

# Distributionally Robust Segmentation of Abnormal Fetal Brain 3D MRI

Lucas Fidon<sup>1</sup>, Michael Aertsen<sup>2</sup>, Nada Mufti<sup>1,3,4</sup>, Thomas Deprest<sup>2</sup>,  
Doaa Emam<sup>4,6</sup>, Frédéric Guffens<sup>2</sup>, Ernst Schwartz<sup>5</sup>, Michael Ebner<sup>1</sup>,  
Daniela Prayer<sup>5</sup>, Gregor Kasprian<sup>5</sup>, Anna L. David<sup>3,4</sup>, Andrew Melbourne<sup>1</sup>,  
Sébastien Ourselin<sup>1</sup>, Jan Deprest<sup>2,3,4</sup>, Georg Langs<sup>5</sup>, and Tom Vercauteren<sup>1</sup>

<sup>1</sup> School of Biomedical Engineering & Imaging Sciences, King’s College London, UK

<sup>2</sup> Department of Radiology, University Hospitals Leuven, Belgium

<sup>3</sup> Institute for Women’s Health, University College London, UK

<sup>4</sup> Department of Obstetrics and Gynaecology, University Hospitals Leuven, Belgium

<sup>5</sup> Department of Biomedical Imaging and Image-guided Therapy Medical University  
of Vienna, Austria

<sup>6</sup> Department of Gynecology and Obstetrics, University Hospitals Tanta, Egypt

**Abstract.** The performance of deep neural networks typically increases with the number of training images. However, not all images have the same importance towards improved performance and robustness. In fetal brain MRI, abnormalities exacerbate the variability of the developing brain anatomy compared to non-pathological cases. A small number of abnormal cases, as is typically available in clinical datasets used for training, are unlikely to fairly represent the rich variability of abnormal developing brains. This leads machine learning systems trained by maximizing the average performance to be biased toward non-pathological cases. This problem was recently referred to as hidden stratification. To be suited for clinical use, automatic segmentation methods need to reliably achieve high-quality segmentation outcomes also for pathological cases. In this paper, we show that the state-of-the-art deep learning pipeline nnU-Net has difficulties to generalize to unseen abnormal cases. To mitigate this problem, we propose to train a deep neural network to minimize a percentile of the distribution of per-volume loss over the dataset. We show that this can be achieved by using Distributionally Robust Optimization (DRO). DRO automatically reweights the training samples with lower performance, encouraging nnU-Net to perform more consistently on all cases. We validated our approach using a dataset of 368 fetal brain T2w MRIs, including 124 MRIs of open spina bifida cases and 51 MRIs of cases with other severe abnormalities of brain development.

## 1 Introduction

The segmentation of fetal brain tissues in MRI is essential for the study of abnormal fetal brain developments [2]. Fetal brain structures segmentation could

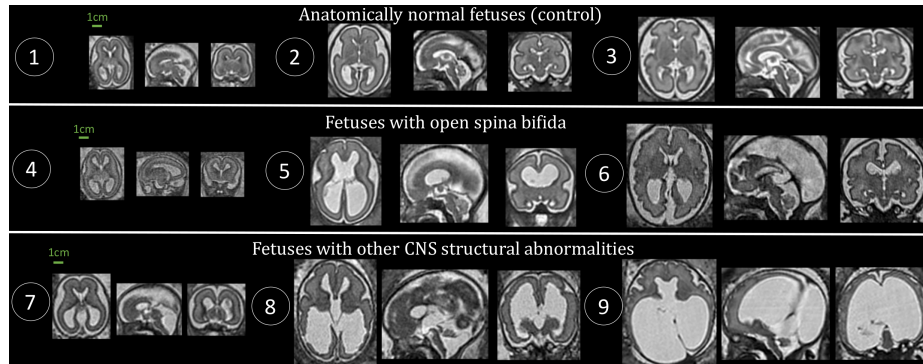


Fig. 1: Illustration of the anatomical variability in fetal brain across gestational ages and diagnostics. 1: Control (22 weeks); 2: Control (26 weeks); 3: Control (29 weeks); 4: Spina bifida (19 weeks); 5: Spina bifida (26 weeks); 6: Spina bifida (32 weeks); 7: Dandy-walker malformation with corpus callosum abnormality (23 weeks); 8: Dandy-walker malformation with ventriculomegaly and periventricular nodular heterotopia (27 weeks); 9: Aqueductal stenosis (34 weeks).

also support the evaluation and prediction of surgery outcome for open spina bifida [1,4,16,21,22]. Accurate and automatic methods for fetal brain segmentation are necessary as manual segmentation is very time-consuming and suffers from high inter- and intra-rater variability. Recently, deep neural network-based methods for fetal brain T2w MRI segmentation have been proposed [7,8,15,18,19]. On average, deep learning currently achieves state-of-the-art segmentation performance. However, those studies do not evaluate specifically the generalization and robustness properties when applied to fetuses with a pathological central nervous system.

