

Geometric Transformation Uncertainty for Improving 3D Fetal Brain Pose Prediction from Freehand 2D Ultrasound Videos

Jayroop Ramesh¹, Nicola Dinsdale¹, Pak-Hei Yeung^{1,2*}, and Ana I.L. Namburete^{1*}

¹ Oxford Machine Learning in NeuroImaging Lab, Department of Computer Science, University of Oxford, Oxford, United Kingdom

² College of Computing and Data Science, Nanyang Technological University, Singapore
`jayroop.ramesh@cs.ox.ac.uk`

Abstract. Accurately localizing two-dimensional (2D) ultrasound (US) fetal brain images in the 3D brain, using minimal computational resources, is an important task for automated US analysis of fetal growth and development. We propose an uncertainty-aware deep learning model for automated 3D plane localization in 2D fetal brain images. Specifically, a multi-head network is trained to jointly regress 3D plane pose from 2D images in terms of different geometric transformations. The model explicitly learns to predict uncertainty to allocate higher weight to inputs with low variances across different transformations to improve performance. Our proposed method, *QAERTS*, demonstrates superior pose estimation accuracy than the state-of-the-art and most of the uncertainty-based approaches, leading to 9% improvement on plane angle (PA) for localization accuracy, and 8% on normalized cross-correlation (NCC) for sampled image quality. *QAERTS* also demonstrates efficiency, containing $5\times$ fewer parameters than ensemble-based approach, making it advantageous in resource-constrained settings. In addition, *QAERTS* proves to be more robust to noise effects observed in freehand US scanning by leveraging rotational discontinuities and explicit output uncertainties.

Keywords: fetal neurosonography, plane localization, uncertainty

1 Introduction

Two-dimensional (2D) ultrasound is the predominant tool used for routine monitoring of fetal growth, due to its cost-effectiveness, flexibility, and portability, which enable US to provide equitable obstetric examinations in low- and middle-income countries (LMICs). Optimizing the utilization of US within LMICs has

*These authors jointly supervised this work.

the potential to reduce the number of congenital abnormalities that go undetected during early pregnancies [16]. However, each 2D US image represents a cross-sectional view of the 3D anatomy, and localising these planes within the 3D fetal brain is difficult. Reliable and accurate localization of standard planes in fetal brain ultrasound (US) is vital to effectively analyze fetal growth and development [6], and automate detection of antenatal structural defects to enable widespread deployment and adoption in limited resource settings [17].

Deep Neural Networks (DNNs) have been utilized to help with automated ultrasound plane detection and localization [4, 8, 14, 22]. Despite their remarkable performance, these approaches are limited by their need for external sensor or large amounts of manually annotated data. This was addressed in [25], where a DNN was used to predict the 3D pose of 2D US fetal brain images using regression within a pre-defined 3D reference coordinate system. However, its accuracy is affected by the differences in acquisition and quality of the input images, limiting the network performance in practice. *AdLocUI* [24] was proposed to overcome the former limitation using unsupervised cycle consistency but does not address input variability. In this work, we aim to tackle the latter challenge using uncertainty-based approaches.

We hypothesize that expanding *PlaneInVol* [25] to predict different geometric transformations representing the same 3D pose for each 2D US fetal brain image along with their respective variances can promote robustness by accounting for input-dependent noise. We thus make the following contributions:

1. First, our proposed model, *QAERTS*, incorporates multi-head components to regress 3D locations through diverse geometric transformations and explicitly regressed variability, inspired by ensembling [13], ensemble diversity [27], multi-task uncertainty aware learning [12] and geometric reprojection error [10].
2. Second, we adapt and extend multiple state-of-the-art uncertainty-based approaches [2, 5, 13] to the task investigated in this work, namely ultrasound plane localization, as baselines for comparison.
3. Thirdly, we benchmark *QAERTS* and compare with baselines to empirically validate that exploiting geometric transformational variations can improve 3D localization quality compared to most baselines with only a $\sim 0.1\%$ increase in parameters. Our code is publicly available at <https://github.com/jayrmh/QAERTS.git>.

