Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

LowFU-HC: A Multi-Task Learning Method for Fetal Head Circumference Estimation in a Low-Resource Setting

YASINE LEHMIANI

yasine.lehmiani@um5r.ac.ma

Mohammed V University

Ayoub Elkhouzari

Université Mohammed VI Polytechnique

Abdelhak Mahmoudi

Mohammed V University

Maria Deprez

King's College London

Research Article

Keywords: Fetal Head Circumference, Ultrasound Imaging, Multi-task Learning, Low-resource Settings, Deep Learning, Medical Image Segmentation

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11733529/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

LowFU-HC: A Multi-Task Learning Method for Fetal Head Circumference Estimation in a Low-Resource Setting

Yasine Lehmiani[1*], Ayoub Elkhouzari[3], Abdelhak Mahmoudi[2], Maria Deprez[4]

[1]LIMIARF, Faculty of Sciences, Mohammed V University in Rabat, Morocco.
[2]ENSIAS, SSlab, Mohammed V University in Rabat, Morocco.
[3]College of Computing, University Mohammed VI Polytechnic, Ben Guerir, Morocco.
[4]School of Biomedical Engineering and Imaging Sciences, King's College London, UK.

*Corresponding author(s). E-mail(s): yasine.lehmiani@um5r.ac.ma;

Abstract

Maternal and perinatal mortality remain significant challenges in low-to-middle-income countries, where access to advanced fetal monitoring technologies is limited. Deep learning models have shown promise in automating fetal biometry; however, their deployment is often hindered by computational constraints and limited data availability. This study aims to develop an efficient and accurate method for fetal head circumference (HC) estimation from 2D ultrasound images, specifically designed for resource-limited settings. In this work, we propose LowFU-HC, a multi-task learning approach combined with a modified loss function for fetal head circumference (HC) estimation in 2D ultrasound images that combines head segmentation and fetal ellipse parameter prediction. During inference, only the encoder and regression layers are retained, ensuring efficiency. compared to some state-of-the-art methods, LowFU-HC achieves an HC error of 2.77 mm, with an inference time of 0.35 seconds on CPU and a memory usage of 370 MB on a 400x400 image. By optimizing model efficiency while maintaining high accuracy, our approach enables reliable HC estimation on CPU-only devices. This makes it a practical solution in low-resource settings, where access to high-performance computing resources is limited.

Keywords: Fetal Head Circumference, Ultrasound Imaging, Multi-task Learning, Low-resource Settings, Deep Learning, Medical Image Segmentation

1 Introduction

Ultrasound imaging, a low-cost, non-invasive modality, has become an indispensable tool in modern obstetrics. It provides a real-time and detailed view of the developing fetus, enabling the monitoring of fetal development, the diagnosis of potential abnormalities, and the guidance of clinical management to ensure maternal and fetal well-being [1]. Fetal biometry estimation, an essential component of prenatal ultrasound, typically includes the evaluation of head circumference (HC) [2], biparietal diameter (BPD) [3], abdominal circumference (AC)[4] and femur length (FL) [5]. These measurements are crucial for estimating gestational age [6], assessing fetal growth [1].

According to the World Health Organization (WHO), a minimum of one ultrasound examination is recommended for every pregnancy. However, in low-income countries, there is a noticeable lack of specialists and expert sonographers

[7]. In addition, the low income of many families restricts pregnant women from undergoing ultrasound examinations. In 2020, almost 95% of all maternal deaths occurred in low- and lower-middle-income countries [8], with Sub-Saharan Africa alone represented approximately 70% of these deaths [9]. The maternal mortality ratio in the African region was estimated at 531 deaths per 100,000 live births [9].

In the past decade, deep learning (DL) has significantly transformed medical imaging, achieving performance comparable to human-level accuracy. In the field of fetal health, Tadepalli et al. [10] explore the potential of AI in enhancing prenatal diagnostics and improving clinical decision-making. Similarly, Lin et al. [11] propose a deep learning fetus framework, which outperforms human experts with 1, 3, and 5 years of ultrasound training in classifying standard and non-standard ultrasound images. In the context of fetal biometry, numerous studies have investigated the use of DL to automate the estimation of fetal biometry measurements [12] [13]. However, to the best of our knowledge, the vast majority of these approaches have primarily focused on the performance of DL models, often neglecting their efficiency and deployability, making them unsuitable for resource-limited settings, especially computational constraints, which are especially prevalent in low-to-middle-income countries, particularly in Africa.

In this paper, we propose the following.

