This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

ALRAU: Atlas Learning from Routine Antenatal Ultrasound

Johannes Tischer[1,2,3][0009−0004−0870−0914], Stephan Anzengruber[4][0000−0002−1531−1333], Martin Mienkina[3,4][0000−0001−6122−2908], Tim Dorittke[3,5][0000−0002−9121−3401], Julia Binder[3,5][0000−0002−5725−9270], Gregor Kasprian[3,6][0000−0003−3858−3347], Georg Langs[1,2,3][0000−0002−5536−6873], and Roxane Licandro[1,2,3][0000−0001−9066−4473]

[1] Computational Imaging Research Lab, ELIA Group, Department of Biomedical Imaging and Image-guided Therapy, Medical University of Vienna (MedUni Wien), Austria johannes.tischer@meduniwien.ac.at
[2] Comprehensive Center for Artificial Intelligence in Medicine, MedUni Wien, Austria
[3] Christian Doppler Laboratory for Mathematical Modelling and Simulation of Next-Generation Medical Ultrasound Devices, Austria
[4] GE Healthcare, Zipf, Austria
[5] Department of Obstetrics and Gynecology, MedUni Wien, Austria
[6] Division of General Radiology, Department of Biomedical Imaging and Image-guided therapy, MedUni Wien, Austria

Abstract. Ultrasonography (US) is the primary modality in prenatal assessment, as it provides insights into soft tissue characteristics, is widely available and cost-effective. However, variations in acoustic impedance across tissues, fetal motion, and probe motion degrade image quality, hindering objective assessment of fetal brain development. Brain atlases offer a standardized reference framework by aligning subjects within a common coordinate system, enabling consistent evaluation across populations. We propose ALRAU, a learning-based framework for constructing an artifact-reduced, time continuous fetal brain atlas from routine abdominal US data. Trained on over 300 volumes spanning 18–28 weeks of gestation, ALRAU achieves competitive image quality and pronounced anatomical definition of fine structures (Peak Signal-to-Noise Ratio: 24.6 dB). Age prediction on the constructed atlases confirms that age-discriminative anatomical features are preserved after registration. The mean absolute difference in predicted gestational age between test subjects and warped templates is 0.3 weeks. Warping the age-matched atlas to unseen subjects achieves robust alignment (Normalized Cross-Correlation: 0.81±0.07, Structural Similarity Index: 0.82±0.03) with perfect deformation invertibility, enabling reliable registration without anatomical distortion. This approach enables real-time, volumetric assessment of fetal brain development and supports informed prenatal decision-making. The model code is available at `https://github.com/cirmuw/ALRAU`.

Keywords: Antenatal Ultrasound · Brain Atlas · Fetal Imaging.

1 Introduction

During pregnancy, Ultrasonography (US) is routinely performed from the 10^{th} Gestational Week (GW) onward to assess maternal health, placental characteristics, and fetal anatomy. Multiple abdominal or transvaginal US acquisitions are performed to assess biometric measures [22,23]. Specifically, neurosonography enables the examination of the central nervous system and detection of anomalies with high accuracy [22]. However, a comprehensive fetal brain assessment is time-consuming (up to 1 h [9]) and exhibits significant inter-operator variability in anomaly detection rates [9,22]. In addition, the heterogeneity of image artifacts disturbs the image intensity profile and confounds intra- and inter-subject comparability.

Brain atlases offer a promising approach to address these limitations by providing a (distinct) population-based anatomical reference that encodes spatial relationships and morphological variability of brain structures [7,8,10,19,24]. The propagation of atlas-based information like anatomical labels or landmarks to an individual fetal brain US enables the automated extraction of standardized biometric measurements and volumetric estimates of brain structures, thereby reducing operator dependence and improving reproducibility [7,8]. However, atlas construction typically relies on high-quality input data and standardized imaging protocols, requiring considerable effort in data curation and selection. Specifically, in US, a substantial number of acquired volumes are discarded because inclusion criteria are not met [18]. This raises a fundamental question: can an anatomically consistent brain atlas be learned directly from routine clinical data?

