Predicting Longitudinal Brain Development via Implicit Neural Representations

Maik Dannecker[1] and Daniel Rueckert[1,2,3]

[1] School of Computation, Information and Technology, and School of Medicine and Health, Technical University Munich, Germany m.dannecker@tum.de
[2] Department of Computing, Imperial College London, United Kingdom
[3] Munich Center for Machine Learning (MCML), Munich, Germany

Abstract. Predicting individualized perinatal brain development is crucial for understanding personalized neurodevelopmental trajectories, however, remains challenging due to limited longitudinal data. While population based atlases model generic trends, they fail to capture subject-specific growth patterns. In this work, we propose a novel approach leveraging Implicit Neural Representations (INRs) to predict individualized brain growth over multiple weeks. Our method learns from a limited dataset of less than 100 paired fetal and neonatal subjects, sampled from the developing Human Connectome Project. The trained model demonstrates accurate personalized future and past trajectory predictions from a single calibration scan. By incorporating conditional external factors such as birth age or birth weight, our model further allows the simulation of neurodevelopment under varying conditions. We evaluate our method against established perinatal brain atlases, demonstrating higher prediction accuracy and fidelity up to 20 weeks. Finally, we explore the method's ability to reveal subject-specific cortical folding patterns under varying factors like birth weight, further advocating its potential for personalized neurodevelopmental analysis.

Keywords: Brain Development · Implicit Neural Representations · MRI

1 Introduction

Longitudinal brain analysis is crucial to understanding intricate growth patterns and precursors of abnormal neurodevelopment. Fetal and neonatal brains assume here a pivotal role, undergoing considerable morphological changes in mere weeks. Normative brain models from MRI—known as atlases—are a common tool to analyze neurodevelopment[3]. These atlases are constructed from large cohorts of subjects and registered to a reference space, achieving a general representation of the brain. To better capture the morphological changes of the perinatal brain over time, atlases for multiple time points, typically every week, are constructed [13, 12, 8, 21, 22]. These population-based atlases are designed to be as generic as possible, keeping bias to individual brains to a minimum. With recently large published datasets, e.g., the developing Human Connectome Project (dHCP) dataset [10, 17], generating unbiased population-based

atlases for the developing brain has become feasible. These atlases, however, fall short when applied to predict *individual* rather than average brain development. Especially for perinatal brains, morphology can vary drastically between individuals. So far, only few methods exist that take into account subject-specific morphology to predict individualized future growth patterns of the perinatal brain [29, 18, 30]. Another shortcoming of established population-based atlases is their inflexibility to external factors that influence brain development, for example, birth age or weight, both known to have strong confounding factors of diminished cortical folding and impaired brain development [2, 7, 13, 26].

Contribution In this work, we present the first atlas to predict individualized growth patterns for fetal and neonatal brains. We demonstrate that our method faithfully forecasts subject-specific development up to 20 weeks into the future or past and operates seamlessly across the event of birth. Integrated conditioning on external factors (e.g., birth age, birth weight) enables simulations of custom *what-if* scenarios. Trained on a limited longitudinal cohort (100 subjects, one follow-up each), our model generates accurate individualized trajectories from a single calibration scan, outperforming population-based atlases.

To summarize, we:

- Predict individualized growth, providing the first atlas delivering subject-specific fetal and neonatal brain trajectories up to ± 20 weeks.
- Span the birth event, bridging prenatal and postnatal domains via shared latent codes that unify fetal and neonatal scans.
- Condition on external factors, using learnable tokens for scan age, birth age, and weight to simulate *what-if* developmental scenarios.
- Operate data-efficiently, achieving high-accuracy predictions with only one follow-up scan per subject on a cohort of 100 subjects.

