International Journal of Applied Mathematics
Volume 38 No. 10s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

DEEP CNN ENSEMBLE MODELS FOR ACCURATE OVARIAN CANCER DETECTION

Snehalatha N, JSS Academy of Technical Education, Bangalore, snehalathan@jssateb.ac.in

Annapoorna V, JSS Academy of Technical Education, Bangalore

Divya N, JSS Academy of Technical Education, Bangalore

H Khushi, J SS Academy of Technical Education, Bangalore

ABSTRACT

Ovarian cancer is still among the most lethal gynaecologic malignancies because it is detected too late and has blurred early manifestations. An early diagnosis is crucial to enhance the cure rates and survival rate. In this paper, we intend to propose a deep convolutional neural network ensemble framework based on transfer learning for the classification of ovarian tumours into three categories, namely Normal, Benign, and Malignant, with high accuracy. This ensemble combines three optimized models: a fine-tuned ResNet50, a VGG16 coupled with an FCNN designed in this work, and a fine-tuned EfficientNet-B0. All three models have been pre-trained on ImageNet and were adapted to the ovarian ultrasound domain by selective fine-tuning of layers and domain-specific optimization techniques.

in order to handle class imbalance and improve diagnostic reliability, class weights and Focal Loss were used (for ResNet50), along with extensive TTA to improve model generalization. Adaptive optimizers, early stopping, and checkpoint mechanisms were employed to ensure the convergence of the networks in a stable manner. Experimental evaluation shows that the proposed ensemble significantly improves the accuracy, sensitivity and specificity of classification as compared to the any standalone model, thereby offering the great potential in serving as an intelligent and reliable decision support system for early ovarian cancer screening. The study emphasizes how combining the transfer learning techniques with ensemble learning models can significantly enhance the performance of AI driven medical diagnostic systems, making them more accurate, reliable, and clinically useful for the real-world diagnosis healthcare applications.

KeyWords – Deep CNN, Transfer Learning, ResNet50, VGG16, EfficientNet-B0, Fully Connected Neural Network, Test-Time Augmentation, Focal Loss, Ensemble Model, Ovarian Cancer Detection, Ultrasound Imaging, Computer-Aided Diagnosis, Early Cancer Detection, AI in Healthcare.

I. INTRODUCTION

Ovarian cancer is a very aggressive type of gynaecologic malignancy, often diagnosed at late stages due to the lack of specific early signs. Early detection and accurate diagnosis are critical since it significantly increases the chances of successful treatment and survival.

Received: July 20, 2025

Ultrasound imaging is a basic diagnostic tool for abnormal ovary conditions, defined by its non-invasive, inexpensive, and easily accessible nature. Manual interpretation of ultrasound scans is largely dependent on the professional skills of radiologists and is prone to errors, potentially leading to misdiagnosis or a delay in treatment. Recent advances in deep learning, especially the CNN architectures, have shown a strong potential for the automation of medical image analysis and improvement in diagnostic consistency.

This paper proposes an ensemble of transfer learning–based deep CNN models to improve the accuracy and reliability in classifying ovarian tumours from ultrasound images. It brings together the three optimized architectures, namely ResNet50,VGG16 with FCNN, and EfficientNet-B0 each model was fine-tuned for the domain-specific learning and balanced performance. ResNet50 mainly adopts the residual connections and the Focal Loss to adapt the well to class imbalance, while VGG16-FCNN reinforces generalization for the images through dropout and the L2 regularization. Meanwhile, EfficientNet-B0 performs lightweight but the powerful feature extraction through the compound scaling. Further improving the prediction consistency, Test-Time Augmentation is used to ensure robust performance in different image orientations and variations.
Artificial intelligence (AI), in deep learning, has become an effective tool in medical image analysis. Convolutional neural networks (CNNs), with the assistance of transfer learning involving architectures like ResNet50 and VGG16, can learn to extract fine features in images procured with the aid of ultrasound so that accurate classification becomes functional in spite of the constraints placed by limited datasets. Deep learning models, though remarkable, do not fare well with noisy or uncertain medical data so that the accuracy of the prediction can be affected.

It realizes superior diagnostic accuracy, sensitivity, and robustness by combining the advantages of the two complementary models in the ensemble. In general, the current study has shown that both the transfer learning and ensemble strategies can boost the AI-assisted diagnosis of ovarian cancer to support early and reliable clinical decisions. This research proposes to develop an effective diagnostic model based on AI that helps clinicians to render judicious and timely decisions and improve the outcome for the patients.

II. PROBLEM STATEMENT AND OBJECTIVES

Ovarian cancer is often detected at an advanced stage because early-stage symptoms are difficult to identify and can be easily disregarded. Despite being a widely used non-invasive diagnostic method, ultrasound imaging's manual interpretation relies heavily on the radiologist's expertise and can lead to inconsistent or inaccurate diagnoses. Traditional machine learning and deep learning techniques have improved image-based diagnosis; however, these methods are hampered by the uncertainty, noise, and small datasets that are typical of medical imaging. Moreover, most existing systems only deal with deep learning, neglecting the reasoning and interpretability elements needed for clinical decision support.

Therefore, there is an urgent need for a clever and reliable diagnostic model that combines the high feature-extraction power of deep transfer learning. Moreover, most existing systems

International Journal of Applied Mathematics
Volume 38 No. 10s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

only deal with deep learning, neglecting the reasoning and interpretability elements needed for clinical decision support.