Datasets used to train deep neural networks typically contain some underrepresented subsets of cases. These cases are not specifically dealt with by the training algorithms currently used for deep neural networks. This problem has been referred to as hidden stratification [17]. Hidden stratification has been shown to lead to deep learning models with good average performance but poor performance on some clinically relevant subsets of the population [17]. While uncovering the issue, the study of [17], which is limited to classification, does not study the cause or propose a method to mitigate this problem. Cases with abnormal fetal brain development are likely to suffer from hidden stratification effects for two reasons: 1) The presence of abnormalities exacerbates the anatomical variability of the fetal brain between 18 weeks and 38 weeks of gestation, as illustrated in Fig. 1; and 2) The prevalence of those diseases is typically below 1/1000 births [1].

In this work, we study the problem of hidden stratification in fetal brain MRI segmentation using deep learning. We claim that the methodology currently used to train deep neural networks, that is maximizing the average performance

across the training volumes, is at the root of the hidden stratification problem. Instead of the average empirical risk, training safe and robust deep learning models requires an asymmetric measure of risk that gives higher weights to the cases for which the algorithm fails (hard examples). Percentiles, also known as value-at-risk, is such a measure of risk that has even been adopted in industry regulations [13]. Given a per-volume fetal brain MRI segmentation metric such as the Dice score and an algorithm, the percentile at 5% is the value of the score below which 5% of the cases fall, i.e. perform worse than the percentile. The percentile relates to hidden stratification effects as it informs us of how badly worst-case examples are performing. Our contributions are four-fold. 1) We empirically show that the state-of-the-art deep learning pipeline nnU-Net [14] trained by maximizing the average segmentation performance leads to clinically significant failures for fetal brain MRI segmentation. 2) We propose to use percentiles of the Dice score on clinically relevant subpopulations as a measure of hidden stratification effects. 3) We propose to train a deep learning network to minimize a percentile of the per-volume loss function. 4) We propose a relaxation of this optimization problem based on distributionally robust optimization that can be solved efficiently in practice. We evaluate the proposed methodology for the automatic segmentation of white matter, ventricles, and cerebellum based on fetal brain 3D T2w MRI. We used a total of 368 fetal brain 3D MRIs including anatomically normal fetuses, fetuses with open spina bifida, and fetuses with other central nervous system pathologies for gestational ages ranging from 19 weeks to 39 weeks. Our empirical results suggests that the proposed training method based on distributionally robust optimization leads to better percentile values for abnormal fetuses. In addition, qualitative results shows that distributionally robust optimization allows to reduce the number of clinically relevant failures of nnU-Net.

## 2 Minimization of a Percentile Loss using Distributionally Robust Optimization

In this section, we study how a deep neural network can be trained to minimize percentiles of the loss function using a distributionally robust optimization (DRO) approach [10].

Standard deep learning training consists in optimizing the parameters  $\theta$  of a deep neural network  $f(\cdot; \theta)$  by minimizing the average per-example loss  $\mathcal{L}$

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(\mathbf{x}_i; \theta), \mathbf{y}_i) \quad (1)$$

Within this empirical risk minimization framework,  $f(\cdot; \theta)$  is typically a Convolutional Neural Network (CNN),  $\mathcal{L}$  is a smooth per-volume loss function, and  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$  is the training dataset.

In our case,  $\mathbf{x}_i$  are the input 3D fetal brain T2w MRI volumes and  $\mathbf{y}_i$  are the ground-truth manual segmentations. This approach is the one used to train

state-of-the-art deep learning methods for segmentation using stochastic gradient descent [14]. Due to the scarcity and the higher anatomical variability of abnormal cases illustrated in Fig. 1, we cannot assume that the set of all possible fetal brain anatomies is sampled uniformly in the training dataset. However, in (1), all brain volumes are given the same weight equal to  $\frac{1}{n}$ .

Instead of the average per-volume loss, for robust and safe segmentation, we argue that it might be more interesting to minimize the percentile  $l_\alpha$  at  $\alpha$  (e.g. 5%) of the per-volume loss function. Formally, this corresponds to the minimization problem

$$\min_{\boldsymbol{\theta}, l_\alpha} l_\alpha \quad \text{such that} \quad \mathbb{P}(\mathcal{L}(f(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y}) \geq l_\alpha) \leq \alpha \quad (2)$$

where  $\mathbb{P}$  is the empirical distribution defined by the training dataset. In other words, if  $\alpha = 0.05$ , the optimal  $l_\alpha^*(\boldsymbol{\theta})$  of (2) for a given value set of parameters  $\boldsymbol{\theta}$  is the value of the loss such that the per-volume loss function is worse than  $l_\alpha^*(\boldsymbol{\theta})$  5% of the time. As a result, training the deep neural network using (2) corresponds to minimizing the percentile of the per-volume loss function  $l_\alpha^*(\boldsymbol{\theta})$ .

Unfortunately, the minimization problem (2) cannot be solved directly using stochastic gradient descent to train a deep neural network. We now propose a tractable upper bound for  $l_\alpha^*(\boldsymbol{\theta})$  and show that it can be solved in practice using distributionally robust optimization [10].

