

FERN: A Fetal Echocardiography Registration Network for 2D-to-3D Alignment

Paula Ramirez

PAULA.RAMIREZ_GILLILAND@KCL.AC.UK

David Lloyd

DAVID.LLOYD@KCL.AC.UK

Jacqueline Matthew

JACQUELINE.MATTHEW@KCL.AC.UK

Reza Razavi

REZA.RAZAVI@KCL.AC.UK

Milou van Poppel

MILOU.VAN_POPPEL@KCL.AC.UK

Andrew King

ANDREW.KING@KCL.AC.UK

Maria Deprez

MARIA.DEPREZ@KCL.AC.UK

School of Biomedical Engineering and Imaging Sciences, King’s College London, London, UK.

Editors: Under Review for MIDL 2025

Abstract

2D Freehand echocardiography remains the primary imaging modality for routine fetal cardiac care, essential in the antenatal detection of Congenital Heart Disease (CHD). However, there is a lack of spatial context which requires 3D imaging. Current 3D methods, such as Spatio-Temporal Image Correlation (STIC), face limitations in success rate, image quality, and ease of use, and come at the cost of lower spatial and temporal resolution compared to 2D acquisitions. This work studies the feasibility of aligning real high spatial and temporal resolution 2D fetal echocardiography into a reference 3D space defined by lower resolution 3D STIC. FERN, a **Fetal Echocardiography Registration Network**, employs transformers for standard fetal echocardiography view alignment. The network is trained on simulated 2D slices derived from 3D volumes at end-diastole, and validated on real 2D acquisitions from fetuses with Coarctation of the Aorta and Right Aortic Arch diagnoses, achieving a mean Euclidean distance of 2.98 ± 1.27 mm on cardiac region-of-interest points between predicted and manually selected planes. Compared to manually aligned planes, improved image similarity to an average atlas is achieved, confirmed by blinded best plane selection. This work demonstrates that high spatial and temporal resolution 2D fetal echocardiography can be integrated into a 3D context provided by lower-resolution 3D acquisitions or fetal cardiac atlases, potentially resulting in a new 3D visualization tool for enhanced CHD diagnosis.

Keywords: Fetal echocardiography, 3D Fetal ultrasound, STIC, Congenital Heart Disease, 3D plane localisation, Fetal cardiology, Fetal cardiac imaging

1. Introduction

Congenital Heart Disease (CHD) is the most common group of malformations in fetuses and infants, with an incidence of around 6 per 1000 live births for moderate and severe forms (Hoffman and Kaplan, 2002). Antenatal diagnosis highly impacts patient prognosis, reducing mortality and morbidity (Tworetzky et al., 2000; Bonnet et al., 1999).

2D fetal echocardiography, the primary modality for antenatal CHD detection, offers high spatial and temporal resolution at a low cost. However, it lacks 3D context, requiring sonographers to mentally reconstruct 3D anatomy. Spatio-Temporal Image Correlation (STIC) (DeVore et al., 2003) offers 3D+time ultrasound for fetal cardiac evaluation via

automatic volume sweeps reordered by the cardiac cycle. However, despite over 20 years of availability, it is underused due to significant technological limitations.

STIC is challenging to acquire, with a limited success rate (Inamura et al., 2020) which may lengthen examination times. STIC is anisotropic, with lower spatial and temporal resolution than 2D acquisitions and susceptibility to motion-induced synchronization errors.

2D Fetal echocardiography acquisitions cover different cross-sectional levels of the cardiac anatomy using a standardized protocol (Carvalho et al., 2023), at an approximately axial position (see Fig. 3), avoiding shadowing from the fetal ribs. The standard views encompassing the cardiac anatomy include: 4 Chamber view (4CH), Left Ventricle Outflow Tract (LVOT), Right Ventricle Outflow Tract (RVOT), 3 Vessel and 3 Vessel and Trachea Views (3VV/3VT). This work endeavours to align these 2D views into an average 3D space using deep learning, to provide 3D context to these high resolution acquisitions during clinical examinations.

1.1. Related Works

Hou et al. (2018) introduced a CNN-based method for 2D to 3D fetal MRI slice localisation into a canonical atlas space, forming the foundation for later approaches using anchor points, point loss, and Fibonacci sphere sampling (Xu et al., 2022; Yeung et al., 2021). Alternative works explored geodesic loss functions (Mohseni Salehi et al., 2019). The most relevant prior work in the literature is PlaneInVol (Yeung et al., 2021), addressing the 3D localisation of 2D Ultrasound brain planes. PlaneInVol utilises a CNN architecture, computing attention across input slices to predict slice transformations. It is trained on dense inputs ($N=32$), and is sensitive to low numbers of input slices (< 10 slices). Unlike a transformer architecture, where positional tokens are an inherent part of the structure, PlaneInVol does not leverage any prior positional information, such as the view type.

Slice-to-Volume Registration Transformer (SVoRT) (Xu et al., 2022) is a transformer-based method for reconstructing fetal brain MRI volumes from stacks of slices using an iterative approach, proceeding from classical methods (Kuklisova-Murgasova et al., 2012). In SVoRT, the slices within a given stack are assumed to have highly correlated positions, due to the acquisition protocols of fetal MRI. This work builds on SVoRT and PlaneInVol, adapting the framework for sparse fetal echocardiography data. Such approaches have not been developed for fetal echocardiography data.

