DK-Net: Semi-Supervised Wavelet-KAN Landmark Localization for Angle of Progression Measurement in Intrapartum Transperineal Ultrasound

Bo Deng[1], Jieyun Bai[1,2,3*], Huijin Wang[1], Shun Long[1], and Shuo Li[4]

[1] College of Information Science and Technology, Jinan University, Guangzhou, China
[2] Department of Cardiovascular Surgery, The first Affiliated Hospital of Jinan University, Jinan University, Guangzhou, China
[3] Auckland Bioengineering Institute, University of Auckland, Auckland, New Zealand
jbai996@aucklanduni.ac.nz
[4] Case Western Reserve University, Cleveland, OH, USA

Abstract. The Angle of Progression (AoP), derived from three anatomical landmarks in intrapartum ultrasound, is a crucial metric for assessing fetal head descent and supporting safe labor management. However, automated AoP measurement is severely impeded by speckle noise, acoustic shadows, heterogeneous acquisition conditions, and scarce expert annotations. To address these challenges, we propose *DK-Net*, a semi-supervised *Wavelet-KAN* framework that combines student–teacher self-ensembling with ultrasound-aware spectral–spline modeling. DK-Net introduces two Kolmogorov–Arnold Network (KAN) driven components: (1) a *domain-adaptive channel recalibration* module that dynamically adjusts feature responses to contrast variations, and (2) a *wavelet-domain spectral refiner* that learns non-linear gating functions on high-frequency subbands to suppress speckle-like perturbations while preserving landmark-defining structures. The network further enforces prediction consistency between weakly and strongly augmented views using confidence-filtered pseudo-labels to effectively leverage unlabeled data. Experiments on the public *Intrapartum Ultrasound Grand Challenge* (IUGC) benchmark demonstrate that DK-Net consistently outperforms representative single-model baselines, enabling accurate landmark localization and reliable AoP estimation for clinical labor monitoring. Our code is publicly available at https://github.com/db0725/DKNet, and the datasets can be accessed on Kaggle (https://www.kaggle.com/aspirexxx) and Zenodo (https://zenodo.org/records/18217137).

Keywords: Intrapartum ultrasound · Angle of Progression · Landmark detection · Semi-supervised learning · Kolmogorov–Arnold Networks

* Corresponding author.

1 Introduction

Labor is a dynamic physiological process in which timely assessment of fetal head descent is crucial for preventing prolonged labor and reducing maternal–fetal risks. Recent intrapartum monitoring frameworks emphasize standardized assessment and decision support throughout labor [3,6]. While digital vaginal examination is widely used, it is invasive and subject to inter-observer variability. Transperineal intrapartum ultrasound provides a more objective and repeatable alternative and is recommended by professional guidelines, particularly before operative interventions [19]. A key ultrasound-derived measurement is the *Angle of Progression* (AoP) [5], an interpretable geometric surrogate of fetal head descent computed from three anatomical landmarks: the two endpoints of the pubic symphysis (PS) contour and the tangency point on the visible fetal skull. However, manual AoP measurement is time-consuming and varies across operators, motivating automated and reproducible solutions [4,1,31,17,2,30].

Developing reliable automation for intrapartum ultrasound is challenging. Ultrasound imaging is affected by speckle-dominated formation, acoustic shadows, and operator- and device-dependent gain/contrast variations [20,13,26]. As a result, landmark-defining boundaries—especially the fetal skull edge—may be weak, partially missing, or entangled with high-frequency speckle patterns. Heatmap regression has become a standard paradigm for landmark localization due to its robustness and simplicity [32,12,9]. In intrapartum ultrasound, recent deep learning efforts have also explored anatomical modeling of the PS and fetal head (e.g., segmentation-based pipelines) under limited annotations [33,8]. Nevertheless, performance often degrades when clean expert labels are scarce, which is common in clinical datasets requiring careful quality control.

