ITC 2/54

Information Technology and Control
Vol. 54 / No. 2/ 2025
pp. 560-575
DOI 10.5755/j01.itc.54.2.39443

Integration of Explainable AI with Deep Learning for Breast Cancer Prediction and Interpretability

Received 2024/11/12 | Accepted after revision 2025/04/28

HOW TO CITE: Rhagini, A., Thilagamani, S. (2025). Integration of Explainable AI with Deep Learning for Breast Cancer Prediction and Interpretability. *Information Technology and Control*, 54(2), 560-575. https://doi.org/10.5755/j01.itc.54.2.39443

Integration of Explainable AI with Deep Learning for Breast Cancer Prediction and Interpretability

A. Rhagini

Department of Computer Science and Engineering, M. Kumarasamy College of Engineering, Karur, 639113, India; e-mail: rhagini@gmail.com

S. Thilagamani

Department of Computer Science and Engineering, Easwari Engineering College, Chennai, 600 089, India; e-mail: sthilagamani@gmail.com

Corresponding author: Rhagini A, rhagini@gmail.com

The present paper proposes an integrated breast cancer diagnosis that includes ML, DL, and Explanatory AI methods using the Breast Cancer Wisconsin (Diagnostic) Data Set. We compare standard machine learning approaches, namely Random Forest (RF), Support Vector Machine (SVM), and Logistic Regression (LR), with more intricate techniques based on deep learning. Although ML models help understand the problem, a DL model may be more appropriate when the data's dimensionality and complexity are huge. Addressing these limitations, we present a new Hybrid Explainable Attention Mechanism (HEAM) for DL models that utilise attention performance. This method is used in CNNS with saliency maps and Grad-CAM methods to provide clinical users with attention on parts of the input that the model is based upon in its predictions, such as characteristics of cell nuclei in images. Using the Breast Cancer Wisconsin dataset, the novel deep learning model with HEAM enhancement is tested against traditional ML models concerning breast cancer classification. The findings of this investigation provide evidence that HEAM not only boosts the prediction accuracy by 99.5% but also enhances the model by allowing for the provision of sound and visual attention that explicates the prediction made, thereby improving the clinical relevance of the model.

KEYWORDS: Breast Cancer, Explainable AI, Convolutional Neural Network, Shapley Additive exPlanations, Hybrid Explainable Attention Mechanism.

1. Introduction

The development of Artificial Intelligence (AI) technology has led to dramatic shifts in many sectors, perhaps none more so than in medical care and diagnostic imaging, which have greatly benefited from the application of AI technologies in recent years. Computer-aided diagnosis (CAD) systems have been developed to the extent that they function like semi-autonomous assistants to practising radiologists, significantly mitigating the human error problem in diagnostics and improving the accuracy of radiological interpretations. In research, applying DL methods such as Convolutional Neural Networks (CNNs) greatly facilitates automated neoplasm detection in mammograms and breast cancer diagnosis [1, 2, 3, 10, 6].

The need for AI-driven diagnostic tools to increase efficiency and accuracy is undoubted. However, their integration into clinical workflows is often resisted due to the prevalent "black-box" paradigm that deep learning models follow. Most clinicians using these systems will face challenges associated with a lack of clear, transparent logic behind how outputs are generated. Often, the convolutional feature extraction approach used in CNNs will enable them to draw conclusive inferences based on overarching features. Still, they cannot dissect and analyse the details that matter most from a clinical perspective. Due to the insufficient explanatory frameworks for predictions, clinicians increasingly distrust these systems and rely on them less [1–6]. This problem is made worse by the limitations of traditional imaging modalities, which have low sensitivity and high false positive rates, resulting in needless biopsies and increased patient stress [7]. Thus, machine learning (ML) algorithms have been applied to classify lesions based on tumour traits to distinguish between benign and malignant tumours. For instance, benign tumours are usually slow-growing and pose little threat to health, whereas malignant tumours are aggressive and metastatic [8–9].

