



HAL
open science

Comprehensive analysis of synthetic learning applied to neonatal brain MRI segmentation

R Valabregue, F Girka, A Pron, F Rousseau, G Auzias

► **To cite this version:**

R Valabregue, F Girka, A Pron, F Rousseau, G Auzias. Comprehensive analysis of synthetic learning applied to neonatal brain MRI segmentation. 2023. hal-04199794

HAL Id: hal-04199794

<https://hal.science/hal-04199794>

Preprint submitted on 8 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Domain

Comprehensive analysis of synthetic learning applied to neonatal brain MRI segmentation

R. Valabregue^a, F. Girka^a, A. Pron^b, F. Rousseau^c, G. Auzias^b

^aCENIR, Institut du Cerveau (ICM) - Paris Brain Institute, Inserm U 1127, CNRS UMR 7225, Sorbonne Université, Paris, France

^bAix-Marseille Université, CNRS, Institut de Neurosciences de la Timone, UMR 7289, Marseille, France

^cIMT Atlantique, LaTIM INSERM U1101, Brest, France

Abstract

Brain segmentation from neonatal MRI images is a very challenging task due to large changes in the shape of cerebral structures and variations in signal intensities reflecting the gestational process. In this context, there is a clear need for segmentation techniques that are robust to variations in image contrast and to the spatial configuration of anatomical structures. In this work, we evaluate the potential of synthetic learning, a contrast-independent model trained using synthetic images generated from the ground truth labels of very few subjects.

We base our experiments on the dataset released by the developmental Human Connectome Project, for which high-quality T1- and T2-weighted images are available for more than 700 babies aged between 26 and 45 weeks post-conception. First, we confirm the impressive performance of a standard Unet trained on a few T2-weighted volumes, but also confirm that such models learn intensity-related features specific to the training domain. We then evaluate the synthetic learning approach and confirm its robustness to variations in image contrast by reporting the capacity of such a model to segment both T1- and T2-weighted images from the same individuals. However, we observe a clear influence of the age of the baby on the predictions. We improve the performance of this model by enriching the synthetic training set with realistic motion artifacts and over-segmentation of the white matter. Based on extensive visual assessment, we argue that the better performance of the model trained on real T2w data may be due to systematic errors in the ground truth. We propose an original experiment combining two definitions of the ground truth allowing us to show that learning from real data will reproduce any systematic bias from the training set, while synthetic models can avoid this limitation. Overall, our experiments confirm that synthetic learning is an effective solution for segmenting neonatal brain MRI. Our adapted synthetic learning approach combines key features that will be instrumental for large multi-site studies and clinical applications.

Keywords: deep learning, segmentation, newborn, brain, MRI, synthetic, dHCP

1. Introduction

1.1. Context

Automated segmentation of perinatal brain MRI remains a challenging task due to massive changes in the global shape of the brain and large variations in image intensity reflecting the rapid tissue maturation that occurs around birth [Kostović et al. \(2019\)](#). Several segmentation methods specifically designed for increased robustness to these factors such as multi-atlas label fusion techniques have been proposed [Makropoulos et al. \(2014\)](#), [Gholipour et al. \(2012\)](#), [Benkarim et al. \(2017\)](#); [Makropoulos et al. \(2018a\)](#); [Li et al. \(2019\)](#). Those methods have been applied to large open datasets like the developing Human Connectome Project (dHCP) [Makropoulos et al. \(2018b\)](#) enabling a better characterization of early brain development [Edwards et al. \(2022\)](#); [Dimitrova et al. \(2021\)](#).

More recently, supervised deep-learning techniques have been introduced as the next generation of segmentation techniques in medical images, showing higher performances and lower computing time than previous approaches. In particular, the UNet architecture [Ronneberger et al. \(2015\)](#) outperformed previous approaches in many different challenges [Isensee et al. \(2021\)](#). Nevertheless, a well-known limitation of supervised learning methods is their strong reduction in performances when applied

to unseen data [Karani et al. \(2018\)](#). This “domain-gap” problem has been identified as a major bottleneck in the field [Pan and Yang \(2010\)](#); [Zhou et al. \(2022\)](#) and an extensive body of literature investigated potential solutions and reported various gains in robustness, depending on the context. While our focus is not to review this large literature, we summarise the main approaches in order to better situate our strategy in the context of perinatal brain MRI segmentation.

A common approach consists in augmenting the training set with synthetic perturbations that explicitly control for the deviation from the initial training dataset [Ilse et al. \(2021\)](#). An obvious advantage is that it avoids the costly solution of getting more training data. In the context of brain development, two aspects of data augmentation corresponding to the two key challenges pointed above can be distinguished: 1) **spatial** augmentation to account for variations in the spatial arrangement of the different tissues and in the shape of specific anatomical structures (e.g., increase in cortical folding with age); 2) **style** (or appearance) augmentation to account for changes in tissue contrast, which can be induced either by variations in the acquisition settings or scanner or by variations related to brain maturation.

The design of those synthetic augmentations can either rely

on a **physics-based** (i.e., with a direct analytic model) or a **learning-based** generative model. Examples of physics-based augmentation strategies combine random affine or nonlinear deformations for spatial augmentation and random gamma perturbations for intensity augmentation [Zhang et al. \(2020\)](#); [Pérez-García et al. \(2021\)](#). The efficiency of this type of approach has been demonstrated in an intra-modality context [Zhao et al. \(2019\)](#); [Zhang et al. \(2020\)](#); [Isensee et al. \(2021\)](#), but it is less efficient for cross-modality [Karani et al. \(2018\)](#). The learning-based augmentation techniques are often referred to as “domain adaptation”. The recent works in this field have focused on the design of unsupervised learning approaches aiming at generating realistic synthetic training sets without requiring manually labeled data in the target external domain. Such techniques either learn a latent space that is common to the original domain where ground truth labels are available and to the target external domain [Kamnitsas et al. \(2017\)](#); [Ganin et al. \(2017\)](#); [Tomar et al. \(2022\)](#) or learn a direct image-to-image translation [Zhang et al. \(2018\)](#). These two approaches are combined in [Chen et al. \(2019\)](#). The use of adversarial generative models for domain adaptation has also been considered in [Chartsias et al. \(2018\)](#). Another key feature of synthetic augmentation techniques is the capacity to train these models using very few manually labeled training data. Indeed, one-shot learning studies propose to reduce the training data to only one template image with corresponding ground truth labels [Tomar et al. \(2022\)](#); [Zhao et al. \(2019\)](#). All those methods alleviate the need for time-consuming and expertise-demanding ground truth segmentation in the target domain since the training of the domain transfer model requires a pool of unlabeled data representative of the target domain. The major drawback of the learning-based approaches is the need to train a new model for any new domain. Of note, the very recent work [Ouyang et al. \(2022\)](#) proposes to leverage this limitation by generating a wide range of contrasts from a single domain dataset using augmentation techniques inspired by the different acquisition processes. Recently, [Billot et al.](#) introduced a method called SynthSeg [Billot et al. \(2023, 2020\)](#) that does not rely on any real MRI data during the training process. We refer to this type of model as “**synthesis-based**”. The key is to avoid the potential bias toward the domain of the training set by introducing a framework allowing to train the models without any real imaging data. A fully synthetic training dataset is generated from a set of real labels maps using physics-based generative models of the correspondence between label maps geometry and underlying intensity distributions. Under the assumption of homogeneous tissues, the image signal is sampled from a Gaussian distribution with different mean and variance for each tissue (MR data is indeed a mixture of Gaussian intensities). The generated signal is then enriched with additional commonly used random transformations (bias field, gaussian noise, and spatial deformation). The approach proposed by [Billot et al.](#) is based on the “domain randomization” concept [Tobin et al. \(2017\)](#); [Tremblay et al. \(2018\)](#), which postulates that the variations across real data from different domains need to be encompassed within the distribution of the generated synthetic data. Therefore, the setting of the parameters in the generative process is key, and the

design of the transforms has to generate large enough variations, without strong constraint on their biological relevance. More specifically in the context of the present work, robustness to variations in image intensity distributions can be favored by randomly sampling the mean and variance of the generated signal; while the robustness to variations in brain size and cortical folding magnitude can be induced by tuning the random deformations. In [Billot et al. \(2023, 2022\)](#), the authors validated their approach on highly heterogeneous data acquired on adults using various settings from clinical practice, demonstrating impressive robustness to challenging variations in image contrast and resolution. They reported higher segmentation accuracy and robustness compared to other methods of domain adaptation. In addition, the authors of [Billot et al. \(2023\)](#) investigated the influence of the size of the training set on the performances of SynthSeg and reported that only a few training examples are sufficient to converge towards its maximum accuracy on a population of adults. In the present work, we assess whether these conceptually appealing features and impressive results on the adult population extend to the context of neonatal brain MRI segmentation.

