Deep Transfer Learning-Based PCOS Detection From Ovarian Ultrasound Images

Janardhan Neha
MTech, Department of Computer Science & Engineering,
BVRIT HYDERABAD College of Engineering for Women (Autonomous),
Bachupally, Hyderabad - 500090.
janardhanneha5@gmail.com

Dr. M. Shanmuga Sundari
Associate Professor, Department of Computer Science and Engineering,
BVRIT HYDERABAD College of Engineering for Women (Autonomous),
Bachupally, Hyderabad - 500090.
sundari.m@bvrithyderabad.edu.in

Abstract*— Polycystic Ovary Syndrome (PCOS) is a common hormone-related condition that may cause irregular menstrual cycles, infertility, obesity, diabetes, and other metabolic complications. Early and accurate diagnosis is important, but manual interpretation of ovarian ultrasound images is time-consuming and subjective. This study develops an automated PCOS detection framework using deep transfer learning. Two ImageNet-pretrained architectures, EfficientNetB0 and InceptionResNetV2, learn discriminative ovarian features and classify ultrasound scans into PCOS and Normal categories. The workflow includes dataset collection, image-quality screening, resizing, normalization, conversion to numerical arrays, augmentation, model training, evaluation, and deployment. Performance is assessed through accuracy, precision, recall, F1-score, AUC, confusion matrices, and classification reports. EfficientNetB0 achieved the strongest overall results with 96.8% accuracy, 96.5% precision, 96.7% recall, and a 96.6% F1-score, while InceptionResNetV2 obtained 96.6%, 96.1%, 96.5%, and 95.9%, respectively. The best-performing model is integrated into a Streamlit web application that accepts an uploaded ovarian ultrasound image and displays the predicted category with a confidence score. The framework provides a fast, consistent, and accessible computer-aided screening tool that can support healthcare professionals in early PCOS identification while reducing manual workload.

***Keywords*— Polycystic Ovary Syndrome, Deep Transfer Learning, Ovarian Ultrasound, EfficientNetB0, InceptionResNetV2, Image Classification, Streamlit.**

I. INTRODUCTION

Polycystic Ovary Syndrome is a widespread endocrine disorder that commonly affects women during the reproductive years. It is generally associated with irregular menstrual cycles, elevated androgen levels, enlarged ovaries, and the presence of multiple follicles or cyst-like structures visible in ultrasound scans. The condition is clinically important because it influences reproductive function and may also be associated with obesity, insulin resistance, type-2 diabetes, cardiovascular complications, and psychological distress. Early identification enables timely clinical evaluation, lifestyle modification, and treatment planning [22].

Diagnosis usually combines medical history, physical examination, hormonal assessment, and pelvic ultrasound imaging. Among these methods, ultrasound is widely used to identify ovarian morphology and follicular patterns. However, manual interpretation depends strongly on the experience of the radiologist or gynecologist. The process can be slow when patient volume is high, and the same scan may be interpreted differently by different clinicians. Probe orientation, patient body structure, contrast variation, speckle noise, and image resolution can further affect visual assessment [19], [25].

These limitations motivate the development of automated image-analysis systems that can process ultrasound images consistently and rapidly. Traditional computer-aided approaches relied on handcrafted texture, shape, or follicle descriptors followed by classical classifiers. Although useful on controlled datasets, their performance often decreased when image conditions changed. Deep learning reduces this dependence on manual feature engineering by learning hierarchical representations directly from image pixels.

Transfer learning is particularly suitable for medical image analysis where annotated data are limited. A pretrained convolutional network provides general visual filters learned from a large source dataset, while task-specific layers are adapted to the target medical problem. This strategy reduces training time, improves convergence, and can provide strong performance without building a deep network entirely from the beginning [28], [30].

The objective of this work is to design an intelligent PCOS detection system that classifies ovarian ultrasound images into PCOS and Normal categories. The project applies EfficientNetB0 and InceptionResNetV2, compares their performance through multiple evaluation measures, and deploys the superior model through a Streamlit interface. The proposed system is intended to assist early screening and clinical decision support; it is not designed to replace professional diagnosis.

The main contributions are: (1) a complete image-processing pipeline for data collection, quality screening, resizing, normalization, and augmentation; (2) a comparative transfer-learning implementation using two strong pretrained architectures; (3) evaluation using accuracy, precision, recall, F1-score, AUC, confusion matrices, and classification reports; and (4) a web-based prediction interface that returns both a class label and confidence score.