2 Related Work

Population Based Atlases First normative perinatal brain models from MRI leveraged group-wise affine registration to align a given population to a common reference space. Notable examples include the fetal atlas by Habas et al. [9] and the pre-term neonatal atlas by Kuklisova-Murgasova et al. [12]. Both are spatio-temporal atlases, capturing anatomy for pre-defined time intervals, typically representing the generic brain for each week from 20 weeks gestational age (GA) to 46 weeks post-menstrual-age (PMA). Follow up work from Serag et al. [21] integrated non-rigid free-form deformations [19] and age-adaptive kernels to balance subject contributions producing sharper atlases with enhanced tissue delineations. Makropoulos et al. [13] and Gholipour et al. [8] followed by developing spatio-temporal atlases from large populations including 420 neonates and 81 fetuses, respectively. The atlases offer brain tissue segmentation at different granularities, from broad cortical and subcortical divisions [21] to detailed maps with over 100 regions [8]. However, these traditional population-based atlases are

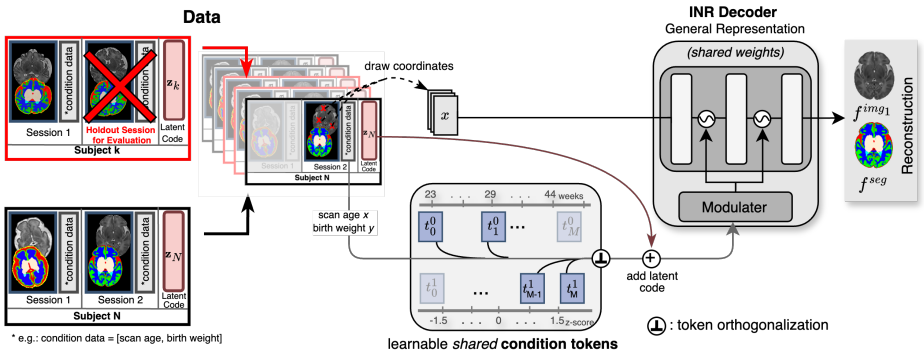

Fig. 1. Individualized Atlas. Paired sessions share a common latent code to encode subject-specific features. Shared condition tokens, sampled and interpolated based on subject conditions (e.g., scan- and birth age), are added to the latent code to modulate the INR decoder, which predicts intensities and tissue labels. During inference, unseen holdout sessions are predicted in a single forward pass via the same process.

limited by their inability to account for external factors known to influence brain development, e.g., birth age or weight [7, 2, 26], and their, by design, averaged representation, restricting their applicability and use to individual subjects.

Longitudinal Prediction Methods Longitudinal brain prediction has been explored in various contexts, including adult and infant brains. Sivera et al. [24] introduced a deformation-based framework for Alzheimer's, conditioned on age and disease state, however, limited to small morphological changes over years. Yoon et al. [28] employed a diffusion model for cardiac and adult brain imaging, using large datasets and predefined age templates for training. Yuan et al. [29] and Rekik et al. [18] modeled infant brain development from 3 to 24 months. Zhang et al. [30] predicted fetal brain growth, conditioned on an initial scan and predicting the subsequent week. Besides the limited prediction horizon, the method lacks evaluation on ground-truth, excludes neonatal brains, and does not incorporate external developmental factors.

To the best of our knowledge, this work presents the first comprehensive and individualized, image-based atlas for perinatal brain development. We propose a condition-aware approach that predicts individualized growth trajectories over time spans up to 20 weeks, bridging the gap between population atlases and personalized modeling.

3 Methodology

Derived from [4], our individual atlas model, outlined in Fig. 1, employs three essential building blocks: latent code, condition tokens, and INR decoder. In the following, we outline each of these blocks in detail.

3.1 Architecture

Latent Code For each of N training subjects, we employ a latent code $\{\boldsymbol{z}_j\}_{i=j}^{N}$, tasked to encode subject specific properties during training. To enforce the model to learn subject specific but not session specific features, a subject's sessions (in this work two) share the same latent code initialized as $\boldsymbol{z} \sim \mathcal{N}(0, 10^{-2})$ [16].