- LowFU-HC, a new method that estimates fetal head circumference (HC) from 2D ultrasound images in a low resource setting. The method implements a multitask learning architecture, where a shared encoder is trained for both the segmentation of the fetal head and the prediction of its ellipse parameters. During inference, only the trained encoder along with the regression layers are used to estimate the ellipse parameters.
- A custom loss function to improve the prediction of fetal head orientation parameter.
- A comparative study with a bench of state-of-the-art methods for HC prediction. Key metrics are reported including inference time on CPU and RAM usage.

Our experimental results demonstrate that this multitask learning-based method, combined with the proposed modified loss function, improves the accuracy of fetal parameter predictions compared to a model trained solely for this purpose. Moreover, it reduces the model's memory footprint and inference time on a CPU, making it well-suited for deployment in resource-constrained environments.

2 Related Work

Recent advancements in prenatal ultrasound imaging have driven the development of sophisticated deep learning models aimed at automating fetal biometry measurements. Researchers have explored diverse methodologies to enhance the accuracy and efficiency of these models, which can be broadly categorized into segmentation-based approaches and direct parameter prediction methods.

Segmentation-based methods typically involve predicting a fetal head mask, followed by applying an ellipse-fitting algorithm to determine fetal head parameters. For instance, Bano et al. [12] proposed a segmentation-based approach in which a multi-class network was trained to segment and then predict contour parameters the fetal head, abdomen, and femur. Similarly, Yasrab et al. [14] proposed a hybrid convolutional neural network (CNN) architecture to automate first-trimester biometry measurements of fetal structures.

Plotka et al. [13] proposed a multi-task deep learning framework called FUVAI for 2D+t spatio-temporal analysis of fetal ultrasound videos. FUVAI automates the recognition of standard planes and biometric measurements of fetal head, abdomen, and femur directly from video recordings. Amini et al. [15] proposed a modified Link-Net model, termed Deep-Link-Net, to segment the fetal head in multi-scale ultrasound images. Their approach integrates skip connections between the encoder and decoder, enhancing segmentation performance by concatenating multi-scale inputs and a weighted loss function to prioritize fetal head pixels, particularly in low-contrast images. Sobhaninia et al. [2] proposed a method with fewer convolutional layers that incorporates a half-scale input image concatenated with the feature map in the encoder's first layer

and reported better performance with reduced computational complexity.

Another class of methods focuses on directly predicting fetal head parameters, bypassing segmentation for better efficiency. Zhang et al. [16] evaluated various CNN architectures for direct estimation of fetal head circumference (HC) from ultrasound images. Their study compared lightweight custom networks, such as CNN 263K (263,000 parameters) and CNN 1M (1,000,000 parameters), alongside pre-trained models like VGG16 and ResNet50, which excel in robust feature extraction. Similarly, Frinotino [17] proposed a hybrid method of two consecutive models: a tiny YOLO model to localize the fetal head by generating a bounding box and centering it, then a U-Net architecture to predict a regression mask that delineates the fetal head.

Despite significant progress in this field, most studies rely on non-African datasets, with limited focus on African populations. Sendra-Balcells et al. [18] introduced an African dataset comprising 125 patients for fetal standard plane classification, collected from five African countries. Slimani et al. [19] presented a larger annotated but private dataset of 172 patients from five Moroccan health centers. Their work leveraged Mask-RCNN architectures to estimate fetal biometry (FB) and amniotic fluid volume (AFV) from ultrasound cine-loops, adhering to ISUOG quality standards [20].

While these works achieve impressive performance, most existing methods are not explicitly designed for environments with low computational resources, such as those relying solely on CPUs. Addressing this gap is crucial to ensure that these models can be effectively deployed in low-resource settings without compromising their performance.

3 Method

3.1 The Proposed Method

Multi-task learning leverages shared representations to improve both generalization and resource efficiency [21, 22]. To address fetal head circumference prediction in low-resource settings, we propose LowFU-HC, illustrated in Fig. 1. Our approach operates in two stages. During training, a shared encoder cooperatively learns two tasks: fetal head segmentation and the direct prediction of five fetal head ellipse parameters (c_x, c_y, a, b, θ). The segmentation task acts as an auxiliary mechanism, forcing the encoder to extract richer, more spatially relevant feature representations. During inference, the segmentation decoder is discarded; we deploy only the trained encoder and the fully connected regression head. This streamlined configuration yields a lightweight model optimized for rapid deployment on CPU-only devices.

To identify the optimal backbone for our shared encoder, we adapted three established architectures: standard U-Net [23], Attention U-Net [24], and Spatial Attention U-Net (SA-UNet) [25]. While Attention U-Net utilizes attention gates just before concatenating feature maps in the decoder, SA-UNet integrates a spatial attention mechanism operating directly between the encoder and decoder to better focus on spatially relevant features. For each variant, we integrated a fully connected decoder branch to regress the ellipse parameters alongside the standard upsampling decoder. We specifically selected these CNN-based architectures over recent transformer models, such as SAM [26], because they achieve highly competitive accuracy while requiring a fraction of the computational overhead and training time [27].