2 Methods

Problem Setup: The pipeline of our proposed work is presented in Fig. 1. During training, N 2D US images $\mathbf{I}_{i=1}^N$, $\mathbf{I} \in \mathbb{R}^{H \times W}$ are sampled from 3D volumes aligned to the 3D anatomical atlas, meaning the true locations are known and no further human annotation is required. We adapt the prediction layer from [25], changing it into a multi-head network, each of which predicts different geometric transformations to express rotational quantities and direct pose predictions [10]. Then, we modify the network to predict learned variances in parallel such that




Fig. 1. Pipeline of our proposed work. During training, 2D slices sampled from aligned 3D volumes are augmented and used to train our proposed uncertainty-aware multi-head model with diverse parameterizations. The feature extractor (cyan) is composed of ten pairs of consecutive 2D convolutional blocks with instance normalization and rectified linear unit (ReLU) activations followed by a maxpooling operation. The generated z embedding is flattened with a adaptive pooling operation, and is propagated through two fully connected layers with a ReLU activation after each layer to a multi-head predictor (black). The trained network can be used to regress the parameters of five different geometric transformations (pink) and their resulting coordinate-wise predictive variances (blue) obtained through a set of independent fully connected layers with no activations along with shared translation and scaling parameters (yellow). The averaged 3D poses and variance obtained from an arbitrary number of 2D images is then used to compute the loss function (green).

the outputs are now Gaussian distributions, and employ maximum likelihood estimation to optimize the model weights with consideration of uncertainty.

Model Architecture: A feature extractor generates a fixed-length feature vector z for each input image, I_i . The feature extractor $f_\theta(x)$ is shared across all inputs and is expected to be invariant to the permutation and number of input images. This supports our task of ultrasound image analysis as freehand image acquisition is operator-dependent and there is variability during the scanning process. The standard direct supervised pose regression model learns parameters $w_r \in \mathbb{R}^{D \times 9}$ to project the learned embedding (D) $z = f_\theta(x)$ into a 3×3 matrix referring to 3 Cartesian coordinates (x, y, z) of 3 reference locations defining the predicted plane, denoted by $\hat{p}^{direct}$. We then compare with the ground truth pose defined by p as follows: $\hat{p}^{direct} = w_r^T f_\theta(x)$. The 3 reference locations correspond to the center, bottom right, and bottom left of the plane as per [25].

Parameterizing Rotational Representations: The intention behind exploiting differing transformations is to separate the learning of rotational representations and variances of transformed reference coordinates, to promote diversity in the loss landscape. We refer to our proposed model as *QAERTS*, where the acronym reflects the geometric transformations used: **Q**uaternions, **A**xis-angles,

Euler angles, Rotation matrices, Translation displacement and Scaling. More formulation details regarding them are found in Supplementary Materials (Table 3). As shown in [7], mapping from the $S0(3)$ space of 3D rotations to angular rotations with representations having four or fewer dimensions is discontinuous. This is because DNNs fit continuous representations with better fidelity [28], and we seek to exploit the rotational discontinuities observed as a measure of the prediction ensemble uncertainty. This is likely to occur as different geometric parameterizations cover different ways of achieving the same 3D transformation and can obfuscate the parallel predictions during training [23]. Thus, we assume that the geometrical inconsistencies that can arise from utilizing the parameterizations of $\mathbf{A}$, $\mathbf{E}$ and $\mathbf{R}$ compared to $\mathbf{Q}$ and *direct pose regression* indicate an input which is *harder* to learn and ambiguous. It should also be noted that large penalties will be avoided, as the loss functions are minimized with respect to the transformed set of points, and not the rotational quantities themselves.

Uncertainty-Aware Learning: Our work builds upon *PlaneInVol* [25] which used the mean-squared-error (MSE) loss. However, MSE does not capture predictive uncertainty, as if the variance is uniform, the Negative Log-Likelihood (NLL) is equivalent to MSE times a constant, but if it is a learnable parameter, higher weight is given to data points with lower variance as per [9, 18]. We extend the base model following [13], and add multiple output heads to predict Gaussian distributions, each having separate branches to predict each geometric transformation and the respective reference coordinate variances across them. This modification leads to the Mean-Variance Estimator (MVE) variant of *PlaneInVol*. Optimal weights are found by maximizing the likelihood, which is equivalent to minimizing the negative-likelihood of the predictive distributions assuming an isotropic Gaussian prior (GNLL). Instead of assuming the aleatoric uncertainty to be the same for all the data samples (homoscedastic), we explicitly model the noise considering heteroscedasticity. This is because subjective sonographer judgment and potential fetal motion relative to the probe placement [24] is likely to affect each frame in 2D US scans differently, as a consequence of difference expertise levels in anatomical understanding.