The Chernoff bound [3] applied to the per-volume loss function and the empirical training data distribution states that for all  $l_\alpha$  and  $\beta > 0$

$$\mathbb{P}(\mathcal{L}(f(\mathbf{x}; \boldsymbol{\theta}), \mathbf{y}) \geq l_\alpha) \leq \frac{\exp(-\beta l_\alpha)}{n} \sum_{i=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i)) \quad (3)$$

To link this inequality to the minimization problem (2), we set  $\beta$  such that

$$\alpha = \frac{\exp(-\beta \hat{l}_\alpha(\boldsymbol{\theta}))}{n} \sum_{i=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i)) \quad (4)$$

$$\iff \hat{l}_\alpha(\boldsymbol{\theta}) = \frac{1}{\beta} \log \left( \frac{1}{\alpha n} \sum_{i=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i)) \right) \quad (5)$$

$\hat{l}_\alpha(\boldsymbol{\theta})$  is therefore an upper bound for  $l_\alpha^*(\boldsymbol{\theta})$ , independently to the value of  $\boldsymbol{\theta}$ . We propose to relax the minimization problem (2) by

$$\min_{\boldsymbol{\theta}} \frac{1}{\beta} \log \left( \sum_{i=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i)) \right) \quad (6)$$

where  $\beta > 0$  is a hyperparameter, and where the term  $\frac{1}{\beta} \log \left( \frac{1}{\alpha n} \right)$  was dropped as being independent of  $\boldsymbol{\theta}$ . While in (6),  $\alpha$  does not appear in the optimization problem directly anymore,  $\beta$  essentially acts as a substitute for  $\alpha$ . The higher the value of  $\beta$ , the higher weights the per-volume losses with a high value will have in (6).

Table 1: **Training and testing dataset details.** Other Abn: other brain structural abnormalities. There is no overlap of subjects between training and testing.

Train/Test	Origin	Condition	Volumes	Gestational age (in weeks)
Training	Atlas [12]	Control	18	[21, 38]
Training	FeTA [18]	Control	5	[22, 28]
Training	UHL and MUV	Control	116	[20, 35]
Training	UHL and MUV	Spina Bifida	28	[22, 34]
Training	UHL and MUV	Other Abn	10	[23, 35]
Testing	FeTA [18]	Control	28	[20, 34]
Testing	FeTA [18]	Spina Bifida	31	[22, 31]
Testing	FeTA [18]	Other Abn	16	[20, 34]
Testing	UHL and MUV	Control	26	[26, 37]
Testing	UHL and MUV	Spina Bifida	65	[19, 33]
Testing	UHL and MUV	Other Abn	25	[21, 40]

We give a proof in the supplementary material that (6) is equivalent to solving the distributionally robust optimization problem

$$\min_{\theta} \max_{\mathbf{q} \in \Delta_n} \left( \sum_{i=1}^n q_i \mathcal{L}(f(\mathbf{x}_i; \theta), \mathbf{y}_i) - \frac{1}{\beta} D_{KL} \left( \mathbf{q} \parallel \frac{1}{n} \mathbf{1} \right) \right) \quad (7)$$

where a new unknown probabilities vector parameter  $\mathbf{q}$  is introduced,  $\frac{1}{n} \mathbf{1}$  denotes the uniform probability vector  $(\frac{1}{n}, \dots, \frac{1}{n})$ ,  $D_{KL}$  is the Kullback-Leibler divergence,  $\Delta_n$  is the unit  $n$ -simplex, and  $\beta > 0$  is a hyperparameter.  $D_{KL}$  measures the dissimilarity between  $\mathbf{q}$  and the uniform probability vector  $\frac{1}{n} \mathbf{1}$  that corresponds to assign the same weight  $\frac{1}{n}$  to each sample. Therefore,  $\frac{1}{\beta}$  controls how much the samples with a relatively high loss value (hard examples) are weighted.

Recently, hardness weighted sampling [10] was introduced as a principled hard example mining method to solve (7). Here, we proved that it can be used to minimize the proposed relaxed minimization (6) of the percentile loss problem.

### 3 Anatomically Abnormal Fetal Brain T2w MRI Dataset

In this section, we give details about the fetal brain 3D MRI data, the labelling protocol, and the pre-processing used in our experiments.

**Public Fetal Brain Datasets** We used the 18 control fetal brain 3D MRI volumes of the spatio-temporal fetal brain atlas<sup>7</sup> [12] for gestational ages ranging from 21 weeks to 38 weeks. We also used 80 volumes from the publicly available

<sup>7</sup> [http://crl.med.harvard.edu/research/fetal\\_brain\\_atlas/](http://crl.med.harvard.edu/research/fetal_brain_atlas/)

FeTA MICCAI challenge dataset<sup>8</sup> [18]. For the 40 MIAL 3D MRIs, corrections of the segmentations were performed by authors MA, LF, and PD to reduce the variability against the published segmentation guidelines that was released with the FeTA dataset [18]. Those corrections were performed as part of our previous work [8] and are publicly available<sup>9</sup>. Brain masks for the FeTA data were obtained via affine registration using two fetal brain atlases<sup>10</sup> [11,12].

