slice of the volumetric image, and after the propagation, we stack the segmentation of all slices to construct the final volumetric segmentation result.

3 Experiments and Results

Experiment Data and Setting. MRI scanning of 6 fetuses in the second trimester were collected. For each fetus we had two volumetric data in different views that were used independently: 1), axial view with slice dimension

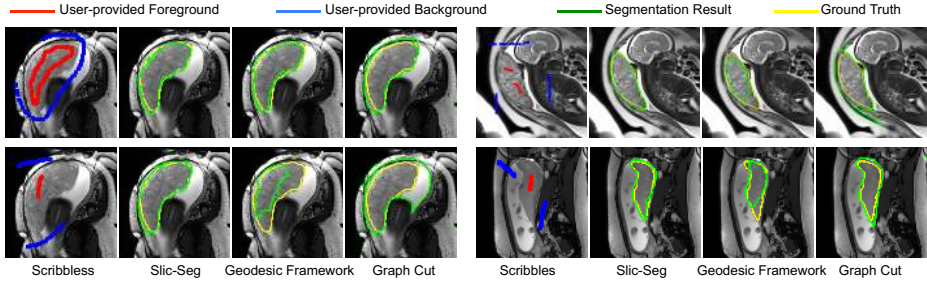


Fig. 3. Visual comparison of segmentation in the start slice by different methods. Upper left: user inputs are extensive, all methods result in a good segmentation. Lower left: user inputs are reduced in the same slice, only Slic-Seg preserves the accuracy. Right: two more examples show Slic-Seg has a better performance than Geodesic Framework and Graph Cut with only a few user inputs.

512×448, voxel spacing 0.7422 mm×0.7422 mm, slice thickness 3mm. 2) sagittal view with slice dimension 256×256, voxel spacing 1.484mm×1.484mm, slice thickness 4mm. A start slice in the middle region of the placenta was selected from each volumetric image, and 8 users provided scribbles in the start slice. A manual ground truth for each slice was produced by an experienced radiologist. The algorithm was implemented in C++ with a MATLAB interface. Parameter setting was: $\lambda=1$, $N=20$, $\alpha=4.8$. We found the segmentation was not sensitive to α in the range of 2 to 15 (see supplementary). The depth of trees was 10.

Results and Evaluation. We compared Slic-Seg with two widely used interactive segmentation methods: Geodesic Framework¹ of Bai and Sapiro [1] and Graph Cut [2]. Fig. 3 shows four examples of interactive segmentation in the start slice. In each subfigure, the same scribbles were used by different segmentation methods. On the left side of Fig. 3, the same slice was used with different scribbles. In the upper left case, scribbles provided by the user almost roughly indicate the boundary of the placenta, and all of the three methods obtain good segmentation results. In the lower left case, scribbles are reduced to a very small annotation set, Geodesic Framework and Graph Cut fail to preserve their performance, but Slic-Seg can still get a rather accurate segmentation. Two more cases on the right of Fig. 3 also show Slic-Seg can successfully segment the placenta using only a few number of scribbles.

In the propagation, the above three methods used the same morphological operations as mentioned previously to automatically generate foreground and background seeds for a new slice. In addition, we compared Slic-Seg with its offline counterpart where only user inputs in the start slice were used for training of an offline RF. Fig. 4 shows an example of propagation by different methods with the same user inputs in the start slice. S_i represents the i th slice following the start slice. In Fig. 4, though a good segmentation is obtained in the start

¹ Implementation from: [http://www.robots.ox.ac.uk/~sim\\$vgg/software/iseg/](http://www.robots.ox.ac.uk/~sim$vgg/software/iseg/)