1.2. Contributions

This study presents and validates FERN, a **F**etal **E**chocardiography **R**egistration **N**etwork for 2D-to-3D alignment, combining the spatial context of 3D imaging with the superior resolution and accessibility of 2D echocardiography. FERN is validated on real CHD cases, laying the foundation for a novel 3D fetal cardiac visualisation tool.

Two novel technical contributions are introduced compared to prior works: (1) the use of a view positional indicator and a transformer architecture for ultrasound plane localization, contrasting with the CNN from Yeung et al. (2021), and (2) the ability to handle sparse inputs consisting of only 1-5 standard view slices in any orientation. In contrast, SVoRT (Xu et al., 2022) requires dense inputs comprising multiple highly correlated slice stacks.

2. Methods

FERN automates the alignment of 2D fetal echocardiography standard views in an average 3D space. Using a Slice-to-Volume Registration Transformer (Xu et al., 2022) (see Appendix A), it localizes 1-6 standard views in a predefined 3D space (Fig. 1). The network is trained in a supervised manner, simulating 2D slices from real 3D STIC volumes. Random transformations for N slices are sampled from each volume (see Sec. 2.2). FERN predicts transformations aligning 2D slices to 3D space using spatial context across all N slices.

The main modifications to SVoRT framework are summarised in Table 1. Due to the sparse input, components designed for the full reconstruction are omitted (Xu et al., 2022).

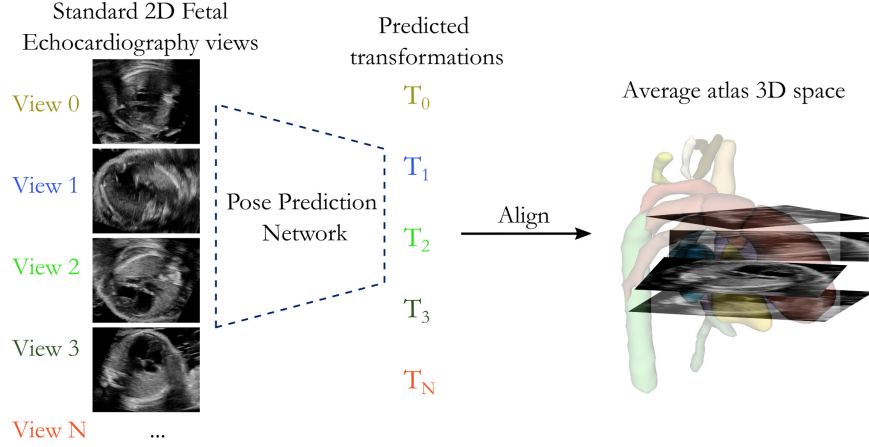


Figure 1: End-diastole frames from standard views of one subject are localised in a standard 3D space via a **F**etal **E**chocardiography **R**egistration **N**etwork. Atlas from Uus et al. (2022b).

Table 1: Comparison of the Original SVoRT Framework (Xu et al., 2022) and FERN

Aspect	SVoRT	FERN
Dataset	Fetal brain MRI	Fetal echocardiography
Transformation Sampling	Any orientation	Approx. standard view
Slice Transformations in a stack	Correlated	Independent
Positional Information	Stack and slice index	View indicator
Input Sparsity	Dense, stacks of slices	Sparse, 1-6 views
Output	3D reconstruction	3D slice transformations

2.1. Loss Functions

Predicted transformations are parameterized by three anchor points in a plane (Hou et al., 2018). The network is trained using an L2 loss between predicted (\hat{P}) and target (P) anchor points (**point loss**). The impact of an **image loss** is also studied via a smooth L1 loss (Appendix B) between the input y_{in} and predicted y_{pred} slices. The total loss function is

$$\mathcal{L}_{localisation} = \|\hat{P}_1 - P_1\|_2^2 + \|\hat{P}_2 - P_2\|_2^2 + \|\hat{P}_3 - P_3\|_2^2 + \lambda \mathcal{L}_{smooth}(y_{in}, y_{pred}), \quad (1)$$

2.2. Training Transformation Sampling and Positional Embedding

FERN is designed for standard fetal echocardiography views, which are taken at an approximate axial position (see Fig. 3). Therefore the slice sampling of the 2D training slices from the 3D STIC volumes should reflect this. Instead of evenly sampling across a sphere (Hou et al., 2018), the spherical coordinates (azimuth θ and polar ϕ angles) in this approach are:

$$\theta = 2\pi \cdot \text{Uniform}(0, 1), \quad (2)$$

$$\phi \sim \mathcal{N}\left(0, \left(\frac{\pi}{12}\right)^2\right), \quad \text{with } \phi \in \left[-\frac{\pi}{4}, \frac{\pi}{4}\right] \text{ rad.} \quad (3)$$

Here, ϕ is sampled from a normal distribution centred at 0 (no through-plane rotation) with $\sigma = 15^\circ$. The full 360° in-plane rotations θ are sampled. Fig. 2 shows random transformations applied to the unit X, Y and Z-axes vectors (collinear to each axis).

Positional embedding uses view type information as prior knowledge (Baumgartner et al., 2017). Slices near the expected location of a standard view share the same positional value (Fig. 3).

