



UNIVERSITY OF LEEDS

This is a repository copy of *A Deep Attentive Convolutional Neural Network for Automatic Cortical Plate Segmentation in Fetal MRI*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/180017/>

Version: Accepted Version

Article:

Dou, H orcid.org/0000-0001-8628-5489, Karimi, D, Rollins, CK et al. (7 more authors) (2021) A Deep Attentive Convolutional Neural Network for Automatic Cortical Plate Segmentation in Fetal MRI. *IEEE Transactions on Medical Imaging*, 40 (4). pp. 1123-1133. ISSN 0278-0062

<https://doi.org/10.1109/tmi.2020.3046579>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

A Deep Attentive Convolutional Neural Network for Automatic Cortical Plate Segmentation in Fetal MRI

Haoran Dou, Davood Karimi, Caitlin K. Rollins, Cynthia M. Ortinau, Lana Vasung, Clemente Velasco-Annis, Abdelhakim Ouaalam, Xin Yang, Dong Ni, Ali Gholipour, *Senior Member, IEEE*

Abstract—Fetal cortical plate segmentation is essential in quantitative analysis of fetal brain maturation and cortical folding. Manual segmentation of the cortical plate, or manual refinement of automatic segmentations is tedious and time-consuming. Automatic segmentation of the cortical plate, on the other hand, is challenged by the relatively low resolution of the reconstructed fetal brain MRI scans compared to the thin structure of the cortical plate, partial voluming, and the wide range of variations in the morphology of the cortical plate as the brain matures during gestation. To reduce the burden of manual refinement of segmentations, we have developed a new and powerful deep learning segmentation method. Our method exploits new deep attentive modules with mixed kernel convolutions within a fully

convolutional neural network architecture that utilizes deep supervision and residual connections. We evaluated our method quantitatively based on several performance measures and expert evaluations. Results show that our method outperforms several state-of-the-art deep models for segmentation, as well as a state-of-the-art multi-atlas segmentation technique. We achieved average Dice similarity coefficient of 0.87, average Hausdorff distance of 0.96 mm, and average symmetric surface difference of 0.28 mm on reconstructed fetal brain MRI scans of fetuses scanned in the gestational age range of 16 to 39 weeks (28.6 ± 5.3). With a computation time of less than 1 minute per fetal brain, our method can facilitate and accelerate large-scale studies on normal and altered fetal brain cortical maturation and folding.

Index Terms—Cortical plate, Automatic segmentation, Fetal MRI, Deep learning, Convolutional neural network, Attention

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This work was supported in part by the National Institutes of Health (NIH) under Grant R01 EB018988, Grant R01 NS106030, Grant K23 NS101120, and Grant K23 HL141602; in part by the Technological Innovations in Neuroscience Award from the McKnight Foundation; in part by awards from the Mend A Heart Foundation; in part by the Pediatric Heart Network; in part by the American Academy of Neurology; in part by the Brain and Behaviour Research Foundation; in part by the National Key Research and Development Program of China under Grant 2019YFC0118300; and in part by the Shenzhen Peacock Plan under Grant KQTD2016053112051497 and Grant KQJSCX20180328095606003. The work of Lana Vasung was supported by the Ralph Schlager fellowship at Harvard University.

H. Dou is with the National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, Guangdong Key Laboratory for Biomedical Measurements and Ultrasound Imaging, School of Biomedical Engineering, Health Science Center, Shenzhen University, Shenzhen, China; and with the Medical UltraSound Image Computing (MUSIC) Lab, Shenzhen University, China; and also with the Computational Radiology Laboratory in the Radiology Department at Boston Children’s Hospital, Harvard Medical School, Boston, Massachusetts, USA (email:douhaoran2017@email.szu.edu.cn).

D. Karimi, and A. Gholipour are with the Computational Radiology Laboratory (CRL) in the Department of Radiology at Boston Children’s Hospital, and Harvard Medical School, Boston, Massachusetts, USA (email:davood.karimi@childrens.harvard.edu, ali.gholipour@childrens.harvard.edu).

C.K. Rollins is with the Department of Neurology at Boston Children’s Hospital, and Harvard Medical School, Boston, Massachusetts, USA.

C.M. Ortinau is with the Department of Pediatrics in Washington University School of Medicine in St. Louis, Missouri, USA.

L. Vasung is with the Department of Pediatrics at Boston Children’s Hospital, and Harvard Medical School, Boston, Massachusetts, USA.

C. Velasco-Annis, and A. Ouaalam are with the CRL in the Department of Radiology at Boston Children’s Hospital, Boston, Massachusetts, USA.

X. Yang, and D. Ni are with the National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, Guangdong Key Laboratory for Biomedical Measurements and Ultrasound Imaging, School of Biomedical Engineering, Health Science Center, Shenzhen University, Shenzhen, China, and also with the MUSIC Lab, Shenzhen University, China (email:nidong@szu.edu.cn)

Corresponding authors: Dong Ni and Ali Gholipour.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or the foundations.

The code for the techniques presented in this study can be found at: <https://github.com/bchimagine>, or <https://github.com/wulalago/FetalCPSeg>

I. INTRODUCTION

FETAL magnetic resonance imaging (MRI) has been established as a reliable method for quantitative evaluation of cortical development in the fetus, and for the diagnosis and analysis of congenital neurological disorders [1]. It is a safe technique that provides much better soft tissue contrast than ultrasound. It also provides advanced mechanisms to image the micro-structure and function of the fetal brain *in-utero* [2]. Hence, it provides information that cannot be obtained by any other imaging technique. Fetal MRI has enabled *in-vivo* mapping and quantitative analysis of cortical maturation in the fetus [3]–[9], and has shed light on patterns and forces that govern cortical folding and expansion [10]–[13]. It has also enabled studies that have quantified altered cortical development in fetuses with congenital disorders of the brain and heart [14]–[18]. These studies require accurate segmentation of the developing fetal brain tissue, in particular the cortical plate, on fetal MRI.

Fetal brain tissue segmentation on MRI has been historically challenged due to the limitations of imaging the fetal brain *in-utero* by MRI. In particular, fetal movements disrupt the spatial encoding that is necessary to acquire real 3D images of the fetal brain. Therefore, fetal MRI is performed through 2D stack-of-slice acquisitions that do not make 3D images with coherent anatomic boundaries in 3D. Additionally, there are technical challenges in acquiring 3D fetal brain MRI scans. Hence, there has been a lack of carefully-labeled 3D (and more importantly spatiotemporal or 4D) images of the fetal brain MRI that can be used as atlases for multi-atlas automatic segmentation, or as training data for learning-based segmentation methods. These issues, however, have been addressed in

the recent years through the development of super-resolution fetal brain MRI reconstruction methods, e.g. [19]–[24]. These methods, in-turn, led to the construction of carefully labelled, spatiotemporal atlases of fetal brain MRI [25]. These developments allowed automatic segmentation of fetal brain tissue, mainly through multi-atlas segmentation methods [25]–[27]. Nonetheless, accurate segmentation of the fetal brain tissue, in particular segmentation of thin structures such as the cortical plate (CP), remains challenging due to several outstanding issues that we discuss next.

A. Fetal Cortical Plate Segmentation

The CP was manually or semi-automatically segmented on MRI in early studies on quantitative analysis of fetal cortical maturation [3], [14], [15]. Manual segmentation of the super-resolved reconstructed fetal brain MRI scans, however, is laborious and time consuming [25]. Therefore, more recent studies relied upon careful manual refinement of automatic (multi-atlas based) fetal brain MRI segmentations [8], [18]. Nonetheless, manual refinement of automatic segmentations is also tedious and difficult on super-resolved 3D images.

Cortical plate segmentation on fetal MRI is particularly challenging as the fetal CP is a very thin ribbon with a thickness that is comparable to the best achievable resolution on fetal MRI scans. As a result, partial voluming effects cause major issues in fetal CP segmentation on MRI. Furthermore, because the tissues are immature, the contrast of the developing white matter (WM) and gray matter (GM) for the fetal (and newborn) brain on MRI is the reverse of the contrast of the WM and GM on MRI of mature brains [28], [29]. This exacerbates the partial voluming effect as partial voluming in the CP and the cerebrospinal fluid (CSF) interface generates a WM-like intensity that confuses automatic segmentation algorithms. This can cause registration errors and inaccurate segmentations that appear as holes and topological errors in thin or highly-folded areas of the fetal cortex. This issue can be seen in Fig. 1, which shows T2-weighted MRI images of the brain of two fetuses scanned at 23 weeks and 35 weeks gestational age as well as T2-weighted brain MRI images of a newborn and an adult. Compared to the adult brain, in which the GM intensity falls between the range of the intensities of the WM and CSF, in the fetal and newborn brains the WM appears brighter than the CP. This is problematic as the partial voluming between CP and CSF resembles WM, which may then result in topological errors in segmentation.

Another factor that makes automatic segmentation of the fetal CP challenging is the substantial variations in fetal brain morphology due to the rapid development of the brain throughout gestation. As Fig. 1 shows, the CP of the fetal brain in the second trimester (e.g. at 23 weeks) has a smooth shape and does not have many sulci. With the rapid growth of the fetus, sulci form and the CP folds, ultimately forming an adult-like, highly-folded area just before birth, which then continues to mature after birth. The large variations in morphology, appearance, and scale of the fetal CP, along with the challenges discussed earlier, make it difficult to develop effective and robust solutions for automatic CP segmentation in fetal MRI.