**Image Acquisition and Preprocessing for the Private Dataset** All images in the private dataset were part of routine clinical care and were acquired at UHL and MUV due to congenital malformations seen on ultrasound.

In total, 93 cases with open spina bifida, 35 cases with other central nervous system pathologies, and 142 cases with other malformations, though with normal brain, and referred as controls, were included. The gestational age at MRI ranged from 19 weeks to 40 weeks. We have started to make fetal brain T2w 3D MRIs publicly available<sup>11</sup>. For each study, at least three orthogonal T2-weighted HASTE series of the fetal brain were collected on a 1.5T scanner using an echo time of 133ms, a repetition time of 1000ms, with no slice overlap nor gap, pixel size 0.39mm to 1.48mm, and slice thickness 2.50mm to 4.40mm. A radiologist attended all the acquisitions for quality control.

The reconstructed fetal brain 3D MRIs were obtained using NiftyMIC [6] a state-of-the-art super resolution and reconstruction algorithm. The volumes were all reconstructed to a resolution of 0.8 mm isotropic and registered to a fetal brain atlas [12]. Our pre-processing improves the resolution, and removes motion between neighboring slices and motion artefacts present in the original 2D slices [6]. We used volumetric brain masks to mask the tissues outside the fetal brain. Those brain masks were obtained using the automatic segmentation method described in [6,20].

**Labelling Protocol.** The labelling protocol used for white matter, ventricles and cerebellum is the same as in [18]. The three tissue types were segmented for our private dataset by a trained obstetrician and medical students under the supervision of a paediatric radiologist specialized in fetal brain anatomy, who quality controlled and corrected all manual segmentations.

**Separation of the Data into Training and Testing** A summary of the number of fetal brain 3D MRIs used at training and testing for each central nervous system condition can be found in Table 1. The training dataset contains a total of 177 cases with a majority of 139 controls and only 38 abnormal cases which is typical in clinical datasets. Five controls from the FeTA dataset were added in the training dataset because we found in preliminary experiments that

<sup>8</sup> DOI: 10.7303/syn25649159

<sup>9</sup> DOI: 10.5281/zenodo.5148611

<sup>10</sup> DOI: 10.7303/syn25887675

<sup>11</sup> <https://www.cir.meduniwien.ac.at/research/fetal/>




Fig. 2: **Qualitative results.** a) Fetus with aqueductal stenosis (34 weeks). b) Fetus with open spina bifida (27 weeks). For those two cases, nnU-Net [14] misses completely the cerebellum and achieves poor segmentation for the white matter and the ventricles. Our nnU-Net-DRO achieves satisfactory segmentation for the cerebellum for the two cases, and for all tissue types for the aqueductal stenosis case.

nnU-Net [14] fails on most of the FeTA data at testing when it is trained using only data from UHL and MUV and the fetal brain atlas [12]. The testing dataset contains 193 volumes with a majority of abnormal cases which is necessary to cover the anatomical variability of abnormal cases in our evaluation.

## 4 Experiments

**Common Deep Learning Pipeline.** We used nnU-Net [14], a generic deep learning pipeline for medical image segmentation, that has been shown to outperform other deep learning pipelines on 23 public datasets without the need to tune the loss function or the deep neural network architecture. Specifically, we used nnU-Net version 2 in 3D-full-resolution mode which is the recommended mode for isotropic 3D MRI data. nnU-Net automatically splits the training data into 5 folds 80% training/20% validation used to train 5 networks for each method. The predicted class probability maps of the 5 models are averaged at inference to improve robustness [14]. We used NVIDIA Tesla V100 GPUs with 16GB of memory. Training each network took from 4 to 6 days.

**Specificities of Each Method.** The baseline consists in using nnU-Net [14] without any modification. Our method, nnU-Net-DRO, also uses nnU-Net. The only difference is that we changed the sampling strategy to use the hardness weighted sampler for DRO [10]. We used the default hyper-parameter values for the hardness weighted sampler, i.e.  $\beta = 100$  with importance sampling and clipping values  $w_{min} = 0.1$  and  $w_{max} = 10$  as described in [10]. No other values were tested. Our implementation of the nnU-Net-DRO training procedure is publicly available at <https://github.com/LucasFidon/HardnessWeightedSampler>. It provides an implementation of the hardness weighted sampler described in [10].

Table 2: **Evaluation of distribution robustness with respect to the pathology (193 3D MRIs).** **WM**: White matter, **Vent**: Ventricles, **Cer**: Cerebellum.  $\mathbf{p}_X$ :  $X^{\text{th}}$  percentile of the Dice score distribution in percentage. Best values are in bold.