Condition Tokens In addition to the latent codes, we introduce learnable condition tokens $\{\boldsymbol{t}_i^{\boldsymbol{c}}\}_{i=1}^{M}$ with $c \in \{0, \ldots, C\}$, where C denotes the number of conditions and M the number of tokens discretizing each condition domain. For instance, the scan age of perinatal brain data typically spans 20 to 45 weeks postmenstrual age (PMA), which we represent with $M = 5$ tokens. The tokens are *shared* across all subjects. During training, when reconstructing a subject j with a given condition value (e.g., scan age of 26 weeks), the corresponding token value is computed via linear interpolation between its two neighboring tokens:

$$\boldsymbol{t}^{\boldsymbol{c}}(v) = w_1 \boldsymbol{t}_i^{\boldsymbol{c}} + w_2 \boldsymbol{t}_{i+1}^{\boldsymbol{c}}, \quad \text{where} \quad w_1 = \frac{v_{i+1} - v}{v_{i+1} - v_i}, \quad w_2 = \frac{v - v_i}{v_{i+1} - v_i}$$

Here, v is the given condition value and v_i, v_{i+1} are the predefined values associated with tokens $\boldsymbol{t}_i^{\boldsymbol{c}}$ and $\boldsymbol{t}_{i+1}^{\boldsymbol{c}}$, respectively. Next, the interpolated token vectors $\boldsymbol{t}^{\boldsymbol{0}}(v^0), \ldots, \boldsymbol{t}^{\boldsymbol{C}}(v^C)$ are orthogonalized via the Gram-Schmidt process to ensure independent contributions of each condition, preventing redundancy and enhancing disentanglement in the representation. Finally, the latent representation $\boldsymbol{z}_j$ of subject j is attuned by adding the sum of the orthogonalized tokens:

$$\tilde{\boldsymbol{z}}_j = \boldsymbol{z}_j + \sum_{c=0}^{C} \hat{\boldsymbol{t}}^c \tag{1}$$

The new code $\tilde{\boldsymbol{z}}_j$ modulates the INR decoder following the approach of [4, 6, 15].

Modulated INR The INR architecture is identical to [4] that follows the modulation approach of [6, 15]. The INR acts as decoder of latent code and tokens and its weights are shared across all subjects. We transform the modulation code $\tilde{\boldsymbol{z}}_j \in \mathbb{R}^D$ into scale ($\boldsymbol{\phi} \in \mathbb{R}^L$) and shift ($\boldsymbol{\psi} \in \mathbb{R}^L$) parameters via a linear modulation layer $\begin{pmatrix} \boldsymbol{\phi} \\ \boldsymbol{\psi} \end{pmatrix} = \mathbf{P}\mathbf{z} + \boldsymbol{\mu}$ where $\boldsymbol{P} \in \mathbb{R}^{2L \times D}$ and $\boldsymbol{\mu} \in \mathbb{R}^{2L}$. Each modulation layer then outputs $\sin\left(\omega_0 \cdot \boldsymbol{\phi} \cdot (\boldsymbol{W}\boldsymbol{x} + \boldsymbol{b}) + \boldsymbol{\psi}\right)$, where $\boldsymbol{W} \in \mathbb{R}^{L \times L}$ and $\boldsymbol{b} \in \mathbb{R}^L$ are the layer weights and biases of the MLP, and ω_0 is a scaling factor to model high-frequency signals [23]. Note, the weights of the INR as well as $\boldsymbol{P}, \boldsymbol{\mu}, \boldsymbol{W}$, and $\boldsymbol{b}$ are shared across all training data.

