

CINA: Conditional Implicit Neural Atlas for Spatio-Temporal Representation of Fetal Brains

Maik Dannecker¹, Vanessa Kyriakopoulou³, Lucilio Cordero-Grande⁴, Anthony N. Price³, Joseph V. Hajnal³, and Daniel Rueckert^{1,2}

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 Centre for the Developing Brain, School of Biomedical Engineering and Imaging Sciences, King's College London, United Kingdom

⁴ Biomedical Image Technologies, ETSI Telecomunicación, Universidad Politécnica de Madrid, Spain

Abstract. We introduce a conditional implicit neural atlas (CINA) for spatio-temporal atlas generation from Magnetic Resonance Images (MRI) of the neurotypical and pathological fetal brain, that is fully independent of affine or non-rigid registration. During training, CINA learns a general representation of the fetal brain and encodes subject specific information into latent code. After training, CINA can construct a faithful atlas with tissue probability maps of the fetal brain for any gestational age (GA) and anatomical variation covered within the training domain. Thus, CINA is competent to represent both, neurotypical and pathological brains. Furthermore, a trained CINA model can be fit to brain MRI of unseen subjects via test-time optimization of the latent code. CINA can then produce probabilistic tissue maps tailored to a particular subject. We evaluate our method on a total of 198 T2 weighted MRI of normal and abnormal fetal brains from the dHCP and FeTA datasets. We demonstrate CINA's capability to represent a fetal brain atlas that can be flexibly conditioned on GA and on anatomical variations like ventricular volume or degree of cortical folding, making it a suitable tool for modeling both neurotypical and pathological brains. We quantify the fidelity of our atlas by means of tissue segmentation and age prediction and compare it to an established baseline. CINA demonstrates superior accuracy for neurotypical brains and pathological brains with ventriculomegaly. Moreover, CINA scores a mean absolute error of 0.23 weeks in fetal brain age prediction, further confirming an accurate representation of fetal brain development.

Keywords: fetal brain imaging · implicit neural representation · conditional brain atlas.

1 Introduction

The spatio-temporal analysis of fetal brain structures from MRI constitutes a critical cornerstone in understanding normal and abnormal growth patterns of

the early developing brain. Fetal brain atlases serve as important instrument for precise characterization of neural development, facilitating numerous applications in medical image analysis including image segmentation and shape analysis. As fetal brains experience rapid development throughout pregnancy, they demand an atlas with high resolution in space and time to be adequately modeled across gestational age (GA). Existent work focuses on classical registration based approaches to map a set of brains of similar GA to a common reference [8,10,12,14,24,25,26]. These approaches, however, come with notable limitations. They rely on compute intensive, non-rigid registration schemes and are restricted to discrete time points and fixed spatial resolution. More critically, the atlases typically only represent a normative brain and fail to adequately represent the development and characteristics of brains in the presence of pathologies [20]. Studying brains with distinct pathologies requires an atlas based on brains exhibiting similar conditions. However, datasets for such specific pathologies are usually limited. Hence, established methods struggle to generate a temporal unbiased atlas of the fetal brain that faithfully represents pathological growth patterns.

Contribution. We propose CINA, a conditional implicit neural atlas based on Implicit Neural Representation (INR) that learns a continuous, spatio-temporal representation of the fetal brain with corresponding tissue probability maps for brain segmentation. Our method facilitates flexible atlas generation of arbitrary resolution in time and space. Additionally, CINA can condition on specific anatomical features during training and therefore interpolate within the anatomical feature space for atlas generation. E.g., in training, CINA conditions on brain Gyrfication Index or ventricular volume and can then create brain atlases with varying degrees of cortical folding or with ventricles of different sizes, respectively. We evaluate CINA on a total of 198 fetal brain MRI, including neurotypical brains from the dHCP dataset [21] and pathological brains with ventriculomegaly (VM) from the FeTA dataset [20]. CINA provides accurate tissue probability maps, especially for pathological brains with VM, and demonstrates impressive age prediction accuracy with a mean absolute error of 0.23 weeks for neurotypical brains.

