



Original contribution

Automatic brain tissue segmentation in fetal MRI using convolutional neural networks

N. Khalili^{a,*}, N. Lessmann^a, E. Turk^{b,c}, N. Claessens^{b,c}, R. de Heus^d, T. Kolk^b, M.A. Viergever^{a,c,e}, M.J.N.L. Benders^{b,c}, I. Išgum^{a,c}

^a Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands

^b Department of Neonatology, Wilhelmina Children's Hospital, University Medical Center Utrecht, Utrecht, The Netherlands

^c Brain Center Rudolf Magnus, University Medical Center Utrecht, Utrecht, The Netherlands

^d Department of Obstetrics, University Medical Center Utrecht, The Netherlands

^e Utrecht University, Utrecht, The Netherlands

ARTICLE INFO

Keywords:

Fetal MRI
Brain segmentation
Intensity inhomogeneity
Deep learning
Convolutional neural network

ABSTRACT

MR images of fetuses allow clinicians to detect brain abnormalities in an early stage of development. The cornerstone of volumetric and morphologic analysis in fetal MRI is segmentation of the fetal brain into different tissue classes. Manual segmentation is cumbersome and time consuming, hence automatic segmentation could substantially simplify the procedure. However, automatic brain tissue segmentation in these scans is challenging owing to artifacts including intensity inhomogeneity, caused in particular by spontaneous fetal movements during the scan. Unlike methods that estimate the bias field to remove intensity inhomogeneity as a preprocessing step to segmentation, we propose to perform segmentation using a convolutional neural network that exploits images with synthetically introduced intensity inhomogeneity as data augmentation. The method first uses a CNN to extract the intracranial volume. Thereafter, another CNN with the same architecture is employed to segment the extracted volume into seven brain tissue classes: cerebellum, basal ganglia and thalami, ventricular cerebrospinal fluid, white matter, brain stem, cortical gray matter and extracerebral cerebrospinal fluid. To make the method applicable to slices showing intensity inhomogeneity artifacts, the training data was augmented by applying a combination of linear gradients with random offsets and orientations to image slices without artifacts. To evaluate the performance of the method, Dice coefficient (DC) and Mean surface distance (MSD) per tissue class were computed between automatic and manual expert annotations. When the training data was enriched by simulated intensity inhomogeneity artifacts, the average achieved DC over all tissue classes and images increased from 0.77 to 0.88, and MSD decreased from 0.78 mm to 0.37 mm. These results demonstrate that the proposed approach can potentially replace or complement preprocessing steps, such as bias field corrections, and thereby improve the segmentation performance.

1. Introduction

Important neurodevelopmental changes occur in the last trimester of pregnancy, i.e., between 30 and 40 weeks of gestation, including volumetric growth, myelination and cortical gyrification [1–3]. Magnetic resonance imaging (MRI) is widely used to non-invasively assess and monitor the developmental status of the fetal brain *in utero* [4,5]. The cornerstone of volumetric and morphologic analysis in fetal MRI is the segmentation of the fetal brain into different tissue classes, such as white and gray matter. Performing this segmentation manually, however, is extremely time-consuming and requires a high level of expertise. The reasons are not only the complex convoluted shapes of the

different tissues, but also the limited image quality due to imaging artifacts. Fetal MR imaging is particularly challenging in this regard because the receiver coils can only be positioned on the maternal body and not closer to the anatomy of interest. Furthermore, movements of the fetus relative to the mother can only to some extent be controlled and predicted. Especially, fetal motion therefore negatively affects the image quality and causes artifacts such as intensity inhomogeneity (Fig. 1).

Because manual annotation is very time consuming and additionally hampered by these artifacts, a reliable automatic tissue segmentation tool would provide a valuable alternative, especially if it could give detailed fetal brain tissue segmentations in the presence of artifacts. To

* Corresponding author.

E-mail address: n.khalili@umcutrecht.nl (N. Khalili).

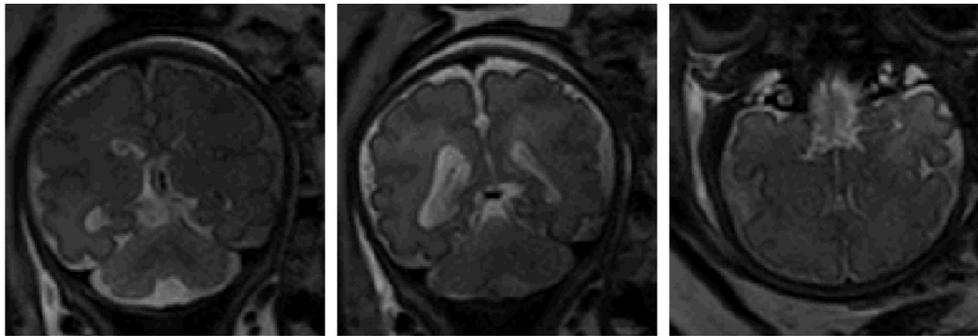


Fig. 1. Examples of slices with intensity inhomogeneity in T2 weighted MRI of the fetus cropped to the brain. Note that the full slices also visualize a larger part of the fetus as well as the maternal body.

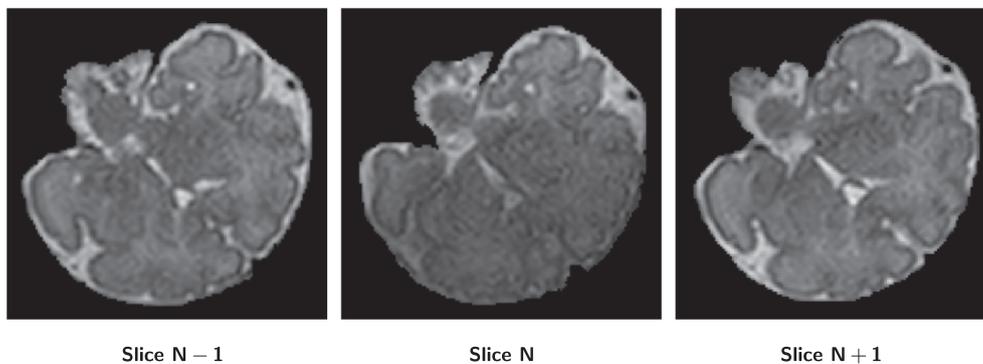


Fig. 2. Example of a fetal T2 weighted MRI with intensity inhomogeneity (middle) and the slice before (left) and after (right) from the same scan, both without artifacts. Structures outside the fetal cranium have been masked out.

cope with imaging artifacts, previous approaches in the literature performed the segmentation in images reconstructed from multiple 2D acquisitions. Most fetal MRI scans are acquired in 2D using single-shot fast spin-echo (SSFSE) sequences [6]. Artifacts such as intensity inhomogeneity may therefore appear only in some slices, e.g., due to movements during acquisition of these slices, but do not have to be present in their immediate neighboring slices as well (Fig. 2). Volumetric reconstruction approaches are typically based on the acquisition of several stacks of 2D slices in axial, sagittal and coronal orientation. These stacks are registered to a common coordinate space so that they can be combined into a single reconstructed 3D volume, thus removing artifacts that affect only some slices and inter-slice inconsistencies [7–11].

For automatic segmentation of fetal brain tissue in reconstructed MR volumes, Habas et al. [12] proposed a method using an atlas-based expectation maximization (EM) model to segment white matter (WM), gray matter, germinal matrix, and extracerebral cerebrospinal fluid (eCSF). Prior to performing the segmentation, another EM model was used for bias field correction. Gholipour et al. [13] proposed a method for segmentation of the ventricles in fetal MRI. As a preprocessing step, in addition to using volumetric reconstructions, intensity inhomogeneity was corrected using the non-parametric entropy maximization method [14]. Initial segmentation was obtained with the use of STAPLE [15], then the final segmentation was derived with a probabilistic shape model that incorporates intensity and local spatial information. Serag et al. [16] proposed an atlas-based brain segmentation method for both neonatal and fetal MRI. Fetal scans were reconstructed into a single 3D brain volume using the slice-to-volume reconstruction method described in [7] and intensity inhomogeneity was removed using the N4 algorithm [17]. Thereafter, the fetal brain scans were segmented into cortex, ventricles and hemispheres.

Deep learning methods have recently been very successful and have often outperformed traditional machine learning and model-based

methods in medical image analysis [18] including brain MRI [19,20]. A major strength of these networks is their ability to extract the features relevant for the tasks directly from the data. There is no need anymore to first derive a set of handcrafted features from the image as input to a classifier or model, the networks rather learn themselves to extract and interpret features relevant to the segmentation task. Therefore, deep learning methods often achieve a better performance than traditional machine learning methods with hand-crafted features. However, CNNs usually require large sets of diverse training data. To enlarge the size of training set and to ensure robustness to expected variability in the data, some studies use data augmentation techniques such as random rotation, random translation and random noise injection [21,22]. We therefore hypothesize that, while artifacts such as intensity inhomogeneity are challenging for traditional approaches and therefore normally require preprocessing of the images, CNNs may be able to adapt and become invariant to such artifacts if they are presented enough examples during training. However, manual segmentation of slices with intensity inhomogeneity is much more cumbersome than segmentation of artifact free slices so that a sizable training database is difficult to obtain. We therefore propose to tackle one of the most common artifacts in fetal MRI, namely intensity inhomogeneity, by randomly adding synthetic intensity inhomogeneity to slices for which a corresponding reference segmentation is available. By only mutating the intensity values but not the orientation or shape of structures in the image, the same reference segmentation can be used as ground truth. This tailored data augmentation strategy affects network training only. At inference time, in contrast to previous methods, no complex preprocessing of the image is required.

