

Accurate and Efficient Fetal Birth Weight Estimation from 3D Ultrasound

Jian Wang^{1,2,3*}, Qiongying Ni^{1,3*}, Hongkui Yu^{5*}, Ruixuan Yao⁶, Jinqiao Ying⁷, Bin Zhang⁵, Xingyi Yang⁵, Jin Peng^{1,3}, Jiongquan Chen^{1,3}, Junxuan Yu^{1,3}, Wenlong Shi⁸, Chaoyu Chen^{1,3}, Zhongnuo Yan^{1,3}, Mingyuan Luo^{1,3}, Gaocheng Cai⁸, Dong Ni^{1,2,3}, Jing Lu^{4(✉)}, and Xin Yang^{1,3(✉)}

¹National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, School of Biomedical Engineering, Shenzhen University, Shenzhen, China

²College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

³Medical Ultrasound Image Computing (MUSIC) Lab, School of Biomedical Engineering, Medical School, Shenzhen University, China

⁴The First Affiliated Hospital of Xiamen University, School of Medicine, Xiamen University, Xiamen, China.

⁵Shenzhen Baoan Women's and Children's Hospital, Shenzhen, China.

⁶University of California, Berkeley, Berkeley, CA, USA

⁷Yongkang Maternal and Child Health Hospital, Yongkang, China

⁸Shenzhen RayShape Medical Technology Co., Ltd, Shenzhen, China

Abstract. Accurate fetal birth weight (FBW) estimation is essential for optimizing delivery decisions and reducing perinatal mortality. However, clinical methods for FBW estimation are inefficient, operator-dependent, and challenging to apply in cases of complex fetal anatomy. Existing deep learning methods are based on 2D standard ultrasound (US) images or videos that lack spatial information, limiting their prediction accuracy. In this study, we propose the first method for directly estimating FBW from 3D fetal US volumes. Our approach integrates a multi-scale feature fusion network (MFFN) and a synthetic sample-based learning framework (SSLF). The MFFN effectively extracts and fuses multi-scale features under sparse supervision by incorporating channel attention, spatial attention, and a ranking-based loss function. SSLF generates synthetic samples by simply combining fetal head and abdomen data from different fetuses, utilizing semi-supervised learning to improve prediction performance. Experimental results demonstrate that our method achieves superior performance, with a mean absolute error of 166.4 ± 155.9 g and a mean absolute percentage error of $5.1 \pm 4.6\%$, outperforming existing methods and approaching the accuracy of a senior doctor. Code is available at: <https://github.com/Qiroy-i/EFW>.

Keywords: Fetal birth weight · 3D ultrasound volumes · Semi-supervised learning

* Jian Wang, Qiongying Ni and Hongkui Yu contribute equally to this work.
Corresponding email: xinyang@szu.edu.cn and lujing810@foxmail.com

1 Introduction

Accurate estimation of fetal birth weight (FBW) prior to delivery is crucial for making optimal delivery decisions (i.e., vaginal birth or cesarean section) and reducing perinatal mortality [14]. As shown in Fig. 1, the standard clinical method for FBW assessment involves a cumbersome procedure: obtaining 2D standard ultrasound (US) planes, measuring biometric parameters, and applying an empirical formula to estimate it [6]. However, each steps poses significant challenges. First, acquiring standard US planes requires radiologists to deeply understand fetal anatomy and accurately locate the 2D plane in 3D space, which is technically demanding and time-consuming [4]. Second, both standard plane acquisition and biometric measurements depend heavily on the radiologist’s experience, leading to intra- and inter-observer variability, which affects the accuracy and reproducibility [10]. Finally, the empirical formulas currently in use are based on regression analysis of a limited set of biometric parameters, failing to account for diverse and complex anatomical variations among different fetuses [12].