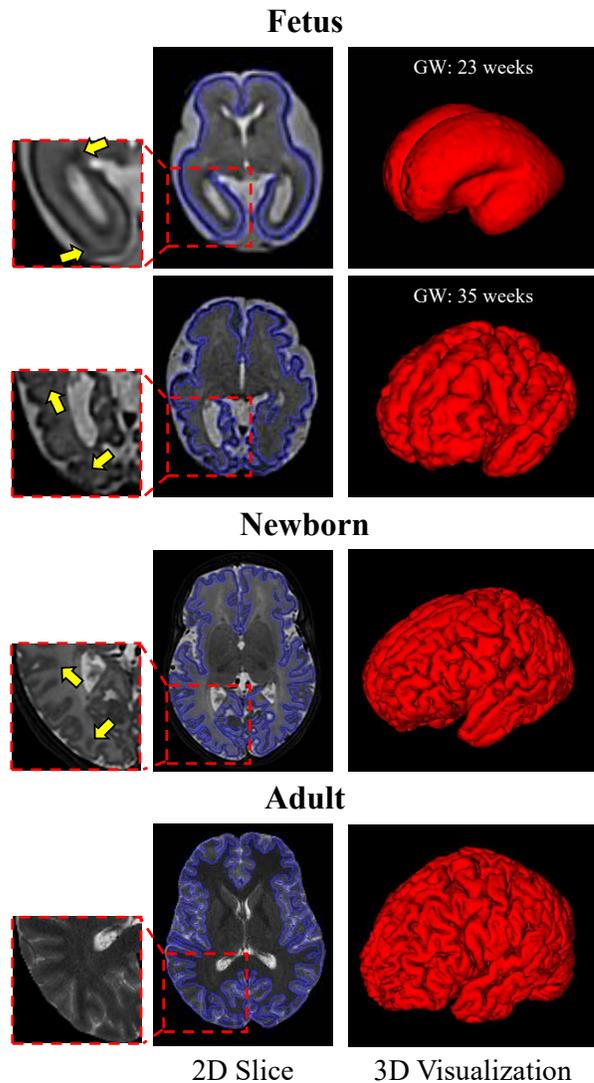


Fig. 1. Comparing the appearance of the fetal brain (in the second and third trimesters) with the brain of a newborn and an adult on T2-weighted MRI. The first column shows an axial slice of brain MRI where the blue line shows the boundary of the cortical plate. The second column is the 3D surface mesh visualization of the CP. Compared to the newborn and adult brains, the fetal brain exhibits wide variations in size and morphology. In addition, the relatively low spatial resolution of fetal MRI, the thin size of the CP, and the reversed contrast of the GM and WM that causes WM-like intensities in the CP-CSF interface (due to partial voluming) (shown by arrows), pose significant challenges for automatic CP segmentation. GW: Gestational Weeks.

B. Related Works

For a review of automatic fetal and neonatal brain MRI segmentation techniques we refer to [27]. In our review of the related works here we focus on techniques that segment the fetal brain *tissue* on 3D reconstructed fetal brain MRI images. Therefore, we do not review techniques that only segment (i.e., extract) the entire fetal brain on the original (stack-of-slice) fetal MRI scans. Those techniques that extract and segment the entire fetal brain on slices of the original MRI scans are useful as a pre-processing step for slice-by-slice motion correction and 3D super-resolution reconstruction of fetal brain MRI. Our work, on the other hand, is focused on post-reconstruction analysis of fetal brain cortical development.

The early works on 3D fetal brain MRI tissue segmentation were based on deformable atlas-to-image registration and propagation of labels from MRI atlases to MRI scans of query subjects. Habas *et al.* [30] first proposed an atlas-based solution to segment the fetal brain tissue with a focus on the germinal matrix. Inspired by [30], they incorporated a depth-based geometric representation model into the conventional atlas segmentation strategy [30] by providing additional anatomical constraints [31]. In their follow-up work, Habas *et al.* constructed 3D statistical MRI atlases of fetal brain tissue distribution [32] and used them to segment fetal brain tissue on 3D reconstructed fetal brain MRI images [26].

Serag *et al.* [33] built a spatiotemporal atlas of the fetal brain from reconstructed fetal brain MRI scans of 80 fetuses scanned in the gestational age (GA) range of 23 to 37 weeks.¹ They used the atlas in a multi-atlas deformable registration and label propagation framework to segment fetal brain tissue on reconstructed fetal brain MRI images. With their atlas and segmentation technique they reported average Dice similarity coefficient (DSC) of 0.84 for automatic CP segmentation.

Gholipour *et al.* [34] developed a multi-atlas multi-shape segmentation technique based on a combination of label fusion from multiple age-group atlases and shape priors. They used their atlas to segment lateral ventricles in fetuses at different ages with different levels of ventriculomegaly. Gholipour *et al.* [25], [35] then developed and distributed fully labeled spatiotemporal (4D) MRI atlases of normal fetal brain growth at 1mm isotropic resolution in the gestational age range of 19-39 weeks.² They used reconstructed images of 81 normal fetuses to construct the atlas and used it for multi-atlas tissue segmentation where they achieved average DSC of 0.9 and 0.84 for CP segmentation (using a leave-one-out strategy) for fetuses scanned at ≤ 34 weeks and > 34 weeks, respectively.

Khalili *et al.* [36] developed a fully convolutional neural network based on the U-Net architecture to segment fetal brain tissue on 3D reconstructed fetal brain MRI scans of 12 fetuses in the 23-35 weeks GA range. They reported average test DSC of 0.835 for CP segmentation, which compared favorably against the average DSCs of 0.82 and 0.84 reported for the different test sets in [26] and [33], respectively. Although, as noted in [36], these values should not be compared directly because they were calculated on different test sets.

C. Contributions

To reduce the burden of manual refinement of segmentation and topology correction, and to improve quantitative analysis of fetal cortical plate development, we aimed to significantly improve automatic segmentation of the cortical plate in fetal MRI. We aimed to address the outstanding issues that were mentioned above, in particular to deal with the substantial variability in the size, shape, and complexity of the thin fetal brain CP. To this end, in this investigation we developed a new network architecture with novel attention modules using mixed kernel convolutions. The attention modules helped our network extract and learn important multi-scale information from the

feature maps in a stage-wise manner. We compared our trained model with several state-of-the-art models including: the 3D U-Net [37] used in fetal brain MRI segmentation [36], the Plane Aggregated U-Net (PAUNet) [38], the Attention U-Net [39], and the Squeeze & Excitation Fully Convolutional Network (SE-FCN) [40]. We compared methods quantitatively based on several performance metrics on held-out test sets. We also compared the results of our method against a multi-atlas segmentation technique based on blind expert evaluations. The results indicate significant improvement in CP segmentation in fetal MRI using our proposed network with attentive learning.

The paper is organized as follows: Section II involves the materials and methods including the details of the fetal MRI dataset, our proposed network architecture, the attention refinement module, loss function, and training. Section III describes the experimental results; Section IV includes a discussion; and Section V includes our concluding remarks.

II. MATERIALS AND METHODS

A. Fetal MRI data

The fetal MRI dataset used in this study consisted of super-resolution reconstructed volumes of 57 fetuses scanned at a gestational age (GA) between 16 and 39 weeks (29.5 ± 5.5). The scans were performed on 3T Siemens Skyra scanners with body matrix and spine coils. These images were reconstructed from slices of T2-weighted single shot fast spin echo scans of each fetus, repeated 3-4 times in each of the axial, coronal, and sagittal planes with respect to the fetal head. The protocol and the research imaging was approved by the institutional review board committee. Written informed consent was obtained from all pregnant women volunteers who participated in the study. The in-plane resolution of the original scans was 1 mm, the slice thickness was 2-3 mm, and the repetition time and echo time were 1500 and 120 ms, respectively.

Approximate fetal brain masks were automatically extracted on the original scans using the model in [41]. They were used along with slice-to-volume registration (for inter-slice motion correction) to reconstruct a super-resolved 3D volumes of the fetal brain using the algorithm described in [22] at an isotropic resolution of 0.8 mm^3 . Final 3D brain masks were generated on the reconstructed images using the Auto-Net [42], and were manually corrected in ITK-SNAP [43] as needed. Brain-extracted reconstructed volumes were then registered to a spatiotemporal fetal brain MRI atlas and the brain tissue and cortical plate were segmented using the procedure described in [25]. To generate reference (ground truth) segmentation data, this procedure involved manual segmentation and refinement of the cortical plate in several rounds in different planes. This procedure took anywhere between 2 to 8 hours depending on the age of the fetus and the quality of the reconstructions.