The main objectives of this study are:

i. **Design and develop transfer-learning deep CNN models, VGG16+FCNN, ResNet50(Fine-tuned), and EfficientNet-B0** for the accurate classification of ovarian tumour using ultrasound images.

ii. **To improve model adaptability and durability** by incorporating advanced data preprocessing, class balancing, and fine-tuning strategies to refine pretrained architectures to the ovarian-cancer images.

iii. **To implement Test Time Augmentation(TTA)** techniques which will improve the credibility and robustness of predictions across different orientations and noisy conditions of ultrasound images.

iv. **To combine single-model outputs into an ensemble diagnostic framework** that employs the strengths of each architecture, leading to increased performance in classification and diagnostic accuracy.

v. **To develop an intelligent, automated framework for early and accurate ovarian cancer detection**, assisting medical experts in making accurate diagnostic decisions using deep learning–based image analysis.

III. LITERATURE SURVEY

The enhancements in artificial intelligence (AI) that have taken place over the last few years have significantly improved the early detection and staging of ovarian cancer, particularly through the combination of deep-learning and machine-learning methods. Some studies provide evidence that the ability of the AI systems could possibly reach or even exceed the conventional methods of diagnosis.[2]

Deep learning algorithms used for ultrasound and histopathological imaging have pointed at a remarkable future ahead. One of the studies, published in the journal Ultrasound in Obstetrics & Gynaecology in 2021, stated that deep neural networks (DNNs) could reach the same level of diagnostic performance as expert radiologists in classifying benign and malignant ovarian tumours.[1] At the same time, a study that was conducted in 2017 and utilized segmentation and edge detection techniques, succeeded in discriminating different cyst types in ultrasound pictures, thereby uncovering the basic capacity of image processing in the field of ovarian diagnostics.[5] A new study employing superior pre-processing techniques along with PCA-based algorithms obtained an AUC score of 0.9136 and generalization accuracy of 89.79%, thus indicating the role of pre-processing as a major factor affecting the robustness of the model.[10] In the field of histopathological imaging, research conducted in 2023 with MobileNetV3 and ResNet50 models got 96.3% and 92.08% accuracy, respectively, whereas a 2025 hybrid ResNet-SE-Net model upgraded the result to 96%, validating the importance of architecture optimization for visual feature learning.[11][13]

Received: July 20, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Besides imaging, machine learning systems developing on the basis of biomarker data have also emerged in early-stage prediction. In a 2024 article on Prediction of Ovarian Cancer using AI Tools, it was identified that through the assistance of biomarkers such as HE4, CA125, and NEU, a Random Forest classifier was able to achieve high precision exceeding 86%.[8] In similar 2022 studies on clinical data and blood-based biomarkers, 91% precision was achieved while emphasizing the predictivity of biochemical characteristics.[12] However, such works also determined that the performance of biomarker-based systems largely depends on the presence and stability of clinical data and, therefore, future systems should integrate both imaging and biomarker modalities.

At the same time, hybrid and fuzzy deep learning systems have continued to extend accuracy frontiers for histopathological images. In Scientific Reports in 2024, researchers proposed a fuzzy deep learning system that integrated ResNet50 and recursive feature elimination to obtain 98.99% accuracy on benchmark data.[3] These kinds of systems not only increase the precision of the diagnosis but also implement the handling of uncertainty, which is very important in terms of clinical variability. Likewise, optimized hybrid networks such as ResNet-SE-Net illustrate how feature refinement with attention boosts malignancy classification performance.[13]

Systematic reviews also solidify AI's therapeutic potential. AI's increasing role in early diagnosis, drug response prediction, and personalized therapy was emphasized in the 2024 OxJournal review and 2025 article Early Diagnosis of Ovarian Cancer: A Comprehensive Review.[4] Multimodal classification review articles showcase how incorporating imaging data with markers and CTs provide more complete diagnostic paradigms. In addition, explainable AI (XAI) methods continue to grow in usage due to transparency, illustrated in the 2024 research utilizing Support Vector Machines and ensemble models with MRMR feature selection that had 89% accuracy.[6][7][9]

The study called MMOTU (Zhao et al., 2023) presents an extensive ovarian tumour ultrasound dataset with multi-modalities that consist of 2D and CEUS images and pixel-wise annotations. Moreover, it suggests the Dual-Scheme Domain-Selected Network (DS2Net) for enhancing unsupervised cross-domain segmentation of ultrasound modalities, which is a solution to the domain shift problem in medical imaging.[14]

Overall, existing literature presents that AI-related systems, in particular deep learning and fuzzy hybrid systems, have attained over 98% diagnostic precision under constrained settings. However, most require multi-institutional validation, data variability, and multimodal input integration to validate generalizability and clinical reasonableness. Such results provide a foundation on which to develop an integrated, interpretable, and durable deep learning system to identify early-stage ovarian cancer.[3]

IV. METHODOLOGY

The research work has presented the application of modern deep learning and hybrid AI techniques which will significantly improve the detection of ovarian cancer through ultrasound imaging in its early stage. Ovarian cancer detection through ultrasound images is

International Journal of Applied Mathematics
Volume 38 No. 10s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

the main topic of concern. The new method that has been put forward connects transfer learning (VGG16, ResNet50) with fuzzy logic and fully connected networks to produce a solid, precise, and explainable classification.