1.2. Contributions

Our work focuses on a comprehensive analysis of synthetic learning approaches for segmenting neonatal brain MRI data. To this end, we first reimplemented the SynthSeg model [Billot et al. \(2023\)](#) within the Pytorch framework relying on the torchio transformations [Pérez-García et al. \(2021\)](#). The code is made available ¹. In contrast to [Billot et al. \(2023\)](#) who demonstrated the robustness of SynthSeg to variations in image resolution and contrast on a very large clinical dataset, we focus here on the potential advantages when applied to perinatal brain MRI, in comparison to a classical UNet trained on real T2w data using a few shot learning strategy. Since the SynthSeg model did not perform as well as expected on neonatal brain MRI data, we propose two solutions to address its limitations, yielding better performances: adding simulated motion augmentation or subdividing the WM tissue into several sub-compartments.

Using our improved synthesis-based model, we then confirm the robustness of the predictions to variations in the contrast of the images, with very consistent predictions from either T1w or T2w images from the same subjects. We also demonstrate another key advantage of the synthetic learning approach: the synthesis-based models learn an unbiased correspondence between the geometry of the labels and image intensities. To quantitatively support this feature, we build a second type of ground truth from the same dataset. We report a much lower influence of the definition of the ground truth on the predictions from the synthesis-based models, compared to a model learned on real MRI data, which reproduces any systematic bias from the ground truth.

The quantitative evaluations are complemented with a careful

¹<https://github.com/romainVala/torchQC/tree/master/segmentation>

visual assessment of the predictions and ground truth. This allows us to better interpret our results, but also to report and discuss the limitations of the dHCP data and segmentation.

2. Material and methods

2.1. Neonatal MRI data and ground truth segmentation

In this work, we evaluate the performance of the SynthSeg approach [Billot et al. \(2023\)](#) on neonatal data using the third release of the publicly available developing Human Connectome Project (dHCP) dataset (<http://www.developingconnectome.org/>) [Edwards et al. \(2022\)](#). The dHCP dataset contains high-quality anatomical MRI scans of 885 neonates (age range from 26 to 45 weeks post-conception) acquired with both T1-weighted (T1w) and T2-weighted (T2w) sequences with a 0.5 mm isotropic resolution on a 3T Philips scanner (see [Edwards et al. \(2022\)](#) for further information about acquisitions).

2.1.1. Ground truth based on volumetric segmentation: *GT_drawEM*

We first used the segmentation in 9 tissues provided by the dHCP consortium [Makropoulos et al. \(2018b\)](#) (**CSF** (Cerebrospinal fluid), **GM** (Cortical gray matter), **WM** (White matter), Background, Ventricles, **Cereb** (Cerebellum), **deepGM** (deep Gray Matter), **Bstem** (Brainstem), **HipAmy** (Hippocampi+Amygdala)). The segmentation is based on the multi-atlas method drawEM [Makropoulos et al. \(2014\)](#) applied to the T2w data. As mentioned by the authors, drawEM is very robust and efficient in most cases but may fail to capture the highly complex shape of the cortical geometry. Extensive quality control was performed prior to the first release of the data, but localized inaccuracies remain. For instance, the authors reported that entire folds may be excluded from the automatic segmentation in 2% of cases. As a consequence, it is important to remind throughout this study (and other works focusing on segmentation using this dataset) that the segmentations provided should be considered as *pseudo* ground truth, although the term “ground truth” is used for simplicity. In this work, we merged the CSF and Ventricle labels into a single class (only for the evaluation) in order to avoid potential perturbations in the performance related to the tedious delineation between these two labels with similar intensity distributions. We refer to these ground truth segmentation maps as **GT_drawEM**.

2.1.2. Ground truth based on surface reconstruction: *GT_surf*

We derived a second type of ground truth from the same images, based on the internal (white) and external (pial) cortical boundaries, represented as surfaces. Both surfaces are provided by the dHCP and are computed using the surface deformation tool introduced in [Schuh et al. \(2017\)](#). Briefly, white-matter (internal) surface extraction is performed by fitting a closed, genus-0, triangulated surface mesh onto the segmentation boundary under constraints incorporating intensity information from the T2w, as well as controlling for surface topology and smoothness. The external (pial) surface is then obtained by deforming the internal surface outwards in order to fit

the tissue boundaries [Makropoulos et al. \(2018b\)](#); [Schuh et al. \(2017\)](#). From these internal and external surfaces surrounding the cortical tissue, we compute partial volume maps of the GM on the same 3D voxel grid as the T2w volume using a surface-based approach [Kirk et al. \(2020\)](#). We then obtain a binary segmentation of the GM by applying a threshold of 0.5. Finally, this different segmentation map for the GM is incorporated into the DrawEM label maps by replacing the original GM label and propagating the adjacent labels to preserve their topology. Compared to the original segmentation maps *GT_drawEM*, all the structures remain identical except WM, GM, and CSF. We denote this second segmentation map as **GT_surf**.

We illustrate these two types of ground truth segmentation maps in [Fig. 1](#), with a plot showing the GM volume ratio (*GT_drawEM*/*GT_surf*) by subject, ordered by age. While the differences might look subtle visually on a single slice, we measured an average 25% increase of GM volume in the *GT_drawEM* compared to *GT_surf*. This ratio is not influenced by the age of the baby. In this study, we use these two different, but both anatomically plausible, pseudo-ground truths to assess the influence of the definition of the segmentation map on the predictions of the models.

2.1.3. Head label

The background label from drawEM maps contains only a thin layer surrounding the CSF, as the segmentations are computed on a brain-masked volume [Makropoulos et al. \(2018b\)](#). Using only these labels for the generative process would limit the application to segment only skull-stripped data. We add other head tissue labels using the MIDA template [Iacono et al. \(2015\)](#), which contains 153 labels segmented from an adult MRI. We extract the extra-brain labels, which we grouped into 9 classes (dura mater, air, eyes, mucosa, muscle, nerves, skin, skull, and vessel) plus the background. These labels are combined with the 9 labels of the drawem9_dseg volume after registering the MIDA template to each subject using the FIRST method from FSL [Jenkinson and Smith \(2001\)](#) for the affine part, and `reg_f3d` from NiftyReg [Modat et al. \(2010\)](#) for the nonlinear part. We perform the label fusion to keep the original labels within the brain unchanged, and missing voxels outside the brain are set to air tissue. All these labels were used for the synthetic data generation but were then grouped into a single class (head) for the target objective.

2.2. Generative model

The key idea of the SynthSeg approach proposed in [Billot et al. \(2023\)](#) consists in generating the entire training set as synthetic images from 3D segmentation labels, meaning that no real image is used for training. This is based on the assumption that the MR signal is homogeneous within each label. We implemented different transforms for simulating variations in tissue intensity, shape variability, and MRI artifacts (bias and noise) using a generative model detailed below. We further enriched the generative model from [Billot et al. \(2023\)](#) by adding simulated motion artifacts and white matter inhomogeneity. The generative model was implemented using PyTorch

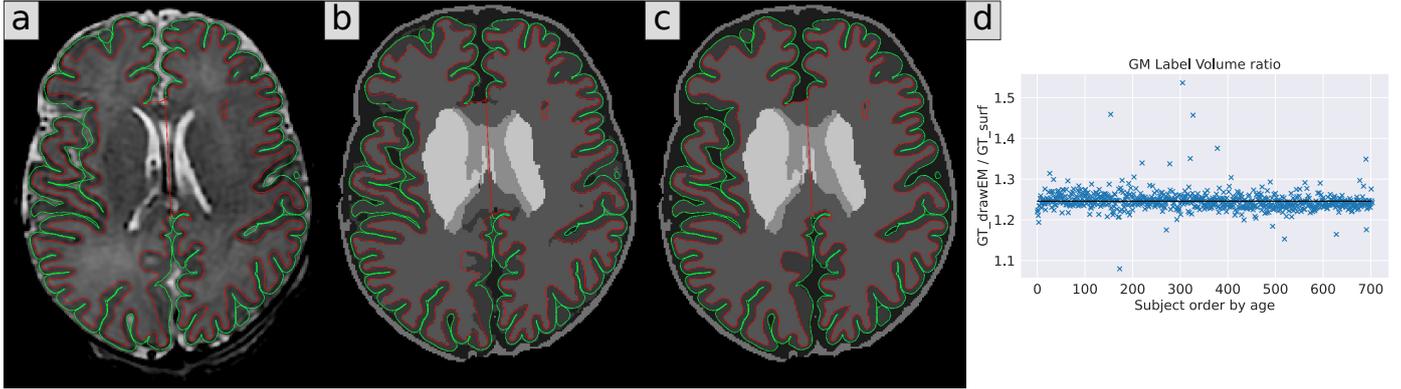


Figure 1: Illustration of the two different ground truths used in our study. We overlay in red and green the internal and external GM surfaces resp. a) original T2w; b) label map from GT_drawEM; c) label map from GT_surf. d) Volume of GM from GT_drawEM divided by the volume of GM from GT_surf for each subject, ordered by age. The black line represents the mean value of 1.25, illustrating a 25% increase of GM volume in GT_drawEM compared to GT_surf.

and TorchIO Pérez-García et al. (2021).