B. Model Architecture

Figure 2 shows a schematic representation of our proposed deep attentive fully convolutional neural network (CNN) architecture for CP segmentation in fetal MRI. Our model consists of a backbone network with an encoder-decoder architecture with forward skip connections from the encoder stages to the

¹<http://brain-development.org/brain-atlases/fetal-brain-atlases/>

²http://crl.med.harvard.edu/research/fetal_brain_atlas/

corresponding decoder stages. This is followed by a stage-wise attention refinement module that leverages mixed kernel convolutions to capture multi-scale contextual information. In this section and the following subsections, we explain the details of our model and its attentive learning that enables learning to segment the complex and variable structure of the cortical plate in fetuses scanned at different ages.

As shown in Fig. 2, our network works on 3D patches of size 64^3 . This patch size was chosen in a trade-off to enable learning multi-scale features while limiting memory requirements. For each 3D input patch, the network outputs a CP probability map of the same size in an end-to-end manner. The network first extracts a series of feature maps with different resolutions. The shallower feature maps (Stage 1) contain high-resolution details needed for accurate delineation of the CP boundary. The deeper features maps contain coarse and high-level information that help predict the overall outline of the CP. Our backbone network uses forward skip connections [37] and convolutions with residual connections [44]. The input of each stage is first processed by two $3 \times 3 \times 3$ convolutional layers followed by batch normalization (BN) [45] and parametric rectified linear unit (PReLU) for activation [46].

A shortcut (i.e., residual) connection is added between the input and the output of every convolutional block in this network. The numbers of feature maps in the encoder part of the network are 16, 32, 64, 128, 128, increasing the number of features as their size shrinks. We limited the number of feature maps at the last encoder stage to 128 to accommodate training on a graphical processing unit (GPU) with 8GB of memory. Every feature map computed by the backbone network is then upsampled using trilinear interpolation to the size of the input patch. Then a convolutional operation with a kernel size of $1 \times 1 \times 1$ is applied to every feature map to create 16 features in each of the feature maps. The feature maps then go through deep supervision modules [47] that improve the gradient flow and encourage learning more useful representations. Similar deep supervision modules are also used on the outputs of the attention modules and merge the resulting feature maps via concatenation. These feature maps go through two convolutional layers followed by BN and PReLU to produce the CP probability map.

C. Stage-wise Attentive Refinement

The goal of our attention modules is to increase the network's ability to capture the multi-scale details of the brain CP. We expect that these modules increase the richness of the information of the multi-scale feature maps learned by the backbone network. Figure 3 shows the architecture of our proposed attention module as it is used in each of the stages of our stage-wise attentional refinement network in Fig. 2. We explain the details of the attention module here.

As shown in Fig. 3, our attention module takes the feature maps generated from each stage of the backbone network and outputs a refined attentive feature map of the same size. The

operation of the attention module at stage i has the following functional form:

$$F'_i = f_i(F_i; \theta) \otimes F_i + F_i. \quad (1)$$

where θ denotes the parameters of the module that include the weights of the convolutional layers; F_i and F'_i respectively denote the input feature maps and the output refined attentive feature maps at stage i ; and \otimes is element-wise multiplication.

In our attention module, the input feature map F_i is first processed with group convolutional blocks. The sizes of the kernels vary from $3 \times 3 \times 3$ to $9 \times 9 \times 9$, but the number of feature maps in each group is fixed at 4. Each convolutional operation is followed by BN and PReLU. This design not only reduces the GPU memory footprint and computation, but also increases the richness of the learned representations as demonstrated by similar designs [48]. The generated feature maps are concatenated and passed through a second, and identical, group of convolutional blocks. A convolution layer with a kernel size of $1 \times 1 \times 1$ is used to merge these multi-scale feature maps into a single feature map. Finally, a sigmoid activation is applied to obtain the attentive map A_i . The attentive map is multiplied element-wise with the input feature map to encourage attention to the relevant locations in the feature map. The refined attentive feature map is added to the input feature map in the spirit of residual connections to reduce the difficulty of learning the attentive map. Note that we use BN on the feature maps from these two branches before adding them together. This is necessary to ensure that the values of the two feature maps are not very differently distributed.

D. Loss Function

We trained our network with the weighted binary cross-entropy loss, L_{wbce} :

$$L_{wbce} = \frac{1}{N} \sum_{i=1}^N [\alpha g_i \log p_i + (1 - g_i) \log (1 - p_i)]. \quad (2)$$

In this equation, N is the number of voxels in the patch. Moreover, p_i and g_i denote the predicted cortical plate probability map and the binary ground truth cortical plate map at voxel i . Lastly, α is the weight hyperparameter, which is set independently for each training mini-batch.

Our overall loss function is the weighted sum of the losses at different points in the network involving supervision:

$$L_{total} = \sum_{i=1}^n w^i L_{signal}^i + \sum_{j=1}^n w^j L_{signal}^j + w^p L_{signal}^p. \quad (3)$$

the above equation, w^i and L_{signal}^i denote the weight and loss for the points of supervision at stage i of the backbone network, as shown in Fig. 2. Similarly, w^j and L_{signal}^j denote the weight and loss for the point of supervision at stage j of the attentive refinement network. The number of stages is denoted with n , and w^p and L_{signal}^p denote, respectively, the weight and loss computed at the final network output. Using cross-validation in this study, we set $w^i = 0.8, 0.7, 0.6, 0.5|_{i=1:4}$, $w^j = 0.8, 0.7, 0.6, 0.5|_{j=1:4}$, and $w^p = 1$.

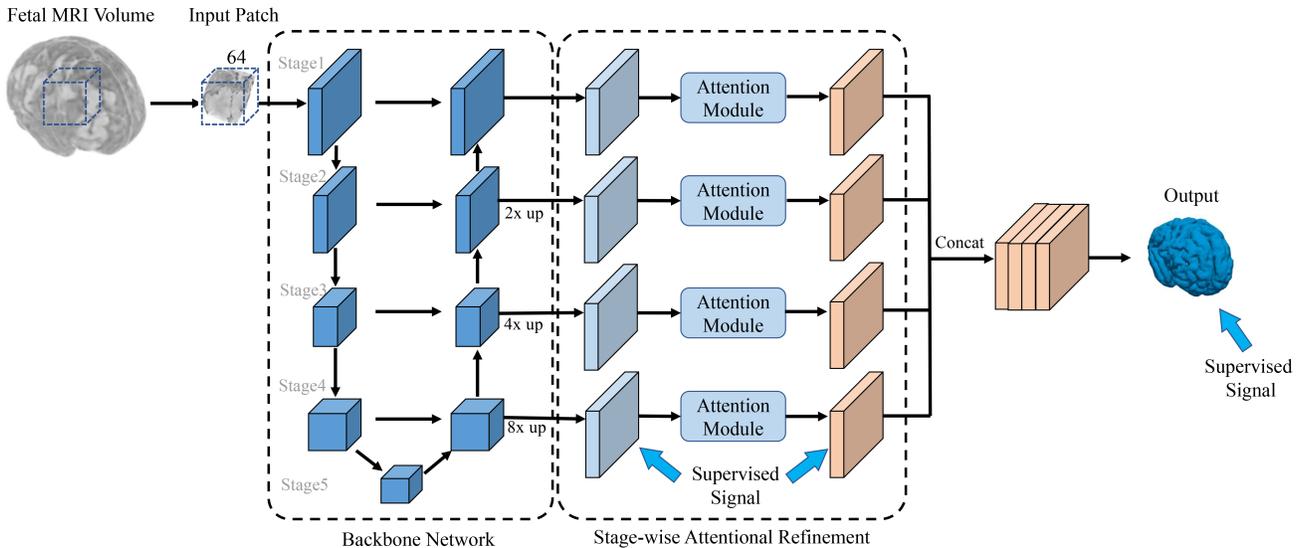


Fig. 2. Schematic illustration of our proposed CNN architecture, which consists of a backbone fully convolutional encoder-decoder network with stagewise forward skip connections, and a stagewise attention refinement module. The network takes sliding, overlapping 3D patches of size $64 \times 64 \times 64$ from the input image, and using training data, learns to generate a 3D segmentation of the input image.

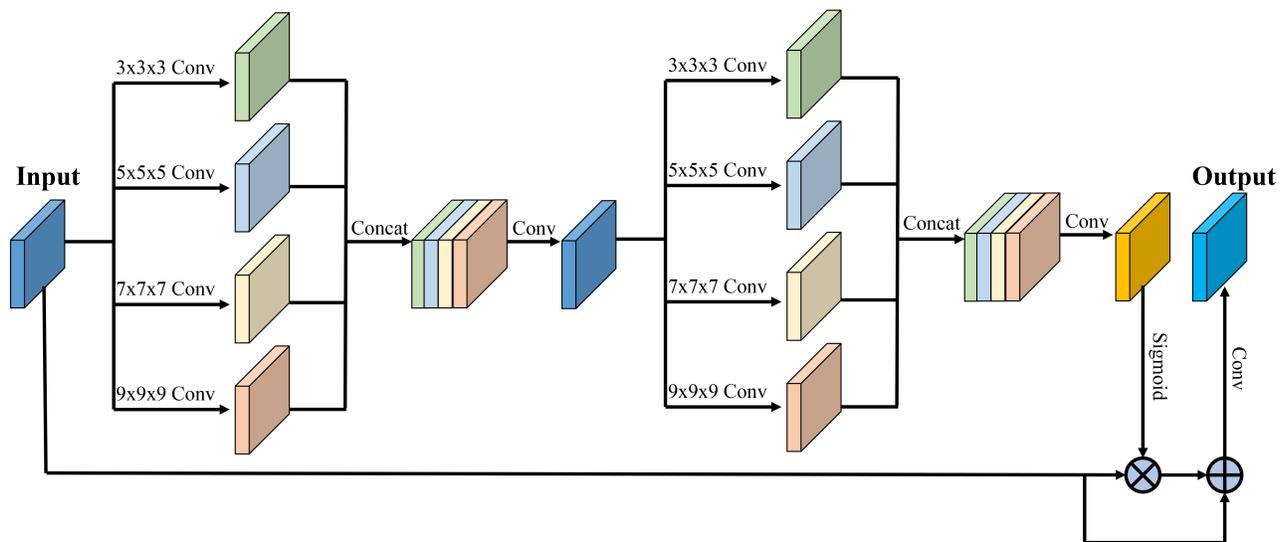


Fig. 3. Schematic illustration of the architecture of our attention module, which takes an input feature map and generates a refined attentive feature map of the same size. The module includes two layers of group convolutional blocks followed by feature concatenation and $1 \times 1 \times 1$ bottleneck convolution blocks. The module uses a sigmoid activation function and a residual connection between its input and output to reduce the difficulty of learning attentive maps.