Modifying Loss: Using the predicted parameters for each rotational representation in orthogonal space, we transform a reference xy grid with respect to the mid-plane and extract the values of the same 3 reference coordinates after the transformation through each combined translation, scaling, and rotational parameterization. Since each head learns a different set of poses and variances for each input, the ensemble of predictions can be denoted by $\hat{p}_i = \{\hat{p}_i^Q, \hat{p}_i^A, \hat{p}_i^E, \hat{p}_i^R, \hat{p}_i^{direct}\}$. Then, the mean $\mu(\hat{p}_i)$ of the implicit ensemble of pose predictions for every input image is computed. As the network also learns the variances for each pose prediction along each head for every independent input denoted by $\hat{s}_i = \{\hat{s}_i^Q, \hat{s}_i^A, \hat{s}_i^E, \hat{s}_i^R, \hat{s}_i^{direct}\}$, the resulting ensemble variance $\hat{\sigma}_i^2 = \mu(\hat{s}_i)$. Following the assumption of $\mathcal{N}(\mu_i, \sigma_i^2)$, the GNLL is minimized wrt the ground truth p_i for every input image i .

The loss function for the heteroscedastic Gaussian Negative Log-Likelihood is:

$$\mathcal{L}_{GNLL} = \sum_i \frac{1}{2} \log(\hat{\sigma}_i^2) + \frac{(\mu(\hat{p}_i) - p_i)^2}{2\hat{\sigma}_i^2} \quad (1)$$

where σ^2 is the average predicted variance by the model and μ is the mean of either the explicit or implicit prediction ensembles.

3 Experimental Setup

Dataset: The 3D ultrasound fetal brain volumes ($160 \times 160 \times 160$ voxels at an isotropic resolution of $0.6mm^3$) were obtained as part of the INTERGROWTH-21st study [20]. The volumes were pre-aligned to a common reference atlas space using [17], and further details are available in [16].

Preprocessing: 4 training and 2 validation 3D volumes selected at 19 gestational weeks.

Artificial 2D Slice Sampling: For each 3D volume in the training epoch, a dense 2D grid along the X and Y axis respectively is generated in the range of -80 to 80, with each xy coordinate initialized to zero. These discrete coordinates will then contain the pixel values of the 2D images, sampled from the 3D volume, and interpolation is performed where the grid does not correspond to a relevant pixel value in the volume. A third dimension representing the z axis is added to localize this sampling space. Intuitively, this implies that 2D planes move along the surface normals tangential to the 3 dimensions of the volume, mimicking the ultrasound probe acquisition process. This process allows for the introduction of random variation through the sampling of infinite combinations of different 2D slices from a 3D volume. Brain segmentation masks [15] were generated to minimize learning with extracranial details from the background.

Data Augmentations: As the 2D images are artificially sampled in a self-supervised manner (training labels derived from the nature of the data itself), we augment sampled slices to be generalizable with freehand 2D acquired videos with random in-plane rotations. They are also subject to random rotations along the x axis and y axis in the range -20 to 20, the z axis (along the surface normal) in the range -90 to 90, and finally random translation also along the z axis between -40 and 60. Lastly, we employ random scaling in the range of 0.75 to 1.8, as well as contrast and intensity modification to introduce noise to the resolution to provide further variations to the sampling process. No normalization of pixel/voxel values were performed.

3.1 Evaluation

Baseline Models: We compared to the commonly deployed approaches [1] representative of each category of approach, namely: Bayesian with Monte-Carlo Dropout (MCD) [5], ensembling with Deep Ensembles (DE) [13] and deterministic methods with Evidential Deep Learning (EDL) [2] to serve as baselines




Fig. 2. Examples from the test set for “high quality” and “low quality” predictions are provided in a)-b) and c)-d) respectively. Input frames (yellow) to the models extracted from 3D volumes are shown in the first column of a) and c). First column of a) and c) show ground-truth planes (green), and predicted (red) are visualized in 3D atlas space along the second to fifth columns of a) and c). Slices sampled (indicated by frame color) from the 3D atlas using the predicted and ground-truth plane locations for each model are shown along second to fifth columns of b) and d).

against our proposed model. Specific implementation details can be found in Supplementary Materials (Table 1). All other training parameters, i.e., combinations of sampled 2D slices, epochs (5000), learning rate (0.0001), optimizer (Adam [25]), patience (20), and batch size (49) were kept consistent between all models. All experiments utilized VGG-16 architecture with random weight initialization as a backbone.

