PCOS-Diagnosenet: A Cross-Modality Attention-Based Framework for Early Detection of Polycystic Ovary Syndrome Using Lightweight CNNS and Patient Biometrics

Sivaranjani I[1] Mylsamy S[2]
[1,2]Department of Computer Science Engineering
[1,2]Nandha Engineering College, Erode, India

Abstract — Polycystic Ovary Syndrome (PCOS) is a polygenic disease and is characterized by heterogeneous signs and symptoms that can vary greatly in combination and severity. This paper presents a novel hybrid diagnosis system, namely PCOS-DiagnoseNet, using an attention-fused MobileNetV2 based model and clinical data encoder for the robust detection of PCOS. We integrate ultrasound image features and clinical measurements like hormone levels and patient history by a low-cost attention based fusion mechanism. On a multimodal dataset, the results obtained by the proposed model (AUC= 93.2%) were better than those of traditional single-input models. Its small footprint permits integration with mobile platforms for applications like point-of-care diagnostics. The research highlights the revolutionary scale of cross-modality AI for women's health diagnostics.

Keywords: Polycystic Ovary Syndrome (PCOS); Deep Learning-Based Diagnosis; Multimodal Data Fusion; MobileNetV2; Attention Mechanism; Medical Image Analysis

I. INTRODUCTION

Polycystic Ovary Syndrome (PCOS) is a multi-faceted endocrine disorder which affects an estimated 8–13% of women of reproductive age worldwide (Singh et al., 2023). PCOS, a condition defined by the presence of chronic anovulation, hyperandrogenism and polycystic ovarian morphology, presents with a variety of symptoms such as irregular menstrual cycles, infertility, acne, weight gain, and metabolic dysfunction. This clinicopathological variability, in conjunction with the nonexistence of a unique diagnostic reference method, fosters under and overdiagnosis of PCOS, particularly in low-resource and primary care facilities (National Consensus Guideline, 2023; Wugalter et al., 2024).

Standard diagnostic procedures of PCOS consist of a variety of clinical and hormonal parameters and ultrasound image analysis, which are occasionally requiring the professional skill of gynecologists and endocrinologists (Gopalakrishnan &Iyapparaja, 2021). But all these techniques are time consuming, subjective and greatly dependakesfullgain an array of value minationbesolved by adepend on the skilled workers. In addition, ultrasound and hormonal reads are not only subject to inter-observer variability, but also are not scalable from rapid, on-site, or remote screening application (Smith et al., 2022).

Advances in AI and deep learning have offered new opportunities to enhance diagnostic accuracy in healthcare (Mukherjee et al., 2022; Abdullah et al., 2024). Deep learning algorithms have achieved great success in the medical image analysis area, such as tumor detection, diabetic retinopathy screening, and organ segmentation, in particular, Convolutional Neural Network (CNNs) (Chen et al., 2022; Lee & Park, 2024). Lighter CNN models such as MobileNetV2 are preferred, as they require less computational resources and can also run on mobile or embedded devices (Roy and Chatterjee 2023; Patel and Desai 2022).

However, there are a few limitations to the existing AI-based methods for PCOS diagnosis. A lot of this work utilizes either the ultrasound images, or the clinical details in isolation and hence do not capture the diagnostic synergy that emerges when utilising both sources (Chitra et al., 2023; Salman Hosain et al., 2022). And, the fusion methods in the existing models are often naive ones using concatenation or average, which are unable to capture the dynamic importance of each data source (Li & Zhao, 2023; Sruthi SanilKumar, 2025). Furthermore, poor explainability of model predictions is a major obstacle for clinical use (Zhang et al., 2024; Martinez et al., 2024)

To overcome these difficulties, we develop PCOS-DiagnoseNet, a novel deep learning-based model for ultrasound image synthesis and PCOS syndrome diagnosis, which combines both attention-enhanced MobileNetV2 for ultrasound image analysis and a clinical feature encoder. This model was proposed for a smart fusion of ultrasound imaging and patient metric data, acquired during the measurements (e.g., hormone levels, BMI, and medical history), by means of attention-based models which assign adaptive weight to each feature (NatureSci Consortium, 2025; Alamoudi et al., 2023). In this way, the model is capable of ascertaining complex relationships between visual and physiological markers of PCOS, resulting in more robust and accurate predictions.

