

# Automated Detection and Segmentation of Fetal Brain Abnormalities Using Real-Time YOLO Deep Learning Architecture

Renuka Devi<sup>1</sup>, Ritesh<sup>2</sup>, Veeresh<sup>3</sup>, Ajay Walikar<sup>4</sup>, Praveen Kyade<sup>5</sup>

<sup>1</sup>Professor, Department of CSE, Government Engineering College, Bidar, Karnataka, India

<sup>2345</sup>Student, B.E, Government Engineering College, Bidar, Karnataka, India

\*\*\*

**Abstract** - Early identification of structural anomalies in fetal brain tissue via prenatal ultrasonography constitutes a fundamental prerequisite for appropriate perinatal intervention strategies and informed parental decision-making. This investigation introduces a specialized deep learning framework leveraging the YOLO (You Only Look Once) real-time object detection paradigm, engineered to autonomously identify and spatially localize multiple categories of fetal central nervous system abnormalities in B-mode ultrasound imagery. The proposed methodology incorporates domain-specific image preprocessing techniques, hierarchical multi-scale feature representation, and customized loss function formulations optimized for medical imaging contexts. Comprehensive evaluation across a clinically-sourced dataset comprising 2,847 annotated prenatal ultrasound studies demonstrates detection accuracy of 96.5%, precision of 95.8%, recall of 94.7%, and mean average precision of 91.8%, significantly surpassing established baseline approaches including Faster R-CNN and Mask R-CNN. The framework successfully identifies diverse pathological entities including ventricular enlargement (ventriculomegaly), agenesis of the corpus callosum, cerebellar developmental deficiency, open neural tube pathology, and choroid plexus cystic formations. The sub-100ms inference latency enables practical deployment as a real-time clinical decision-support modality, facilitating standardized, operator-independent assessment protocols during routine prenatal screening. This contribution advances the intersection of deep learning methodologies and obstetric diagnostic imaging, establishing a reproducible, scalable, and clinically-translatable framework for enhancing detection sensitivity and reducing inter-observer variability in fetal neurosonography.

**Key Words:** Fetal brain abnormality detection, YOLO, deep learning, medical image segmentation, computer-aided diagnosis, prenatal diagnostics

## 1. INTRODUCTION

Prenatal diagnosis of fetal brain structural anomalies is a crucial element of obstetric imaging. Early detection of abnormalities such as ventriculomegaly, corpus callosum agenesis, cerebellar hypoplasia, and neural tube defects supports appropriate maternal counseling and perinatal management. Conventional ultrasound-based diagnosis depends heavily on expert interpretation, which introduces

operator-dependent variability and limits consistency. Recent advances in deep learning, particularly real-time object detection methods like YOLO, offer an opportunity to automate and standardize diagnostic workflows. This study develops an optimized YOLO-based framework designed for rapid and accurate detection of fetal brain abnormalities in prenatal ultrasound images.

## 1.1 Research Motivation and Contributions

Clinical motivation arises from limitations of current diagnostic practices, including prolonged examination time, inter-observer variation, and risks of diagnostic oversight. The primary contributions of this research include: specialized YOLO optimization tailored for medical ultrasound preprocessing, extensive comparative evaluation against Faster R-CNN, Mask R-CNN, and traditional machine-learning baselines, rigorous validation using a dataset of 2,847 annotated clinical ultrasound examinations, and development of a clinically oriented integration framework for obstetric diagnostic workflows.

## 1.2 Technological Evolution and Motivation

Recent advancements in convolutional neural network (CNN) architectures and single-stage object detection methodologies have catalyzed a paradigm shift in medical image analysis. The YOLO framework, first introduced by Redmon et al. and subsequently refined through multiple iterations (YOLOv3-v8), processes images in a unified, end-to-end manner rather than generating intermediate region proposals, thereby achieving substantial computational efficiency gains while maintaining or improving detection accuracy metrics.

Previous applications of deep learning to fetal imaging have predominantly employed two-stage detectors or standard CNN classification pipelines, which, while achieving reasonable accuracy levels (typically 85-92%), exhibit computational overhead unsuitable for real-time clinical workflows or require extensive post-processing. The emergence of optimized YOLO variants addresses these constraints through integrated architectural innovations including feature pyramid networks, spatial pyramid pooling, and enhanced loss functions. This investigation hypothesizes that thoughtfully adapted YOLO architecture, coupled with preprocessing optimizations specific to

obstetric ultrasound signal characteristics, can deliver clinical-grade detection performance suitable for integration into existing obstetric ultrasound workstations and departmental imaging networks.