Related Work. Initial work employed pairwise affine registration with kernel regression across age to construct an atlas from MRI of preterm neonates [12]. Follow up work extended the registration model to non-rigid deformations using free-form deformation models [14,22,25] or diffeomorphic registration [24]. First work on a fetal brain atlas used polynomial fitting and non-rigid deformation to construct a spatio-temporal atlas of 20 neurotypical fetal brains [10]. This was followed by an FFD approach from [26] constructing a fetal brain atlas from MRI of 80 healthy fetuses spanning a GA of 22 to 38 weeks. In 2017, Gholipour et al. [8] constructed a spatio-temporal MRI atlas with corresponding brain tissue segmentation from 81 healthy fetuses of 21-37 weeks GA. However, these models are built from neurotypical populations and therefore inadequate to model pathological growth patterns in fetal brains.

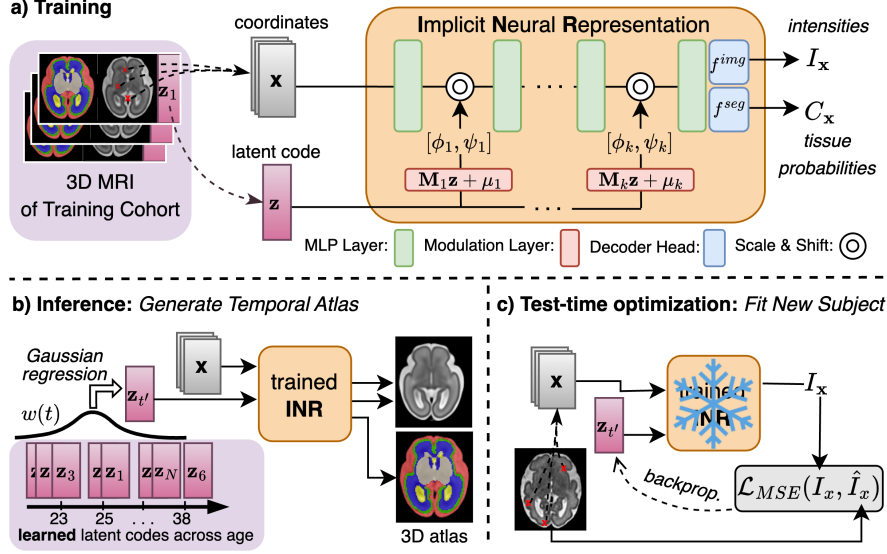


Fig. 1. CINA comprises three phases. **a)** In the training phase, CINA learns to reconstruct 3D MRI with tissue segmentations of N subjects. Each subject is assigned its unique *trainable* latent code \mathbf{z} . Thus, the INR learns general domain knowledge while subject specific information is represented in \mathbf{z} . **b)** CINA has learned a spatio-temporal representation of the fetal brain. To generate an atlas for time t , we regress a new \mathbf{z}_t from the trained latents $\{\mathbf{z}_i\}_{i=1}^N$. Then, a single forward pass yields the atlas with tissue probability maps. **c)** To fit the atlas to a new subject of age t' , we perform test-time optimization on the MRI intensities, learning a new latent \mathbf{z}'_t while freezing the INR. Finally, we get the atlas and probability maps, tailored to the subject, via a forward pass while conditioning on \mathbf{z}'_t .