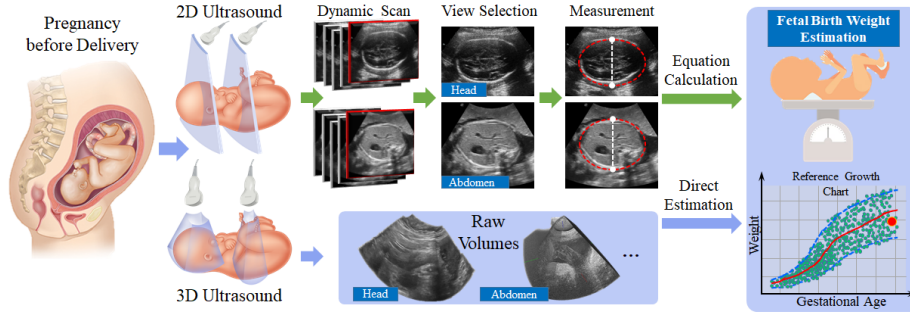


Fig. 1. Workflow for conventional 2D US-based and our 3D US-based FBW estimation.

To address these clinical challenges, several artificial intelligence-based approaches have been proposed, broadly categorized into three types: (1) Some studies use manually measured parameters and maternal information, applying machine learning methods to predict FBW [3, 11, 21]. (2) Other studies leverage deep learning to predict parameters from standard US planes or videos, followed by applying empirical formulas to estimate FBW [1, 17, 18]. (3) A few studies employ deep learning models to directly predict FBW from standard US videos combined with clinical information [15, 16, 19]. However, to date, no approaches have been developed to comprehensively address all three key challenges. Furthermore, due to the lack of spatial information in 2D US imaging, methods based on it struggle to accurately capture fetal anatomical structures, limiting prediction accuracy. In contrast, as shown in Fig. 1, 3D fetal US imaging provides more comprehensive anatomical information and reduces reliance on the operator’s technique, making it a promising alternative for estimating FBW [7, 9, 13].

However, its relatively lower resolution and higher noise levels pose challenges for manual biometric measurements and empirical formula-based estimation. Moreover, while 3D US has not yet achieved widespread clinical adoption due to factors like cost and accessibility, it is increasingly available in advanced obstetric settings. This trend highlights the growing need for automated methods that can fully leverage 3D US data and bypass the limitations of handcrafted features and empirical modeling.

In this study, we propose the first method for directly estimating FBW from 3D fetal US volumes. The key challenges lie in extracting informative features from large, anatomically complex 3D data under sparse supervision, and addressing data scarcity given the wide prediction space. To tackle these, we introduce a multi-scale feature fusion network (MFFN) that integrates head and abdomen volumes using channel attention and Mamba-based spatial attention. We further design a ranking-based loss based on inter-fetal FBW relationships to improve generalization. To alleviate limited sample size, we propose a synthetic sample-based learning framework (SSLF), which generates new fetal samples by mixing head and abdomen volumes from different fetuses and learns from them via semi-supervised training. The experimental results demonstrate that our method achieves superior performance, exceeding existing methods and approaching the accuracy of a senior doctor.

2 Methodology

As illustrated in Fig. 2, our approach consists of a multi-scale feature fusion network (MFFN) and a synthetic sample-based learning framework (SSLF). In MFFN, we integrate channel attention and spatial attention to extract and fuse multi-scale features. Additionally, we design a ranking-based loss function to model the FBW relationships between fetuses. In SSLF, we design a simple yet effective sample synthesis method to expand the dataset and then apply semi-supervised learning to enhance FBW estimation.

2.1 Multi-scale Feature Fusion Network

Due to the significant variations in the spatial scales of fetal anatomical structures, features extracted from a single scale are inadequate for accurate FBW prediction. Thus, we design a multi-scale feature fusion network (Fig.2(a)).