Furthermore, previous methods focused on segmenting the brain into the three main tissue classes: WM, cortical gray matter and ventricles. However, characteristics of other tissue classes, such as cerebellum (CB) and brain stem (BS), are important to understand and predict healthy or aberrant brain development in preterm infants of

similar gestational age as fetuses [23]. The cerebellum is particularly of clinical interest as it is one of the fastest growing brain regions during the last trimester of pregnancy [24].

Another challenge for segmentation of the fetal brain in MRI is the large field of view of these scans. Since the fetus is scanned *in utero*, the images also visualize parts of the maternal and the fetal body, and not only the head of the fetus as would be the case in regular brain MRI. Similar to previous publications [25, 26], we therefore propose to first automatically segment the intracranial volume (ICV) of the fetus to identify the region of interest. A number of studies proposed segmentation of the ICV in fetal MRI [27–30]. Following our previous work [29], we segment the ICV directly in the entire image to fully automatically detect a region of interest.

The method we propose performs segmentation of fetal and brain tissues. The method first identifies the ICV from the fetal MRI slices using a convolutional neural network. Subsequently, the identified volume is segmented by another 2D convolutional neural network. Note that the proposed approach is applied to 2D slices of images reconstructed in a standard way, i.e. without reconstruction to high resolution volumes. The contribution of this paper is twofold: First, we propose data augmentation technique that synthesizes intensity inhomogeneity artifacts to improve the robustness against these artifacts. Second, the fetal brain segmentation is performed into seven classes: CB, basal ganglia and thalami (BGT), ventricular cerebrospinal fluid (vCSF), WM, BS, cortical gray matter (cGM) and eCSF in contrast to previous methods which focused on WM, cGM and cerebrospinal fluid only.

The remainder of this paper is organized as follows: in Section 2 the data set used for the method development and evaluation is described, in Section 3 the method for fetal brain segmentation and the simulation of intensity inhomogeneity are described, in Section 4 the evaluation method is given. The performed experiments and their results are presented in Section 5, followed by a discussion of the method and the results in Section 6. Our conclusions are given in the final section.

2. Data

2.1. Fetal MRI dataset

This study includes T2-weighted MR scans of 12 fetuses (22.9–34.6 weeks post menstrual age). Images were acquired on a Philips Achieva 3T scanner at the University Medical Center (UMC) Utrecht, the Netherlands, using a turbo fast spin-echo sequence. Repetition time (TR) was set to 2793 ms, echo time (TE) was set to 180 ms and the flip angle to 110°. The acquired voxel size was $1.25 \times 1.25 \times 2.5 \text{ mm}^3$, the reconstructed voxel size was $0.7 \times 0.7 \times 1.25 \text{ mm}^3$, and the reconstruction matrix was $512 \times 512 \times 80$. The images were not reconstructed using high resolution framework. We included images acquired in coronal plane as they appeared less affected by fetal motion in comparison with images acquired in axial or sagittal plane. The orientation of the fetal brain was determined using fast survey scanning. During the scan, the mother was lying on her left side to reduce the chance of inferior vena cava syndrome. The local ethical board approved the study and parental informed consent was obtained.

The reference standard was defined by manual annotation of all scans in 2D slices by a trained medical student. The brain was segmented into seven tissue classes: CB, BGT, vCSF, WM, BS, cGM and eCSF. Annotation was accomplished by manual pixel-wise painting of the brain tissues in each coronal image slice using an in-house developed software. The labeling of each of the seven classes was indicated by a color overlay (Fig. 3). The software allowed the user to zoom-in, zoom-out and scroll through the slices during the manual segmentation. The manual segmentation protocol was identical to the protocol described by Išgum et al. [25] for neonatal brain tissue segmentation. The ICV was defined as the union of all manually segmented tissue classes.

In total 15 slices (1.5%), 7 in the training set and 8 in the test set, were too distorted by severe motion artifacts to be manually annotated. In total 126 of the remaining, manually annotated slices (26.2%), 32 in the training set and 94 in the test set, were identified as affected by intensity inhomogeneity that hampered manual annotation.

2.2. Neonatal MRI dataset

While we propose a method aimed specifically at segmentation of fetal MRI scans, the proposed segmentation approach and especially the data augmentation technique that simulates intensity inhomogeneity (detailed in the following section) might be useful for brain segmentations in MRI scans. Therefore, to evaluate the applicability of this technique to a different MRI data set, we additionally included nine brain MR scans of preterm born infants.

T2-weighted MR images were acquired on a Philips Achieva 3T scanner at the University Medical Center Utrecht, the Netherlands. The images were made with Turbo Field Echo (TFE) and Turbo Spin Echo (TSE) sequences with TR set to 4847 ms and TE set to 150 ms. The scans were acquired at 40 weeks of post menstrual age in the coronal plane. The acquired voxel size was $0.35 \times 0.35 \times 1.2 \text{ mm}^3$ and the reconstruction matrix was $512 \times 512 \times 110$. The reference standard was defined by manual annotation of an expert into seven tissue types (CB, BGT, vCSF, WM, BS, cGM and eCSF) in MR scans without visible intensity inhomogeneity artifacts. Manual annotation was performed using the same protocol as described above for fetal MRI. These scans are part of the NeoBrainS12 segmentation challenge [31] and did not show intensity inhomogeneity artifacts.

The remaining four scans showed intensity inhomogeneity and they were used as a test set. However, manual reference segmentations in these scans were not available and, consequently, we evaluated the segmentation performance in this set only by visual inspection.

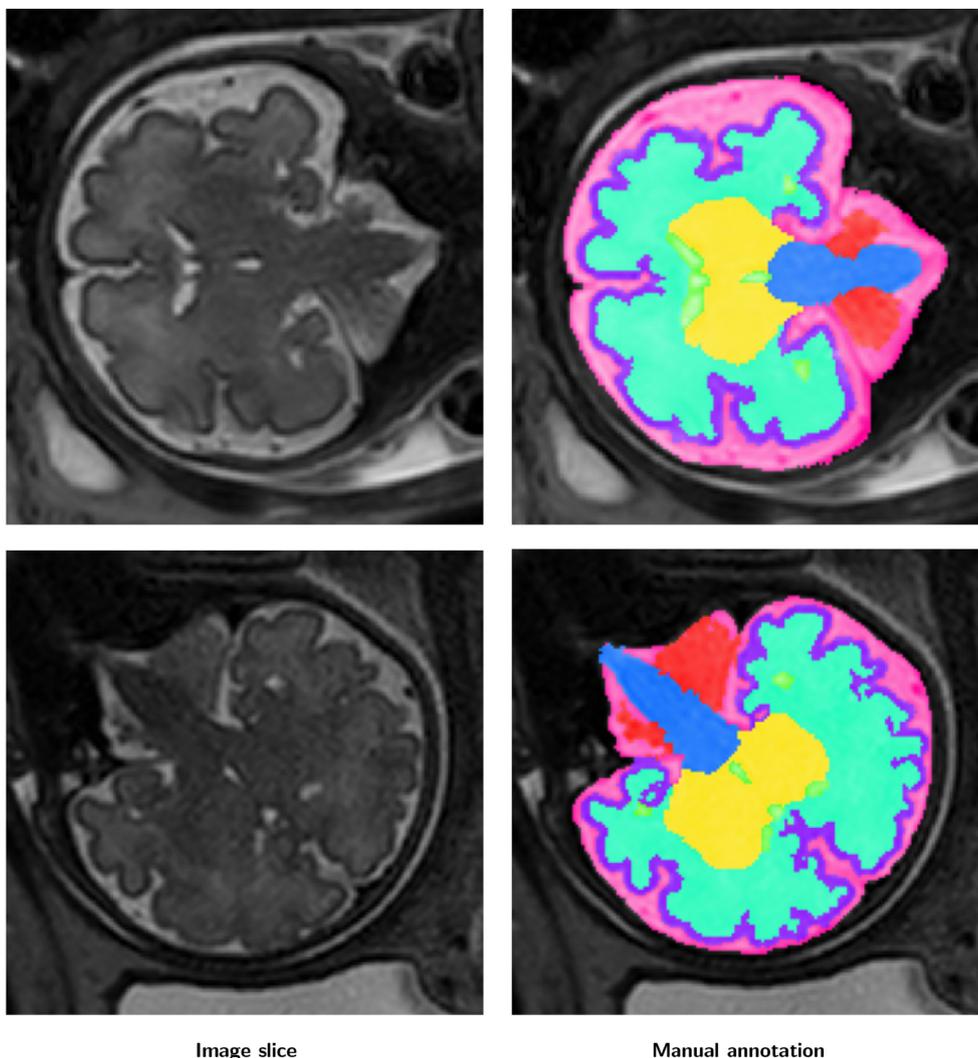
3. Method

To simplify the brain tissue segmentation and allow the segmentation method to focus on the fetal brain only, the fetal ICV is first automatically extracted. Subsequently, the identified ICV is automatically segmented into seven tissue classes. An overview of this pipeline is shown in Fig. 4. The same network architecture, described in Section 3.1, was used for ICV extraction and brain tissue segmentation.

3.1. Brain segmentation

Extraction of ICV and its subsequent segmentation into seven brain tissue classes are achieved with two fully convolutional networks (FCN) with identical U-net architecture [32] trained with 2D slices. Each network is trained independently to perform its specific task. The U-net architecture consists of a contracting path and an expanding path. The contracting path consists of repeated 3×3 convolutions followed by rectified linear units (ReLU). A 2×2 max pooling downsamples the features. The number of feature channels doubles after every two convolutional layers. In the expansion path, an up-sampling is followed by a 2×2 convolution which halves the number of feature channels. The results are concatenated with the corresponding contraction path and convolved by two 3×3 convolutional layers followed by a ReLU. At the final layer, one 1×1 convolutional layer maps each component of the feature vector to the desired number of classes. Batch normalization [33] is applied after all convolutional layers to allow for faster convergence. The network architecture is illustrated in Fig. 5.