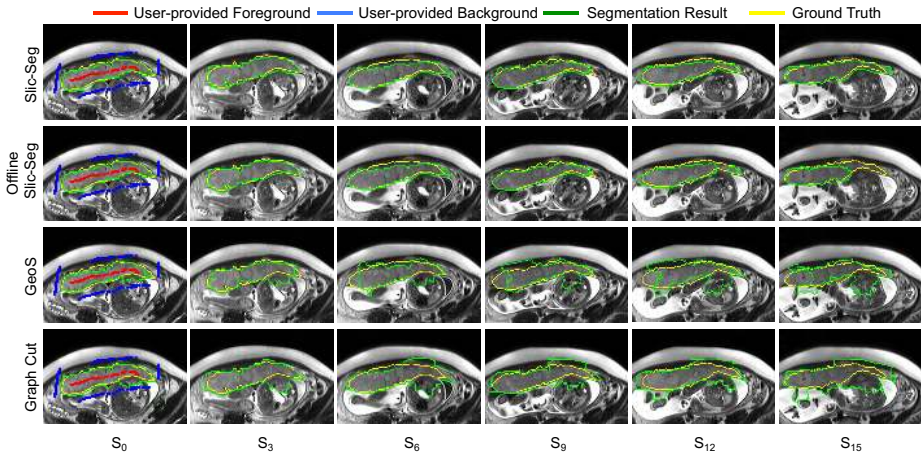


Fig. 4. Propagation of different methods with the same start slice and scribbles. S_i represents the i th slice following the start slice. User provided scribbles in S_0 are extensive and all methods have a good segmentation in that slice. However, during the propagation, only Slic-Seg keeps a high performance. (More slices are shown in the supplementary video).

slice due to an extensive set of scribbles, the error of Geodesic Framework and Graph Cut become increasingly large during the propagation. In a slice that is close to the start slice (e.g. $i \leq 9$), offline Slic-Seg can obtain a segmentation comparable to that of Slic-Seg. When a new slice (e.g. $i \geq 12$) is further away from the start slice, offline Slic-Seg fails to track the placenta with high accuracy. In contrast, online Slic-Seg has a stable performance during the propagation.

Quantitative evaluation was achieved by calculating the Dice coefficient and symmetric surface distance (SSD) between segmentation results and the ground truth. Fig. 5 shows the Dice coefficient and SSD for each slice in one volumetric image (the same image as used in Fig. 4). For each slice, we use error bars to show the first quartile, median and the third quartile of the Dice coefficient and SSD. Fig. 5 shows that Slic-Seg has a better performance in the start slice and during the propagation than offline Slic-Seg, Geodesic Framework and Graph Cut. The less dispersion of Slic-Seg indicates its less variability between users. Fig. 6 shows the evaluation results on data from all the patients. We present Dice and SSD in both the start slice and the whole image volume.

Discussion. The experiments show that Slic-Seg using RF, CRF and segmentation propagation has better performances in the start slice and during propagation than Geodesic Framework and Graph Cut. This is due to the fact that the last two methods use low level appearance features to model placenta and background, which may not be accurate enough in fetal MRI images with poor quality. In contrast, the RF in our method uses mid-level features of multiple aspects including intensity, texture and wavelet coefficients, which may provide

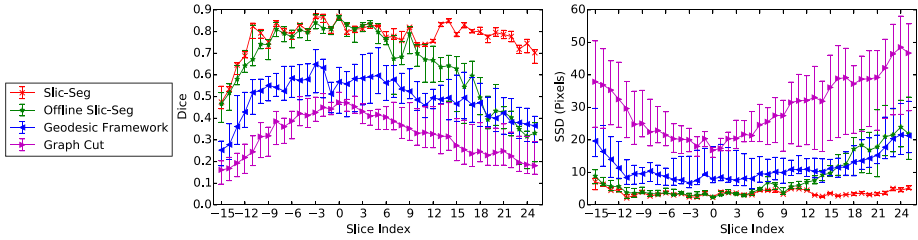


Fig. 5. Evaluation on one image volume in terms of Dice (left) and SSD (right) in each slice (evaluation on other image volumes can be found in the supplementary). Each error bar shows the median, first quartile and third quartile across all the 8 users, each of which segmented the image twice with different scribbles. Note that Slic-Seg has a high accuracy in the start slice and during the propagation, with low variability among different users.