1. Data Classification

The acquired MMOTU (Microscopic Multi-class Ovarian Tumour Ultrasound) dataset comprised raw ultrasound images and their respective annotation masks in a colour-oriented display format. In the annotation imagery, specific colours are utilized per type of tissue and tumour class. For the purpose of supervised learning, however, the dataset was first pre-processed such that the images were automatically classified per the categories of Benign, Malignant, and Normal from the available colour information in the annotation files.

Colour-Based Label Mapping

The MMOTU dataset employs certain RGB colour codes in the annotation images so as to indicate varied classes:

Malignant neoplasm → Teratoma [0, 0, 64],

HighGradeSerousCystadenoma [0, 0, 128],

Normal ovarian tissue → Gray ([64, 64, 64])

Benign tumour → Any other shade area apart from the above two RGB triplets were employed as the reference values so that each picture can be classified. The script translates each annotation file into a NumPy array and performs a pixel-wise matching in order to identify the presence of such colours.

Image–Label Pairing

Each annotation file (e.g., image_001.PNG) was paired with its respective raw image (image_001.JPG). The files were skipped if no matching picture was available for the input label in order to maintain the integrity of the data.

Classification Logic

The class identification was based on the existence of pixels with a unique colour in the annotation:

- When the annotation consisted of Teratoma [0, 0, 64], HighGradeSerousCystadenoma [0, 0, 128], then the corresponding ultrasound imagery was annotated with Malignant
- If the annotation contained grey pixels in the form ([64, 64, 64]), then it was assigned Normal.
- In the absence of detection of these two colours, the picture was rated Benign.

This rule-based classification provides an automatic and reproducible labelling scheme compatible with the colouring coding system in the dataset.

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Relocation and File Structuring

Follows the process of classifying files, script uses the shutil library to move each detected picture into its respective category folder. The process essentially restructured the dataset into separate directories, thus making the later stages like:

- Training and validation partition
- Image segmentation or training model for classification

Error Handling and Logging

The system includes multiple safety verification mechanisms:

- Verification that the image and the annotation files exist.
- Exception handling for unreadable or corrupted image files.
- Console messages to log the progress and notify about missing or mismatched files.

2. Data Preprocessing and Augmentation

All the ultrasound images also went through a defined preprocessing pipeline prior to model training to maintain consistency, noise elimination, and the best possible input quality for the deep learning models. Preprocessing was done by using OpenCV and NumPy libraries in Google Colab while the datasets were directly accessed from Google Drive for effective cloud processing.

The unprocessed ultrasound images were acquired from various sources and archived in a specific directory. In order to standardize the dataset, a completely automated preprocessing script developed in Python was executed Each image underwent the subsequent sequential procedures:

Loading and checking images

All the ultrasound images were imported from dataset folder and checked for corrupted or unreadable files in to prevent runtime errors when training.

Grayscale Conversion:

Ultrasound analysis does not need colour information; therefore, all images were processed to grayscale. This facilitated a 3-dimensional data reduction and allowed the model to be more sensitive to texture patterns based on intensity that are significant for ovarian structures.

Resizing and Padding:

The ultrasound images were resized to a central resolution of 224×224 pixels in order to conform to the input dimensions of the pre-trained models (VGG16 and ResNet50) which required these dimensions. To keep the aspect ratio intact, each image was first uniformly scaled and then black borders were added around it to reach the target size without any distortion.

Normalization:

Pixel values were transformed through a linear operation to a scale of 0-1 by the divisor of 255.0. Thus, the portals were opened for equal intensity distribution which led to the improvement of model convergence and heightened numerical stability throughout the training process.

Batch Processing and Warehousing:

The complete dataset was batch processed and stored afterwards in a systematic output folder for further training. For the training process, the data augmentation techniques rotating, flipping, and contrast variation were employed for the purpose of balancing the classes to improve the overall generalization. The preprocessing pipeline consequently established consistency among all input images and reduced data anomalies, thereby rendering the dataset resilient for the training of all ResNet50, VGG16 + FCNN and EfficientNet-B0 models.

3. Model Architecture

The research utilizes deep transfer learning models, particularly VGG16 and ResNet50, fused with fully connected and EfficientNet-B0, to allow accurate multi-class classification of the ovarian ultrasound images. The models are built to encode hierarchical feature representations while at the same time boosting the robustness and interpretability of the decision in the early detection of cancer.

A. Model 1: Fine-Tuned ResNet50 Architecture

The first proposed model in this ensemble framework is ResNet50, which is a 50-layer deep convolutional neural network pre-trained on the ImageNet dataset. Due to its residual skip connections, ResNet50 effectively avoids the vanishing-gradient problem while training deep feature hierarchies for the recognition of complex images.

The transfer learning is applied in such a way that the convolutional layers of pre-trained ResNet50 were kept as a feature extractor and the upper layers were replaced and tuned for the three categories, viz., Normal, Benign, and Malignant ovarian tumours. To adapt this network to the ovarian ultrasound domain, the last 60 layers of the backbone were unfrozen for fine-tuning of the network, allowing selective updating of deeper features while preserving generalized image representations.
The layout of the network includes the Global Average Pooling layer that is subsequently followed by Batch Normalization, two fully connected layers of 256 and 64 neurons respectively using ReLU activation with L2 regularization to prevent overfitting of the ultrasound image dataset. To enhance generalization and reduce neuron co-adaptation,

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

dropout layers with rates 0.5 and 0.3 were added successively. Finally, this is followed by the SoftMax output layer that predicts probabilities for the three target classes.