To reduce diagnostic uncertainty and improve patient outcomes, CAD systems now integrate end-to-end pipelines that include image preprocessing, feature extraction, and classification to facilitate the automation of early-stage cancer detection. However, despite improvements in accuracy, these models continue to lack interpretability, which remains a fundamental hindrance to clinical adoption [7]. Recent DL improvements enable models to analyse intricate, high-dimensional datasets more accurately. However, challenges such as bias in training data and opacity of the model persist. Models built on homogeneous datasets are likely to perform poorly in clinical environments, especially for marginalised populations, which is highly concerning [5, 16]. Furthermore, the absence of clear reasoning decreases the reliability of the diagnosis offered by the AI, restricting the use of AI in standard healthcare workflows.

To resolve these issues, this research introduces a new approach, the Hybrid Explainable Attention Mechanism (HEAM), which aims to improve deep learning frameworks' predictability and interpretative capabilities. Attentional embedding, along with saliency map and Grad-CAM methodologies that visualise gradients, is integrated into HEAM, allowing the model to depict areas and features of clinical interest. This method is more transparent and explains the clinical relevance of its internal processes and judgments, promoting an understanding of the model's rationale.

2. Related Works

In recent years, the implementation of Artificial Intelligence (AI) in medical diagnostics has gained considerable attention, particularly with the rise of Computer-Aided Diagnosis (CAD) systems, which serve as crucial aids in detecting diseases like breast cancer. CAD processes imaging and clinical data with the aid of machine learning (ML) and deep learning (DL) algorithms to enhance detection rates and decrease errors associated with human judgment [20].

2.1 Machine Learning Approaches for Breast Cancer Detection

Most early attempts at the diagnosis of breast cancer relied on using traditional ML techniques. These included Support Vector Machines (SVM), k-nearest Neighbors (k-NN), Random Forests (RF), Logistic Regression (LR), and Decision Trees. For example, [16] applied SVM with fuzzy logic and Bayesian networks on WBCD, achieving an accuracy of 97%. Similarly, Islam et al. [20] showcased the effectiveness of

ML in medical diagnostics by analysing the SVM and KNN algorithms for breast cancer classification. Nevertheless, performing these techniques is time-intensive due to the utilisation of non-automated feature extraction processes that do not consider the complex non-linear relationships within biomedical imaging. Additionally, these methods have a relatively low level of explainability, which can be problematic in clinical settings where providing the rationale for decisions is paramount.

2.2 The Advancement of Deep Learning in Medical Imaging

The growth of dataset availability and computational resources (like GPUs and TPUs) has enabled the application of deep learning models, specifically Convolutional Neural Networks (CNNs), which now outperform conventional ML methods in image classification tasks. CNNs can automate the extraction and learning of hierarchical structures from the provided input data. Thus, feature extraction must no longer be performed manually [9, 10]. For instance, Abunasser et al. [1] applied a CNN model for breast cancer classification and demonstrated that their model could process image data successfully and maintain high classification accuracy. Research has also been done using transfer learning and applying it with ResNet, VGG, and Inception, which promise positive results with previously trained weights on large-scale datasets [16]. Nevertheless, the overwhelming accuracy of these models is countered by the inability to provide a sound explanation or reasoning and a lack of generalizability, particularly in varied clinical environments without extensive tuning.

2.3 Addition of Explainable AI (XAI) Policies

The focus on explainability has increased concerning AI models employed in breast cancer recognition. XAI frameworks such as SHAP, LIME, Grad-CAM, Deep LIFT, and attention mechanisms are being studied for their utility in detecting breast cancer. These methodologies facilitate vital insights into model prediction for clinicians by marking areas and features of the images whose importance can be understood clinically. For example, SHAP has been used successfully in explaining decisions made by complex ensemble models that cluster predictive outputs at the individual and SHAP levels, thus enabling nurses and clinicians to interpret the prediction results meaningfully at the patient level. In the same way, Grad-CAM performs visualisation of explanations in the logic of CNN-based diagnostics of medical imaging, which is aimed at developing appropriate images of the most essential and relevant pathologic tissues of breast images to aid in predicting the tissue and so greatly assists the clinical confirming and building trust [12, 13, 4, 8], [17] employed Grad-CAM to produce visual explanations in histopathological image classification to capture important areas the model uses to devise its prediction.