3.2 The Loss Function

The total loss function for the multitask learning framework is expressed as a combination of three loss components:

$$L_{\text{total}} = \lambda_{\text{Dice}} \cdot L_{\text{Dice}} + \lambda_{\text{MSE}} \cdot L_{\text{MSE}} + \lambda_{\theta} \cdot L_{\theta} \quad (1)$$

where $\lambda_{\text{Dice}}, \lambda_{\text{MSE}} and \lambda_{\theta}$ are weighting factors of the Dice loss, MSE loss, and rotation loss, respectively. These factors satisfy the following constraints:

$$\lambda_{\text{Dice}} + \lambda_{\text{MSE}} + \lambda_{\theta} = 1, \quad \text{and} \quad \lambda_{\text{Dice}}, \lambda_{\text{MSE}}, \lambda_{\theta} > 0$$

The first term is the dice loss (L_{Dice}) of the segmentation task. The second term is the Mean Squared Error (MSE) loss (L_{MSE}) of the parameters of the predicted ellipse $\hat{c}_x, \hat{c}_y, \hat{a}$ and $\hat{b}$ compared to the ground truth c_x, c_y, a and b:

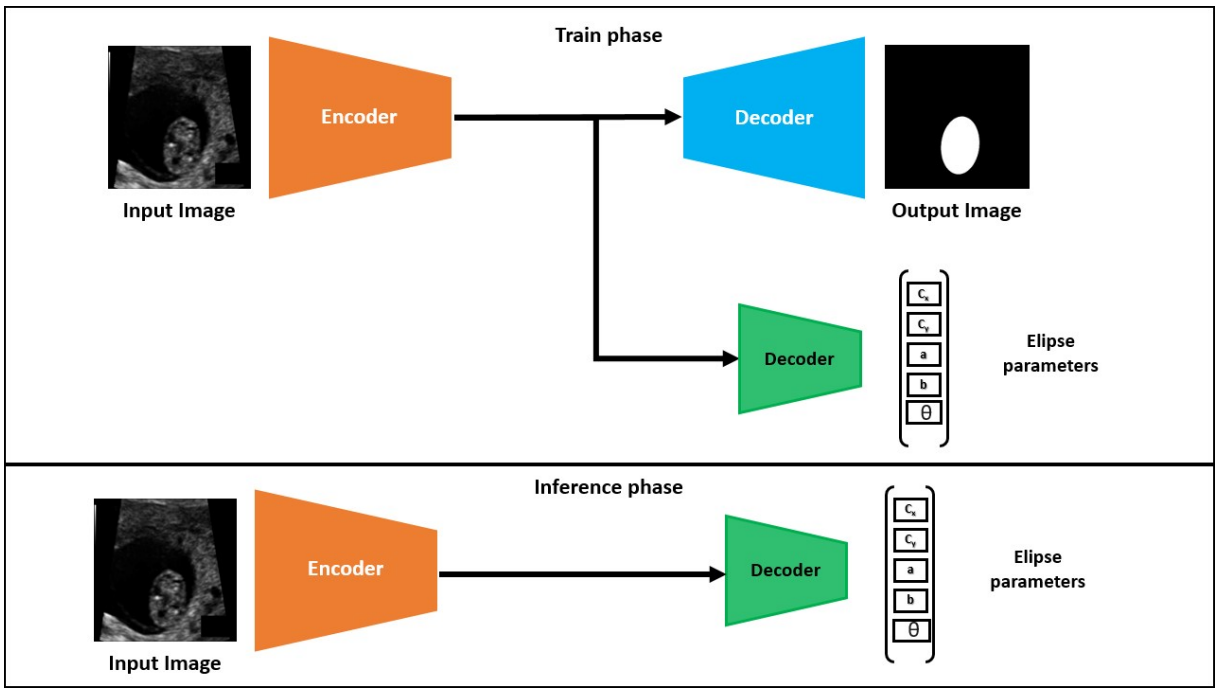

Fig. 1: Overview of the proposed approach. During training, the network jointly learns fetal head segmentation and parameter prediction. During inference, the segmentation decoder is omitted, relying exclusively on the trained encoder and regression network to maximize efficiency.

$$L_{\text{MSE}} = \frac{1}{4} \cdot \text{MSE}((\hat{c}_x, \hat{c}_y, \hat{a}, \hat{b}), (c_x, c_y, a, b)) \quad (2)$$

The orientation angle θ cannot be effectively learned using the MSE loss function because it yields to misleading results for angles. For example, if the ground truth is 0 and the prediction is 2π, the MSE loss incorrectly calculates the difference as $(360°)^2$, even though the two ellipses are aligned. Similarly, if the ground truth is 0 and the prediction is $\pi/2$, the loss would be $(90°)^2$, penalizing the model less than in the first case, even though the ellipses do not overlap .