Semi-supervised learning (SSL) provides a practical way to leverage abundant unlabeled frames. Teacher–student self-ensembling methods such as Mean Teacher [14] and uncertainty-aware variants [12] have shown strong performance in medical image analysis. More broadly, recent SSL approaches combine weak/strong augmentation, confidence-based pseudo-labeling, and consistency regularization [24,11,23]. However, generic SSL recipes can be suboptimal for intrapartum ultrasound due to speckle, shadowing, and acquisition-dependent contrast shifts, motivating ultrasound-aware feature modulation and refinement.

The IUGC benchmark report [3] has shown that recent challenge methods substantially advance automated AoP measurement by optimizing performance in the benchmark setting. However, top-ranked solutions are often leaderboard-oriented pipelines that may involve dedicated preprocessing, model selection or ensembling, test-time augmentation, and post-processing. Although these strategies can improve final benchmark scores, they also increase inference complexity and may complicate practical deployment in real-time intrapartum ultrasound assistance. This distinction is important because clinical decision-support systems require not only accurate measurements, but also stable latency, simple inference procedures, and reproducible behavior across deployment environments. Rather than optimizing a full competition pipeline, our goal is to study whether ultrasound-aware calibration can improve a clean single-model semi-supervised

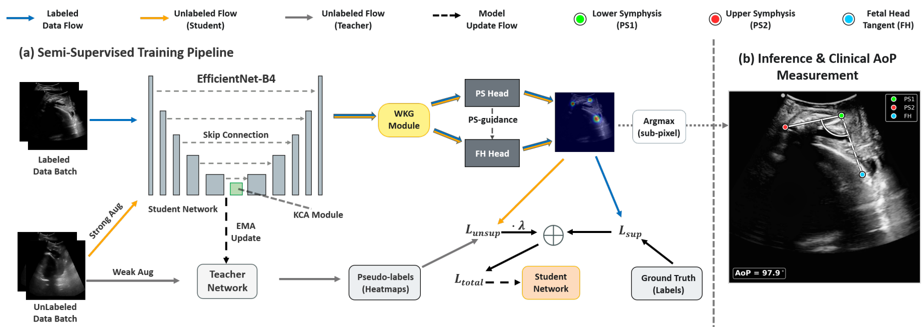

Fig. 1. Overview of DK-Net. (a) Semi-supervised training based on Mean Teacher. **(b)** Inference: landmark coordinates are recovered via argmax with sub-pixel refinement.

framework under a controlled protocol. This motivates a framework that maintains competitive AoP accuracy while emphasizing efficiency, interpretability, and deployment simplicity, especially under limited expert annotations.

Following this motivation, we propose *DK-Net*, an ultrasound-aware semi-supervised framework for AoP landmark localization. DK-Net follows a Mean Teacher pipeline with weak-to-strong consistency and confidence-based pseudo-label filtering, but introduces two ultrasound-specific innovations based on *Kolmogorov–Arnold Networks* (KAN) [16]. Recent progress has explored integrating KAN into U-Net style backbones for medical vision tasks [15], motivating us to leverage KAN's spline-parameterized non-linear mappings for more expressive and interpretable feature modulation. First, a KAN-driven *domain-adaptive channel recalibration* module adjusts feature responses to contrast variations conditioned on simple feature statistics. Second, inspired by wavelet-CNN designs for image restoration [7] and multiresolution analysis [21], we introduce a *wavelet-domain soft-thresholding refiner* that performs patch-wise KAN gating on high-frequency subbands to suppress speckle-dominated responses while preserving landmark-defining structures. Furthermore, because AoP is geometrically anchored by the PS axis, DK-Net incorporates PS-guided interaction to stabilize fetal head tangent localization when the skull boundary is degraded by shadows.

In summary, we propose DK-Net, a semi-supervised landmark localization framework for automated AoP measurement on intrapartum transperineal ultrasound. It adopts a PS-guided interaction head to stabilize fetal head tangent localization under shadowing, a KAN-based domain-adaptive channel recalibration module (KCA) to mitigate acquisition-induced contrast shifts, and a wavelet-domain refinement module (WKG) to suppress speckle noise while preserving critical landmark features. Evaluated on the public IUGC benchmark with paired bootstrap tests, DK-Net outperforms representative single-model baselines and maintains real-time inference capability.