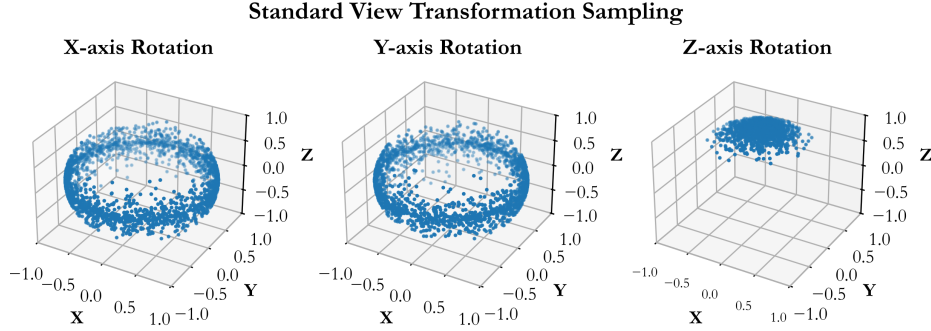


Figure 2: Standard view transformation sampling applied to unit axis vectors, reflecting in-plane rotations (X and Y-axis), and through-plane (Z-axis).

2.3. Training

The network is trained on dense inputs of up to 94 slices, with a 30% chance of randomly switching to 1–5 slices at expected standard view locations, enabling it to learn full spatial context while handling sparse inputs during inference. Data augmentation includes random noise, contrast adjustments, affine transformations, and heart and thorax masking to ensure robustness to shadows and extracardiac features.

All experiments are trained for 800000 iterations, using an AdamW optimizer (Loshchilov and Hutter, 2019). Each experiment is repeated five times to improve reliability and account for variability caused by random weight initialization, stochastic training processes,




Figure 3: Schematic diagram of the coarse standard view positions: a range of slices are assigned the same position. Atlas from [Uus et al. \(2022b\)](#).

and data augmentation. Model ensembling is also studied, where the predicted points of all five networks are averaged. The batch size=8, with a linearly decaying learning rate initialised at 2×10^{-4} . For the image loss, $\lambda = 1000$ (Eqn. 1).

2.4. Ablation study

An ablation study is performed to evaluate the impact of SVoRT components. The **Baseline** uses 3 SVT iterations ([Xu et al., 2022](#)) with view positional embedding (PE) and image loss. Components are individually removed: PE, image loss, reduced number of iterations, and the transformer, replaced by a CNN ([Yeung et al., 2021](#)).

2.5. Training Dataset

The training dataset consists of 3D US (STIC) acquired with a Phillips X6-1 matrix probe. Of an initial 128 volumes from 85 subjects, only 31 volumes from 19 subjects (including 7 for testing) were used due to poor image quality, highlighting the limitations of current 3D acquisitions. The dataset includes 15 Right Aortic Arch (RAA), 8 Coarctation of the Aorta (CoA), and 1 Tetralogy of Fallot case. Despite the small dataset, random slice sampling with through-plane and in-plane rotations and translations greatly increases data variability.

End-diastole 3D STIC frames were rigidly registered to the atlas space (see Appendix C), manually quality-checked, and resampled to $94 \times 94 \times 94$.

2.6. Testing Datasets

The testing dataset includes 24 real 2D scans from 7 subjects, each subject having a paired 3D STIC (*Paired real 2D*) and 54 unpaired 2D images from 15 cases (*Unpaired real 2D*),

comprising 10 CoA and 12 RAA cases. The 2D paired scans, acquired with a Phillips C9-2 curvilinear probe, include standard views manually aligned to STIC volumes, which were registered to the reference 3D space (Appendix C, D). Paired T2-weighted MRI reconstructions (Uus et al., 2020, 2022a) for these cases were also aligned to their respective 3D STIC volumes. The average gestational age = 31.17 ± 1.24 weeks. An **average atlas** was created from high-quality 3D STIC volumes (Appendix E) and used as an evaluation reference.

2.7. Evaluation

Five random in-plane transformations are applied to each 2D image in the *Paired real 2D* dataset, generating 120 test instances ($24 \text{ slices} \times 5 \text{ transformations}$). Each experiment is repeated five times and evaluated using Euclidean Distance on a heart ROI mask (ED_{mask}), RMSE of translation components ($RMSE_{\text{trans}}$), and Geodesic Distance (GD, Appendix F). Statistical significance is evaluated using a Wilcoxon signed-rank test against the Baseline. Sec. 3.2 includes a similarity analysis, comparing the real 2D scans to the planes selected from their paired 3D volume (STIC and MRI), given either the manual or predicted transformations, including Multi-Scale Structural Similarity Index (MS-SSIM (Wang et al., 2003)), Normalised Cross-Correlation (NCC) and Normalised Mutual Information (NMI).

For the *Unpaired real 2D* dataset, 2D slices are passed through the network at their acquired orientation, with predicted transformations used to extract corresponding atlas slices, enabling a similarity comparison (Sec. 3.3). Results for the *Paired real 2D* dataset are also included, comparing acquired 2D slices to slices extracted from the atlas using the predicted or manually selected transformation.

A blinded qualitative analysis (Sec. 3.4) evaluates predicted and manually selected 3D US (STIC) and atlas planes for feature consistency and presence compared to the 2D input scan, selecting the best plane.