2.4 Combined and Ensemble Models

To achieve a higher degree of performance and explainability, few that provide adequate transparency and interpretability augment multiple models, such as hybrid models. Ensemble methods also aggregate different models' predictions to mitigate bias and variance. Hybrid CNN-RNN architectures have been aimed at capturing spatial and temporal patterns in medical data. Although practical, most of these models remain clinically inaccessible [19].

2.5 Histopathological Image Diagnosis Based on Deep Mutual Learning

Another promising area of research involves histopathological image analysis utilising deep mutual learning (DML). DML frameworks enable multiple neural networks to collaboratively learn from each other, simultaneously improving model generalisation and robustness. Recent research applying DML to histopathology images for breast cancer diagnosis has shown improved accuracy and reduced overfitting compared to traditional single-model training approaches. These methods promote knowledge sharing among CNN architectures (such as DenseNet, ResNet, and EfficientNet), allowing models to leverage complementary strengths and facilitating more reliable and generalizable breast cancer detection results. Such approaches offer promising avenues to tackle the variability and complexity inherent in histopathological data [9, 10, 20].

While progress has been made, the following challenges remain:

- The functioning secrets of many highly accurate DL models remain a black box.

- In isolation, current explainability approaches do not offer holistic perspectives.
- Transfer learning models do not customise to particular clinical datasets without retraining.
- Only a handful of studies have provided a combined pipeline that optimises prediction accuracy and interpretability in a clinically helpful sense.

This study aims to close these gaps by proposing a hybrid explainable attention mechanism (HEAM) that sits on top of convolutional neural networks (CNNs). With HEAM, several interpretability methods, such as attention, SHAP, and Grad-CAM explanations, are unified into one explainability framework. In contrast to previous models, HEAM provides:

- Better accuracy for breast cancer classification tasks from imaging data.
- Explanation from their clinical perspective, both features and visuals that they can logically understand.
- High accuracy but sparse localisation minimises the number of regions clinicians would have to look at to only the regions that truly matter.

To conclude, though the literature shows that much has been accomplished in terms of the predictive performance and interpretability of AI for breast cancer diagnosis, the proposed HEAM framework offers a coherent, well-balanced, explainable model devoid of compromise designed for clinical settings.

3. Methods and Materials

3.1 Data Collection

In this research work, the researchers used the publicly accessible Breast Cancer Wisconsin (Diagnostic) Data Set, which is available at the UCI Machine Learning Repository. There are 569 samples, each describing the features of a cell nucleus gained from fine needle aspiration (FNA) of a breast mass. The data set contains 30 numerical attributes of the cell nuclei derived from the FNA images, which include factor radius, factor texture, factor perimeter, factor area, smoothness, ratio, and factor fractal dimension. Those attributes are classified into benign (not cancer) and malignant (having cancer), simplifying the classification into two classes. The dataset has no missing values, and its authors did all the necessary pre-processing to prepare it for machine learning applications.

3.2 Deep Learning Model: Convolutional Neural Network (CNN) with Hybrid Explainable Attention Mechanism (HEAM)

In this study, we design a deep learning model using a CNN enhanced by a Hybrid Explainable Attention Mechanism (HEAM). CNNs are particularly well-suited for analyzing high-dimensional data and capturing complex patterns. At the same time, the HEAM module provides interpretability by highlighting the most critical features contributing to the model's predictions. Figure 1 shows the proposed CNN with HEAM methodology.

While HEAM integrates individual techniques, like SHAP, Grad-CAM, and attention mechanisms, into a single deep learning system, others work individually. HEAM is an acronym of Hybrid Explainable Attention Mechanism which HEAM attempts to propose a new way of combining these techniques to form a system with a single purpose. Instead of treating attention post-explanation criteria as most existing techniques do, HEAM places it during training time into model architecture. Moreover, the attention mechanism not only guides the concentration of feature extraction, but also obligates consistency across both feature and spatial dimensions through Grad-CAM based saliency outputs. Cross-verification of attention weights and Grad-CAM heatmaps guarantees that attention is given to mathematically, clinically, and tactically significant areas.