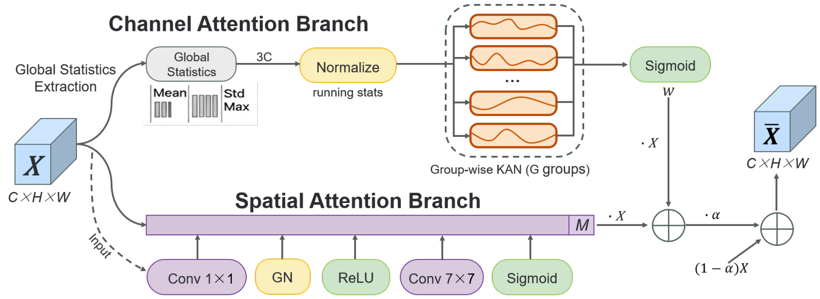

(a) Domain-adaptive channel calibration (KCA)

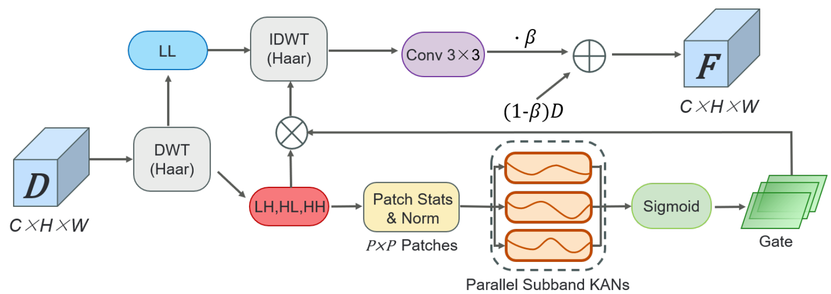

(b) Wavelet-domain adaptive feature refinement (WKG)

Fig. 2. Detailed architectures of the KAN-driven modules. **(a) KCA** uses grouped KANs to learn non-monotonic channel-wise recalibration. **(b) WKG** applies patch-wise KAN soft-thresholding to the high-frequency subbands of a Haar DWT, suppressing speckle-like perturbations while preserving coarse structures in the LL subband.

2 Methods

2.1 Problem Formulation and Overall Architecture

Given a transperineal ultrasound I, we localize three landmarks—pubic symphysis endpoints $\mathbf{p}_1, \mathbf{p}_2$ and fetal head tangent $\mathbf{h}$—to compute the Angle of Progression: $\text{AoP} = \arccos\left(\frac{(\mathbf{p}_2-\mathbf{p}_1)\cdot(\mathbf{h}-\mathbf{p}_1)}{\|\mathbf{p}_2-\mathbf{p}_1\|\|\mathbf{h}-\mathbf{p}_1\|}\right)$ [10]. We formulate this via heatmap regression, where targets are 2D Gaussians $Y_k(\mathbf{x}) = \exp(-\|\mathbf{x}-\boldsymbol{\mu}_k\|^2/2\sigma^2)$. Final coordinates are recovered via sub-pixel spatial argmax during inference (Fig. 1**(b)**).

As shown in Fig. 1**(a)**, DK-Net employs an EfficientNet-B4 [9] encoder–decoder backbone. We integrate two novel modules: a domain-adaptive channel calibrator (**KCA**, Sec. 2.2) at the bottleneck $X \in \mathbb{R}^{C_b \times h_b \times w_b}$, and a wavelet-domain soft-thresholding refiner (**WKG**, Sec. 2.3) at the decoder output $D \in \mathbb{R}^{C \times h \times w}$, yielding refined features F.

PS-Guided FH Prediction. Because the fetal skull edge is often obscured by acoustic shadows, we explicitly use the PS predictions as a structural prior. Specifically, we extract a spatial guidance map $G = \max(\sigma(\hat{Y}_1^{ps}), \sigma(\hat{Y}_2^{ps}))$ from the predicted PS heatmaps. The FH heatmap is then predicted via an interaction head that fuses the refined features and guidance: $\hat{Y}^{fh} = g_{fh}([F; G])$.

