

Original Article

Explainable Deep Tabular Learning with SHAP-Guided Feature Optimization and Squeeze-Excitation Enhanced TabNet for Polycystic Ovary Syndrome Classification

V. Lakshmi^{1*}, B. Pushpa²

^{1,2}Department of Computer and Information Science, Annamalai University, Annamalainagar, Tamilnadu, India.

*Corresponding Author : lakshmiphda@gmail.com

Received: 02 March 2026

Revised: 17 June 2026

Accepted: 24 June 2026

Published: 29 August 2026

Abstract - Polycystic Ovary Syndrome (PCOS) is one of the most prevalent reproductive endocrinopathies in women of reproductive age. This requires diagnostic systems that reliable and interpretable. This paper presents a deep tabular framework that like to maintain efficiency in the classification of Polycystic Ovarian Syndrome (PCOS) through effective data pre-processing, synthetic balancing of data, and optimizing for significant features. Initially, the clinical dataset is preprocessed by iterative imputation, outlier handling, feature normalization and synthetic balancing for statistical consistency and robust model learning. The study uses Conditional Tabular GAN (CTGAN) to augment the sample diversity and address class imbalance by generating additional minority samples. Moreover, a Transformer-based feature extractor captures associations between features to encode better reference feature embeddings. Then, performing Shapley Additive Explanations (SHAP)-guided feature selection using an XGBoost explainer to determine what predictors are most important in the model performance, hence improving both interpretability and computational efficiency. The TabNet-SE ahead of running the baseline model, i.e. Random Forest (RF) and XGBoost, in addition both conducting of these steps with the complete feature set and also SHAP-selected features. The experimental results show that TabNet-SE consistently outperforms all baseline models in terms of accuracy, precision, recall and F1-score when trained on the set of features optimized with SHAP. Model performance was further corroborated with confusion matrix, Receiver Operating Characteristic (ROC) curve, and other evaluation metrics. This framework not only facilitates accurate diagnosis of PCOS but also enhances interpretability, thus bridging the gap between clinical requirements and model transparency.

Keywords - Conditional Tabular GAN (CTGAN), Explainable Deep Learning, Medical data classification, Polycystic Ovary Syndrome (PCOS), SHAP-based Feature Selection, Squeeze-Excitation Network, TabNet-SE, Transformer feature extraction.

1. Introduction

PCOS is a complex hormonal and metabolic disorder that affects many women of reproductive age. Clinically, it is expressed by different features including oligomenorrhea/amenorrhea (irregular cycles), hyperandrogenism (increased levels of androgen hormones), insulin resistance, and follicular morphology resembling that of polycystic ovaries. This diversity makes it difficult to diagnose and treat rapidly [1]. Standard diagnostic methods using Rotterdam or NIH criteria often rely on limited biochemical and ultrasound data. This leads doctors to make alternative decisions and delay treatment for patients [2]. With the development of AI and deep learning to deal with clinical data, new great methodologies can be used for scanned PCOS detection. Models trained with PCOS-specific data have shown the ability to accurately assess risk of developing PCOS by providing demographic, biochemical and hormonal characteristics [3].

The data is very flawed, riddled with inconsistencies and a considerable amount of missing information. There's also a large class imbalance between healthy and affected, i.e. sick people. Any medical learning process requires verification of data preparation. Iterations of the missing value using model-based methods helps retain complex relationships between features while filtering out noisy or wrong data [4]. Iterative estimators help us maintain the relationships between multiple variables better than just using a basic mean or median imputation. This is essential for detecting hormonal and metabolic patterns.

Furthermore, restricting the outer limits of our lab results to the 1st and 99th Percentile reduces the influence very high or low values may have on our results but also ensures that valid clinical samples were retained. Class imbalance is a common problem in diagnostic modeling. In small clinical datasets, existing oversampling or SMOTE techniques can



distort the relationships between features. In contrast, CTGAN generates realistic samples of the minority class by preserving the underlying conditional distributions. Research shows that the CTGAN data makes for better recall and calibration of diseases including PCOS [5]. Image feature representation plays a vital role in classification model performances. Transformer-based encoders with attention mechanisms can capture complex relationships between clinical indicators. Attention-guided feature extraction allows us the flexibility to consider these interactions as biochemical, physical and menstrual factors are not linearly separable.

TabNet-style attention techniques and lightweight transformer encoders achieve quite a high level of accuracy and interpretability on several structured biomedical datasets [6, 7]. Explainability is a huge issue with respect to medicine. SHAP values quantify how each feature contributes to an individual prediction, allowing for deeper insights at the global and local levels. By using SHAP, you can select the most prominent variables to make feature selections. Selecting features that jointly account for 90-95% of the total SHAP importance prevents redundancy and improves interpretability and efficiency of downstream deep learning models [8].

Some deep tabular learners (such as TabNet) use sequential attention masks that encourage certain features to be weighted more heavily for each prediction. The SE gating mechanism helps in scaling the feature responses and puts more emphasis on critical diagnostic features. The SE-augmented TabNet provides intuitive feature weights that correlate with the SHAP importance rankings and achieves competitive performance against tree-based models such as Random Forest and XGBoost.

Accuracy is just one of many evaluation metrics. Discrimination ability was assessed by using confusion matrices, Receiver Operating Characteristic curves, and area under the curve. Calibration plots and Brier scores evaluate probability accuracy whereas DeLong's test statistically validates differences in receiver operating characteristic Area Under the Curve (AUC) between models. The addition of Monte-Carlo dropout variance gives more uncertainty estimates, which have implications for clinical risk assessment.

In spite of more advancement in the prediction of PCOS by machine learning and deep learning models still, there are various limitations. The majority of methods use typical machine learning classifiers or separate deep learning architectures, lack the combination of explainable feature optimization and adaptive attention based tabular learning in a single framework. Moreover, a lot of studies do not well deal with class imbalance, feature redundancy and interpretability, limiting their usability on real world clinical decision support. Classifications by presently deployed models also offer little direct information about the feature-level reasoning for

predictions which broadens opacity in sensitive health care applications. To overcome these shortcomings, the current work proposes an explainable deep tabular learning framework that integrates CTGAN-based synthetic balancing of underrepresented phenotypes; Transformer-based feature extraction linking context to phenotype; SHAP analysis for improving discriminatory features across multi-phenotype datasets; and a Squeeze-Excitation boosted TabNet classifier enabling accurate PCOS diagnosis in an interpretable manner. This proposed framework is a robust alternative to improving discriminative learning while maintaining model transparency for a trustworthy and clinically interpretable specialized diagnostic system (PCOS classification).