Problem definition: Let $\{(x_i^h, x_i^a, s_i^h, s_i^a, g_i, y_i) \mid i \in 1, 2, \dots, N\}$ denote a mini-batch randomly sampled from the 3D fetal US dataset. Here, $x_i^h, x_i^a \in \mathbb{R}^{1 \times D \times H \times W}$ represent the 3D US volumes of the fetal head and abdomen, where D , H , and W are the depth, height, and width of the volumes, respectively. $s_i^h, s_i^a \in \mathbb{R}^3$ are the voxel spacing in the depth, height, and width directions for the head and abdomen. $g_i \in \{0, 1, 2, 3\}$ is the time interval (in days) from the 3D US scan to delivery. $y_i \in [0, 1]$ denotes the normalized true FBW, and N is the batch size. During clinical inference, g_i is provided by the doctor as the time interval from the 3D US scan to the estimated delivery. To incorporate spacing and time

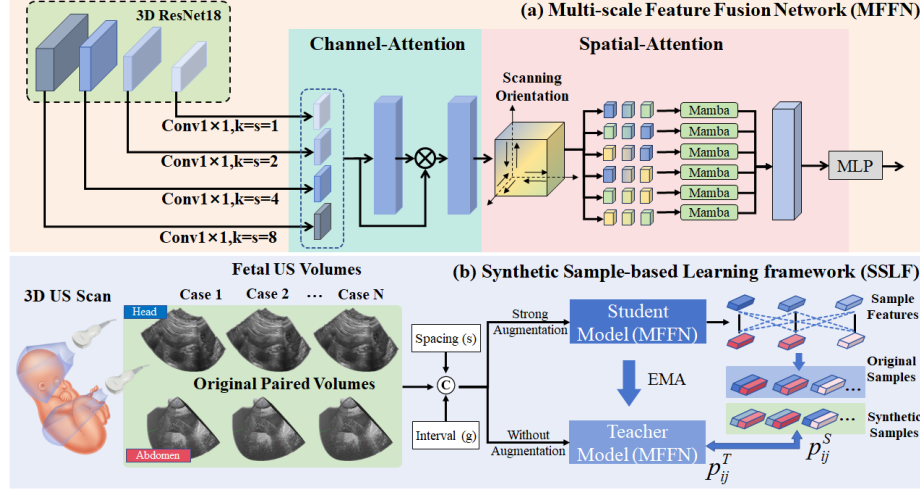


Fig. 2. Overview of the proposed FBW estimation method.

interval information, we expand the dimensions of s_i^h and g_i and concatenate them along the channel dimension with x_i^h , resulting in a five-channel head volume $X_i^h \in \mathbb{R}^{5 \times D \times H \times W}$. Similarly, we obtain abdominal volume data X_i^a . The objective of this study is to predict FBW y_i given the input volumes (X_i^h, X_i^a) .

Feature Fusion: We designed a shared network for both head and abdomen data, so their forward computation processes are identical. For simplicity, we describe the computation process using head data as an example. 3DResNet18 [23] is adopted as the backbone for feature extraction. The features from four intermediate layers, denoted as $Z_i^{h^1}, Z_i^{h^2}, Z_i^{h^3}, Z_i^{h^4}$, are extracted. These features are spatially downsampled by factors of 4, 8, 16, and 32, with corresponding channel dimensions of 64, 128, 256, and 512. To enable multi-scale feature fusion at a uniform spatial resolution, we first apply $1 \times 1 \times 1$ convolutions to project all feature maps to a consistent channel dimension of 128. Then, four sets of convolutions are applied to downsample the features to the same spatial resolution, with kernel sizes and strides set to 8, 4, 2, and 1, respectively. This process aligns features across all four scales both spatially and channel-wise. Finally, we concatenate the features along the channel dimension, resulting in the fused representation $Z_i^h \in \mathbb{R}^{512 \times \frac{D}{32} \times \frac{H}{32} \times \frac{W}{32}}$.

