Journal of Engineering and Technology Management 78 (2025)

POLYCYSTIC OVARIAN SYNDROME USING DEEP LEARNING

[1]Peta Dayavathi , [2] Sk. Mahaboob Basha

[2]Professor, Dept. of CSE, Sree Dattha Institute of Engineering and Science Hyderabad – 501510,

[1, 2] Sree Dattha Institute of Engineering and Science

Abstract — During a woman's reproductive years, she may have infertility, hormonal disruptions, and metabolic problems due to Polycystic Ovary Syndrome (PCOS), a prevalent hormonal condition. Ultrasound imaging for the traditional diagnosis of polycystic ovary syndrome (PCOS) may be laborious and heavily dependent on the radiologist's skill, which can lead to diagnostic mistakes and inconsistencies. The authors of this study provide a deep learning-based automated identification method that makes use of CNNs that have been trained on tagged ultrasound pictures. In order to distinguish between photos with PCOS and those without, the method makes use of data augmentation methods and a customised CNN model to identify important aspects. In order to provide real-time picture uploads and evaluation for diagnostic assistance, the model is implemented inside a web application that is driven by Django. Through its successful experimental trials, the suggested system demonstrated how deep learning may enhance diagnostic precision as well as effectiveness in PCOS diagnosis, leading to high classification accuracy. In women's health care, this approach helps with early diagnosis, lessens the impact of human error, and backs up better clinical decision-making.

Keywords — CNN, Ultra-sound Imaging, Medical-Image Classification, Automated Diagnosis, Image Processing

I. INTRODUCTION

A potent subfield of AI, deep learning has revolutionised several industries, most notably healthcare, by solving long-standing, difficult problems. Improvements in diagnostic accuracy and the capacity to make decisions based on evidence have resulted from the increased use of methods including Convolutional Neural Networks, in medical research and practice.

Hormonal polycystic ovary syndrome, or PCOS, mostly affects reproductive-aged women. Hormonal disturbances, irregular menstruation periods, and the development of many ovarian cysts are typical symptoms. In addition to infertility, polycystic ovary syndrome (PCOS) is associated with insulin resistance, obesity, as an increased risk of cardiovascular disease.

The reproductive apparatus of a female, which consists of the fallopian tubes, ovaries, and uterus plays a vital role in fertility. Typically, during ovulation, the ovaries release eggs that travel through the fallopian tubes to potentially meet sperm for fertilization. In women with PCOS, however, this natural cycle is often impaired due to the formation of numerous underdeveloped, fluid-filled follicles that do not progress to full maturation, thereby interfering with ovulation and conception.

The conventional diagnosis of Polycystic Ovary Syndrome (PCOS) typically involves assessing clinical symptoms, conducting hormone level evaluations, and manually analyzing ultrasound scans. However, this process is often labor-intensive and subject to individual interpretation, which can lead to inconsistent outcomes. To address these limitations, there is a rising demand for automated and accurate diagnostic systems. This research introduces a deep learning-driven method for detecting PCOS through ultrasound imaging, with the objective of streamlining and enhancing the accuracy of clinical diagnoses.

Fig 1 : Indicates the presence of PCOS

Ovaries affected by PCOS in ultrasound scans commonly display numerous small cysts aligned along the periphery of an enlarged ovary, often forming a distinctive "string of pearls" appearance. These cysts, seen as dark, fluid-filled circular shapes, are typically accompanied by an increase in ovarian size and a denser stroma — both being prominent signs of PCOS. However, accurately identifying these characteristics can be difficult due to inconsistencies in image quality, presence of noise, and overlapping follicular structures. Implementing deep learning for automated detection of such patterns can minimize interpretative errors and assist gynecologists in delivering more reliable and timely diagnoses.