E. Implementation and Training

We implemented our network in PyTorch [49] and trained it in an end-to-end manner using Adam optimizer [50] with a weight decay of $1e-4$. We used an initial learning rate of $1e-3$ and reduced it by a factor of 0.1 after 3k and 4.5k training iterations. Learning stopped after 6k iterations. We used random 3D flip and 90 degree rotations for augmentation of input patches during training. We also implemented and trained alternative networks for comparison with the same training data and with similar hyper-parameters for optimization and data augmentation. We trained all models on an Nvidia GTX-1080Ti GPU with a mini-batch size of 4. Training our network

with 6k iterations took approximately 24 hours.

III. EXPERIMENTS AND RESULTS

In this section, we first describe the alternative, state-of-the-art deep learning models that we have compared and the evaluation criteria that we used to compare the results. Then we describe the results of the comparisons as well as expert evaluations in comparing the results of our deep learning technique with multi-atlas segmentation. We discuss the results of the comparisons in detail in Section IV.

A. Alternative Techniques

To demonstrate the advantages of our proposed method in fetal brain CP segmentation, we compared it with four alternative deep learning models, including two non-attentive models and two attentive models.

Non-attentive models:

- 3D UNet [37]: is based on the well-known U-Net architecture [51]. It has proven to be a powerful segmentation CNN that has achieved success in multiple segmentation tasks, e.g. [41], [52]–[55]. The 3D UNet leverages skip connections to strengthen the fusion of multi-level features to generate refined segmentation results.
- Plane Aggregated U-Net (PAUNet) [38]: segments the CP with a 2D UNet with dense blocks [56] from each plane and constructs a final prediction via plane aggregation. We used the mean of predictions for plane aggregation.

Attentive models:

- Attention UNet [39]: integrates a self-attention gating module into the U-Net model, which allows attention coefficients to be more specific to local regions.
- Squeeze & Excitation Fully Convolutional Network (SE-FCN) [40]: uses spatial and channel-wise squeeze and excitation attention to improve segmentation performance.

In the ablation studies, we investigated the impact of our proposed attentive module and deep supervision mechanisms. First, to demonstrate the advantages of our stage-wise attentive refinement module, we compared our framework with a model in which we removed the attention module and concatenated all feature maps from the backbone network with different resolutions to form the final prediction. We refer to this model as 3D Deeply Supervised Residual Network (DSRNet). Then, to investigate the efficacy of deep supervision, we trained our model with different supervision strategies. Specifically, we compared training with supervision at: 1) only the final network output, 2) both the final output and the multi-scale features, and 3) the attentive maps instead of attentive features in an approach similar to [57].

We also compared the results of our deep learning model with a state-of-the-art multi-atlas segmentation (MAS) technique as described in [25]. This technique uses diffeomorphic deformable registration [58] between each atlas and the query image to propagate labels from the atlas to the query image and then fuses the label maps using probabilistic STAPLE [59].

B. Evaluation Criteria

To compare segmentation results quantitatively on the test set, we used three evaluation metrics: 1) the Dice similarity coefficient (Dice), 2) the 95 percent Hausdorff Distance (95HD), and 3) the Average symmetric Surface Distance (ASD) [60]–[62]. The Dice is used to estimate the spatial overlap between the prediction and ground truth, 95HD is an outlier-robust measure based on the Hausdorff distance between two boundaries, and ASD is the average of all distances from points in the surface of the prediction to the ground truth.

MAS was used to obtain the very initial segmentations for our fetal MRI datasets that experts used as guidance in segmentations [25]. Even after significant time-consuming manual

segmentations and refinements, we were unable to ensure that the quantitative comparison of our method with MAS based on the overlap and distance metrics was unbiased. Therefore, for comparing our method with MAS, we performed blind expert evaluations. For this purpose, on a set of 15 fetal MRI scans from the test set, three experts independently compared segmentation results of our method and MAS. This was done in a randomized and blind fashion. Specifically, the experts visually compared the two segmentation results overlaid on the grayscale anatomical image of each fetus in different planes. They indicated if either of the segmentations A or B was better or if there were about the same quality.

C. Evaluation Results

As summarized in Table I, our method outperformed other methods based on almost all evaluation metrics. Our method achieved mean Dice of 0.87, 95HD of 0.96, and ASD of 0.28. The best value of each performance measure has been highlighted by boldface text in the table. According to these results our method generated better segmentations than both 3D UNet and DSRNet according to both Dice and ASD, and it outperformed PAUNet based on Dice. Our method showed an average of 5% and 0.12 mm improvement in Dice and 95HD over PAUNet, respectively. Moreover, our method achieved superior segmentation accuracy compared with the two competing attentive models (Attention U-Net and SE-FCN) in terms of Dice, 95HD, and ASD. Comparing the results of the DSRNet with those of the 3D UNet shows the improvement achieved by our backbone network (DSRNet) that exploited residual learning with deep supervision. Comparing our method with DSRNet shows the contribution of our attentive refinement to the improvement in performance.

In Table I, we have also reported the number of parameters and running time for all compared models. We observe that by adding 2.6% extra parameters, our attention module could improve the performance by 3% (from 0.84 by DSRNet to 0.87 by Ours) in terms of Dice. Although our method takes longer than the other methods in inference, all the test times are reasonable for the segmentation of post-acquisition reconstructed images, as these run times are an order of magnitude smaller than the time it takes to reconstruct images.

To investigate if the differences between methods were statistically significant, we conducted paired *t*-tests between the results of our method and those of 3D UNet, PAUNet, Attention UNet, SE-FCN and DSRNet. This test was carried out for all performance measures (Dice, 95HD, and ASD). In our paired *t*-tests, the significance level was set as 0.05. The *p*-values for the paired *t*-tests is summarized in Table II. Overall, the results of the comparisons and tests in Tables I and II indicate that our method performed best among all the techniques that we implemented and examined.

Table III shows the results of our experiments to investigate the effectiveness of our deep supervision strategy. In these experiments, we compared our deep supervision strategy with alternative training strategies that apply supervision at different network layers, as explained above. We can observe from Table III that our deep supervision strategy (backbone + attentive

TABLE I

COMPARING SEGMENTATION RESULTS OF DIFFERENT METHODS (3D UNET, PAUNET, ATTENTION UNET, SE-FCN, DSRNET, AND OUR METHOD) BASED ON QUANTITATIVE PERFORMANCE MEASURES (DICE, 95HD, AND ASD) AND COMPUTATIONAL COST (NUM OF PARAMS, AND RUN TIME). THE BEST VALUE FOR EACH EVALUATION MEASURE HAS BEEN SHOWN IN BOLDFACE TEXT.

Method	Dice \uparrow	95HD(mm) \downarrow	ASD(mm) \downarrow	Num of Params	Run time(s)
3D UNet	0.81 \pm 0.06	1.30 \pm 0.55	0.54 \pm 0.18	2.51M	5.70
PAUNet [38]	0.82 \pm 0.09	1.08 \pm 0.51	0.28\pm0.13	0.76M	5.03
Attention U-Net [39]	0.85 \pm 0.08	1.00 \pm 0.58	0.45 \pm 0.32	3.01M	8.63
SE-FCN [40]	0.82 \pm 0.07	1.31 \pm 0.82	0.45 \pm 0.26	2.59M	8.70
DSRNet	0.84 \pm 0.06	1.12 \pm 0.48	0.40 \pm 0.17	2.68M	12.89
Ours	0.87\pm0.06	0.96\pm0.38	0.28\pm0.14	2.75M	45.14

features + output) achieved the highest segmentation accuracy. The results suggest that supervision at the attentive features level (SAF) achieves slightly better segmentation accuracy than supervision at the attention map level (SAM). Feature maps at different channels contain different representations and require different attention maps. SAM imposes the same attention coefficients on feature maps of different channels, which is not as accurate as SAF for different channels.