The main contributions of this work were summarized as follows: (i) robust pre-processing and balancing framework for minimizing information loss in incomplete and imbalanced clinical PCOS data using a CTGAN, (ii) contextual clinical representation using a Transformer-based feature encoder, (iii) integration of SHAP-guided XGBoost feature optimization for interpretable feature selection, and (iv) TabNet-SE architecture design that considers Squeeze-Excitation attention to facilitate adaptive deep tabular learning. The experimental results show that the proposed TabNet-SE model using SHAP-based selected features outperforms the baseline Random Forest and XGBoost models in terms of classification accuracy.

2. Related Works

Advanced technologies in the machine learning and deep learning area have recently achieved remarkable performance on automated Polycystic Ovary Syndrome (PCOS) diagnosis based on clinical and physiological data. Agirsoy et al. [9] introduced a new non-invasive machine learning framework for PCOS classification based on clinical attributes and showed that data-driven prediction models could assist in early diagnosis.

Likewise, Shanmugavadeivel et al. [10] introduced an improved predictive analytics and classification framework by making use of clinical data as well as other feature optimization strategies for reliable prediction in PCOS identification. Yadav et al. [11] performed comparative analysis of various machine learning algorithms for PCOS diagnosis and concluded that ensemble methods invariably have better accuracy than conventional ones when it comes to modeling these complex clinical patterns. Oviya Graselina et al. [12] expanded on a machine learning-based PCOS classification system to emphasize the importance of consolidated preprocessing and feature engineering towards the advancement of diagnostic efficacy. In particular, Rahman et al. [13] proposed an online machine learning framework for early prediction of PCOS, indicating the real-world practicality and application potential of intelligent clinical decision-support systems in healthcare environments.

To enhance the explainability and interpretability of PCOS in medical prediction systems, attention has been focused on some studies that use Explainable Artificial Intelligence (XAI) methods to diagnose PCOS. An optimized feature selection and explainable AI based machine learning framework for detection of PCOS was proposed by Elmannai et al. [14], highlighting the significance of interpretable decisions made during deployment in clinical practice. Similarly, Pooja et al. [15] built an interpretable ensemble learning framework based on combined feature selection methods that improved interpretability of model predictions and increased its utility for diagnostics. Chen et al. [16] performed a feature-driven machine learning study for biomarker discovery and immune cell infiltration analysis in PCOS, highlighting the importance of integrating multi-omics data to develop a disease progression paradigm, improve diagnosis and understand treatment. Emanuel et al. [17] applied machine learning methodologies to analyze the treatment sentiment of three PCOS-related online forums, showing that AI techniques for extracting clinical implications from unstructured human-generated textual data can be adapted to clinical medicine at large.

Class imbalance and data scarcity in the medical field continue to present significant challenges in machine learning systems for healthcare. Synthetic tabular data generation approaches have received significant attention in the recent years to mitigate these restrictions. Papadaki et al. [18] conducted a recent review on synthetic tabular data generation, and emphasized the importance of better generalization and model robustness with enhanced privacy protection. Pezoulas et al. [19] carried out an extensive literature review of synthetic healthcare data generation techniques, and evaluated their utility at addressing the challenges associated with small and imbalanced medical datasets. The inclusion of synthetic healthcare tabular data can improve model performance while ensuring diversity during training and reducing overfitting [20].

For medical tabular data generation, Hernandez et al. [21] proposed an elaborate evaluation methodology and emphasized that distributional fidelity and clinical consistency is imperative. In 2024, Alqulaity and Yang [22] introduced a modified conditional GAN model for generating high quality synthetic healthcare tabular data with reported well-representation of the minority class samples. In addition, Kang et al. [23] proposed a Tabular Transformer GAN framework that directly addresses the task of heterogeneous healthcare data distribution modelling with outcomes of exposure effect-focused generative modeling for medical tabular analytics. Specialized deep learning architectures for structured tabular learning have also demonstrated comparable performance to state-of-the-art methods on some healthcare prediction tasks. A hybrid CNN-GRU-TabNet framework for mortality prediction was proposed by Fahim et al. [24] illustrating the

feasibility of a TabNet-based attentive learning mechanism to enhance both feature interpretability and classification performance. Lim et al. [25] applied machine learning methods based on radial pulse wave characteristics to PCOS diagnosis and suggested that representations of physiological features can serve as promising non-invasive indicators for disease diagnosis. In [26], Panjwani et al. proposed an optimized machine learning framework for early PCOS detection, and highlighted the significance of rigorous preprocessing and optimized feature selection in increasing predictive accuracy.

Existing studies have shown significant improvement in PCOS diagnosis using machine learning and deep learning techniques; however, drawbacks are still prevalent. However, most of the existing frameworks rely on traditional classifiers and do not provide an integrated framework that is connected with various data balancing paradigms, contextual feature extraction, explainable feature selection, as well as tabular attention mechanisms in a single plugged-in manner. Moreover, incorporating deep tabular learning architectures with explainable AI has received limited attention for clinically interpretable PCOS diagnosis. Motivated by these limitations, this work presents an explainable deep tabular learning framework that incorporates CTGAN-based class balancing, Transformer-based contextual feature extraction as well as SHAP guided optimization of the predictive features and a Squeeze-Excitation auto context-designed TabNet model aiming to provide robust, interpretable, and accurate classification of PCOS.

3. Proposed Methodology

The proposed classification methodology for PCOS involves the integration of preprocessing, synthetic feature generation, explainable feature optimization, and deep tabular learning for a more robust and interpretable diagnostic approach. Clinical data first pass through a preprocessing pipeline for normalization and consistency. CTGAN offers a solution to any imbalance in split classes by synthesizing realistic samples. These features are then passed through a Transformer-based encoder that accounts for inter-feature dependencies to extract deep contextual features. Following this, features that could be learned better, while still maintaining distinct-interpretability (v-NC) properties, are preferred according to an XGBoost explainer through SHAP-based feature selection. For the classification setting, we present a new Squeeze-Excitation enhanced TabNet (TabNet-SE) model that incorporates key features of channel-wise attention and sparse feature selection to accentuate informative patterns while reducing redundant information. Accuracy, precision, recall, F1-score, AUC metrics, confusion matrix and ROC analyses are used to examine the efficiency and reliability of the framework. These tests demonstrate its robust diagnostic performance and interpretability. Figure 1 depicts the overall workflow of our approach.




Fig. 1 The overall workflow of the proposed model

3.1. Pre-Processing Pipeline

The initial component of the Explainable Deep Tabular Learning for PCOS classification is the pre-processing pipeline which prepares clinical data. As such, it guarantees that the data is uniform, debiased and statistically normalized to facilitate strong model learning. The first is to cover five steps: (1) cleaning and encoding, (2) imputing missing values based on behavior of a suitable model trained on data with the same features, (3) detection of outliers and winsorization, (4) scaling each feature to normalize it and parenthesis (5), balancing by generating synthetic data using CTGAN. Below, we outline and explain each step.