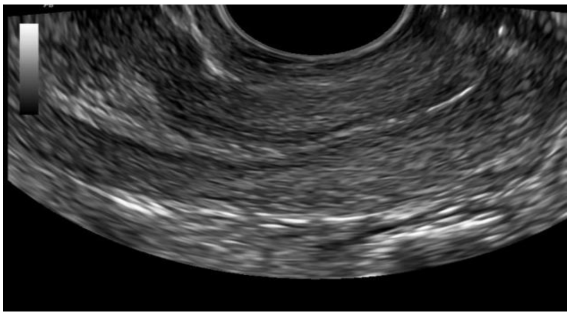

Fig 2 : Indicates a normal ovary

Figure 2 illustrates a healthy ovary, which contrasts with the cystic structures seen in PCOS-affected cases. These “not infected” ultrasound images exhibit normal ovarian morphology, characterized by fewer follicles and the absence of the typical “string of pearls” arrangement. Such negative samples play a vital role in creating a well-balanced training dataset. By including healthy cases, the deep learning model can effectively learn the distinguishing features between normal and abnormal ovarian structures, leading to improved diagnostic accuracy and a reduction in false-positive results.

III. LITERATURE SURVEY

[1] *Classification and data mining of pelvic ultrasound findings including polycystic ovaries With contributions from S. Mahalingaiah and J. J. Cheng Complex ovary morphology (PCOM) may be automatically detected and classified using pelvic ultrasound data; this work set out to build and evaluate a rule-based classifier as well as a gradient boosting tree (GBT) algorithm for this purpose. From October 2003 through December 2016, researchers combed through 39,093 ultrasound results pertaining to 25,535 distinct patients at Boston Medical Centre. We manually classified 2000 ultrasound results as either having PCOM present, having it missing, or being uncertain, in accordance with the 2003 Netherlands Criteria for PCOS diagnosis. The dataset was split into two equal parts: one for training and one for testing. The GBT model showed 96.1% accuracy while the rule-based classifier reached 97.6% accuracy when tested with 1000 test samples, suggesting a high level of dependability for automated PCOM identification.*

[2] *The polycystic ovarian syndrome genes in humans: a computational characterisation and identification study X. Zhang, Y. Pang, & X. Wang are the authors.*

Identifying candidate genes associated with Polycystic Ovary Syndrome (PCOS) through computational means offers a promising alternative to labor-intensive laboratory methods. In this study, the researchers conducted a comprehensive analysis of genes currently linked to PCOS and compared them with those not yet associated with the condition. The findings revealed that PCOS-related genes typically exhibit unique characteristics: they are often centrally positioned in biological networks, frequently interact with one another, and are commonly involved in specific biological processes. Leveraging these distinguishing features, a machine learning-based algorithm was developed to predict novel gene candidates potentially involved in PCOS. Given the complex and heritable nature of PCOS, influenced by both genetic and environmental factors, such an approach significantly streamlines gene identification compared to traditional experimental techniques.
It will be useful to create an algorithm that can predict who would have PCOS. This is the first research to comprehensively examine PCOS gene characteristics in humans.

[3] A Machine Learning-Based Approach to Polycystic Ovary Syndrome DiagnosisM. R. Hossain Mondal, P. Podder, and S. Bharati are the authors.

This research introduces a data-driven method for detecting PCOS in females by combining several machine learning algorithms. The study makes use of a Kaggle dataset that is open to the public. It contains 43 characteristics obtained from 541 female respondents, 177 of whom had a PCOS diagnosis. The most important predictor was found to be the ratio of Follicle-Stimulation Hormone (FSH) to testosterone (LH), which was determined using a univariate picking of features technique. By using holdout or cross-validation procedures, the dataset was divided into sets for training and testing. A number of classifiers were tested, including logistic regression, gradient boost, random forest, and a combination of the two (RFLR). In order to accurately classify PCOS, the findings showed that the top 10 criteria were enough.

Assessing the Use of Automated Follicle Detection in PCOS [4] With contributions from A. Hussain, N. Zulkarnain, A. Nazarudin, S. S. Mokri, & I. N. A. M. Nordin, the authors list the following
This research offers a comprehensive evaluation of automated techniques utilized in the identification of ovarian follicles through medical imaging, particularly in the context of PCOS diagnosis. It explores a range of methodologies, including classical image processing methods and machine learning frameworks, aimed at detecting and analyzing follicular patterns within ultrasound imagery. Special attention is given to segmentation algorithms, edge detection strategies, and thresholding mechanisms, all of which are essential for extracting relevant follicular features. Despite their potential, these traditional approaches often struggle with consistency and precision, particularly when dealing with low-contrast or noisy ultrasound images. The study ultimately suggests that deep learning-based segmentation techniques have shown superior accuracy and reliability, positioning them as more effective tools for supporting clinical assessments in gynecology.