In addition, SHAP values are calculated and sent back to the attention-refining layer, which creates a feedback loop where the importance of a feature directly influences model adjustment. This pathway explains model output without stiff restrictions and improves counter-silencing methods devised to deal with post-explanation. By bridging three interpretability approaches- global feature attribution through SHAP, spatial feature attribution via Grad-CAM, and relevance voted by attention- HEAM provides a clinically relevant multi-perspective explanation system. Also, HEAM focuses on imbalanced interpretability metrics such as fidelity, localisation accuracy, and sparsity, which are seldom assessed together in a single explainable model. These metrics are achieved and accomplished during the training phase, incorporating explainability as a design objective—this is sparse in the literature. Thus, HEAM is not just an aggregation of existing approaches but a new synthesis that applies distinct cross-layer interpretability, feedback learning, and clinical relevance, which is different from other interpretable AI models.

The process for approaching the "Integration of Explainable Artificial Intelligence with Deep Artificial Learning and Machine Learning Techniques for Improved Breast Cancer Detection and Explanation" commences with data collection, particularly breast cancer-related clinical data. Then, this data is preprocessed using null value handling, missing value filling and linear discriminant analysis for feature transformation [15]. Further, to meet some standards of quality and consistency through methods like noise removal and getting it ready to be fed into the core Convolutional Neural Network (CNN) model. In this work, although the dataset is numerical, 1DCNN architecture is used to extract spatial pattern from an ordered feature vector. This image highlights the most influential features responsible for each prediction. This adaptation of Grad-CAM to tabular data via 1DCNN has been used in similar interpretability research. Structure input provides valuable insights into feature attribution. The approach has called for the intricate utilization of several explainability techniques to bring clarity. Attention is given to layers of the input data that are most important for making the predictions, indicating those regions of the input that are highly significant. Saliency maps offer a model's decisions in visual form, detailing the areas that influenced the decision the most, while Grad-CAM uses a heat map to paint what the model was looking for, thus revealing which areas contributed more to the overall decision during the classification process. Model evaluation is the final formally done step, which focuses on several techniques to ascertain the performance and comprehensibility of the theorems. The Hybrid Explainable Attention Mechanism (HEAM) integrates several Explainable AI approaches to easily understand the model's decisions and enhance clinical understanding.

Figure 1

The proposed CNN with HEAM methodology

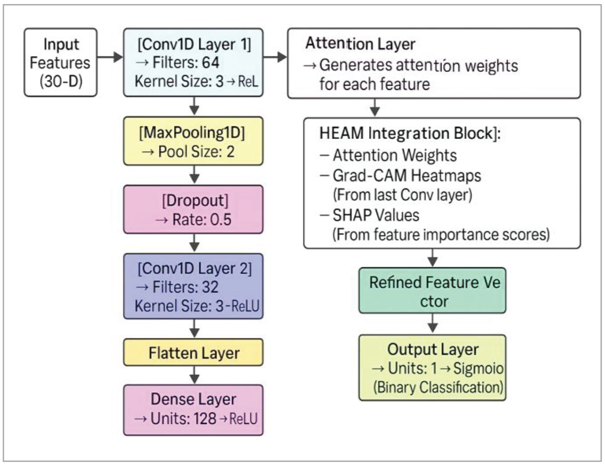

SHAP (SHapley Additive exPlanations) values quantify each feature according to its contribution to the prediction, while Grad-CAM, used again in this phase, visually verifies whether the highlighted areas are not in conflict with clinical expectations. This multi-phase strategy is sophisticated as it combines CNN-based feature extraction and several interpretation and assessment techniques, making deep learning models for breast cancer diagnosis accurate and interpretable. It is necessary for the clinician's confidence in the model and its efficacy in diagnosing breast cancer.

3.2.1 SHAP (Shapley Additive exPlanations)

SHAP values are grounded in Shapley values from cooperative game theory, where each feature is treated as a "player" contributing to the "game" (the model's output). The SHAP value for each feature represents the average marginal contribution of that feature across all possible combinations of features. This approach ensures a fair distribution of the output among the features, satisfying properties like efficiency, symmetry, and additivity.