To address this, we adapt the loss function proposed by [28], defined as:

$$L = 1 - \cos(\hat{\theta} - \theta) \quad (3)$$

However, directly using this cosine-based loss can lead to issues. For instance, it may approach its maximum value of 2 when the predicted and ground truth rotations differ by π, even though an ellipse rotated by π remains aligned with its original state. To resolve this, we propose a new loss function to measure the rotation error (L_θ):

$$L_\theta = 1 - \cos((\hat{\theta} - \theta) \bmod \pi) \quad (4)$$

By applying the modulo π operation, our modified loss function ensures that the rotational difference is treated symmetrically, addressing the alignment issue caused by rotations of π.

3.3 Inference

During inference, we only use the encoder and the prediction layers to predict the five ellipse parameters which reduces inference time. The predicted $\hat{a}$ and $\hat{b}$ can be converted to millimeter size value by using the provided pixel-to-millimeter parameter. The following Ramanujan formula is then applied to estimate the HC value:

$$HC \approx \pi \left[3(\hat{a} + \hat{b}) - \sqrt{(3\hat{a} + \hat{b})(\hat{a} + 3\hat{b})} \right] \quad (5)$$

3.4 Data

The dataset [29] features a collection of 3,832 high-resolution ultrasound images of fetal heads, each

with dimensions of 959×661 pixels. The dataset is reproduced from two publicly available datasets: the HC18 dataset [30] and the Fetal Plane dataset [31]. The HC18 dataset comprises 999 images of fetal heads with corresponding masks, and the Fetal Plane dataset includes 12,400 ultrasound images designed for the classification of fetal planes, such as the brain, abdomen, and thorax, where from this dataset, the authors selected only the images containing fetal heads and labeled specific anatomical regions alongside the fetal head: the brain, cavum septum pellucidum (CSP) and lateral ventricles (LV). A pixel size in millimeters (mm) is associated with each image. Forty images does not contain heads annotations, thus we were left with a total of 3,792 annotated images. From the fetal head binary masks, we extracted the five true ellipse parameters (c_x, c_y, a, b and θ) using least square ellipse fitting.

3.5 Evaluation Metrics

For evaluating the performance of the model during the inference stage, we used the Mean Absolute Error (MAE) as the metric for the parameters (c_x, c_y, a, b) and the Head Circumference (HC). For the rotation angle θ, we used its loss function L_θ as a metric. Additionally, the Hausdorff Distance (HD) was employed to measure the distance between the predicted ellipse, generated from the predicted parameters, and the ground-truth ellipse, generated from the ground truth parameters. Finally, we use the Dice score to further assess the model's performance and investigate the relationship between the fetal head segmentation task and its parameter prediction tasks.

3.6 Experimental Setting

The data was split into 65% for training, 20% for validation and 15% for testing. We trained the architecture with the three models for 100 epochs with a batch size of 12 and a learning rate of 0.0001. The validation set allowed us to fine-tune the models and end up with the following loss functions weighting factors $\lambda_{\text{Dice}} = 0.1$, $\lambda_{\text{MSE}} = 0.45$, and $\lambda_\theta = 0.45$. The training was performed on a Kaggle free GPU environment. To test it in a low-resource environment, the model was deployed on a no-GPU machine with 8 GB of RAM and a CPU of 4 cores.

4 Results

4.1 Comparison of the Three Architectures

Table 1 presents the results of the proposed approach using the three architectures: U-Net, Attention U-Net, and Spatial Attention (SA) U-Net. For the center coordinates c_x and c_y, the SA U-Net yields the lowest errors, with values of **0.54 mm** and **0.47 mm**, respectively, whereas the Attention U-Net achieves the highest errors for these parameters. When evaluating the a and b, the U-Net and SA U-Net exhibit similar performance, sharing the same error value for the major axis. Regarding the HC error, the SA U-Net demonstrates the best performance, achieving the lowest error of **2.77 mm**, followed by the U-Net. The Attention U-Net records the highest HC error, with a value of **3 mm**. Finally, for the θ parameter, all architectures perform similarly, with an error of approximately **0.01 rad**. For the segmentation task, the Attention U-Net achieved the highest Dice score of **94.56**, followed very closely by SA U-Net with a score of **94.2**.

Given these observations, we adopt the SA U-Net architecture in our multi-task approach, leveraging the proposed loss function and the hyperparameter settings described in the experimental setup.