Random contrast. The first step consists in generating an MRI volume from the label set by sampling the intensity of the voxels of each tissue from different Gaussian distributions, resulting in a synthetic MRI with a random contrast. The mean and the standard deviation of the Gaussian distribution of each tissue are sampled independently from uniform distributions $U[0,1]$ and $U[0.02, 0.1]$, respectively as in Billot et al. (2023) ($U[a, b]$ is the uniform distribution in the interval $[a, b]$). We implemented this process within TorchIO with the RandomLabelsToImage transform.

Shape variability. To generate synthetic MRIs with variations in brain anatomy from a limited number of subjects, we apply affine and non-linear deformations to the label set with nearest-neighbors interpolation. We use a composition of the following transforms from TorchIO: RandomAffine (scaling factor $U[0.9, 1.1]$ rotation $U[-20^\circ, 20^\circ]$, translation $U[-10\text{mm}, 10\text{mm}]$), and RandomElasticDeformation (12 control points and a max displacement of 8mm).

MRI artifacts. We further augment the synthetic dataset by adding an intensity bias field and a global Gaussian noise. We use RandomBiasField, which simulates spatial intensity inhomogeneity with a polynomial function of order 3 and a maximum magnitude of 0.5, and RandomNoise which add a Gaussian random noise with 0 mean and a standard deviation sampled from $U[5e-3, 0.1]$.

Motion simulation (Mot). We then extend the data augmentation beyond Billot et al. (2023) by adding a RandomMotion transform to simulate subject motion during the MRI acquisition. We use our own implementation of the motion simulation introduced in Reguig et al. (2022), which allows us to simulate a realistic time course of rigid head motion. As shown in Fig. 2, the motion simulation induces inhomogeneities in the different tissues because of the tissue mixing in the k-space induced by motion. We use a maximum displacement sampled from $U[3, 8]$ (in mm) for the translation and $U[3, 8]$ (degrees) for the rotation. Note that we force the background signal intensity to zero when generating motion artifacts to avoid mixing background intensity with the motion process.

White Matter inhomogeneities (Inh). Inhomogeneities in the WM tissue are expected during the developmental period covered, and are visible in the MRI data. We adapted the method proposed by Billot Billot et al. (2023) to account for such variations within a label: we subdivide WM into smaller sub-regions by clustering the T2w intensities, within the WM mask, using the Expectation Maximization algorithm Dempster et al. (1977). We choose N regions ($N \subseteq [2, 3, 4, 5, 6]$) in order to represent the inhomogeneities with different levels of granularity. Each subregion is then considered as a distinct tissue in the generative model (thus with a different random intensity) but they are regrouped for the segmentation objective in order to predict the whole WM. (see Fig. 2 c). Note that the term ‘transform’ is used for simplicity but is not adapted here since it is only a fixed modification of the input labels. Note also that real T2w data is used only for the generation of the label maps, but no real data is used for generating the synthetic images.

Finally, an intensity normalization is performed to scale the min and max signal intensity between 0 and 1 for each synthetically generated dataset. This generative model is used to produce synthetic training sets based on the same ground truth segmentation maps for the following 4 synthesis-based models:

Synth: SynthSeg method (same as Billot) with the following augmentation: random contrast, shape variability (affine and non-linear) and MRI artifacts (intensity Bias and noise)

SynthMot: Synth enriched with motion simulation with a probability of 0.5

SynthInh: Synth with extra labels within the white matter to simulate inhomogeneities (with a probability of 0.5)

SynthMotInh: Synth with a combination of both WM inhomogeneities and simulated motion augmentations with a probability of 0.5 each

DataT2 (baseline): The performance of the 4 models based on synthetic training sets are compared with a baseline model defined as a UNet trained on real dHCP T2w acquisitions from the same 15 subjects. We apply the same data augmentation as for the synthesis-based models except for random contrast and Motion. We also add a random gamma augmentation to

simulate slight variations in the intensity distribution

2.3. Training and backbone architecture of the models

Training and testing sets. The final sample of data used in this work was composed by selecting the images from the 709 scanning sessions of the dHCP with both T1w and T2w available among the 885 scanning sessions. 5 were excluded due to failure to generate the GT_surf ground truth segmentation maps. Among the 176 sessions for which only the T2w was available (without the T1w acquisition), we selected 15 sessions uniformly distributed across the entire age range to generate the synthetic training sets for the models. In total, the data from 719 subjects from the dHCP data were used in this study: 15 for the training set and 704 subjects for the test set. We used the unprocessed T1w and T2w (no brain mask or bias field correction).

Network architecture. The network architecture used for all methods was the well-established 3D Unet architecture [Ronneberger et al. \(2015\)](#) with residual skip connections. We used five levels, each separated with either a max-pooling for the encoder path or an upsampling operation for the decoder part. All levels contained 3 convolution layers, with $3 \times 3 \times 3$ kernels. Every convolutional layer was followed by a batch normalization, a ReLU activation function, and a 10% dropout layer, except for the last one, which was only followed by a softmax. The first block contained 24 feature maps and this number was doubled after each max-pooling and halved after each upsampling. This led to a total of 21,6 million parameters.

Training. All the models were trained with patches of size 128^3 . For each generated volume, we randomly selected 8 patches, sampled from a uniform distribution with the same probability of containing each structure. All models were trained with the average dice loss. Note that thanks to the fully convolutional nature of the Unet architecture, the inference was performed on the entire volume at native resolution. We used a batch size of 4 and the Adam optimizer with a learning rate of $1e-4$. The training was stopped after 240 000 iterations. The training of each model took 6 days on an NVIDIA tesla V100 GPU (<http://www.idris.fr/>).

2.4. Quantitative measures and qualitative evaluation

We report the binary dice, which is commonly used for segmentation evaluation, defined as $dice = 1 - 2 * (X * Y) / (X^2 + Y^2)$, where X is the binary prediction of a given tissue and Y is the ground truth label (already binarized). We also computed the average surface distance from MONAI [Cardoso et al. \(2022\)](#), but do not report this measure since it is fully consistent with the Dice score. We report the distribution of the Dice score across individuals separately for the different labels, as well as the distribution of the average of the Dice across all labels. In order to assess the potential effect of the age of the babies on the predictions, we also report the distribution of the Dice averaged over all structures computed in four age groups: 29 subjects in [26, 32[; 96 in [32, 36[; 183 in [36, 40[; 394 in [40 45]. To assess the robustness of the prediction to changes in image contrast, we computed for each tissue type the Pearson correlation

between the volumes obtained from the predicted segmentation from either the T1w or T2w images, across the 704 subjects of the test set. An ideal, fully contrast-independent segmentation technique would produce almost identical segmentations from either T1w and T2w images and thus get a correlation value close to 1.

Visual assessment is critical to complement quantitative measures and better interpret the results of segmentation tools but is time-consuming and expertise-demanding. As a tradeoff, we focused our visual assessment on the GM, which is the most challenging anatomical structure to segment, and thus appropriate for assessing the variations in performances across the methods. We describe and illustrate our observations in combination with the quantitative measures for each of our experiments in the next section.

3. Experiments and Results

We designed three different experiments in order to address the following questions: 1) What are the performances of SynthSeg and our enriched versions compared to training on real data? 2) Is the high robustness with respect to variations in image contrast reported in [Billot et al. \(2023\)](#) confirmed on neonatal MRI data? 3) How does the definition of the ground truth segmentation maps affect the performances? For each experiment, we provide both quantitative and qualitative assessments allowing us to interpret potential variations in the performances across the models. This extensive visual assessment enabled us to identify different types of limitations in the segmentation provided by the dHCP. We report in section 3.4 our observations that we believe are important for future studies on this widely used dataset.

3.1. Experiment #1: evaluation of synthesis-based approaches on dHCP T2w dataset

The aim of this experiment was to assess the performances of synthesis-based methods on neonatal brain MRI data and compare them to a learning strategy on real data in the absence of domain shift and with high-quality data. To this end, we used the following experimental setup:

- Training set: 15 subjects, ground truth = GT_DrawEM, data used for DataT2: T2w
- Testing set: 704 subjects, ground truth = GT_DrawEM, prediction for all methods on T2w

As can be seen in all of the plots of Fig. 3, the performance of the methods are ranked in a consistent order across all structures and across age groups. DataT2 is performing best with an average dice of 96 (and above 95 for all structures except CSF and Hip-Amy with respectively 94 and 92). The performance of the Synth model proposed by [Billot et al. \(2023\)](#) is lower than the DataT2 model by 9 dice points on average (87). The differences are smaller (between 5 and 6 points) for WM/bstem/cereb/deepGM, but larger (10 points) for GM and Hip-Amy and even 20 points for CSF+ventricle.