II. RELATED WORK

Research on automated PCOS analysis has progressed from handcrafted image-processing methods to convolutional and transfer-learning systems. Earlier techniques extracted follicular shape, texture, and boundary information using segmentation, SIFT descriptors, particle swarm optimization, competitive neural networks, support vector machines, or artificial neural networks. These approaches demonstrated that ovarian morphology could be quantified computationally, but their dependence on carefully designed features made them sensitive to noise, probe angle, and changes in acquisition conditions [17], [18], [21], [23]–[27].

Convolutional neural networks later enabled automatic feature discovery from ultrasound images. CNN-based methods learn low-level edges and textures in early layers and increasingly complex follicular or ovarian structures in deeper layers. Cahyono et al. applied CNNs to PCO classification, while Boddu et al. and Hosain et al. developed deep architectures for direct PCOS recognition from ultrasound images [9], [12], [20]. These studies supported the transition from manual descriptors to end-to-end image classification.

Recent research has emphasized transfer learning because ovarian ultrasound datasets are usually smaller than general computer-vision datasets. Pretrained networks can be fine-tuned to recognize follicle arrangement, ovarian texture, and cystic patterns with reduced computational effort. Fusion and ensemble methods have also been investigated to combine complementary representations and improve classification stability [1], [7], [8]. Lightweight models and data-balancing strategies further address limited resources and imbalanced medical data [5].

Segmentation-oriented studies attempt to locate follicles or cystic structures before or alongside classification. Feature-fusion attention networks and multilevel-thresholding methods improve sensitivity to small regions, while dedicated follicle networks provide more targeted ovarian analysis [2], [4], [6], [13]. Such localization can make the diagnostic process more interpretable, but it may require additional annotations and more complex training pipelines.

Other work has explored machine-learning prediction from structured clinical data and electronic medical records. These studies demonstrate the value of combining imaging with hormone levels, metabolic indicators, and patient history; however, multimodal information is not always available in image repositories [10], [11], [14]–[16]. For practical ultrasound screening, an image-only system remains useful when it can operate reliably and provide rapid predictions.

The literature identifies several continuing limitations. High-quality annotated ovarian datasets are scarce, many experiments use institution-specific samples, and real-world variability is not always represented. Complex hybrid or transformer-based models may demand high-performance hardware, and only a limited number of systems provide a user-friendly interface for real-time use. Interpretability and clinical validation also remain important concerns. The present framework addresses practical usability by using standardized preprocessing, comparing an efficient model with a deeper residual-inception model, and integrating the selected model into an accessible web application.

III. MATERIALS AND METHODS

The proposed method follows a modular workflow beginning with ovarian ultrasound image collection and ending with real-time prediction. Images are reviewed, organized into two classes, standardized, augmented, and supplied to pretrained convolutional networks. The learned feature representations are passed through custom classification layers, and model performance is measured on previously unseen images. EfficientNetB0 and InceptionResNetV2 are trained under the same evaluation framework so that their results can be compared fairly.

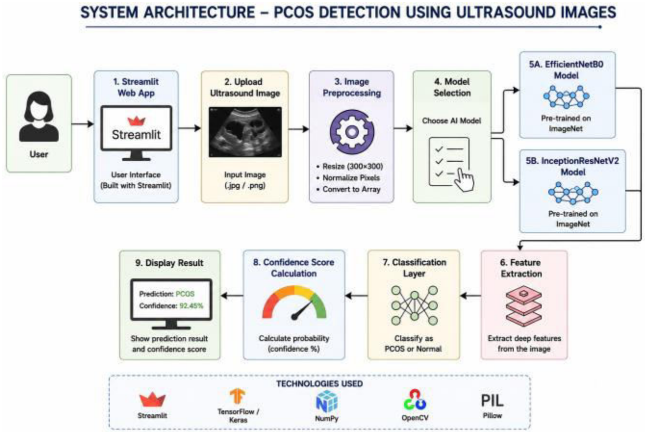

Fig. 1. System architecture of the proposed PCOS detection framework.

Figure 1 presents the complete architecture. The user interacts with the Streamlit web application and uploads an ovarian ultrasound image. The preprocessing module resizes the image, normalizes its pixels, and converts it into the array format expected by the model. The user can select EfficientNetB0 or InceptionResNetV2. The chosen network extracts deep features, the classification layer predicts PCOS or Normal, a confidence score is calculated, and the result is displayed through the interface.

A) Dataset Collection

The dataset consists of ovarian ultrasound images collected from publicly available medical repositories and trusted research datasets. Each image is assigned to one of two categories: PCOS or Normal. Before model development, the collected data undergo an initial quality assessment. Duplicate files, corrupted images, incomplete scans, and severely blurred samples are removed so that only meaningful images contribute to learning.