Channel Attention: High-level features typically capture global semantics but with lower resolution, while low-level features focus on fine-grained details at higher resolution. Given their different roles, it is necessary to weigh them appropriately. To achieve this, we apply channel attention to the fused features Z_i^h . We use a bottleneck MLP with the same input and output dimensions, where the activation function is ReLU and the dimensionality reduction factor is set to 16, denoted as $MLP(\cdot)$. The channel attention mechanism is formulated as:

$$Z_i^h = Z_i^h \cdot \sigma(MLP(GAP(Z_i^h))), \quad (1)$$

where $\sigma(\cdot)$ denotes the sigmoid activation function and $GAP(\cdot)$ represents global average pooling over the spatial dimensions. This attention mechanism assigns independent weights to channels at different scales, improving FBW prediction.

Spatial Attention: Since not all regions in the volume correspond to fetal anatomy, and different structures contribute differently to FBW prediction, it is essential for the model to focus on relevant fetal regions. Inspired by the Mamba architecture [5], we design a spatial attention mechanism that employs six Mamba models, each corresponding to a unique scanning orientation: (1) left-to-right, (2) right-to-left, (3) top-to-bottom, (4) bottom-to-top, (5) front-to-back, and (6) back-to-front. Specifically, six spatial sequences $Z_i^{h_1}, Z_i^{h_2}, Z_i^{h_3}, Z_i^{h_4}, Z_i^{h_5}, Z_i^{h_6} \in \mathbb{R}^{512 \times L}$ are generated by applying different flattening methods to the input features, where $L = \frac{D}{32} \times \frac{H}{32} \times \frac{W}{32}$. The sequences are processed by their corresponding Mamba models, and their outputs are fused using an averaging operation, yielding the final feature map $Z_i^h \in \mathbb{R}^{512 \times L}$. This fusion captures spatial information from multiple perspectives, improving FBW prediction.

Loss Function: After the above three steps, we obtain the head and abdomen features Z_i^h and Z_i^a . These features are concatenated along the channel dimension and subjected to global average pooling over the spatial dimensions, resulting in the fused head-abdomen representation $Z_{i,i}^{ha} \in \mathbb{R}^{1024}$. A single fully connected layer followed by a sigmoid activation function is applied to predict the FBW p_i . The regression loss is defined as:

$$Loss_{reg} = \frac{1}{N} \sum_{i=1}^N MSE(p_i, y_i), \quad (2)$$

where $MSE(\cdot)$ denotes the mean squared error between the predicted FBW p_i and the ground truth y_i . To further enhance the model's ability to capture relative weight differences among fetuses, we introduce a ranking-based loss function. This loss encourages the predicted order of FBWs to align with the ground truth order. The ranking loss is formulated as:

$$Loss_{rank} = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \max[0, -(p_i - p_j)] \cdot \tau(y_i - y_j), \quad (3)$$

where $\tau(\cdot)$ equal to 1 if the argument is greater than 0 and 0 otherwise.

2.2 Synthetic Sample-based Learning Framework

In this study, a key challenge we face is the conflict between limited sample size and nearly infinite possible predictions, making conventional supervised learning models prone to severe overfitting. To address this issue, we propose an synthetic sample-based learning framework, as shown in Fig. 2(b).

Sample synthesis: Here, the sample refers to the fetus rather than the head or abdominal volume. We hypothesize that the model predicts FBW by estimating fetal size from both head and abdominal volumes. Due to the anatomical

independence between head and abdominal regions, there is little or no direct correlation between their volumes within the same fetus. Based on this observation, we propose a simple method to synthesize new samples from different fetuses. Specifically, we create synthetic samples by combining the head volume of one fetus with the abdominal volume of another. Let $\{(X_i^h, X_j^a) | i, j \in 1, 2, \dots, N, i \neq j\}$ denote the set of synthetic samples within a mini-batch, where N is the batch size. For each mini-batch, this method can generate $N(N-1)$ synthetic samples, greatly expanding the dataset without any complicated operations.