Dynamic computation of class weights was performed in order to handle the problem of class imbalance, while a custom Focal Loss function with parameters $\gamma = 2$ and $\alpha = 0.25$ was adopted to drive the network's learning toward the hard-to-classify and minority malignant cases. The network was then trained using the Adam optimizer set with a learning rate of 3×10^{-5}, early stopping, adaptive learning-rate reduction, and checkpoint callbacks to ensure stable and efficient convergence.

Furthermore, TTA techniques were used in both the training and the evaluation processes. A variety of augmented copies for each image were generated, such as random rotations, flips, and brightness modifications, which allowed the model to extract invariant and strong features from different directions and visual conditions, thus the fine-tuned ResNet50's total generalization capability was improved significantly.

B. Model 2:Fine-Tuned EfficientNet-B0 Architecture

The second model EfficientNet-B0 is a lightweight CNN enhanced with proportional scaling that balances network depth, width, and resolution. Pre-trained on ImageNet it provides highly efficient feature extraction with minimal execution cost, the EfficientNet-B0 architecture can be extended for medical image classification tasks such as ovarian tumour detection.

In this model, the pre-trained EfficientNet-B0 was used through the timm library and replaced its original classification head with a single fully connected layer containing three output neurons corresponding to classes Benign, Malignant and Normal respectively. Backbone layers were retained in order to preserve general visual features while fine-tuning the classifier against the ultrasound ovarian cancer dataset.

All the images were resized to 224×224 pixels, and was normalized by ImageNet mean–standard deviation parameters, and was loaded PyTorch's ImageFolder and DataLoader utilities. Training images used a batch size of 32, splitting the data into 70%, 15% for validation and another 15% for testing. For the purpose of ensuring stable convergence and to prevent overfitting, the Adam optimizer was used with a learning rate of 1×10^{-4} and weight decay of 1×10^{-5}.

Class imbalance problem was solved by calculating class weights based on the inverse frequency of samples in each class and integrated this into the Cross-Entropy Loss function. This was trained for 20 epochs with early stopping and model checkpointing using improvements in validation accuracy for optimal model parameters.

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The training pipeline included per-class accuracy monitoring for both the training and validation phases, therefore allowing balanced learning across all the tumour categories. After training, the best models were saved for interpretation, embedding into the ensemble framework.

C. Model 3: VGG16 with Fully Connected Neural Network

The third ensemble model comprises a VGG16 convolutional backbone together with a custom FCNN for classifying ovarian tumours into categories. VGG16 is a pre-trained 16-layer deep CNN on ImageNet, which has been proved efficient in extracting low and mid-level features through its uniformly applied 3 × 3 convolutional filters and max-pooling layers.

In this model, transfer learning was done on the pre-trained VGG16 model for the ovarian ultrasound dataset. The early convolutional layers were frozen to retain generic visual representations, while the last four layers were unfrozen to adapt to domain-specific fine-tuning. The convolutional outputs were flattened and normalized using Batch Normalization in order to further stabilize the gradient updates and improve convergence.

The classification head was designed with three fully connected layers having 512, 256, and 64 neurons respectively, and the ReLU activation function was used along with L2 weight decay of 1×10^{-3} as a means of regularization to prevent overfitting. The model's generalization strength was improved by using dropout layers with rates of 0.5, 0.4, and 0.3 after each dense layer. The last SoftMax layer provided class probabilities for the three categories of tumours: Normal, Benign, and Malignant.

The network is compiled with the Adam optimizer, with a learning rate of 5×10^{-5}, and trained using categorical cross-entropy loss. In handling class imbalance, class weights were computed dynamically based on the frequency distribution of training samples. Training was further optimized using early stopping, learning-rate reduction on plateau, and model checkpointing, ensuring that efficient, stable convergence would be preserved along with the best weights.

Additionally, TTA was integrated to improve the stability of the model and the reliability of the predictions in the model. Through the controlled application of rotations and flips to each test image and taking the mean of their predictions as the final output, this method was able to lower the influence of image orientation and noise that helped getting better results.

4. Model Evaluation and Prediction

The evaluation and prediction phase evaluates the diagnostic performance of the proposed deep learning models, namely VGG16 + Fully Connected Neural Network (FCNN), Fine-

International Journal of Applied Mathematics
Volume 38 No. 10s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Tuned ResNet50, and EfficientNet-B0, to determine whether they can generalize and perform on totally unseen ovarian ultrasound images. While each has an architecture that is specifically different, they use a common evaluation and prediction pipeline for methodological consistency and to compare the results between the different models.

Model Framework and Transfer Learning

All three models have been developed in a transfer learning fashion where pre-trained convolutional backbones, initialized on ImageNet, are adapted to the domain of ovarian ultrasound. The pre-trained layers act like feature extractors with capabilities of capturing fine textural and morphological features of the ovarian tissue.

- VGG16 + FCNN: Feature extraction with VGG16 is followed by a fully connected network that is designed with dense, dropout, and batch normalization layers to enhance generalization.
- ResNet50: It used residual connections for stability in the gradient, fine-tuning higher layers to enhance discrimination of malignant features. Further, the Focal Loss function with $\gamma = 2.0$ and $\alpha = 0.25$ was applied to focus on hard-to-classify minority samples.
- EfficientNet-B0: Scaling up network depth, width, and input resolution together results in much higher accuracies at lower computational cost.