4.2 Suitability for Clinical Deployment

Table 2 represents the comparison of various state-of-the-art methods for fetal head estimation. It shows that our method LowFU-HC achieved the highest performance in almost all metrics with an Intersection over Union (IoU) of **94.2 ± 6**, outperforming other approaches including Slimani et al. [19] (**91** ±.), Plotka et al. [13] (**90.5 ± 3.12**), and Sobhaninia et al. [2] (**92.46**). For HC estimation, LowFU-HC yielded a mean error of **2.77 ± 2.69 mm**, which is less than the mean errors obtained by the segmentation-based methods including Plotka et al. [13] (**2.9 ± 1.2 mm**) and Slimani et al. [19] (**6.7 ± 6.9 mm**). It is worth to mention that these methods estimate fetal head parameters by fitting ellipses to the predicted binary masks. On the other hand, compared to direct regression-based methods where

Table 1: Results of the LowFU-HC using three models on the test set.

Model	$\mathbf{c_x}$ [mm]	$\mathbf{c_y}$ [mm]	**b** [mm]	**a** [mm]	θ [rad]	**HC** [mm]	**IoU**
U-Net	0.57 ± 0.54	0.53 ± 0.60	0.53 ± 0.46	0.52 ± 0.49	0.01 ± 0.02	2.8 ± 2.47	93.61 ± 6
Attention U-Net	0.60 ± 0.58	0.56 ± 0.67	0.55 ± 0.49	0.56 ± 0.51	0.01 ± 0.021	3.04 ± 2.68	$\mathbf{94.56 \pm 5}$
SA U-Net	$\mathbf{0.54 \pm 0.56}$	$\mathbf{0.47 \pm 0.58}$	$\mathbf{0.52 \pm 0.46}$	$\mathbf{0.52 \pm 0.53}$	0.01 ± 0.025	$\mathbf{2.77 \pm 2.69}$	94.2 ± 6

Table 2: Comparison of LowFU-HC with some State-of-the-Art Methods for Fetal Head Estimation.

Model	IoU	HC [mm]	HD [mm]	Memory [MB]	Time [s]
Slimani et al. [19]	$91 \pm$ NA	6.7 ± 6.9	NA	NA	NA
Plotka et al. [13]	90.5 ± 3.12	2.9 ± 1.2	NA	NA	NA
Zhang et al. [16]	NA	4.52 ± 4.27	NA	NA	NA
Amini et al. [15]	93.6 ± 2.8	2.12 ± 2.9	1.89 ± 6.3	NA	NA
Sobhaninia et al. [2]	92.46	2.22	3.53	NA	NA
Fiorntino et al. [17]	97.76 ± 1.32	1.90 ± 1.77	1.32 ± 0.73	NA	0.7
LowFU-HC (Our)	$\mathbf{94.2 \pm 6}$	$\mathbf{2.77 \pm 2.69}$	$\mathbf{1.57 \pm 1.46}$	370	0.35

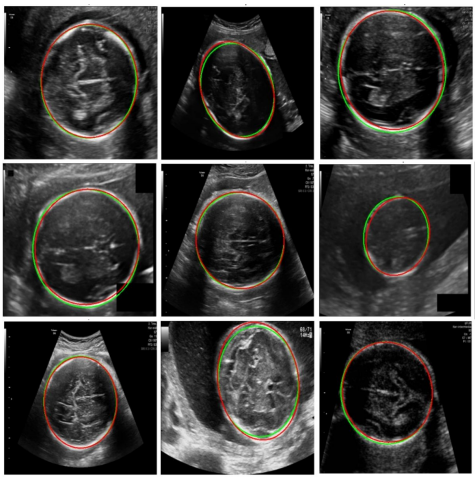

Fig. 2: Visualization of the model's predictions. Where the predicted ellipse in green and the ground truth ellipse in red.

the parameters are predicted direclty without the need for segmentation mask, such as Zhang et al. [16] who reported an HC error of $\mathbf{4.52 \pm 4.27}$ **mm**, our LowFU-HC method acheived lower HC error of ($\mathbf{2.77 \pm 2.69}$ **mm**).

Similarly, our method demonstrated superior accuracy in Hausdorff Distance (HD), achieving an error of $\mathbf{1.57 \pm 1.46}$ **mm**, outperforming Amini et al. [15] ($\mathbf{1.89 \pm 6.3}$ **mm**) and Sobhaninia et al. [2] (**3.53 mm**), and close to the one by Fiorontino et al. [17]. In terms of computational efficiency, our approach uses only 370MB of memory and an inference time of 0.35 seconds for a 400x400 image in a 8 GB of RAM and a 4-core CPU machine. In contrast, the method by Fiorontino et al. [17] achieved an inference time of 0.7 s on a MacBook Pro equipped with Intel Iris Plus Graphics (1536 MB) and 16 GB of RAM. Figure 2 represents the visualization of the prediction results of some example images where the green ellipse is the predicted one and the red ellipse is the ground truth.