3.1.1. Data Cleaning and Categorical Encoding

Given a raw dataset

$$D = \{(x_i, y_i) | x_i \in \mathbb{R}^n, y_i \in \{0,1\}\}, i = 1, 2, \dots, N \quad (1)$$

Where x_i represents a feature vector of dimension n and y_i denotes the binary label (Normal = 0, PCOS = 1), the preprocessing begins with the removal of redundant and missing structural data. Columns with non-informative identifiers like "Unnamed" fields are Removed.

Categorical variables $x_{ij} \in \mathcal{C}$ are transformed into numerical form using Label Encoding, where each unique category $c_k \in \mathcal{C}$ is mapped to an integer:

$$x'_{ij} = LE(x_{ij}) = k, \forall x_{ij} = c_k, k \in \{0, 1, \dots, K-1\} \quad (2)$$

This ensures that the categorical features can be mathematically processed by further imputation and learning models.

3.1.2. Model-Based Missing Value Imputation

If the features in the datasets have missing values, it can cause bias in the model training and lead to poor generalization. Assume that the dataset is divided as follows:

$$X = [x_1, x_2, \dots, x_m]^T \in \mathbb{R}^{N \times m} \quad (3)$$

where some entries are missing, i.e., $x_{ij} = NaN$.

The method uses a model-based iterative imputation strategy with Extra Trees Regressor for recovering the missing elements.

This technique is a random forest of decision trees. The imputation process estimates each feature x_j as a function of all other features X_{-j} :

$$\hat{x}_j = f_j(X_{-j}) + \epsilon_j \quad (4)$$

Where f_j is a regression model learned iteratively and ϵ_j denotes residual noise.

At each iteration t , the imputed value is updated as:

$$x_j^{(t)} = \begin{cases} x^{(t-1)}, & \text{if not missing} \\ f_j^{(t-1)}, & \text{if missing} \end{cases} \quad (5)$$

until convergence or until the maximum iteration count T is reached.

This framework is more stable and less biased than mean or median substitution options. It keeps the relationships of several features.

3.1.3. Outlier Detection and Winsorization

Isolation Forest Random partitioning is an unsupervised anomaly detection algorithm. It detects trolls, or data that are markedly different from the general spread of samples.

Each sample x_i receives an anomaly score defined as:

$$s(x_i, n) = 2^{-\frac{E(h(x_i))}{c(n)}} \quad (6)$$

Where $E(h(x_i))$ denotes the expected path length of x_i in the isolation tree ensemble, $c(n)$ is the average path length of a successful search in a binary tree, and n is the total number of samples. Samples with $s(x_i)$ close to 1 are considered outliers.

To keep the data size, the detected outliers are not directly removed but instead Winsorization is applied to preserve extreme values within specified percentiles:

$$x'_{ij} = \begin{cases} P_1(x_j), & x_{ij} < P_1(x_j) \\ P_{99}(x_j), & x_{ij} > P_{99}(x_j) \\ x_{ij}, & \text{Otherwise} \end{cases} \quad (7)$$

Where $P_1(x_j)$ and $P_{99}(x_j)$ denote the 1st and 99th percentiles of feature x_j . This operation mitigates the impact of extreme outliers on the dataset while preserving its statistical soundness.

3.1.4. Normalization of Numerical Features

The normalization of features plays a significant role in stabilizing the optimization process during deep learning. A Z-score normalization (Standard Scaling) is applied to each numerical feature x_j :

$$x'_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j} \quad (8)$$

Where μ_j and σ_j represent the mean and standard deviation of the j^{th} feature:

$$\mu_j = \frac{1}{N} \sum_{i=1}^N x_{ij}, \quad \sigma_j = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_{ij} - \mu_j)^2} \quad (9)$$

This helps ensure that all of the features are centered around zero and have unit variance, which helps propagate gradients evenly when training the model.

3.1.5. Synthetic Data Balancing using Conditional Tabular GAN (CTGAN)

Most PCOS datasets are imbalanced. They have considerably more normal samples than cases with PCOS. CTGAN generates realistic synthetic minority samples in order to overcome this problem.

While traditional oversampling techniques (such as SMOTE) treat each row in a tabular dataset independently, CTGAN uses adversarial training to model complex nonlinear relationships within the rows of your tabular data.

The CTGAN framework consists of a Generator (G) and a Discriminator (D), both optimized through a minimax game:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x|c)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z, c)|c))] \quad (10)$$

Where z is a random noise vector, c represents the conditioning label (class), and p_{data} is the real data distribution.

The generator learns to produce synthetic samples $G(z, c)$ that resemble real minority-class data $x_{minority}$, improving the class balance:

$$D_{balanced} = D_{majority} \cup G(z, c = 1) \quad (11)$$

This enables the output conditionally generated synthetic data to retain class semantics and feature interdependence.

Transformation of Varied, Incomplete and Imbalanced Clinical PCOS Data into Statistical Reliable Feature Space using the preprocessing pipeline with the help of model-based imputation, outlier winsorization, scaling and CTGAN-based balancing, pipeline ensures to have the data ready for downstream deep tabular learning.

Not only these steps help to improve the interpretability by reducing bias cause by conventional mathematical reasoning but they also serve as a good starting point for TabNet-SE classification stage to achieve better classification performance and transparency. Figure 2 depicts the steps in preprocessing pipeline.




Fig. 2 Preprocessing pipeline flow diagram

3.2. Feature Extraction Process

The proposed framework is based on a feature extraction mechanism that employs the Transformer model to extract contextually rich clinical features from preprocessed tabular data. In contrast to feature selection or PCA-based linear transformation, the attention transformative module captures inter-feature dependencies and contextual correlations with self-attention and multi-head transformations. It creates expressive embeddings to enhance model interpretability and discriminatory power.

3.2.1. Input Representation

Let the preprocessed clinical dataset be represented as

$$X = [x_1, x_2, \dots, x_N]^T \in \mathbb{R}^{N \times d} \quad (12)$$

Where N denotes the number of patient samples and d denotes the number of clinical features (e.g., hormonal levels, BMI, insulin, follicle count). Each feature vector $x_i = [x_{i1}, x_{i2}, \dots, x_{id}]$ corresponds to a normalized and imputed data instance.

A learnable linear projection is applied to project each feature into a high-dimensional latent space such that the model learns nonlinear feature relations:

$$E = XW_e + b_e \quad (13)$$

Where $W_e \in \mathbb{R}^{d \times d_e}$ and $b_e \in \mathbb{R}^{d_e}$. Here, d_e is the embedding dimension, typically chosen as 64 or 128. By integrating feature values through this projection, it guarantees that the subsequent Transformer encoding can be performed within a continuous and differentiable representation.