Related Work. Fetal brain atlas construction has been extensively studied in magnetic resonance imaging [6], with approaches ranging from traditional group-wise registration [10,24] to advanced deep-learning models that capture continuous spatiotemporal anatomical variation [7,25,26].

In contrast, the construction of a US fetal brain atlas remains largely unexplored. Namburete et al. [18,19], employs diffeomorphic demons–based registration that integrates both intensity and edge features. This approach enables group-wise alignment of 3D US volumes and allows for the generation of a population-based representation from the 14^{th} to the 31^{st} week of gestation. However, their atlas is constructed at discrete time points using exclusively high-quality abdominal US volumes with Crown–Rump Length (CRL) corrected gestational age (GA) measurements, which are not reflective of most retrospective datasets. Cabezas et al. [5] constructed a 2D fetal brain atlas at 20 weeks gestation by aligning 104 images to a common space using ellipse fitting and affine registration. Hereby generated probabilistic maps of six anatomical structures via population averaging. The resulting atlas is limited to a 2D representation and relies on affine registration without additional deformable refinement. In addition, it requires manual landmark annotation (n=20) per image. Bastiaansen et al. [4] presented an embryonic brain atlas spanning 8–12 GW. They constructed a conditional, continuous atlas by optimizing a deep-learning-based image registration framework [3] that encourages image similarity while constraining allowed deforma-

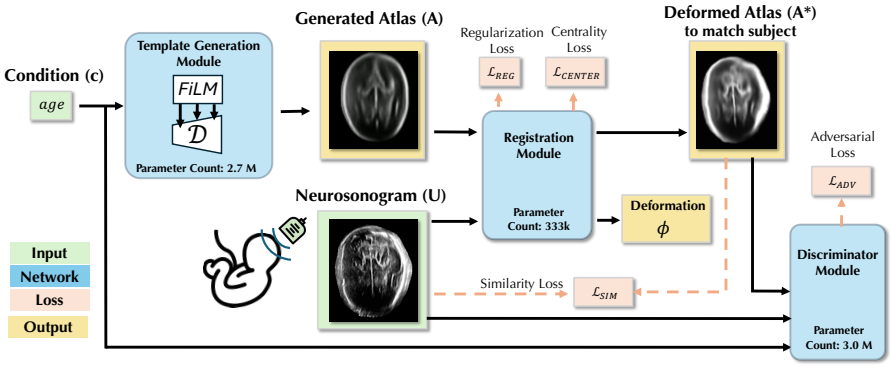

Fig. 1: ALRAU network architecture for conditional US brain atlas learning

tions.

Contribution. We introduce ALRAU, a template generation framework that constructs a conditional, continuous US atlas of the developing fetal brain from routine imaging data. To model anatomical variability, we train ALRAU on a diverse, multi-center dataset of 350 fetal brain neurosonograms from 10 countries across 5 continents, covering 18 to 28 GWs. The atlas depicts sharp representations of the fetal brain, capturing both global structures and subtle anatomical variations. We evaluate ALRAU on a total of 117 fetal brain US, by quantitatively assessing the accuracy of subject-to-atlas alignment and demonstrating its robustness across gestation.

2 Methodology

ALRAU learns a gestational-age-conditioned brain atlas from a population of fetal US volumes without a predefined template. It couples a generative atlas module, a diffeomorphic registration module, and a projection discriminator (see Fig. 1).

Template Generation Module. The generator synthesises an age-specific atlas A from a single scalar condition c (gestation age, normalized between 0 and 1). Therefore, c is first encoded by a four-layer Multi-Layer Perceptron (MLP) into a 64-dimensional vector. A small learnable seed tensor at $1/8$ resolution is then progressively decoded to full resolution through five residual blocks, followed by three upsampling decoder blocks. Every convolutional block is modulated by Feature-wise Linear Modulation (FiLM) layers driven by the conditioning vector, so that age controls the template's appearance and shape throughout the network [21]. The output head emits an intensity template A.