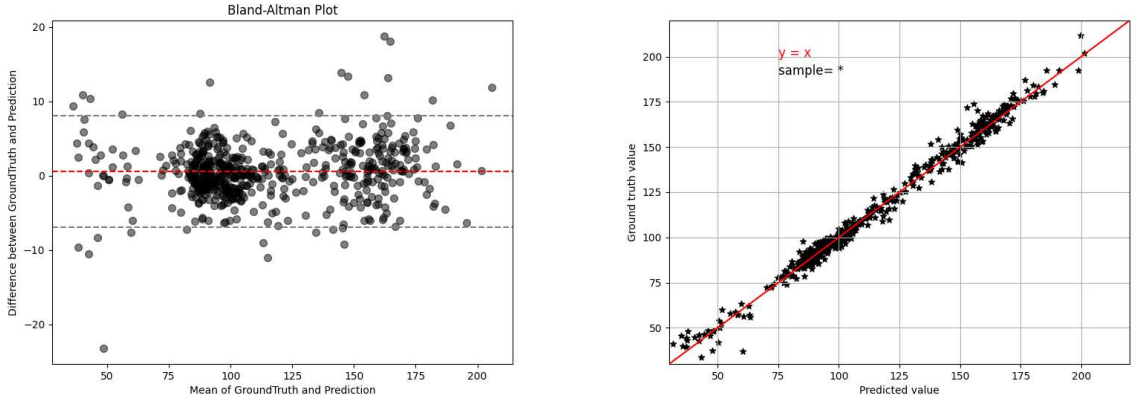

(a) (b)

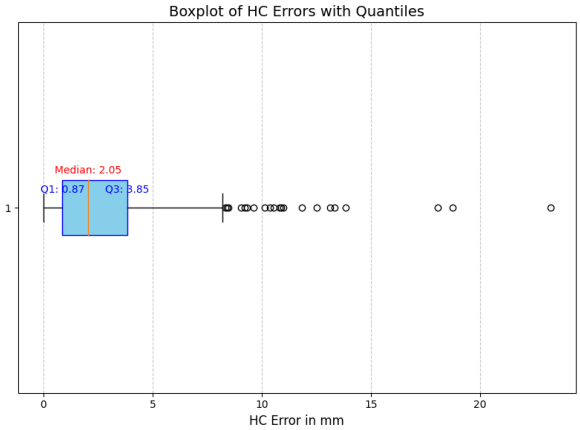

(c)

Fig. 3: (a) Box plot of HC error in millimeters using our approach with the SA-UNet model. (b) Bland-Altman plot illustrating the agreement between ground truth and predicted HC values, with most differences lying within the 95% limits of agreement. (c) Correlation plot showing the linear relationship between ground truth and predicted HC values, with a correlation coefficient of 0.99.



Fig. 4: Visualization of the failed model's predictions. Where the predicted ellipse in green and the ground truth ellipse in red.

5 Discussion

From Table 1, we observe that our method achieves the highest performance when using the

SA-UNet. The SA-UNet incorporates a spatial attention mechanism, which allows the encoder to focus more effectively on the spatial regions of interest, particularly the pixels corresponding to the fetal head. Additionally, the segmentation decoder serves as an auxiliary task, further enhancing the encoder's ability to concentrate on the fetal head region where the segmentation and parameter prediction are relatively cooperative tasks. Thus, during training, we avoid the phenomenon of negative transfer, where one task adversely affects the other's performance. Furthermore, the study by [32] demonstrates [33]that grouping related tasks can help both tasks achieve better performance. Figure 3a presents the Bland-Altman plot for HC, showing that the model is unbiased, with errors centered around 0. Moreover, most errors fall within the 95% confidence interval, which is ±10 mm. Figure 3b illustrates the correlation plot between the predicted and ground truth HC, with a Pearson correlation coefficient of 0.993. The strong linear relationship observed in the plot indicates that the predicted values closely align with the ground truth measurements. Figure 3c represents the box plot of the HC error using the SA-UNet. The median HC error is 2.05 mm, as represented by the red line inside the box, indicating that half of the predictions have an error below this value. Additionally, according to Loughna et al. [6], the fetal HC ranges from approximately 60 mm at 12 weeks of gestational age (GA) to 350 mm at term. Therefore, an HC error of 2.77 mm is too small to significantly impact the estimation of an accurate GA. This demonstrates that the proposed approach is feasible for clinical applications. Furthermore, the approach is deployable in a low-resource setting, as the model only requires 370 MB of memory and can run on a machine with 8 GB of RAM without a GPU. The prediction time is approximately 0.35 seconds for a 400x400 image, ensuring both accuracy and deployability, making it suitable for low-resource settings.