The cleaned images are stored in class-specific directories and divided into training, validation, and testing subsets. The training subset is used to update the model weights, the validation subset supports learning control and model selection, and the testing subset provides an independent estimate of generalization. This separation prevents evaluation on images already used to optimize the network.

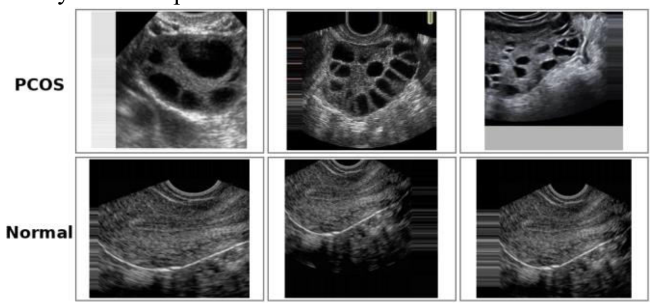

Fig. 2. Representative PCOS and Normal ovarian ultrasound images dataset

B) Visualization

Representative images from both classes are visually inspected before training. This step helps verify label correctness, identify unusable samples, and understand variations in contrast, orientation, ovarian shape, follicle distribution, and ultrasound noise. The PCOS examples show multiple follicular regions and altered ovarian morphology, while the Normal images provide comparison samples for binary classification.

Visualization also supports assessment of class organization and data quality. Because the images originate from different acquisition conditions, they vary in brightness, resolution, and probe orientation. These differences justify the use of uniform resizing, pixel normalization, and augmentation so that the models learn disease-related patterns rather than acquisition-specific artifacts.

C) Pre-processing

Image preprocessing standardizes the raw ultrasound data before feature extraction. Every image is resized to 224 × 224 pixels, matching the configured input resolution of the transfer-learning models. The image is then converted into a multidimensional numerical array and its pixel intensities are rescaled to the range 0–1. Normalization improves numerical stability and helps the optimizer converge efficiently.

The processed images are loaded through batch-wise data generators. The implementation uses a batch size of 32 and trains for up to 20 epochs. Training, validation, and testing generators maintain the class labels and provide the same basic normalization. The test generator does not shuffle samples so that predictions can be aligned correctly with the ground-truth labels when computing confusion matrices and classification reports.

Medical image datasets are often limited, so data augmentation is applied during training to reduce overfitting and improve robustness. The transformations include random rotation, horizontal flipping, zooming, width shifting, and height shifting. These operations produce realistic variations while preserving the clinically meaningful ovarian structures. Validation and testing images are normalized but are not randomly augmented, ensuring consistent performance measurement.

After preprocessing, the images are supplied to the pretrained feature-extraction backbone. The final ImageNet classification layers are removed and replaced with task-specific layers. Global average pooling compresses spatial feature maps into a compact vector, dropout reduces co-adaptation, and dense layers learn the binary decision boundary. A Softmax output layer returns probabilities for PCOS and Normal.

Model training uses the Adam optimizer and categorical cross-entropy loss. During each epoch, forward propagation produces class probabilities, the loss measures the difference between predictions and actual labels, backpropagation calculates gradients, and Adam updates the trainable parameters. Early stopping monitors validation loss and restores the best weights when improvement stops. Model checkpointing saves the strongest version for final evaluation and deployment.

D) Algorithms

EfficientNetB0: EfficientNetB0 is a compact convolutional architecture that balances network depth, width, and image resolution using compound scaling [28]. Rather than increasing one dimension alone, the architecture scales these dimensions together to obtain strong accuracy with comparatively fewer parameters. In the proposed system, the ImageNet-pretrained backbone extracts ovarian texture, follicle arrangement, shape irregularities, cystic regions, and other morphological characteristics. Global average pooling, dropout, and dense layers transform these features into PCOS and Normal probabilities.

The model is suitable for real-time medical image analysis because it provides efficient feature extraction, faster inference, and reduced computational requirements. Its compact design also lowers the risk of overfitting on a limited dataset. The predicted probability is converted into a confidence percentage and presented to the user with the final class label.

InceptionResNetV2: InceptionResNetV2 combines multi-scale Inception modules with residual connections [29]. Parallel convolution branches process visual patterns at different receptive-field sizes, enabling the network to represent both fine follicular details and broader ovarian morphology. Residual connections support gradient flow through the deep architecture and improve training stability. In this implementation, the pretrained backbone is followed by global average pooling and custom dense layers for binary classification.

InceptionResNetV2 offers strong feature-learning capacity and can detect subtle visual differences, but its deeper structure requires more computation than EfficientNetB0. Evaluating both networks under the same preprocessing and testing conditions therefore provides a meaningful comparison between a lightweight compound-scaled model and a deeper multi-scale residual model.