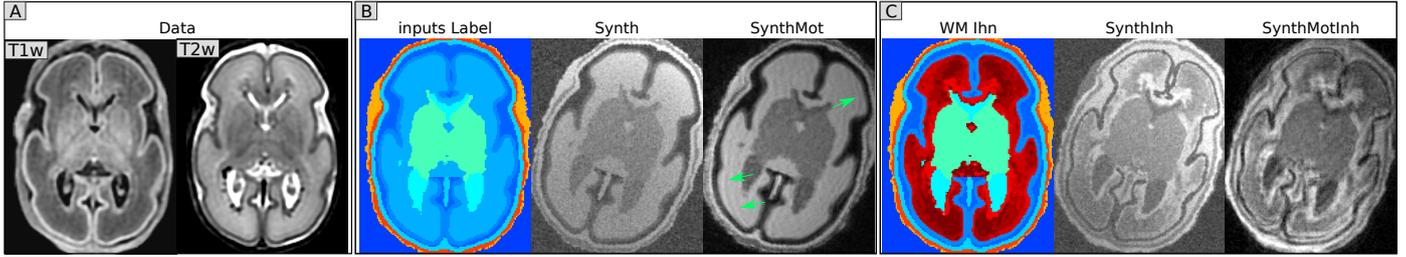


Figure 2: Illustration of the synthetic datasets obtained from one individual data. A) the T1w and T2w MRI data from this subject; B) and C) show respectively the 2 different label maps (in color) used as input of the generative model (without and with additional labels in the WM to simulate inhomogeneities). For each label map, we show an illustrative example of augmented synthetic images, corresponding to the 4 synthesis-based models. Green arrows indicate subtle artifacts induced by motion simulation.

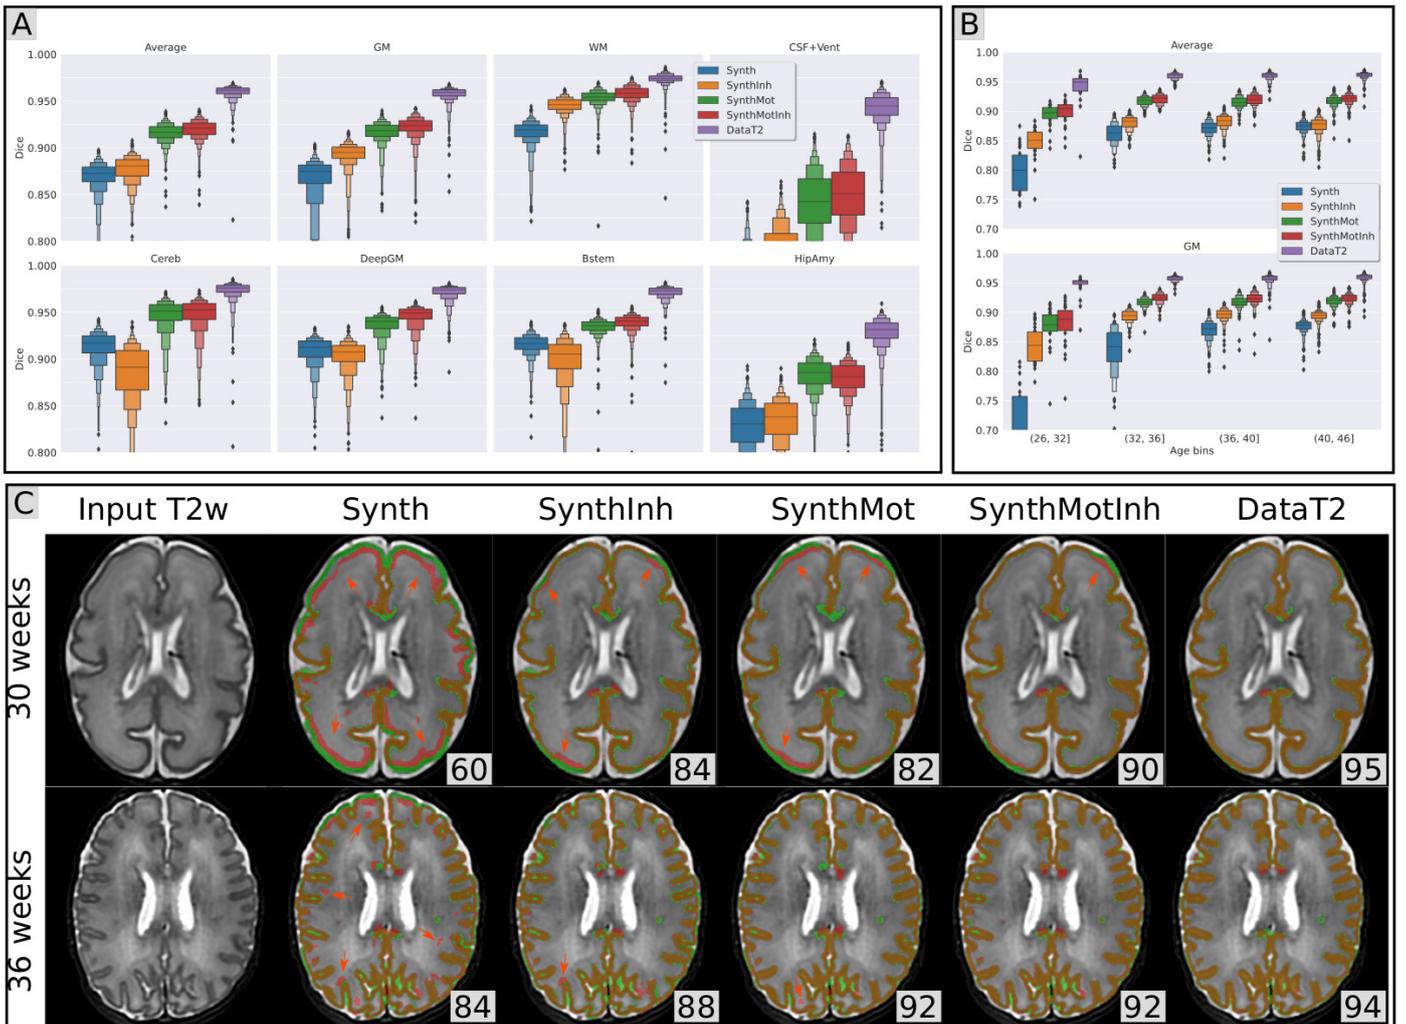


Figure 3: A) Distribution across the 704 subjects of the test set of the Dice score for each structure, for the 5 models. B) Distribution of the Dice averaged across all structures, computed in four age groups: 29 subjects in [26, 32]; 96 in [32, 36]; 183 in [36, 40]; 394 in [40, 45]. C) illustrations of the predicted GM for 2 subjects with different cortical folding magnitudes related to their age: sub-CC00657XX14 (30 weeks) and sub-CC00570XX10 (36 weeks). The numbers correspond to the dice score computed for the slice shown. Ground truth label (GT_drawEM) is shown in green, and the predictions in red. Red arrows indicate local errors in the predictions

Adding the motion augmentation greatly improves the performance, for all structures and age ranges. The difference between the SynthMot model and the DataT2 model is reduced by a factor of 2 for all structures, with an average dice of 92 (only 4 points difference with DataT2). Adding white matter

inhomogeneity in the SynthIhn model is also beneficial compared to the Synth, but the gain is mitigated: we observed an improvement for GM but not for Cereb DeepGM and Bsteam. On the other hand the SynthMotIhn model with both, motion and WM inhomogeneity, performs best (after dataT2) with a

slight improvement compared to SynthMot. Regarding the effect of age, we observe a drop in performance for all methods for the younger group (below 32 weeks). While the performance loss for DataT2 is limited, the Synth model shows the largest decrease in performance related to age, with a drop of 9 points. Adding motion augmentation and white matter inhomogeneities in SynthMotInh clearly mitigated this drop in the performance of the synthesis-based approach. The visual assessment showed that results are very consistent among subjects, with noticeable differences across methods on the first age bin. As illustrated in Panel C of Fig. 3, the predicted GM from the Synth model shows large errors with shifts of the GM prediction within the WM for the younger subjects. Those errors are largely fixed with the enriched synthesis-based models (SynthInh / SynthMot / SynthMotInh). The loss in performance of the synthesis-based models (even SynthMotInh) compared to DataT2 observed on Panel A and B corresponds to much more localized but clear errors. For older subjects, the types of errors are different. The errors of the Synth model are restricted only to regions where marked signal inhomogeneities are present in the WM. Here also, those local errors are largely fixed with other synthetic alternatives. The lower dice scores compared to DataT2 are mainly due to subtle errors along the boundary between GM and WM or CSF. Careful visual inspection showed that the SynthMotInh prediction better follows image contrast than the GT_drawEM. Regarding the model DataT2 (evaluated on T2w in this experiment) is performing almost perfectly at all ages: the predictions strictly follow the ground truth. The slight differences with GT_drawEM are due to local errors of the ground truth, even for the first age bin. Overall, our observations are:

- DataT2 is highly accurate at every age even when trained on only 15 subjects.
- Synth does not perform well, especially on younger subjects, with large regions of GM shifted within the WM.
- The enriched synthesis-based models enable to fix most of the errors but local errors still occur especially for the younger subjects.
- For older subjects, the predictions from SynthMotInh are visually accurate and better follow the underlying image contrast than the GT_drawEM.
- DataT2 does not make any obvious error (except for one subject), it reproduces the same tissue boundary as the ground truth and seems more robust to image noise than GT_drawEM. We further examine the anatomical relevance of the predictions relative to GT_drawEM in section 3.4 below.