TABLE II

p -VALUES OF PAIRED t -TESTS BETWEEN THE RESULTS OF EACH METHOD AND OUR METHOD FOR THE THREE SEGMENTATION PERFORMANCE MEASURES USED IN THIS STUDY (DICE, 95HD, AND ASD).

Metrics	Dice	95HD	ASD
3D U-Net vs. Ours	10^{-5}	10^{-3}	10^{-10}
PAUNet vs. Ours	10^{-3}	0.21	0.92
Attention UNet vs. Ours	0.12	0.70	10^{-3}
SE-FCN vs. Ours	10^{-3}	0.01	10^{-4}
DSRNet vs. Ours	0.02	0.09	10^{-3}

Figure 4 illustrates the feature maps before and after attentive refinement and the corresponding attentive map. The attentive maps highlight the regions that are important for the task (i.e cortical plate). Figure 5 shows axial slices of two representative cases with CP segmentation overlaid on the grayscale image using different methods, compared to the ground truth. In terms of the topology of CP segmentations, these results show that 1) the PAUNet and the UNet under- and over-segmented different regions of the CP depending on contrast and partial voluming; 2) the Attention U-Net under-segmented the CP; 3) the SE-FCN over-segmented the CP; and 4) our attentive model generated segmentations that were better than our DSRNet and were the most similar (among all models) to the ground truth. Visual comparison of the DSRNet results and the results of our method shows the extents of the differences in CP segmentation associated with the differences in quantitative measures reported in Tables I and II.

Figure 6 shows 3D surface error maps of the CP segmentation generated by each method with reference to the ground truth for a representative test subject. The average values of the performance metrics (Dice, 95HD, and ASD) for this case have also been shown for each method. This shows how the extents of the differences in these values correspond to 3D surface map and topological errors in CP segmentation. We can observe that our method achieved the most accurate segmentation results among all the compared methods. Our method generated topologically more accurate CP segmentations than the

other methods. In particular, there exist many holes in the segmentation results of UNet, PAUNet, DSRNet, and Attention U-Net in areas with significant partial voluming. This occurred due to the inability of those techniques in leveraging the useful information of multi-scale features. The PAUNet generated low surface errors in many areas, but completely failed in other areas. The SE-FCN generated high surface errors in many areas. Our DSRNet, overall, performed much better than the other competing methods and its performance significantly improved through deep attentive learning as can be seen in the results of our method in the lower right side of Fig. 6.

D. Expert Evaluation Results

Table IV shows the expert evaluation results of CP segmentation using our method against MAS (multi-atlas segmentation) for 15 test subjects. We separated the test subjects to 9 and 6 based on the original reconstructed image quality. This was done because we expected the segmentation to be more challenging in low-quality reconstructions due to the effects of residual motion artifacts and partial voluming. The last row of the table shows the results of majority voting among the three experts, where a three-way tie went to equal. This evaluation shows that our method was chosen most often as the more accurate segmentation method by each expert and also based on majority voting. Our method outperformed MAS regardless of the quality of the input reconstructed image, which shows its robust and accurate performance. It should also be noted that our method generated the CP segmentation for every test subject in less than 1 minute. MAS, on the other hand, took approximately 20 minutes because of the need for deformable registration between each atlas and the subject.

IV. DISCUSSION

In this paper we presented a deep attentive convolutional neural network for cortical plate segmentation on fetal MRI. Fetal cortical plate segmentation in MRI is very challenging due to the rapid changes and variations in the microstructure and shape of the transient fetal brain compartments. The relatively low spatial resolution of fetal MRI compared to the thin structure of the CP makes the task more challenging. Due to the large variations in CP structure and severe partial voluming effects, automatic segmentation methods struggle to extract multi-scale contextual information from images. To alleviate those issues, our proposed network is equipped with stage-wise attentive refinement modules with mixed kernel convolutions

TABLE III
SUPERVISED LEARNING BASED ON DIFFERENT FEATURES.

Supervision				Dice \uparrow	95HD(mm) \downarrow	ASD(mm) \downarrow
Backbone	Attentive Feature	Attention Map	Output			
			✓	0.85 \pm 0.06	1.14 \pm 0.57	0.49 \pm 0.18
✓			✓	0.86 \pm 0.06	1.12 \pm 0.43	0.46 \pm 0.38
✓		✓	✓	0.86 \pm 0.05	1.07 \pm 0.34	0.32 \pm 0.12
✓	✓		✓	0.87\pm0.06	0.96\pm0.38	0.28\pm0.14

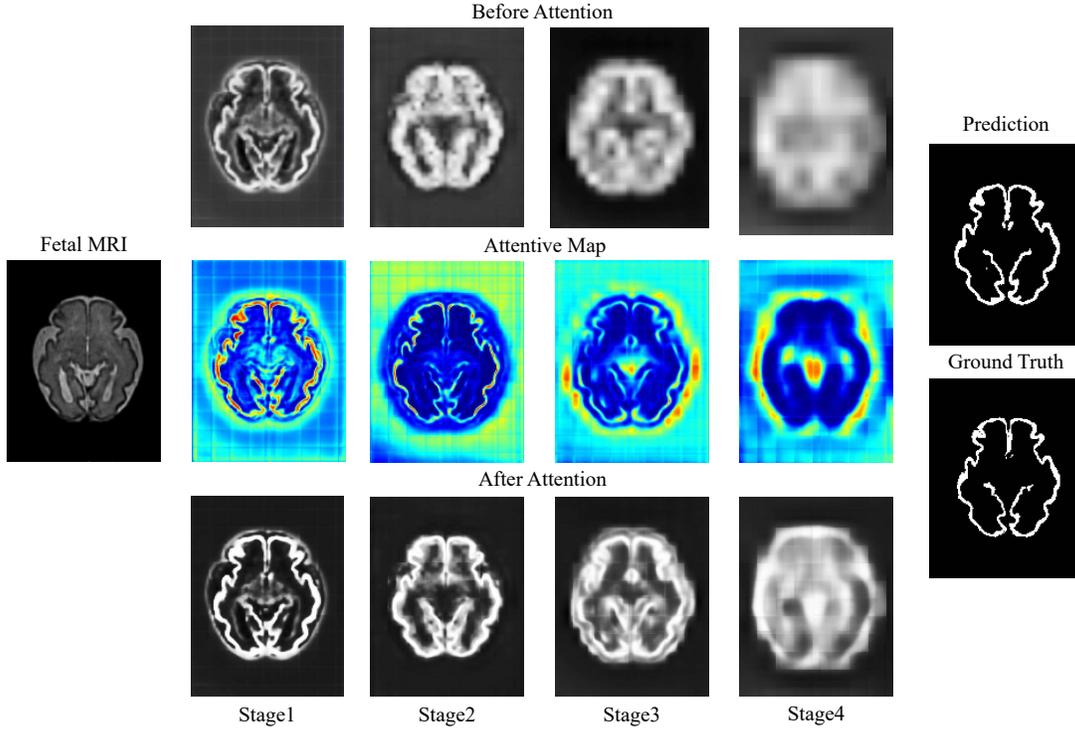


Fig. 4. The illustration of the feature maps in the attention module. The first column shows an original 3D reconstructed fetal brain MRI image. The second to the fifth columns show the feature maps in the attention module from stage1 to stage4. The first row to the third row shows the input features, the attentive maps, and the output features of the attention module, respectively. The third column shows the prediction of the network and the ground truth.

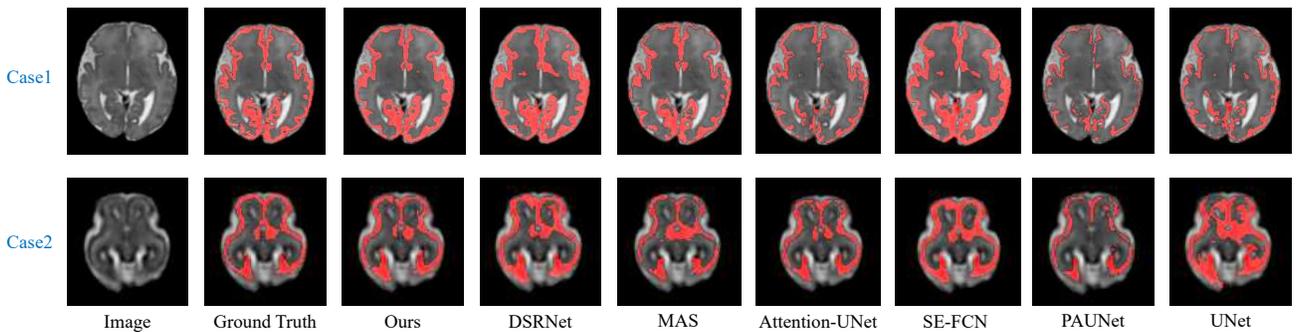


Fig. 5. Visual comparison of the segmentation results obtained from different methods (ours, DSRNet, MAS, PAUNet, and 3D UNet) compared to the ground truth for two representative test cases. As supported by the results in Table I, our method generated segmentations that were most similar to the ground truth.

to capture multi-scale information. Besides, the patch-wise input provides sufficient training data in comparison to the image-wise input, which alleviates the risk of over-fitting [63].