Semi-supervised learning: Although the previous steps synthesize a large number of new samples to expand the dataset, they lack real FBW labels. To address this, we leverage semi-supervised learning by treating synthetic samples as unlabeled data. Inspired by the MeanTeacher framework [22], we design a semi-supervised learning setup consisting of two models: the teacher model and the student model, both of which share the same architecture (i.e., MFFN). During training, the student model’s parameters are updated via backpropagation, while the teacher model’s parameters are updated using an exponential moving average (EMA) of the student model’s parameters, as follows:

$$\theta_T^t = m\theta_T^{t-1} + (1 - m)\theta_S^t, \quad (4)$$

where θ_T and θ_S denote the parameters of the teacher and student models, respectively. t denotes the iteration step, and m is the momentum coefficient that controls the update rate of the teacher model.

In each training batch, there are N head and N abdomen data samples, resulting in N^2 possible combinations. Of these, N combinations correspond to real fetal samples (with labels), while the remaining $N(N-1)$ combinations correspond to synthetic fetal samples (without labels). For these combinations, let $\{p_{ij}^T | i, j \in 1, 2, \dots, N\}$ and $\{p_{ij}^S | i, j \in 1, 2, \dots, N\}$ represent the predictions from the teacher and student models, respectively. We treat the teacher Model’s predictions as pseudo-labels to supervise the student Model. The loss function for the semi-supervised learning phase is defined as:

$$Loss_{semi} = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j=1}^N MSE(p_{ij}^S, p_{ij}^T), i \neq j. \quad (5)$$

Finally, the total loss function is the weighted sum of three components:

$$Loss_{total} = Loss_{reg} + \alpha Loss_{rank} + \beta Loss_{semi}, \quad (6)$$

where α and β are the loss weighting coefficients. With this approach, our model can effectively learn from synthetic samples generated by combining fetal head data with abdomen data from different fetuses.

3 Experiments and Results

3.1 Datasets and Implementations

With ethics committee approval, we collected a dataset of 491 pregnancies by two senior doctors, each with 17 years of experience. Each case includes one

head, abdomen, and femur ultrasound scan acquired within 72 hours before delivery, using a GE HealthCare E8 system with an RAB6-D probe. Prior to delivery, a senior doctor estimated fetal birth weight (FBW) using the Hadlock formula based on manual biometric measurements. Post-delivery, the true FBW was recorded as ground truth. The mean true FBW was 3229.3 ± 467.8 g, ranging from 1000 g to 4610 g. Voxel spacing was extracted from metadata. The dataset was randomly split into training, validation, and test sets at a 7:1:2 ratio, and all volumes were resized to $160 \times 128 \times 96$ while preserving aspect ratio.

We trained the model for 200 epochs with a batch size of 16 using the Adam optimizer. The learning rate started at $1e^{-4}$, warmed up for 5 epochs, and then decayed via cosine annealing. Data augmentation included rotation, flipping, scaling, and contrast and brightness adjustments to improve generalizability. Hyperparameters were set as follows: $\alpha = 0.001$, β increased linearly from 0 to 0.2, and momentum m from 0.99 to 0.9999. These values were selected via grid search on the validation set. All experiments ran on an RTX 4090 GPU server.

3.2 Results

To evaluate our method, we compared it with several existing approaches, including Hadlock [6], INTERGROWTH-21st [20], BabyNet [16], BabyNet++ [15], 3DCNN-1 [24], 3DCNN-2 [2], and DSCNN [8]. Hadlock and INTERGROWTH-21st are widely used clinical methods based on manual biometric measurements and empirical formulas; we used the pre-delivery estimate by a senior doctor as the Hadlock reference. BabyNet and BabyNet++ predict FBW from 2D ultrasound videos, with the latter also incorporating biometric and maternal data; we adopted the reported results from their original studies. 3DCNN-1, 3DCNN-2, and DSCNN are recent 3D regression models for other tasks, which we re-implemented following their published protocols. All methods were evaluated using mean absolute error (MAE), root mean square error (RMSE), and mean absolute percentage error (MAPE).