3. Results

3.1. Transformation Accuracy on Paired Real 2D dataset

Table 2 presents mean error metrics for the *Paired real 2D* dataset. Training with both coarse PE and image loss improves localisation accuracy, while using 3 iterations shows no significant improvement over 1 iteration. However, all average ED and translation errors remain under 4 mm, the maximum inter-subject longitudinal displacement in the manual alignment of 4CH views of the test set. As expected, model ensembling boosts performance significantly, decreasing ED errors by 8.7%, albeit increasing computational cost fivefold. Detailed ensemble results for other ablations are in G, showing similar trends.

The CNN architecture yields higher errors likely due to lacking positional information, while the transformer, despite fewer parameters, trains ~ 2 hours faster via task parallelization. The chosen model is the **ensemble transformer with 1 SVT iteration, PE, and image loss**. Competitive results are also achievable without PE.

3.2. Similarity to Paired 3D Acquisition (STIC)

Fig. 4 presents similarity metrics for the *Paired real 2D* dataset, comparing 2D slices to those from the 3D dataset (STIC/MRI) using predicted or manual transformations. Comparable

Table 2: Mean results on manually aligned 2D data. ED = Euclidean Distance, GD = Geodesic Distance. P-values (compared to baseline): * < 0.05, ** < 0.005, *** < 0.0005, **n.s.**=non-significant

Experiment	ED _{mask} (mm)	RMSE _{trans} (mm)	GD (deg.)	Params.
3 Iters. (Baseline)	3.264 ± 1.384	1.4 ± 0.756	9.749 ± 4.917	74M (≈ 38.12 h)
No PE	3.748 ± 1.964***	1.7 ± 1.024***	10.657 ± 5.802*	
No Image Loss	3.548 ± 1.486**	1.576 ± 0.802***	10.146 ± 5.309 n.s.	
CNN	3.996 ± 1.555***	1.875 ± 0.921***	10.177 ± 5.663 n.s.	17M (≈ 15.38 h)
1 Iter.	3.198 ± 1.405 n.s.	1.357 ± 0.714 n.s.	9.847 ± 5.345 n.s.	32M (≈ 13.4 h)
1 Iter. Ensemble	2.979 ± 1.269	1.284 ± 0.677	9.074 ± 4.691	150M

median and mean values are seen across metrics, with no significant difference between manual alignment and GT, indicating competitive network performance.

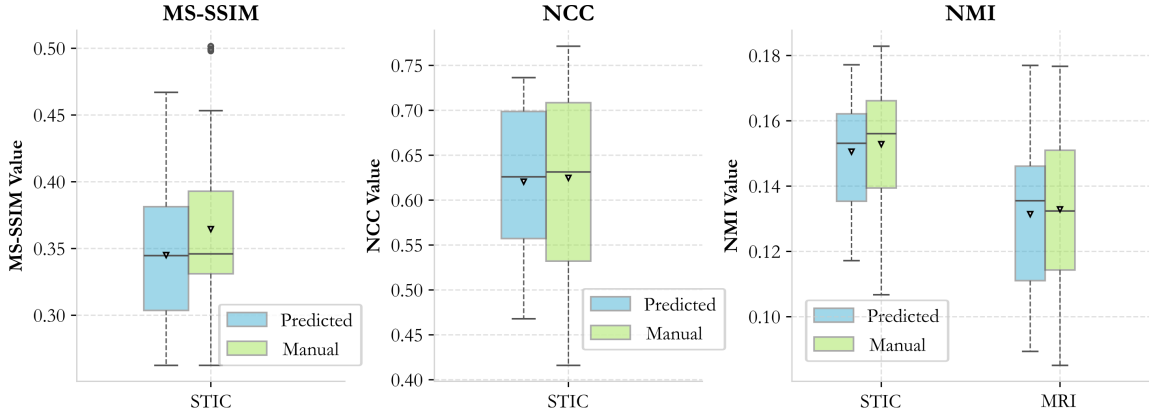


Figure 4: Similarity metrics comparing predicted/manual STIC slices to the 2D.

3.3. Similarity to Average Atlas

Fig. 5 compares 2D slices to their corresponding slices in the average atlas, extracted using the predicted transformations. Both *Unpaired real 2D* and *Paired real 2D* datasets are included, with manually aligned results for the latter. Predicted slices show higher mean and median metrics than manually aligned ones (higher=better), with significant differences (tested for paired cases), indicating FERN’s superior alignment to the average 3D space.

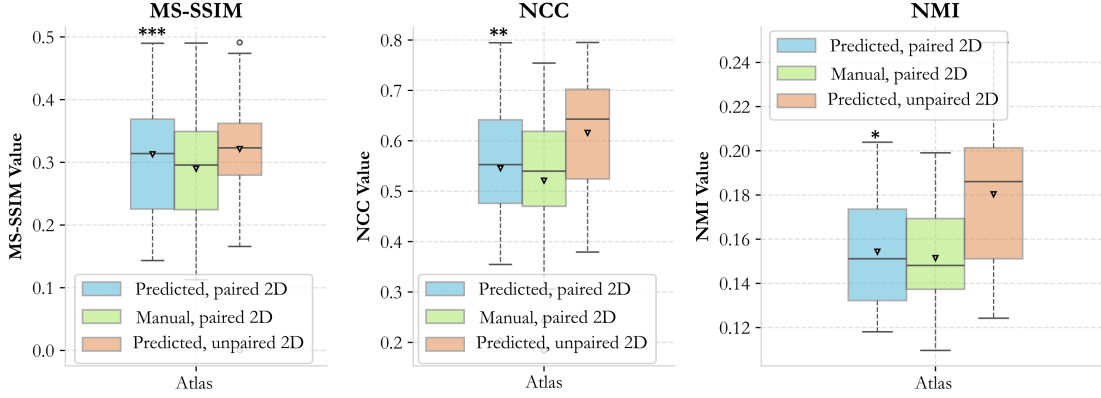


Figure 5: Similarity metrics comparing atlas and 2D US slices.