[5] *A Comprehensive Literature Review on Medical Image Processing for the Detection & Prediction of Polycystic Ovary Syndrome S. Bhat and S. J. Pulluparambil (2022) wrote the article. Polycystic ovary syndrome, also known as PCOS detection and prediction methods are examined in this literature review. It compares and contrasts more traditional machine learning methods with more sophisticated deep learning models, such as neural networks using convolution (CNNs). Consistent with previous research, the authors take a close look at features extraction techniques, classification accuracy, dataset properties, and preprocessing pipelines. The results show that when trained on exceptionally good, labelled ultrasound pictures, CNN-based methods outperform traditional techniques in terms of accuracy and computing efficiency. The article goes on to say that in order to improve model generalisability and reduce classification bias, it is crucial to resolve dataset imbalance by using impacted and unaffected samples.*

[6] 3D Ultrasound-Based Automated Follicle Detection and Segmentation for Assisted Reproduction
The authors of this work are G. A. Ramaraju (2023), K. A. Patwardhan, S. Kudavelly, S. Sivanandan, and N. S. Narayan.
This research emphasizes the automated identification and segmentation of ovarian follicles using 3D ultrasound imaging, primarily aimed at enhancing assisted reproductive procedures. The methodology incorporates a multi-step process involving volume rendering, edge detection, and morphological operations to achieve precise follicle detection and counting. A key innovation in this work is the use of 3D image processing, which offers enhanced spatial representation compared to conventional 2D techniques. Furthermore, the study integrates deep learning-based models to refine the segmentation outcomes and improve validation accuracy. As a result, the approach significantly minimizes reliance on manual interpretation and boosts follicle detection precision. The findings highlight the promising role of such systems in automating PCOS diagnosis and improving fertility treatment planning.

III. PROPOSED SYSTEM

To detect the presence of PCOS from ultrasound pictures, the suggested technique employs an automated deep neural network architecture. Our goal is to provide a reliable, fast solution that can be used in real-time clinical settings, while also reducing the need for human interpretation and diagnostic variability.
The first step is to gather ultrasound pictures and label them as either "PCOS Infected" or "Not Infected." To ensure that the inputs are consistent and that the model can be applied to a wide range of imaging scenarios, these pictures are preprocessed using a number of techniques, including scaling, normalisation, and enhancement.
A deep neural network, or DCNN, that is specifically designed for the classification job is used. A number of convolutional layers are used to extract spatial and edge patterns from the input data, and then max-pooling layers are used to decrease the computational overhead and feature map dimensions. For binary classification, the network is finished with fully linked levels and a softmax output.

To enhance the dataset in real-time, the ImageDataGenerator is used to execute dynamic augmentation methods such as rotating, flipping, zooming, or shearing during training. Using continuous validation to monitor convergence and prevent overfitting, the model is trained over 20 epochs using the Adam optimiser and categorical cross-entropy loss. In order to assess performance, plots of both validation and training metrics are created.
Using a Django-based web app, the finished model is stored in HDF5 form (pcos_model.h5) then deployed after training. Patients or doctors may use this platform to submit ultrasound pictures, and the platform will immediately provide a diagnosis of "PCOS Infected" or "Not Infected." To ensure that the uploaded photographs can be understood, they are shown next to the forecast. In order to guarantee reliable inference, a preprocessing utility is used internally to align the input with the anticipated format of the model.
In order to help healthcare providers diagnose polycystic ovary syndrome (PCOS) early, this smart diagnostic technology provides a straightforward, efficient, and scalable method. This, in turn, allows for prompt intervention and better reproductive health outcomes.