Table 1 summarizes the comparison results. Our method achieves strong performance, with a MAE of 166.4 ± 154.9 g, RMSE of 227.3 ± 227.3 g, and MAPE of $5.1 \pm 4.6\%$. Notably, it is highly competitive with Hadlock—widely used in clinical practice—and INTERGROWTH-21st, a global fetal growth standard. The closest method, BabyNet++, leverages standard US videos, manual biometric measurements, and maternal information. In contrast, our approach relies solely on 3D US volumes, requiring minimal scanning expertise and no manual input. We also evaluated the two key components of our method. MFFN alone achieved a MAE of 172.0 ± 164.6 g, RMSE of 238.1 ± 237.1 g, and MAPE of $5.3 \pm 4.9\%$, demonstrating its ability to extract features effectively from 3D volumes under sparse supervision. When combined with SSLF, multiple 3D regression models showed improved performance, indicating its adaptability and potential for broader application.

To evaluate the effectiveness of each component in our method, we conducted a series of ablation experiments. Specifically, we assessed the impact of single-site

Table 1. Comparison results of different methods.

Method	MAE(g)↓	RMSE(g)↓	MAPE(%)↓
Hadlock [6]	160.2 ± 118.5	199.3 ± 197.9	5.0 ± 3.7
INTERGROWTH-21 st [20]	182.9 ± 137.0	228.5 ± 220.9	5.7 ± 4.2
BabyNet [16]*	254.0 ± 230.0	341.0 ± 215.0	7.5 ± 6.6
BabyNet++ [15]*	179.0 ± 19.0	203.0	5.1 ± 0.6
3DCNN-1 [24]	268.7 ± 245.2	363.8 ± 358.7	8.2 ± 7.1
3DCNN-2 [2]	253.1 ± 222.1	336.8 ± 336.2	7.9 ± 6.9
DSCNN [8]	264.7 ± 239.1	356.7 ± 355.0	8.3 ± 7.8
$MFFN_{(Ours)}$	172.0 ± 164.6	238.1 ± 237.1	5.3 ± 4.9
3DCNN-1 + SSLF	264.7 ± 227.8	349.3 ± 348.9	8.2 ± 7.1
3DCNN-2 + SSLF	248.3 ± 212.2	326.6 ± 325.1	8.1 ± 6.7
DSCNN + SSLF	220.0 ± 218.3	310.0 ± 309.1	6.8 ± 6.5
$MFFN + SSLF_{(Ours)}$	166.4 ± 154.9	227.3 ± 227.3	5.1 ± 4.6

*Results reported from original publication; model not retrained on our dataset due to input modality mismatch.

Table 2. Experimental results of ablation study with different configurations of key components of our model. 3DR18: 3DResNet18, WS: Weight-sharing, FF: Feature-fusion, CA: Channel attention, SA: Spatial attention, RL: Ranking-based loss

Model Architecture	Input	Results	
3DR18 WS FF CA SA RL SSLF	Head Abdomen	MAE(g)	MAPE(%)
✓	✓	283.1 ± 229.0	9.0 ± 7.5
✓	✓	252.8 ± 191.8	8.0 ± 6.4
✓	✓	222.4 ± 193.2	6.7 ± 5.1
✓ ✓	✓	200.6 ± 178.1	6.1 ± 5.1
✓ ✓ ✓	✓	189.7 ± 173.9	5.9 ± 5.4
✓ ✓ ✓ ✓	✓	191.0 ± 166.1	5.9 ± 5.0
✓ ✓ ✓ ✓ ✓	✓	177.2 ± 175.2	5.9 ± 5.0
✓ ✓ ✓ ✓ ✓ ✓	✓	172.0 ± 164.6	5.3 ± 4.9
✓ ✓ ✓ ✓ ✓ ✓ ✓	✓	166.4 ± 154.9	5.1 ± 4.6