The 3D UNet and PAUNet are powerful models. However, they do not have specific multi-scale architecture design and only use single scale kernels to generate feature maps. As shown in Figure 5 and 6, the subtle and highly variable

boundaries of the fetal CP were much more effectively learned and predicted through the multi-scale attentive modules that exploited different kernel widths in our network. Compared with the 3D UNet and PAUNet, our ablation study method (DSRNet) achieved better segmentation of the cortical regions by concatenating multi-scale feature maps. However DSRNet still lost some details in the subtle boundary of the fetal CP.

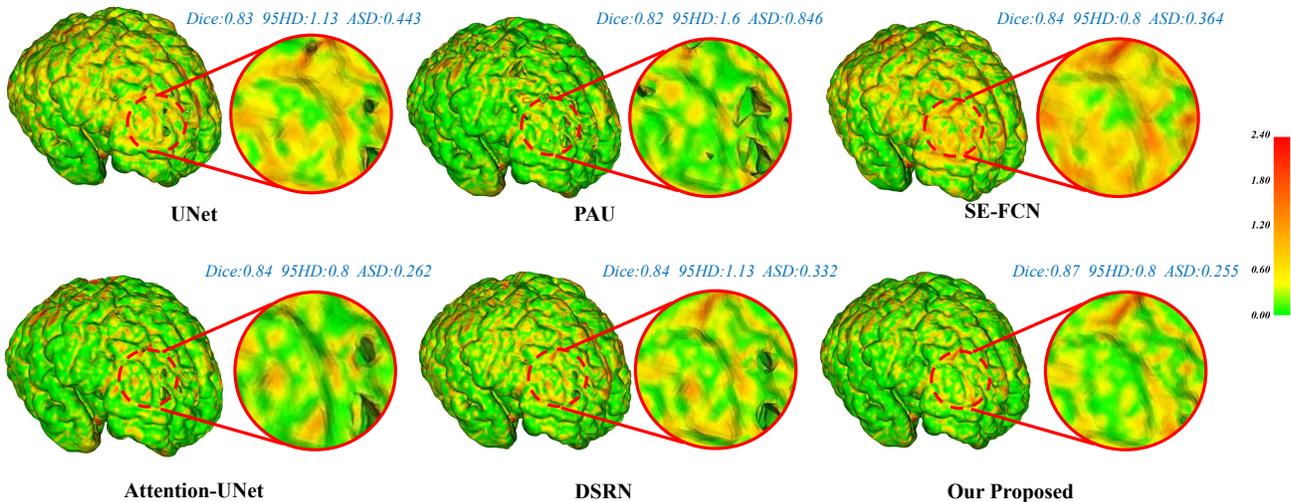


Fig. 6. Visualization of the cortical surface error maps on a representative test case along with the values of the performance measures for different methods (3D UNet, PAUnet, SE-FCN, Attention U-Net, DSRNet, and our method). Our method generated the lowest average error, whereas other methods generated holes in the CP surface.

TABLE IV
EXPERT EVALUATION RESULTS OF COMPARING CP SEGMENTATIONS USING OUR METHOD AGAINST MULTI-ATLAS SEGMENTATION (MAS).

Voter	Good Reconstruction			Bad Reconstruction		
	Ours	MAS	Equal	Ours	MAS	Equal
Expert1	9/9	0/9	0/9	6/6	0/6	0/6
Expert2	5/9	1/9	3/9	3/6	1/6	2/6
Expert3	9/9	0/9	0/9	4/6	2/6	0/6
Majority	9/9	0/9	0/9	4/6	0/6	2/6

Our method exploited group convolutional blocks with multiple kernel sizes to extract multi-scale information and generate attentive maps that refined the feature maps at different scales. Using the residual spatial attention mechanism, our attentive network efficiently and effectively aggregated complementary multi-scale information for accurate representation of the fetal CP regions. Given the multi-scale information selected by the stage-wise attentive refinement module, features of highly-variable anatomy were captured accurately and represented effectively during the learning process.

Prior studies have proposed a number of different attention mechanisms for such applications as image classification [64] and image segmentation [65]. Those attention mechanisms have different forms, such as spatial attention [66], channel-wise attention [64], and self-attention [67]. However, none of the existing attention mechanisms was suitable for our application. Specifically, channel-wise attention lacks the spatial attention, and self-attention is not suitable for 3D images due to its excessive memory requirements. Our comparisons also showed that our proposed model was superior to prior attentive models for medical image segmentation [39], [40]. We attribute this to the superior ability of our model to learn multi-scale representations of the cortical plate compared to those other models. Our spatial attention mechanism effectively refined the feature maps from input 3D image patches at different scales to improve segmentation performance.

The idea of using convolutional operations with multiple kernel sizes has been promoted by various network architectures, including Inception [68], MixNet [69], and PSPNet [70]. All of these architectures extract multi-scale feature maps using multi-branch convolutional operations with different kernel sizes. These convolutional operations are interlaced with pooling operations to increase the richness of the learned representations. Our attention module is based on similar principles. It consists of multiple convolutional blocks with different kernel sizes to generate multi-scale attention maps.

We used 3D convolutional kernels to build our 3D CNN. 3D CNNs are demanding in terms of GPU memory usage. A common practice to reduce memory requirements in 3D CNNs is to use large 3D patches as inputs. However, if the patches are too small compared to the image structures, models may not learn global context information. In our setting, we were able to use a large input patch size of $64 \times 64 \times 64$ to fit our model on a GPU with 11GB of memory. This patch size covered about $50 \times 50 \times 50 \text{ mm}^3$ physical volume size, which was sufficiently large compared to the size of the fetal brain.

Atlas-based segmentation techniques are also challenged by the same issues that make learning-based segmentation of the fetal CP challenging. By using age-matched atlases to segment a test subject, these techniques may handle the variability in cortical folding and maturation levels relatively well. However, they are more sensitive to artifacts and partial voluming effects. These techniques rely on deformable registration between each atlas and the test image to propagate the CP label from the atlas to the test subject image. The registration may fail or may be inaccurate in thin and folded areas of the CP where partial voluming disrupts the normal range of the intensity values and features that are used to inform deformable registration. In MAS, registration errors translate into segmentation errors. Another issue with MAS methods is their computational cost and time, which is high due to the required optimization-based deformable registration between multiple atlases and the subject at test time. Once

our model is trained, it generates CP segmentation of a test case in less than 1 minute which is an order of magnitude faster than MAS. More importantly, our method generated more accurate segmentations than MAS according to expert evaluations (Table IV). We attribute this to the superior ability of our model in capturing multi-scale information and learning the context from training samples.

Cortical plate segmentation on reconstructed fetal MRI is a fundamental first step in quantitative analysis of fetal cortical maturation. To evaluate the generalization ability of our model with respect to the fetal MRI image reconstruction algorithms, we reconstructed five fetal MRI scans with NiftyMIC [24], and tested them with our model which was trained with images reconstructed using the algorithm by Kainz et al. [22]. The results have been shown in the Supplementary Figure S1 and Table S1. For this comparison, because our ground truth segmentations were made on images reconstructed using [22], we had to register and resample the NiftyMIC reconstructions to the reconstructions by [22] for each subject. We note that resampling degrades image quality. Moreover, any registration errors would lead to reduced Dice and increased ASD computations due to misalignment between the ground truth and the test CP segmentations. Despite these factors, the results of this experiment show that our model performed well in the new (NiftyMIC reconstruction) domain.

V. CONCLUSION

In this study, we developed a deep attentive neural network with mixed kernel convolutions for automatic cortical plate segmentation in fetal MRI. The key feature of our technique is to leverage a residual spatial attention mechanism to capture multi-scale information. We used two group convolutional blocks with mixed kernels to generate attentive maps to adaptively select and learn important information through the attention mechanism in a deep learning model that leveraged multi-stage deep supervision in both feature extraction and attention modules. Through this novel architecture and the attention mechanism we developed the first 3D fully convolutional deep neural network model for automatic cortical plate segmentation in fetal MRI. Our model outperformed several state-of-the-art deep learning models as well as a multi-atlas segmentation method.