However, the approach does exhibit some limitations, as shown in Figure 4. Errors are observed in images where the fetal head is not centered, leading to inaccuracies in predicting the center, even though the HC error is as low as 0.11 mm. Where centering the head is recommended to obtain the best standard planes, ensuring more accurate predictions. Another challenging scenario involves twin pregnancies, which account for only 3.7% of the dataset. Unfortunately, due to limited computational resources, it was not feasible to augment the data size adequately. These limitations highlight the need for further improvements to enhance the model's generalization capabilities and robustness against data imbalance.

6 Conclusion

In this paper, we proposed LowFU-HC, a model to estimate fetal Head Circumference (HC) in 2D ultrasound images that leverages multi-task learning to predict both the fetal head mask and its ellipse parameters. The model relies solely on parameter regression during inference making it deployable in low-resource settings. This multi-task learning approach significantly improved the prediction accuracy and reduced HC measurement errors. The results demonstrate that our approach outperforms recent methods in the literature in terms of both performance and inference time, making it highly suitable for deployment in non-GPU and low-resource environments. Future work will focus on improving the proposed method to be more robust against data imbalance, which is a common challenge in low-resource settings. Additionally, we aim to enhance the model's generalization to out-of-domain data, making it more suitable for clinical applications, while also improving the model's efficiency.

Declarations

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.
Conflict of Interest: The authors declare no competing interests. **Ethics Approval & Consent:** This study did not require ethical approval as it involved the retrospective analysis of anonymized, publicly available datasets (HC18 and Fetal Plane datasets). Consent to participate and consent for publication are not applicable. **Availability of Data and Materials:** The dataset used in this study is publicly available in the Data in Brief repository [29]. This dataset reproduces and annotates images from the HC18 [30] and Fetal Plane [31] datasets (DOI: https://doi.org/10.1016/j.dib.2023.109708).
Code Availability: The source code for the

proposed LowFU-HC method and experiments is available at: https://github.com/yasinelehmiani/LowFU-HC/tree/main. **Authors' Contributions: Y. Lehmiani:** Conceptualization, Methodology, Software, Writing – Original Draft, Data Curation. **A. Elkhouzari:** Validation, Visualization. **A. Mahmoudi:** Supervision, Conceptualization, Writing – Review & Editing. **M. Deprez:** Supervision, Conceptualization, Writing – Review & Editing.

References

[1] Widyawardana Jatmiko, Irfan Habibie, M. Arif Ma'sum, et al. Automated telehealth system for fetal growth detection and approximation of ultrasound images. *International Journal of Smart Sensing and Intelligent Systems*, 8(1):697–719, 2015.

[2] Zahra Sobhaninia et al. Localization of fetal head in ultrasound images by multiscale view and deep neural networks. *Proceedings of the 2020 25th International Computer Conference (CSICC)*, 1, 2020.

[3] Jee-Youn Hong et al. Can a difference in gestational age according to biparietal diameter and abdominal circumference predict intrapartum placental abruption? *Journal of Clinical Medicine*, 10(11):2413, 2021.

[4] K. Pressman, L. Odibo, J.R. Duncan, and A.O. Odibo. Impact of using abdominal circumference independently in the diagnosis of fetal growth restriction. *Journal of Ultrasound in Medicine*, 41:157–162, 2022.

[5] V. D'Ambrosio, F. Vena, C. Marchetti, et al. Midtrimester isolated short femur and perinatal outcomes: A systematic review and meta-analysis. *Acta Obstetricia et Gynecologica Scandinavica*, 98:11–17, 2019.

[6] Pamela Loughna, Lyn S Chitty, Tessa Evans, and Trish Chudleigh. Fetal size and dating: charts recommended for clinical obstetric practice. *Ultrasound*, 17(3):160–166, 2009.

[7] J. M. Carrera. Obstetric ultrasounds in africa: is it necessary to promote their appropriate use? *Donald School Journal of Ultrasound in Obstetrics and Gynecology*, 5:289–296, 2011.

[8] World Health Organization. Maternal mortality. https://www.who.int/news-room/fact-sheets/detail/maternal-mortality, 2024. Accessed: 2024.

[9] World Health Organization Regional Office for Africa. Maternal mortality: The urgency of a systemic and multisectoral approach in mitigating maternal deaths in africa. https://files.aho.afro.who.int/afahobckpcontainer, 2023. Accessed: 2024.