Both networks were trained using stochastic gradient descent with back propagation. We optimized both networks using the Adam optimizer with Nesterov momentum [34, 35] using a fixed learning rate of 0.0001. Standard data augmentation techniques, namely random flipping and rotation, were used during training to increase the variation of the training data and to mimic different orientations of the fetal brain.



CB BGT vCSF WM BS cGM eCSF

Fig. 3. Examples of manual reference segmentation in coronal fetal MRI, showing slices cropped to the region of interest (first column) and the slices overlaid with the manual segmentations (second column). (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

The slices were flipped in horizontal and vertical direction with 50% probability and were rotated with a rotation angle randomly chosen between 0 and 360°.

To first identify the intracranial area in the image, a U-net was trained by minimizing the cross-entropy between network output and manual segmentation for all pixels in each slice in each training batch. The convolutional neural network was trained with batches of 12 slices in each iteration. Given that the network performs voxel classification, ICV segmentation may result in small isolated clusters of false positive voxels outside the ICV. These were removed by discarding 3D connected components smaller than 3 cm³. This threshold was empirically

chosen in this study. It was chosen large enough to remove small false positive clusters of voxels, and also small enough to prevent removing any parts of the brain that are not fully connected in 3D. The latter is often the case in scans with substantial motion artifacts where the signal is lost in one slice and consequently, the ICV segmentation is not fully connected in 3D.

The segmented intracranial fetal areas were further segmented into seven tissue classes using another, separately trained, U-net. Each pixel in the image was classified as either CB, BGT, vCSF, UWM, BS, cGM, eCSF or background. In contrast to the network for ICV segmentation, this network was trained by maximizing the Dice coefficient between

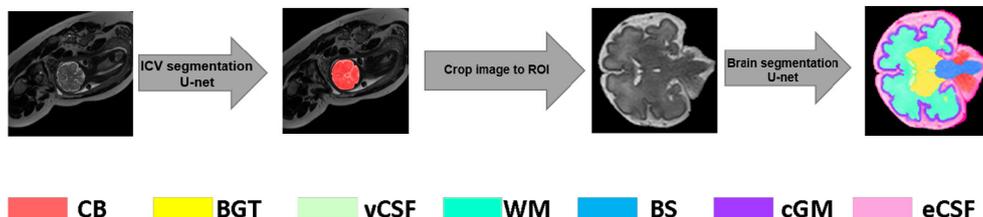


Fig. 4. Proposed pipeline for automatic tissue segmentation method. The fetal ICV is automatically segmented in the original slice. Thereafter, the image is automatically cropped to the region of interest (ROI) so that the tissue segmentation can be restricted to the fetal ICV only.

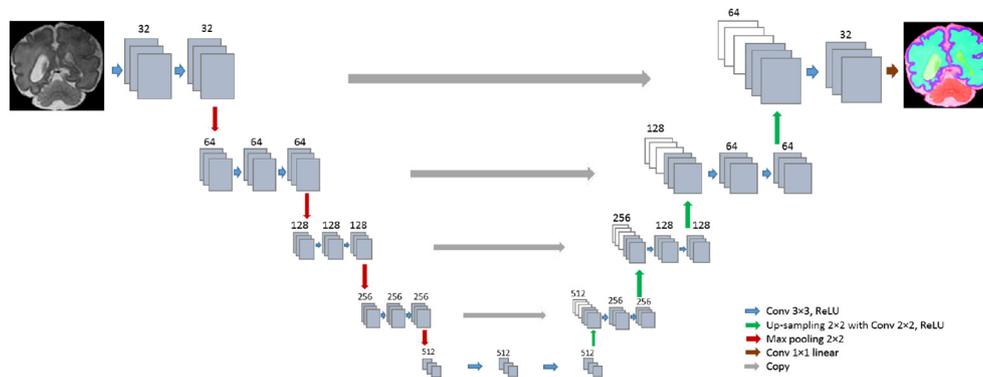


Fig. 5. Network architecture: The network consists of a contraction path and an expansion path. The contraction path consists of repeated convolution layers followed by max pooling, and the expansion path consists of convolution layers followed by upsampling.

network output and manual segmentation. This was done to achieve robustness against an imbalance of samples from the different classes. This network was trained with batches of 18 slices in each iteration. We have implemented the network in Keras, an open-source neural-network library written in Python [36].

3.2. Intensity inhomogeneity augmentation (IIA)

To make the network segmenting brain tissue classes robust to intensity inhomogeneity artifacts, we trained this network with slices containing simulated intensity inhomogeneity artifacts. The artifacts were simulated by applying a combination of linear gradients with random offsets and orientations to a slice without intensity inhomogeneity artifacts (I):

$$Z = I \times ((X + x_0)^2 + (Y + y_0)^2), \quad (1)$$

where X and Y are 2D matrices with integer values from zero to the size of the image in x and y direction, respectively. The offsets x_0 and y_0 control the balance between the x and y components and were randomly chosen from different ranges (x_0 : [43, 187]; y_0 : [-371, 170]). The optimal ranges were found with a random hyperparameter search. Additionally, the gradient patterns were randomly rotated between 0 to 360° to mimic intensity inhomogeneity in various directions. These random components, offsets and rotation, result in inhomogeneity patterns that allow the network to become invariant to the location and orientation of regions with low and decreasing contrast. The intensities in both the original slices as well as the slices with simulated intensity inhomogeneity were normalized to the range [0,1023] before feeding them to the network. Fig. 6 shows examples of two slices from a fetal MRI scan with added synthetic intensity inhomogeneity.

4. Evaluation

The automatic brain tissue segmentation was evaluated by means of the Dice coefficient (DC) for volume overlap and the mean surface distance (MSD) between manual reference segmentation and automatically obtained segmentation. In the fetal MRI scans, these metrics were calculated in 2D, i.e., per slice, and were then averaged across all slices. In the neonatal MRI scans, following previous work [26], these metrics were calculated in 3D.

5. Experiments and results

In our experiments, we first evaluated the overall segmentation performance of the proposed pipeline with respect to the different tissue classes. To evaluate the influence of the proposed intensity inhomogeneity augmentation technique with the standard augmentation techniques, the segmentation performance before and after applying intensity inhomogeneity augmentation was evaluated. Furthermore, we

evaluated whether this augmentation technique is able to generalize to different data, i.e., whether it leads to similar performance improvements in neonatal brain segmentation. The fetal MRI dataset was randomly divided into a training and a test set. Each set contained 6 scans. The neonatal MRI dataset was divided into a training set with 3 scans and a test set with the remaining 6 scans for which the manual reference segmentations of two scans were available. The training and test sets are listed in Table 1.

5.1. Segmentation performance

The performance of the proposed method using standard augmentation and the proposed IIA as described in Section 3 was evaluated. Slices with intensity inhomogeneity artifacts resulting from image acquisition were excluded from the training data. The average performance in the six test scans is listed in Table 2 for each of the seven tissue classes (Set 1). The average DC ranged from 0.80 for CB to 0.94 for eCSF and the average MSD ranged from 0.62 mm for CB to 0.18 mm for BS. Furthermore, to evaluate whether IIA improves the performance differently in slices with artifacts than in artifact free slices, we compared the performance on slices with clearly visible intensity inhomogeneity artifacts and slices without visible artifacts. These results are also listed in Table 2. As shown, the automatic segmentations were less accurate on slices with intensity inhomogeneity strong enough to hamper manual annotation compared with slices without visible intensity inhomogeneity. Figs. 7 and 8 illustrate the segmentation performance on slices with intensity inhomogeneity and without visible intensity inhomogeneity, respectively. Note that these results were obtained with networks trained without any slices with intensity inhomogeneity artifacts resulting from the image acquisition, but slices with simulated intensity inhomogeneity were used for training. Excluding slices with intensity inhomogeneity from the training set is a more realistic training scenario because manual segmentation of such slices is cumbersome and reference segmentations for such slices might therefore not be available. However, extending the training set with image slices affected by real intensity inhomogeneity artifacts but in which manual annotation was still feasible could potentially further improve the performance. We therefore trained the networks also including slices with intensity inhomogeneity artifacts resulting from image acquisition (Set 2). The quantitative results are listed in Table 3. As shown in the table, in this experiment, we also separately evaluated the performance and impact of IIA in slices with visible and without visible intensity inhomogeneity artifacts.