Recently, neural network based method for conditional atlas creation [3,4,5] can provide a more flexible brain representation. Implicit Neural Representations (INRs) promise a further direction for conditional spatio-temporal atlas construction. INRs model data with a multilayer perceptron (MLP) as continuous functions of coordinates and comprise a resolution agnostic representation of the data [17,18,23,28,30]. Typically, an INR is (over-)fit on a specific data instance without exploiting any task-specific knowledge as prior, limiting the representation. However, multiple strategies to fit an INR to an entire dataset across a domain exist [2,6,15,16,19,29]. Here, the INR represents domain knowledge, whereas instance specific information is encoded in a instance specific learnable latent code. During training, the latent code can either be directly fed to the INR by concatenation to the input coordinates [16,19] or mapped to the weights of the INR layers using a hypernetwork [28,29]. Other work proposes to linearly map the latent code to shift and scale parameters which then modulate the output of the INR layers [6,15]. This approach demonstrates enhanced reconstruction quality during inference and fast training. The latent code itself can be optimized

through an encoder or auto-decoder approach [19,27,28]. In the auto-decoder approach, latent code is randomly initialized, usually Gaussian distributed close to zero, and optimized during training by backpropagation through the INR, end-to-end [19]. During inference on an unseen signal instance, a new initialized latent is fit to the signal while freezing the INR. The auto-decoder setup is fully resolution agnostic and fast to train, but slower for inference due to the latent code optimization.

2 Methodology

CINA is modeled as an auto-decoder approach, consisting of a fully-connected network with sinusoidal activation functions following [28]. To condition on latent code, we employ modulated layers as in [6,15]. The latent code $\mathbf{z} \in \mathbb{R}^D$ is mapped to scale $\phi \in \mathbb{R}^L$ and shift $\psi \in \mathbb{R}^L$ parameters through a linear modulation layer as $\begin{pmatrix} \phi \\ \psi \end{pmatrix} = \mathbf{M}\mathbf{z} + \boldsymbol{\mu}$, with $\mathbf{M} \in \mathbb{R}^{2L \times D}$ and $\boldsymbol{\mu} \in \mathbb{R}^{2L}$. The output of an INR layer with employed modulation equates to $\sin((\omega_0 \cdot \phi \cdot \mathbf{W}\mathbf{x} + \mathbf{b}) + \psi)$, where $\mathbf{W} \in \mathbb{R}^{L \times L}$ and $\mathbf{b} \in \mathbb{R}^L$ denote weights and biases of the INR layer, and ω_0 denotes a multiplicative term to better fit high frequency signals [28]. Different to [6], we do not want the modulation shift ψ to be scaled by ω_0 , as the latent code represents temporal changes of the brain and should predominantly capture low-frequency features to enforce a smooth latent representation. Similar to [19], each subject of the training cohort has its unique latent code \mathbf{z} which is initialized close to zero as $\mathbf{z} \sim \mathcal{N}(0, 10^{-2})$. This helps the model to learn a compact representation of \mathbf{z} , facilitating smooth interpolation within the latent code.

Training. Fig. 1 a) depicts the training process. We maximize the joint log posterior over all N training subjects, optimizing the respective latent codes $\{\mathbf{z}_i\}_{i=1}^N$ and the model parameters θ :

$$\arg \min_{\theta, \{\mathbf{z}_i\}_{i=1}^N} = \sum_{i=1}^N \left(\sum_{\mathbf{x} \in \mathbf{X}_i} \mathcal{L}_{MSE} \left(f_{\theta}^{img}(\mathbf{x}, \mathbf{z}_i), \hat{I}_{\mathbf{x}} \right) + \mathcal{L}_{CE} \left(f_{\theta}^{seg}(\mathbf{x}, \mathbf{z}_i), \hat{C}_{\mathbf{x}} \right) \right) \quad (1)$$

Here $\hat{I}_{\mathbf{x}}$ and $\hat{C}_{\mathbf{x}}$ denote the groundtruth intensity and tissue label of voxel \mathbf{x} , f_{θ}^{img} and f_{θ}^{seg} denote the decoder heads of the INR, and \mathcal{L}_{MSE} and \mathcal{L}_{CE} denote the mean-squared error and the cross entropy loss. We found that no explicit regularization for \mathbf{z} is necessary during training to enforce a compact latent representation.