Image Preprocessing and Normalization

All ultrasound images were resized to a fixed size of 224×224 pixels. For each backbone-specific preprocessing function, the normalization of the pixel intensity distributions and the alignment of the data with the ImageNet standards were applied. Therefore, the input scaling was consistent, the variations due to lighting were minimized, and the impact of ultrasound artifacts was decreased, which eventually led to more reliable feature extraction.

Test-Time Augmentation (TTA)

TTA was applied during prediction to increase the model's robustness and to decrease the output variance. Each test image was augmented through slight rotations and horizontal flips to create multiple copies of that image. The final class prediction was made by averaging the probabilities of all the augmented outputs.
This ensemble-like inference strategy not only increased model stability but also ensured reliable performance across different imaging orientations.

Received: July 20, 2025

Prediction Process

The trained model's convolutional layers generated a deep feature map for each pre-processed test image during inference, where every single processed test image was passed through the model's convolutional layers step by step. The Normal, Benign, and Malignant ovarian tumour types were distinguished using these features that had captured the most important spatial and textural attributes together with others nearby ones.

Then the feature maps get flattened and passed on to the fully connected classifier which was responsible for generating class probabilities through SoftMax activation. The model's last prediction was determined by selecting the class with the highest probability as according to the generated probabilities.

Evaluation Metrics

The performance of the models was evaluated following standard statistical metrics which were: Accuracy, Precision, Recall, F1-Score, and Area Under the Receiver Operating Characteristic Curve. The confusion matrix was employed to visualize true versus predicted labels, while ROC curves were plotted for each class as part of the evaluation of discriminative capability. Using these metrics, we will get a thorough understanding of the predictive reliability of each model on unseen ultrasound image dataset.

The definitive assessment and forecasting model that this framework combines encompasses transfer learning, Focal Loss optimization, and TTA-based inference, thus making it a very reliable method for detecting ovarian cancers through diagnostics.

On the one hand, VGG16 + FCNN provides very good baseline accuracy, ResNet50 promotes the learning of cancer features by means of residual optimization, and EfficientNet-B0 claims to be both fast and scalable in the case of computation. All in all, these networks reinforce the ensemble with their combined accuracy and ability to cope with various datasets.

5. Design and Implementation

The web UI was implemented as responsive single page application with HTML5, CSS3, and JavaScript resulting in minimalistic and interactive user interface. Model information, dataset statistics, and measurement metrics are all packed into a JSON manifest, which on runtime dynamically loads and fills out UI components like metric panels, cards, and inline ROC curve plots with SVG. Even more interactivity and effects are achieved through Particles.js, Typed.js, VanillaTilt.js, and ScrollReveal.js to develop an interactive user interface without server assistance, which further makes simple deployment on static hosting possible.

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

For real-time prediction, the interface is coupled with a Flask-based backend to accept ovarian ultrasound images to automatically classify from users. The backend is in touch with a prediction module to perform preprocessing, model inference, and outputting class predictions and confidence scores. Expected label and confidence scores are shown in real-time so results output as JSON for possible integration into other systems. Solid error handling is used to handle unsupported or empty uploads accurately. This solution fills the gap between deep learning model development and clinical deployment, with a simple-to-use tool for ovarian cancer early diagnosis.

The flowchart (Figure 1) provides a vivid picture of the overall process that has been proposed for classifying the ovarian tumour, step by step. After the input data is classified and pre-processed first, the steps will include changing the data into grayscale, reshaping, normalizing, and processing in batches. After this, the data will be sent through three various but equally powerful deep learning architectures: Fine-Tuned ResNet50, Fine-Tuned EfficientNet-B0, and VGG16 with Fully Connected Neural Network (FCNN). The predictions and evaluations are carried out subsequently, based on the outputs from the models using transfer learning technologies along with the strategies of test-time augmentation and the corresponding performance metrics. At long last, the design and implementation phase backs up the model outcomes, thus leading to the final results and analysis.

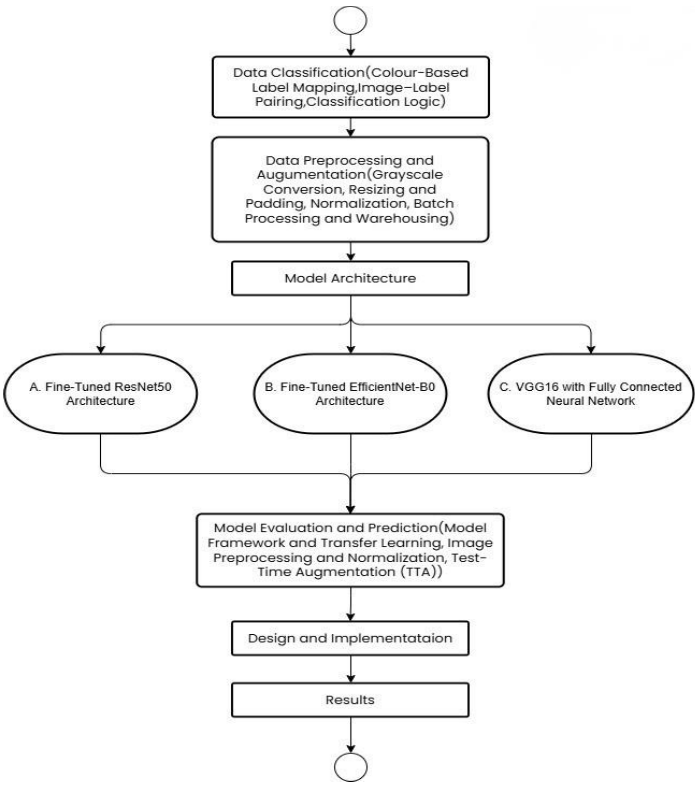

Figure 1: *Overall Workflow of the Proposed Ovarian Tumour Classification Framework*