Testing Set: A batch of 32 images sampled from each of the 7 3D volumes at 19 weeks of gestational age were used to quantitatively assess performance of different models. Images extracted from 2D free-hand US videos of three patients are also considered for qualitative analysis at 21 weeks of gestational age.

Metrics: With the 3D anatomical atlas as a reference, we evaluate localization accuracy and sampled image quality with the following metrics. These are Euclidean distance (**ED**) and Plane angle (**PA**) between all coordinates of the predicted and ground-truth planes measured in radians, mean squared error (**MSE**) between predicted and ground truth reference points, as well as Normalized Cross-Correlation (**NCC**) and Structural Similarity (**SSIM**) between images sampled with predicted planes and ground truth planes.

Table 1. Mean results ($\pm$ standard deviation) for quantitative metrics during inference on the testing set. Best scores are in **bold**, and second best scores are underlined. Upward arrows (*uparrow*) indicates higher value is better and downward arrows ($\downarrow$) indicate lower value is better.

Model	ED $\downarrow$	PA $\downarrow$	MSE $\downarrow$	NCC $\uparrow$	SSIM $\uparrow$	Parameters
<i>PlaneInVol</i>	0.34 $\pm$ 0.25	0.42 $\pm$ 0.25	277.19 $\pm$ 245.08	0.62 $\pm$ 0.24	0.57 $\pm$ 0.27	$\sim$ 35.88M
MVE	0.34 $\pm$ 0.25	0.40 $\pm$ 0.25	224.83 $\pm$ 241.28	0.63 $\pm$ 0.25	0.60 $\pm$ 0.25	$\sim$ 35.89M
EDL	0.34 $\pm$ 0.24	0.42 $\pm$ 0.24	227.90 $\pm$ 236.49	0.61 $\pm$ 0.25	0.56 $\pm$ 0.28	$\sim$ 35.90M
MCD 0.1	<u>0.31$\pm$0.24</u>	0.45 $\pm$ 0.25	<u>203.47$\pm$238.51</u>	0.57 $\pm$ 0.23	0.52 $\pm$ 0.28	$\sim$ 35.89M
MCD 0.25	0.34 $\pm$ 0.24	0.43 $\pm$ 0.24	229.54 $\pm$ 235.79	0.49 $\pm$ 0.23	0.44 $\pm$ 0.29	$\sim$ 35.89M
DE	0.29$\pm$0.23	<u>0.40$\pm$0.27</u>	186.42$\pm$223.21	0.73$\pm$0.24	0.63$\pm$0.24	$\sim$ 141.83M
QAERTS	0.32 $\pm$ 0.24	0.39$\pm$0.24	215.10 $\pm$ 241.52	<u>0.67$\pm$0.22</u>	<u>0.61$\pm$0.25</u>	$\sim$ 35.92M

4 Results and Discussion

Relying on a single regression prediction, and not predicting a set of distributions, leads to a larger localization error, as absolute pose regression usually generates outliers very far from the actual scanning angles due to the instability of training regression models [10]. The experiments confirm our hypothesis that expanding the output layer with multiple parameterizations can reduce generalization error of the overall prediction [9].

Quantitative Comparisons with Baselines: From Table 1, we report that the best-performing model is DE, *QAERTS* and MVE across **NCC** and **SSIM** reflecting high sampled image quality. Lowest scores of **ED**, **PA** and **MSE** are once again obtained with DE, with *QAERTS* and MCD in terms of plane coordinate similarity. Without explicit ensembling, we can observe improvements of 4.457% in **ED**, 9.268% in **PA**, 25.225% in **MSE**, 7.556% in **NCC**, and 5.241% in **SSIM** scores between the base model and *QAERTS*, showing advantage of modifying the scoring criteria of the *PlaneInVol* model. Interestingly, performance decreases as the dropout rate increases and both variants are among the weakest in our experiments. As DE is effectively an explicit ensembling of five MVE models, it surpasses the *PlaneInVol* model significantly with **ED**, **PA**, **MSE**, **NCC**, and **SSIM** scores of 14.664%, 6.002%, 39.155%, 16.470% and 9.90% respectively, however, this comes at a cost with $5\times$ more parameters.




Fig. 3. Examples from freehand 2D US videos across three patients. The first column in a), b) and c) shows a frame from each video that were input to each model. Second to fourth column in a), b) and c) show corresponding slices sampled from the 3D atlas, using the predicted plane locations. The pink, yellow and cyan arrows indicate the anatomical structures of lateral ventricles, choroid plexus and Sylvian fissure.