3.4. Qualitative Analysis

The blinded preferred plane selection results are included in Fig. 6, where the closest plane to the 2D is selected. These suggest a similar trend to the similarity metrics reported in Sec. 3.2 and 3.3, where performance between 3D STIC manual or predicted planes is roughly matched, and network alignment to the average atlas space is more accurate than manual alignment. Examples of this are also depicted in Fig. 6 for atlas alignment and 3D STIC.

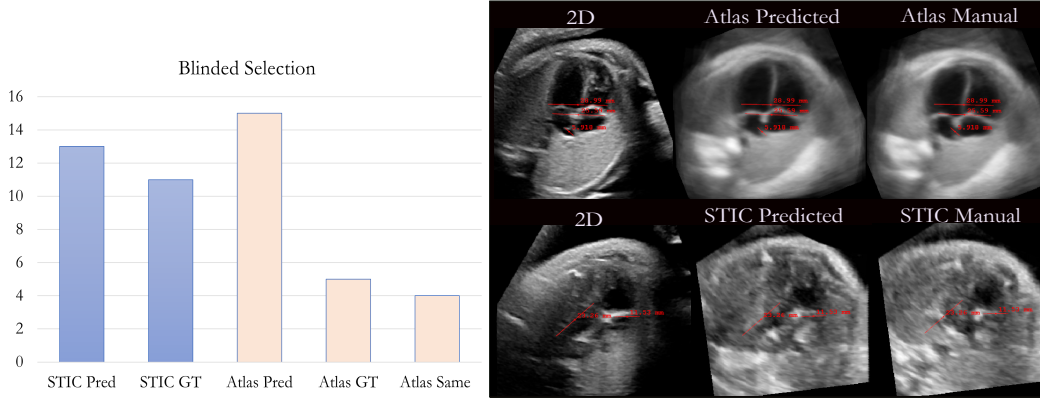


Figure 6: Histogram of the blindly selected preferred slice for the 3D US (STIC) and average atlas based on anatomical similarity to the input 2D slice, with example cases included.

4. Conclusion

This study demonstrates the feasibility of aligning sparse standard fetal echocardiography views into an average 3D space using a transformer network, allowing the inclusion of a view positional indicator to improve localisation confidence. FERN enables the combination of high resolution 2D fetal echocardiography with the low resolution 3D acquisitions or atlases, to provide 3D context during fetal cardiac examinations.

Acknowledgments

We would like to acknowledge funding from the EPSRC Centre for Doctoral Training in Smart Medical Imaging (EP/S022104/1).

We thank everyone who was involved in the acquisition and examination of the datasets and all participating mothers. This work was supported by the Rosetrees Trust [A2725], the Wellcome/EPSRC Centre for Medical Engineering at King’s College London [WT 203148/Z/16/Z], the Wellcome Trust and EPSRC IEH award [102431] for the iFIND project, the NIHR Clinical Research Facility (CRF) at Guy’s and St Thomas’ and by the National Institute for Health Research Biomedical Research Centre based at Guy’s and St Thomas’ NHS Foundation Trust and King’s College London. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

References

- Christian F. Baumgartner, Konstantinos Kamnitsas, Jacqueline Matthew, Tara P. Fletcher, Sandra Smith, Lisa M. Koch, Bernhard Kainz, and Daniel Rueckert. SonoNet: Real-Time Detection and Localisation of Fetal Standard Scan Planes in Freehand Ultrasound. *IEEE Transactions on Medical Imaging*, 36(11):2204–2215, November 2017. ISSN 0278-0062, 1558-254X. doi: 10.1109/tmi.2017.2712367. URL <https://ieeexplore.ieee.org/document/7974824/>. Publisher: Institute of Electrical and Electronics Engineers (IEEE).
- Damien Bonnet, Anna Coltri, Gianfranco Butera, Laurent Fermont, Jerome Le Bidois, Jean Kachaner, and Daniel Sidi. Detection of Transposition of the Great Arteries in Fetuses Reduces Neonatal Morbidity and Mortality. *Circulation*, 99:916–918, February 1999. doi: 10.1161/01.cir.99.7.916. URL <https://pubmed.ncbi.nlm.nih.gov/10027815/>.
- J. S. Carvalho, R. Axt-Fliedner, R. Chaoui, J. A. Copel, B. F. Cuneo, D. Goff, L. Gordin Kopylov, K. Hecher, W. Lee, A. J. Moon-Grady, H. A. Mousa, H. Munoz, D. Paladini, F. Prefumo, E. Quarello, J. Rychik, B. Tutschek, M. Wiechec, and S. Yagel. ISUOG Practice Guidelines (updated): fetal cardiac screening. *Ultrasound in Obstetrics & Gynecology*, 61(6):788–803, June 2023. ISSN 0960-7692, 1469-0705. doi: 10.1002/uog.26224. URL <https://obgyn.onlinelibrary.wiley.com/doi/10.1002/uog.26224>.
- G. R. DeVore, P. Falkensammer, M. S. Sklansky, and L. D. Platt. Spatio-temporal image correlation (STIC): new technology for evaluation of the fetal heart. *Ultrasound in Obstetrics & Gynecology*, 22(4):380–387, October 2003. ISSN 0960-7692, 1469-0705. doi: 10.1002/uog.217. URL <https://obgyn.onlinelibrary.wiley.com/doi/10.1002/uog.217>.
- Julien I.E Hoffman and Samuel Kaplan. The incidence of congenital heart disease. *Journal of the American College of Cardiology*, 39(12):1890–1900, June 2002. ISSN 07351097. doi: 10.1016/S0735-1097(02)01886-7. URL <https://linkinghub.elsevier.com/retrieve/pii/S0735109702018867>.
- Benjamin Hou, Bishesh Khanal, Amir Alansary, Steven McDonagh, Alice Davidson, Mary Rutherford, Jo V. Hajnal, Daniel Rueckert, Ben Glocker, and Bernhard Kainz. 3-D Re-