Semi-Supervised Training. We train DK-Net using the Mean Teacher paradigm [14] (Fig. 1**(a)**). The student f_θ learns from strongly augmented views, while the EMA teacher $f_{\theta'}$ generates pseudo-labels $\hat{Y}'$ from weakly augmented views. To prevent error amplification, teacher predictions are filtered by peak confidence thresholds (τ_{ps} for PS and τ_{fh} for FH). Accepted teacher peaks are converted into Gaussian pseudo-heatmaps $\tilde{Y}$, and the student is trained with a masked consistency loss:

$$\mathcal{L}_{unsup} = \frac{\sum_{k=1}^{K} m_k \, \|\hat{Y}_k - \tilde{Y}_k\|_2^2}{\sum_{k=1}^{K} m_k + \epsilon}, \tag{1}$$

where $m_k = \mathbb{I}(\max \hat{Y}'_k > \tau_k)$ is a per-channel mask and ϵ avoids division by zero. The overall objective is $\mathcal{L} = \mathcal{L}_{sup} + \lambda \mathcal{L}_{unsup}$, where $\mathcal{L}_{sup}$ is the supervised heatmap regression loss.

2.2 Domain-Adaptive Channel Calibration (KCA)

To mitigate acquisition-induced contrast shifts, KCA (Fig. 2**(a)**) extracts a descriptor $\mathbf{s} \in \mathbb{R}^{3C_b}$ containing the channel-wise mean, std, and max of the bottleneck feature X. Normalized by running statistics, channels are split into G groups, and a lightweight KAN per group predicts channel weights:

$$\mathbf{w}^{(g)} = \sigma\Big(\phi^{(g)}_{\mathrm{KAN}}\big(\mathbf{s}^{(g)}\big)\Big) \in (0,1)^{C_b/G}, \quad g = 1, \ldots, G. \tag{2}$$

In parallel, a lightweight spatial branch yields an attention mask $M \in [0,1]^{1 \times h_b \times w_b}$ via consecutive 1×1 and 7×7 convolutions. The calibrated output is fused via a learnable residual gate α:

$$\bar{X} = (1-\alpha)\, X + \alpha \,\big(X \odot \mathbf{w} + X \odot M\big). \tag{3}$$

2.3 Wavelet-Domain Adaptive Feature Refinement (WKG)

For frequency-aware speckle suppression, WKG (Fig. 2**(b)**) applies a Haar DWT to decoder features D, decomposing it into coarse structure (LL) and high-frequency subbands $S \in \{\mathrm{LH}, \mathrm{HL}, \mathrm{HH}\}$. For each subband, we partition coefficients into non-overlapping $P \times P$ patches (we use $P = 16$ in all experiments). We extract channel-averaged, absolute-value statistics (mean, std, max) to form a patch descriptor $\bar{\mathbf{t}}_{i,j}$. A subband-specific KAN predicts a soft-thresholding sigmoid gate for each patch:

$$g^S_{i,j} = \sigma\big(\psi^S_{\mathrm{KAN}}(\bar{\mathbf{t}}_{i,j})\big) \in (0,1), \tag{4}$$

which is broadcast spatially within the patch. We gate the high-frequency subbands to suppress noise, perform inverse DWT, and fuse the result back into the feature stream using a learnable parameter β:

$$F = (1-\beta)\, D \;+\; \beta\, \mathrm{Conv}_{3\times 3}\Big(\mathrm{IDWT}\big(\mathrm{LL},\, g^{\mathrm{LH}} \odot \mathrm{LH},\, g^{\mathrm{HL}} \odot \mathrm{HL},\, g^{\mathrm{HH}} \odot \mathrm{HH}\big)\Big). \tag{5}$$

Table 1. Quantitative comparison on the IUGC 2025 test set ($N = 501$). We report landmark MRE (px) and AoP MAE (°) with 95% CI; p-values are from paired bootstrap tests (20,000 resamples). AD-MT [29] is adapted to our heatmap regression setting.