Method	CNS	ROI	Dice Score (%)					
			Mean	Std	$\mathbf{p}_{50}$	$\mathbf{p}_{25}$	$\mathbf{p}_{10}$	$\mathbf{p}_5$
(baseline) nnU-Net	Controls (54 cases)	<b>WM</b>	<b>93.9</b>	<b>2.9</b>	<b>94.1</b>	<b>91.5</b>	<b>90.6</b>	<b>89.3</b>
		<b>Vent</b>	87.8	6.8	89.7	82.1	78.1	<b>76.8</b>
		<b>Cer</b>	<b>94.5</b>	3.2	<b>94.6</b>	92.4	<b>90.7</b>	<b>89.8</b>
	Spina Bifida (98 cases)	<b>WM</b>	89.9	7.9	92.5	89.1	79.9	73.4
		<b>Vent</b>	90.6	10.6	93.0	88.6	84.8	80.7
		<b>Cer</b>	78.2	28.7	<b>89.8</b>	<b>84.2</b>	13.9	<b>0.0</b>
	Other Abn. (41 cases)	<b>WM</b>	<b>90.3</b>	9.8	<b>92.7</b>	<b>89.7</b>	<b>82.7</b>	70.1
		<b>Vent</b>	87.1	7.3	87.1	82.5	77.7	75.2
		<b>Cer</b>	89.7	14.7	<b>92.8</b>	89.4	85.1	81.6
(ours) nnU-Net-DRO	Controls (54 cases)	<b>WM</b>	93.8	3.0	93.9	91.2	90.1	89.2
		<b>Vent</b>	<b>87.9</b>	<b>6.7</b>	<b>89.9</b>	<b>82.6</b>	<b>78.3</b>	76.7
		<b>Cer</b>	94.4	<b>3.1</b>	<b>94.6</b>	<b>92.6</b>	<b>90.7</b>	89.5
	Spina Bifida (98 cases)	<b>WM</b>	<b>90.3</b>	<b>7.5</b>	<b>92.9</b>	<b>89.2</b>	<b>81.5</b>	<b>73.7</b>
		<b>Vent</b>	<b>90.9</b>	<b>10.3</b>	<b>93.2</b>	<b>89.2</b>	<b>85.1</b>	<b>81.7</b>
		<b>Cer</b>	<b>79.7</b>	<b>27.6</b>	89.7	84.1	<b>40.4</b>	<b>0.0</b>
	Other Abn. (41 cases)	<b>WM</b>	<b>90.3</b>	<b>9.5</b>	92.5	89.6	82.5	<b>72.0</b>
		<b>Vent</b>	<b>87.5</b>	<b>7.1</b>	<b>87.5</b>	<b>82.7</b>	<b>80.4</b>	<b>76.7</b>
		<b>Cer</b>	<b>90.6</b>	<b>10.5</b>	<b>92.8</b>	<b>89.8</b>	<b>85.5</b>	<b>82.9</b>

**Evaluation Method.** We evaluate the quality of the automatic fetal brain MRI segmentations using the Dice score [5,9]. We are particularly interested in measuring the statistical risk of the results as a way to evaluate the robustness of the different methods. To this end, in addition to the mean and standard deviation, we also report the percentiles of the Dice score at 50%, 25%, 10%, and 5%. In Table 2, we report those quantities for the Dice scores of the three tissue types white matter, ventricular system, and cerebellum.

For each method, nnU-Net is trained 5 times using different train/validation splits and different random initializations. The 5 same splits, computed randomly, are used for the two methods. The results in Table 2 are for the ensemble of the 5 3D U-Nets. Ensembling is known to increase the robustness of deep learning methods for segmentation [14]. It also makes the evaluation less sensitive to the random initialization and to the stochastic optimization.

**Evaluation of nnU-Net and nnU-Net-DRO.** Quantitative evaluation of nnU-Net and nnU-Net-DRO for the three different central nervous system conditions control, spina bifida, and other abnormalities can be found in Table 2.



For spina bifida and other brain abnormalities, the proposed nnU-Net-DRO achieves same or higher mean Dice scores and lower standard deviations than nnU-Net [14] for the three tissue types. For controls, the mean Dice scores and standard deviation of nnU-Net-DRO and nnU-Net differ by less than 0.1 percentage points (pp) for the three tissue types.

The comparison of the percentiles of the Dice score allows us to compare methods at the tail of the Dice scores distribution where segmentation methods reach their worst-case performance. For spina bifida, nnU-Net-DRO achieves higher values of percentiles than nnU-Net for the white matter (+0.6pp for  $\mathbf{p}_{10}$ ), for the ventricular system (+1.0pp for  $\mathbf{p}_5$ ), and for the cerebellum (+26.5pp for  $\mathbf{p}_{10}$ ). And for other brain abnormalities, nnU-Net-DRO achieves higher values of percentiles than nnU-Net for the white matter (+1.9pp for  $\mathbf{p}_5$ ), for the ventricular system (+1.5pp for  $\mathbf{p}_5$  and +2.7pp for  $\mathbf{p}_{10}$ ), and for the cerebellum (+1.3pp for  $\mathbf{p}_5$ ). All the other percentile values differ by less than 0.5pp of Dice score between the two methods. This suggests that nnU-Net-DRO achieves better worst case performance than nnU-Net for abnormal cases.