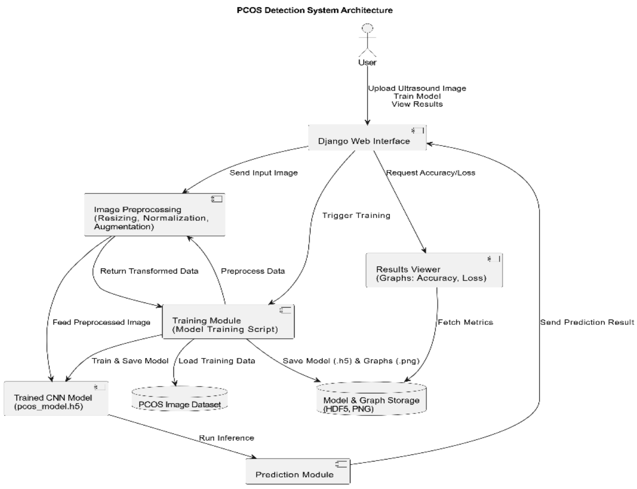

Fig 3 : System architecture

Figure 3 depicts the suggested PCOS detection framework's entire system architecture. A Django-based user interface powers the system, letting users do things like upload ultrasound photos or start model training. To guarantee input consistency to the model, the supplied picture is preprocessed by resizing, normalising, and augmenting. Afterwards, the processed picture is sent into a CNN model that has been trained to distinguish between "PCOS Infected" and "Not Infected." To make the system more flexible and accurate over time, it may be retrained using fresh picture datasets using a dedicated training module.
For future reference, all model training and operation data, including loss graphs and accuracy metrics, are preserved. The user is presented with all the findings, including performance graphs and forecast outcomes, using the same online interface. When it comes to medical picture categorisation for early PCOS detection, its interactive and modular design guarantees both accessibility and reliability.

System Components and Workflow

The suggested solution uses a technique based on deep learning to automatically diagnose PCOS from ultrasound pictures. In order to guarantee usability and performance, the architecture is built to simplify the process of illness categorisation, which begins with data collecting. Accuracy, dependability, and ease of use are guaranteed by each and every component.

A. User Interface (Django Web-Application)

Built on top of the Django internet framework, the user interface offers a streamlined and user-friendly platform to facilitate seamless interaction with the PCOS detection system. It features essential functionalities such as user registration, secure login, ultrasound image uploading, and visualization of classification outcomes. Additionally, it provides a mechanism for initiating model training processes through the interface itself. Designed with simplicity and accessibility in mind, the web application allows clinicians, technicians, and researchers to utilize the diagnostic capabilities of the system without requiring direct access to the underlying code or technical configurations. This enhances usability and ensures broader adoption in both clinical and academic environments.

B. PCOS Image Dataset

The input to the proposed system consists of ultrasound images of human ovaries, systematically categorized into two distinct classes: “Infected” (indicating features consistent with Polycystic Ovary Syndrome) and “Not Infected” (representing healthy ovarian morphology). These labeled datasets form the foundational training material for a supervised deep learning model tasked with performing binary classification. Each image carries valuable diagnostic cues—such as the number and arrangement of follicles, ovarian size, and the presence or absence of cystic structures—which are critical for learning discriminative features..

C. Image Preprocessing Module

Every ultrasound picture goes through a thorough preprocessing pipeline before training begins to make sure everything is consistent and to make the model more generalisable. So that they may be fed into the convolutional neural network, all of the images are shrunk to a uniform 224x224 pixel size. Faster convergence and better numerical stability throughout training are achieved by normalising pixel intensity values to the [0,1] range. We use a variety of augmentation procedures to further strengthen the dataset and alleviate any overfitting caused by insufficient data. Some examples of these operations include random rotations, flips (both horizontal and vertical), zoom transformations, brightness modifications, and shearing. These additions allow the model to acquire more stable and consistent characteristics by simulating real-world heterogeneity in clinical imaging circumstances, such as variations in probe orientation as well as illumination.