3.2. Experiment #2: robustness to variations in image contrast

In the second experiment, we used the same trained model and changed the evaluation. Our aim was to assess the robustness of the models with respect to variations in the contrast of the image of the test set relative to the training set. We took the T1w image from the same individuals as an extreme change

in image contrast relative to the T2w. We used the following experimental setup:

- Training set: same as Exp.#1 (15 subjects, ground truth = **GT_DrawEM**, data used for DataT2: T2w)
- Testing set: 704 subjects, ground truth = GT_DrawEM, prediction for all methods on **T1w**

We show on Fig. 4 detailed results for the best synthesis-based model from Experiment #1 SynthMotInh and report the results for all the methods in a supplementary csv file². All panels show that DataT2, as expected, failed to predict on T1w inputs, with an average Dice below 10. The histograms within the GM shown on Panel E confirm that the dataT2 model learned the correspondence between the labels and the intensity distribution: the predicted GM from the T1w image (in blue) corresponds to voxels that have the same intensity range as the intensity in the GM from T2w image. On the contrary, the histograms from Panel D show that the SynthMotInh model gives very consistent segmentations from both T2w and T1w inputs. Indeed, the intensity distributions within the predicted GM from T1w and T2w images are very different and match the ground truth distribution well in both cases. Panel B shows the strong linear correlation between the volumes of the GM obtained from the SynthMotInh predictions from the T1w and T2w images across the 704 subjects of the test set. The Pearson correlation is above 0.99. This plot also shows a slight deviation of the data compared to $y=x$, indicative of a slightly larger estimated GM volume on T1w compared to T2w (5% on average).

The visual assessment was critical for this experiment, as illustrated in Panel C. First, we observe that the Synth model suffers from the same limitations as in Exp #1. Regarding the SynthMotInh model, we observe that the predictions on T1w are visually as good as the one from T2w, which is consistent with the high correlation shown on Panel B, but inconsistent with the drop of 4 points of Dice on average across all structures shown on Panel A. More specifically, robustness is excellent for DeepGM, Bstem, Cereb, and Hip-Amy with a decrease in performance of only 0.7, 1.7, 2.4, and 2 points of Dice respectively. The Dice is however lower for GM, WM, and CSF+Ventricle (loss of 5.4, 4, and 10 resp.). The same trends are observed for the other synthesis-based models.

Careful visual assessment enabled us to observe that the disagreement between prediction and GT is mostly due to residual misregistration between T1w and T2w images from the dHCP dataset. This is visible on Panel C for SynthMotInh, with a slight shift in the location of GM especially in the left posterior region of the brain. The impact of such residual misregistration on the Dice scores is stronger for external tissues (GM, WM, and CSF+Ventricles) than for deep structures, as observed in Panel A. We report further observations regarding the impact of misregistration on our evaluation in subsection 3.4. Overall, our observations are:

²https://github.com/romainVala/Synthetic_learning_on_dHCP

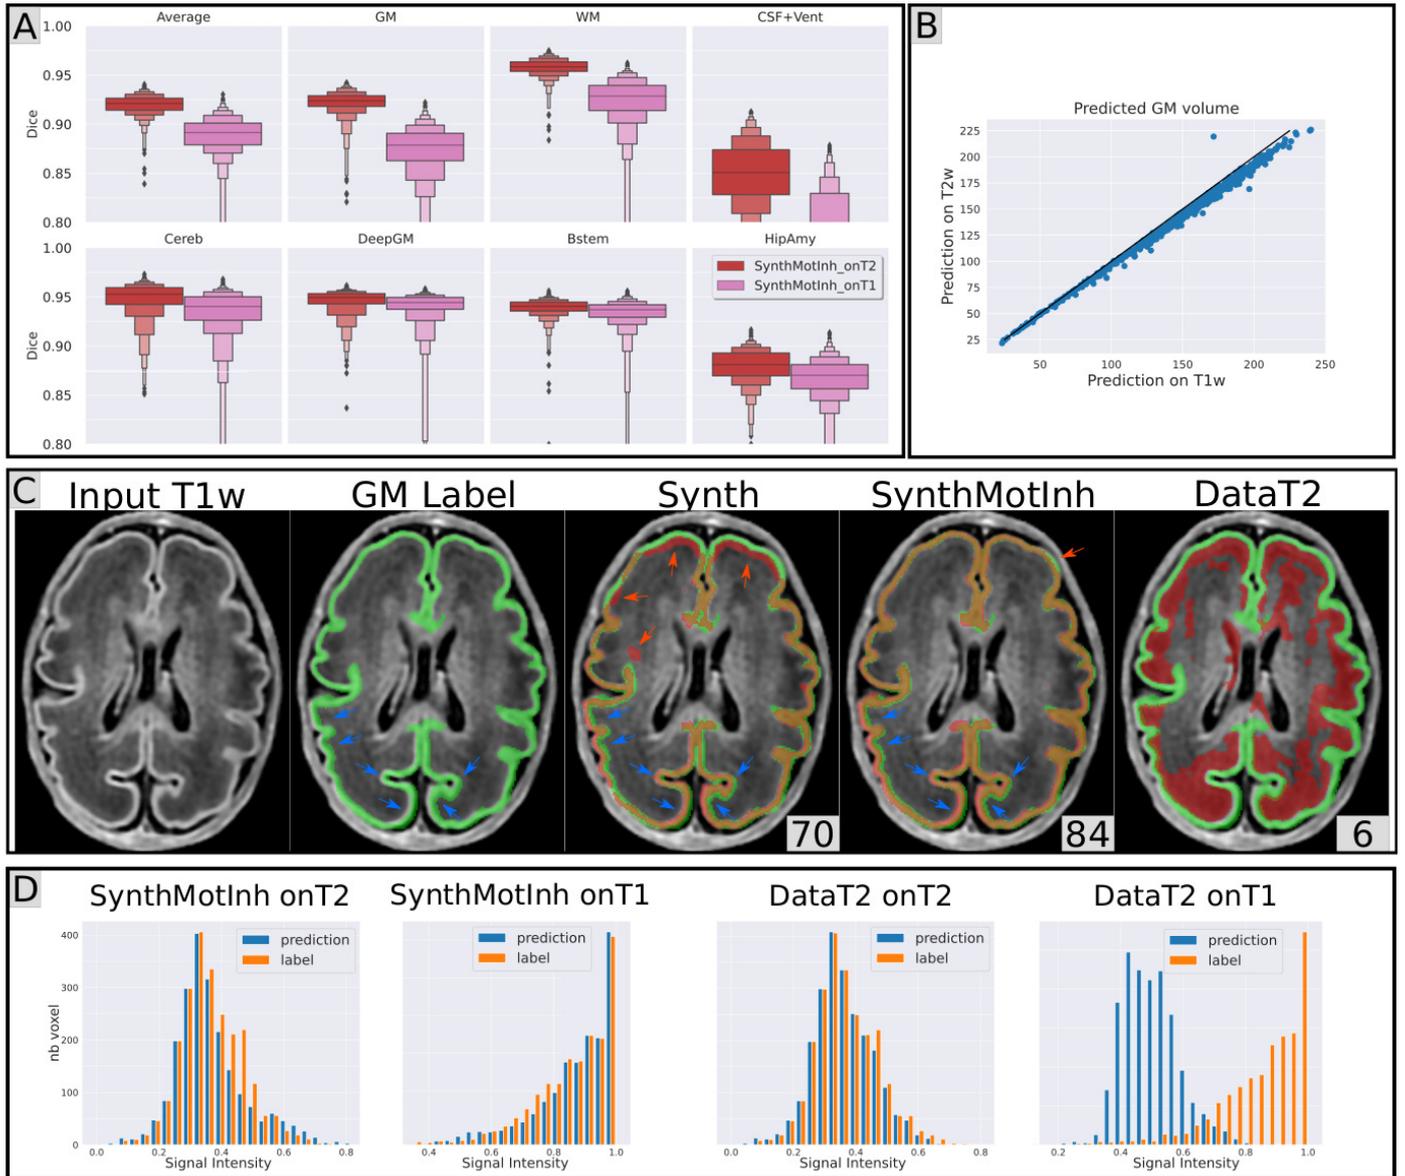


Figure 4: A) distribution across all subjects of the Dice score computed between GT_drawEM and the predictions of the SynthMotInh model from either T2w or T1w, for the different structures. B) scatter plot of the GM volume computed from prediction of the SynthMotInh model. On the y-axis predictions are made from the T2w volumes and on the x-axis from the T1w volumes. C) illustration of the visual observations across the different methods. Ground truth label (GT_drawEM) is shown in green, and the prediction in red. Blue arrows indicate regions with visible misalignment of the GT with regard to the T1w image. Red arrows indicate local errors in the predictions. D) histograms of the intensities of the T1w and T2w images within the predicted GM label (in blue) and GM GT (in orange).