- construction in Canonical Co-Ordinate Space From Arbitrarily Oriented 2-D Images. *IEEE Transactions on Medical Imaging*, 37(8):1737–1750, August 2018. ISSN 0278-0062, 1558-254X. doi: 10.1109/TMI.2018.2798801. URL <https://ieeexplore.ieee.org/document/8295121/>.
- Noboru Inamura, Takeshi Taniguchi, Tatsumi Yamada, Tomohiko Tanaka, Ken Watanabe, Ken Kumagaya, Yoichiro Ishii, Tsubura Ishii, George Hayashi, Nobuko Shiono, and Tomoaki Ikeda. The Evaluation of Fetal Cardiac Remote Screening in the Second Trimester of Pregnancy Using the Spatio-Temporal Image Correlation Method. *Pediatric Cardiology*, 41(5):979–984, June 2020. ISSN 0172-0643, 1432-1971. doi: 10.1007/s00246-020-02346-7. URL <https://link.springer.com/10.1007/s00246-020-02346-7>.
- Maria Kuklisova-Murgasova, Gerardine Quaghebeur, Mary A. Rutherford, Joseph V. Hajnal, and Julia A. Schnabel. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical Image Analysis*, 16(8):1550–1564, December 2012. ISSN 13618415. doi: 10.1016/j.media.2012.07.004. URL <https://linkinghub.elsevier.com/retrieve/pii/S1361841512000965>.
- Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, January 2019. URL <http://arxiv.org/abs/1711.05101>. arXiv:1711.05101 [cs].
- Seyed Sadegh Mohseni Salehi, Shadab Khan, Deniz Erdogmus, and Ali Gholipour. Real-Time Deep Pose Estimation With Geodesic Loss for Image-to-Template Rigid Registration. *IEEE Transactions on Medical Imaging*, 38(2):470–481, February 2019. ISSN 0278-0062, 1558-254X. doi: 10.1109/TMI.2018.2866442. URL <https://ieeexplore.ieee.org/document/8443391/>.
- Julia A. Schnabel, Daniel Rueckert, Marcel Quist, Jane M. Blackall, Andy D. Castellano-Smith, Thomas Hartkens, Graeme P. Penney, Walter A. Hall, Haiying Liu, Charles L. Truwit, Frans A. Gerritsen, Derek L. G. Hill, and David J. Hawkes. A Generic Framework for Non-rigid Registration Based on Non-uniform Multi-level Free-Form Deformations. In Gerhard Goos, Juris Hartmanis, Jan Van Leeuwen, Wiro J. Niessen, and Max A. Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, volume 2208, pages 573–581. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001. ISBN 978-3-540-42697-4 978-3-540-45468-7.
- Wayne Tworetzky, Doff B McElhinney, V Mohan Reddy, Michael M Brook, Frank L Hanley, and Norman H Silverman. Improved Surgical Outcome After Fetal Diagnosis of Hypoplastic Left Heart Syndrome. *Circulation*, 103:1269–1273, November 2000. doi: 10.1161/01.cir.103.9.1269. URL <https://pubmed.ncbi.nlm.nih.gov/11238272/>.
- Alena Uus, Tong Zhang, Laurence H. Jackson, Thomas A. Roberts, Mary A. Rutherford, Joseph V. Hajnal, and Maria Deprez. Deformable Slice-to-Volume Registration for Motion Correction of Fetal Body and Placenta MRI. *IEEE Transactions on Medical Imaging*, 39(9):2750–2759, September 2020. ISSN 0278-0062, 1558-254X. doi: 10.1109/TMI.2020.2974844. URL <https://ieeexplore.ieee.org/document/9001020/>.