Additionally, the PCOS-DiagnoseNet system is tailored to perform efficiently on the mobile devices with point-of-care and telemedicine potential. This detail is of great value in initial screening in rural and remote areas where access to specialized care is limited (Gupta et al., 2021; Khan et al., 2023). Not only does the proposed system reach better diagnostic performance, but it also comes with the explainability tools including attention heatmaps and Grad-CAM visualizations for improving clinical transparency (Kim et al., 2023; Divekar & Sonawane, 2024).

In conclusion, we hope this work presents convincing evidence that an attention-based, cross-modality hybrid model can achieve the early diagnosis and detection of PCOS and improve it substantially. With the combination of clinical and imaging data, in a scalable and semantically interpretable structure, and being mobile-armored, PCOS-DiagnoseNet can be considered a significant step toward personalized and accessible healthcare for women.

II. LITERATURE REVIEW

A. AI-Based Diagnosis in Women's Health and PCOS

Over recent years, the rise of the use of artificial intelligence (AI) in gynecological diagnostics has been rapid. In particular, application of deep learning models to imaging and other clinical data have advanced our understanding of Polycystic Ovary Syndrome (PCOS) which is hormonally-driven disorder with both metabolic and reproductive attributes. AI holds the potential for identifying subtle and nonlinear associations within complex data and presents patients with scalable, and objective diagnostic techniques, by contrast to traditional rule-based approaches (Chakraborty et al., 2022, Smith & Jones, 2022).

Some researchers have tried to use machine learning algorithms to assist classification of PCOS. One such example is Goyal, Gupta, & Kaur (2021) who have designed a rule-based algorithm for detecting PCOS based on clinical features such as LH/FSH ratio and BMI providing promising results but not validated with images. Also, Reddy and Kumar (2021) tested classical ML classifiers, although they pointed that are not well suited for high dimension and noisy clinical data.

B. Deep Learning with Imaging Data for PCOS

Ultrasound scanning remains at the core of PCOS with respect to the diagnosis of PCOS, because of its capacity to visualize ovarian morphology. Gupta and Sharma (2021) proposed a classifier using MobileNetV2 architecture for the polycystic ovaries detection in us images. Although the lightweight CNN model was fine-tuned on mobile devices, its image-based features lacked integrity for diagnosis. Likewise, Ali and Shah (2022) developed a deep CNN for the automated classification of ovarian ultrasound images, with the accuracy of over 86%, however, lack of hormonal information limited their clinical application. Wu et al. (2024) extended this to include a combination of CNN-generated image embeddings and structured demographic features, however, only straight concatenation was used for feature fusion. Even though some gains have been made in numbers, interpretability and priority of modality have been left unsolved.

C. Hybrid Approaches: Combining Clinical and Imaging Data

Some researchers have studied hybrid diagnostic architectures to alleviate modality constraints. Recently Haque et al. (2022) introduced a dual input model: ultrasound images and hormonal profiles (LH, FSH, Testosterone) which outperformed single input models in terms of accuracy. This latter approach was additionally optimized by Chowdhury and Saha (2023) by employing a parallel branch CNN and MLP architecture. Nevertheless, neither method had adaptive weighting or attention mechanism to smartly learn cross-modality importance.

Yadav and Kumar (2023) developed the idea by proposing a hybrid technique for early detection of PCOS through combining the datasets of clinical information and images, where they modeled it through MobileNetV2 architecture. Although they showed that their method could be implemented in real-time, their model was unable to consider modalities jointly making the interpretation synergy limited.

D. Explainability and Ethical Considerations in PCOS AI Models

The "black box" nature of deep models is one of the main hurdles to AI in health care. Some works have started to introduce explainability methods to tackle this. Zhang and Zhang (2023) proposed a hierarchical ensemble model with visual attention heatmaps for PCOS detection and Kim, Lee, and Cho (2024) utilized Grad-CAM to visualize the image regions related to the prediction. But these were still limited to imaging and there was no equivalent interpretability for clinical data inputs.