3.2.2. Transformer Encoder for Contextual Feature Learning

The feature extraction consists of N layers based on the original Transformer Encoder. Its self-attention mechanism captures dependencies over long ranges of features. The encoder is made up of L layers, with each layer containing a

Multi-Head Self-Attention and Feed-Forward Network, along with residual normalization.

Multi-Head Self-Attention (MHSA)

The self-attention mechanism computes relevance between every pair of features from the input embedding matrix E by projecting it onto three different matrices:

$$Q = EW_Q, \quad K = EW_K, \quad V = EW_V \quad (14)$$

Where $W_Q, W_K, W_V \in \mathbb{R}^{d_e \times d_k}$ are learnable projection matrices corresponding to queries, keys, and values, respectively.

The attention score between two feature tokens i and j is given by:

$$Attention(i, j) = \frac{\exp\left(\frac{Q_i \cdot K_j^T}{\sqrt{d_k}}\right)}{\sum_{m=1}^d \exp\left(\frac{Q_i \cdot K_m^T}{\sqrt{d_k}}\right)} \quad (15)$$

which measures the importance of feature j when encoding feature i .

The attention output for the i^{th} feature is computed as:

$$Z_i = \sum_{j=1}^d Attention(i, j) V_j \quad (16)$$

For multi-head attention, h attention heads are computed in parallel, projected independently:

$$MHSA(E) = Concat(Z_1, Z_2, \dots, Z_h) W_O \quad (17)$$

Where $W_O \in \mathbb{R}^{hd_k \times d_e}$ aggregates the outputs across heads.

Such an attention mechanism enables the model to leverage different features of these feature relations. This

involves the hormonal correlations and metabolic dependencies for PCOS diagnosis.

3.2.3. Feed-Forward Network (FFN) and Residual Normalization

Following the MHSA layer, patterned features are fed through a two-layer Feed-Forward Network:

$$FFN(Z) = \text{ReLU}(ZW_1 + b_1)W_2 + b_2 \quad (18)$$

Where $W_1 \in \mathbb{R}^{d_e \times d_f}$, $W_2 \in \mathbb{R}^{d_f \times d_e}$, and $d_f > d_e$.

Residual connections and layer normalization are both used by the model in order to stabilize the learning process and prevent vanishing gradients:

$$H^{(l)} = \text{LayerNorm}\left(E^{(l-1)} + \text{MHSA}(E^{(l-1)})\right) \quad (19)$$

$$E^{(l)} = \text{LayerNorm}\left(H^{(l)} + \text{FFN}(H^{(l)})\right) \quad (20)$$

$E^{(l)}$, the output of the last Transformer layer, are embeddings that have been refined in the sense that each feature has access to contextual information from all other features in your dataset.

3.2.4. Contextual Feature Pooling

In order to have a compact and discriminative representation oriented for classification, global average pooling is used across the contextualized embeddings:

$$F = \frac{1}{d} \sum_{i=1}^d E_i^{(l)} \quad (21)$$

Where $F \in \mathbb{R}^{d_e}$ represents the contextual feature vector of a patient.

This pooling approach aggregates information across all clinical dimensions. It minimizes duplication of data and upholds essential correlations.

3.2.5. Explainable Feature Attribution

The model provides interpretable analysis in the feature extraction to foster comprehensibility of the proposed framework. It contains attention visualization and SHAP-based attribution. The Transformer-based encoder inherently provides attention weights (α_{ij}), which quantify how strongly feature j influences the contextual representation of feature i . These weights are computed as:

$$\alpha_{ij} = \frac{\exp\left(\frac{Q_i \cdot K_j}{\sqrt{d_k}}\right)}{\sum_{m=1}^d \exp\left(\frac{Q_i \cdot K_m}{\sqrt{d_k}}\right)} \quad (22)$$

Where Q_i and K_j represent the query and key vectors of features i and j , respectively, and d_k is the scaling dimension. The attention coefficients from the model provide importance scores indicating how informative each clinical attribute (LH, FSH, AMH, insulin and BMI) is in determining outcomes for PCOS.

Figure 3 shows the average attention weights across all samples, which emphasizes the primary clinical predictors assembled by the model. The high attention values indicate the implicit prioritization that is done by the Transformer encoder; i.e., which features play a more important role in model decisions.




Fig. 3 Average attention weight distribution across clinical features

This work proposes SHAP, as a complement to attention-based interpretability for semantic insight in model behavior at local and global scale. SHAP value of a given feature representing its individual contribution in the model comes from cooperative game theory rules which evaluate average marginal contribution it directly makes to model outputs with regards to its complexity.

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N|-|S|-1)!}{|N|!} [f(S \cup \{i\}) - f(S)] \quad (23)$$

Where N represents the set of all features, and $f(S)$ denotes the model prediction using the subset S . A higher $|\phi_i|$ indicates a greater influence of feature i on the classification outcome.

Figure 4: visualizes the average contribution of each feature in affecting a given model prediction using the SHAP summary plot. This allows for clarification of the order of clinical determinates in PCOS classification.




Fig. 4 SHAP-based interpretability analysis showing the contribution and importance ranking of extracted clinical features for PCOS classification

The combination of the attention map (Figure 3) and SHAP summary (Figure 4) forms a dual interpretability framework. They use attention-based and a SHAP for the model based and feature centered explanation. This combination ensures that the proposed feature extraction process increases classification accuracy while maintaining interpretability of the clinical reasoning for both health practitioners and researchers.

Moreover, the SHAP analysis enhances clinical interpretability by clearly signaling which specific hormonal and metabolic attributes played significant roles in predicting PCOS at feature-level for transparent decision making.

3.3. Classification Process

The classification phase of the proposed PCOS diagnostic framework provides a mechanism to convert the extracted and enhanced tabular features into a discernable and accurate prediction of disease status. Concerning the classification, there are two main nodes of experimentation (a) the whole set of extracted features and (b) SHAP based selected feature subset. To evaluate how well feature optimization has improved performance, next it compares baseline to optimized performance. Three models have been evaluated, which include the proposed Squeeze-Excitation enhanced TabNet (TabNet-SE) model and two traditional baselines in Random

Forest (RF) and XGBoost (XGB). Which the entire process can be described mathematically and orderly as such.

3.3.1. SHAP-Based Feature Selection

The model makes use of a SHAP-based XGBoost Explainer to eliminate redundancy and retain only the most discriminative features. Let f_{XGB} denote the trained XGBoost model. For each input x_i , the SHAP framework decomposes the model prediction into feature contributions:

$$f_{XGB}(x_i) = \phi_0 + \sum_{j=1}^d \phi_j^{(i)} \quad (24)$$

Where $\phi_j^{(i)}$ represents the SHAP value of feature j for sample i , and ϕ_0 is the model bias.