- Alena U. Uus, Irina Grigorescu, Milou P.M. Van Poppel, Johannes K. Steinweg, Thomas A. Roberts, Mary A. Rutherford, Joseph V. Hajnal, David F.A. Lloyd, Kuberan Pushparajah, and Maria Deprez. Automated 3D reconstruction of the fetal thorax in the standard atlas space from motion-corrupted MRI stacks for 21–36 weeks GA range. *Medical Image Analysis*, 80:102484, August 2022a. ISSN 13618415. doi: 10.1016/j.media.2022.102484. URL <https://linkinghub.elsevier.com/retrieve/pii/S1361841522001311>.
- Alena U. Uus, Milou P.M. Van Poppel, Johannes K. Steinweg, Irina Grigorescu, Paula Ramirez Gilliland, Thomas A. Roberts, Alexia Egloff Collado, Mary A. Rutherford, Joseph V. Hajnal, David F.A. Lloyd, Kuberan Pushparajah, and Maria Deprez. 3D black blood cardiovascular magnetic resonance atlases of congenital aortic arch anomalies and the normal fetal heart: application to automated multi-label segmentation. *Journal of Cardiovascular Magnetic Resonance*, 24(1):71, January 2022b. ISSN 10976647. doi: 10.1186/s12968-022-00902-z. URL <https://linkinghub.elsevier.com/retrieve/pii/S1097664723012139>.
- Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, pages 1398–1402, Pacific Grove, CA, USA, 2003. IEEE. ISBN 978-0-7803-8104-9. doi: 10.1109/ACSSC.2003.1292216. URL <http://ieeexplore.ieee.org/document/1292216/>.
- Junshen Xu, Daniel Moyer, P. Ellen Grant, Polina Golland, Juan Eugenio Iglesias, and Elfar Adalsteinsson. SVoRT: Iterative Transformer for Slice-to-Volume Registration in Fetal Brain MRI. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, volume 13436, pages 3–13. Springer Nature Switzerland, Cham, 2022. ISBN 978-3-031-16445-3 978-3-031-16446-0. doi: 10.1007/978-3-031-16446-0_1. URL https://link.springer.com/10.1007/978-3-031-16446-0_1. Series Title: Lecture Notes in Computer Science.
- Pak-Hei Yeung, Moska Aliasi, Aris T. Papageorghiou, Monique Haak, Weidi Xie, and Ana I.L. Namburete. Learning to map 2D ultrasound images into 3D space with minimal human annotation. *Medical Image Analysis*, 70:101998, May 2021. ISSN 13618415. doi: 10.1016/j.media.2021.101998. URL <https://linkinghub.elsevier.com/retrieve/pii/S136184152100044X>.

Appendix A. SVoRT Architecture

Figure 7 details the transformer architecture (Xu et al., 2022) used for predicting the transformations for each slice.

Appendix B. Smooth L1 Loss

The smooth L1 Loss used as image loss is given by

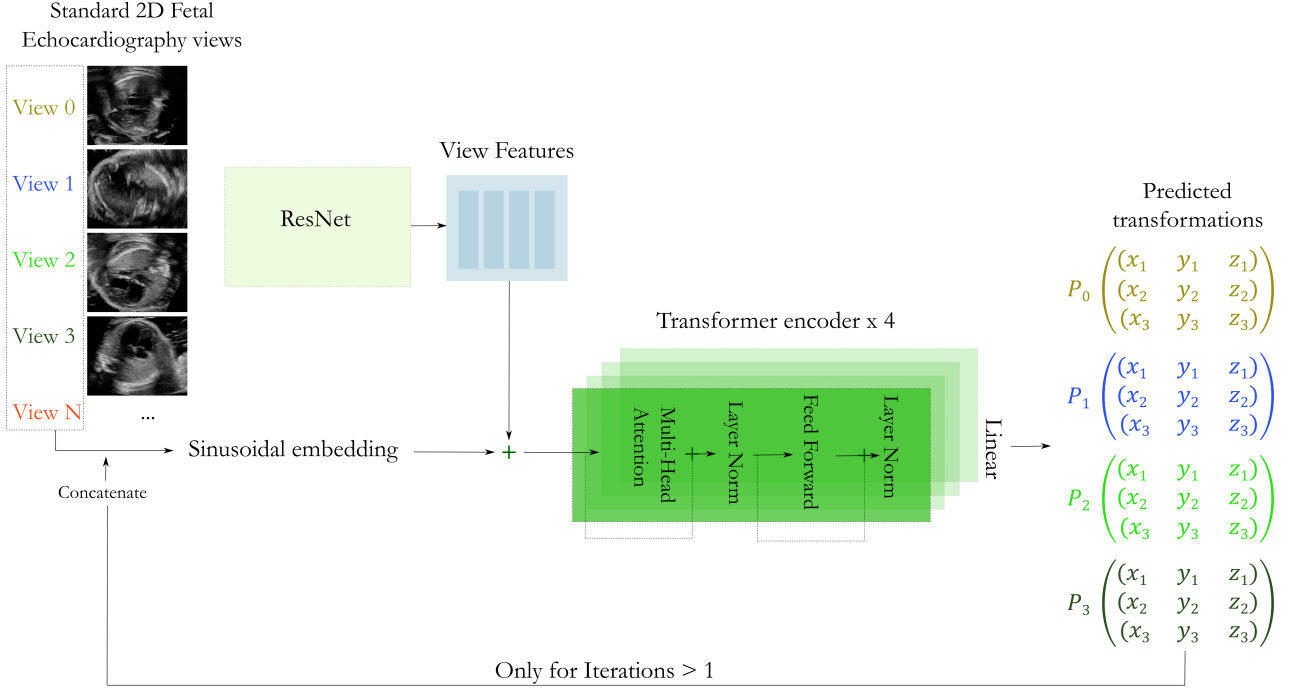


Figure 7: Simplified SVoRT architecture (Xu et al., 2022) used in FERN.

$$\mathcal{L}_{1smooth}(y_{in}, y_{pred}) = \begin{cases} \frac{0.5(x_n - y_n)^2}{\beta}, & \text{if } |x_n - y_n| < \beta \ (\beta = 0.01) \\ |x_n - y_n| - 0.5 \cdot \beta, & \text{otherwise} \end{cases} \quad (4)$$

where y_{in} are the input slices, and y_{pred} are the predicted slices (simulated from the training STIC volumes).