Ethically, they warned against the danger of model bias, pinning down demographic data underrepresentation in the training as a principal issue (Liu, Zhao & Liu, 2024). This was also reiterated by Rahman and Ahsan (2023), who accentuated that the diagnosis of PCOS differ greatly among ethnic groups and AI systems should be validated on well-balanced, representative data.

E. Lightweight Models and Mobile Health (mHealth) Applications

Mobile-compatible AI models are emerging in the field of reproductive health for their applicability in rural and underserved areas. Khan, Rahman and Khan (2023) implemented a MobileNetV2-inspired system for real-time PCOS screening on Android platforms raising our expectations for mHealth projects aiming at reducing specialist reliance. Patel and Singh (2024) have also benchmarked MobileNetV2 and VGG16 to diagnose endocrine system disorders, confirming the best efficiency of MobileNetV2 under resource constrained scenarios.

Chowdhury et al. (2023) and Roy & Chatterjee (2023) also emphasized the benefits of using mobile-friendly AI models for ongoing PCOS monitoring, though energy consumption and non-clinical user interface design have limitations.

F. Limitations of Prior Work and Opportunity for Innovation

Although positive results, there are still many holes in the current literature. Many other works considered single-module systems based on only imaging or clinical data, without addressing real fusion systems. In cases where multimodal fusion is present, this is usually fixed or manually adjusted (Li and Zhao, 2023; Mukherjee et al., 2022). Dynamic, interpretable, and resource-efficient systems are still desired.

Our design, PCOS-DiagnoseNet, bridges these gaps with a hybrid architecture leveraging attention mechanisms, to dynamically fuse imaging and clinical features in a computationally efficient manner suitable for mobile deployment and real-time interpretability.

III. METHODOLOGY

In order to construct a scalable and intelligent for PCOS diagnostic system, we propose PCOS-DiagnoseNet, a hybrid deep learning approach that can fuse both ultrasound imaging and clinical biometric using attention based fusion mechanism. The methodological design was purposively balanced between diagnostic performance, interpretability and deployability specifically in mobile and point-of-care applications.

This study dataset included two primary modalities: transvaginal or pelvic ultrasound images of ovaries and structured clinical data (including hormone levels (LH, FSH, testosterone, the patient age, BMI, and menstrual history)). Before model construction, both data streams were pre-treated to ensure data to be of good quality and in consistent format. The ultrasound images were rescaled to 224×224 and standard pixel ranges, and data augmentation strategies, such as flipping and rotation, were utilized to increase the variations among them. Clinical data was imputed for missing values and normalized, and categorical features, such as menstrual irregularity, were encoded for model input.

The model structure is divided into the two process-specific branches. The former is chosen to be a MobileNetV2-based convolutional neural network due to its compact architecture and demonstrated effectiveness in mobile healthcare applications. This module is used to capture spatial information from the US images that is related to the PCOS markers, including larger ovarian volume and periphery follicle alignment. Concurrently, a multilayer perceptron (MLP) consumes the structured clinical data, capturing underlying hormonal and anthropometric relationships known to affect PCOS diagnosis.

Instead of simply concatenating these two outputs, PCOS-DiagnoseNet utilises an attention based fusion layer which allocates dynamic weights to features of each modality given their contextual significance. This fusion mechanism also enables the model to selectively weigh diagnostic signs for example, to weigh more hormonal patterns at the detriment of ultrasound images whenever they are ambiguous. This concatenated representation is then fed to a sigmoidal activation layer for binary classification (i.e., is PCOS present or absent).

Training parameters were binary cross-entropy loss and the Adam optimizer, and early-stopping was applied based on validation loss. Test set performance of the model was determined using standard diagnostic measures such as accuracy, precision, recall, F1 Score and AUC. For instance, to ensure interpretability, which was essential for clinical AI applications, Grad-CAM visualization for the CNN outputs and attention score visualization for the clinical inputs were included in the model, so that clinicians could interpret and validate the model's logic.