- The DataT2 model can not generalize to other contrasts;
- The synthesis-based models perform equally well on both modalities;
- Remaining differences in Dice are mostly due to misregistration between T1w and T2w and not to segmentation errors.

3.3. Experiment #3: Influence of the definition of the ground truth

The aim of this experiment was to assess the influence of the definition of the ground truth on the performance of the models. We computed another ground truth for the GM derived from the cortical surfaces GT_surf, as explained in section 2.1.2 (all

other tissues remain the same, except WM and CSF). We re-run the same experiments with this new ground truth for two models only: SyntMotInh and DataT2. We used the following experimental setup:

- Training & testing #1: ground truth = GT_DrawEM, prediction on T2w compared to GT_DrawEM (same model as Exp#1)
- Training & testing #2: ground truth = GT_Surf, prediction on T2w compared to GT_Surf
- The influence of the change in the ground truth is assessed by comparing the predictions from the two training sessions.

The results from this experiment for the GM are reported on

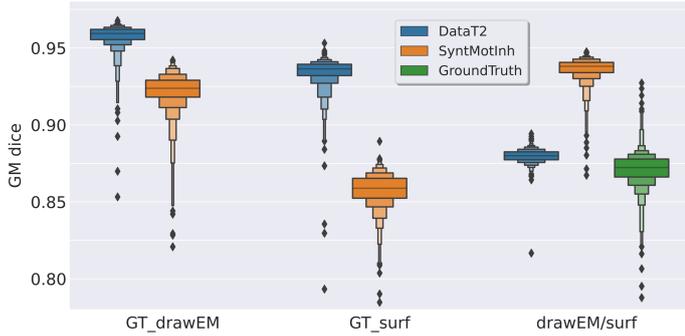


Figure 5: Dice computed between predictions and ground truth. The columns GT_DrawEM show the dice obtained when the models are trained and evaluated using the GT_DrawEM ground truth. (same values as Fig. 2.) The columns GT_surf show the measures when the models are trained and evaluated with the GM derived from the surfaces. For the columns drawEM/surf, the Dice is computed between the predictions of the model trained with GT_drawEM and the predictions of models trained with surf GT_surf. For comparison, we also show the Dice between the two ground truths (in green).

Fig. 5. The first observation is the clear impact of the ground truth on the performance of the models. We observe a drop of 2 points for DataT2 when using GT_surf instead of GT_DrawEM (from 95 to 93). The drop is even larger for the SyntMotInh model: 6 points (from 92 to 86). On the other hand, when we look at the consistency between the predictions from the two training sessions (column drawEM/surf), we observe a dice of 88 for the DataT2 model, which is very close to the dice between the two ground truths. In contrast, the Dice value of 94 for the SyntMotInh model indicates that the predictions are much more consistent, showing a lower influence of the type of ground truth used in the training set.

Visual assessment enabled us to better interpret the drop of 2 points of Dice for Data T2 model trained (and evaluated) with GT_surf compared to the same model trained on GT_drawEM. Indeed, this drop in performance is not due to a more difficult task or less accurate predictions, but to local inaccuracies in GT_surf. We observed focal errors in GT_surf that are related to bad positioning of the white or pial surfaces, probably due to the balance between topology correction and data attachment terms in the surface deformation algorithm Schuh et al. (2017). These observations also explain the spread in the distribution of Dice computed between the two ground truths (in green on Fig. 5). Note also that the spread is reduced for the prediction of DataT2 models (third column in blue compared to green), demonstrating a better robustness of DataT2 predictions compared to the both ground truths.

In summary, the DataT2 model learns very well the ground truth whatever its definition, even with only 15 subjects. The synthesis-based model is less impacted by a change in the definition of the ground truth labels used for the generative model.

3.4. Detailed assessment of image quality and anatomical relevance of the segmentation provided by the dHCP

The anatomical validity of ground truth is rarely discussed in the deep learning literature, but it plays an important role. In this work, we used the segmentation provided by the dHCP

consortium as one of our ground truths (GT_drawEM). These segmentation maps were obtained using automated image processing as described in Makropoulos et al. (2018b). The segmentation pipeline was optimized for instance by modeling additional tissue classes to account for inhomogeneity in WM. Such segmentations are of great value and our study would not be feasible without such material, as well as many other publications. In this section, we report additional observations from our extensive visual assessment that might serve for future works based on this dataset. We visualized systematically all the images corresponding to potential outliers, i.e. for which the Dice score was far from the mean. For comparison, we also visualized randomly picked images to assess the average performance.

As illustrated in Panel A of Fig. 6, the anatomical relevance of the predictions from DataT2 is better than GT_drawEM. We identify three types of outliers: Outlier type 1 (N=23) corresponds to obvious, relatively large errors of the GT_drawEM, despite the high quality of the T2w image; Outlier type 2 (N=11) corresponds to lower quality T2w data for which GT_drawEM was affected by artifacts, while the prediction from the DataT2 model looks much more anatomically relevant; Outlier type 3 (N=1) corresponds to the only subject for which the prediction from DataT2 shows obvious errors with False positive GM prediction near the ventricle. Note that this subject was classified as pathological by the dHCP consortium (radiological_score=5), which suggests that the visually enlarged lateral ventricles likely correspond to very large variations with respect to the normal brain configuration.

In Panel B of Fig. 6, we report our observations relative to the quality of the T1w versus T2w data from the dHCP. We visually checked the 33 outliers from the distribution of Dice score obtained for the SyntMotInh predictions of GM from T1w images and we identified three types: Outlier type 1 (N=11) corresponds to high-quality images but with obvious misregistration between T1w and T2w. Note that we show in panel B 1) the outlier with the highest dice score (above 0.8) but we observed 8 subjects with a dice lower than 0.5 due to obvious misregistration; Outlier type 2 (N=12) corresponds to subjects with bad quality T1w images, for which a low overlap with the GT is expected; Outlier type 3 (N=10) corresponds to very young subjects, for which the SyntMotInh model was less accurate (on both contrasts) than for the older ones. We do not illustrate this configuration since it is already shown on Fig. 3 and 4. The misalignment and quality issues from outlier types 1 and 2 explain the loss of Dice in Exp #2 since the GT_drawEM has been defined from the T2w volumes only.

Overall, the visual assessment of the outlier showed cases with large errors in the GT. Either because of failure of the drawEM pipeline, or because of mis-coregistration issues. In addition, for both experimental examples we also found outliers due to the poor image quality.