E) Integration of Streamlit Framework

The trained models are integrated into a Streamlit application to provide an interactive deployment environment. The interface includes information about PCOS, a home page, a model-performance dashboard, a prediction page, and a result-analysis page. Users select a trained model, upload a supported image file, and request prediction without interacting with the underlying Python code.

The deployment pipeline applies the same resizing, normalization, and array conversion used during training. The selected network produces class probabilities, the highest probability determines the prediction, and the application displays the model name, class label, and confidence score. The dashboard summarizes accuracy, precision, recall, F1-score, and AUC, while the analysis page presents confusion matrices and comparative graphs.

The interface is designed for research, education, and computer-aided screening. It reduces the need for manual feature engineering and provides results within a short time. A caution message clarifies that the output supports but does not replace professional medical assessment.

The proposed workflow replaces the conventional sequence of ultrasound acquisition, repeated manual inspection, specialist interpretation, and delayed reporting with an automated first-level analysis. After the image is uploaded, the system executes preprocessing, feature extraction, classification, confidence calculation, and result display without requiring the user to define texture or shape descriptors. This reduces variation caused by manual feature selection and enables the same processing sequence to be applied to every image.

Deep feature extraction is performed by the convolutional backbone of each pretrained network. Early layers respond to edges, gradients, and local ultrasound texture, while deeper layers encode more complex patterns related to ovarian boundaries, follicle clustering, cystic regions, and spatial morphology. Freezing the pretrained base during the initial learning stage preserves stable visual filters, and the custom classification head adapts the representation to the two project classes.

Benchmarking is conducted with the same testing data and the same metric definitions for both models. Predictions are converted from class probabilities to class labels and compared with the true labels. Scikit-learn utilities are used to calculate accuracy, precision, recall, F1-score, confusion matrices, and classification reports. AUC values from the project dashboard provide an additional measure of class separability across decision thresholds.

Model selection considers more than the highest accuracy. Precision and recall are examined to verify balanced behavior, the F1-score summarizes the trade-off between them, and the confusion matrix reveals the types of errors. Computational complexity, prediction speed, and deployment suitability are also considered because a clinically useful prototype should provide timely results on commonly available hardware.

The implementation is organized into independent modules for dataset preparation, preprocessing, feature extraction, classification, prediction, and result visualization. This modular arrangement allows each component to be tested separately while preserving a clear end-to-end workflow. An uploaded image passes through the modules in a fixed order, which improves reproducibility and ensures that training and real-time prediction use compatible transformations.

Validation performance is reviewed after every epoch to monitor learning progress and detect overfitting. Early stopping prevents unnecessary training when validation loss no longer improves, while model checkpointing stores the strongest weights instead of only the final epoch. The independent test set is then used once for final metric calculation. This separation between optimization, model selection, and testing supports a more reliable comparison of the two architectures.

The software environment uses Python with TensorFlow and Keras for transfer-learning development, NumPy for array processing, Scikit-learn for evaluation measures, Matplotlib for graphs, and Streamlit for deployment. These tools connect the training notebook, saved model files, metric visualization, and web interface into a single implementation pipeline. The design can operate on commonly available computing hardware while still supporting GPU acceleration when available.

IV. EXPERIMENTAL RESULTS

The two transfer-learning models are evaluated on the independent testing set. Accuracy measures the proportion of correctly classified PCOS and Normal images. Precision measures how many images predicted as PCOS are actually PCOS, recall measures how many actual PCOS images are identified, and the F1-score provides a balanced combination of precision and recall. TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN} \quad (6)$$

Precision: Precision evaluates the fraction of predicted phishing samples that are correctly identified:

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive} \quad (7)$$

Recall: Recall measures the ability of the model to identify the phishing samples that are actually present:

$$Recall = \frac{TP}{TP + FN} \quad (8)$$

F1-Score: F1-Score combines Precision and Recall through their harmonic mean:

$$F1\ Score = 2 \times \frac{Recall \times Precision}{Recall + Precision} \times 100 \quad (9)$$

TABLE I. PERFORMANCE EVALUATION

DL Model	Accuracy	Precision	Recall	F1-Score
EfficientNetB0	0.968	0.965	0.967	0.966
InceptionResNetV2	0.966	0.961	0.965	0.959

Table I shows that both models achieved strong and closely matched classification performance. EfficientNetB0 obtained 0.968 accuracy, 0.965 precision, 0.967 recall, and a 0.966 F1-score. InceptionResNetV2 obtained 0.966 accuracy, 0.961 precision, 0.965 recall, and a 0.959 F1-score. EfficientNetB0 therefore produced the highest value for each of the four reported measures.