Atlas-to-Subject Alignment. The registration network predicts a dense deformation field ϕ that aligns template A to a subject volume U [3]. A U-Net CNN encodes the concatenated image pair (A,U) and outputs a velocity field v. The velocity field is integrated via the scaling-and-squaring method to produce

the deformation field ϕ. Subsequently, it is used to warp the generated atlas using a spatial transformer network ($A \circ \phi$) [12].

Discriminator Module: A projection conditional discriminator [17] provides adversarial feedback on warped templates, using spectral normalization and age conditioning through learned embeddings. Its training uses a least-squares GAN objective with one-sided label smoothing and a lazy R1 gradient penalty for stability [16].

Loss Function. The total loss is:

$$\mathcal{L}_{TOTAL} = \mathcal{L}_{SIM} + \alpha\,\mathcal{L}_{REG} + \beta\,\mathcal{L}_{CENTER} + \gamma\,\mathcal{L}_{ADV} \tag{1}$$

where the local normalized cross correlation (NCC) and $L1$ loss acts as similarity metric $\mathcal{L}_{SIM} = \mathrm{NCC}(U, A \circ \phi) + L1(U, A \circ \phi)$ between the training samples and the deformed atlas. $\mathcal{L}_{REG}$ enforces smoothness of the deformation ϕ via L2 gradient penalty. The global centrality loss $\mathcal{L}_{CENTER}$ penalizes the squared magnitude of the mean deformation field, encouraging the average deformation across all subjects to remain near zero so the learned atlas stays unbiased and spatially centered. Finally, the discriminator provides adversarial feedback $\mathcal{L}_{ADV}$ enforcing realistic, sharp templates.

3 Materials

Dataset. The dataset comprises abdominal 3D US scans of the fetal brain from 421 pregnant women, totaling 550 acquisitions, covering 18 to 28 weeks of gestation (see Fig. S1a). All volumes were acquired during routine clinical exam and reflected typical US image quality with common imperfections such as acoustic shadowing, speckle noise, and restricted fields of view. The voxel spacing ranges from 0.056 to 0,711 mm ($\overline{x} = 0.296$ mm, $\sigma = 0.064$ mm). The data originates from 10 different countries across five continents (see Fig. S1b) and 2 US system. Detailed information in the Appendix (see Tab. S2 and S3).

Preprocessing. The US volumes are resampled to an isotropic voxel-spacing of 0.5 mm and subsequently center-cropped to a cubic size of 224^3 voxels. Apart from the imaging data, no additional information was available. Following common practice in MRI atlas construction, we operate on brain-extracted images to focus on the brain [7,25,26]. To obtain 3D brain masks, we use embeddings from FetalCLIP [15], a vision-language foundation model for fetal US, as input features to a UNETR segmentation network [11]. Since FetalCLIP was trained on 2D mid-axial slices, the UNETR is trained in a supervised, slice-wise manner using a single manually annotated volume. To improve generalization beyond this limited supervision, extensive patch-based data augmentation is applied, including rotation, translation, scaling, intensity perturbations, and elastic deformations, until convergence. As a second step, we generate GA labels. Here, the zero-shot capabilities of FetalCLIP were not suitable for our approach, particularly at early GW (<20GW) [15]. Therefore, we estimate GA from the biparietal diameter (BPD) and head circumference (HC) [20] based on the measurements of the 3D brain masks, retaining only samples for which $\hat{GA}_{BPD} - \hat{GA}_{HC} < 1$ GW.