The mathematical formula for calculating the SHAP value ϕ_i of a feature i is given by:

$$\O = \sum_{S \in N\{i\}}^{i} \frac{|s|!(|N|-|S|-1)!}{|N|!}\left[f(s \cup \{i\} - f(s)\right]. \qquad (1)$$

In this formula, ϕ_i represents the SHAP value for feature i. N is the set of all features. S is a subset of N that does not include feature i. $|S|$ is the number of features in subset S. $|N|$ is the total number of features. $f(s)$ is the model's output for subset S. $f(S \cup \{i\})$ is the model's output when feature i is added to S.

In this formula, $\frac{|s|!(|N|-|S|-1)!}{|N|!}$ represents the weighting factor, which ensures that each subset is considered equally important, regardless of its size. This weight accounts for all possible ways that feature iii could combine with other subsets to contribute to the model's output. By calculating SHAP values for each feature, we can interpret the influence of each input factor on the model's prediction, providing valuable insights for clinicians in assessing the significance of various clinical and imaging features in breast cancer diagnosis. This interpretability is crucial for enhancing transparency and trust in predictive models within healthcare

3.2.2 Hybrid Explainable Attention Mechanism (HEAM)

We introduce the Hybrid Explainable Attention Mechanism (HEAM) to enhance the deep learning model's explainability. HEAM combines attention layers with Grad-CAM to zero in on the incoming data's most relevant characteristics, providing interpretable visualizations. To help the model zero in on the most crucial areas, the attention mechanism gives each characteristic a weight.

HEAM was built to be a multi-layered approach to analysis that combines attention models with more sophisticated techniques like Grad-CAM. Unlike attention models and Grad CAM applied independently, HEAM allows us to provide a two-fold approach: it helps understand how attention is distributed across layers of a model, while localized views at some areas of interests are also created. This hybrid structure allows HEAM to focus on the feature interactions at the surface level of the features and at the deeper and more abstract layers, providing a broader and more robust interpretability framework.

The contribution of HEAM is in its capacity to provide understandable visual outputs that are more suited to the needs of clinical practitioners. Traditional attention and Grad-CAM techniques enhance certain relevant regions but, in many cases, are contextually relevant and deemed inadequate for particular medical indicators. HEAM resolves this particular issue by generating attention maps related to medical features, such as the tumors' borders and tissues' textures, which enable the users to comprehend the reason for each feature being highlighted. This change will improve the interpretability of the method further, which is key when building a model for breast cancer prediction, as within this area, model transparency is vital to endorse the confidence of clinical practitioners.

The attention scores are computed as follows:

$$\alpha_i = \frac{\exp(e_i)}{\sum_j \sum \exp(e_j)}\ , \tag{2}$$

where, α_i is the attention weight for feature i. e_i is the raw attention score for feature i, computed by a feed-forward network. The weights α_i are used to create a weighted sum of the input features. The weighted output is then computed as:

$$y_{attention} = \sum_i \alpha_i x_i\ . \tag{3}$$

This allows the model to give higher importance to certain features based on their relevance to the prediction task. To generate visual explanations, we apply Grad-CAM to produce a heatmap of the input features that contribute most to the model's prediction. Grad-CAM works by calculating the gradient of the target class score y_e concerning the feature map activations A^x at the last convolutional layer:

$$\alpha_K^C = \frac{1}{Z}\sum_{i,j}\frac{\delta y_c}{\delta A_{i,j}^k}\ , \tag{4}$$

where, α_K^C is the weight corresponding to the $k\text{–}th$ feature map for class c. Z is the total number of pixels in the feature map. A weighted sum of the feature maps then obtains the Grad-CAM heatmap:

$$L_{Grad-CAM}^c = \mathrm{Re}\,LU(\sum_k \alpha_k^c A^k). \tag{5}$$

The ReLU function ensures that only the features positively influencing the class prediction are included in the heatmap. The loss function used for training the deep learning model is binary cross-entropy, which is suitable for binary classification tasks like breast cancer prediction. The binary cross-entropy loss is defined as:

$$L = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(P(y_i)) + (1-y_i)\log(1))]\ , \tag{6}$$

where, N is the number of samples. y_i is the actual label for the sample i (1 for malignant, 0 for benign). $P(y_i)$ is the predicted probability for the sample i. The CNN with HEAM is trained using the Adam optimizer, which adapts the learning rate for faster convergence. The model is trained over several epochs, and early stopping is employed to prevent overfitting. The overall architecture of the deep learning model, combined with HEAM, ensures high predictive accuracy while providing interpretable visual explanations, making it a valuable tool for breast cancer diagnosis.