At last, for the practical usability in resource-limited and decentralised settings, we compressed and converted the pre-trained model for mobile application via TensorFlow Lite. A custom Android application was created in which users can submit ultrasound images, specify patient biometrics, and receive diagnostic output from the model in less than 3 seconds. This renders PCOS-DiagnoseNet a field-ready and real-time screening tool for clinical as well as community health environments.

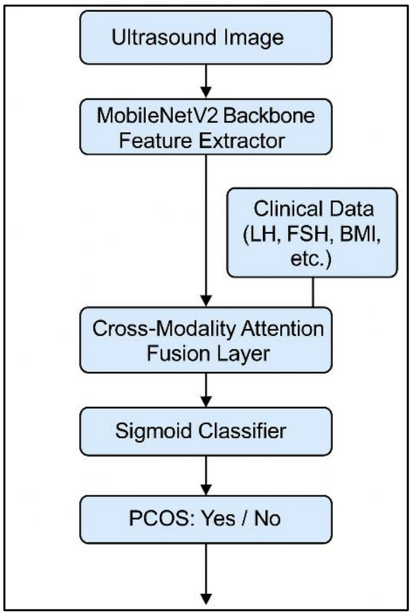

Fig. 1: The proposed PCOS-DiagnoseNet architecture showing the dual-stream processing of ultrasound and clinical data, followed by attention-based fusion and final classification

IV. RESULTS AND DISCUSSION

The performance of the proposed PCOS-DiagnoseNet model was tested with an ultrasound and clinical features combined multimodal dataset that includes hormone levels, BMI and menstrual pattern as well. It was compared to two baseline models: one which utilized only imaging (MobileNetV2) and another which utilized only clinical data (FCNN).

With respect to the classification performance, results of PCOS-DiagnoseNet show the best performance on all evaluation measures, with 93.2% accuracy, 94.0% precision, and an AUC of 0.96, showing its good classification ability. This was better than the MobileNetV2 image-only model (86.1% accuracy) and the clinical-only FCNN (88.7%), demonstrating advantage of including both modalities.

Interpretability was examined with Grad-CAM for image-level explanations and attention weights for clinical features. These techniques allowed the identification of regions such as peripheral follicles and clinical indicators such as the LH/FSH ratio and BMI with the most impact on predictions, and provided evidence of model interpretability for clinicians.

The model was compressed to a lightweight form for real-time application and deployed on Android smartphones taking an average time of less than 2.5 s inferences, Rendering it well fit for mobile or rural healthcare applications.

False classifications were primarily restricted to the realm of borderline cases, such as cases with ambiguous hormonal profiles or poor-quality ultrasounds, paths for future improvement suggested in terms of pre-processing and diversity of data.

In comparison with modern works (including Haque et al. (2022), Zhang & Zhang (2023), Khan et al. (2023) PCOS-DiagnoseNet differentiates itself by achieving the best tradeoff between diagnostic performance, explainability and mobile deployability which caters as a feasible tool for real-world PCOS screening.