Appendix C. Anatomical Landmarks

The following anatomical landmarks were used for rigid alignment of the 3D STIC volumes used for training:

- Inferior vena cava at the level of the diaphragm
- Descending aorta at the level of the diaphragm
- Descending aorta at the level of the left atrium
- Descending aorta at the level of the left atrium
- Left ventricle apex, transverse plane
- Ascending aorta at the level of the right pulmonary artery
- Right ventricular outflow tract (RVOT) just before bifurcation

- Superior vena cava at the level of the 3 vessel view (3VV), right side only
- Crux of the heart
- Carina
- Spine at the level of the three vessel view

Appendix D. Paired real 2D test set

The *Paired real 2D* dataset consists of 7 unseen cases. Each case consists of a 3D STIC volume, 2D standard view scans and paired MRI volumes (Kuklisova-Murgasova et al., 2012; Uus et al., 2020). The standard view scans were inspected and a high-quality End-Diastole frame was selected.

The corresponding STIC volume slices were examined, and the closest slice to the 2D frame was selected. If the standard view was not visible in the STIC volume it was excluded, leading to variation in the number of 2D views across the 7 cases (24 slices across 7 subjects).

Corresponding landmarks were recorded for the paired 2D and 3D slices. The 2D slice was then rigidly registered to the matching 3D slice (2D-to-2D registration), allowing to generate a pseudo-volume of 2D slices, by placing the registered 2D slices at the selected 3D slice index location.

Random in-plane transformations can then be applied to the slices in this volume, and a 3D ground truth transformation recovered, used to quantify network performance.

A limitation of this method is the assumption that the selected 2D slice aligns perfectly with an axial slice in the 3D volume, whereas a slight tilt (through-plane rotation) may be present in practice. For this reason, an additional qualitative assessment was performed.

The MRI scans were rigidly aligned to the STIC which is in the average atlas space using the same landmarks as described in Appendix C.

Appendix E. Average Atlas

The average atlas was created using 20 high-quality pre-aligned 3D STIC volumes. The atlas construction steps include:

1. Selection of high-quality initial target image.
2. Rigid, Affine and Free-form deformation of all volumes using MIRTk (Schnabel et al., 2001) to target image. A control point spacing of 6 voxels is used.
3. Spatial average of the registered volumes.
4. Registration bias correction, by averaging all transformations and applying the inverse average transform to the spatial average image.
5. Inspection of results and quality control.
6. Rerun steps 2-5 for 3 iterations, using the corrected average image as updated target image.

Fig. 8 depicts the resulting atlas across all three dimensions.




Figure 8: 3D US atlas.

Appendix F. Evaluation: Quantitative Metrics

The Euclidean Distance within the masked heart region is calculated as

$$\text{ED}_{\text{mask}} = \frac{1}{N} \sum_{j=1}^N \|\hat{P}_j - P_j\|^2,$$

where N are the total number of points within the mask, \hat{P}_j are the predicted points, i.e. the points in the mask after transformed using the predicted matrix; and P_j are the GT points, obtained using the manually aligned transformation matrix.

The Geodesic Distance is calculated as

$$\text{GD} = \arccos\left(\frac{\text{Tr}(R) - 1}{2}\right)$$

where R is the rotation matrix from the predicted plane to the target, i.e. the composition of the inverse GT rotation matrix with the predicted rotation matrix.

The RMSE of the translation components using

$$\text{RMSE}_{\text{trans}} = \sqrt{\frac{1}{3} [(\hat{t}_x - t_x)^2 + (\hat{t}_y - t_y)^2 + (\hat{t}_z - t_z)^2]},$$

where $\hat{t}_x, \hat{t}_y, \hat{t}_z$ are the predicted translation components, and t_x, t_y, t_z are the GT translation components for each slice.

Appendix G. Model Ensembling

A similar performance trend is found compared to non-ensemble, with a decrease in the significance of the results, likely influenced by the reduction in data points after ensembling. The results for 1 Iteration achieving comparable performance to using 3 Iterations remain, despite more than doubling the number of parameters for the latter.

Table 3: Mean \pm standard deviation test set results on manually aligned 2D data, after ensembling models from five training rounds.

Experiment	ED _{mask} (mm)	RMSE _{trans} (mm)	GD (deg.)	Params.
3 Iters. (Baseline)	2.951 \pm 1.231	1.275 \pm 0.714	8.712 \pm 4.028	74M \times 5 (\approx 38.12 h \times 5)
No PE	3.298 \pm 1.562	1.508 \pm 0.785	9.574 \pm 4.939	
p-value	0.101	0.646	0.128	
No Image Loss	3.118 \pm 1.305	1.379 \pm 0.74*	9.099 \pm 4.363	
p-value	0.107	0.019	0.240	
1 Iter.	2.979 \pm 1.269	1.284 \pm 0.677	9.074 \pm 4.691	32M \times 5 (\approx 13.4 h \times 5)
p-value	0.643	0.921	0.277	
CNN	3.599 \pm 1.245**	1.666 \pm 0.749**	9.145 \pm 5.127	17M (\approx 15.38 h)
p-value	0.000650	0.00124	0.966	