### 1.3 Research Contributions and Scope

The primary contributions of this research comprise: (i) systematic optimization of YOLO detection architecture parameters for the specific acoustic and morphological characteristics of prenatal neurosonography; (ii) comprehensive comparative benchmarking against established state-of-the-art methodologies (Faster R-CNN, Mask R-CNN, standard SVM and Random Forest baselines); (iii) rigorous validation utilizing a substantial clinically-annotated dataset with multi-observer consensus labeling; (iv) component-level ablation studies quantifying the marginal contribution of each architectural and preprocessing innovation; and (v) development of an integrated clinical decision-support framework addressing practical deployment requirements including inference latency, model interpretability, and regulatory compliance considerations.

## 2. Literature Review and State of Knowledge

### 2.1 Deep Learning Paradigms in Obstetric and Fetal Imaging

The application of artificial intelligence methodologies to prenatal diagnostic imaging has expanded dramatically over the past five years. Burgos-Artizzu et al. established foundational work in automated fetal anatomical plane classification using CNNs, demonstrating that standard architectures (DenseNet, ResNet) achieve >92% accuracy in distinguishing standard scanning planes required for comprehensive fetal assessment. Zegarra et al. extended this work to three-dimensional pose estimation from transperineal ultrasound, employing customized CNN architectures that attain excellent discrimination between distinct fetal head positions during labor, achieving accuracies exceeding 94%. Concurrent investigations have targeted specific anomaly detection tasks. Shivaprasad et al. applied MobileNetV2 transfer learning to identify fetal brain abnormalities, achieving 90% accuracy on a smaller dataset of 1,200 images; however, this approach required post-hoc segmentation refinement and exhibited computational demands unsuitable for integration into existing clinical scanning workflows. The systematic review conducted by Ramirez and colleagues synthesized 140+ publications addressing deep learning applications across diverse fetal ultrasound modalities, highlighting that while accuracy metrics have improved substantially, challenges persist regarding generalization across different ultrasound systems, clinical populations, and imaging protocols.

### 2.2 Object Detection Architectures and Comparative Performance

Two primary architectural families dominate medical image object detection applications: two-stage detectors (R-CNN variants, Mask R-CNN) and single-stage detectors (YOLO, SSD). Two-stage methodologies first generate candidate regions of interest via a region proposal network, subsequently performing classification and refinement within proposed regions. While this approach achieves high precision, computational costs typically range from 30-100ms per image, limiting real-time capability.

Single-stage detectors simultaneously predict category membership and bounding box coordinates across the entire image, yielding 3-10x inference speedups. Recent YOLO iterations incorporate feature pyramid networks, enabling multi-scale detection crucial for identifying small anatomical structures within volumetric ultrasound data. A comprehensive medical imaging review synthesizing 60 high-quality YOLO applications documented detection accuracies ranging from 0.75-0.98 across diverse modalities including CT, MRI, mammography, and ultrasound, with particular advantages when detecting small-scale targets or requiring real-time inference.

### 2.3 Medical Ultrasound-Specific Challenges and Preprocessing Strategies

B-mode ultrasound acquisition introduces characteristic signal degradations including multiplicative speckle noise, signal-dependent intensity variation, acoustic shadowing, and reverberation artifacts. Simple Gaussian filtering inadequately addresses these phenomena due to their multiplicative nature and tissue-dependent characteristics. Bilateral filtering, which preserves edge features while attenuating noise, has demonstrated superior performance in obstetric ultrasound preprocessing pipelines, improving downstream CNN performance by 2-4% when applied appropriately. Adaptive histogram equalization enhances local contrast variation, particularly beneficial for regions of interest characterized by limited dynamic range. These preprocessing innovations collectively address the signal-to-noise ratio limitations that constrain classical computer vision approaches and establish superior feature representation foundations for deep learning architectures.

## 3. Technical Framework and METHODOLOGY

### 3.1 System Architecture and Dataset Preparation

The proposed system consists of ultrasound preprocessing, YOLO-based abnormality detection, and clinical integration modules. Preprocessing incorporates bilateral filtering for speckle noise reduction, adaptive histogram equalization for intensity normalization, and geometric standardization to 640×640 resolution. The dataset includes 2,847 prenatal ultrasound examinations annotated by expert radiologists,

covering classes such as ventriculomegaly, corpus callosum agenesis, cerebellar hypoplasia, neural tube defects, choroid plexus cysts, and normal cases. Data were divided into training (80%), validation (10%), and testing (10%) sets.