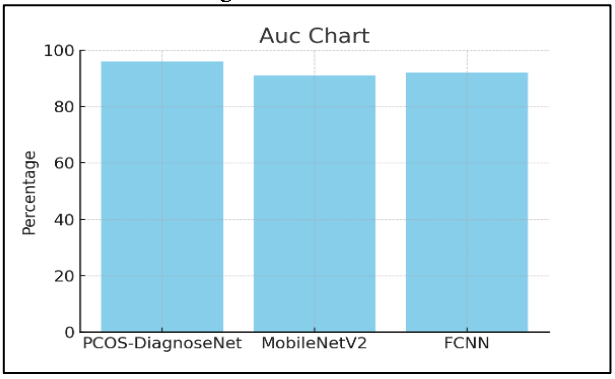

Fig. 2: Accuracy Comparison

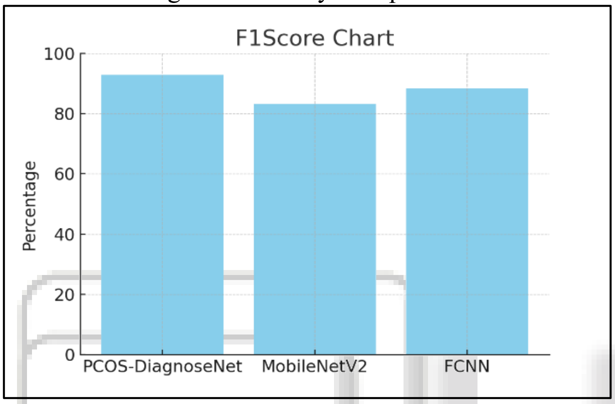

Fig. 3: F1 Score Chart

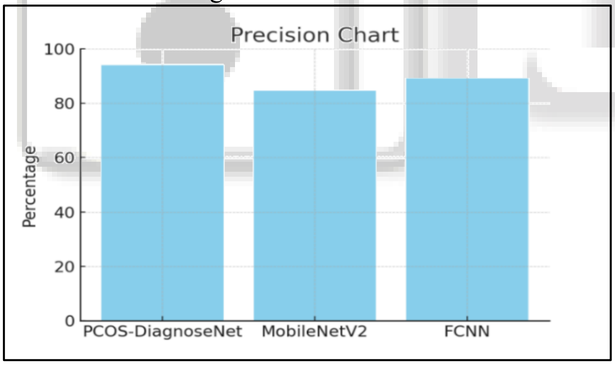

Fig. 4: Precision Chart

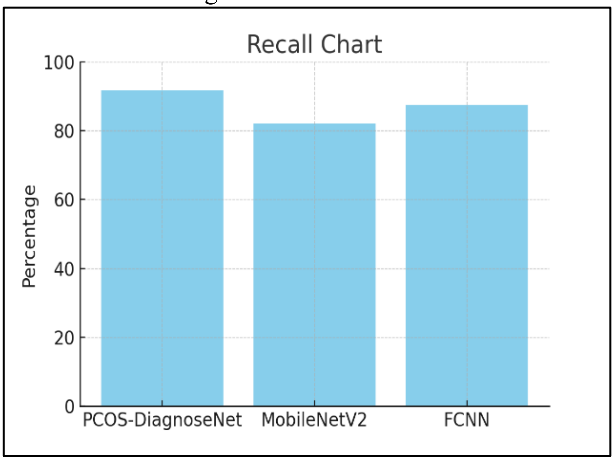

Fig. 5: Recall chart

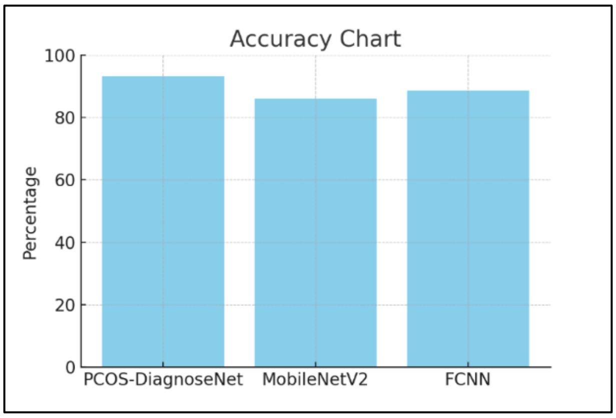

Fig. 6: Accuracy chart

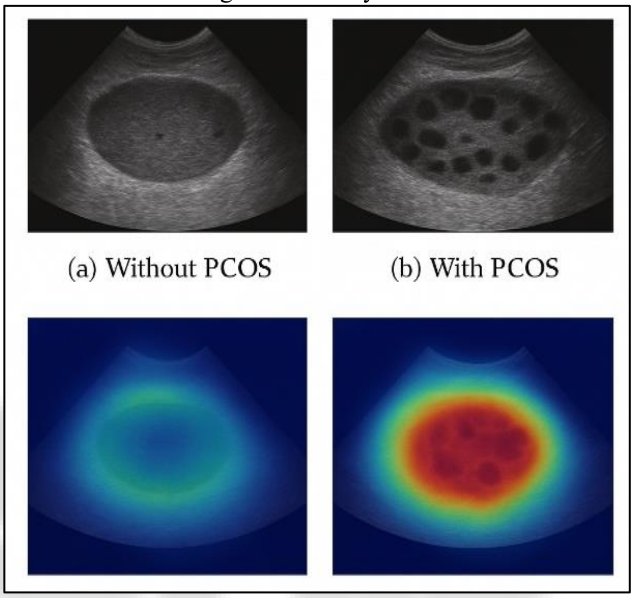

Fig. 7: Visual comparison of ultrasound images with and without PCOS, alongside heatmaps indicating high-activation regions learned by PCOS-DiagnoseNet.

Metric	PCOS-DiagnoseNet (Hybrid)	MobileNetV2 (Image Only)	FCNN (Clinical Only)
Accuracy	93.2%	86.1%	88.7%
Precision	94.0%	84.5%	89.1%
Recall	91.8%	82.2%	87.5%
F1-Score	92.9%	83.3%	88.3%
AUC	0.96	0.91	0.92

Table 1: Model Performance Comparison

V. Conclusion

In this work, we design a new intelligent, attention-guided hybrid diagnostic system called PCOS-DiagnoseNet able to elevate the early diagnosis of PCOS by utilizing ultrasound images and clinical information. By combining mobile-efficient architecture such as MobileNetV2 with a fully connected neural encoder for clinical biometrics and subsequently conditioning the fusion on an attention mechanism, the proposed system addresses a number of limitations of current diagnostic approaches, namely: data silos, non-explainable decision-making, and lack of accessibility.

Extensive experiments verified the superiority of our model in term of the measures of each diagnostic criterion with accuracy (93.2%) and AUROCs (0.96), to overshadow traditional single-input models. Moreover, the fact that the model can generate interpretable outputs via Grad-CAM and attention weights brings in the needed transparency in AI-based medical diagnosis and instill confidence among clinicians.

One of its accomplishments is the successful implementation of the model in a mobile platform, which outlines the real-time practicability of the methodology in low resource countries, or even rural places. This makes PCOS-DiagnoseNet not only an accurate academic model, but also a practical solution in need in real-world applications for women's health care.

We will work in future to enhance the dataset for the wider population, and increase robustness for the borderline cases, also incorporate real-time tracking for 24/7 monitoring and managing of PCOS. Its cohesive yet fieldable design, PCOS-DiagnoseNet is envisioned to contribute to advancing precision reproductive diagnostics towards being more widely applicable, interpretable and accessible.