Method	**Landmark Error (px) ↓**				**AoP Measurement (°) ↓**		
	PS1	PS2	FH	Mean (MRE)	**MAE**	**95% CI**	**p-value**
Stacked Hourglass [22]	8.21	10.20	23.82	14.08	4.76	[4.35, 5.21]	0.0473
ViTPose [28]	9.19	10.48	22.02	13.90	4.79	[4.37, 5.25]	0.0135
AD-MT [29]	7.52	10.46	22.65	13.54	4.84	[4.41, 5.29]	0.0109
SimpleBaseline [27]	8.30	10.14	25.32	14.59	4.94	[4.49, 5.42]	0.0045
HRNet-W18 [25]	8.07	10.63	23.35	14.02	5.03	[4.51, 5.61]	0.0010
Baseline (U-Net)	8.53	10.39	23.14	14.02	5.05	[4.57, 5.56]	0.0029
DK-Net (Ours)	**7.79**	**9.43**	**21.27**	**12.83**	**4.30**	**[3.89, 4.77]**	–

Table 2. Ablation study on the held-out test set ($N = 501$). We compare KAN-based modules (ours) against MLP counterparts under identical architectures. Best results in **bold**.

Method	KCA	WKG	MRE (px)↓	AoP MAE (°)↓
Baseline (no KCA/WKG)	–	–	14.02	5.05
MLP-based modules (SE-style attention / MLP gating)				
+ KCA_{MLP}	✓	–	13.80	4.92
+ WKG_{MLP}	–	✓	14.06	4.92
+ Both_{MLP}	✓	✓	13.09	4.49
KAN-based modules (ours)				
+ KCA_{KAN}	✓	–	13.83	4.51
+ WKG_{KAN}	–	✓	13.66	4.63
+ DK-Net (ours)	✓	✓	**12.83**	**4.30**

Both α and β are initialized to small values to encourage approximately identity behavior at the onset of training.

3 Experiments and Results

Dataset and Evaluation. We evaluate our framework on the large-scale public benchmark from the Intrapartum Ultrasound Grand Challenge (IUGC 2024/2025) [18,7,32,26]. The dataset comprises IRB-approved transperineal intrapartum ultrasound videos acquired from 24 medical centers. Following the benchmark protocol, we use 1,580 labeled and 5,497 unlabeled images for training, 95 images for validation, and report final results on the strictly held-out IUGC 2025 test set ($N = 501$), which is not seen during training. Images are resized to 512×512. We report landmark mean radial error (MRE, px) and AoP mean absolute error (MAE, °).

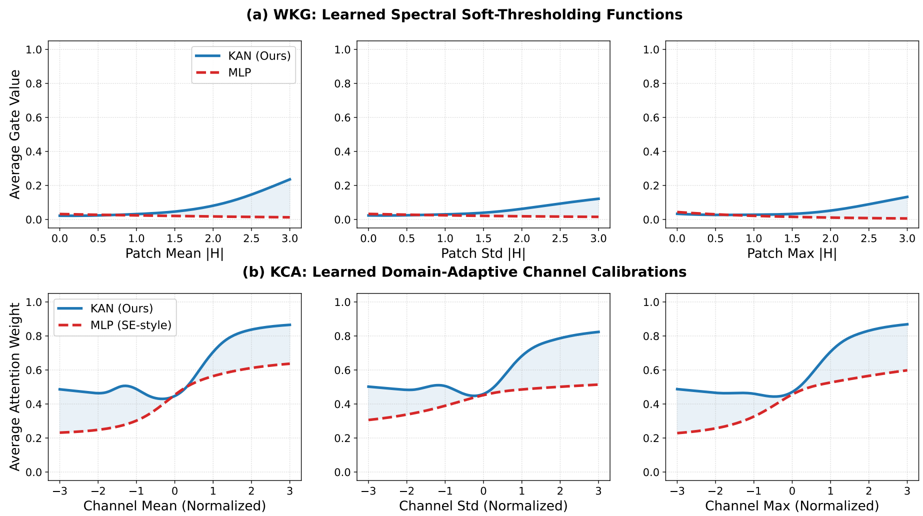

Fig. 3. Visualization of learned functions. **(a)** In WKG, KAN learns a dynamic soft-thresholding curve to preserve anatomical edges. **(b)** In KCA, KAN captures steeper, non-linear channel recalibration profiles than the standard SE-style MLP.