Finally, we compared the performance of the proposed fetal brain tissue segmentation method with the performance of previous methods (Table 4). The performance of the proposed method was comparable to the performance of other methods, even though it performs a finer segmentation into seven tissue classes instead of only four [12] or

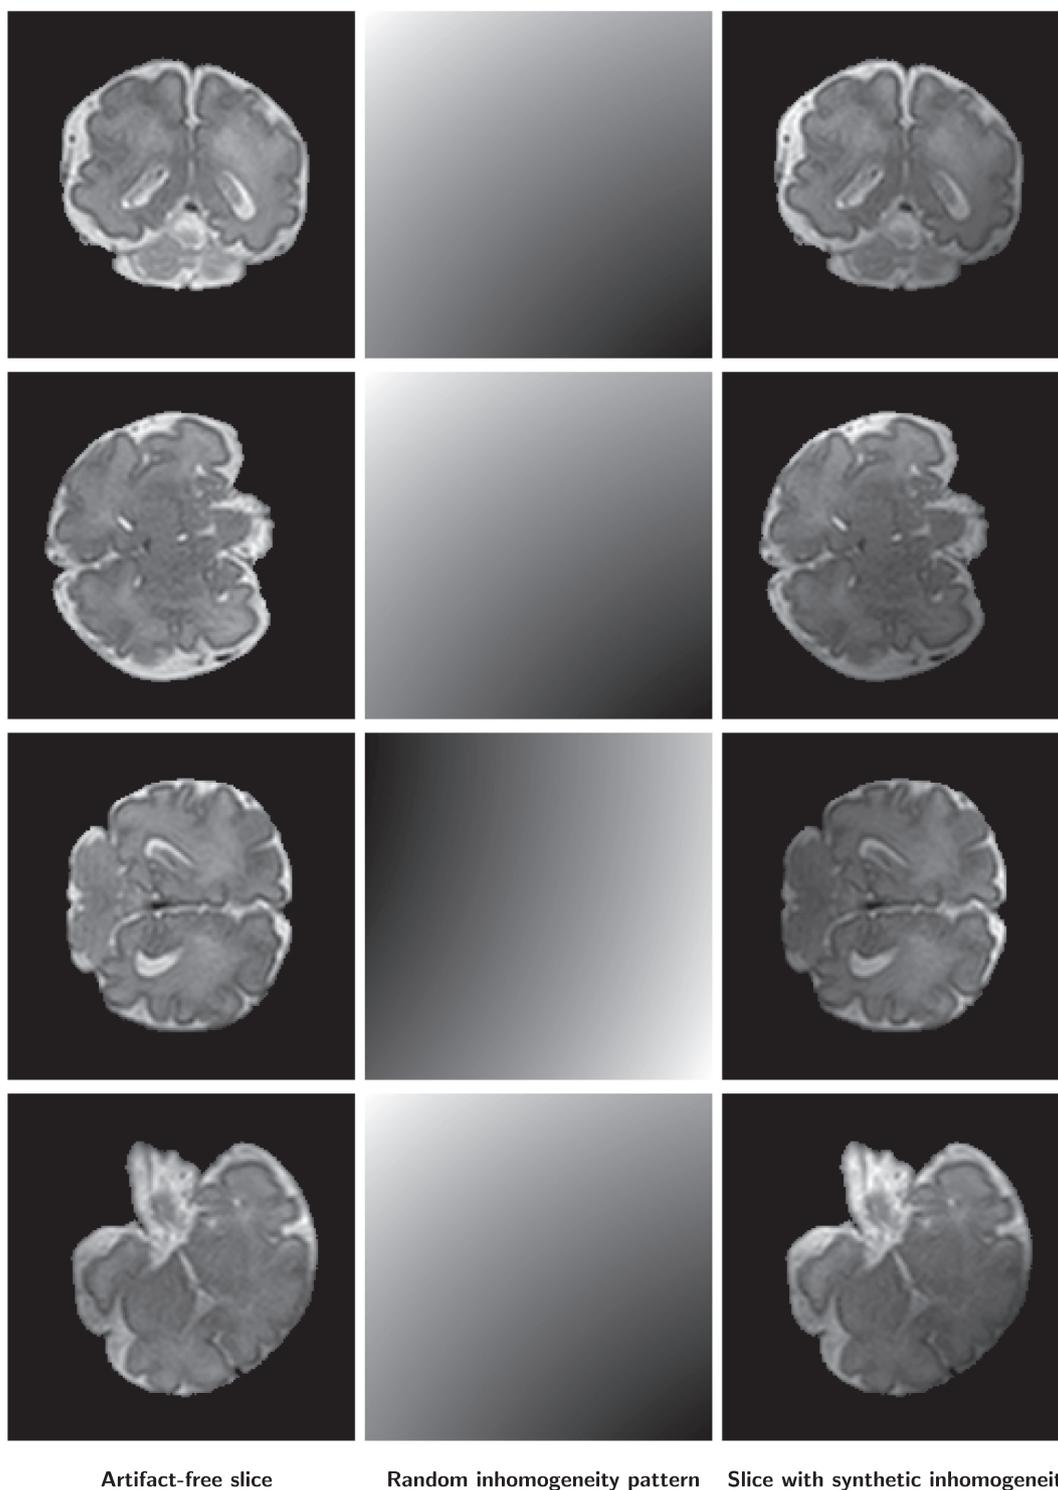


Fig. 6. Example of coronal slices with simulated intensity inhomogeneity. Original slice (first column), simulated intensity inhomogeneity pattern (second column) and the image slice after adding the synthetic intensity inhomogeneity artifact (third column).

three [16] tissues. The performance of previous methods is taken from the literature. Hence, the methods have been evaluated using different data set and thus this comparison can provide an indication only.

5.2. Comparison of data augmentation techniques

To evaluate the influence of the proposed IIA as well as the influence of the standard data augmentation techniques used in this study (random flipping and random rotation), the following experiments

using fetal scans were performed. In this experiment using fetal scans, all slices with intensity inhomogeneity were removed from the training set (Set 1). First, a network was trained without any data augmentation to serve as baseline for the comparison. Second, a network was trained using random flipping of the training slices to augment the training data. Third, a network was trained using random flipping and random rotation of the training slices. Finally, as presented in Section 5.1, a network was trained using random flipping and random rotation, and additionally with IIA, i.e., randomly simulated intensity

Table 1

The fetal MRI dataset was randomly divided into a training and a test set each containing 6 scans. We defined two training sets with fetal data: Set 1 contained 6 scans, but without those slices in which intensity inhomogeneity was visible. Set 2 contained the entire slices of 6 scans. The neonatal MRI dataset was divided into a training set with three scans and a test set with the six remaining scans (Set 3). The test set in all three sets contained 6 scans.

	Training set	Training set description
Set 1 (fetal)	6 scans	Excluding slices with intensity inhomogeneity artifacts
Set 2 (fetal)	6 scans	Including slices with intensity inhomogeneity artifacts
Set 3 (neonatal)	3 scans	Intensity inhomogeneity not visible in any slice

Table 2

Performance of fetal brain tissue segmentation into seven tissue classes when the network is trained on slices without intensity inhomogeneity resulting from image acquisition. The network is evaluated on the entire test set and additionally on only the slices with intensity inhomogeneity artifacts (94 slices) and on the slices without visible intensity inhomogeneity artifacts (357 slices). The segmentation performance with intensity inhomogeneity augmentation (IIA) used during training is compared with the performance of the same network without IIA. The results are expressed as the mean Dice coefficient (DC) and the mean surface distance (MSD) in mm.

			CB	BGT	vCSF	WM	BS	cGM	eCSF	Mean
All test slices	IIA	DC	0.802	0.889	0.875	0.922	0.930	0.829	0.943	0.884
		MSD	0.620	0.414	0.470	0.384	0.181	0.318	0.188	0.368
	Without IIA	DC	0.688	0.807	0.724	0.849	0.850	0.672	0.820	0.773
		MSD	0.995	0.726	1.331	0.875	0.253	0.769	0.549	0.785
Slices with II	IIA	DC	0.694	0.901	0.807	0.899	0.947	0.782	0.877	0.844
		MSD	0.813	0.372	0.807	0.565	0.286	0.418	0.416	0.525
	Without IIA	DC	0.467	0.572	0.485	0.744	0.704	0.492	0.678	0.592
		MSD	1.724	1.340	3.742	1.656	0.629	1.275	1.000	1.624
Slices without II	IIA	DC	0.754	0.914	0.854	0.918	0.926	0.829	0.923	0.874
		MSD	0.428	0.406	0.729	0.426	0.159	0.351	0.241	0.392
	Without IIA	DC	0.719	0.837	0.802	0.902	0.933	0.758	0.870	0.832
		MSD	0.536	0.502	0.675	0.595	0.208	0.446	0.403	0.481

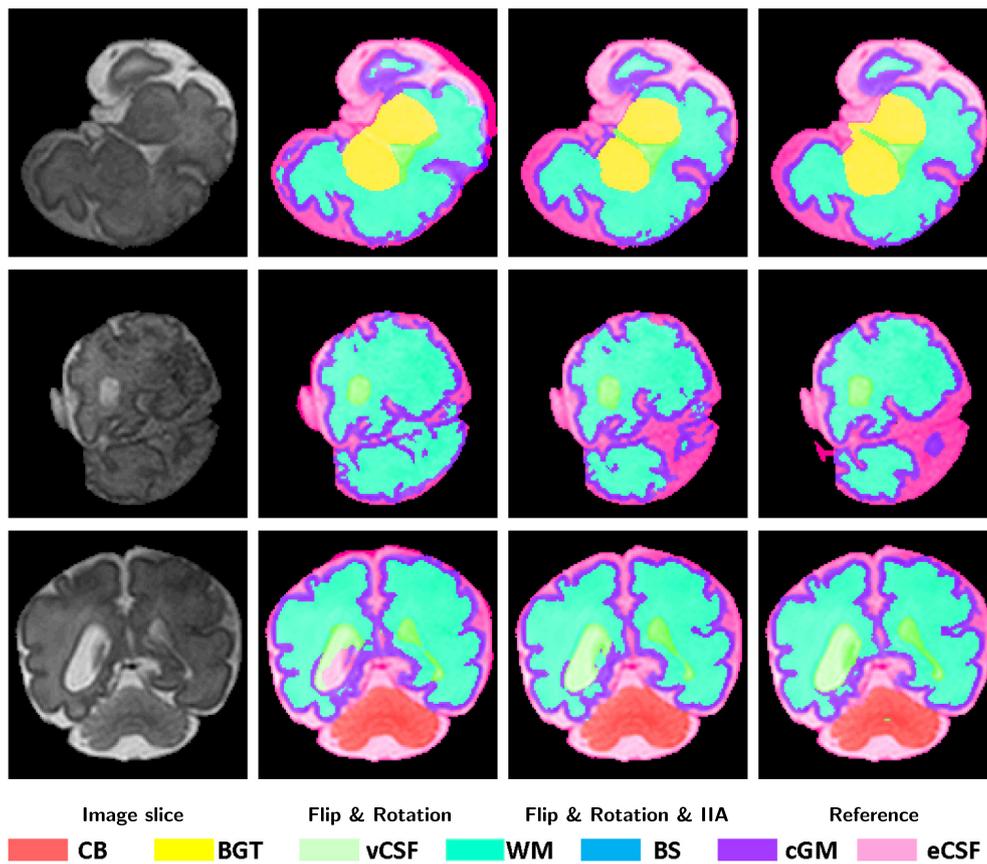


Fig. 7. Examples of automatic brain tissue segmentation in slices with visible intensity inhomogeneity. A slice from T2-weighted fetal MRI scan with visible intensity inhomogeneity (first column); segmentation obtained with network only using flipping and rotation augmentation (second column); segmentation obtained with network using IIA (third column); manual reference segmentation (fourth column).