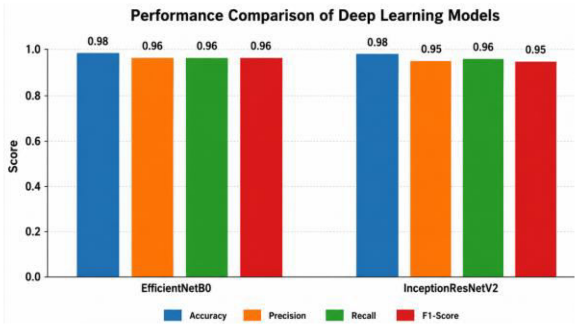

Fig. 3. Combined comparison of accuracy, precision, recall, and F1-score for both models.

Figure 3 combines the four evaluation measures in one grouped bar chart, following the comparison approach used in the sample conference paper. The difference between the networks is small, but EfficientNetB0 remains consistently higher. The project dashboard further reports an AUC of 97.10% for EfficientNetB0 and 96.10% for InceptionResNetV2. AUC summarizes class-separation ability across decision thresholds and supports the selection of EfficientNetB0.

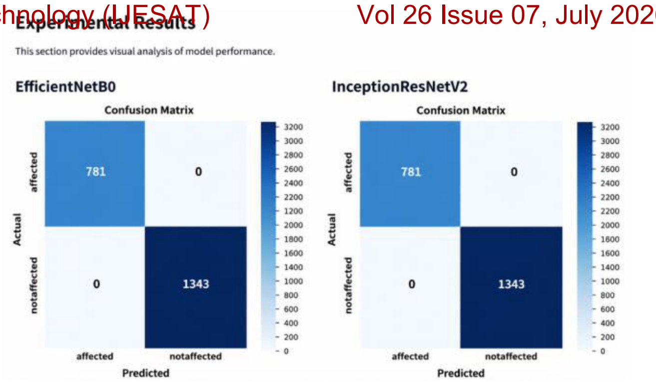

Fig. 4. Confusion matrices for EfficientNetB0 and InceptionResNetV2.

The confusion matrices indicate strong separation between PCOS and Normal images, with only a small number of incorrect classifications. The classification reports likewise show balanced precision and recall. Together, these results indicate that the models did not obtain high accuracy simply by favoring one class. EfficientNetB0 provided the strongest overall balance and lower computational complexity, so it was selected for final deployment.

The results also show that the performance difference between the two networks is modest. InceptionResNetV2 provides strong multi-scale representation, but EfficientNetB0 obtains slightly higher values for every reported classification measure. Its compound-scaled structure is therefore more suitable for the project objective of achieving accurate prediction with efficient computation. The selected model can be loaded more quickly and can return predictions with lower processing overhead in the Streamlit application.

Data preprocessing and augmentation contributed to stable testing performance. Uniform resizing and normalization reduced inconsistency among images, while rotation, flipping, zooming, and shifting exposed the models to realistic variations in orientation and acquisition. Because augmentation was applied only to training data, the reported testing metrics reflect performance on unchanged unseen images rather than artificially transformed evaluation samples.

The experimental evidence supports the use of the framework as a screening aid. High accuracy indicates overall correctness, while closely matched precision and recall reduce concern that the model systematically misses affected images or overpredicts PCOS. Nevertheless, the results are based on the available project dataset and should be interpreted as prototype performance rather than clinical certification. Independent multi-center validation remains necessary before routine medical use.

The metric pattern is internally consistent. EfficientNetB0 has the highest accuracy and also retains high precision, recall, and F1-score, indicating that the improvement is not limited to a single measure. InceptionResNetV2 remains a strong alternative, but its lower F1-score shows a slightly weaker balance between positive predictive reliability and sensitivity. The small numerical gap also emphasizes the importance of using computational efficiency and deployment speed as secondary model-selection criteria.

The reported AUC values reinforce the threshold-based results. EfficientNetB0 reaches 97.10%, while InceptionResNetV2 reaches 96.10%, showing that both networks separate the two classes effectively over a range of possible operating thresholds. In a practical screening interface, the confidence score can help users recognize uncertain cases. Images with low confidence should be referred for careful expert review rather than treated as definitive predictions.

Error analysis remains essential because ultrasound images can differ substantially in contrast, probe angle, patient positioning, and visible follicular structure. Incorrect predictions may arise when image quality is poor or when ovarian appearance is visually similar across classes. The preprocessing pipeline reduces these variations but cannot remove every ambiguity. Therefore, the system output is presented together with a caution statement and is positioned as decision support rather than an autonomous diagnostic conclusion.