REFERENCES

- [1] F. V. Coakley, O. A. Glenn, A. Qayyum, A. J. Barkovich, R. Goldstein, and R. A. Filly, "Fetal MRI: a developing technique for the developing patient," *American Journal of Roentgenology*, vol. 182, no. 1, pp. 243–252, 2004.
- [2] A. Gholipour, J. A. Estroff, C. E. Barnewolt, R. L. Robertson, P. E. Grant, B. Gagoski et al., "Fetal MRI: a technical update with educational aspirations," *Concepts in Magnetic Resonance Part A*, vol. 43, no. 6, pp. 237–266, 2014.
- [3] C. Clouchoux, D. Kudelski, A. Gholipour, S. K. Warfield, S. Viseur, M. Bouyssi-Kobar et al., "Quantitative in vivo MRI measurement of cortical development in the fetus," *Brain Structure and Function*, vol. 217, no. 1, pp. 127–139, 2012.
- [4] V. Rajagopalan, J. Scott, P. A. Habas, K. Kim, J. Corbett-Detig, F. Rousseau et al., "Local tissue growth patterns underlying normal fetal human brain gyrfication quantified in utero," *Journal of neuroscience*, vol. 31, no. 8, pp. 2878–2887, 2011.
- [5] J. A. Scott, P. A. Habas, K. Kim, V. Rajagopalan, K. S. Hamzelou, J. M. Corbett-Detig 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.
- [6] J. Corbett-Detig, P. Habas, J. A. Scott, K. Kim, V. Rajagopalan, P. McQuillen et al., "3d global and regional patterns of human fetal subplate growth determined in utero," *Brain Structure and Function*, vol. 215, no. 3–4, pp. 255–263, 2011.
- [7] R. Wright, V. Kyriakopoulou, C. Ledig, M. A. Rutherford, J. V. Hajnal, D. Rueckert et al., "Automatic quantification of normal cortical folding patterns from fetal brain MRI," *NeuroImage*, vol. 91, pp. 21–32, 2014.
- [8] L. Vasung, C. K. Rollins, H. J. Yun, C. Velasco-Annis, J. Zhang, K. Wagstyl et al., "Quantitative in vivo MRI assessment of structural asymmetries and sexual dimorphism of transient fetal compartments in the human brain," *Cerebral Cortex*, vol. 30, no. 3, pp. 1752–1767, 2020.
- [9] L. Vasung, C. K. Rollins, C. Velasco-Annis, H. J. Yun, J. Zhang, S. K. Warfield et al., "Spatiotemporal differences in the regional cortical plate and subplate volume growth during fetal development," *Cerebral Cortex*, 2020.
- [10] C. D. Kroenke and P. V. Bayly, "How forces fold the cerebral cortex," *Journal of Neuroscience*, vol. 38, no. 4, pp. 767–775, 2018.
- [11] X. Wang, C. Studholme, P. L. Grigsby, A. E. Frias, V. C. C. Carlson, and C. D. Kroenke, "Folding, but not surface area expansion, is associated with cellular morphological maturation in the fetal cerebral cortex," *Journal of Neuroscience*, vol. 37, no. 8, pp. 1971–1983, 2017.
- [12] C. Llinares-Benadero and V. Borrell, "Deconstructing cortical folding: genetic, cellular and mechanical determinants," *Nature Reviews Neuroscience*, vol. 20, no. 3, pp. 161–176, 2019.
- [13] S. Rana, R. Shishegar, S. Quezada, L. Johnston, D. W. Walker, and M. Tolcos, "The subplate: a potential driver of cortical folding?" *Cerebral Cortex*, vol. 29, no. 11, pp. 4697–4708, 2019.
- [14] C. Clouchoux, A. Du Plessis, M. Bouyssi-Kobar, W. Tworetzky, D. McElhinney, D. Brown et al., "Delayed cortical development in fetuses with complex congenital heart disease," *Cerebral cortex*, vol. 23, no. 12, pp. 2932–2943, 2013.
- [15] K. Im, A. Guimaraes, Y. Kim, E. Cottrill, B. Gagoski, C. Rollins et al., "Quantitative folding pattern analysis of early primary sulci in human fetuses with brain abnormalities," *American Journal of Neuroradiology*, vol. 38, no. 7, pp. 1449–1455, 2017.
- [16] T. Tarui, N. Madan, N. Farhat, R. Kitano, A. Ceren Tanritanir, G. Graham et al., "Disorganized patterns of sulcal position in fetal brains with agenesis of corpus callosum," *Cerebral Cortex*, vol. 28, no. 9, pp. 3192–3203, 2018.
- [17] O. M. Benkarim, N. Hahner, G. Piella, E. Gratacos, M. A. G. Ballester, E. Eixarch et al., "Cortical folding alterations in fetuses with isolated non-severe ventriculomegaly," *NeuroImage: Clinical*, vol. 18, pp. 103–114, 2018.
- [18] C. M. Ortinau, C. K. Rollins, A. Gholipour, H. J. Yun, M. Marshall, B. Gagoski et al., "Early-emerging sulcal patterns are atypical in fetuses with congenital heart disease," *Cerebral Cortex*, vol. 29, no. 8, pp. 3605–3616, 2019.
- [19] 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.
- [20] C. Studholme, "Mapping fetal brain development in utero using magnetic resonance imaging: the big bang of brain mapping," *Annual review of biomedical engineering*, vol. 13, pp. 345–368, 2011.
- [21] 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.
- [22] B. Kainz, M. Steinberger, W. Wein, M. Kuklisova-Murgasova, C. Malamatiniou, K. Keradren 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.
- [23] A. Alansary, M. Rajchl, S. G. McDonagh, M. Murgasova, M. Damodaram, D. F. Lloyd et al., "Pvr: Patch-to-volume reconstruction for large area motion correction of fetal MRI," *IEEE transactions on medical imaging*, vol. 36, no. 10, pp. 2031–2044, 2017.
- [24] M. Ebner, G. Wang, W. Li, M. Aertsen, P. A. Patel, R. Aghwane et al., "An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI," *NeuroImage*, vol. 206, p. 116324, 2020.
- [25] A. Gholipour, C. K. Rollins, C. Velasco-Annis, A. Oualam, A. Akhondi-Asl, O. Afacan et al., "A normative spatiotemporal MRI

- atlas of the fetal brain for automatic segmentation and analysis of early brain growth," *Scientific reports*, vol. 7, no. 1, p. 476, 2017.
- [26] P. A. Habas, K. Kim, F. Rousseau, O. A. Glenn, A. J. Barkovich, and C. Studholme, "Atlas-based segmentation of developing tissues in the human brain with quantitative validation in young fetuses," *Human brain mapping*, vol. 31, no. 9, pp. 1348–1358, 2010.
- [27] A. Makropoulos, S. J. Counsell, and D. Rueckert, "A review on automatic fetal and neonatal brain MRI segmentation," *NeuroImage*, vol. 170, pp. 231–248, 2018.
- [28] H. Xue, L. Srinivasan, S. Jiang, M. Rutherford, A. D. Edwards, D. Rueckert *et al.*, "Automatic segmentation and reconstruction of the cortex from neonatal MRI," *Neuroimage*, vol. 38, no. 3, pp. 461–477, 2007.
- [29] N. I. Weisenfeld and S. K. Warfield, "Automatic segmentation of newborn brain MRI," *Neuroimage*, vol. 47, no. 2, pp. 564–572, 2009.
- [30] P. A. Habas, K. Kim, F. Rousseau, O. A. Glenn, A. J. Barkovich, and C. Studholme, "Atlas-based segmentation of the germinal matrix from in utero clinical MRI of the fetal brain," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2008, pp. 351–358.
- [31] P. A. Habas, K. Kim, D. Chandramohan, F. Rousseau, O. A. Glenn, and C. Studholme, "Statistical model of laminar structure for atlas-based segmentation of the fetal brain from in utero mr images," in *Medical Imaging 2009: Image Processing*, vol. 7259. International Society for Optics and Photonics, 2009, p. 725917.
- [32] P. A. Habas, K. Kim, J. M. Corbett-Detig, F. Rousseau, O. A. Glenn, A. J. Barkovich *et al.*, "A spatiotemporal atlas of mr intensity, tissue probability and shape of the fetal brain with application to segmentation," *Neuroimage*, vol. 53, no. 2, pp. 460–470, 2010.
- [33] A. Serag, V. Kyriakopoulou, M. Rutherford, A. Edwards, J. Hajnal, P. Aljabar *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.
- [34] 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.
- [35] A. Gholipour, C. Limperopoulos, S. Clancy, C. Clouchoux, A. Akhondi-Asl, J. A. Estroff *et al.*, "Construction of a deformable spatiotemporal MRI atlas of the fetal brain: evaluation of similarity metrics and deformation models," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2014, pp. 292–299.
- [36] N. Khalili, N. Lessmann, E. Turk, N. Claessens, R. de Heus, T. Kolk *et al.*, "Automatic brain tissue segmentation in fetal MRI using convolutional neural networks," *Magnetic resonance imaging*, vol. 64, pp. 77–89, 2019.
- [37] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3d u-net: learning dense volumetric segmentation from sparse annotation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2016, pp. 424–432.
- [38] J. Hong, H. J. Yun, G. Park, Y.-H. Park, J.-M. Lee, and K. Im, "Fetal cortical plate segmentation using 2d u-net with plane aggregation."
- [39] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa *et al.*, "Attention u-net: Learning where to look for the pancreas," *arXiv preprint arXiv:1804.03999*, 2018.
- [40] A. G. Roy, N. Navab, and C. Wachinger, "Recalibrating fully convolutional networks with spatial and channel "squeeze and excitation" blocks," *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 540–549, 2018.
- [41] S. S. M. Salehi, S. R. Hashemi, C. Velasco-Annis, A. Ouaalam, J. A. Estroff, D. Erdogmus *et al.*, "Real-time automatic fetal brain extraction in fetal MRI by deep learning," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 2018, pp. 720–724.
- [42] S. S. M. Salehi, D. Erdogmus, and A. Gholipour, "Auto-context convolutional neural network (auto-net) for brain extraction in magnetic resonance imaging," *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2319–2330, 2017.
- [43] P. A. Yushkevich, A. Pashchinskiy, I. Oguz, S. Mohan, J. E. Schmitt, J. M. Stein *et al.*, "User-guided segmentation of multi-modality medical imaging datasets with itk-snap," *Neuroinformatics*, vol. 17, no. 1, pp. 83–102, 2019.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [45] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [46] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [47] S. Xie and Z. Tu, "Holistically-nested edge detection," in *Proceedings of the IEEE Int. conference on computer vision*, 2015, pp. 1395–1403.
- [48] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1492–1500.
- [49] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [50] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [51] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [52] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016, pp. 565–571.
- [53] S. S. M. Salehi, D. Erdogmus, and A. Gholipour, "Tversky loss function for image segmentation using 3d fully convolutional deep networks," in *International Workshop on Machine Learning in Medical Imaging*. Springer, 2017, pp. 379–387.
- [54] G. Du, X. Cao, J. Liang, X. Chen, and Y. Zhan, "Medical image segmentation based on u-net: A review," *Journal of Imaging Science and Technology*, 2020.
- [55] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: A survey," *arXiv preprint arXiv:2001.05566*, 2020.
- [56] S. Jégou, M. Drozdal, D. Vazquez, A. Romero, and Y. Bengio, "The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 11–19.
- [57] G. Wang, J. Shapey, W. Li, R. Dorent, A. Demetriadias, S. Bisdas *et al.*, "Automatic segmentation of vestibular schwannoma from t2-weighted mri by deep spatial attention with hardness-weighted loss," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 264–272.
- [58] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, "A reproducible evaluation of ants similarity metric performance in brain image registration," *Neuroimage*, vol. 54, no. 3, pp. 2033–2044, 2011.
- [59] A. Akhondi-Asl and S. K. Warfield, "Simultaneous truth and performance level estimation through fusion of probabilistic segmentations," *IEEE transactions on medical imaging*, vol. 32, no. 10, pp. 1840–1852, 2013.
- [60] H.-H. Chang, A. H. Zhuang, D. J. Valentino, and W.-C. Chu, "Performance measure characterization for evaluating neuroimage segmentation algorithms," *Neuroimage*, vol. 47, no. 1, pp. 122–135, 2009.
- [61] A. M. Mendrik, K. L. Vincken, H. J. Kuijf, M. Breeuwer, W. H. Bouvy, J. De Bresser *et al.*, "Mrbrains challenge: online evaluation framework for brain image segmentation in 3t MRI scans," *Computational intelligence and neuroscience*, vol. 2015, p. 1, 2015.
- [62] V. Yeghiazaryan and I. Voiculescu, "An overview of current evaluation methods used in medical image segmentation," *Tech. Rep. RR-15-08, Department of Computer Science*, 2015.
- [63] M. H. Hesamian, W. Jia, X. He, and P. Kennedy, "Deep learning techniques for medical image segmentation: Achievements and challenges," *Journal of digital imaging*, vol. 32, no. 4, pp. 582–596, 2019.
- [64] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [65] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille, "Attention to scale: Scale-aware semantic image segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3640–3649.
- [66] S. Woo, J. Park, J.-Y. Lee, and I. So Kweon, "Cbam: Convolutional block attention module," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 3–19.