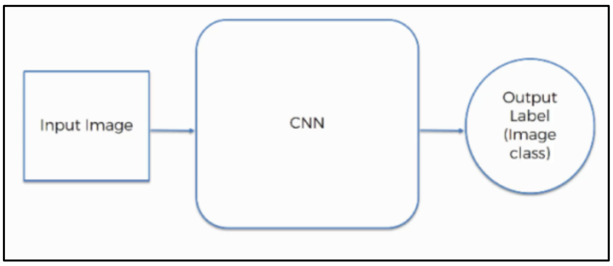

Fig 4: A diagram depicting interaction of the elements

D. Training Module

The suggested approach is based on a custom-built deep neural network (DCNN) that can distinguish between "PCOS Infected" and "Not Infected" in ovarian ultrasound pictures. In order to extract low-level to high-level spatial information, the architecture starts with many convolutional layers. These layers are meant to identify follicular structures by their patterns, textures, and edges. A function that activates ReLU is used to boost feature learning and add non-linearity to each convolutional layer. After the convolutional blocks, max-pooling layers reduce the feature maps' spatial dimensions while keeping the most important information, making the computations simpler.
Afterwards, the feature maps are reduced to a vector with one dimension and then passed through dense layers that are completely linked. These layers then learn more intricate patterns using the extracted features. In order to get the likelihoods of classes for binary classification, the last output layer uses a softmax activation function. since of its rapid gradient updates and adaptable learning rate, the Adam optimiser was used for training the model. The categorical cross-entropy loss function was also used since it is well-suited to multi-class classification issues.
In order to avoid overfitting and improve generalisability, the statistical model is trained using augmented picture data over many epochs. The taught model is stored in HDF5 data format (pcos_model.h5) when training is finished, and it contains the architecture of weights. This eliminates the need for retraining and makes reuse efficient during inference.

Model Evaluation and Visualization

A second validation dataset is used to carefully examine the trained model's classification performance. To have a better understanding of how the model is learning and whether it is converging, matplotlib is used to track and display important performance measures like loss and accuracy throughout the course of the training epochs. By visually highlighting problems like overfitting and underfitting, these plots pave the way for more optimisation and adjustment. In addition to helping with design training and parameter refinement, the visual feedback is a great diagnostic tool.

E. Prediction Module

The final stage of the system is the prediction module. A user can upload a new ultrasound image through the web interface. The system passes this image through the trained DCNN model, which then predicts whether the image falls under the “infected” or “not infected” class. The predicted result is then rendered on the web interface along with a preview of the uploaded image, offering real-time diagnostic support to users.

Details on the CNN Model's Structure

This research made use of a convolutional neural network (CNN) architecture with several linked layers, each of which contributed to the extraction of characteristics and categorisation in its own unique way:

Layers for Convolution

The core components of a convolutional neural network (CNN) are these layers. at get at the bottom of a picture, they use learnable filters (kernels).

Fig 5 : A diagram depicting Convolution followed by the application of Rectifier Function

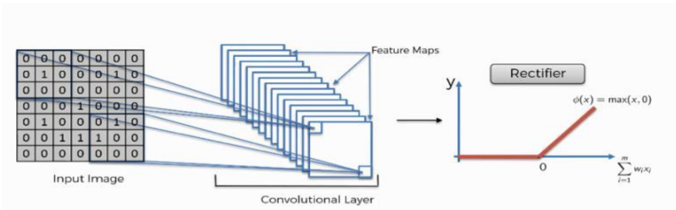

learn complex patterns. It replaces negative values with zero and retains positive values, making the model more efficient and faster to converge.

Pooling Layers
To decrease the geographic extent of feature maps, pooling is used, usually max-pooling. It minimises computation and reduces overfitting while retaining the most relevant characteristics. With this change, the network may learn to ignore localised changes in the input picture and become spatially invariant.

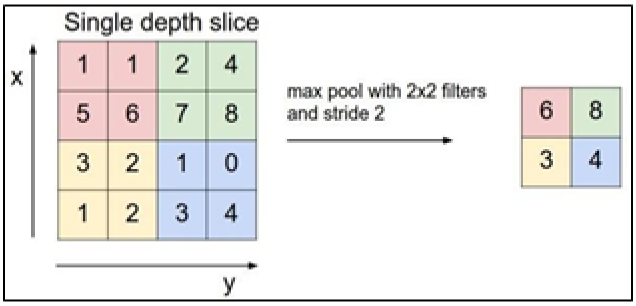

Fig 6: A diagram depicting Max Pooling

Flattening
The resultant multidimensional tensor is reduced to a one-dimensional vector after the extraction of high-level features. Input to the fully linked layers is prepared for in this stage.