### 3.2 Image Preprocessing and Normalization

Raw ultrasound data undergo a hierarchical preprocessing sequence designed to enhance feature detectability while minimizing computational distortion:

**Stage 1 - Noise Reduction:** Bilateral filtering (kernel radius: 5 pixels, intensity sigma: 0.1, spatial sigma: 1.0) addresses multiplicative speckle noise characteristic of ultrasound while preserving anatomical boundary sharpness. This technique outperforms linear filtering approaches by maintaining feature contrast during noise suppression.

**Stage 2 - Intensity Standardization:** Adaptive histogram equalization (tile size: 8×8 blocks, clip limit: 2.0) normalizes grayscale intensity distributions across heterogeneous ultrasound systems and transducers, enhancing visibility of low-contrast anatomical features. This preprocessing step proved particularly beneficial for detecting subtle ventricular enlargement and subtle cerebellar hypoplasia manifestations.

**Stage 3 - Geometric Normalization:** All preprocessed images undergo bilinear interpolation to a standard resolution of 640×640 pixels, compatible with YOLO input specifications. This geometric standardization accommodates scanner-dependent acquisition parameters while maintaining aspect ratio integrity and avoiding excessive interpolation artifacts.

**Stage 4 - Channel Augmentation:** Single-channel grayscale ultrasound data undergo expansion to three-channel format through replicated channel stacking, enabling compatibility with standard CNN architectures designed for RGB imagery. The preprocessing pipeline was validated on randomly selected subsets of the dataset through qualitative assessment by radiologists and quantitative metrics including contrast-to-noise ratios before and after transformation.

### 3.3 YOLO Architecture Configuration and Customization

The detection framework employs an optimized YOLO variant configured with the following architectural specifications:

**Backbone Network:** CSPDarknet (Cross Stage Partial Darknet) configuration with 53 convolutional layers, incorporating residual skip connections and cross-stage connection mechanisms. This architecture balances feature extraction capacity with computational efficiency.

**Neck Structure:** Feature Pyramid Network (FPN) with four detection scales (P3, P4, P5, P6), enabling hierarchical feature fusion and multi-scale target detection. This multi-scale approach proves essential given the variable size presentations of fetal brain anomalies, from millimeter-scale lesions to anomalies spanning centimeter dimensions.

**Detection Heads:** Three detection heads corresponding to different feature scales (large, medium, small targets), each performing simultaneous classification (category membership among 7 abnormality classes plus normal tissue) and bounding box coordinate regression.

### Loss Function Composition:

The unified loss function combines four weighted components:

$$L_{total} = \lambda_1 L_{CioU} + \lambda_2 L_{obj} + \lambda_3 L_{cls} + \lambda_4 L_{dfl}$$

Where:

- $L_{CioU}$  represents Complete Intersection over Union loss for bounding box localization ( $\lambda_1 = 7.5$ )
- $L_{obj}$  denotes objectness prediction loss using binary cross-entropy ( $\lambda_2 = 1.0$ )
- $L_{cls}$  specifies classification loss across anomaly categories ( $\lambda_3 = 1.0$ )
- $L_{dfl}$  constitutes distribution focal loss addressing class imbalance ( $\lambda_4 = 1.5$ )

The weighted combination addresses the class imbalance problem inherent in medical imaging (normal cases typically constitute 60-70% of datasets) through distribution focal loss, while simultaneously maintaining precise localization through CIOU optimization.

### 3.4 Training Configuration and Data Augmentation Strategy

Model training employed stochastic gradient descent with Nesterov momentum (momentum coefficient: 0.937), initial learning rate 0.01 with cosine annealing schedule, batch size 32, and 100 epochs. Early stopping mechanisms halted training if validation loss failed to improve over 20 consecutive epochs, preventing overfitting on the relatively limited ultrasound dataset.

Data augmentation operations applied stochastically during training included:

- Random rotation ( $\pm 15^\circ$ ) to accommodate various fetal positioning scenarios

- Random horizontal/vertical flipping (probability: 0.5 each)
- Random brightness/contrast adjustment ( $\pm 15\%$  range) simulating ultrasound system parameter variations
- Random Gaussian blur (kernel size 1-3, sigma 0.1-2.0) introducing realistic acoustic noise variations
- Random cropping and repositioning (95-100% original size) enhancing spatial invariance

Augmentation parameters were carefully selected to ensure realistic variations reflecting genuine clinical acquisition variability, rather than introducing unphysical distortions. The augmentation strategy was validated through visual inspection of augmented image samples by radiologists to ensure clinical plausibility.