Fig. 6. Evaluation on data from all the 6 patients (each having 2 orthogonal datasets) in terms of Dice (left) and SSD (right) in the start slice and the whole image volume. Each of 8 users segmented these images twice with different scribbles. Note Slic-Seg and offline Slic-Seg get the same result in the start slice. Slic-Seg has a high performance with less variability in both the start slice and the whole image volume. The p value between Slic-Seg and offline Slic-Seg on the image volumes is 0.0043 for Dice, and 0.0149 for SSD.

a better description of the differences between the placenta and background. Because the appearance of the placenta in a remote slice could be different from that in the start slice, the offline RF that only uses user-provided scribbles for training may give a poor prediction after propagating along several slices. The online RF that accepts sequentially obtained training data addresses this problem and is adaptive to the appearance change, which leads to a more robust segmentation during the propagation. The short error bars of Slic-Seg in Fig. 5 and Fig. 6 indicate that the performance of this method has a low variability among different users. Though our method requires user interactions only in the start slice, it could allow user corrections with some additional scribbles when the segmentation propagates to terminal slices. Since it needs fewer user interactions and allows the training data to be expanded efficiently, the segmentation can be conveniently improved with little additional user efforts.

4 Conclusion

We present an interactive, learning-based method for the segmentation of the placenta in fetal MRI. Online RF is used to efficiently learn from mid-level features describing placenta and background, and it is combined with CRF for labelling. The slice-by-slice segmentation only requires user inputs in a start slice, and other slices are segmented sequentially and automatically to get a volumetric segmentation. Experiments show that the proposed method achieves high accuracy with minimal user interactions and less variability than traditional methods. It has a potential to provide an accurate segmentation of the placenta for fetal surgical planning. In the future, we intend to combine sparse volumetric data in different views for a 3D segmentation.

Acknowledgements. This work was supported through an Innovative Engineering for Health award by the Wellcome Trust [WT101957]; Engineering and Physical Sciences Research Council (EPSRC) [NS/A00027/1], the EPSRC (EP/H046410/1, EP/J020990/1, EP/K005278), the National Institute for Health Research University College London Hospitals Biomedical Research Centre (NIHR BRC UCLH/UCL High Impact Initiative), a UCL Overseas Research Scholarship and a UCL Graduate Research Scholarship.

References

1. Bai, X., Sapiro, G.: A Geodesic Framework for Fast Interactive Image and Video Segmentation and Matting. *IJCV* 82(2), 113–132 (2008)
2. Boykov, Y., Jolly, M.P.: Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In: *ICCV 2001*, vol. 1, pp. 105–112, July 2001
3. Criminisi, A., Sharp, T., Blake, A.: GeoS: Geodesic image segmentation. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008, Part I*. LNCS, vol. 5302, pp. 99–112. Springer, Heidelberg (2008)
4. Deprest, J.A., Flake, A.W., Gratacos, E., Ville, Y., Hecher, K., Nicolaides, K., Johnson, M.P., Luks, F.I., Adzick, N.S., Harrison, M.R.: The Making of Fetal Surgery. *Prenatal Diagnosis* 30(7), 653–667 (2010)
5. Gholipour, A., Estroff, J.A., Warfield, S.K., Member, S.: Robust Super-Resolution Volume Reconstruction From Slice Acquisitions: Application to Fetal Brain MRI. *IEEE TMI* 29(10), 1739–1758 (2010)
6. Habas, P.A., Kim, K., Corbett-Detig, J.M., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C.: A Spatiotemporal Atlas of MR Intensity, Tissue Probability and Shape of the Fetal Brain with Application to Segmentation. *NeuroImage* 53, 460–470 (2010)
7. Kainz, B., Malamateniou, C., Murgasova, M., Keraudren, K., Rutherford, M., Hajnal, J.V., Rueckert, D.: Motion corrected 3D reconstruction of the fetal thorax from prenatal MRI. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) *MICCAI 2014, Part II*. LNCS, vol. 8674, pp. 284–291. Springer, Heidelberg (2014)

8. Saffari, A., Leistner, C., Santner, J., Godec, M., Bischof, H.: On-line random forests. In: ICCV Workshops 2009 (2009)
9. Schroff, F., Criminisi, A., Zisserman, A.: Object class segmentation using random forests. In: BMVC 2008, pp. 54.1–54.10 (2008)
10. Xu, C., Prince, J.L.: Snakes, Shapes, and Gradient Vector Flow. IEEE TIP 7(3), 359–369 (1998)