- [67] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-attention generative adversarial networks," *arXiv preprint arXiv:1805.08318*, 2018.
- [68] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov *et al.*, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [69] M. Tan and Q. V. Le, "Mixconv: Mixed depthwise convolutional kernels," *CoRR*, *abs/1907.09595*, 2019.
- [70] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.

Supplementary Material

A Deep Attentive Convolutional Neural Network for Automatic Cortical Plate Segmentation in Fetal MRI

Haoran Dou, Davood Karimi, Caitlin K. Rollins, Cynthia M. Ortinau, Lana Vasung, Clemente Velasco-Annis, Abdelhakim Oualam, Xin Yang, Dong Ni, Ali Gholipour

I. EFFECT OF DOMAIN SHIFT DUE TO DIFFERENT RECONSTRUCTION ALGORITHMS ON SEGMENTATION.

We investigated the effect of domain shift due to different reconstruction algorithms on the performance of our segmentation model. Given our model trained on fetal MRI images reconstructed with the method by Kainz et al. [22], we evaluated the performance of our model on images reconstructed with 1) Kainz et al.’s method [22], and 2) the *NiftyMIC* method [24]. Images from five subjects were included in this experiment. Because our ground truth (manual) segmentations were in the original domain (i.e., reconstructions with [22]), we had to register *NiftyMIC* reconstructions to Kainz et al.’s reconstructions and resample them. The segmentation results with the two reconstruction methods are shown in Fig. S1 for two representative test cases. Despite the domain shift, as shown in Table S1, our model performed very well, achieving average Dice of 0.838 and ASD of 0.338 mm. This experiment shows that our model performed robustly in the presence of domain shift due to the use of different reconstruction algorithms.

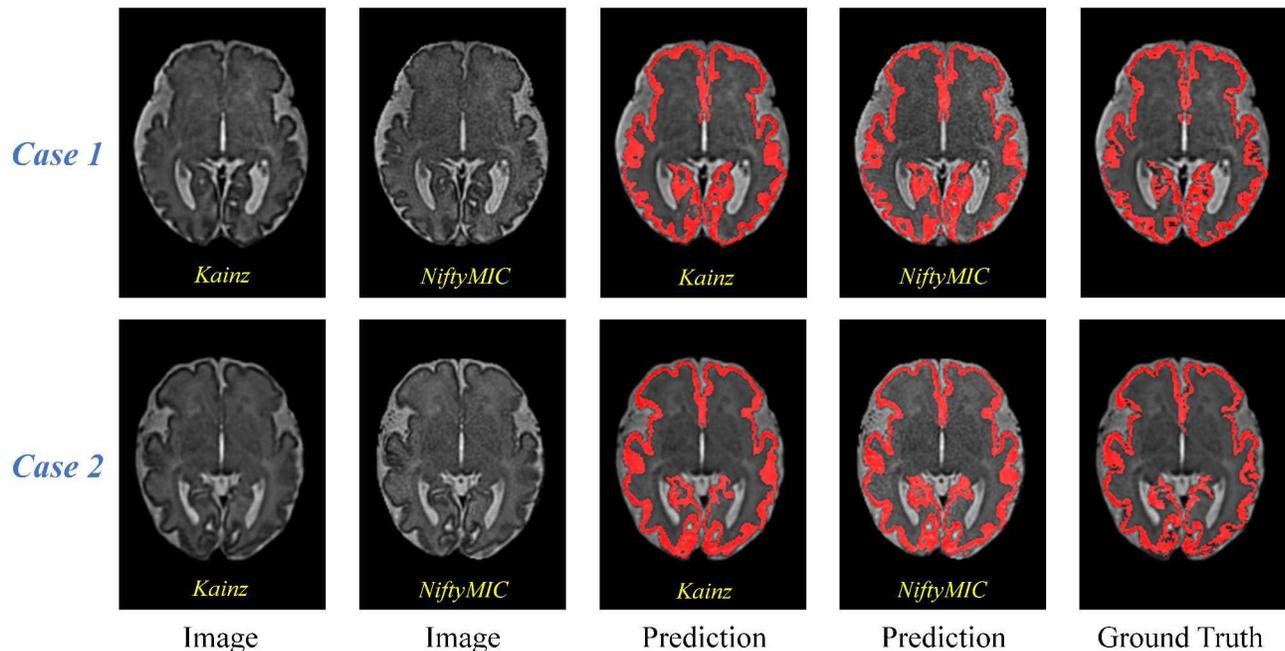


Fig. S1. Visualization of the segmentations obtained from our model trained and tested with ground truth data on reconstructions from [22] (Kainz et al.) on reconstructions obtained from *NiftyMIC*, for two representative test cases. Since the ground truth was in the original domain (i.e., on Kainz reconstructions), we registered and resampled the *NiftyMIC* reconstructions to Kainz reconstructions prior to segmenting them.

TABLE S1

SEGMENTATION ACCURACY OF THE PROPOSED DEEP LEARNING MODEL WHEN THE DOMAIN WAS SHIFTED DUE TO THE USE OF A DIFFERENT RECONSTRUCTION ALGORITHM: NIFTYMIC WAS USED TO RECONSTRUCT IMAGES AS OPPOSED TO THE ALGORITHM BY KAINZ ET AL. [22] WHICH WAS USED IN THE TRAINING DATA. BECAUSE THE GROUND TRUTH SEGMENTATIONS WERE IN THE ORIGINAL DOMAIN (KAINZ ET AL. RECONSTRUCTIONS), WE HAD TO REGISTER AND RESAMPLE THE NIFTYMIC RECONSTRUCTIONS TO THE KAINZ RECONSTRUCTIONS. WE NOTE THAT ANY INACCURACY IN REGISTRATION RESULTS IN REDUCED DICE AND INCREASED ASD FOR THE NIFTYMIC METHOD. DESPITE THESE ISSUES, THE RESULTS SHOW THAT OUR MODEL GENERALIZED WELL TO NIFTYMIC RECONSTRUCTIONS.

Reconstruction method	Dice \uparrow	ASD(mm) \downarrow
Kainz	0.887 \pm 0.002	0.226 \pm 0.008
NiftyMIC	0.838 \pm 0.012	0.338 \pm 0.027