Atlas Construction. After training, the INR has learned a general representation of the target domain, i.e., the fetal brain, having pushed all individual information into the subject specific latent codes $\{\mathbf{z}_i\}_{i=1}^N$. This is in particular the case for GA and confirmed by Fig. 3. Illustrated in Fig. 1 b), this allows us to generate a brain atlas for any given time point t by using a time regressed latent \mathbf{z}_t defined as

$$\mathbf{z}_t = \sum_{i=1}^N w(t, t_i) \mathbf{z}_i, \quad (2)$$

where w represents a Gaussian kernel with mean t and a standard-deviation σ , empirically set to $\sigma = 0.35$ for all experiments. Finally, by conditioning on the regressed latent \mathbf{z}_t , we get the atlas with intensities $I_{\mathbf{x}}$ and tissue probability maps $C_{\mathbf{x}}$ for voxels \mathbf{X} from the trained INR via a forward pass:

$$\text{INR}_{\theta}(\mathbf{X}|\mathbf{z}_t) = \left(f_{\theta}^{img}(\mathbf{X}, \mathbf{z}_t), f_{\theta}^{seg}(\mathbf{X}, \mathbf{z}_t) \right) = (I_{\mathbf{x}}, C_{\mathbf{x}}). \quad (3)$$

Atlas matching to new subjects. To match the atlas to a new subject of gestational age t' , we optimize a new latent code \mathbf{z}'_t to the MRI intensities $\hat{I}(\mathbf{X})$, while freezing the INR parameters θ , minimizing

$$\mathbf{z}_{t'} = \arg \min_{\mathbf{z}} \sum_{\mathbf{x} \in \mathbf{X}} \mathcal{L} \left(f_{\theta}^{img}(\mathbf{X}, \mathbf{z}), \hat{I}(\mathbf{x}) \right) + \frac{1}{\sigma^2} \|\mathbf{z}\|_2^2, \quad (4)$$

where $\mathbf{z}_{t'}$ is initialized according to Eq. 2 or as $\mathbf{z}_{t'} \sim \mathcal{N}(0, 10^{-2})$ should the subject's age be unknown. To enforce a compact representation of \mathbf{z} and to prevent overfitting, we apply L_2 regularization for inference. We can also use secondary information for latent code optimization, e.g., tissue segmentation, to gain a richer representation. In this work, however, we only use the scan intensities to demonstrate the methods capability to infer accurate tissue probability maps from the optimized latent code. We generate these probability maps via Eq. 3.

2.1 Conditioning on Anatomical Characteristics

CINA *implicitly* conditions on the patterns of anatomical variations that are present in the training population, e.g., gestational age, and encodes these variations in the subject specific latent code \mathbf{z} . However, we can also force CINA to *explicitly* condition on specific brain anatomy to create an atlas tailored to a sub-population, e.g., subjects with ventriculomegaly (VM), by quantifying the desired brain anatomy by vector $\mathbf{y} \in \mathbb{R}^E$ and concatenating \mathbf{y} as additional dimension(s) to the latent code as $(\frac{\mathbf{y}}{\mathbf{z}}) \in \mathbb{R}^{D+E}$. During training we keep \mathbf{y} fixed. For inference on a new subject j , we make \mathbf{y}_j trainable and initialize it as $\mathbf{y}_j = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i$. We demonstrate the procedure on two examples: volume of the lateral ventricles (LV) and the degree of cortical folding.

Lateral Ventricular Volume. Conditioning on the LV volume is particularly helpful to adequately model brains with VM. Whereas diagnosis of VM usually involves the width of the atrial diameter, there also exists a linear relationship with the LV volume [13]. As the FeTA dataset provides segmentations of the LV but not the width of the atrial diameter, we decided to condition on the volume of the LV which is easily computed from the segmentation.

Cortical Folding. We further demonstrate the effectiveness of this approach on a more complex shape by conditioning on the cortex, specifically on the degree of cortical folding. We quantify the degree of folding with the Gyrfication Index [31] using the cortical gray matter segmentation provided with FeTA.