Implementation Details. Optimized with AdamW and mixed-precision training on an RTX 3090 Ti GPU, the student is supervised using a foreground-weighted heatmap MSE with ultrasound-aware augmentations (e.g., speckle and acoustic shadow simulation). Mean Teacher provides confidence-filtered pseudo-labels (τ_{ps}=0.75, τ_{fh}=0.82). DK-Net contains 18.48M parameters and achieves real-time inference at 40.1 FPS (batch size 1, 512×512).

Comparison with Baselines. Tab. 1 reports a controlled single-model comparison under the same semi-supervised setting. DK-Net achieves the lowest AoP MAE of **4.30°** (95% CI: [3.89, 4.77]) and an MRE of **12.83** px, and paired bootstrap tests on per-image AoP errors (20,000 resamples; $N = 501$) indicate improvements over all baselines ($p < 0.05$), including the strong ViTPose transformer baseline [28] ($p = 0.0135$) and the recent AD-MT strategy [29] ($p = 0.0109$). Notably, AoP MAE is not strictly determined by the average Euclidean landmark error (MRE): AoP depends on the relative geometry between the PS axis and the FH point, and errors parallel to the defining rays (or correlated shifts of PS endpoints) may increase MRE while having limited effect on the measured angle. DK-Net also reduces FH localization error (21.27 px vs. 23.14 px for the baseline U-Net), which is important under speckle and acoustic shadowing.

Ablation Study. Tab. 2 shows that combining KCA and WKG yields the best performance (AoP MAE 4.30° vs. 5.05° without KCA/WKG). Individually, WKG improves robustness to high-frequency perturbations, whereas KCA improves

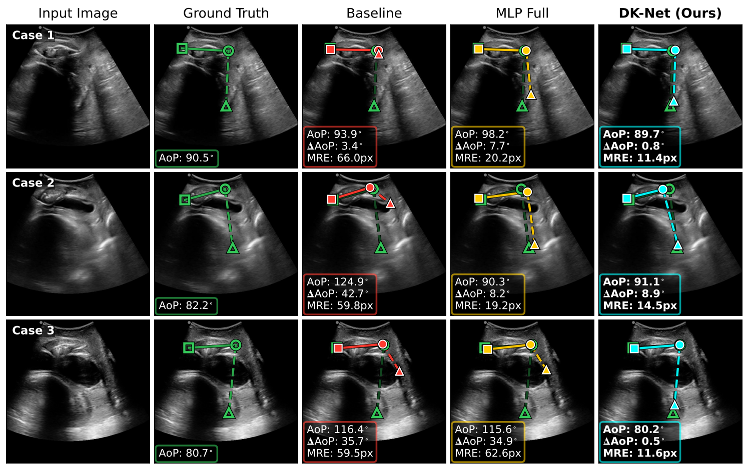

Fig. 4. Qualitative comparison on three challenging test cases. DK-Net (cyan) accurately localizes the landmarks despite heavy speckle noise and acoustic shadows.

robustness to acquisition variation via channel recalibration. Notably, the full model reduces the FH landmark error from 23.14 px (baseline) to 21.27 px, which is important since FH boundaries are often corrupted by speckle and shadowing. Finally, KAN-based modules outperform their MLP counterparts (Both$_{\text{MLP}}$: 4.49° vs. DK-Net: 4.30°), supporting the advantage of spline-based non-linear calibration for ultrasound features.

This observation suggests that KAN is not used as a generic replacement for CNNs or Transformers, but as a task-specific statistics-to-calibration function learner. In KCA, KAN maps channel-level statistics to adaptive feature recalibration weights, helping compensate for acquisition-dependent contrast variations. In WKG, KAN maps local high-frequency wavelet statistics to data-dependent gating functions, which can suppress speckle-dominated responses while preserving structured PS/FH boundary cues. Since AoP is determined by the relative geometry between the PS axis and the FH point, these calibration effects are particularly useful for stabilizing FH localization under shadowing and noise.