Received: July 20, 2025

V. RESULTS AND ANALYSIS

The performance of the proposed deep learning models was evaluated on the labelled ovarian ultrasound image dataset. The three trained and tested models were Fine-Tuned ResNet50, VGG16 coupled with a Fully Connected Neural Network (FCNN), and EfficientNet-B0. These models were used to classify the images into three categories: benign, malignant, and normal. For the evaluation, accuracy, precision, recall, F1-score, and Area Under the ROC Curve (AUC) are used; these combined give a comprehensive view regarding classification performance. The confusion matrices, ROC curves, and classification reports for each model that depict their predictive ability and generalization skill are presented in the following figures and tables.

Model 1: Fine-Tuned ResNet50 Architecture

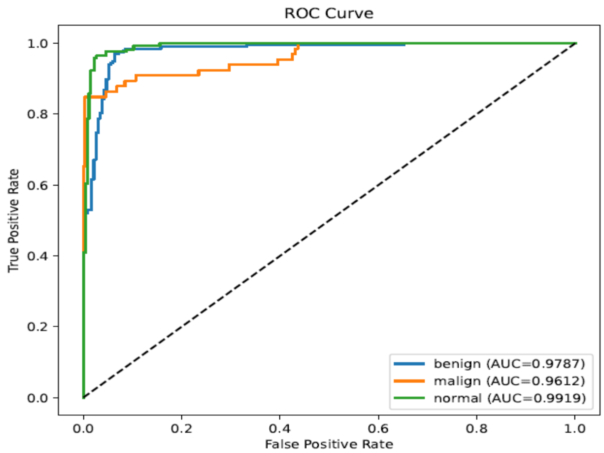

Figure 2: *ROC Curves of the Fine-Tuned ResNet50 Model*

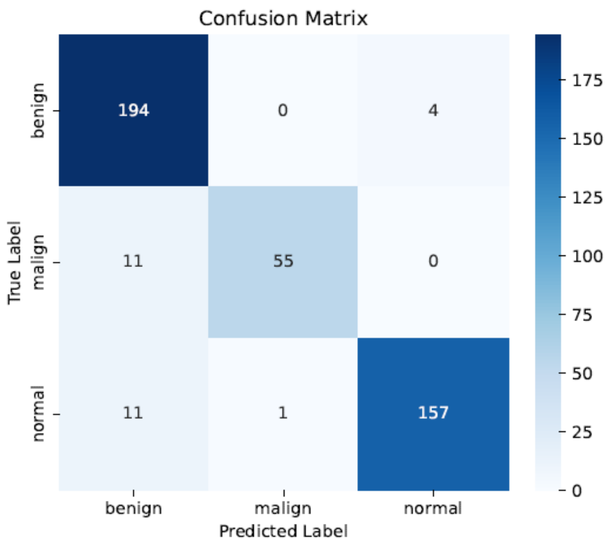

Figure 3: *Confusion Matrix of the Fine-Tuned ResNet50 Model*

Received: July 20, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The ROC curves (Figure 2) depict the model's ability to distinguish among the three tumour classes. The **AUC values** of **0.9787 (benign)**, **0.9612 (malignant)**, and **0.9919 (normal)** highlight the model's excellent classification capability and strong generalization performance. These results validate the effectiveness of the **ResNet50 architecture with transfer learning and fine-tuning** in achieving robust and reliable ovarian tumour classification.

The confusion matrix (Figure 3) represents the classification performance of the fine-tuned ResNet50 model across three ovarian tumour categories: **benign**, **malignant**, and **normal**. The model demonstrates high prediction accuracy, correctly identifying 194 benign, 55 malignant, and 157 normal cases. Only a few misclassifications occurred, indicating that the model effectively captures discriminative features from ultrasound images. The use of **Focal Loss** helped mitigate class imbalance and improved the detection of minority malignant cases.