Overall, the evaluation confirms that the proposed workflow achieves its principal goals: automated image handling, reliable binary classification, balanced performance, efficient model selection, and direct integration with a user-facing application. These results justify deploying EfficientNetB0 in the prototype while retaining InceptionResNetV2 for comparative analysis and future refinement.

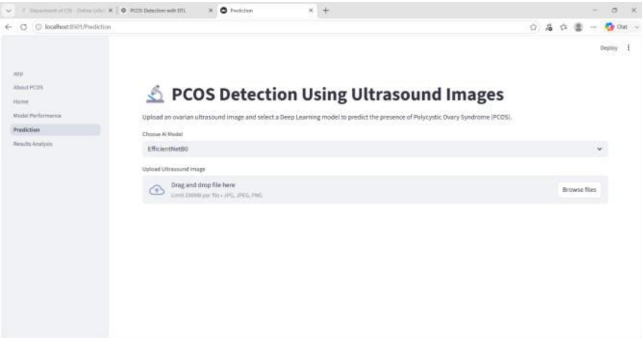

Fig. 5. Streamlit page for model selection and ovarian ultrasound image upload.

The prediction interface allows a user to choose EfficientNetB0 or InceptionResNetV2 and upload an ovarian ultrasound image. The application verifies the file, converts it to the expected input format, and applies the same preprocessing steps used during model development. This maintains consistency between offline evaluation and real-time inference.

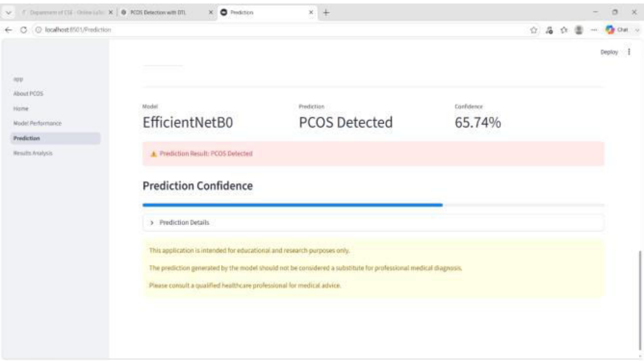

Fig. 6. PCOS prediction generated by EfficientNetB0 with a 65.74% confidence score.

Figure 6 shows a representative prediction generated by the deployed application. The screen displays the selected model, the predicted category "PCOS Detected," and a 65.74% confidence score. A progress indicator and caution message make the output easier to interpret and emphasize that the application is a decision-support prototype.

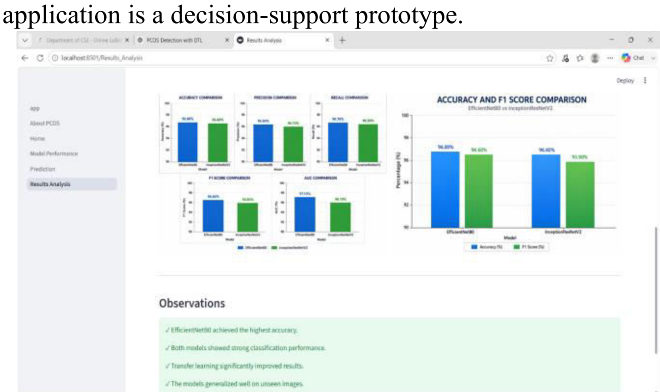

Fig. 7. Result-analysis page showing comparative metric graphs.

The result-analysis page consolidates model performance into confusion matrices and metric visualizations. Presenting both predictive output and experimental evidence in the same application improves transparency and allows users to review why EfficientNetB0 was recommended for deployment. The combined findings demonstrate that transfer learning, standardized preprocessing, and augmentation can provide stable PCOS classification from ovarian ultrasound images.

V. CONCLUSION

This work developed a complete deep transfer-learning framework for automatic PCOS detection from ovarian ultrasound images. The pipeline includes dataset collection, image-quality review, resizing, normalization, numerical conversion, data augmentation, feature extraction, model training, testing, and web deployment. EfficientNetB0 and InceptionResNetV2 automatically learned ovarian texture, follicle distribution, shape irregularities, and cystic patterns without handcrafted features.

Both models achieved high performance, but EfficientNetB0 produced the best overall results with 96.8% accuracy, 96.5% precision, 96.7% recall, a 96.6% F1-score, and a 97.10% AUC. InceptionResNetV2 achieved 96.6% accuracy, 96.1% precision, 96.5% recall, a 95.9% F1-score, and a 96.10% AUC. EfficientNetB0 was selected because it combined slightly stronger predictive performance with lower computational complexity and faster inference.