3.2 Training and Inference

The model is trained by optimizing the latent codes $\{\boldsymbol{z}_j\}_{j=1}^{N}$, the condition tokens $\{\boldsymbol{t}_i^{\boldsymbol{c}}\}_{i=1}^{M}$, and the network parameters θ. Given N training subjects, we maximize

the joint log posterior over all subjects by minimizing the reconstruction loss for both image intensity and segmentation:

$$\underset{\theta,\{\boldsymbol{z}_j\}_{j=1}^{N},\{\boldsymbol{t}_i^c\}_{i=1}^{M},c\in\{0,\ldots,C\}}{\arg\min} \sum_{j=1}^{N}\sum_{\boldsymbol{x}\in\boldsymbol{X}_j}\Big[\mathcal{L}_{MSE}\left(f_\theta^{img}(\boldsymbol{x},\tilde{\boldsymbol{z}}_j),\hat{I}_{\boldsymbol{x}}\right)+\mathcal{L}_{CE}\left(f_\theta^{seg}(\boldsymbol{x},\tilde{\boldsymbol{z}}_j),\hat{C}_{\boldsymbol{x}}\right)\Big]. \tag{2}$$

Here, $\tilde{\boldsymbol{z}}_j$ is the modulated latent code for subject j, computed as in Eq. 1. The functions f_θ^{img} and f_θ^{seg} denote the INR decoder heads for image reconstruction and segmentation, respectively, with $\hat{I}_{\boldsymbol{x}}$ and $\hat{C}_{\boldsymbol{x}}$ being the ground-truth intensity and tissue label of voxel $\boldsymbol{x}\in\boldsymbol{X}_j$. The loss functions $\mathcal{L}_{MSE}$ and $\mathcal{L}_{CE}$ represent the mean squared error and cross-entropy loss, ensuring accurate intensity and tissue label predictions. The optimization updates the network parameters θ, latent codes $\{\boldsymbol{z}_j\}_{j=1}^{N}$, and the shared condition tokens $\{\boldsymbol{t}_i^{\boldsymbol{c}},\boldsymbol{t}_{i+1}^{\boldsymbol{c}}\}$. Inference is done with customized or *target* condition data to sample the corresponding tokens which are added to the subject's optimized latent code. The customized or *target* brain is reconstructed by the INR decoder via a forward pass.

4 Experimental Setup

4.1 Dataset

We use the dHCP (developing Human Connectome Project) dataset for our studies. Imaging includes T2-weighted sequences, following the protocol described in [14, 17]. Most scans include quality-controlled, automatically segmented tissue maps of brain regions [14]. The dataset contains almost 300 fetal and over 800 neonatal scans. For this work we required paired data, i.e., subjects with at least two scans of different time points. dHCP contains 22 subjects with ≥ 2 fetal scans per subject, 79 subjects with a fetal- and a follow-up neonatal scan, and 99 subjects with ≥ 2 neonatal scans. Discarding subjects with missing segmentation files or radiology score > 3, indicating "Incidental findings with possible/likely clinical significance [...]" left us with 70 subjects for training, 15 subjects for validation, and 30, subjects for testing, totaling 115 subjects with 230 scans.

4.2 Pre-Processing and Hyperparameters

All subjects were rigidly aligned to the same reference space, masked, skull-stripped, and brain intensities normalized to $[0, 1]$. Due to resolution-agnostics of the INR, we do not resample the scans to the same resolution, preserving image details by avoiding interpolation. All experiments used the same set of hyperparameters optimized on the validation set. We configured the INR with 9 hidden layers, 1024 units each, and modulation applied to layers 1, 3, 5, 7, and 9. The latent code and token dimension was set to 128. Number of tokens was set to $M = 5$ as a larger number of tokens promoted overfitting. Learning rates were set to $l_r = 5e^{-5}$. Batch size, i.e., number of coordinates, was set to 25 000. We trained for 10 epochs requiring ~4 hours on a NVIDIA A6000 GPU.