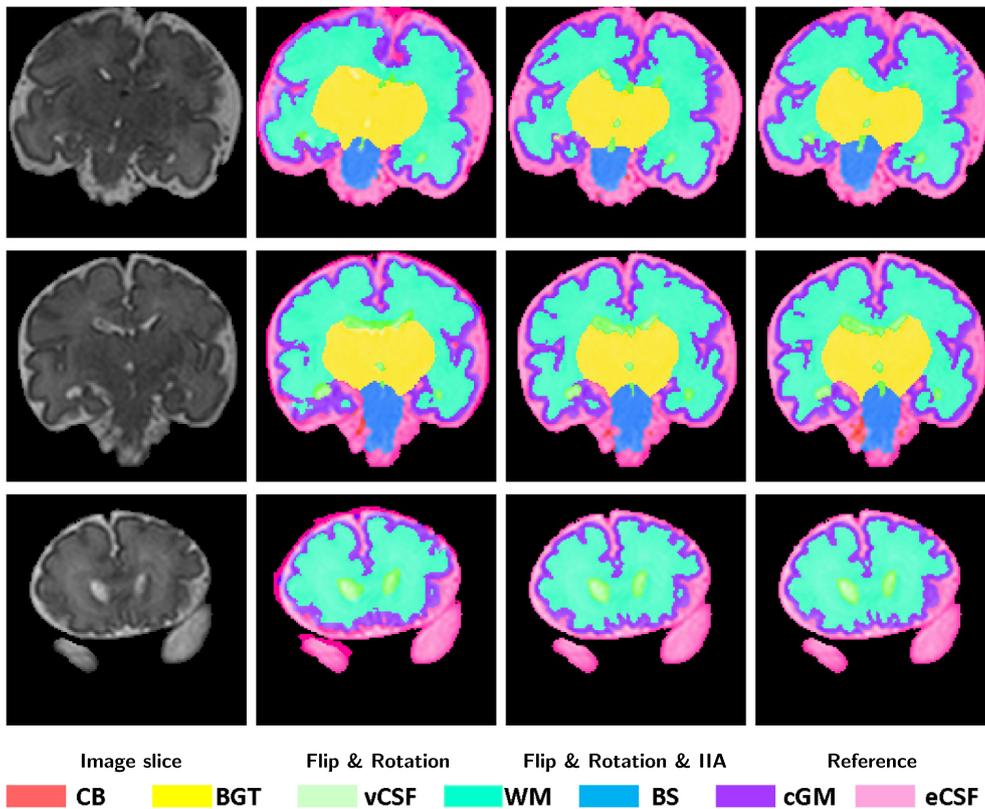


Fig. 8. Examples of automatic brain tissue segmentation in fetal images acquired in the slices without visible intensity inhomogeneity. A slice from T2-weighted fetal MRI scan with visible intensity inhomogeneity (first column); segmentation obtained with network using only flipping and rotation augmentation (second column); segmentation obtained with network additionally using IIA (third column); manual reference segmentation (fourth column).

Table 3

Performance of fetal brain tissue segmentation into seven tissue classes when the network is trained with slices with intensity inhomogeneity resulting from the image acquisition. The network is evaluated on the entire test set and additionally on only slices with intensity inhomogeneity (94 slices) and on the slices without intensity inhomogeneity (357). The segmentation performance with intensity inhomogeneity augmentation (IIA) in the training is compared with the performance of the same network without IIA in the training. The results are expressed as the mean Dice coefficient (DC) and the mean surface distance (MSD) in mm.

			CB	BGT	vCSF	WM	BS	cGM	eCSF	Mean
All test slices	IIA	DC	0.794	0.931	0.874	0.919	0.946	0.835	0.944	0.892
		MSD	0.715	0.428	0.434	0.381	0.195	0.307	0.182	0.377
	Without IIA	DC	0.778	0.887	0.851	0.923	0.931	0.821	0.940	0.876
		MSD	0.771	0.501	0.521	0.440	0.193	0.356	0.195	0.425
Slices with II	IIA	DC	0.719	0.906	0.816	0.901	0.934	0.790	0.878	0.849
		MSD	0.735	0.426	0.703	0.613	0.278	0.398	0.413	0.509
	Without IIA	DC	0.722	0.877	0.842	0.905	0.941	0.801	0.885	0.853
		MSD	0.951	0.535	0.705	0.521	0.267	0.351	0.397	0.532
Slices without II	IIA	DC	0.749	0.905	0.884	0.927	0.926	0.856	0.934	0.883
		MSD	0.489	0.438	0.388	0.343	0.150	0.272	0.212	0.327
	Without IIA	DC	0.731	0.923	0.888	0.933	0.933	0.850	0.932	0.884
		MSD	0.520	0.402	0.363	0.360	0.141	0.290	0.215	0.327

Table 4

Segmentation performance of the proposed method and of other methods evaluated with Dice coefficient. Performance of previous methods is taken from the literature. Hence, this comparison can be used as indication only.

	CB	BGT	vCSF	WM	BS	cGM	eCSF
Proposed method	0.794	0.931	0.874	0.919	0.946	0.835	0.944
Hebas et al. [12]	–	–	0.900	0.900	–	0.820	–
Serag et al. [16]	–	–	0.920	0.900	–	0.840	–

inhomogeneity. In this last experiment, all slices were manipulated with IIA during the training. Results of this last experiment are listed in Table 2.

The achieved average DC and MSD for each scan in the test set are shown in Fig. 9. The performance improved in all scans the more data augmentation was used and especially further improved when IIA was

added in addition to the standard data augmentation techniques. IIA largely reduced performance differences between different scans, which standard augmentation techniques were not able to achieve. The performance improvement was particularly large in a scan in which nearly all slices showed intensity inhomogeneity (yellow marker in Fig. 9). With IIA, the segmentation performance reached an accuracy comparable to that achieved on scans with fewer artifacts.

Additionally, the achieved average DC and MSD for each tissue class are shown in Fig. 10. Overall, all augmentation methods improved segmentation performance. Adding augmentation based on random flipping and rotation of the image slices improved the segmentation of CB, BS, vCSF, BGT noticeably but not the segmentation of WM, cGM and eCSF. IIA was able to further improve the segmentation performance in all tissue classes including WM, cGM and eCSF. Overall, CB segmentation performance benefited the most from all augmentations where the DC increased from 0.3 to 0.8. Furthermore, augmentation