The trained model was integrated into a Streamlit application that supports model selection, ultrasound image upload, immediate prediction, confidence display, performance review, and result analysis. The system demonstrates how deep learning can reduce manual effort and support consistent early screening. It is intended as a computer-aided support tool and should be used together with clinical evaluation by qualified healthcare professionals.

The present study is limited to binary classification of static ovarian ultrasound images from the available project dataset. It does not combine image predictions with hormone tests, metabolic indicators, patient history, or longitudinal follicle development. Variations between hospitals, ultrasound machines, and patient populations may also affect generalization. These limitations do not reduce the value of the prototype, but they define the validation steps required before broader clinical adoption.

Future work can expand the dataset using multi-center ultrasound images representing diverse patient populations and imaging devices. The framework may be extended to detect additional gynecological conditions, incorporate hormonal profiles and metabolic indicators, and analyze ultrasound video sequences rather than isolated frames. Explainable AI methods such as Grad-CAM, LIME, or SHAP can highlight image regions that influence predictions and improve clinical trust. Further validation against expert interpretation and prospective clinical data would strengthen reliability and practical applicability.

Additional development can optimize the selected network for portable and low-resource environments, including lightweight deployment on clinical workstations or cloud-supported interfaces. Cross-validation on larger annotated cohorts, systematic hyperparameter tuning, and comparison with expert readings can provide stronger evidence of generalization. Privacy-preserving strategies such as federated learning may also support collaborative model improvement across hospitals without directly exchanging sensitive patient images.

A future clinical version may also provide batch processing, secure result storage, confidence-threshold alerts, and standardized reports that can be reviewed alongside laboratory findings and patient history. Such extensions would transform the current prototype into a broader decision-support workflow while preserving the central requirement that final interpretation remains under qualified medical supervision.

REFERENCES

[1] A. Ghosh, K. Srinivasan, et al., “EffiDenseGenOp: Ensemble Transfer Learning With Hyperparameter Tuning Using Genetic Algorithm Optimization for PCOS Detection From Ultrasound Sonography Images,” IEEE Access, 2025.

[2] P. Moral, D. Mustafi, A. Mustafi, and S. K. Sahana, “Feature Fusion Context Attention Gate UNet for Detection of Polycystic Ovary Syndrome,” Scientific Reports, 2025.

[3] “PCOS Detection Using Deep Learning Techniques,” Proc. IEEE Int. Conf. Networking, Multimedia and Information Technology (NMITCON), 2025, doi: 10.1109/NMITCON65824.2025.11188955.

[4] P. Moral, D. Mustafi, A. Mustafi, and S. K. Sahana, “CystNet: An AI-Driven Model for PCOS Detection Using Multilevel Thresholding of Ultrasound Images,” Scientific Reports, vol. 14, 2024.

[5] R. Ahmad, L. A. Maghrabi, I. A. Khaja, et al., “SMOTE-Based Automated PCOS Prediction Using Lightweight Deep Learning Models,” Diagnostics, vol. 14, no. 19, 2024.

[6] S. Sowmiya, S. Umapathy, et al., “F-Net: Follicles Net—An Efficient Tool for the Diagnosis of Polycystic Ovary Syndrome Using Deep Learning Techniques,” PLOS ONE, vol. 19, no. 8, 2024.

[7] P. Chitra, P. Sharmila, et al., “Classification of PCOS in Ultrasound Images Using Deep Learning Methods,” Proc. IEEE 2nd Int. Conf. IoT, Communication and Automation Technology (ICICAT), 2024.

[8] E. Alamoudi et al., “A Deep Learning Fusion Approach to Diagnose Polycystic Ovary Syndrome (PCOS),” Applied Computational Intelligence and Soft Computing, 2023.

[9] A. K. M. S. Hosain, M. H. K. Mehedi, and I. E. Kabir, “PCONet: A Convolutional Neural Network Architecture to Detect Polycystic Ovary Syndrome from Ovarian Ultrasound Images,” Proc. IEEE ICEET, 2022.

[10] A. Makhdoomi, N. Jan, P. Palak, and N. Goel, “Machine Learning Techniques for Medical Images in PCOS,” Proc. 4th Int. Conf. Artificial Intelligence and Speech Technology (AIST), IEEE, 2022.

[11] S. A. Suha and M. N. Islam, “An Extended Machine Learning Technique for Polycystic Ovary Syndrome Detection Using Ovary Ultrasound Image,” Scientific Reports, vol. 12, 2022.