Model 2: Fine-Tuned EfficientNet-B0 Architecture

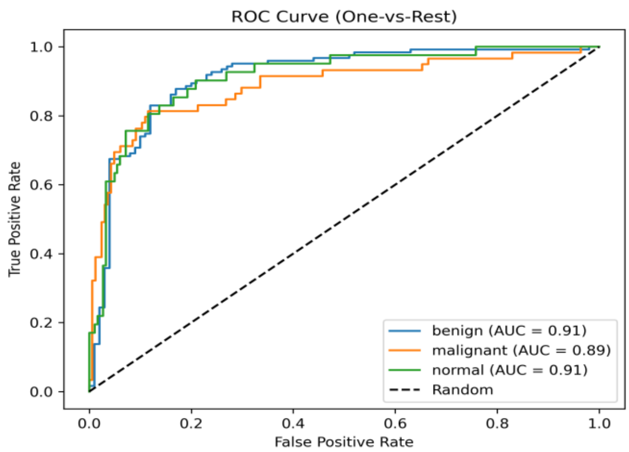

Figure 4: *ROC Curves of the EfficientNetB0 Model*

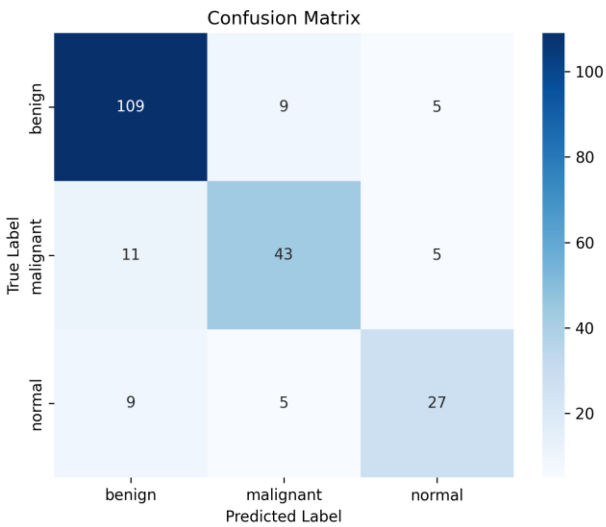

Figure 5: *Confusion Matrix of the EfficientNetB0 Model*

Received: July 20, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The ROC curves (Figure 4) represent the one-vs-rest classification performance for each tumour class. The model attained **AUC scores of 0.91 (benign)**, **0.89 (malignant)**, and **0.91 (normal)**, confirming its high diagnostic accuracy and strong separability across categories. The slightly lower AUC for malignant tumours indicates potential for further refinement through targeted data augmentation or class-weight optimization. Overall, the high AUC values validate **EfficientNetB0** as a lightweight yet highly effective model for ovarian tumour classification.

The confusion matrix (Figure 5) illustrates the classification performance of the proposed **EfficientNetB0** architecture on the ovarian tumour dataset. The model accurately classified 109 benign, 43 malignant, and 27 normal samples, demonstrating robust discriminative capability across tumour categories. While minor misclassifications occurred between benign and malignant cases, the overall distribution shows a strong predictive performance. The balanced learning achieved through EfficientNetB0's compound scaling and efficient feature extraction contributes to its reliable generalization on unseen data.

Model 3: VGG16 with Fully Connected Neural Network

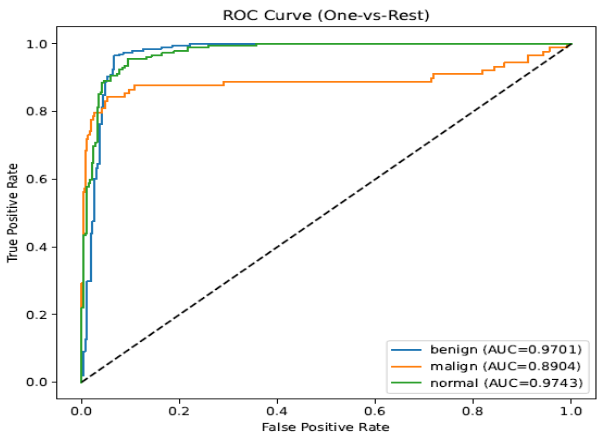

Figure 6: *ROC Curves of the VGG16 + FCNN Model*

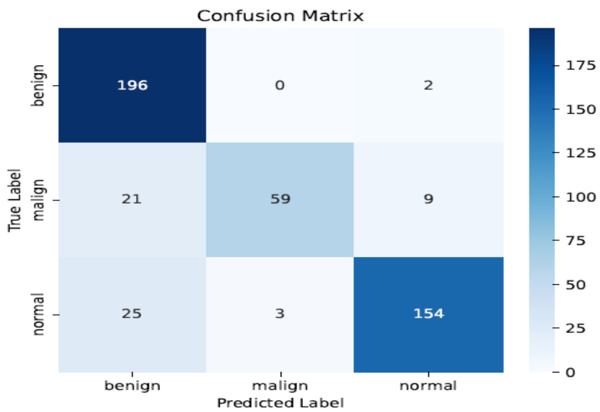

Figure 7: *Confusion Matrix of the VGG16 + FCNN Model*

Received: July 20, 2025

The ROC curves (Figure 6) illustrate the one-vs-rest classification performance for each tumour type. The model achieved **AUC scores of 0.9701 (benign)**, **0.8904 (malignant)**, and **0.9743 (normal)**, confirming a strong overall diagnostic capability. The slightly lower AUC for malignant cases suggests scope for further optimization, while the overall high AUC values validate the model's effectiveness in reliable ovarian cancer classification.

The confusion matrix (Figure 7) depicts the classification performance of the proposed VGG16 + FCNN architecture on the ovarian tumour dataset. The model correctly identified **196 benign**, **59 malignant**, and **154 normal** samples, with relatively few misclassifications across categories. This indicates that the integration of fully connected layers with the fine-tuned VGG16 backbone enhances feature generalization and improves tumour discrimination accuracy.

Table 1 compares the performance of three deep learning models—ResNet50, EfficientNet-B0, and VGG16 + FCNN—by presenting their accuracy rates in ovarian tumour classification. The ResNet50 model was the most accurate with 93.76% and the best overall model, which was attributed to the employed residual connections that helped with feature extraction, minimized the issue of vanishing gradients, and, thus, contributed to great generalization. The VGG16 + FCNN model got 87.21% accuracy, and it was characterized by high recall for benign tumours and even performance across classes, which was the result of the incorporation of fully connected layers that helped in feature abstraction and discrimination. The EfficientNet-B0 model came in with 80.00% accuracy, which was granted by its excellence in computational efficiency and a small footprint; thus, it could be hardly seen as less precise and less recall but still a candidate for real-time or resource-limited scenarios. Ultimately, the best diagnostic performance was that of ResNet50, whereas the EfficientNet-B0 model provided a perfect balance between accuracy and efficiency.