4.3 Baselines

1. **BD-Atlas (fetal & neonatal) [21, 13]:** 4D fetal and neonatal brain atlas built from T2w MRI of 80 fetuses and 420 neonates, respectively. Tissue probability maps include cerebrospinal fluid (CSF), cortical gray matter (cGM), lateral ventricles (LV), and brain hemispheres (BH). Both BD-Atlases are publicly available at brain-development.org.
2. **Deepali [20]:** An open-source GPU-accelerated registration library which we used to construct fetal and neonatal atlases from the dHCP data, using a group-wise registration approach following [25].

4.4 Training and Evaluation Setup

Additionally to the 70 training subjects, we included one session of each of the 30 test subjects, referred to as *known* sessions, for training. After training, the model predicts the 30 unseen holdout-sessions, referred to as *target* sessions, via a single forward pass. The prediction was rigidly aligned to the ground truth [1]. Predictions with the atlas baselines were performed using two methods:

- M1: An atlas corresponding to the same developmental stage (fetal or neonatal) as the *target* brain and of similar scan age was rigidly aligned and compared to the target brain, i.e., the ground truth.
- M2: The scan from the *known* session was deformed to match the atlas of the *target* scan age and then rigidly aligned and compared to the *target* brain.

5 Experiments and Results

5.1 Longitudinal Brain Prediction

We evaluate brain prediction accuracy using the 30 holdout sessions of 30 test subjects. For 15 subjects the model has seen the first session and is tasked with predicting the follow-up session (prediction of future states). For the remaining 15 subjects, the model has seen the follow-up session and predicts the first session (prediction of past states). Test subjects include fetal-fetal, fetal-neonate, and neonate-neonate sessions. We evaluate image similarity via PSNR and SSIM, and segmentation accuracy via dice score (DSC) of cGM and mean dice score ($\overline{DSC}$) over CSF, cGM, LV, and BH. Finally, we compare the predicted brain volume to the ground-truth brain volume VOL_{Δ}. Quantitative (Table 1) and qualitative results (Fig. 2) underscore the shortcomings of population atlases in representing the individual perinatal brain development. While our method's predictions adequately capture subject specific morphological characteristics, baselines represent average brain anatomy while ignoring individual mophology.

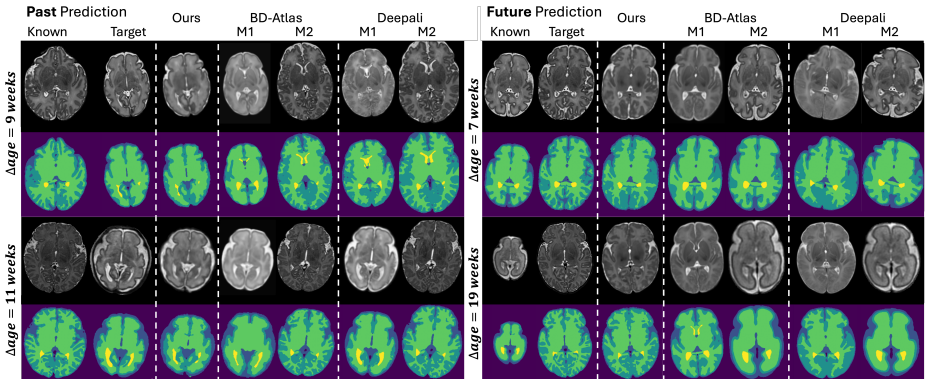

Fig. 2. Qualitative results for predicting **past** time points (left) and **future** time points (right). While the baselines can only represent average brain anatomy, our method adequately captures individual anatomical characteristics, even for prediction from fetal to neonatal state or vice versa.

Table 1. Quantitative evaluation between predicted and ground-truth brains. Metrics include PSNR, SSIM, Dice score for cortical gray matter ($DSC\,cGM$), mean Dice score ($\overline{DSC}$) over all labels, and brain volume difference (VOL_Δ) in cm^3. $MEAN_{\pm STD}$ over 30 subjects. Best sores in **bold**.