[10] Kalyan Tadepalli, Abhijit Das, Tanushree Meena, and Sudipta Roy. Bridging gaps in artificial intelligence adoption for maternal-fetal and obstetric care. *Computer Methods and Programs in Biomedicine*, 263:108682, 2025.

[11] Qi Lin et al. How much can ai see in early pregnancy: A multi-center study of fetus head characterization in week 10-14 in ultrasound using deep learning. *Computer Methods and Programs in Biomedicine*, 226:107170, 2022.

[12] Sophia Bano, Brian Dromey, Francisco Vasconcelos, et al. Autofb: Automating fetal biometry estimation from standard ultrasound planes. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 1:403–413, 2021.

[13] S. Płotka, A. Klasa, A. Lisowska, et al. Deep learning fetal ultrasound video model match human observers in biometric measurements. *Physics in Medicine & Biology*, 67(4):045012, 2022.

[14] Robail Yasrab et al. Automating the human action of first-trimester biometry measurement from real-world freehand ultrasound. *Ultrasound in Medicine & Biology*, 50(6):805–816, 2024.

[15] Moloud Amini. Head circumference measurement with deep learning approach based on multi-scale ultrasound images. *Multimedia Tools and Applications*, 81:1–20, 2022.

[16] Jing Zhang, Caroline Petitjean, Pierre Lopez, and Samia Ainouz. Direct estimation of fetal head circumference from ultrasound images based on regression cnn. *Medical Imaging with Deep Learning*, 121:914–922, 2020.

[17] Maria Chiara Fiorentino et al. A regression framework to head-circumference delineation from us fetal images. *Computer Methods and Programs in Biomedicine*, 198:105771, 2021.

[18] C. Sendra-Balcells et al. Generalisability of fetal ultrasound deep learning models to low-resource imaging settings in five african

countries. *Scientific Reports*, 13:2728, 2023.

[19] S. Slimani, S. Hounka, A. Mahmoudi, et al. Fetal biometry and amniotic fluid volume assessment end-to-end automation using deep learning. *Nature Communications*, 14(1):7047, 2023.

[20] L. J. Salomon, Z. Alfirevic, F. Da Silva Costa, et al. Isuog practice guidelines: ultrasound assessment of fetal biometry and growth. *Ultrasound in Obstetrics & Gynecology*, 53(6):715–723, 2019.

[21] Yiyue Zhou, Hao Chen, Yang Li, et al. Multi-task learning for segmentation and classification of tumors in 3d automated breast ultrasound images. *Medical Image Analysis*, 70:101918, 2021.

[22] Min Wang, Xiaolei An, Zhi Pei, et al. An efficient multi-task synergetic network for polyp segmentation and classification. *IEEE Journal of Biomedical and Health Informatics*, 28(3):1228–1239, 2024.

[23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. 9351:234–241, 2015.

[24] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, et al. Attention u-net: Learning where to look for the pancreas. *Medical Imaging with Deep Learning (MIDL)*, 1:1–10, 2018.

[25] Changlu Guo et al. Sa-unet: Spatial attention u-net for retinal vessel segmentation. *2020 25th International Conference on Pattern Recognition (ICPR)*, 1:1236–1242, 2021.

[26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, et al. Segment anything. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1:3992–4003, 2023.

[27] Yasine Lehmiani, Sami Belattar, and Abdelhak Mahmoudi. A comparative study of sam's performance against three u-net architectures on retinal vessels segmentation task. *2023 14th International Conference on Intelligent Systems: Theories and Applications (SITA)*, 1:1–7, 2023.

[28] Xinyu Zhou, Cong Yao, He Wen, et al. East: An efficient and accurate scene text detector. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1:5551–5560, 2017.

[29] Mahmood Alzubaidi et al. Large-scale annotation dataset for fetal head biometry in ultrasound images. *Data in Brief*, 47:109708, 2023.

[30] Thomas L. A. van den Heuvel, Dagmar de Bruijn, Chris L. de Korte, and Bram van Ginneken. Automated measurement of fetal head circumference using 2d ultrasound images. Dataset, 2018. https://doi.org/10.17632/xyz1234.2.

[31] X. P. Burgos-Artizzu, D. Coronado-Gutierrez, B. Valenzuela-Alcaraz, et al. FETAL_PLANES_DB: Common maternal-fetal ultrasound images. Dataset, 2020.

[32] Trevor Standley, Amir R. Zamir, Dawn Chen, et al. Which tasks should be learned together in multi-task learning? *International Conference on Machine Learning (ICML)*, 1:9120–9132, 2020.

[33] Md Zahangir Alom, Theus Aspiras, Tarek M. Taha, and Vijayan K. Asari. Skin cancer segmentation and classification with nabla-n and inception recurrent residual convolutional networks. arXiv preprint arXiv:1904.11126, 2019.