Mean absolute SHAP importance for every feature over all samples:

$$I_j = \frac{1}{N} \sum_{i=1}^N |\phi_j^{(i)}| \quad (25)$$

Features are subsequently ordered in a descending manner with respect to I_j and the subset of features covering 95% cumulative importance selected as:

$$F_{selected} = \left\{ j: \frac{\sum_{k=1}^j I_k}{\sum_{k=1}^d I_k} \leq 0.95 \right\} \quad (26)$$

Thus, two feature matrices are defined:

- Full feature matrix $X_{all} \in \mathbb{R}^{N \times d}$
- Selected feature matrix $X_{sel} \in \mathbb{R}^{N \times d'}$, where $d' < d$

Thus, there were 64 features extracted from the preprocessed data. Among these, 47 features were selected using SHAP-based XGBoost Explainer. These selected features divided according to their metrics for predicting PCOS are depicted in Figure 5.




Fig. 5 Ranking of key features selected by SHAP-Based XGBoost explainer

The model framework is not an independence assumption since both the Transformer based encoder and TabNet-SE architecture are built to learn inter-feature relationships in a contextual, attention-driven fashion. In addition, to avoid data leak when optimizing features, restrict SHAP based feature importance estimation and feature selection only to the training subset after dataset splitting, meaning that the test data remains fully unseen during any stage of feature optimization and model building.

3.3.2. Proposed TabNet-SE Model

The TabNet-SE model is based on the original TabNet architecture, with an additional Squeeze-Excitation (SE) block. The addition allows for feature recalibration and increases interpretability.

Forward Propagation

Given input $x \in \mathbb{R}^{d'}$:

$$h_1 = \phi_1(W_1x + b_1), \quad h_2 = \phi_2(W_2h_1 + b_2) \quad (27)$$

Where ϕ_1 and ϕ_2 denote nonlinear activation functions (GELU), and W_i, b_i are trainable parameters.

Squeeze-Excitation Attention

The SE mechanism first computes a global descriptor:

$$z = \frac{1}{C} \sum_{i=1}^C h_{2,i} \quad (28)$$

which is followed by two fully connected layers to generate feature-wise attention weights:

$$s = \sigma(W_2\delta(W_1z)) \quad (29)$$

Where δ is ReLU, and σ is the sigmoid function.

The recalibrated feature output is then:

$$\tilde{h}_2 = s \odot h_2 \quad (30)$$

The adaptive weight enables the model to amplify the features most significant for PCOS forecast.

Prediction Layer

The last few layers do a transformation (often linear, often with non-linearity) and give probability estimates:

$$h_3 = \phi_3(W_3\tilde{h}_2 + b_3) \quad (31)$$

$$\hat{y} = \sigma(W_4h_3 + b_4) \quad (32)$$

Where $\hat{y} \in [0,1]$ represents the predicted probability of PCOS. The architecture of the proposed TabNet-SE model is depicted in Figure 6. It shows that the powerful feature transformation is performed sequentially, which is followed by attention-based Squeeze-Excitation block and decision refinement layers for efficient PCOS classification.




Fig. 6 Architecture of the proposed TabNet-SE model

3.3.3. Baseline Models

Hence, two baseline models of Random Forest and XGBoost were used to compare analysis:

Random Forest (RF)

An ensemble of T decision trees $\{h_t\}_{t=1}^T$ produces:

$$\hat{y}_{RF} = \frac{1}{T} \sum_{t=1}^T h_t(x) \quad (33)$$

Extreme Gradient Boosting (XGBoost)

Constructs additive weak learners f_t iteratively

$$\hat{y}_{XGB}^{(t)} = \hat{y}_{XGB}^{(t-1)} + \eta f_t(x) \quad (34)$$

optimizing a regularized objective:

$$\mathcal{L}^{(t)} = \sum_i l(y_i, \hat{y}_i^{(t)}) + \Omega(f_t) \quad (35)$$

With $\Omega(f_t) = \gamma T + \frac{1}{2} \sum \|w_t\|^2$ controlling model complexity.

3.3.4. Training and Optimization

The TabNet-SE model minimizes Binary Cross-Entropy (BCE) loss:

$$\mathcal{L}_{BCE} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (36)$$

with AdamW optimizer for regularized convergence:

$$\theta_{t+1} = \theta_t - \eta \frac{\nabla_{\theta_t} \mathcal{L}}{\sqrt{v_t + \epsilon}} + \lambda \theta_t \quad (37)$$

The model implements early stopping, which occurs when the validation AUC fails to increase over a number of epochs.

In the classification phase, the model has shown that SHAP guided selection enhances interpretability of features. It also shows how the introduction of adaptive attention to deep tabular reasoning through TabNet-SE enables it to outperform traditional ensemble learners.

Finally, dual evaluation in both feature spaces is applied to confirm the stability and explainability of the proposed diagnostic model for PCOS. The below algorithm describes the complete proposed framework.

Algorithm: Explainable Deep Tabular Learning for PCOS Classification (TabNet-SE)
Input: Preprocessed clinical dataset (CSV), labels, feature selection params, model hyperparameters, output directories
Output: Trained TabNet-SE, RF, XGBoost models; Selected features; Evaluation plots and metrics
1. Data Preparation and Class Balancing a. Load preprocessed dataset b. Verify label column c. Split features and labels d. Standardize numerical features using training statistics e. If class imbalance present, train CTGAN on dataset with conditioning on label f. Generate synthetic minority samples to balance classes g. Merge synthetic samples and real data h. Store the preprocessed dataset 2. Transformer-based Feature Extraction a. Apply Transformer feature encoder on balanced data to obtain contextual embeddings for each sample b. Pool embeddings to obtain feature vectors c. Store the extracted features and labels 3. SHAP-based Feature Selection a. Train XGBoost on extracted features b. Compute SHAP values and mean absolute SHAP importance I_j c. Select features covering 95% cumulative importance d. Save selected feature list 4. Model Training (TabNet-SE) a. Build two datasets: X_{all} and X_{sel} b. Perform split (80/20) c. For each train/val split: instantiate TabNet-SE architecture with SE block, train using AdamW with OneCycleLR, apply label smoothing, use early stopping on validation AUC d. Save best model weights and compute probabilities on test set 5. Baseline Training (RF & XGBoost) a. Train RandomForest and XGBoost on both X_{all} and X_{sel} with identical splits b. Save models and compute probabilities on test set 6. Results Visualization and Evaluation a. Compute confusion matrix, accuracy, precision, recall, F1, AUC, Brier score, calibration curve b. Perform DeLong test to compare AUCs between TabNet-SE and baselines c. Generate SHAP summary plot and attention bar chart d. Store all metrics and plots in the results directory.