It is worth noting that the Dice scores decrease for the white matter and the cerebellum between controls and spina bifida and abnormal cases. It was expected due to the higher anatomical variability in pathological cases. However, the Dice scores for the ventricular system tend to be higher for abnormal cases than for controls. This can be attributed to the large proportion of pathological cases with enlarged ventricles because the Dice score values tend to be higher for larger region of interests.

As can be seen in the qualitative results of Table 2, there are cases for which nnU-Net predicts an empty cerebellum segmentation while nnU-Net-DRO achieves satisfactory cerebellum segmentation. There were no cases for which the converse was true. Robust segmentation of the cerebellum for spina bifida is particularly relevant for the evaluation of fetal brain surgery for open spina bifida [1,4,21]. Additional qualitative results in the supplementary material illustrates 5 other cases for which nnU-Net-DRO outperforms nnU-Net.

## 5 Conclusion

The high anatomical variability of the developing fetal brain across gestational ages and pathologies hampers the robustness of deep neural networks trained by maximizing the average per-volume performance. Specifically, it limits the generalization of deep neural networks to abnormal cases for which few cases are available during training. In this paper, we propose to mitigate this problem by training deep neural networks to minimize a percentile of the per-volume performance rather than the average. To allow to do this in practice, we propose to train deep neural networks with Distributionally Robust Optimization (DRO) and we show that the DRO objective is a relaxation of the per-volume loss percentile. We have validated the proposed training method on a multi-centric dataset of 368 fetal brain T2w 3D MRIs with various diagnostics. nnU-Net trained with DRO achieved improved segmentation results for pathological

cases as compared to the unmodified nnU-Net, while achieving similar segmentation performance for the neurotypical cases. Our results suggest that nnU-Net trained with DRO is more robust to anatomical variabilities than the original nnU-Net.

**Acknowledgments** This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement TRABIT No 765148. This work was supported by core and project funding from the Wellcome [203148/Z/16/Z; 203145Z/16/Z; WT101957], and EPSRC [NS/A000049/1; NS/A000050/1; NS/A000027/1]. TV is supported by a Medtronic / RAEng Research Chair [RCSRF1819\7\34].

## References

1. Aertsen, M., Verduyckt, J., De Keyzer, F., Vercauteren, T., Van Calenbergh, F., De Catte, L., Dymarkowski, S., Demaerel, P., Deprest, J.: Reliability of MR imaging-based posterior fossa and brain stem measurements in open spinal dysraphism in the era of fetal surgery. *American Journal of Neuroradiology* **40**(1), 191–198 (2019)
2. Benkarim, O.M., Sanroma, G., Zimmer, V.A., Muñoz-Moreno, E., Hahner, N., Eixarch, E., Camara, O., Gonzalez Ballester, M.A., Piella, G.: Toward the automatic quantification of in utero brain development in 3D structural MRI: A review. *Human brain mapping* **38**(5), 2772–2787 (2017)
3. Chernoff, H., et al.: A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* **23**(4), 493–507 (1952)
4. Danzer, E., Joyeux, L., Flake, A.W., Deprest, J.: Fetal surgical intervention for myelomeningocele: lessons learned, outcomes, and future implications. *Developmental Medicine & Child Neurology* **62**(4), 417–425 (2020)
5. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
6. Ebner, M., Wang, G., Li, W., Aertsen, M., Patel, P.A., Aughwane, R., Melbourne, A., Doel, T., Dymarkowski, S., De Coppi, P., et al.: An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. *NeuroImage* **206**, 116324 (2020)
7. Fetit, A.E., Alansary, A., Cordero-Grande, L., Cupitt, J., Davidson, A.B., Edwards, A.D., Hajnal, J.V., Hughes, E., Kamnitsas, K., Kyriakopoulou, V., et al.: A deep learning approach to segmentation of the developing cortex in fetal brain mri with minimal manual labeling. In: *Medical Imaging with Deep Learning*. pp. 241–261. PMLR (2020)
8. Fidon, L., Aertsen, M., Emam, D., Mufti, N., Guffens, F., Deprest, T., Demaerel, P., David, A.L., Melbourne, A., Ourselin, S., et al.: Label-set loss functions for partial supervision: Application to fetal brain 3d mri parcellation. arXiv preprint arXiv:2107.03846 (2021)
9. Fidon, L., Li, W., Garcia-Peraza-Herrera, L.C., Ekanayake, J., Kitchen, N., Ourselin, S., Vercauteren, T.: Generalised wasserstein dice score for imbalanced multi-class segmentation using holistic convolutional networks. In: *International MICCAI Brainlesion Workshop*. pp. 64–76. Springer (2017)