Method	PSNR	SSIM	$DSC\,cGM$	$\overline{DSC}$	VOL_Δ (cm^3)
BD-Atlas M1	$18.57_{\pm 1.04}$	$0.34_{\pm 0.20}$	$0.48_{\pm 0.07}$	$0.55_{\pm 0.05}$	$302_{\pm 101}$
BD-Atlas M2	$17.64_{\pm 1.25}$	$0.62_{\pm 0.03}$	$0.48_{\pm 0.09}$	$0.54_{\pm 0.08}$	$171_{\pm 63}$
Deepali M1	$17.01_{\pm 1.08}$	$0.69_{\pm 0.08}$	$0.54_{\pm 0.09}$	$0.58_{\pm 0.09}$	$51_{\pm 28}$
Deepali M2	$17.39_{\pm 1.56}$	$0.61_{\pm 0.06}$	$0.45_{\pm 0.10}$	$0.52_{\pm 0.07}$	$176_{\pm 48}$
Ours	$\mathbf{20.66}^*_{\pm 1.51}$	$\mathbf{0.75}^*_{\pm 0.08}$	$\mathbf{0.62}^*_{\pm 0.05}$	$\mathbf{0.65}^*_{\pm 0.05}$	$\mathbf{45}_{\pm 19}$

*Value significantly better than the second best ($p < 0.05$, paired t-test).

5.2 Trajectory Prediction

Provided with a subject's conditional data and a single MRI scan, our model predicts the future and past growth trajectory over several weeks via extrapolation of the condition token domain—including scan age and birth age—in the corresponding direction. Fig. 3 illustrates the predicted growth trajectory of a subject. Note how the model realistically transitions from fetal brain appearance to neonatal, bridging the event of birth faithfully.

We can further predict trajectories conditioning on other secondary properties, for example the subject's birth weight z-score, i.e., the factor of standard-deviations from the mean weight. Several work in literature have found cortical folding, cortical surface area and gyrification index, to be lower for neonates of lower birth weight z-scores [5, 7, 11]. As depicted in Fig. 4, this correlation is captured by our model, with decreased cortical folding if conditioned on lower birth weight.

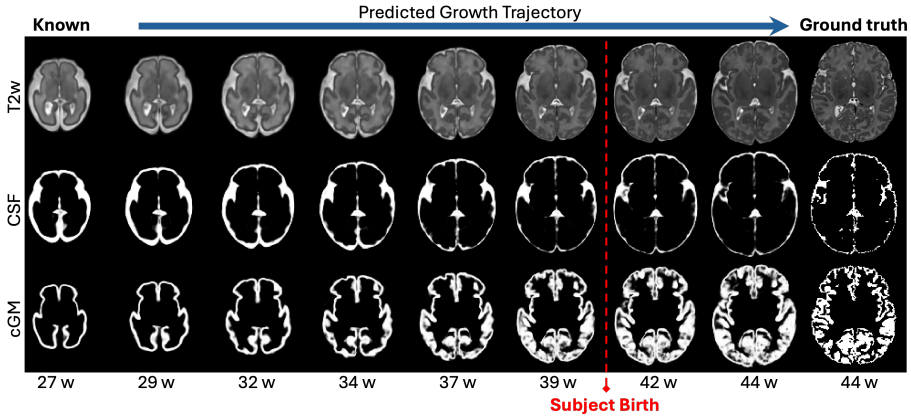

Fig. 3. Growth trajectory of the individual brain with CSF, and cGM probability maps. Red line indicates subject's time of birth. Left column shows the *known* time point, right column shows the *target* time point, i.e., the ground truth.