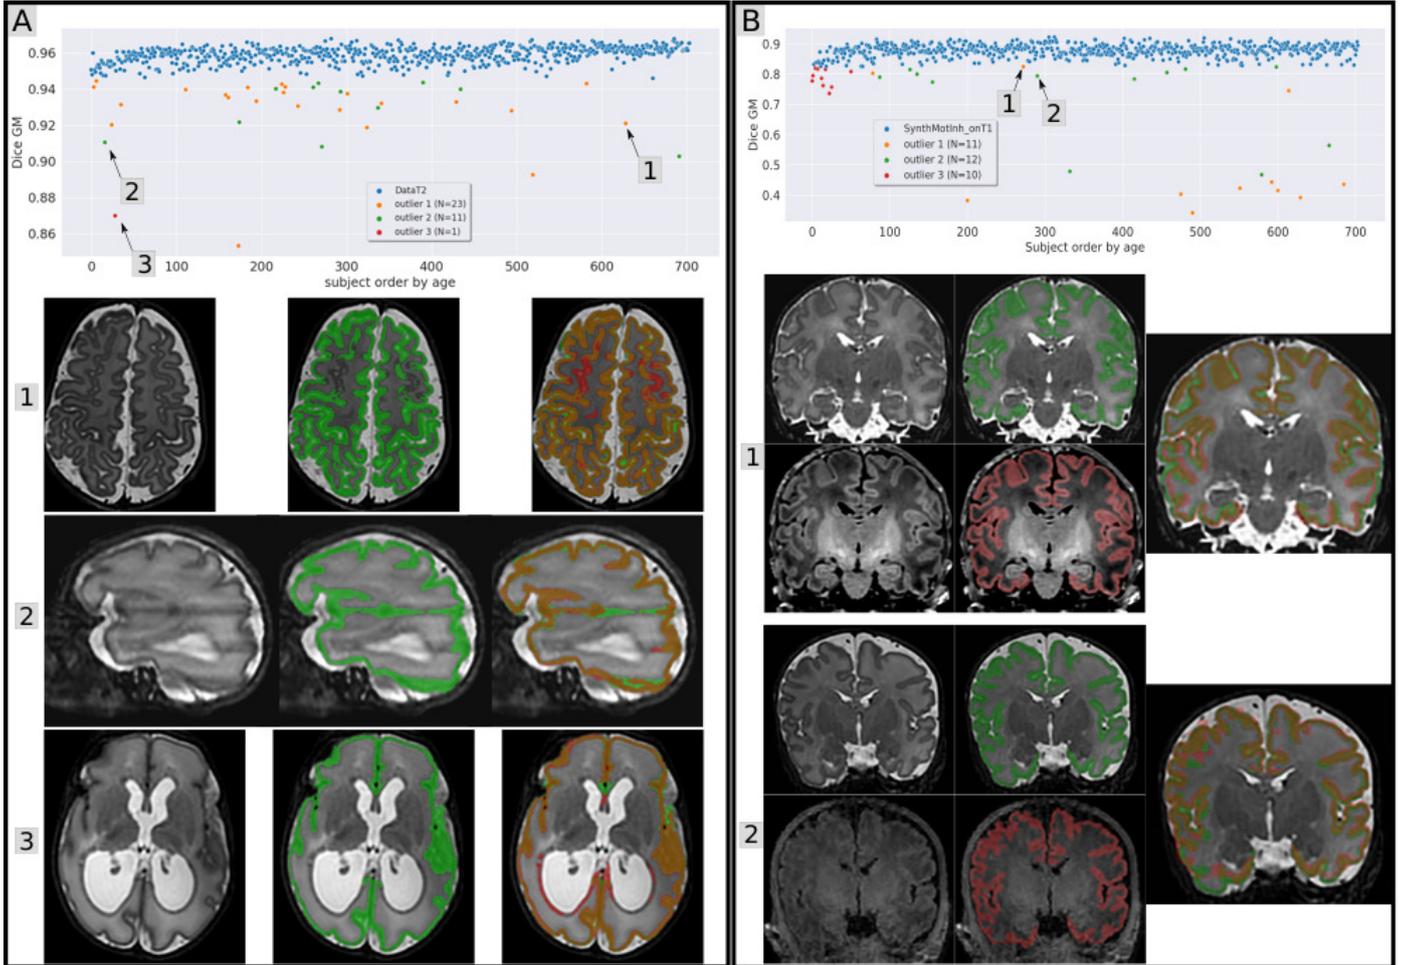


Figure 6: A) Dice score of DataT2 model (evaluated on T2w) for all subjects ordered by the age of the subjects. We visually checked the 35 outliers, grouped them into three categories, and provided an example for each type. (The examples shown below are indicated by a black arrow). B) Dice score of the SynthMotInh GM predictions (evaluated on T1w) for all subjects, after excluding the 35 outliers from panel A. We visually checked the 33 outliers and grouped them into three categories. Chosen examples are shown by a black arrow. Ground truth labels (GT_drawEM) are shown in green, and the predictions in red.

4. Discussion

4.1. Synthetic learning models require adaptations to perform on neonatal MRI

The synthetic approach proposed by Billot et al. (model Synth) performed worse than we anticipated. However, this low performance only affects the GM of the youngest subjects (i.e. those with unfolded GM). The method is more effective for older subjects, which agrees with its good performance for the adult brain Billot et al. (2023, 2022). However, this finding is unexpected since the GM segmentation might appear simpler to perform on brains without convolution. Furthermore those errors occur in regions with high image quality, with a clear contrast between GM and surrounding tissues. We hypothesize that this failure mode is due to a too simple generative model that cannot account for the variations in intensity within the immature WM, which are much larger in this dataset compared to adult brains. Our results show that adding motion augmentation in the generative process significantly improves the performance, as shown by the SynthMot results. The motion simulation mixes different tissues and can lead to local-

ized artifacts (illustrated on Fig. 2) that are qualitatively similar to the inhomogeneities in the white matter induced by maturation, with a spatial pattern in layers propagating inward from the GM Pogledic et al. (2020). Therefore, the gain in performances might be interpreted as a positive side effect rather than an anatomically relevant data augmentation. Anyway, these observations confirm the statement from Billot et al. (2023); Tobin et al. (2017); Tremblay et al. (2018) that in the domain randomization approach, the key is to generate enough variations to cover the expected variations from real data, even if the generative process does not perfectly model the real data generation. In addition to data augmentation with simulated motion, we also explored the alternative solution of explicitly modeling heterogeneity in WM by decomposing it into subregions (between 2 and 6). Our results suggest that this approach is less efficient than the proposed motion augmentation but still beneficial. Further work is needed to fully understand the potential of this strategy, as performance gains could be achieved by refining the number of subregions or by incorporating additional anatomical priors.

4.2. Robustness to variations in image contrast

Our results confirm that the DataT2 model did learn a correspondence between the spatial location and the underlying intensity in the image, as expected. Improving the generalization properties of the deep learning models is a very active topic, with various strategies investigated in parallel such as e.g. [Tomar et al. \(2022\)](#); [Zhao et al. \(2019\)](#); [Ouyang et al. \(2022\)](#). In this article we did not enter into evaluating these methods since in the context of early brain development, variations in image contrast have a biological meaning and might not be considered as a domain adaptation problem. Indeed, robustness to variations in image contrast is a key feature in this context, which is different from generalization between two predefined domains. In our experiments, we consider the variations in contrast between T1w and T2w as a prototypical, extreme change. We confirm the contrast-agnostic properties of synthesis-based models with highly consistent predictions from either T1w or T2w images. We argue that the residual difference of 5 points of Dice for GM between predictions from T1w and T2w is not due to inaccurate predictions. We identified three main factors explaining this finding. First, residual misregistration between the T1w and T2w images from the same subjects are clearly present in the dHCP dataset. Second, different artifacts might affect the two acquisitions, impacting the predictions differently but inducing systematically a reduction of the Dice score. Third, the T1w and T2w image contrasts may be affected differently by brain maturation at the cellular level [Croteau-Chonka et al. \(2016\)](#). Despite these uncontrolled sources of variance, the very high correlations across tissue volumes computed from T1w versus T2w confirm the robustness of the proposed SynthMotInh segmentation to variations in image contrast.

The residual mis-registration issues we report in Section 3.4 are somewhat inconsistent with the dHCP image processing pipeline description [Makropoulos et al. \(2018b\)](#). The authors noted that gradient non-linearity correction was not necessary, and reported that rigid co-registration was effective. However, this study was conducted on the first release of the dHCP dataset, which contains 465 subjects, whereas we included 704 subjects from the third release. Through our careful visual assessment, we observed large and obvious residual mis-registration errors for at least 11 subjects. Therefore, we believe that this dataset is affected by mis-registration for a larger proportion of individuals. Our results suggest that comparing the segmentation predicted from a synthesis-based model from T1w and T2w is effective to detect mis-registration. This could be used to guide the registration between these two modalities. This is consistent with the recent study by [Iglesias](#) which introduced a robust multicontrast affine registration approach based on synthetic learning segmentations [Iglesias \(2023\)](#).

4.3. Robustness of synthesis-based approaches relative to the definition of the ground truth

In general, the quantitative evaluation of a supervised segmentation method, such as dataT2 in this work, measures the network’s ability to learn the ground truth from the training set (i.e. to learn the mapping between an image and the corresponding label map), regardless of the quality of the ground

truth. The high DICE scores observed for dataT2 in Exp #1 confirms this capacity to learn very accurately the ground truth. Consequently, the predictions are highly dependent on the ground truth, and any systematic bias affecting the ground truth would be learnt by such models. The bias would in turn affect the predictions. This was clearly confirmed by the result of Exp #3 where the predictions of DataT2 models (GT_drawEM and GT_surf) show the same dice score (0.85) as the labels. In contrast, the synthetic approach does not learn the relationship between image intensities and segmentation maps. By design, the synthetic framework relies on a generative model to simulate images, which induces an exact correspondence between the boundaries of the structures of interest and image intensities in the training set. This could explain the robustness of the predictions from the SynthMotInh method with respect to variations in the ground truth (see Exp #3). Synthesis-based models are thus less biased by the quality of the ground truth than classical supervised models, which constitutes a new path to define boundaries between adjacent tissues, a major bottleneck for quantitative analysis.