7. Output

- Save trained models
- Record random seeds, package versions, and training hyperparameters

8. End

4. Results and Discussion

The performance of the proposed model was evaluated on a clinical dataset of Polycystic Ovary Syndrome (PCOS), which is freely accessible to public. The dataset consisted of 364 normal samples and 177 PCOS samples [27]. Conditional Tabular Generative Adversarial Network was applied to increase the minority class synthetically with 251 PCOS samples. The purpose was to handle class imbalance and enhance model generalization. Feature extraction Feature extraction was performed with a Transformer-based feature encoder that generated 64 features for each clinical sample in the pre-processed dataset. Then, a SHAP-based XGBoost Explainer was employed for feature selection that identified

and selected 47 important features out of the total 64 extracted features. The TabNet-SE classifier that has been described in its proposed form, along with two baseline models - the size and dispersion of both models, were trained and tested using both the full feature set, as well as the selected key features. To conduct systematic experimentation, the complete dataset was divided into training and test sets in 80:20 ratios.


Models were evaluated with standard metrics of accuracy, precision, recall, F1-score and the ROC-curve analysis. All of these measurements indicate how well the models can reliably discriminate between PCOS and non-PCOS cases. Table 1 presents the dataset's fine-grained class distribution.

Table 1. Detail on clinical dataset


Classes	No. of Actual Samples	No. of Augmented Samples
Normal	364	364
PCOS	177	251
Total Samples	541	615

Figure 7 provides a comparison of the different classifiers, using full features. The confusion matrix is shown in Figure 7(a). The proposed TabNet-SE classifier identifies 62 instances as Normal and 46 instances as PCOS, which indicates high discriminative power between classes. The ROC curves for TabNet-SE and the baseline methods based on Random Forest and XGBoost are presented in Figure 7(b). By observing as it is shown in Figure 7(b), TabNet-SE had the highest value of ROC for all classes with a significant difference, which indicates that the classification performance is better and has strong generalization compared to other models.

The performance of the compared models evaluated with the selected feature set is shown in Figure 8. In Figure 8(a), demonstrate the confusion matrix of proposed TabNet-SE model that correctly classify 67 instance as Normal and 48 instance as PCOS. It shows that it is better at discriminating with optimized features. Figure 8(b) presents a comparison of the ROC curves for the proposed TabNet-SE model with baseline classifiers, Random Forest and XGBoost. The TabNet-SE model achieved the highest ROC value in all classes as seen from the figure. This means that it has more accuracy for classification and uses the features properly, and provide better generalization than baseline models.




(a) Confusion matrix of proposed TabNet-SE method for complete features




(b) ROC curves comparison against proposed and baseline models for complete features

Fig. 7 Classification analysis of proposed and baseline approaches under complete feature set during testing.



(a) Confusion matrix of proposed TabNet-SE method for selected features



(b) ROC curves comparison against proposed and baseline models for selected features

Fig. 8 Classification analysis of proposed and baseline approaches under selected feature set during testing.

The calibration curve of proposed TabNet-SE is shown in Figure 9. The closer the two-dimensional points are to the diagonal line, the better is a model in terms of producing probability estimates that are finely calibrated for PCOS classification.

The closeness of our calibration curve to the perfect diagonal reference line demonstrates that we have reliable confidence consistency across the spectrum well. These results further support the strong consistency of this model for real-world clinical applications.




Fig. 9 Calibration curve of the proposed TabNet-SE model

The model uncertainty demonstrated by the MC dropout predictive standard histogram clearly shows that the variability in predictions from stochastic dropout sampling is due to variance in a realisation of MC dropout.

At inference time, the dropout layers that are typically disabled will be turned on. This means that the model will be then fed multiple times with the same input to create slightly different probability estimates. The standard deviation of these predictions represents epistemic uncertainty, or how confident the model is in its outputs. In Figure 10, the MC Dropout predictive uncertainty histogram based on the TabNet-SE is shown. This figure shows the uncertainty

distribution of all test samples. As we can observe, the values are concentrated towards lower standard deviation levels further confirming that the proposed model is producing stable and confident predictions with very little uncertainty (or confidence) in these predictions.

This contributes to its robustness and reliability for the classification of PCOS. Higher standard deviations of a few suggest borderline cases which are always harder to interpret. This study reveals that it is suitable and highly robust predictive model for TabNet-SE achieve clinically actionable results.




Fig. 10 MC dropout predictive uncertainty histogram of the proposed TabNet-SE model

The performance comparison of the proposed TabNet-SE model with baseline models, RF and XGBoost, is given in Table 2 and Figure 11. This includes using both the complete feature set for diagnosing PCOS and the feature set selected by SHAP. As can be seen from the accuracy, precision, recall, and F1-score statistics across all of the models, the models trained with a subset of features out-performed their counterparts trained on the full feature set. The highest accuracies of 93.50%, precision of 88.89% recall of 96.00% and F1-score of 92.31% were obtained with a TabNet-SE model on the selected feature set by outperforming all the baseline models as shown later in this section. The results indicate the usefulness of SHAP-based feature selection on enhancing model performance to discriminate among classes and indicates strength and adaptability to model building supporting the efficacy of TabNet-SE architecture for robust classification of PCOS.

The effectiveness of CTGAN-based class balancing, Transformer-driven contextual feature extraction, SHAP-guided feature optimization and attention-based deep tabular learning results in the superior performance of the developed TabNet-SE framework. With the help of SHAP, the model avoids redundant attributes and keep clinically explainable features to make it more interpretable and easier to learn. The Transformer encoder captures complex inter-feature relationships. In addition, the enhanced TabNet architecture uses Squeeze-Excitation, in which an attention mechanism is employed whereby discriminative clinical patterns are adaptively focused on through feature reweighting. Through these several mechanisms, the proposed framework boosts generalizability and reduces predictive uncertainty, while also improving classifier robustness, all of which collectively allow the proposed frame to perform better than existing state-of-the-art PCOS classification methods.




Fig. 11 Performance comparison of models across evaluation metrics

Table 2. Comparative performance analysis of models on complete and selected feature sets

Metrics	Complete Feature Set			Selected Feature Set		
	XGBoost (%)	RF (%)	TabNet-SE (%)	XGBoost (%)	RF (%)	TabNet-SE (%)
Accuracy	78.05	85.37	87.80	86.99	90.24	93.50
Precision	75.56	84.78	80.70	80.36	83.93	88.89
Recall	68.00	78.00	92.00	90.00	94.00	96.00
F1-Score	71.58	81.25	85.98	84.91	88.68	92.31

The experimental study confirms that CTGAN-based balancing, Transformer-driven contextual feature extraction, SHAP-guided feature optimization and TabNet-SE classification is a highly effective scheme to enhance the predictive performance and interpretability. Beyond enabling a better classification accuracy, the proposed framework exhibits stable uncertainty behavior and reliability in probability prediction, thus demonstrating its clinical interpretability potential for PCOS decision-support systems.