Baselines. To ensure a comprehensive comparison, we select representative methods from distinct registration paradigms. As a traditional deformable registration framework, we employ ANTs [2] on 32 CPU cores using Cross Correlation (CC) as the similarity metric and a pyramidal four-level coarse-to-fine setup. In addition, we used FireANTs [13] with default settings, except for 20 iterations of template construction using the greedy algorithm and the same pyramidal setup as ANTs. For learning-based approaches, we evaluate CINeMA [7], which leverages implicit neural representations to model continuous spatial transformations. CINeMA is trained for 1 epoch and implemented with a latent space of size 256x3x3x3, 5 hidden layers with 1024 units each, and modulation layers applied to hidden layers 1, 3, and 5. Finally, adaptation to new subjects during testing was performed by optimizing the latent code over 10 epochs [7].
Implementation Details. The dataset was stratified by GA, where for each GA, 60% of subjects were allocated to the training set, and 20% each to the validation and test sets (see Fig. S1). The network is implemented using PyTorch v2.5.1 and CUDA v12.1. The model was trained with a batch size of 1 on an NVIDIA A100 GPU and optimized with ADAM [14] (l_r=1e-4). Larger batch sizes are architecturally supported. We use a batch size of 1 to maximize spatial resolution within the available memory. The training of ALRAU is performed in three stages. For the first 10 epochs, $\mathcal{L}_{SIM}$ is solely L1 (see Eq. 1) to establish an initial template, essentially a linear average. This step is important to stabilize later NCC training, which relies on normalized intensities. In the second step (30 epochs), the contribution of the L1 term gradually fades out, while NCC ramps up. In the final stage, the adversarial weight increases, while the $\mathcal{L}_{SIM}$ operates on NCC only. This combination drives the network from learning global, blurry representations to capturing local, sharp details. Training was stopped early if the validation performance plateaued ($\Delta \leq 1\,\%$) over 40 consecutive epochs. The hyperparameters of Eq. 1 are set to $\alpha = 1.5$, $\beta = 0.5$, $\gamma = 0.05$. The window size of the localized NCC is set to 15 (image size=224) or 9 (image size=128).

4 US Fetal Brain Atlas

The constructed atlases in axial orientation are visualized for all baselines and our approach in Fig. 2. Quantitative assessment employed training time (t_N), inference time (t_i). Template image quality was assessed using the Entropy Focus Criterion (EFC) [1] and Peak Signal-to-Noise Ratio (PSNR) after warping the atlas into the subject space. To assess whether the top-performing approach significantly outperformed all competing methods, we conducted one-sided Wilcoxon signed-rank tests against each competitor for every metric ($\alpha < 0.05$) (see Tab. 1). Across all methods and GA the atlases capture the increase in brain size and morphological complexity, with the Choroid Plexus (CP), Cavum Septi Pellucidi (CSP), Frontal Horns (FH), Occipital Horns (OH), and Atria (At) consistently resolvable (see Fig. 2 and Fig. 4). However, the reconstruction quality is systematically poorer in the hemisphere proximal to the transducer, reflecting the reverberation artifact inherited from the US samples. The

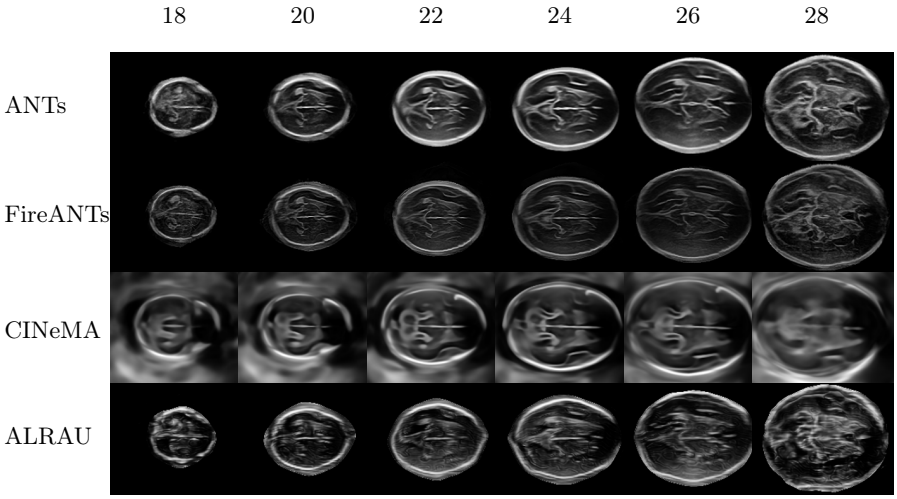

Fig. 2: Grid plot illustrating the learned US fetal brain templates generated by the baseline approaches and ALRAU for GAs ranging from 18 to 28 weeks.