3 Experiments and Results

Datasets. In this study we use the two fetal brain datasets, dHCP [21] and FeTA [20]. From the dHCP dataset, we use T2 weighted 3D MRI of 138 fetal brains showing no abnormalities. Each 3D volume comes with an isotropic resolution of 0.5 mm and a quality controlled FreeSurfer [7,9] segmentation of 16 anatomical regions. The age range at the time of scan was between 22 and 38 weeks gestational age (GA). Further details on acquisition and pre-processing parameters can be found under [21].

The FeTA dataset contains T2 weighted 3D MRI of 80 neurotypical and pathological fetal brains. All scans are of 0.5 mm isotropic resolution and come with ground truth segmentations for 7 anatomical regions annotated by experts. For more details on acquisition and annotation protocol, we refer the reader to [20]. Of the 80 scans we discarded 20 scans due to insufficient quality leaving 30 normal and 30 abnormal cases within a age range of 21 to 36 weeks GA. The histogram in the supplementary material shows the detailed distributions of the two datasets.

Preprocessing. We rigidly aligned all brain volumes to MNI space. Next, we harmonized the label maps of dHCP and FeTA data by merging smaller regions so that we ended up with the six anatomical regions: Cerebrospinal fluid (CSF), cortical Gray Matter (cGM), White Matter (WM), Lateral Ventricles (LV), Cerebellum (CB), and Brainstem (BS).

Training Setup. We trained two CINA models. One on a population of 128 healthy fetal subjects from the dHCP dataset [21] and one on a population of 25 healthy and 25 pathological subjects from the FeTA dataset [20]. We configured CINA with 5 hidden layers of size 512, modulation layers of size 330 after the first, third, and fifth hidden layer, and latent code size of $\mathbf{z} = 330$. We trained the model for 200 epochs using ADAM [11] optimizer with a learning rate of $2e^{-4}$ and decay of 0.98 every epoch. Training took 2.5 h on an A6000 GPU.

3.1 Atlas Based Brain Segmentation and Age Prediction

This section evaluates the representation fidelity of CINA and compares it with the fetal brain atlas from the Computational Radiology Laboratory (CRL) at Boston Children’s Hospital [8]. We test brain age prediction and segmentation. For segmentation, we register the CRL atlas to test subjects using ANTs [1]. Next, we project the CRL segmentation on the test brain and compute the dice coefficient between the projected and groundtruth segmentation. CINA does not require a registration. Instead, we get the segmentation by fitting a new latent code $\mathbf{z}_t \sim \mathcal{N}(0, 10^{-2})$ to the test subject as described in section 2 and Fig. 1 c). Note, we only fit the MRI intensities, not the ground truth segmentation. We get the predicted segmentation via Eq. 3. The inference with CINA took on average 10 seconds per subject (GPU: NVIDIA A6000) compared to 8 seconds for the traditional registration approach using ANTs (CPU: AMD 5975WX 32-Cores). Table 1 presents the results. Both methods show good accuracy across all tissue classes for neurotypical cases cGM represents the most challenging




Fig. 2. a) CINA and the CRL Atlas are matched with a target subject. The CRL atlas fails to adequately represent the enlarged ventricles. b+c) Interpolation on the conditioned anatomy for a fixed time point. Note, top left and right brains start to show corruptions as these extremes were never encountered during training (extrapolation).