[12] V. Boddu, R. Yalavarthi, and K. Vineesha, “Detection of Polycystic Ovarian Syndrome Using Convolutional Neural Networks,” Int. J. Current Research and Review, vol. 13, no. 6, pp. 155–159, 2021.

[13] B. Rachana et al., “Detection of Polycystic Ovarian Syndrome Using Follicle Recognition Technique,” Global Transitions Proceedings, 2021.

[14] S. Bharati, P. Podder, and M. R. H. Mondal, “Diagnosis of Polycystic Ovary Syndrome Using Machine Learning Algorithms,” Proc. IEEE Region 10 Symposium (TENSYMP), 2020.

[15] K. Maheswari, T. Baranidharan, S. Karthik, and T. Sumathi, “Modelling of F3I Based Feature Selection Approach for PCOS Classification and Prediction,” J. Ambient Intelligence and Humanized Computing, 2020.

[16] J. J. Cheng and S. Mahalingaiah, “Data Mining Polycystic Ovary Morphology in Electronic Medical Record Ultrasound Reports,” Fertility Research and Practice, vol. 5, Art. 13, 2019.

[17] C. Gopalakrishnan and M. Iyapparaja, “Detection of Polycystic Ovary Syndrome from Ultrasound Images Using SIFT Descriptors,” Bonfring Int. J. Software Engineering and Soft Computing, vol. 9, no. 2, pp. 26–30, 2019.

[18] R. M. Dewi and U. N. Wisesty, “Classification of Polycystic Ovary Based on Ultrasound Images Using Competitive Neural Network,” J. Physics: Conference Series, vol. 971, 2018.

[19] L. J. Brattain, B. A. Telfer, M. Dhyani, J. R. Grajo, and A. E. Samir, “Machine Learning for Medical Ultrasound: Status, Methods, and Future Opportunities,” Abdominal Radiology, vol. 43, pp. 786–799, 2018.

[20] B. Cahyono, M. S. Mubarok, and U. N. Wisesty, “An Implementation of Convolutional Neural Network on PCO Classification Based on Ultrasound Image,” Proc. 5th Int. Conf. Information and Communication Technology (ICoICT), IEEE, 2017.

[21] O. R. Isah, A. D. Usman, and A. M. S. Tekanyi, “A Hybrid Model of PSO Algorithm and Artificial Neural Network for Automatic Follicle Classification,” Int. J. Bioautomation, vol. 21, no. 1, 2017.

[22] R. Azziz, E. Carmina, Z. Chen, et al., “Polycystic Ovary Syndrome,” Nature Reviews Disease Primers, vol. 2, Art. 16057, 2016.

[23] B. Purnama, U. N. Wisesty, F. Nhita, A. Gayatri, and T. Mutiah, “A Classification of Polycystic Ovary Syndrome Based on Follicle Detection of Ultrasound Images,” Proc. 3rd Int. Conf. Information and Communication Technology (ICoICT), IEEE, pp. 396–401, 2015.

[24] E. Setiawati and A. B. W. Tjokorda, “Particle Swarm Optimization on Follicles Segmentation to Support PCOS Detection,” Proc. 3rd Int. Conf. Information and Communication Technology (ICoICT), IEEE, 2015.

[25] I. O. Rabiu, A. D. Usman, and A. M. S. Tekanyi, “A Review on Computer Assisted Follicle Detection Techniques and PCOS Diagnostic Systems,” Int. J. Computer Trends and Technology, vol. 28, no. 1, pp. 41–45, 2015.

[26] B. Purnama, U. N. Wisesty, F. Nhita, A. Gayatri, and T. Mutiah, “A Classification of Polycystic Ovary Syndrome Based on Follicle Detection of Ultrasound Images,” Proc. 3rd Int. Conf. Information and Communication Technology (ICoICT), IEEE, pp. 396–401, 2015.

[27] S. S. Deshpande and A. Wakankar, “Automated Detection of Polycystic Ovarian Syndrome Using Follicle Recognition,” Proc. IEEE Int. Conf. Advanced Communications, Control and Computing Technologies (ICACCCT), pp. 1341–1346, 2014.

[28] M. Tan and Q. V. Le, “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,” Proc. Int. Conf. Machine Learning, pp. 6105–6114, 2019.

[29] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, “Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning,” Proc. AAAI Conf. Artificial Intelligence, vol. 31, no. 1, 2017, doi: 10.1609/aaai.v31i1.11231.

[30] S. J. Pan and Q. Yang, “A Survey on Transfer Learning,” IEEE Trans. Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2010, doi: 10.1109/TKDE.2009.191.