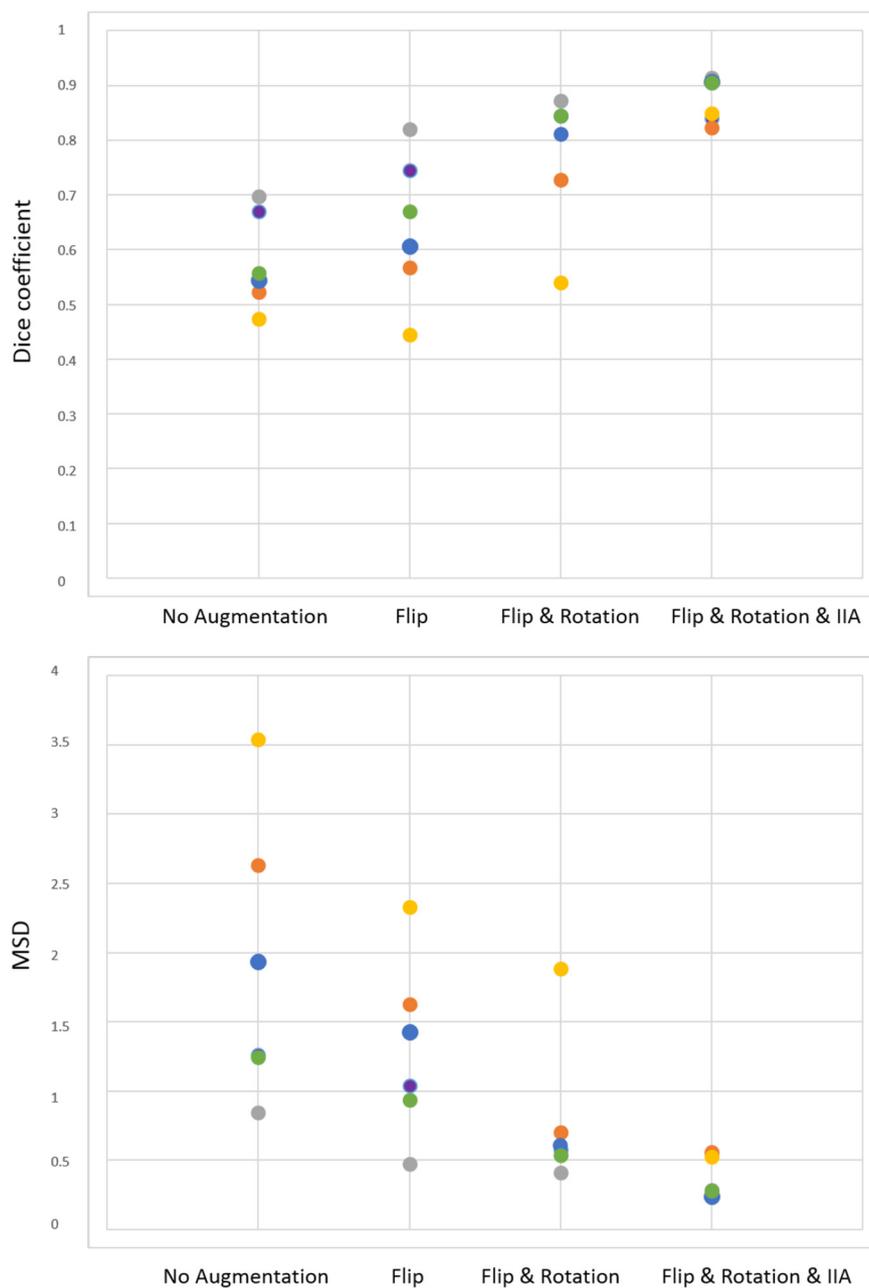


Fig. 9. Performance of automatic fetal brain tissue segmentation into seven tissue classes when the network is trained without any augmentation, with flipping slices as augmentation, with flipping and rotating slices as augmentation, and with flipping, rotating and IIA. The results are expressed as the mean Dice coefficient and the mean surface distance (MSD) in mm. Each marker corresponds to the mean performance across all tissue classes in one of the six test scans. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

based on random rotation severely influenced performance in segmentation of BGT, BS and CB.

5.3. Impact of extent of intensity inhomogeneity augmentation

To assess the impact of the proportion of slices in each training batch with simulated intensity inhomogeneity on the network performance, we trained the network with different percentage of artifact free slices to slices with simulated intensity inhomogeneity. In this experiment, all slices with intensity inhomogeneity were removed from the training set (Set 1). We varied number of slices with added synthetic intensity inhomogeneity from 0% to 100%. Fig. 11 shows the obtained results. Training with only simulated slices led to slightly worse performance compared with a mix of original and manipulated slices. Even

having only 20% of the slices with synthetic intensity inhomogeneity already improved the performance substantially. There was, however, no marked improvement for larger percentages of slices with synthetic intensity inhomogeneity, the performance was comparable for percentage of 20% to 100%.

5.4. Evaluation of IIA on neonatal brain segmentation

Intensity inhomogeneity is an artifact that occurs in various of MRI scans, albeit often to lesser extent than in fetal MRI. To assess whether the proposed augmentation technique-IIA-is also able to improve the performance of a segmentation task in other MR images, we trained the brain tissue segmentation network to perform segmentation in neonatal brain MRI scans. Following previous work [26] to limit the number of

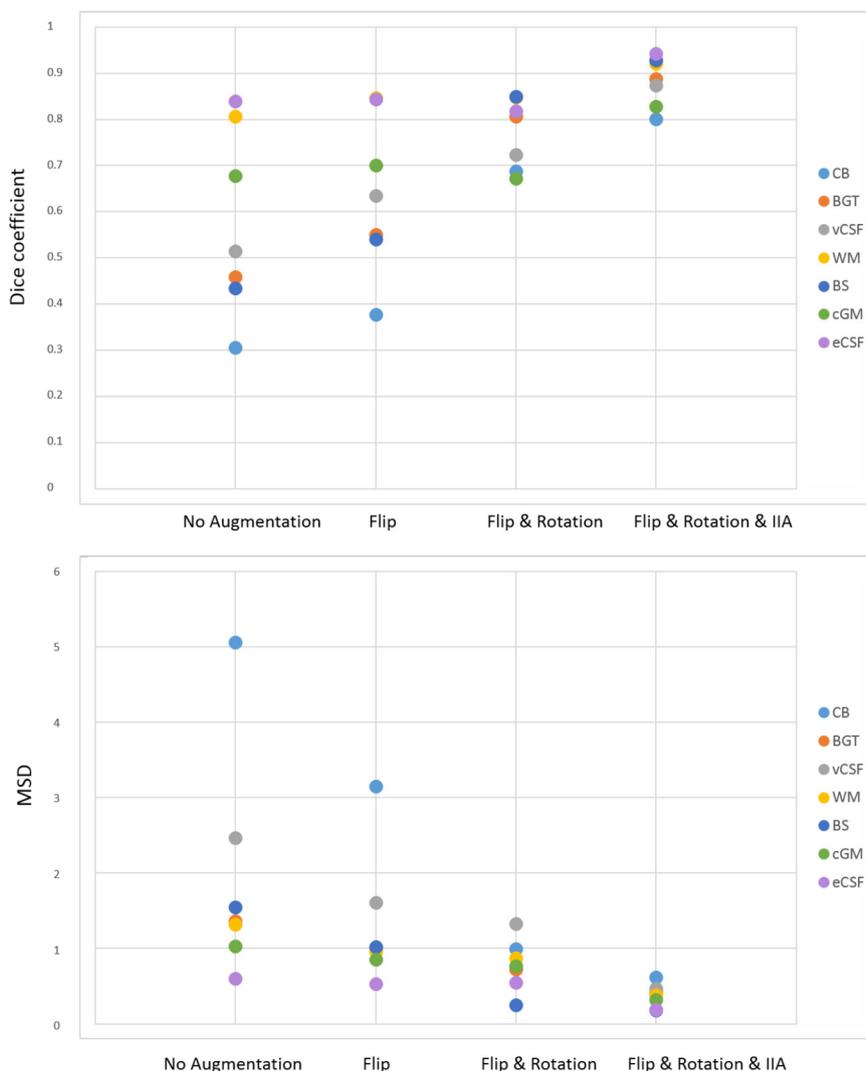


Fig. 10. Performance of automatic fetal brain tissue segmentation for each of the seven tissue classes when the network is trained without any augmentation, with flipping augmentation, with flipping and rotating augmentation, and with flipping, rotating and IIA. The results are expressed as the mean Dice coefficient and the mean surface distance (MSD) in mm. Each marker corresponds to the mean performance across all six test scans for one of the seven tissue classes.

voxels considered in the classification, brain masks were generated with BET [37]. The network was trained with three scans and tested with the remaining two scans (Set 3). Like in the previous experiments, the network was first trained using standard data augmentation, i.e., using random flipping and rotation of the training slices as described in Section 3, and subsequently the network was trained additionally with IIA. The obtained segmentation results are listed in Table 5. As the two scans in the test set did not show any intensity inhomogeneity artifacts, we additionally qualitatively evaluated the segmentation performance on four scans with visible intensity inhomogeneity artifacts for which manual reference segmentations were not available. We illustrate segmentation results in the four scans without reference standard in Fig. 12. In [26] these scans were not analyzed due to the presence of artifacts. Visual inspection of the results in these scans reveals that the segmentation was more accurate when IIA was used, particularly in BGT, BS, vCSF and eCSF.

6. Discussion

We presented a pipeline for automatic segmentation of the fetal brain into seven tissue classes in MRI. The method consists of two fully convolutional networks with identical U-net architectures. The first network extracts ICV and the second network performs segmentation of

the brain into seven tissue classes. The results demonstrate that segmentation using the proposed data augmentation with simulated intensity inhomogeneity artifacts leads to accurate segmentations of the brain tissue classes. Moreover, we demonstrated that the method performs accurate segmentation while trained using manual reference segmentation only in slices without artifacts that occur during image acquisition. In other words, we showed that the proposed data augmentation is able to compensate for the lack of training data in which performing manual annotations is cumbersome.

Using the proposed data augmentation technique, we were able to achieve state-of-the-art segmentation performance with a substantially lower number of training scans. Our method was trained with only 6 fetal scans while previous methods used 20 [13] and up to 80 fetal scans [38]. Given that manual annotation of a fetal brain MR scan into 7 tissue classes requires about 40 h, reducing the number of the manually annotated training scans substantially reduces the required manual annotation effort and associated costs.

Intensity inhomogeneity is a frequently occurring artifact in MRI and often hampers automatic image analysis due to diminishing contrast between different tissues. We demonstrated that the described method based on convolutional neural networks can become more robust to these artifacts by training with data augmented with simulated random intensity inhomogeneities. This can potentially replace or