and multi-site volumes, weight-sharing in multi-site volume fusion, feature fusion, channel attention, spatial attention, ranking-based loss function, and SSLF. The results, presented in Table 2, demonstrate that weight-sharing in multi-site volume fusion is crucial for accurate FBW estimation. Additionally, simple multi-scale feature fusion enhances accuracy, confirming its necessity. While channel attention increases the average MAE by 1.3g, it reduces the standard deviation by 7.8g, making it beneficial overall. Spatial attention further improves performance, decreasing the average MAE by 13.8g. The ranking-based loss function aligns the FBW relationships between fetuses with the groundtruth, effectively boosting prediction accuracy. Finally, SSLF significantly expands the dataset and improves performance through a semi-supervised learning framework.

4 Conclusion

In this study, we present a novel method for predicting FBW directly from 3D fetal US volumes. To address challenges in feature extraction and sample size limitations, our method designs a multi-scale feature fusion network and a synthetic sample-based learning framework. Our results outperform existing methods and close to clinical accuracy. Rather than replacing 2D workflows, our method offers a complementary tool for enhanced decision-making. Future work will extend this framework to jointly estimate biometric parameters and FBW, improving automation and interpretability.

Acknowledgments. This work was supported by the grant from National Natural Science Foundation of China (82201851, 12326619, 62171290); Guangxi Province Science Program (2024AB17023); Yunnan Major Science and Technology Special Project Program (No.202402AA310052); Yunnan Key Research and Development Program (202503AP140014); Science and Technology Planning Project of Guangdong Province (2023A0505020002); Frontier Technology Development Program of Jiangsu Province (No. BF2024078).

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

References

1. Bano, S., Dromey, B., Vasconcelos, F., Napolitano, R., David, A.L., Peebles, D.M., Stoyanov, D.: Autofb: automating fetal biometry estimation from standard ultrasound planes. In: Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VII 24. pp. 228–238. Springer (2021)
2. Cali, R.J., Bhatt, R.R., Thomopoulos, S.I., Gadewar, S., Gari, I.B., Chattopadhyay, T., Jahanshad, N., Thompson, P.M.: The influence of brain mri defacing algorithms on brain-age predictions via 3d convolutional neural networks. In: 2023 45th Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC). pp. 1–6 (2023). <https://doi.org/10.1109/EMBC40787.2023.10340740>
3. Camargo-Marín, L., Guzmán-Huerta, M., Piña-Ramirez, O., Perez-Gonzalez, J.: Multimodal early birth weight prediction using multiple kernel learning. *Sensors* **24**(1), 2 (2023)
4. Chen, H., Dou, Q., Ni, D., Cheng, J.Z., Qin, J., Li, S., Heng, P.A.: Automatic fetal ultrasound standard plane detection using knowledge transferred recurrent neural networks. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I 18. pp. 507–514. Springer (2015)
5. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
6. Hadlock, F.P., Harrist, R., Sharman, R.S., Deter, R.L., Park, S.K.: Estimation of fetal weight with the use of head, body, and femur measurements—a prospective study. *American journal of obstetrics and gynecology* **151**(3), 333–337 (1985)