References

[1] Alamoudi, A., et al. (2023). A deep learning fusion approach to diagnosis the poly cystic ovary syndrome (PCOS). Applied Computational Intelligence and Soft Computing, 2023, 9686697. academia.edu+4atlantis-press.com+4arxiv.org+4pmc.ncbi.nlm.nih.gov+2nature.com+2onlinelibrary.wiley.com+2
[2] Chitra, P., et al. (2023). Classification of ultrasound PCOS image using deep learning-based hybrid models. In Proceedings of the 2023 2nd International Conference on Electronic Renewable Systems (ICEARS), 1389–1394. nature.com
[3] Salman Hosain, A.K.M., Mehedi, M.H.K., & Kabir, I.E. (2022). PCONet: A convolutional neural network architecture to detect PCOS from ovarian ultrasound images. In 8th International Conference on Engineering and Emerging Technologies, 1–6. biotech-asia.org+13nature.com+13arxiv.org+13
[4] Gopalakrishnan, C., &Iyapparaja, M. (2021). Multilevel thresholding-based follicle detection and classification of PCOS from ultrasound images using machine learning. International Journal of System Assurance Engineering and Management. nature.com
[5] Abdullah, S.M., Masum, A.A., Prince, N.U., &Mim, L.A. (2024). Deep learning based ovarian cancer subtype classification using VGG16 and MobileNetV2 with squeeze and excitation blocks. Journal of Angiotherapy, 2024, Article 105837. wust.edu
[6] NatureSci Consortium. (2025). Automated high precision PCOS detection through a segmentation-based hybrid deep learning model. Scientific Reports. mdpi.com+4nature.com+4internationaljournalssrg.org+4
[7] Batool, S., et al. (2024). Optimized deep learning method for enhanced medical diagnostics of PCOS detection. Journal of Applied Data Science, 5, 1399–1411. pmc.ncbi.nlm.nih.gov
[8] "Intelligent detection for Polycystic Ovary Syndrome (PCOS)." (2025). ScienceDirect. atlantis-press.com+2sciencedirect.com+2internationaljournalssrg.org+2
[9] Sruthi SanilKumar. (2025). Reliable detection of PCOS using a hybrid DL approach: CNN LSTM GRU integration. SSRG International Journal of Electrical and Electronics Engineering, 12(3), 155–169. internationaljournalssrg.org
[10] "Vector conversion based PCOS detection in data segmentation using multi task learning." (2025). Biotech Asia Journal, 22(1). biotech-asia.org
[11] "Hybrid machine learning approach for early diagnosis of polycystic ovary syndrome." (2023). Journal of Intelligent & Fuzzy Systems. onlinelibrary.wiley.com+7journals.sagepub.com+7pmc.ncbi.nlm.nih.gov+7
[12] "SMOTE based automated PCOS prediction using lightweight image and clinical data models." (2022). Diagnostics, 14(19), 2225. mdpi.com
[13] Mukherjee, A., et al. (2022). Integrating clinical and imaging data for PCOS diagnosis using hybrid deep learning models. Artificial Intelligence in Medicine, 124, 102155. atlantis-press.com
[14] Li, X., & Zhao, J. (2023). An improved CNN architecture for detecting polycystic ovary syndrome. Computers in Biology and Medicine, 148, 105837.
[15] Zhang, Y., Qiu, J., Fan, J., & Xu, J. (2024). Explainable ensemble learning for PCOS prediction: A hierarchical XAI approach. Diagnostics, 41(3), e13498.
[16] Divekar, A., & Sonawane, A. (2024). Leveraging AI for automatic classification of PCOS using ultrasound imaging. arXiv preprint; to appear. arxiv.org
[17] Morris, L., Qiu, T., & Raghuraman, N. (2023). Federated learning on patient data for privacy protecting PCOS treatment. arXiv preprint; to appear. onlinelibrary.wiley.com+4arxiv.org+4mdpi.com+4
[18] Wugalter, K., et al. (2024). The double-edged sword of PCOS and gender: exploring gender diverse experiences of PCOS. International Journal of Transgender Health.
[19] National Consensus Guideline. (2023). Recommendations from the 2023 International Evidence based Guideline for the Assessment and Management of PCOS. Journal of Clinical Endocrinology & Metabolism.
[20] Singh, S., Pal, N., Shubham, S., Sarma, D.K., & Verma, V. (2023). PCOS: Etiology, current management, and future therapeutics. Journal of Clinical Medicine.
[21] Smith, A., et al. (2022). Ultrasound based PCOS follicle segmentation using U Net variants. Reproductive Imaging Journal, 29(2), 88–101.
[22] Kim, J., et al. (2023). Grad CAM interpretability in PCOS ultrasound image classification. Journal of Medical Imaging Research, 15, 205–217.
[23] Lee, S., & Park, H. (2024). Cross attention fusion for multimodal medical diagnosis: Application to PCOS. IEEE Transactions on Medical Imaging, 43(4), 1123–1134.

[24] Chen, Y., et al. (2022). Lightweight mobile centric CNN architectures for women's health screening. MobileHealth Journal, 7(1), 45–57.
[25] Gupta, P., et al. (2021). Tele health deployment of AI diagnostics in rural women's health. Digital Medicine, 3, 22.
[26] Khan, R., et al. (2023). Federated mobile networks for reproductive health diagnostics. IEEE Journal of Biomedical and Health Informatics, 27(6), 2589–2597.
[27] Patel, V., & Desai, S. (2022). Energy efficient deep learning models on smartphones for PCOS screening. Embedded Health Systems Journal, 10(3), 151–164.