10. Fidon, L., Ourselin, S., Vercauteren, T.: Distributionally robust deep learning using hardness weighted sampling. arXiv preprint arXiv:2001.02658 (2020)
11. Fidon, L., Viola, E., Mufti, N., David, A., Melbourne, A., Demaerel, P., Ourselin, S., Vercauteren, T., Deprest, J., Aertsen, M.: A spatio-temporal atlas of the developing fetal brain with spina bifida aperta. Open Research Europe (2021)
12. Gholipour, A., Rollins, C.K., Velasco-Annis, C., Ouaalam, A., Akhondi-Asl, A., Afacan, O., Ortinau, C.M., Clancy, S., Limperopoulos, C., Yang, E., et al.: A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific reports* **7**(1), 1–13 (2017)
13. Holton, G.: *Value at Risk: Theory and Practice*. Academic Press (2003)
14. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
15. Khalili, N., Lessmann, N., Turk, E., Claessens, N., de Heus, R., Kolk, T., Viergever, M., Benders, M., Išgum, I.: Automatic brain tissue segmentation in fetal MRI using convolutional neural networks. *Magnetic resonance imaging* **64**, 77–89 (2019)
16. Mufti, N., Aertsen, M., Ebner, M., Fidon, L., Patel, P., Rahman, M.B.A., Bracke-nier, Y., Ekart, G., Fernandez, V., Vercauteren, T., et al.: Cortical spectral matching and shape and volume analysis of the fetal brain pre-and post-fetal surgery for spina bifida: a retrospective study. *Neuroradiology* pp. 1–14 (2021)
17. Oakden-Rayner, L., Dunnmon, J., Carneiro, G., Ré, C.: Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. In: *Proceedings of the ACM conference on health, inference, and learning*. pp. 151–159 (2020)
18. Payette, K., de Dumast, P., Kebiri, H., Ezhov, I., Paetzold, J.C., Shit, S., Iqbal, A., Khan, R., Kottke, R., Grethen, P., et al.: An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific Data* **8**(1), 1–14 (2021)
19. Payette, K., Moehrlen, U., Mazzone, L., Ochsenbein-Kölbl, N., Tuura, R., Kottke, R., Meuli, M., Jakab, A.: Longitudinal analysis of fetal MRI in patients with prenatal spina bifida repair. In: *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis*, pp. 161–170. Springer (2019)
20. Ranzini, M., Fidon, L., Ourselin, S., Modat, M., Vercauteren, T.: MONAI-fbs: MONAI-based fetal brain MRI deep learning segmentation. arXiv preprint arXiv:2103.13314 (2021)
21. Sacco, A., Ushakov, F., Thompson, D., Peebles, D., Pandya, P., De Coppi, P., Wimalasundera, R., Attilakos, G., David, A.L., Deprest, J.: Fetal surgery for open spina bifida. *The Obstetrician & Gynaecologist* **21**(4), 271 (2019)
22. Zarutskie, A., Guimaraes, C., Yepez, M., Torres, P., Shetty, A., Sangi-Haghpeykar, H., Lee, W., Espinoza, J., Shamshirsaz, A., Nassr, A., et al.: Prenatal brain imaging for predicting need for postnatal hydrocephalus treatment in fetuses that had neural tube defect repair in utero. *Ultrasound in Obstetrics & Gynecology* **53**(3), 324–334 (2019)

— Supplemental Document —

## Distributionally Robust Segmentation of Abnormal Fetal Brain 3D MRI

Lucas Fidon,<sup>1</sup> Michael Aertsen,<sup>2</sup> Nada Mufti,<sup>1,3,4</sup> Thomas Deprest,<sup>2</sup>  
Doaa Emam,<sup>4,6</sup> Frédéric Guffens,<sup>2</sup> Ernst Schwartz,<sup>5</sup> Michael Ebner,<sup>1</sup>  
Daniela Prayer,<sup>5</sup> Gregor Kasprian,<sup>5</sup> Anna L. David,<sup>3,4</sup> Andrew Melbourne,<sup>1</sup>  
Sébastien Ourselin,<sup>1</sup> Jan Deprest,<sup>2,3,4</sup> Georg Langs,<sup>5</sup> and Tom Vercauteren<sup>1</sup>

<sup>1</sup>School of Biomedical Engineering & Imaging Sciences, King’s College London, UK

<sup>2</sup>Department of Radiology, University Hospitals Leuven, Belgium

<sup>3</sup>Institute for Women’s Health, University College London, UK

<sup>4</sup>Department of Obstetrics and Gynaecology, University Hospitals Leuven, Belgium

<sup>5</sup>Department of Biomedical Imaging and Image-guided Therapy Medical University of Vienna, Austria

<sup>6</sup>Department of Gynecology and Obstetrics, University Hospitals Tanta, Egypt

### 1 Proof of the equivalence of equations (6) and (7)

In the DRO optimization problem of equation (7), the optimal  $\mathbf{q}$  for any  $\boldsymbol{\theta}$  has the closed-form formula [10, see p.4 or Appendix 11.1]

$$\forall \boldsymbol{\theta}, \quad \mathbf{q}^*(\boldsymbol{\theta}) = \text{softmax}((\beta \mathcal{L}(h(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i))_{i=1}^n)$$