4.1. Statistical Significance and Error Analysis

In order to provide a more thorough proof of the robustness of the proposed framework, DeLong's test for comparing AUC was applied to compare the TabNet-SE model vs baseline classifiers. DeLong's test is the most widely used non-parametric statistical method that compares if the difference between ROC curves is statistically significant. Let AUC_1 and AUC_2 be the AUCs of two competing models. The null hypothesis is that both classifiers have the same discriminative capability:

$$H_0: AUC_1 = AUC_2$$

while the alternative hypothesis is defined as:

$$H_1: AUC_1 \neq AUC_2$$

The corresponding z-score is computed as:

$$z = \frac{AUC_1 - AUC_2}{\sqrt{\text{Var}(AUC_1) + \text{Var}(AUC_2) - 2 \text{Cov}(AUC_1, AUC_2)}} \quad (38)$$

Where, AUC_1 is Area Under Curve of the first model, AUC_2 is Area Under Curve of the second model, $\text{Var}(AUC_1)$ is Variance of the first AUC, $\text{Var}(AUC_2)$ is Variance of the second AUC, and $\text{Cov}(AUC_1, AUC_2)$ is Covariance between the two AUC estimates.

This z-score represents the covariance-based variance estimated using DeLong's formulation. The p-value is subsequently derived from the standard normal distribution.

Comparison of the statistical significance between the proposed TabNet-SE model and baseline methods is shown in Table 3. The resulting p-values are below the common significance level of 0.05 thereby validating that the performance gain achieved by our proposed TabNet-SE framework is not a random variation. Error analysis was also performed based on the confusion matrix observations and uncertainty estimation using Monte Carlo dropout analysis. The majority of misclassifications were seen in borderline clinical cases, suggesting overlap between hormonal and metabolic characteristics in the normal and PCOS categories. However, the suggested TabNet-SE model displayed lower prediction uncertainty with substantial reduced false predictions than baseline models, showing more robustness and generalizability in real-life clinical decision-support scenarios.

Table 3. Statistical significance analysis using DeLong's test

Model Comparison	AUC (Model 1)	AUC (Model 2)	z-score	p-value	Significance
TabNet-SE vs RF	0.778	0.721	2.41	0.0159	Significant
TabNet-SE vs XGBoost	0.778	0.632	3.87	0.0001	Significant
RF vs XGBoost	0.721	0.632	1.98	0.0476	Significant

4.2. Ablation-Based Statistical Analysis

Ablation study is performed to highlight the contributions of each major component in explainable deep tabular learning framework for PCOS classification. First, the model establishes the reference performance levels with respect to baseline XGBoost and Random Forest models trained using the entire feature set. It is then followed by a CTGAN-based balancing and Transformer-driven contextual feature extraction technique that together with the previous stage significantly augment the representational capacity of features ensuring better generalization of the clinical data under

imbalanced conditions. This model is further optimized via SHAP-based feature selection, producing a compact representation of the data by eliminating many redundant and less-informative attributes from the complemented full set of clinical variables while retaining clinically significant features to provide more interpretable learning. Finally, the model proposed a Squeeze-Excitation enhanced TabNet architecture that can achieve the best classification accuracy of 93.50% by adaptively recalibrating the discriminative feature channels using attention-driven learnings. The ablation findings underlines that the sequential incorporation of CTGAN

balancing, feature extraction via Transformers, feature optimization using SHAP, and SE enhancement in TabNet strongly contribute to the robustness, interpretability and

overall predictive capability of the DNN-based PCOS classification framework. The Ablation Results in Table 4 illustrate the detailed workings of these methods.

Table 4. Ablation study of the proposed framework for PCOS classification

Configuration	CTGAN Balancing	Transformer Feature Extraction	SHAP Feature Selection	SE-Enhanced TabNet	Accuracy (%)
XGBoost (All Features)	✓	✓	✗	✗	78.05
Baseline RF (All Features)	✓	✓	✗	✗	85.37
CTGAN + Transformer + TabNet (All Features)	✓	✓	✗	✗	87.80
CTGAN + Transformer + XGBoost (Selected Features)	✓	✓	✓	✗	86.99
CTGAN + Transformer + RF (Selected Features)	✓	✓	✓	✗	90.24
Proposed TabNet-SE (Selected Features)	✓	✓	✓	✓	93.50

All experiments were realised using Python with PyTorch and Scikit-learn. The stratified 80:20 train-test split was performed, using a fixed random seed of 42 allow results to be reproducible. Experiments were performed in a systems having Intel Core i5 CPU, 16 GB RAM and NVIDIA GPU acceleration. The purpose of the current study was to conduct a controlled comparative evaluation of all models in similar experimental conditions, thus external validation and cross-validation strategies were not employed.

4.3. Scalability and Clinical Deployment

In practical situation for development, the proposed framework is computationally efficient because it works on structured clinical tabular data. The lightweight TabNet-SE architecture with SHAP-based interpretability allows scalable integration in clinical decision-support systems for early PCOS diagnosis and assessment of risk. Furthermore, since the calibrated probability estimations are also uncertain aware predictions of unseen cases with respect to seen data on specific exhibits related to pediatric digital health and post-acne growth, they can result in great reliability for real-world healthcare applications. The settings of such a framework make it implementable for pervasive clinical application and an intelligent healthcare platform deployed in resource limited environments.

5. Conclusion and Future Scope

The proposed research establishes a simple and interpretable deep learning framework for classifying PCOS using clinical data. The proposed framework demonstrates superior diagnostic efficiency and interpretability via a joint pipeline of preprocessing, transformer-based feature extraction, SHAP-guided feature selection, followed by a new TabNet-SE classifier. It solves the class imbalance with CTGAN augmentation while honing the feature space through SHAP based XGBoost explainer selection. While the TabNet-

SE is a strong performer in comparison to baseline models, Random Forest, and XGBoost achieving 93.50% accuracy, 88.89% precision, 96.00% recall and F1-score of 92.31%. The model's reliability and transparency is supported by attention weights, SHAP summaries and Monte Carlo dropout uncertainty analyses, and has demonstrated consistent strong performance across routine clinical prediction tasks.

Future studies will expand upon this study by including ultrasound imaging data containing structural and morphological characteristics of the ovaries in addition to clinical features. As a result, this knowledge would enable the framework to do an analysis of PCOS with greater depth as it could utilize both visual and tabular data. Advanced convolutional and attention mechanisms will be studied for determining the relevant features of images, combined with explainable AI approaches to ensure interpretability in driving Indian state automobile policies. Bringing together different kinds of data will supply tailored insights into the various phenotypes of PCOS which will aid in early detection, treatment planning and clinical decision-making in real-world healthcare environments.

Declarations

- Ethics approval and consent to participate
Not applicable. No human participants, animals or sensitive data were involved in this study that would require ethics approval.
- Consent for publication
Not applicable. This manuscript does not include any identifiable personal data.
- Availability of data and material
The dataset we used in this study is publicly available on the Kaggle platform. Upon request from the

corresponding author, we will provide specific preprocessing steps and code.