Qualitative Comparisons with Baselines: Fig. 2 show examples of “high quality” and “low-quality prediction” from the testing set. In practice, the former is likely when the scan contains frames concentrated around the center of the fetal brain with rich structural content, whereas the latter occurs due to partial or no recovery of fetal brain information from the peripheral regions. We use predicted locations to sample the corresponding slices from the 3D atlas to which the 3D training volumes were pre-aligned. Better prediction accuracy is obtained when the sampled slices match with the corresponding input images. From these plots, we can surmise that the additional parameterizations contributes to improved performance, and approach quality provided by explicit ensembling in the case of DE. From visualizing the predicted planes in the 3D atlas with their corresponding ground-truths show better overlap in *QAERTS* when compared to *PlaneInVol*, but is not as close as DE.

We assess the integrity of the models in the absence of ground-truth locations by examining the presence of anatomical landmarks in freehand 2D US in Fig. 3. This is evident from the increased resolution of the lateral ventricles (pink arrow), structural fidelity of the choroid plexus (yellow arrow) and clarity of the Sylvian fissure (blue arrow) on the distal hemisphere with *QAERTS* compared

to *PlaneInVol*. DE continues to perform better in terms of accuracy relative to the input, *QAERTS* is nearly similar, and *PlaneInVol* is the furthest.

Parameter Efficiency: As mentioned previously, improved performance of DE comes at a computational cost, requiring approximately $5\times$ the parameters and $5\times$ the inference latency as any other model to produce these results. MCD requires $5\times$ more inference latency to perform explicit ensembling at test time. Our proposed model *QAERTS* provides considerable improvement in localization quality while maintaining parameter efficiency needing only 41,040 (0.114%) more parameters than *PlaneInVol*. Further ablation experiments are in the Supplementary Materials (Table 2).

Loss Landscape: A single low-loss region is likely to contain redundant solutions since the functions are almost identical [11], and the modifications we propose in terms of different parameterizations at the output head for allow for enhanced diversity in the loss landscape relative to the single model baselines. We hypothesize this is a possible reason for performance improvement with *QAERTS*, however, with DE the independently trained explicit ensemble of MVE models better estimates the true marginal likelihood [13] [19].

Future Work: One tangential approach to improve the proposed work is through adapting deterministic uncertainty models for our regression task of plane location prediction by incorporating distance awareness and sensitivity and smoothness in the latent space [21]. A second approach would be to learn continuous representations by performing self-supervised pre-training using contrastive learning (CL) and then fine-tune for our downstream tasks [3] [26].

5 Conclusion

DNNs are highly parameterised, leading to a wide spectrum of possible functions that can effectively capture the data generation process. We have shown that ensemble methods lead to more consistent results, and that our proposed *QAERTS* gives competitive improvement while requiring $5\times$ fewer parameters than other DE variants. Learning with geometrical parameterizations with proper scoring rules in the form of an NLL considering variance implies more retention of frames due to better pose predictions that capture more structural information. This can facilitate enhanced utility for equitable healthcare in low-resource settings by assisting sonographers in scanning guidance.

Acknowledgments. A.I.L.N. and N.K.D. are supported by the Bill and Melinda Gates Foundation. P.H.Y is funded by the EPSRC Impact Acceleration Account Award and the Presidential Postdoctoral Fellowship from the Nanyang Technological University.