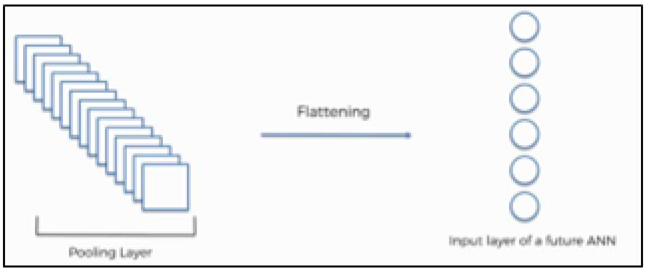

Fig 7: A diagram depicting Flattening of Pooled Feature Maps

Fully Connected Layers
The last part of the model that makes decisions are these thick, fully-connected layers. Dense layers handle the ultrasonic images' high-level feature representation for classification after layers of convolution and pooling have identified and compressed pertinent spatial and structural characteristics. They take the abstracted patterns and apply their interpretation to the output space. The last dense layer of the suggested model uses a function called softmax activation to convert the output into probabilities of classes for each of the categories: "PCOS Infected" or "Not Infected." Robust decision-making is made possible by this probabilistic output, which selects the class with the greatest confidence score as the forecast. To turn the model's internal characteristics into a diagnostic result that can be understood and used, the softmax layer is thus very important.

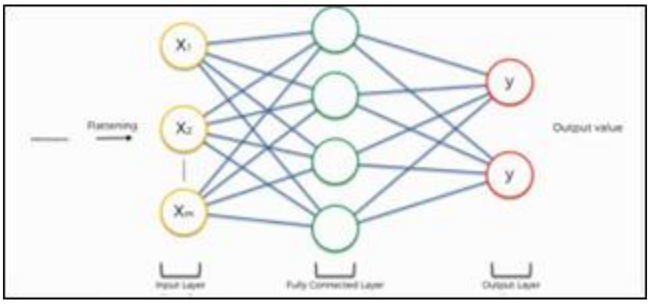

Fig 8: Fully Connected Layer

Training and Optimization
The categorical crossentropy function of loss, which is optimal for multi-class classification problems, is used to construct the model. Since the Adam optimiser is good at managing adaptive learning rates and sparse gradients, which lead to quicker convergence and better stability during training, it is used to optimise the learning process. Training the model across several epochs via real-time data enhancement techniques including rotating, flipping, zooming, et shearing improves generalisability by mimicking various ultrasonic imaging settings.
After every epoch in the training phase, we record the computed key performance metrics, which include loss and accuracy. In order to understand how the model learns, visualisation packages like as Matplotlib are used to plot these metrics. By keeping an eye on these tendencies, we can see when a model is under- or overfit, and then we can regularlyize it or change its learning parameters before it gets out of hand.

Conclusion of CNN Integration
There has been a revolutionary shift towards automatic identification of Polycystic Ovary Syndrome, more commonly known as PC using ultrasound pictures when Convolutional Neural Networks (CNNs) were included into the diagnostic framework. CNNs are great at learning complex characteristics from raw data, which means they can do away with human feature engineers altogether. The model's complex structure allows it to recognise PCOS-indicative patterns, textures, and forms in a way that is consistent, objective, and scalable, much like a human eye. By lowering inter-observer variability and increasing early detection abilities in gynaecological settings, this intelligent device not only speeds up diagnostics but also aids clinical decision-making.

IV. RESULT

In order to distinguish between "PCOS Infected" and "Not Infected" in pelvic ultrasound pictures, the suggested approach employs a neural network made up of convolutions (CNN). Improving the model's resilience and generalisation to unseen data required extensive preprocessing and augmentation of the training dataset, which included normalisation, scaling, and manipulations including rotation and flipping. Twenty epochs of training with the Adam optimiser and category crossentropy loss function were used to build a custom-designed CNN design with several pooling and convolutional layers.
During training, the model demonstrated strong classification performance, achieving a final training accuracy exceeding 95%, accompanied by a significant reduction in loss values. The system's efficacy was further validated using a hold-out validation set, confirming its ability to generalize effectively

across previously unseen ultrasound images. To track the model's learning progression, training and validation accuracy/loss curves were plotted after each epoch. These graphs revealed a steady convergence pattern with increasing accuracy and decreasing error, indicating stable and efficient learning throughout the training phase.