5. Conclusion

In this work, we confirmed that synthetic learning for data segmentation offers key advantages compared to the classical strategy based on real data. In the context of newborn brain MRI, specific signal inhomogeneities affect the performance of the previously proposed synthetic approach. Our enriched generative model with motion simulation greatly improved the predictions, enabling contrast-agnostic segmentation of neonatal brain MRI. In addition, the synthesis-based models are not biased toward the specific intensity distributions of the training set, and the relationship between the geometry of the tissue and the intensity distribution is better controlled than with real data. Furthermore, synthesis-based models can be trained using a very limited amount of manual segmentation examples. All these features will provide critical performance improvements in the perspective of large multi-site studies and clinical applications.

Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocation 2022-AD011011735R3 made by GENCI. The research leading to these results has also been supported by the ANR A4CHILD Project, Grant ANR-19-CHIA-0015, the ANR SulcalGRIDS Project, Grant ANR-19-CE45-0014, the ERA-NET NEURON MULTI-FACT Project, Grant ANR-21-NEU2-0005 and the ANR HINT Project, Grant ANR-22-CE45-0034 funded by the French National Research Agency. dHCP Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

- Benkarim, O. M., Sanroma, G., Zimmer, V. A., Muñoz-Moreno, E., Hahner, N., Eixarch, E., Camara, O., Gonzalez Ballester, M. A., and Piella, G. (2017). Toward the automatic quantification of in utero brain development in 3D structural MRI: A review. *Human brain mapping*, 38(5):2772–2787.
- Billot, B., Greve, D., Van Leemput, K., Fischl, B., Iglesias, J. E., and Dalca, A. V. (2020). A learning strategy for contrast-agnostic MRI segmentation. *arXiv preprint arXiv:2003.01995*.

- Billot, B., Greve, D. N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A. V., and Iglesias, J. E. (2023). SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis*, 86:102789.
- Billot, B., Magdamo, C., Arnold, S. E., Das, S., and Iglesias, J. (2022). Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets. *arXiv preprint arXiv:2209.02032*.
- Cardoso, M. J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al. (2022). Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*.
- Chartsias, A., Joyce, T., Giuffrida, M. V., and Tsaftaris, S. A. (2018). Multi-modal MR Synthesis via Modality-Invariant Latent Representation. *IEEE Transactions on Medical Imaging*, 37(3):803–814.
- Chen, C., Dou, Q., Chen, H., Qin, J., and Heng, P.-A. (2019). Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 865–872. Issue: 01.
- Croteau-Chonka, E. C., Dean III, D. C., Remer, J., Dirks, H., O’Muircheartaigh, J., and Deoni, S. C. (2016). Examining the relationships between cortical maturation and white matter myelination throughout early childhood. *NeuroImage*, 125:413–421.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data Via the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1977.tb01600.x>.
- Dimitrova, R. et al. (2021). Phenotyping the Preterm Brain: Characterizing Individual Deviations From Normative Volumetric Development in Two Large Infant Cohorts. *Cerebral Cortex*, 31(8):3665–3677.
- Edwards, A. et al. (2022). The Developing Human Connectome Project Neonatal Data Release. *Frontiers in Neuroscience*, 16:886772.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2017). Domain-Adversarial Training of Neural Networks. In Csurka, G., editor, *Domain Adaptation in Computer Vision Applications*, pages 189–209. Cham. Series Title: Advances in Computer Vision and Pattern Recognition.
- Gholipour, A., Akhondi-Asl, A., Estroff, J. A., and Warfield, S. K. (2012). Multi-atlas multi-shape segmentation of fetal brain MRI for volumetric and morphometric analysis of ventriculomegaly. *NeuroImage*, 60(3):1819–1831.
- Iacono, M. I., Neufeld, E., Akinagbe, E., Bower, K., Wolf, J., Vogiatzis Oikonomidis, I., Sharma, D., Lloyd, B., Wilm, B. J., and Wyss, M. (2015). MIDA: a multimodal imaging-based detailed anatomical model of the human head and neck. *PLoS one*, 10(4):e0124126.
- Iglesias, J. E. (2023). A ready-to-use machine learning tool for symmetric multi-modality registration of brain MRI. *Scientific Reports*, 13(1):6657.
- Ilse, M., Tomczak, J. M., and Forré, P. (2021). Selecting data augmentation for simulating interventions. In *International Conference on Machine Learning*, pages 4555–4562.
- Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J., and Maier-Hein, K. H. (2021). nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211.
- Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical image analysis*, 5(2):143–156.
- Kamnitsas, K., Baumgartner, C., Ledig, C., Newcombe, V., Simpson, J., Kane, A., Menon, D., Nori, A., Criminisi, A., and Rueckert, D. (2017). Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In *Information Processing in Medical Imaging: 25th International Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings 25*, pages 597–609.
- Karani, N., Chaitanya, K., Baumgartner, C., and Konukoglu, E. (2018). A lifelong learning approach to brain MR segmentation across scanners and protocols. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 476–484.
- Kirk, T. F., Coalson, T. S., Craig, M. S., and Chappell, M. A. (2020). Tolerone: Surface-Based Partial Volume Estimation. *IEEE Transactions on Medical Imaging*, 39(5):1501–1510. Conference Name: IEEE Transactions on Medical Imaging.
- Kostović, I., Sedmak, G., and Judaš, M. (2019). Neural histology and neurogenesis of the human fetal and infant brain. *Neuroimage*, 188:743–773.
- Li, G., Wang, L., Yap, P.-T., Wang, F., Wu, Z., Meng, Y., Dong, P., Kim, J., Shi, F., Rekić, I., Lin, W., and Shen, D. (2019). Computational neuroanatomy of baby brains: A review. *NeuroImage*, 185:906–925.
- Makropoulos, A., Counsell, S. J., and Rueckert, D. (2018a). A review on automatic fetal and neonatal brain MRI segmentation. *NeuroImage*, 170:231–248.
- Makropoulos, A., Gousias, I. S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J. V., Edwards, A. D., Counsell, S. J., and Rueckert, D. (2014). Automatic whole brain MRI segmentation of the developing neonatal brain. *IEEE transactions on medical imaging*, 33(9):1818–1831.
- Makropoulos, A., Robinson, E. C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S. J., Steinweg, J., Vecchiato, K., and Passerat-Palmbach, J. (2018b). The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *NeuroImage*, 173:88–112.
- Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., Fox, N. C., and Ourselin, S. (2010). Fast free-form deformation using graphics processing units. *Computer methods and programs in biomedicine*, 98(3):278–284.
- Ouyang, C., Chen, C., Li, S., Li, Z., Qin, C., Bai, W., and Rueckert, D. (2022). Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Transactions on Medical Imaging*, 42(4):1095–1106.
- Pan, S. J. and Yang, Q. (2010). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Pogledic, I., Schwartz, E., Mitter, C., Baltzer, P., Milos, R.-I., Gruber, G. M., Brugger, P. C., Hainfellner, J., Bettelheim, D., and Langs, G. (2020). The subplate layers: the superficial and deep subplate can be discriminated on 3 tesla human fetal postmortem MRI. *Cerebral Cortex*, 30(9):5038–5048.
- Pérez-García, F., Sparks, R., and Ourselin, S. (2021). TorchIO: a Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Computer Methods and Programs in Biomedicine*, 208:106236.
- Reguig, G., Lapert, M., Lehericy, S., and Valabregue, R. (2022). Global displacement induced by rigid motion simulation during MRI acquisition. *arXiv preprint arXiv:2204.03522*.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241.
- Schuh, A., Makropoulos, A., Wright, R., Robinson, E. C., Tusor, N., Steinweg, J., Hughes, E., Grande, L. C., Price, A., and Hutter, J. (2017). A deformable model for the reconstruction of the neonatal cortex. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 800–803.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30.
- Tomar, D., Bozorgtabar, B., Lortkipanidze, M., Vray, G., Rad, M. S., and Thiran, J.-P. (2022). Self-supervised generative style transfer for one-shot medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1998–2008.
- Tremblay, J., Prakash, A., Acuna, D., Brophy, M., Jampani, V., Anil, C., To, T., Cameracci, E., Boochoon, S., and Birchfield, S. (2018). Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 969–977.
- Zhang, L., Wang, X., Yang, D., Sanford, T., Harmon, S., Turkbey, B., Wood, B. J., Roth, H., Myronenko, A., Xu, D., and Xu, Z. (2020). Generalizing Deep Learning for Medical Image Segmentation to Unseen Domains via Deep Stacked Transformation. *IEEE Transactions on Medical Imaging*, 39(7):2531–2540.
- Zhang, Z., Yang, L., and Zheng, Y. (2018). Translating and segmenting multi-modal medical volumes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9242–9251.
- Zhao, A., Balakrishnan, G., Durand, F., Guttag, J. V., and Dalca, A. V. (2019). Data augmentation using learned transformations for one-shot medical image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8543–8553.
- Zhou, K., Liu, Z., Qiao, Y., Xiang, T., and Loy, C. C. (2022). Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.