Interpretability of KAN Modules. Fig. 3 shows that KAN learns more expressive non-linear responses than the MLP counterpart in both WKG (a) and KCA (b), supporting flexible ultrasound feature adaptation.

Qualitative Results. Fig. 4 visualizes predictions on three challenging cases where the fetal skull is severely blurred by heavy speckle and acoustic shadows. Consequently, Baseline and MLP models suffer from severe localization drift, yielding catastrophic angular errors (ΔAoP $> 30°$ in Cases 2 and 3). In contrast, DK-Net aligns more closely with the expert ground truth and substantially reduces angular errors.

4 Conclusion

We proposed DK-Net, a semi-supervised framework for automated AoP measurement. By integrating KAN-driven channel recalibration (KCA) and wavelet-domain spectral soft-thresholding (WKG) within a Mean Teacher pipeline, DK-Net robustly mitigates ultrasound speckle noise and contrast shifts. Validated on a large-scale multi-center dataset, our method achieves clinically meaningful accuracy, demonstrating strong potential for reliable clinical intrapartum monitoring.

Acknowledgments This work was supported in part by the Guangzhou Science and Technology Planning Project under Grant 2025B03J0127, 2024B03J1283, and 2024B03J1289.

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

References

1. Bai, J., Kang, X., Wang, W., et al.: A multimodal model in the prediction of the delivery mode using data from a digital twin-empowered labor monitoring system. Digital Health **10**, 20552076241304934 (2024)
2. Bai, J., Sun, Z., Yu, S., et al.: A framework for computing angle of progression from transperineal ultrasound images for evaluating fetal head descent using a novel double branch network. Frontiers in Physiology **13**, 940150 (2022)
3. Bai, J., Tang, Y., Liu, X., et al.: Iugc: A benchmark of landmark detection in end-to-end intrapartum ultrasound biometry. Medical Image Analysis **110**, 103960 (2026)
4. Bai, J., Tang, Y., Zhou, Z., et al.: Fugc: Benchmarking semi-supervised learning methods for cervical segmentation. IEEE Transactions on Medical Imaging **45**(6), 2950–2960 (2026)
5. Bai, J., Zhou, Z., Ou, Z., et al.: Psfhs challenge report: Pubic symphysis and fetal head segmentation from intrapartum ultrasound images. Medical Image Analysis **99**, 103353 (2025)
6. Bai, J., Zhou, Z., Tang, Y., et al.: Beyond benchmarks of iugc: Rethinking requirements of deep learning method for intrapartum ultrasound biometry from fetal ultrasound videos. Medical Image Analysis **111**, 104043 (2026)
7. Chen, G., Bai, J., Ou, Z., et al.: Psfhs: intrapartum ultrasound image dataset for ai-based segmentation of pubic symphysis and fetal head. Scientific Data **11**(1), 436 (2024)
8. Chen, S., Wang, H., Long, S., et al.: Ultrasound video segmentation of pubic symphysis and fetal head for angle of progression measurement. In: Proceedings of the 6th ACM International Conference on Multimedia in Asia. pp. 1–8 (2024)
9. Chen, Y., Deng, B., Peng, Z., et al.: Comt: Co-training mean teachers semi-supervised training framework for cervical segmentation. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2025)