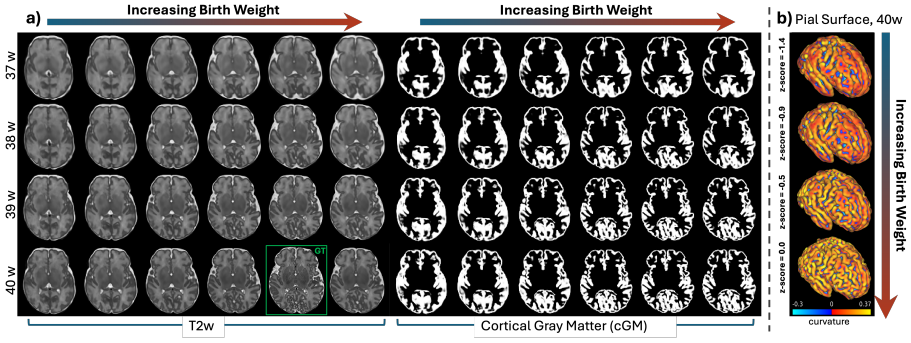

Fig. 4. Simulated individual brain development, conditioned on age (rows, weeks 37 to 40) and birth weight (columns, z-scores -1.5 to 0). b) Pial surface reconstructions, via [14, 27] for the brain of 40 weeks conditioned on birth weight. Top to bottom shows increasing birth weight from -1.4 z-scores to 0 z-scores, i.e., mean birth weight.

6 Conclusion

We presented the first approach for predicting individualized perinatal brain growth using implicit neural representations. Unlike traditional population-based atlases modeling generic trends, our method captures subject-specific growth trajectories with high accuracy requiring only a single calibration scan for predictions. This makes it highly practical in clinical and research settings where longitudinal imaging is scarce. Beyond age progression, our approach allows conditional predictions, incorporating factors such as birth age or birth weight. This enables simulations of neurodevelopment under different conditions, offering in-

sights into how external factors influence individual brain development, relevant, among others, for studying preterm birth effects.

In this work we have only predicted future or past brain development of subjects included in training, limiting the methods applicability. Future work will explore test-time optimization of new unseen subjects for prediction, following established approaches in literature [16, 4]. Finally, integrating implicit cortical surface representation instead of relying on external tools could further strengthen the method.

Acknowledgments. This study was supported by the ERC (Deep4MI - 884622), and by the ERA-NET NEURON Cofund (MULTI-FACT - 8810003808).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Avants, B.B., Tustison, N., Song, G.: Advanced normalization tools (ants). Insight j **2**(365), 1–35 (2009)
2. Bouyssi-Kobar, M., et al.: Third trimester brain growth in preterm infants compared with in utero healthy fetuses. Pediatrics **138**(5) (2016). https://doi.org/10.1542/peds.2016-1640
3. Ciceri, T., et al.: Fetal brain mri atlases and datasets: A review. NeuroImage **292**, 120603 (2024). https://doi.org/10.1016/j.neuroimage.2024.120603
4. Dannecker, M., Kyriakopoulou, V., Cordero-Grande, L., Price, A.N., Hajnal, J.V., Rueckert, D.: Cina: Conditional implicit neural atlas for spatio-temporal representation of fetal brains. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 181–191 (2024)
5. Dubois, J., et al.: Primary cortical folding in the human newborn: an early marker of later functional development. Brain **131**(Pt 8), 2028–41 (2008). https://doi.org/10.1093/brain/awn137
6. Dupont, E., Kim, H., Eslami, S.M.A., Rezende, D.J., Rosenbaum, D.: From data to functa: Your data point is a function and you can treat it like one. In: 39th International Conference on Machine Learning (ICML) (2022)
7. Engelhardt, E., Inder, T.E., Alexopoulos, D., Dierker, D.L., Hill, J., Van Essen, D., Neil, J.J.: Regional impairments of cortical folding in premature infants. Ann Neurol **77**(1), 154–62 (2015). https://doi.org/10.1002/ana.24313
8. Gholipour, A., et al.: A normative spatiotemporal mri atlas of the fetal brain for automatic segmentation and analysis of early brain growth. Sci Rep **7**(1), 476 (2017). https://doi.org/10.1038/s41598-017-00525-w
9. Habas, P.A., et al.: A spatiotemporal atlas of mr intensity, tissue probability and shape of the fetal brain with application to segmentation. Neuroimage **53**(2), 460–470 (2010)
10. Hughes, E.J., et al.: A dedicated neonatal brain imaging system. Magnetic Resonance in Medicine **78**(2), 794–804 (2017). https://doi.org/10.1002/mrm.26462
11. Kersbergen, K.J., et al.: Relation between clinical risk factors, early cortical changes, and neurodevelopmental outcome in preterm infants. Neuroimage **142**, 301–310 (2016). https://doi.org/10.1016/j.neuroimage.2016.07.010