Table 1: Quantitative evaluation of atlas construction performance for *ALRAU* and three baselines. t_N denotes the training time, and t_i denotes the total inference time for the test dataset. Significant better results ($\alpha \leq 0.05$) compared to all competitors are indicated with *.

Method	t_N [h] (↓)	t_i [h] (↓)	EFC (↓)	SSIM (↑)	PSNR [dB] (↑)	NCC (↑)	Def. [mm] (↓)
ANTs	226.8	20.48	0.24 ± 0.05	$\mathbf{0.85 \pm 0.07}^*$	$\mathbf{26.6 \pm 2.35}^*$	$\mathbf{0.90 \pm 0.04}^*$	$\mathbf{0.69 \pm 0.28}^*$
FireANTs	3.75	0.17	0.22 ± 0.05	0.69 ± 0.06	21.7 ± 1.79	0.57 ± 0.10	2.04 ± 0.35
CINeMA	**0.33**	4.16	0.50 ± 0.05	0.61 ± 0.10	20.6 ± 1.86	0.48 ± 0.12	-
ALRAU	17.5	**0.17**	$\mathbf{0.22 \pm 0.03}$	0.82 ± 0.03	24.6 ± 1.90	0.81 ± 0.07	2.56 ± 1.02

optimization-based baselines ANTs and FireANTs produce the highest-quality reconstructions with sharp folding patterns and strong anatomical fidelity. The additional Laplacian filtering by FireANTs yields cleaner boundaries at a reduced construction time of 3.75 h. CINeMA reaches the fastest construction time through continuous modeling across GA (0.33 h), but over-smooths finer structures, evident in its elevated EFC and reduced PSNR (see Tab. 1). The continuous representation trades boundary detail for temporal smoothness. ALRAU balances both, retaining the continuous GA modeling of CINeMA while preserving boundary sharpness comparable to the optimization-based methods (EFC 0.22, PSNR 24.6 dB), at an inference cost of only 0.17 h. In practical terms, this makes age-specific atlas generation feasible on demand without the per-cohort re-optimization required by ANTs and FireANTs, suitable for downstream segmentation and biometry.

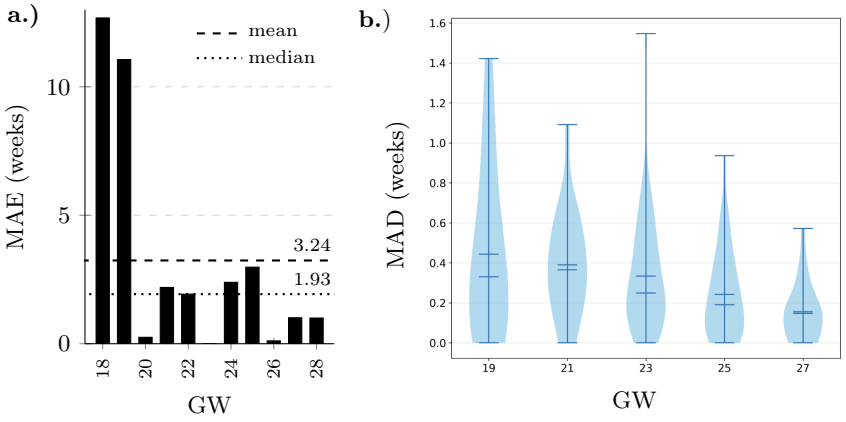

Fig. 3: Age-discriminative feature preservation performance. **a.)** Mean Absolute Error (MAE) between template prediction condition and FetalCLIP zero-shot predicted GA is reported **b.)** Violinplot of Mean Absolute Difference (MAD) between test data based GA predictions and deformed template GA predictions.

4.1 Age Prediction

To assess age-discriminative feature preservation in ALRAU, we employed FetalCLIP in a zero-shot setting to predict GA at discrete classes. Two complementary evaluations were performed. First, the central slice of each age-specific template is extracted and evaluated using FetalCLIP. Second, GA predictions are obtained for the test dataset (n=117) and each age-matched template warped into the corresponding subject space.