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

References

1. Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U.R., et al.: A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion* **76**, 243–297 (2021)
2. Amini, A., Schwarting, W., Soleimany, A., Rus, D.: Deep evidential regression. *Advances in Neural Information Processing Systems* **33**, 14927–14937 (2020)
3. Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., Ballas, N.: Self-supervised learning from images with a joint-embedding predictive architecture. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15619–15629 (2023)
4. Chen, H., Ni, D., Qin, J., Li, S., Yang, X., Wang, T., Heng, P.A.: Standard plane localization in fetal ultrasound via domain transferred deep neural networks. *IEEE journal of biomedical and health informatics* **19**(5), 1627–1636 (2015)
5. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *international conference on machine learning*. pp. 1050–1059. PMLR (2016)
6. Gallery, V.V.: Isuog practice guidelines (updated): sonographic examination of the fetal central nervous system. part 1: performance of screening examination and indications for targeted neurosonography. *Ultrasound Obstet Gynecol* **56**, 476–484 (2020)
7. Hempel, T., Abdelrahman, A.A., Al-Hamadi, A.: 6d rotation representation for unconstrained head pose estimation. In: *2022 IEEE International Conference on Image Processing (ICIP)*. pp. 2496–2500. IEEE (2022)
8. Huang, W., Bridge, C.P., Noble, J.A., Zisserman, A.: Temporal heartnet: towards human-level automatic analysis of fetal cardiac screening video. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 341–349. Springer (2017)
9. Hüllermeier, E., Waegeman, W.: Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning* **110**, 457–506 (2021)
10. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5974–5983 (2017)
11. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems* **30** (2017)
12. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7482–7491 (2018)
13. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems* **30** (2017)
14. Li, Y., Khanal, B., Hou, B., Alansary, A., Cerrolaza, J.J., Sinclair, M., Matthew, J., Gupta, C., Knight, C., Kainz, B., et al.: Standard plane detection in 3d fetal ultrasound using an iterative transformation network. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 392–400. Springer (2018)
15. Moser, F., Huang, R., Papież, B.W., Namburete, A.I., 21st Consortium, I., et al.: Bean: Brain extraction and alignment network for 3d fetal neurosonography. *NeuroImage* **258**, 119341 (2022)

16. Namburete, A.I., Papież, B.W., Fernandes, M., Wyburd, M.K., Hesse, L.S., Moser, F.A., Ismail, L.C., Gunier, R.B., Squier, W., Ohuma, E.O., et al.: Normative spatiotemporal fetal brain maturation with satisfactory development at 2 years. *Nature* **623**(7985), 106–114 (2023)
17. Namburete, A.I., Xie, W., Yaqub, M., Zisserman, A., Noble, J.A.: Fully-automated alignment of 3d fetal brain ultrasound to a canonical reference space using multi-task learning. *Medical image analysis* **46**, 1–14 (2018)
18. Nix, D.A., Weigend, A.S.: Estimating the mean and variance of the target probability distribution. In: *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*. vol. 1, pp. 55–60. IEEE (1994)
19. Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B., Snoek, J.: Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems* **32** (2019)
20. Papageorgiou, A.T., Kennedy, S.H., Salomon, L.J., Altman, D.G., Ohuma, E.O., Stones, W., Gravett, M.G., Barros, F.C., Victora, C., Purwar, M., et al.: The intergrowth-21st fetal growth standards: toward the global integration of pregnancy and pediatric care. *American journal of obstetrics and gynecology* **218**(2), S630–S640 (2018)
21. Postels, J., Segu, M., Sun, T., Sieber, L., Van Gool, L., Yu, F., Tombari, F.: On the practicality of deterministic epistemic uncertainty. *arXiv preprint arXiv:2107.00649* (2021)
22. Ryou, H., Yaqub, M., Cavallaro, A., Roseman, F., Papageorgiou, A., Noble, J.A.: Automated 3d ultrasound biometry planes extraction for first trimester fetal assessment. In: *Machine Learning in Medical Imaging: 7th International Workshop, MLMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 17, 2016, Proceedings 7*. pp. 196–204. Springer (2016)
23. Salehi, S.S.M., Khan, S., Erdogmus, D., Gholipour, A.: Real-time deep pose estimation with geodesic loss for image-to-template rigid registration. *IEEE transactions on medical imaging* **38**(2), 470–481 (2018)
24. Yeung, P.H., Aliasi, M., Haak, M., 21st Consortium, I., Xie, W., Namburete, A.I.: Adaptive 3d localization of 2d freehand ultrasound brain images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 207–217. Springer (2022)
25. Yeung, P.H., Aliasi, M., Papageorgiou, A.T., Haak, M., Xie, W., Namburete, A.I.: Learning to map 2d ultrasound images into 3d space with minimal human annotation. *Medical Image Analysis* **70**, 101998 (2021)
26. Zha, K., Cao, P., Son, J., Yang, Y., Katabi, D.: Rank-n-contrast: Learning continuous representations for regression. *Advances in Neural Information Processing Systems* **36** (2024)
27. Zhou, X., Liu, H., Pourpanah, F., Zeng, T., Wang, X.: A survey on epistemic (model) uncertainty in supervised learning: Recent advances and applications. *Neurocomputing* **489**, 449–465 (2022)
28. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5745–5753 (2019)