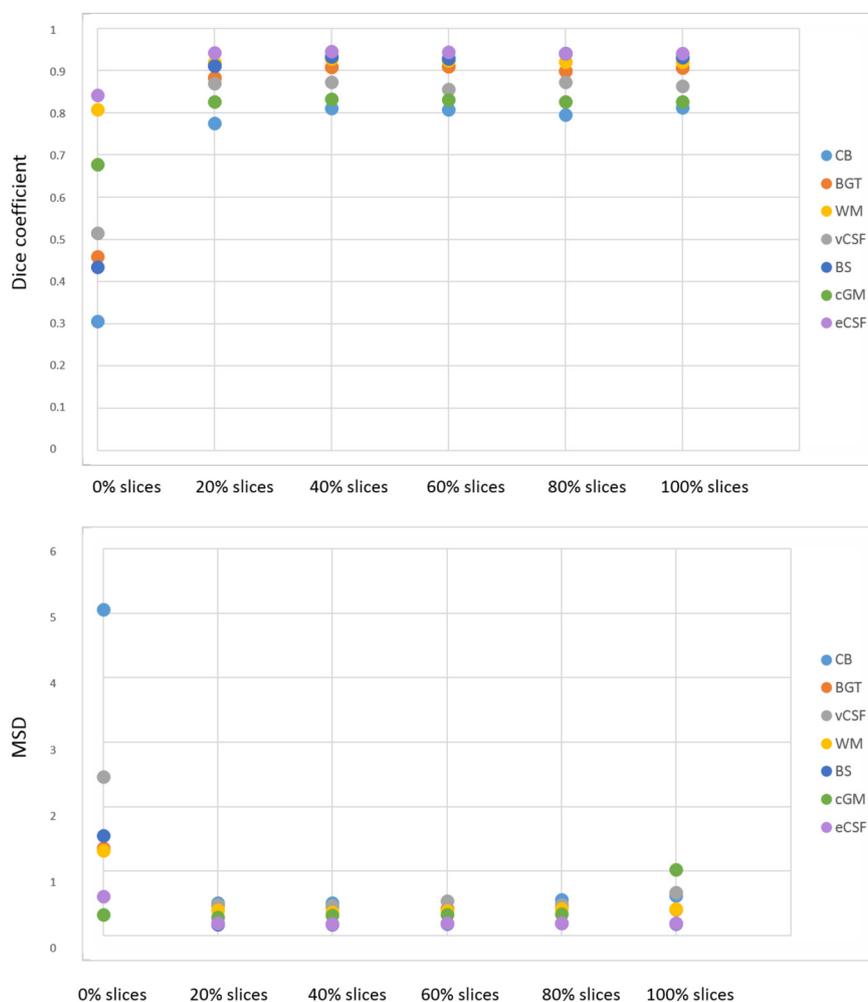


Fig. 11. Performance of fetal brain tissue segmentation into seven tissue classes with different proportion of slices with synthetic intensity inhomogeneity in each training batch. The results are expressed as the mean Dice coefficient (DC) and the mean surface distance (MSD) in mm. Each marker corresponds to the mean performance across all six test scans for a different tissue class.

Table 5

Performance of neonatal brain segmentation into seven tissue classes. The segmentation performance with IIA is compared with the performance of the same network without IIA. The results are expressed as the mean Dice coefficient (DC) and the mean surface distance (MSD) in mm.

			CB	BGT	vCSF	WM	BS	cGM	eCSF	Mean
All test slices	With IIA	DC	0.857	0.883	0.776	0.825	0.765	0.500	0.600	0.744
		MSD	0.919	0.677	0.751	0.385	0.510	0.375	0.625	0.643
	Without IIA	DC	0.858	0.854	0.737	0.819	0.776	0.510	0.596	0.736
		MSD	1.465	1.359	0.617	0.393	1.383	0.372	0.635	0.889

complement preprocessing steps, such as bias field corrections or volumetric reconstructions that would require acquisition of additional MR data. Simulating artifacts instead of manually annotating compromised data for training supervised methods is beneficial as manual reference annotations can be obtained more easily and with higher accuracy for artifact-free data.

Standardly used data augmentation techniques, random flipping and rotation of the image slices, improved the segmentation performance in CB, BS, BGT, and vCSF considerably but had little impact on other tissue types. Introducing training data with simulated intensity inhomogeneity further improved the segmentation performance in all tissue classes, including WM and cGM, which are often challenging to separate due to low inter-tissue contrast when intensity inhomogeneity artifacts are present. Moreover, a frequent mistake of the automatic segmentation method when training without IIA was mistaking vCSF

for eCSF and vice versa. Using IIA helped to overcome this issue in many cases, presumably by forcing the network not to focus on the intensity values only but additionally on other intensity invariant information such as shape and context.

Moreover, evaluating augmentation techniques per scan showed that the MSD reduces with adding random flipping and rotating augmentation in all scans and DC improves in all scans except one (indicated with a yellow marker in Fig. 9). Retrospective visual inspection revealed that this scan has intensity inhomogeneity in nearly all slices. Even though the intensity inhomogeneity is not severe in all slices, the automatic segmentation was still severely affected. Experiments show that adding IIA in the training helped to overcome this issue and increased the segmentation performance in all scans.

Furthermore, our experiments illustrate that training the network with IIA increases the segmentation performance even in slices without

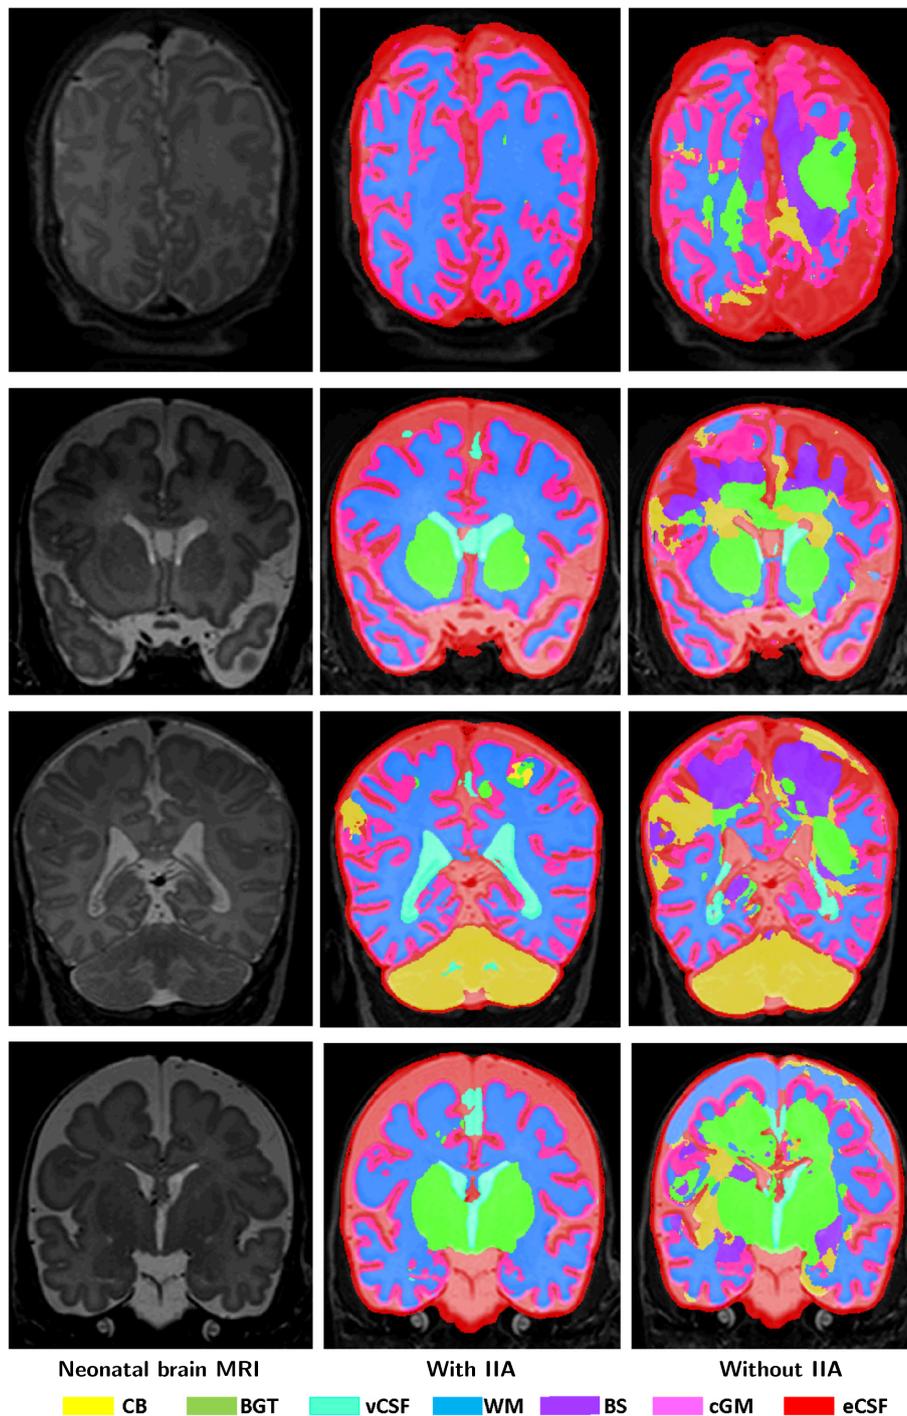


Fig. 12. Examples of brain tissue segmentation in neonatal MRI with intensity inhomogeneity artifacts. A slice from T2-weighted fetal MRI scan (first column); segmentation obtained with network using rotation, flipping and IIA (second column); segmentation obtained with network only using flipping and rotation of the slices as training data augmentation (third column).

visible intensity inhomogeneity. IIA makes the network more robust to intensity variations in MRI, forcing the network not to focus only on the tissue intensity for assigning a label.

In the current study, the segmentation method was evaluated on fetal brain MRI acquired in the coronal plane. Since the presented method is entirely supervised it can be readily applied to fetal MRI acquired in axial or sagittal plane if manually annotated training data is available.

We relied on the U-net architecture for ICV and brain tissue segmentation in fetal MRI. However, the proposed IIA can be used for data

augmentation regardless of the network architecture or even with supervised methods not based on convolutional neural networks. We did not evaluate IIA with other architectures but it would likely improve the segmentation performance of other networks with different architecture as supervised CNNs regularly profit from large and diverse training data.

Additionally, we evaluated IIA on neonatal MRI. The visual inspection shows a substantial improvement on slices with artifacts. In the images without visible artifacts the quantitative results showed slight improvement when IIA is applied. The performance of the

segmentation in the neonatal brain scans is lower than obtained by our previous method [26] as here presented network was not specifically adjusted for segmentation of neonatal brain tissues. However, the results clearly demonstrate the benefit of training with IIA. IIA could be readily applied to the segmentation method presented in [26] enabling the multi-scale CNN to segment neonatal brain scans with intensity inhomogeneity artifacts as shown in Fig. 12.