In Fig. 3a the Mean Absolute Error (MAE) between template label and predicted GA is reported. It is observable that the MAE in GA prediction varies across GW, with younger GW (<20 GW) exhibiting higher prediction errors, reflecting FetalCLIP's underrepresentation of early GAs in its training embeddings. For this reason, we decided to obtain predictions for both, the original test subject and the warped template. The Mean Absolute Difference (MAD) is calculated between the corresponding predictions, isolating anatomical differences and registration-driven variation, while reducing the uncertainty of FetalCLIP. GA predictions derived from the warped atlases closely matched those GA predictions obtained from the corresponding test data subjects, with the MAD decreasing as GA increases (see Fig. 3b). The close agreement at older GAs (>25 GWs) indicates that the registration successfully transfers the age-discriminative anatomical detail of the templates into subject space. This is in contrast to the increasing complexity of the developing brain, as it increases in total brain volume and in cortical morphology changes. One possibility is that mid-gestation and later stages, such as the onset of cortical folding, are more

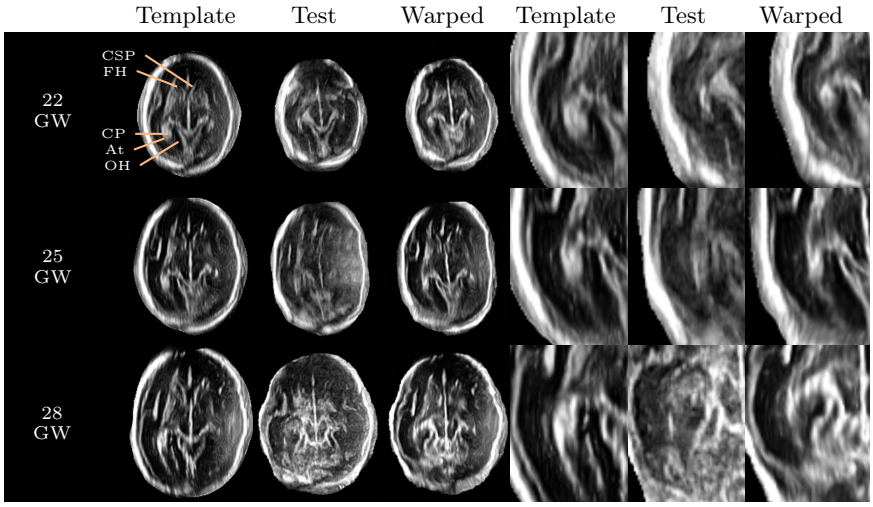

Fig. 4: Qualitative results of the age-specific fetal brain atlas (left column), an unseen test subject (middle column), and the atlas warped into the subject space (right column). On the right, the zoomed versions, respectively.

anatomically distinct, making their age-defining features easier to preserve and recover through prediction.

4.2 Atlas-to-Subject Alignment

Registration quality was assessed using NCC, structural similarity (SSIM), and average deformation (avg. Def) between the templates and the test dataset. ALRAU achieves NCC of 0.81 and SSIM of 0.82, approaching ANTs' performance (NCC=0.90, SSIM=0.85), while maintaining perfect invertibility with zero non-positive Jacobian determinants (JacDet). Average deformation was 2.56 ± 1.02 mm, higher than ANTs (0.69 ± 0.28 mm) but reflecting distinct registration characteristics rather than inferior alignment. The JacDet to evaluate local invertibility is 0.00 ± 0.00 for all approaches. Qualitative inspection of atlas-subject pairs (Fig. 4) reveals that SSIM and NCC metrics have inherent limitations: in artifact-rich regions where image information is absent, atlas information is necessarily preserved unchanged, artificially reducing these similarity metrics. Conversely, regions with large deformations are successfully matched when sufficient image quality enables reliable feature correspondence, demonstrating that deformation magnitude alone does not indicate registration failure. The substantial deformations observed in ALRAU predominantly occur in high-quality anatomical regions where folding patterns differ between template and subject, indicating successful adaptive alignment rather than registration error. We have presented ALRAU, a learning-based framework for constructing high-quality, artifact-reduced fetal brain atlases directly from routine clinical US data.