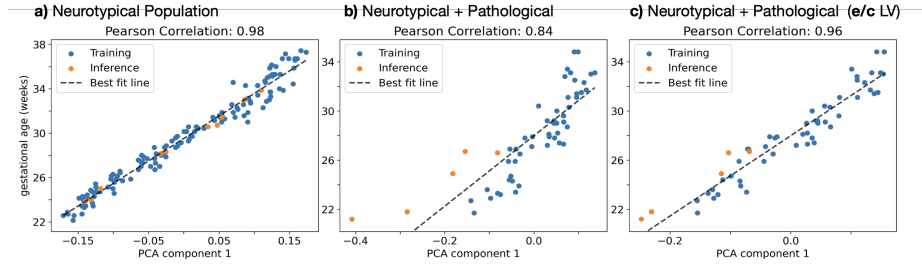


Fig. 3. Learned representation of gestational age (GA). The first PCA component of a latent code is mapped against the corresponding subject’s GA. Blue and orange represent latent codes learned during training and inference. a) A nearly perfect encoding of GA can be observed for the neurotypical population. b) For a mixed population, containing brains with ventriculomegaly, the correlation degrades. c) Explicitly conditioning (e/c) on ventricular volume improves the encoding. Note, in this setup, CINA has never seen a subject’s GA.

class due to its complex topology. For pathological cases with VM, i.e. brains showing enlarged lateral ventricles (LV), the CRL atlas shows clear deficits for the LV. CINA on the other hand demonstrates better modeling of the LV with a notably higher dice score. Moreover, if we model the LV volume explicitly, denoted as **e/c**, we observe further improvements. Fig. 2 a) visualizes a case with VM together with the predictions of CINA and the CRL atlas. Unlike the CRL atlas, CINA adequately captures the LV. Using the CRL atlas we predict a subject’s brain age by mapping it to the atlas of similar brain volume and then interpolating the age. For CINA, we regress the age from the learned latent code, utilizing the learned age correlation presented in Fig. 3.

Table 1. Dice score and mean absolute error of GA (MAE-GA) of CINA and the CRL atlas. Mean values, with standard deviation in parentheses, for four of six brain regions of 10 healthy fetal subjects from dHCP[21] and 5 pathological fetal subjects from FeTA with VM[20]. **e/c** denotes explicit conditioning of the LV volume as described in section 2.1. \overline{DSC} denotes mean dice of all *six* regions. (All results in Supplementary Table.

Method	CSF	cGM	LV	\overline{DSC}	MAE-GA
Neurotypical Brains (dHCP)					
CRL [8]	0.83 (0.03)	0.65 (0.04)	0.66 (0.06)	0.76 (0.03)	1.09 (0.59)
CINA	0.86 (0.04)	0.69 (0.08)	0.79 (0.06)	0.83 (0.04)	0.23 (0.21)
Pathological Brains with Ventriculomegaly (FeTA)					
CRL	0.74 (0.05)	0.43 (0.10)	0.41 (0.16)	0.60 (0.12)	1.14 (0.82)
CINA	0.77 (0.08)	0.50 (0.14)	0.73 (0.20)	0.67 (0.18)	2.61 (1.34)
CINA (e/c LV)	0.79 (0.08)	0.52 (0.21)	0.81 (0.07)	0.70 (0.18)	1.24, (0.44)




Fig. 4. Temporal fetal brain atlas from 23 to 38 GA with corresponding probability maps of the cortical grey matter.

3.2 Atlas Construction

CINA implicitly represents a spatial- and time-continuous atlas. Contrary to 3D image based atlases that are only available for fixed, discrete time points and resolution, CINA can generate atlases for any time point and resolution. Fig. 4 presents a temporal atlas generated with CINA for GA 23 to 38. Moreover, CINA can generate the temporal atlases while simultaneously conditioning on learned anatomical characteristics, visualized by Fig. 2 b+c) for LV volume and Gyrfication Index of the cortex for a specific time point.

4 Conclusion

We have introduced CINA, a novel conditional implicit neural atlas designed for the continuous spatial and temporal generation of 3D fetal brain atlases. CINA can model key anatomical features, including ventricular volume or degree of cortical folding. This facilitates accurate representation of neurotypical brains, but also of pathological cases where classical atlases often fall short. Moreover, CINA achieves a high compression rate, encapsulating the representation of hundreds of brain states across varying resolutions, gestational age, and anatomical expressions, within a mere 30 MB footprint. This compact yet versatile format enables end-users to generate detailed brain atlases in real-time, tailored to their specific needs.