### 3.5 Dataset Composition and Labeling Protocol

The investigation utilized a clinically-sourced dataset comprising 2,847 prenatal ultrasound examinations acquired across multiple institutions using diverse ultrasound equipment (GE Voluson, Siemens ACUSON, Philips HD11). Gestational ages ranged from 18-34 weeks. All images were independently reviewed and annotated by three board-certified maternal-fetal medicine specialists using bounding box annotations for abnormality localization. Regions of disagreement underwent consensus review with senior radiologist arbitration, achieving  $>95\%$  inter-observer concordance.

#### Dataset composition by category:

- Ventriculomegaly (n=642): Lateral ventricular size  $>10$ - $12$ mm beyond gestational age normative
- Corpus callosum agenesis (n=289): Complete or partial absence of corpus callosum echogenicity
- Cerebellar anomalies (n=198): Cerebellar hypoplasia, agenesis, or vermian hypoplasia
- Neural tube defects (n=165): Open spinal dysraphism or anencephalic presentations
- Choroid plexus cysts (n=224): Single or bilateral cystic lesions within choroid plexus tissue
- Other CNS anomalies (n=156): Arachnoid cysts, hydrocephalus, Dandy-Walker spectrum
- Normal fetal brains (n=1,173): Anatomically normal neurosonographic findings

The dataset underwent stratified random partitioning into training (n=2,277, 80%), validation (n=285, 10%), and test (n=285, 10%) subsets, maintaining proportional representation of each abnormality category across partitions.

## 4. Experimental Design and Performance Evaluation Metrics

**4.1 Evaluation Framework** Comprehensive performance assessment employed established object detection metrics and medical imaging-specific evaluation protocols:

#### Standard Detection Metrics:

- **Accuracy (Acc):** Proportion of correctly classified images among all predictions
- **Precision (P):** Ratio of true positive detections to all positive predictions; calculated per class and macro-averaged
- **Recall (R):** Ratio of true positive detections to all actual positive instances; reflects sensitivity
- **F1-Score:** Harmonic mean of precision and recall ( $F1 = 2PR/(P+R)$ )
- **Intersection over Union (IoU):** Ratio of intersection to union of predicted and ground-truth bounding boxes; thresholds at 0.5 and 0.75
- **Mean Average Precision (mAP):** Average precision across all classes, computed at multiple IoU thresholds

#### Medical Imaging-Specific Metrics:

- **Sensitivity (True Positive Rate):** Capability to correctly identify abnormal cases
- **Specificity (True Negative Rate):** Capability to correctly identify normal cases
- **Diagnostic accuracy:** Overall correctness across abnormal and normal categories
- **False Positive Rate:** Incorrectly flagged normal cases requiring secondary review

### 4.2 Comparative Baseline Models

Performance evaluation incorporated rigorous comparison against established methodologies:

1. **Faster R-CNN:** Two-stage region-based CNN with ResNet-50 backbone, 50-layer depth, region proposal network

2. **Mask R-CNN:** Extended Faster R-CNN with segmentation branch, instance-level object characterization
3. **Support Vector Machine (SVM):** Traditional machine learning baseline using HOG (Histogram of Oriented Gradients) features
4. **Random Forest:** Ensemble-based classifier using 100 trees, HOG feature representation

All baseline models underwent identical preprocessing, data partitioning, and hyperparameter optimization protocols to ensure fair comparative assessment.

## 5. RESULTS AND ANALYSIS

The YOLO-based framework achieved significant improvements over baseline models.

**Table -1:** Performance Metrics Comparison

Metric	YOLO Framework	Faster R-CNN Baseline
Accuracy (%)	96.5	91.2
Precision (%)	95.8	92.1
Metric	YOLO Framework	Faster R-CNN Baseline
Recall (%)	94.7	89.3
F1-Score	0.952	0.904

### 5.1 Detailed Performance Analysis

The YOLO system achieved a 5.3% increase in accuracy compared to Faster R-CNN. High precision reflects reduced false-positive detections, while strong recall indicates reliable identification of true abnormalities. A mean IoU of 0.918 demonstrates accurate spatial localization aligned closely with expert-annotated ground truth.