By injecting this in equation (7), we obtain

$$\begin{aligned} & \min_{\boldsymbol{\theta}} \max_{\mathbf{q} \in \Delta_n} \left( \sum_{i=1}^n q_i \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i) - \frac{1}{\beta} D_{KL} \left( \mathbf{q} \parallel \frac{1}{n} \mathbf{1} \right) \right) \\ & \min_{\boldsymbol{\theta}} \left( \sum_{i=1}^n q_i^*(\boldsymbol{\theta}) \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i) - \frac{1}{\beta} \sum_{i=1}^n q_i^*(\boldsymbol{\theta}) \log \left( \frac{\exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i))}{\frac{1}{n} \sum_{j=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_j; \boldsymbol{\theta}), \mathbf{y}_j))} \right) \right) \\ & \min_{\boldsymbol{\theta}} \left( \sum_{i=1}^n q_i^*(\boldsymbol{\theta}) \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i) - \sum_{i=1}^n q_i^*(\boldsymbol{\theta}) \frac{1}{\beta} \log(\exp(\beta \mathcal{L}(f(\mathbf{x}_i; \boldsymbol{\theta}), \mathbf{y}_i))) \right. \\ & \quad \left. + \frac{1}{\beta} \left( \sum_{i=1}^n q_i^*(\boldsymbol{\theta}) \right) \times \log \left( \frac{1}{n} \sum_{j=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_j; \boldsymbol{\theta}), \mathbf{y}_j)) \right) \right) \end{aligned}$$

Since the first two terms cancel each other and  $\sum_{i=1}^n q_i^*(\boldsymbol{\theta}) = 1$ , we obtain

$$\min_{\boldsymbol{\theta}} \frac{1}{\beta} \log \left( \sum_{j=1}^n \exp(\beta \mathcal{L}(f(\mathbf{x}_j; \boldsymbol{\theta}), \mathbf{y}_j)) \right) - \frac{1}{\beta} \log(n)$$

which is equivalent to the optimization problem (6) because the term  $\frac{1}{\beta} \log(n)$  above and the term  $\frac{1}{\beta} \log(\alpha n)$  in (6) are independent of  $\boldsymbol{\theta}$  ■

## 2 Additional qualitative results

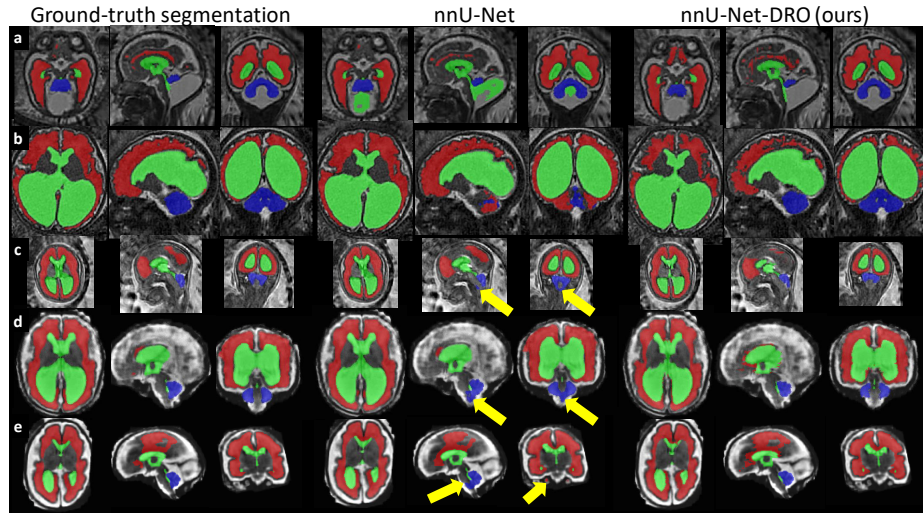


Fig. 1: **Additional qualitative results.** a) Fetus with Blake’s pouch cyst (29 weeks). b) Fetus with tuberous sclerosis complex (34 weeks). c) Fetus with open spina bifida (22 weeks). d) Fetus with open spina bifida (31 weeks). e) Fetus with open spina bifida (28 weeks). For case a), a large part of the Blake’s pouch cyst is wrongly included in the ventricular system segmentation by nnU-Net [14]. This is not the case for the proposed nnU-Net-DRO. For case b), nnU-Net fails to segment the cerebellum correctly and a large part of the cerebellum is segmented as part of the white matter. In contrast, our nnU-Net-DRO correctly segment cerebellum and white matter for this case. For cases c) d) and e), we have added yellow arrows pointing to large parts of the brainstem that nnU-Net wrongly included in the cerebellum segmentation. nnU-Net-DRO does not make this mistake. We emphasise that the segmentation of the cerebellum for open spina bifida is essential for studying and evaluating the effect of surgery in-utero.