Trained on a diverse, multi-center dataset spanning 18–28 GWs, ALRAU demonstrates robust atlas-to-subject alignment with preserved age-discriminative features. The method opens new avenues for automated, standardized prenatal assessment, with potential to reduce inter-operator variability in fetal neurosonography and support more objective clinical decision-making. Future work needs to address label uncertainty and the limitations of intensity-based evaluation metrics such as SSIM and NCC by incorporating additional verification of anatomical correspondence using landmarks. Furthermore, robust alignment of 3D volumes acquired across different orientations should be considered. These extensions will further bridge the gap between research frameworks and clinical implementation, ultimately enabling real-time volumetric biometrics in the prenatal setting.

5 Conclusion

Acknowledgments. The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged. J.T. would like to thank GE Healthcare for providing the dataset used in this work and for their valuable support and collaboration.

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

References

1. Atkinson, D., Hill, D., Stoyle, P., Summers, P., Keevil, S.: Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion. IEEE TMI **16**(6) (1997). https://doi.org/10.1109/42.650886
2. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. NeuroImage **54** (2011). https://doi.org/10.1016/j.neuroimage.2010.09.025
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: A learning framework for deformable medical image registration. IEEE Transactions on Medical Imaging **38** (2019). https://doi.org/10.1109/TMI.2019.2897538
4. Bastiaansen, W.A., Rousian, M., Koning, A.H., Niessen, W.J., de Bakker, B.S., Steegers-Theunissen, R.P., Klein, S.: The 4D human embryonic brain atlas: Spatiotemporal atlas generation for rapid anatomical changes. Computerized Medical Imaging and Graphics **128** (2026). https://doi.org/10.1016/j.compmedimag.2026.102702
5. Cabezas, M., Diez, Y., Martinez-Diago, C., Maroto, A.: A benchmark for 2d foetal brain ultrasound analysis. Scientific Data **11**(1), 923 (2024). https://doi.org/10.1038/s41597-024-03774-3
6. Ciceri, T., Casartelli, L., Montano, F., Conte, S., Squarcina, L., Bertoldo, A., Agarwal, N., Brambilla, P., Peruzzo, D.: Fetal brain MRI atlases and datasets: A review. NeuroImage **292** (2024). https://doi.org/10.1016/j.neuroimage.2024.120603
7. Dannecker, M., Sideri-Lampretsa, V., Starck, S., Mihailov, A., Milh, M., Girard, N., Auzias, G., Rueckert, D.: CINeMA: Conditional implicit neural multi-modal atlas for a spatio-temporal representation of the perinatal brain. IEEE Transactions on Medical Imaging (2025). https://doi.org/10.1109/TMI.2025.3605194
8. Dey, N., Ren, M., Dalca, A.V., Gerig, G.: Generative adversarial registration for improved conditional deformable templates. In: ICCV (2021). https://doi.org/10.48550/arXiv.2105.04349
9. Gupta, V., Adams, J., Heiser, T., Iruretagoyena, J., Hoppe, K., Antony, K.: Detailed fetal anatomic ultrasound examination duration and association with body mass index. Obstetrics & Gynecology **134** (2019). https://doi.org/10.1097/AOG.0000000000003489
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) (2010). https://doi.org/10.1016/j.neuroimage.2010.06.054
11. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H., Xu, D.: UNETR: Transformers for 3D medical image segmentation (2021), https://arxiv.org/abs/2103.10504
12. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: NeurIPS (2015). https://doi.org/10.48550/arXiv.1506.02025
13. Jena, R., Chaudhari, P., Gee, J.C.: Fireants: Adaptive riemannian optimization for multi-scale diffeomorphic registration. Nature Communications (2024)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017), https://arxiv.org/abs/1412.6980
15. Maani, F., Saeed, N., Saleem, T.J., Farooq, Z., Alasmawi, H., Diehl, W., Mohammad, A., Waring, G., Valappil, S., Bricker, L., Yaqub, M.: FetalCLIP: A visual-