- **Competing interests**
The authors have no competing interests in relation to this work.
- **Funding**
This research was not supported by any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.
- **Authors' contributions**
B. Pushpa: The study design, the data preprocessing, experimental analysis and visualizations were done by her.

V. Lakshmi: Model architecture design; writing of the manuscript; statistical evaluation.

Both authors read and approved the final version of the manuscript.

Acknowledgements

The authors would like to thank their respective institutions for providing funding to conduct this research; the support has translated into valuable guidance, feedback and resources that made all the difference.

Conflicts of Interest

The authors have reported no conflicts of interest related to the present study.

References

- [1] Francisco J. Barrera et al., "Application of Machine Learning and Artificial Intelligence in the Diagnosis and Classification of Polycystic Ovarian Syndrome: A Systematic Review," *Frontiers in Endocrinology*, vol. 14, pp. 1-12, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Zahra Zad et al., "Predicting Polycystic Ovary Syndrome with Machine Learning Algorithms from Electronic Health Records," *Frontiers in Endocrinology*, vol. 15, pp. 1-14, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [3] Abrar Alamoudi et al., "A Deep Learning Fusion Approach to Diagnose the Polycystic Ovary Syndrome (PCOS)," *Applied Computational Intelligence and Soft Computing*, pp. 1-15, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] Nikos Fazakis et al., "Iterative Robust Semi-Supervised Missing Data Imputation," *IEEE Access*, vol. 8, pp. 90555-90569, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] Lei Xu et al., "Modeling Tabular data using Conditional GAN," *arxiv preprint*, 2019. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [6] Serkan Ö. Arik, and Tomas Pfister, "TabNet: Attentive Interpretable Tabular Learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 8, pp. 6679-6687, 2021. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [7] Chiranjibi Shah, Qian Du, and Yan Xu, "Enhanced TabNet: Attentive Interpretable Tabular Learning for Hyperspectral Image Classification," *Remote Sensing*, vol. 14, no. 3, pp. 1-21, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Scott M. Lundberg, and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems*, vol. 30, 2017. [[Google Scholar](#)]
- [9] Mehtap Agirsoy, and Matthew A. Oehlschlaeger, "A Machine Learning Approach for Non-Invasive PCOS Diagnosis from Ultrasound and Clinical Features," *Scientific Reports*, vol. 15, no. 1, pp. 1-18, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [10] Kogilavani Shanmugavadeivel et al., "Optimized Polycystic Ovarian Disease Prognosis and Classification using AI based Computational Approaches on Multi-Modality Data," *BMC Medical Informatics and Decision Making*, vol. 24, no. 1, pp. 1-22, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] Neha Yadav, Ranjith Kumar A, and Sagar Dhanraj Pande, "Comparative Analysis of Polycystic Ovary Syndrome Detection using Machine Learning Algorithms," *EAI Endorsed Transactions on Pervasive Health and Technology*, vol. 10, pp. 1-5, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Oviya Graselina S et al., "Development of a Machine Learning Model to Classify Polycystic Ovarian Syndrome," *Technology and Health Care*, vol. 33, no. 3, pp. 1478-1488, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [13] Md Mahbubur Rahman et al., "Empowering Early Detection: A Web-based Machine Learning Approach for PCOS Prediction," *Informatics in Medicine Unlocked*, vol. 47, pp. 1-16, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [14] Hela Elmannai et al., "Polycystic Ovary Syndrome Detection Machine Learning Model based on Optimized Feature Selection and Explainable Artificial Intelligence," *Diagnostics*, vol. 13, no. 8, pp. 1-21, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [15] Pooja Balagouda Patil et al., "Explainable Ensemble-based Machine Learning Model for Polycystic Ovary Syndrome Detection using Hybrid Feature Selection," *International Journal of Information Technology*, pp. 1-15, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Wenxiu Chen et al., "Development of Machine Learning Models for Diagnostic Biomarker Identification and Immune Cell Infiltration Analysis in PCOS," *Journal of Ovarian Research*, vol. 18, no. 1, pp. 1-16, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]

- [17] Rebecca H.K. Emanuel et al., “What do users in a Polycystic Ovary Syndrome (PCOS) Forum Think About the Treatments They Tried: Analysing Treatment Sentiment using Machine Learning,” *Physical and Engineering Sciences in Medicine*, vol. 48, no. 2, pp. 723-741, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [18] Eugenia Papadaki, Aristidis G. Vrahatis, and Sotiris Kotsiantis, “Exploring Innovative Approaches to Synthetic Tabular Data Generation,” *Electronics*, vol. 13, no. 10, pp. 1-20, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [19] Vasileios C. Pezoulas, “Synthetic Data Generation Methods in Healthcare: A Review on Open-Source Tools and Methods,” *Computational and Structural Biotechnology Journal*, vol. 23, pp. 2892-2910, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] M. Kannan et al., “An Enhancement of Machine Learning Model Performance in Disease Prediction with Synthetic Data Generation,” *Scientific Reports*, vol. 15, no. 1, pp. 1-21, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [21] Mikel Hernandez et al., “Comprehensive Evaluation Framework for Synthetic Tabular Data in Health: Fidelity, Utility and Privacy Analysis of Generative Models with and without Privacy Guarantees,” *Frontiers in Digital Health*, vol. 7, pp. 1-18, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [22] Malak Alqulaity, and Po Yang, “Enhanced Conditional GAN for High-Quality Synthetic Tabular Data Generation in Mobile-based Cardiovascular Healthcare,” *Sensors*, vol. 24, no. 23, pp. 1-20, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [23] Ha Ye Jin Kang, Minsam Ko, and Kwang Sun Ryu, “Tabular Transformer Generative Adversarial Network for Heterogeneous Distributions in Healthcare,” *Scientific Reports*, vol. 15, no. 1, pp. 1-12, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [24] Ahmed Fahim et al., “A Hybrid Deep Learning Framework based on CNN-GRU-TabNet for the Predictive Modeling of COVID-19 Mortality,” *Engineering Technology and Applied Science Research*, vol. 15, no. 5, pp. 28057-28062, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [25] Jiekee Lim et al., “Machine Learning Classification of Polycystic Ovary Syndrome based on Radial Pulse Wave Analysis,” *BMC Complementary Medicine and Therapies*, vol. 23, no. 1, pp. 1-15, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [26] Bharti Panjwani et al., “Optimized Machine Learning for the Early Detection of Polycystic Ovary Syndrome in Women,” *Sensors*, vol. 25, no. 4, pp. 1-30, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [27] Kaggle, PCOS Dataset, Kaggle, 2020. [Online]. Available: <https://www.kaggle.com/datasets/shreyasvedpathak/pcos-dataset>