In this study, 2D analysis was applied since fetal MR images has inter-slice motion due to the 2D MR acquisition. Additionally, in a few slices severe motion artifacts occurred. These slices were excluded from the training and test set as manually segmenting them for evaluation would be hardly feasible. Generating such slices that are heavily affected by motion artifacts in fetal MRI could be an interesting direction for future work.

We have trained the proposed method with representative data, i.e. in a supervised manner. Training with non-representative, in addition to the representative data, using transfer learning would allow increasing the training sample size. This could be addressed in future work and it could potentially further improve segmentation performance.

7. Conclusion

We presented an automatic method for brain tissue segmentation in fetal MRI into seven tissue classes using convolutional neural networks. We demonstrated that the proposed method learns to cope with intensity inhomogeneity artifacts by augmenting the training data with synthesized intensity inhomogeneity artifacts. This can potentially replace or complement preprocessing steps, such as bias field corrections, and help to substantially improve the segmentation performance.

Acknowledgments

This study was sponsored by the Research Program Specialized Nutrition of the Utrecht Center for Food and Health, through a subsidy from the Dutch Ministry of Economic Affairs, the Utrecht Province and the Municipality of Utrecht. Furthermore, we thank Nienke Heuvelink for her help with creating the manual fetal brain segmentations.

References

- Ivica Kostović, Nataša Jovanov-Milošević. The development of cerebral connections during the first 20–45 weeks gestation. *Semin. Fetal Neonatal Med.* 2006;11(6):415–22.
- Karl Zilles, Nicola Palomero-Gallagher, Katrin Amunts. Development of cortical folding during evolution and ontogeny. *Trends Neurosci.* 2013;36(5):275–84.
- Marie Schaer, Meritxell Bach Cuadra, Nick Schmansky, Bruce Fischl, Jean-Philippe Thiran, Stephan Eliez. How to measure cortical folding from MR images: a step-by-step tutorial to compute local gyrification index. *J. Vis. Exp.* 2012(59).
- Twickler DianeM, Taylor Reichel, McIntire DonaldD, Magee KevinP, Ramus RonaldM. Fetal central nervous system ventricle and cisterna magna measurements by magnetic resonance imaging. *Am. J. Obstet. Gynecol.* 2002;187(4):927–31.
- Deborah Levine, Barnes PatrickD, Robertson RichardR, Geoffrey Wong, Melita TejasS. Fast MR imaging of fetal central nervous system abnormalities. *Radiology* 2003;229(1):51–61.
- Yamashita Y, Namimoto T, Abe Y, Takahashi M, Iwamasa J, Miyazaki K. MR imaging of the fetus by a haste sequence. *AJR Am. J. Roentgenol.* 1997;168(2):513–9.
- Shuzhou Jiang, Hui Xue, Alan Glover, Mary Rutherford, Daniel Rueckert, Hajnal JosephV. MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies. *IEEE Trans. Med. Imaging* 2007;26(7):967–80.
- Ali Gholipour, Estroff JudyA, Warfield SimonK. Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. Med. Imaging* 2010;29(10):1739–58.
- Maria Kuklisova-Murgasova, Gerardine Quaghebeur, Rutherford MaryA, Hajnal JosephV, Schnabel JuliaA. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Med. Image Anal.* 2012;16(8):1550–64.
- Bernhard Kainz, Markus Steinberger, Wolfgang Wein, Maria Kuklisova-Murgasova, Christina Malamateniou, Kevin Keraudren. Fast volume reconstruction from motion corrupted stacks of 2D slices. *IEEE Trans. Med. Imaging* 2015;34(9):1901–13.
- Michael Ebner, Guotai Wang, Li Wenqi, Michael Aertsen, Patel PremalA, Rosalind Aughwane. An automated localization, segmentation and reconstruction framework for fetal brain MRI. *International Conference on Medical Image Computing and Computer-Assisted Intervention.* Springer; 2018. p. 313–20.
- Habas PiotrA, Kio Kim, Francois Rousseau, Glenn OritA, AJames Barkovich, Colin Studholme. Atlas-based segmentation of developing tissues in the human brain with quantitative validation in young fetuses. *Hum. Brain Mapp.* 2010;31(9):1348–58.
- Ali Gholipour, Alireza Akhondi-Asl, Estroff JudyA, Warfield SimonK. Multi-atlas multi-shape segmentation of fetal brain MRI for volumetric and morphometric analysis of ventriculomegaly. *NeuroImage* 2012;60(3):1819–31.
- Ali Gholipour, Estroff JudyA, Barnewolt CarolE, Connolly SusanaA, Warfield SimonK. Fetal brain volumetry through MRI volumetric reconstruction and segmentation. *Int. J. Comput. Assist. Radiol. Surg.* 2011;6(3):329–39.
- Warfield SimonK, Zou KellyH, Wells WilliamM. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans. Med. Imaging* 2004;23(7):903–21.
- Serag A, Kyriakopoulou V, Rutherford MA, Edwards AD, Hajnal JV, Aljabar P. A multi-channel 4D probabilistic atlas of the developing brain: application to fetuses and neonates. *Ann. Br. Mach. Vis. Assoc.* 2012;2012(3):1–14.
- Tustison NicholasJ, Avants BrianB, Cook PhilipA, Yuanjie Zheng, Alexander Egan, Yushkevich PaulA. N4ITK: improved N3 bias correction. *IEEE Trans. Med. Imaging* 2010;29(6):1310–20.
- Geert Litjens, Thijs Kooi, BabakEhteshami Bejnordi, ArnaudArindraAdiyoso Setio, Francesco Ciompi, Mohsen Ghahfarooi. A survey on deep learning in medical image analysis. *Med. Image Anal.* 2017;42:60–88.
- Zeynettin Akkus, Alfiia Galimzianova, Assaf Hoogi, Rubin Daniell, Erickson BradleyJ. Deep learning for brain MRI segmentation: state of the art and future directions. *J. Digit. Imaging* 2017;30(4):449–59.
- Antonios Makropoulos, Counsell SerenaJ, Daniel Rueckert. A review on automatic fetal and neonatal brain MRI segmentation. *NeuroImage* 2018;170:231–48.
- YuDong Zhang, Chichun Pan, Xianqing Chen, Fubin Wang. Abnormal breast identification by nine-layer convolutional neural network with parametric rectified linear unit and rank-based stochastic pooling. *J. Comput. Sci.* 2018;27:57–68.
- ShuiHua Wang, Lv YiDing, Yuxiu Sui, Shuai Liu, SuJing Wang, YuDong Zhang. Alcoholism detection by data augmentation and convolutional neural network with stochastic pooling. *J. Med. Syst.* 2018;42(1).
- Marine Bouyssi-Kobar, Plessis AdrÉJ du, Robert McCarter, Marie Brossard-Racine, Jonathan Murnick, Laura Tinkleman. Third trimester brain growth in preterm infants compared with in utero healthy fetuses. *Pediatrics* 2016;138(5).
- John Dobbins, Jean Sands. Quantitative growth and development of human brain. *Arch. Dis. Child.* 1973;48(10):757–67.
- Ivana Išgum, Benders ManonJNL, Brian Avants, MJorge Cardoso, Counsell SerenaJ, EldaFischi Gomez. Evaluation of automatic neonatal brain segmentation algorithms: the NeoBrainS12 challenge. *Med. Image Anal.* 2015;20(1):135–51.
- Pim Moeskops, Viergever MaxA, Mendrik AdrienneM, Vries LindaS de, Benders ManonJNL, Ivana Išgum. Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Transactions on Medical Imaging* 2016;35(5):1252–61.
- Jérémie Anquez, Angelini ElsaD, Isabelle Bloch. Automatic segmentation of head structures on fetal MRI. *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on.* 2009. p. 109–12.
- Martin Rajchl, Lee MatthewCH, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai. Deepcut: object segmentation from bounding box annotations using convolutional neural networks. *IEEE Trans. Med. Imaging* 2017;36(2):674–83.
- Khalili N, Moeskops P, Claessens NHP, Scherpenzeel S, Turk E, de Heus R. Automatic segmentation of the intracranial volume in fetal MR images. *Fetal, Infant and Ophthalmic Medical Image Analysis.* Cham: Springer International Publishing; 2017. p. 42–51.
- Salehi SSM, Hashemi SR, Velasco-Ann C, Oualam A, Estroff JA, Erdogmus D. Real-time automatic fetal brain extraction in fetal MRI by deep learning. 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI). 2018. p. 720–4.
- NeoBrainS12. <http://neobrain12.isi.uu.nl/reference.php>.
- Olaf Ronneberger, Philo Fischer, Thomas Brox. U-net: convolutional networks for biomedical image segmentation. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI).* Springer; 2015. p. 234–41.
- Sergey Ioffe, Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning (ICML).* 2015. p. 448–56.
- Kingma D, JBa Adam. A method for stochastic optimisation. *International Conference on Learning Representations (ICLR).* 2015.
- Timothy Dozat. Incorporating nesterov momentum into Adam. 4th International Conference on Learning Representations (ICLR), Workshop. 2016.
- François Chollet. *Keras*. <https://keras.io>; 2015.
- Smith StephenM. Fast robust automated brain extraction. *Hum. Brain Mapp.* 2002;17(3):143–55.
- Ahmed Serag, Manuel Blesa, Moore EmmaJ, Rozalia Pataky, Sparrow SarahA, Wilkinson AG. Accurate learning with few atlases (ALFA): an algorithm for MRI neonatal brain extraction and comparison with 11 publicly available methods. *Sci. Rep.* 2016;6:23470.