language foundation model for fetal ultrasound image analysis. arXiv preprint arXiv:2502.14807 (2025)
16. Mescheder, L., Geiger, A., Nowozin, S.: Which training methods for gans do actually converge? In: International Conference on Machine Learning (ICML). pp. 3481–3490 (2018)
17. Miyato, T., Koyama, M.: cgans with projection discriminator. ArXiv **abs/1802.05637** (2018), https://api.semanticscholar.org/CorpusID:3633127
18. Namburete, A., Papież, B., Fernandes, M., Wyburd, M., Hesse, L., Moser, F., Cheikh Ismail, L., Gunier, R., Squier, W., Ohuma, E., Carvalho, M., Jaffer, Y., Wu, Q., Lambert, A., Winsey, A., Restrepo-Méndez, M., Bertino, E., Purwar, M., Kennedy, S.: Normative spatiotemporal fetal brain maturation with satisfactory development at 2 years. Nature **623** (2023). https://doi.org/10.1038/s41586-023-06630-3
19. Namburete, A.I.L., van Kampen, R., Papageorghiou, A.T., Papież, B.W.: Multichannel groupwise registration to construct an ultrasound-specific fetal brain atlas. In: MICCAI DATRA PIPPI. Springer International Publishing, Cham (2018)
20. Papageorghiou, A.T., Ohuma, E.O., Altman, D.G., Todros, T., Ismail, L.C., Lambert, A., Jaffer, Y.A., Bertino, E., Gravett, M.G., Purwar, M., Noble, J.A., Pang, R., Victora, C.G., Barros, F.C., Carvalho, M., Salomon, L.J., Bhutta, Z.A., Kennedy, S.H., Villar, J.: International standards for fetal growth based on serial ultrasound measurements: the fetal growth longitudinal study of the intergrowth-21st project. The Lancet **384**(9946), 869–879 (2014). https://doi.org/https://doi.org/10.1016/S0140-6736(14)61490-2, https://www.sciencedirect.com/science/article/pii/S0140673614614902
21. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.C.: FiLM: Visual reasoning with a general conditioning layer. In: AAAI (2018). https://doi.org/10.48550/arXiv.1709.07871
22. Salomon, L., Alfirevic, Z., Berghella, V., Bilardo, C., Chalouhi, G., Costa, F.D.S., Hernandez-Andrade, E., Malinger, G., Munoz, H., Paladini, D., Prefumo, F., Sotiriadis, A., Toi, A., Lee, W., on behalf of the ISUOG Clinical Standards Committee: ISUOG practice guidelines (updated): performance of the routine mid-trimester fetal ultrasound scan. Ultrasound Obstet Gynecol **59** (2022)
23. Salomon, L., Alfirevic, Z., Bilardo, C., Chalouhi, G., Ghi, T., Kagan, K., Lau, T., Papageorghiou, A., Raine-Fenning, N., Stirnemann, J., Suresh, S., Tabor, A., Timor-Tritsch, I., Toi, A., Yeo, G.: ISUOG practice guidelines: performance of first-trimester fetal ultrasound scan. Ultrasound in Obstetrics & Gynecology **41** (2013). https://doi.org/10.1002/uog.12342
24. Serag, A., Aljabar, P., Ball, G., Counsell, S., Boardman, J., Rutherford, M., Edwards, D., Hajnal, J., Rueckert, D.: Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. NeuroImage **59** (2012). https://doi.org/10.1016/j.neuroimage.2011.09.062
25. Starck, S., Sideri-Lampretsa, V., Kainz, B., Menten, M., Mueller, T., Rueckert, D.: Diff-def: Diffusion-generated deformation fields for conditional atlases. arXiv preprint arXiv:2403.16776 (2024)
26. Tischer, J., Kienast, P., Stümpflen, M., Kasprian, G., Langs, G., Licandro, R.: Conditional fetal brain atlas learning for automatic tissue segmentation. In: Perinatal, Preterm and Paediatric Image Analysis (2025). https://doi.org/10.1007/978-3-032-05997-0_4