### 5.2 Ablation Study

Ablation experiments showed accuracy drops when key components were removed: elimination of the Feature Pyramid Network reduced accuracy to 93.8%, removal of data augmentation to 91.5%, and omission of preprocessing to 92.1%. These findings confirm the contribution of each design component toward achieving the final performance.

### 5.3 Ablation Study Findings

**Table-2:** Ablation study quantifying component contributions to overall detection accuracy

Configuration Variant	Accuracy (%)	Accuracy Reduction (%)
Complete Framework (Baseline)	96.5	—
Without Feature Pyramid Network	93.8	2.7
Without Data Augmentation	91.5	5.0
Without Preprocessing Pipeline	92.1	4.4
Without CloU Loss Function	94.9	1.6
Without Bilateral Filtering	94.2	2.3

### 5.4 Receiver Operating Characteristic (ROC) Analysis

ROC curve analysis examining detection threshold variation demonstrated area-under-curve (AUC) of 0.968, indicating excellent discrimination capability across confidence threshold ranges. Clinically relevant operating points at 95% sensitivity achieved specificity of 94.8%, while 99% sensitivity maintained specificity at 91.2%.

### 5.5 Per-Class Performance Characterization

Category-specific performance varied according to lesion morphology and acoustic visibility:

- **Ventriculomegaly:** Sensitivity 97.3%, Specificity 96.1% (n=642 cases). Consistently high detection reflects distinctive geometric enlargement patterns
- **Corpus Callosum Agenesis:** Sensitivity 94.8%, Specificity 95.7% (n=289). Midline absence features provide reliable detection cues
- **Cerebellar Anomalies:** Sensitivity 91.2%, Specificity 93.4% (n=198). Subtle dimensional changes present greater detection difficulty

- Neural Tube Defects: Sensitivity 95.6%, Specificity 97.2% (n=165). Pronounced structural disruption enables reliable detection
- Choroid Plexus Cysts: Sensitivity 89.3%, Specificity 96.8% (n=224). Detection complexity relates to small lesion size and distinguishing from developmental variants
- Other CNS Anomalies: Sensitivity 88.7%, Specificity 94.1% (n=156). Heterogeneous morphologies present greater variability
- Normal Brain Detection: Specificity 98.6% (n=1,173). Excellent discrimination between normal and pathological presentations

## 6. DISCUSSION

### 6.1 Clinical Implications and Diagnostic Value

The proposed YOLO-based detection framework addresses critical clinical limitations inherent in conventional prenatal neurosonography. Real-time inference capability (68ms per image) enables concurrent display of detection annotations during live ultrasound scanning, providing objective second-reader functionality and reducing operator-dependent oversight risk. The 96.5% accuracy and 94.7% recall metrics represent substantial improvements over reported inter-observer agreement variability ( $\kappa = 0.78-0.85$ ) in conventional prenatal brain assessment, suggesting potential for standardization and diagnostic enhancement across diverse clinical settings.

Particular advantages manifest in resource-limited environments characterized by shortage of specialized maternal-fetal medicine expertise. Deployment via web-based clinical workstations or mobile applications could extend diagnostic capabilities to secondary and primary care settings, facilitating appropriate triage and referral pathways for pregnancies requiring specialized management.

### 6.2 Technical Considerations and Feature Representation

The framework's superior performance relative to two-stage detectors derives from multiple integrated mechanisms: unified end-to-end detection architecture eliminating region proposal generation inefficiencies; (ii) multi-scale feature pyramid enabling hierarchical representation of anatomically variable lesion sizes; (iii) domain-optimized loss function weighting addressing class imbalance through distribution focal loss; and (iv) ultrasound-specific preprocessing enhancing signal-to-noise characteristics prior to deep feature extraction.

The 2.7% accuracy reduction observed upon Feature Pyramid Network elimination demonstrates that multi-scale

representation proves essential for reliably detecting small lesions (choroid plexus cysts, early cerebellar hypoplasia) alongside large-scale anomalies (open neural tube defects). This aligns with theoretical understanding of feature pyramid benefits in detecting objects spanning wide dimensional ranges.

### 6.3 Comparative Analysis with Existing Literature

Reported performance levels exceed published results from analogous investigations. Shivaprasad et al.'s MobileNetV2-based classification approach achieved 90% accuracy on a smaller dataset but required post-hoc segmentation refinement without real-time capability. A recent systematic review of YOLO applications across 60 medical imaging studies documented average detection accuracy of  $0.87 \pm 0.08$  across diverse modalities, positioning this investigation's 96.5% accuracy in the upper performance percentile. The significant accuracy advantage potentially reflects: larger training dataset size (2,847 vs. 1,200-2,000 in comparable studies); specialized preprocessing optimization for ultrasound modality; and customized loss function formulation addressing medical imaging-specific challenges.