Model	Class	Precision	Recall	F1-Score	Accuracy (%)
ResNet50	Benign	0.89	0.97	0.93	**93.76**
	Malignant	0.98	0.83	0.90	
	Normal	0.97	0.92	0.95	
EfficientNet-B0	Benign	0.84	0.89	0.87	**80.00**
	Malignant	0.75	0.73	0.74	
	Normal	0.73	0.66	0.69	
VGG16 + FCNN	Benign	0.80	0.98	0.89	**87.21**
	Malignant	0.95	0.66	0.78	
	Normal	0.93	0.84	0.88	

Table 1: *Performance comparison of deep learning models for ovarian tumour classification based on precision, recall, F1-score, and overall accuracy.*

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

VI. CONCLUSION

The main intended objective of this research was to develop a highly precise classification of ovarian tumours by means of an ensemble framework of deep CNNs, which rely on transfer learning, applied to the ultrasound images. The proposed approach included three optimized models - Fine-Tuned ResNet50, VGG16 with Fully Connected Neural Network (FCNN), and EfficientNet-B0 - that were able to perfectly recognize and assess the unique spatial and textural characteristics of ovarian lesions.

To be more precise, the experiments pointed out that the peak accuracy of the ResNet50 model which was 93.76% could be considered as a strong indication of its remarkable feature extraction and generalization capabilities. On the other hand, the VGG16 + FCNN model scored 87.21% in accuracy, which is a sign of providing a dependable and stable performance, whereas the EfficientNet-B0 model, having an accuracy of 80.00%, was very efficient in computational terms and thus, ready for real-time deployment. These findings suggest that CNN architectures driven by transfer learning can provide accurate and consistent tumour classification even when the medical imaging data is limited.

To sum it up, the proposed framework can be considered a useful tool in computer-aided ovarian cancer diagnosis, and the clinicians would get an efficient and trustworthy decision-support tool. In the future, the research will be directed toward enlarging the dataset, using attention mechanisms for better feature localization, and applying explainable AI (XAI) methods for increased transparency and clinical interpretability.

VII. EFERENCES

[1] S. Timmerman *et al.*, “Ultrasound image analysis using deep neural networks to discriminate benign and malignant ovarian tumours — a comparison to subjective expert assessment,” *Ultrasound in Obstetrics & Gynaecology*, 2021.

[2] A. Sharma *et al.*, “Prediction of ovarian cancer using artificial intelligence tools,” *Health Science Reports*, 2024.

[3] R. Kumar and P. Singh, “A deep learning approach for ovarian cancer detection and classification based on fuzzy deep learning,” *Scientific Reports*, Nov. 2024.

[4] L. Green and A. Patel, “Literary review of artificial intelligence in ovarian cancer: transforming diagnosis, treatment, and future advancements,” *Ox Journal*, 2024.

[5] S. Banerjee *et al.*, “Classification and detection of ovarian cysts in ultrasound images,” in *Proc. IEEE Int. Conf. on Emerging Innovations (ICEI)*, 2017.

[6] P. Reddy and D. Mehta, “Innovative approach towards early prediction of ovarian cancer: machine learning-enabled XAI techniques,” *Heliyon*, 2024.

[7] J. Chen *et al.*, “Artificial intelligence in ultrasound diagnoses of ovarian cancer: a systematic review and meta-analysis,” *Cancers*, 2024.

International Journal of Applied Mathematics
Volume 38 No. 10s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[8] M. Wilson and S. Gupta, “Early diagnosis of ovarian cancer: a comprehensive review of the advances, challenges, and future directions,” *Diagnostics*, 2025.

[9] V. Das and R. Nair, “Classification of ovarian cancer with multimodal data using AI technology – a review,” in *Proc. IEEE Int. Conf. on Intelligent and Innovative Computing (I4C)*, 2024.

[10] K. Thomas and J. Lee, “Enhanced preprocessing and fast PCA feature extraction algorithm for ultrasound images of ovarian diseases,” in *Proc. IEEE Int. Conf. on Signal Processing (ICSP)*, 2024.

[11] N. Roy and H. Zhang, “Enhancing ovarian cancer detection: a deep learning approach with MobileNetV3 and ResNet50,” in *Proc. IEEE Int. Conf. on Intelligent Image Processing (ICIIP)*, 2023.

[12] F. Alvarez *et al.*, “Early-stage detection of ovarian cancer based on clinical data using machine learning approaches,” *Journal of Personalized Medicine*, 2022.

[13] D. Kapoor and A. Hussain, “Optimizing SE-Net for accurate detection of benign and malignant ovarian cancer,” in *Proc. IEEE Int. Conf. on Emerging Applications of Robotics and Systems (ICEARS)*, 2025.

[14] Q. Zhao, S. Lyu, W. Bai, L. Cai, B. Liu, G. Cheng, M. Wu, X. Sang, M. Yang, and L. Chen, “MMOTU: A Multi-Modality Ovarian Tumor Ultrasound Image Dataset for Unsupervised Cross-Domain Semantic Segmentation,” *arXiv preprint arXiv:2307.06799v4 [cs.CV]*, 2023.