7. Hamza, H.A., Shabana, A.A.E., Abu Shar, A.M., Egiz, M.N.: Accuracy of fetal thigh circumference and fractional thigh volume in prediction of fetal birth weight. *The Egyptian Journal of Fertility and Sterility* **29**(1), 3–13 (2025)
8. Jafrasteh, B., Lubián-López, S.P., Benavente-Fernández, I.: A deep sift convolutional neural networks for total brain volume estimation from 3d ultrasound images. *Computer Methods and Programs in Biomedicine* **242**, 107805 (2023)
9. Kang, L., Wu, Q.Q., Sun, L.J., Gao, F.Y., Wang, J.J.: Predicting fetal weight by three-dimensional limb volume ultrasound (avol/tvol) and abdominal circumference. *Chinese Medical Journal* **134**(9), 1070–1078 (2021)
10. Li, J., Gao, Z., Wang, C., Pu, B., Li, K.: A rule-guided interpretable lightweight framework for fetal standard ultrasound plane capture and biometric measurement. *Neurocomputing* **621**, 129290 (2025)
11. Lu, Y., Fu, X., Chen, F., Wong, K.K.: Prediction of fetal weight at varying gestational age in the absence of ultrasound examination using ensemble learning. *Artificial intelligence in medicine* **102**, 101748 (2020)
12. Ma, J., Cheng, D., Zhang, Z., Cai, B., Xu, X.: Evaluating the accuracy of sonographic fetal weight estimations using the hadlock iv formula in a chinese population. *Quantitative Imaging in Medicine and Surgery* **13**(6), 3726 (2023)
13. Metwally, S., Alsherbiny, M., Elsayed, M., Abuelezz, S.: Three-versus two-dimensional sonographic biometry for predicting birth weight at term. *Benha Journal of Applied Sciences* **7**(10), 79–83 (2022)
14. Milner, J., Arezina, J.: The accuracy of ultrasound estimation of fetal weight in comparison to birth weight: A systematic review. *Ultrasound* **26**(1), 32–41 (2018)
15. Płotka, S., Grzeszczyk, M.K., Brawura-Biskupski-Samaha, R., Gutaj, P., Lipa, M., Trzciński, T., Işgum, I., Sánchez, C.I., Sitek, A.: Babynet++: Fetal birth weight prediction using biometry multimodal data acquired less than 24 hours before delivery. *Computers in Biology and Medicine* **167**, 107602 (2023)
16. Płotka, S., Grzeszczyk, M.K., Brawura-Biskupski-Samaha, R., Gutaj, P., Lipa, M., Trzciński, T., Sitek, A.: Babynet: residual transformer module for birth weight prediction on fetal ultrasound video. In: *International conference on medical image computing and computer-assisted intervention*. pp. 350–359. Springer (2022)
17. Płotka, S., Klasa, A., Lisowska, A., Seliga-Siwecka, J., Lipa, M., Trzciński, T., Sitek, A.: Deep learning fetal ultrasound video model match human observers in biometric measurements. *Physics in Medicine & Biology* **67**(4), 045013 (2022)
18. Płotka, S., Włodarczyk, T., Klasa, A., Lipa, M., Sitek, A., Trzciński, T.: Fetal-net: Multi-task deep learning framework for fetal ultrasound biometric measurements. In: *International Conference on Neural Information Processing*. pp. 257–265. Springer (2021)
19. Płotka, S.S., Grzeszczyk, M.K., Szenejko, P.I., Żebrowska, K., Szymecka-Samaha, N.A., Łęgowik, T., Lipa, M.A., Kosińska-Kaczyńska, K., Brawura-Biskupski-Samaha, R., Işgum, I., et al.: Deep learning for estimation of fetal weight throughout the pregnancy from fetal abdominal ultrasound. *American journal of obstetrics & gynecology MFM* **5**(12), 101182 (2023)
20. Stirnemann, J., Villar, J., Salomon, L., Ohuma, E., Ruyan, P., Altman, D., Nosten, F., Craik, R., Munim, S., Cheikh Ismail, L., et al.: International estimated fetal weight standards of the intergrowth-21st project. *Ultrasound in Obstetrics & Gynecology* **49**(4), 478–486 (2017)
21. Tao, J., Yuan, Z., Sun, L., Yu, K., Zhang, Z.: Fetal birthweight prediction with measured data by a temporal machine learning method. *BMC Medical Informatics and Decision Making* **21**, 1–10 (2021)

22. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* **30** (2017)
23. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 6450–6459 (2018)
24. Wang, Y., Wen, J., Xin, J., Zhang, Y., Xie, H., Tang, Y.: 3dcnn predicting brain age using diffusion tensor imaging. *Medical & Biological Engineering & Computing* **61**(12), 3335–3344 (2023)