### 6.4 Limitations and Reliability Considerations

Several constraints merit acknowledgment: First, the dataset derives from specific geographic regions and ultrasound equipment types (GE, Siemens, Philips), potentially limiting generalization to other scanner brands or different imaging protocols. Secondly, the framework addresses two-dimensional B-mode ultrasound; incorporation of 3D volumetric or 4D temporal ultrasound data could enhance diagnostic capability but substantially increases computational demands. Thirdly, rare anomalies (Dandy-Walker malformation, severe holoprosencephaly) appear underrepresented in the dataset, potentially limiting detection sensitivity for these conditions.

Furthermore, the framework performs object detection and localization but does not provide quantitative biometric measurements (ventricular size in millimeters, cerebellar diameter) essential for definitive clinical diagnosis. Integration with auxiliary measurement modules could enhance clinical utility. Finally, deployment requires appropriate regulatory clearance, clinical validation across diverse populations, and establishment of standardized protocols for operator training and quality assurance.

## 7. CONCLUSION

This investigation introduces an optimized YOLO-based deep learning framework specifically engineered for automated detection and spatial localization of fetal brain structural anomalies within prenatal ultrasound imagery. Rigorous evaluation across a clinically-sourced dataset of 2,847 annotated examinations demonstrates detection accuracy of 96.5%, precision of 95.8%, recall of 94.7%, and mean average

precision of 91.8%, substantially exceeding established baseline methodologies and comparable published investigations. The framework's sub-100ms inference latency enables practical real-time clinical deployment as an objective decision-support mechanism, contributing toward operator-independent diagnostic standardization and enhanced detection consistency across diverse clinical environments.

This work advances medical imaging informatics through systematic architectural optimization, rigorous comparative validation, and explicit attention to clinical implementation requirements. The demonstrated performance provides compelling evidence supporting further development, prospective clinical validation, and regulatory pathway navigation toward real-world obstetric deployment. As healthcare systems worldwide increasingly embrace AI-enabled diagnostic tools, this investigation exemplifies methodologically rigorous approaches toward responsible, evidence-based artificial intelligence integration in perinatal medicine.

## REFERENCES

- [1] [1] Salomon, L. J., Alfirevic, Z., Bilardo, C. M., et al. (2016). ISUOG consensus statement on the role of ultrasound in screening for fetal abnormalities. *Ultrasound in Obstetrics & Gynecology*, 47(3), 313–318. <https://doi.org/10.1002/uog.16006>
- [2] [2] Grandjean, H., Larroque, D., & Levi, S. (1999). The performance of prenatal ultrasonography in the detection of fetal abnormalities in high- and low-risk women. *American Journal of Obstetrics and Gynecology*, 181(2), 346–353.
- [3] [3] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 779–788.
- [4] [4] Ultralytics YOLOv8 Development Team. (2023). YOLOv8: The latest advancement in the YOLO series. Retrieved from <https://github.com/ultralytics/ultralytics>
- [5] [5] Reis, D., Kupec, J., Hong, J., & Dagli, A. (2023). Real-time flying object detection with YOLOv8. *Journal of Robotics and Autonomous Systems*, 15(4), 89–102.
- [6] [6] Burgos-Artizzu, X. P., Coronado-Gutierrez, D., Valenzuela-Alcaraz, B., et al. (2020). Evaluation of deep convolutional neural networks for automated fetal echocardiography screening. *Nature Communications*, 11, 1–12. <https://doi.org/10.1038/s41467-020-17967-y>
- [7] [7] Zegarra, R. R., Munoz, J., Salazar, A., et al. (2024). A deep learning approach to identify the fetal head position from transperineal ultrasound during labor. *Ultrasound in Obstetrics & Gynecology*, 63(4), 512–521.
- [8] [8] Shivaprasad, S., Kumar, A., & Patil, A. (2023). Deep learning-based fetal brain anomaly detection using MobileNetV2 transfer learning. *IEEE Transactions on Biomedical Engineering*, 70(8), 2187–2198.
- [9] [9] Ramirez Zegarra, R. (2023). Use of artificial intelligence and deep learning in fetal ultrasound imaging: A comprehensive review. *Obstetrical & Gynecological Survey*, 78(6), 312–328.