10. Chen, Z., Lu, Y., Long, S., et al.: Fetal head and pubic symphysis segmentation in intrapartum ultrasound image using a dual-path boundary-guided residual network. IEEE Journal of Biomedical and Health Informatics **28**(8), 4648–4659 (2024)
11. Chen, Z., Ou, Z., Lu, Y., et al.: Direction-guided and multi-scale feature screening for fetal head–pubic symphysis segmentation and angle of progression calculation. Expert Systems with Applications **245**, 123096 (2024)
12. Chen, Z., Ou, Z., Lu, Y., et al.: Uncertainty-fetal head and pubic symphysis segmentation with enhanced multi-scale features and sparse visual graph attention. Expert Systems with Applications p. 128998 (2025)
13. Deng, B., Tang, Y., Li, J., et al.: Baseline method of the foundation model challenge for ultrasound image analysis. In: 2026 IEEE 23rd International Symposium on Biomedical Imaging (ISBI). pp. 1–4. IEEE (2026)
14. Jiang, J., Wang, H., Bai, J., et al.: Intrapartum ultrasound image segmentation of pubic symphysis and fetal head using dual student-teacher framework with cnn-vit collaborative learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 448–458. Springer Nature Switzerland (2024)
15. Li, C., Liu, X., Li, W., Wang, C., Liu, H., Liu, Y., Chen, Z., Yuan, Y.: U-kan makes strong backbone for medical image segmentation and generation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 39, pp. 4652–4660 (2025)
16. Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., Hou, T.Y., Tegmark, M.: Kan: Kolmogorov-arnold networks. arXiv preprint arXiv:2404.19756 (2024)
17. Lu, Y., Zhi, D., Zhou, M., et al.: Multitask deep neural network for the fully automatic measurement of the angle of progression. Computational and Mathematical Methods in Medicine **2022**, 5192338 (2022)
18. Lu, Y., Zhou, M., Zhi, D., et al.: The jnu-ifm dataset for segmenting pubic symphysis-fetal head. Data in Brief **41**, 107904 (2022)
19. Luo, Y., Long, S., Wang, H., et al.: Edge awareness network with large kernel attention for small target segmentation from intrapartum ultrasound images. In: 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 5876–5883. IEEE (2025)
20. Luo, Y., Long, S., Wang, H., et al.: Dstcs: Dual-student teacher framework with segment anything model for semi-supervised pubic symphysis-fetal head segmentation. In: 2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 8262–8266. IEEE (2026)
21. Mallat, S.: A theory for multiresolution signal decomposition: the wavelet representation. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(7), 674–693 (1989)
22. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision. pp. 483–499. Springer (2016)
23. Ou, Z., Bai, J., Chen, Z., et al.: Rtseg-net: a lightweight network for real-time segmentation of fetal head and pubic symphysis from intrapartum ultrasound images. Computers in Biology and Medicine **175**, 108501 (2024)
24. Qiu, R., Zhou, M., Bai, J., et al.: Psfhsp-net: an efficient lightweight network for identifying pubic symphysis-fetal head standard plane from intrapartum ultrasound images. Medical & Biological Engineering & Computing **62**(10), 2975–2986 (2024)
25. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5693–5703 (2019)

26. Tang, Y., Zhou, Z., Lu, Y., et al.: Heatmap regression for automated angle of progression measurement: The baseline method for the iugc2025. In: Intrapartum Ultrasound Grand Challenge. pp. 105–117. Springer Nature Switzerland (2025)
27. Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. In: Proceedings of the European conference on computer vision (ECCV). pp. 466–481 (2018)
28. Xu, Y., Zhang, J., Zhang, Q., Tao, D.: Vitpose: Simple vision transformer baselines for human pose estimation. Advances in neural information processing systems **35**, 38571–38584 (2022)
29. Zhao, Z., Wang, Z., Wang, L., Yu, D., Yuan, Y., Zhou, L.: Alternate diverse teaching for semi-supervised medical image segmentation. In: European Conference on Computer Vision. pp. 227–243. Springer (2024)
30. Zhou, M., Wang, C., Lu, Y., et al.: The segmentation effect of style transfer on fetal head ultrasound image: a study of multi-source data. Medical & Biological Engineering & Computing **61**(5), 1017–1031 (2023)
31. Zhou, M., Yuan, C., Chen, Z., et al.: Automatic angle of progress measurement of intrapartum transperineal ultrasound image with deep learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 406–414. Springer International Publishing (2020)
32. Zhou, Z., Lu, Y.: Baseline method at the intrapartum ultrasound grand challenge 2024. In: Intrapartum Ultrasound Grand Challenge. pp. 1–10. Springer Nature Switzerland (2024)
33. Zhou, Z., Lu, Y., Bai, J., et al.: Segment anything model for fetal head-pubic symphysis segmentation in intrapartum ultrasound image analysis. Expert Systems with Applications **263**, 125699 (2025)