12. Kuklisova-Murgasova, M., et al.: A dynamic 4d probabilistic atlas of the developing brain. Neuroimage **54**(4), 2750–63 (2011). https://doi.org/10.1016/j.neuroimage.2010.10.019
13. Makropoulos, A., et al.: Regional growth and atlasing of the developing human brain. Neuroimage **125**, 456–478 (2016). https://doi.org/10.1016/j.neuroimage.2015.10.047
14. Makropoulos, A., et al.: The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. Neuroimage **173**, 88–112 (2018)
15. Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14214–14223 (2021)
16. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
17. Price, A.N., et al.: The developing human connectome project (dhcp): fetal acquisition protocol. In: Proceedings of the annual meeting of the ISMRM. vol. 244 (2019)
18. Rekik, I., Li, G., Wu, G., Lin, W., Shen, D.: Prediction of infant mri appearance and anatomical structure evolution using sparse patch-based metamorphosis learning framework. In: Patch-Based Techniques in Medical Imaging. pp. 197–204. Springer International Publishing, Cham (2015)
19. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast mr images. IEEE transactions on medical imaging **18**(8), 712–721 (1999)
20. Schuh, A., Qiu, H., Research, H.: deepali: Image, point set, and surface registration in pytorch (0.6.3) (2024). https://doi.org/10.5281/zenodo.14197094
21. Serag, A., et al.: Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. Neuroimage **59**(3), 2255–65 (2012). https://doi.org/10.1016/j.neuroimage.2011.09.062
22. Serag, A., et al.: A multi-channel 4d probabilistic atlas of the developing brain: application to fetuses and neonates. Annals of the BMVA **2012**(3), 1–14 (2012)
23. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. Advances in neural information processing systems **33**, 7462–7473 (2020)
24. Sivera, R., Delingette, H., Lorenzi, M., Pennec, X., Ayache, N.: A model of brain morphological changes related to aging and alzheimer's disease from cross-sectional assessments. NeuroImage **198**, 255–270 (2019). https://doi.org/https://doi.org/10.1016/j.neuroimage.2019.05.040
25. Starck, S., et al.: Using uk biobank data to establish population-specific atlases from whole body mri. Communications Medicine **4**(1), 237 (2024)
26. Taoudi-Benchekroun, Y., et al.: Predicting age and clinical risk from the neonatal connectome. NeuroImage **257**, 119319 (2022). https://doi.org/10.1016/j.neuroimage.2022.119319
27. Uus, A.U., et al.: Bounti: Brain volumetry and automated parcellation for 3d fetal mri (2023). https://doi.org/10.7554/elife.88818.1
28. Yoon, J.S., Zhang, C., Suk, H.I., Guo, J., Li, X.: SADM: Sequence-Aware Diffusion Model for Longitudinal Medical Image Generation, p. 388–400. Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-34048-2_30

29. Yuan, X., Cheng, J., Hu, D., Wu, Z., Wang, L., Lin, W., Li, G.: Longitudinally consistent individualized prediction of infant cortical morphological development. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 447–457. Springer Nature Switzerland, Cham (2024)
30. Zhang, K., Chen, G., Huang, S., Zhu, F., Ding, Z., Shen, D.: Development-driven diffusion model for longitudinal prediction of fetal brain mri with unpaired data. IEEE Transactions on Medical Imaging pp. 1–1 (2024). https://doi.org/10.1109/TMI.2024.3496860