Upon deployment, the model successfully classified newly uploaded ultrasound images, accurately identifying the presence or absence of PCOS with high confidence. This underscores the potential of CNNs as reliable tools for automating PCOS detection, particularly in clinical settings where radiological expertise may be limited or subject to variability.

The developed framework not only achieves high diagnostic accuracy but also significantly reduces clinicians' workload by enabling fast, consistent, and objective analysis. These results advocate for the adoption of deep learning-based diagnostic support systems in gynecology, particularly to facilitate early detection and effective treatment planning for Polycystic Ovary Syndrome.

V . CONCLUSION

This paper introduces a method that uses deep learning to automatically diagnose PCOS in ultrasound pictures. A CNN (Convolutional Neural Network) that has been specifically trained on a dataset that has been improved using preprocessing and augmentation approaches is used by the solution. Faster along with more consistent diagnosis results are made possible by the suggested system's automated ultrasound picture categorisation into affected by PCOS and healthy instances, which decreases the need for human analysis and professional radiological interpretation.

During training and validation, the CNN model achieved excellent accuracy in differentiating between infected from non-infected ovarian structures, demonstrating great performance overall. With this skill, patients may be diagnosed early and treated promptly, which might improve their prognosis and slow the onset of PCOS-related problems. Augmentation methods were used to improve generalisability and increase feature variety in order to circumvent problems caused by imbalanced datasets. By improving the speed and precision of PCOS identification, the findings confirm that deep learning, and specifically CNN-based approaches, may be a significant diagnostic tool for radiologists and gynaecologists.

Possible future improvements might include studying hybrid models that blend convolutional neural networks (CNNs) with more conventional machine learning methods, including sophisticated neural architectures (such as ResNet and EfficientNet) for more reliable assessment, and using cross-validation. The model might be more suited for inclusion into real-world healthcare settings if explainable AI (XAI) approaches were to be used to increase its transparency and clinician confidence.

VI . REFERENCES

[1] J. J. Cheng and S. Mahalingaiah, "Data mining and classification of polycystic ovaries in pelvic ultrasound reports," Univ. of Massachusetts Medical Center, Boston, MA, 2021.

[2] X. Zhang et al., "Computational characterization and identification of human polycystic ovary syndrome genes," *Bioinformatics*, vol. 36, no. 7, 2020.

[3] S. Bharati et al., "Diagnosis of Polycystic Ovary Syndrome Using Machine Learning Algorithms," *Int. J. Eng. & Technol.*, vol. 7, no. 4, 2021.

[4] A. A. Nazarudin et al., "Review on Automated Follicle Identification for Polycystic Ovarian Syndrome," *Bull. Electr. Eng. Inf.*, vol. 9, no. 2, pp. 588–593, 2023.

[5] S. J. Pulluparambil and S. Bhat, "Medical Image Processing: Detection and Prediction of PCOS – A Systematic Literature Review," *Int. J. Health Sci. Pharm.*, vol. 5, no. 2, pp. 80–98, 2022.

[6] N. S. Narayan et al., "Automated Detection and Segmentation of Follicles in 3D Ultrasound for Assisted Reproduction," in *Medical Imaging: CAD*, vol. 10575, pp. 431–436, 2023.

[7] A. K. M. S. Hosain et al., "PCONet: A Convolutional Neural Network Architecture to Detect PCOS from Ovarian Ultrasound Images," arXiv, Oct. 2022.

[8] A. Divekar and A. Sonawane, "Leveraging AI for Automatic Classification of PCOS Using Ultrasound Imaging," arXiv, Dec. 2024.

[9] V. Venkatesh et al., "Advancing PCOS Diagnosis: Capsule Network-Based Classification using Ultrasound Images," *J. Innovative Image Processing*, vol. 7, no. 1, pp. 227–229, Mar. 2025.

[10] K. V. Raghavender et al., "Early Stage Detection of PCOS Using Deep Learning," in *ICDSAI*, pp. 215–223, May 2024.