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. Bauer, M., Dupont, E., Brock, A., Rosenbaum, D., Schwarz, J.R., Kim, H.: Spatial functa: Scaling functa to imagenet classification and generation. *arXiv preprint arXiv:2302.03130* (2023)
3. Cheng, J., Dalca, A.V., Zöllei, L.: Unbiased atlas construction for neonatal cortical surfaces via unsupervised learning. In: *Medical Ultrasound, and Preterm, Perinatal and Paediatric Image Analysis, PIPPI 2020, MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings 1*. pp. 334–342. Springer (2020)
4. Dalca, A., Rakic, M., Gutttag, J., Sabuncu, M.: Learning conditional deformable templates with convolutional networks. *Advances in neural information processing systems* **32** (2019)
5. Dey, N., Ren, M., Dalca, A.V., Gerig, G.: Generative adversarial registration for improved conditional deformable templates. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 3929–3941 (2019)
6. Dupont, E., Kim, H., Eslami, S., Rezende, D., Rosenbaum, D.: From data to functa: Your data point is a function and you can treat it like one. *arXiv preprint arXiv:2201.12204* (2022)
7. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
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. Glasser, M.F., et al.: The minimal preprocessing pipelines for the human connectome project. *Neuroimage* **80**, 105–124 (2013)
10. Habas, P.A., Kim, K., Corbett-Detig, J.M., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C.: A spatiotemporal atlas of mr intensity, tissue probability and shape of the fetal brain with application to segmentation. *Neuroimage* **53**(2), 460–470 (2010)

11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
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. Ma, H.L., Zhao, S.X., Lv, F.R., Zhang, Z.W., Xiao, Y.H., Sheng, B.: Volume growth trend and correlation of atrial diameter with lateral ventricular volume in normal fetus and fetus with ventriculomegaly: A strobe compliant article. *Medicine* **98**(26) (2019)
14. Makropoulos, A., Aljabar, P., Wright, R., Huning, B., Merchant, N., Arichi, T., Tusor, N., Hajnal, J.V., Edwards, A.D., Counsell, S.J., Rueckert, D.: Regional growth and atlasing of the developing human brain. *Neuroimage* **125**, 456–478 (2016). <https://doi.org/10.1016/j.neuroimage.2015.10.047>
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. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4460–4470 (2019)
17. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
18. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)* **41**(4), 1–15 (2022)
19. 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)
20. Payette, K., et al.: An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific Data* **8**(1), 167 (2021). <https://doi.org/10.1038/s41597-021-00946-3>, <https://doi.org/10.1038/s41597-021-00946-3>
21. Price, A.N., et al.: The developing human connectome project (dhcp): fetal acquisition protocol. In: *Proceedings of the annual meeting of the International Society of Magnetic Resonance in Medicine (ISMRM)*. vol. 244 (2019)
22. 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)
23. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18507–18516 (2023)
24. Schuh, A., et al.: Construction of a 4d brain atlas and growth model using diffeomorphic registration. In: *Spatio-temporal Image Analysis for Longitudinal and Time-Series Image Data: Third International Workshop, STIA 2014, Held in Conjunction with MICCAI 2014, Boston, MA, USA, September 18, 2014, Revised Selected Papers 3*. pp. 27–37. Springer (2015)

25. 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>
26. 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)
27. Sitzmann, V., Chan, E., Tucker, R., Snavely, N., Wetzstein, G.: Metasdf: Meta-learning signed distance functions. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *NeurIPS*. vol. 33, pp. 10136–10147. Curran Associates, Inc. (2020)
28. 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)
29. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems* **32** (2019)
30. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems* **33**, 7537–7547 (2020)
31. Zilles, K., Armstrong, E., Schleicher, A., Kretschmann, H.J.: The human pattern of gyrification in the cerebral cortex. *Anatomy and embryology* **179**, 173–179 (1988)