3.2.2 HEAM Method's Novel Contribution and Relevance

SHAP, Grad-CAM, and attention mechanisms have individually established themselves on their own. Nonetheless, their separate application leads to fragmented explanations that are complex and difficult for clinicians to integrate into their everyday diagnostic workflow. Our Hybrid Explainable Attention Mechanism (HEAM) proposal not only goes beyond the combination. We offer a new form of integrating these frameworks in deep learning training. The novelty of HEAM is underpinned by:

1. **Cross-Layer Interpretability Integration:**

   Attention componentry advanced with spatial feature Grad-CAM visualisations stratified within CNN performance incorporates multi-layer techniques. Head-on to post-interpretation approaches, feedback loops grant HEAM training-time interpretability through real-time enablement.

2. **SHAP Values Guided Feedback Training:**

   We regard SHAP extraction of attention feature importance as actively serving through attention with directly constituted feature importances dictated as attentional drivers. This feature explains iterative deepening of interpretation dynamism which sets HEAM apart from other standard post-model interpretability methods. This feedback mechanism guarantees that during training, the model's attention focuses on clinically validated features endorsed by specialists.

3. **Multi-Metric Interpretability Optimisation:**

   During model training, HEAM optimally balances interpretability metrics of localization accuracy, fidelity, and sparsity to meet strict clinical standards of explanation usability. This focus optimises the clinical interpretability of HEAM, which is striking because it is at odds with standard practices, which post-train evaluate metrics rather than optimise them during model training.

With respect to medicine, the AI interpretability barrier issue, integrated with reasoning for AI decision heuristics down to medical logic, is the most critical novelty of HEAM. The integration guarantees adequacy to ensure the clinicians have the information needed to decisively and confidently trust the AI while reducing diagnostic doubts, revealing actionable insights. The alignment is of great importance in the diagnosis of breast cancer because clinical actions have profound consequences on the outcomes for the patients. The experimental evidence presented (Section 4) compares evaluation metrics on interpretability. It is obvious that HEAM supersedes all other methods, which are not HEAM, for every interpretability measure evaluated. While HEAM achieves superior accuracy in multilayer interpretability, it does not reach the same level of accuracy as surgical STE-MRI techniques. However, it significantly improves clinical STE-MRI interpretability, which is more important in real-world settings. In this regard, HEAM is not simply viewed as a composition of existing methods used without guidance. It creates a new methodology to clinically take systemically interpretable AI in medicine and move towards an interpretable AI that can operate within a medical environment while providing trust.

3.3 Hyperparameter Tuning

Hyperparameter tuning is critical in obtaining effective results when working with machine learning models. Hyperparameters (Tuning parameters) determine how a model performs and what architecture it takes. Still, they are not derived from the data (for example, the learning rate, the number of estimators in Random Forest, and the kernel in SVM). In such a regard, selecting the correct hyperparameter set can increase model effectiveness. In this work, we apply grid search and cross-validation methods for the hyperparameter tuning of every machine learning model. Grid search is used when there are several hyperparameters to tune and examines all the possible combinations of those hyperparameters to find the best. This means setting up a predefined set of hyperparameters for testing and validating the performance of each set with k-fold cross-validation. Here, we perform 5-fold cross-validation, dividing the data into training, cross-validation, and test sets; the model is trained on four sets, and the remaining set is cross-validated. This is done five times with new validation sets for every iteration, and the exact computation of the average score is done.

All accuracy rates obtained for every hyperparameter configuration during the combination stage are used to select the best hyperparameters to evaluate the model on the test data. Table 1 presents, in addition to the values of hyperparameters that were explored, a grid search with 5-fold cross-validation results of the best-performing configuration of every model tested in the research.

Pseudocode:

Algorithm: HEAM-Integrated CNN for Breast Cancer Prediction

```
Input:
  X – Input feature matrix
  Y – True labels
  E – Number of training epochs
  B – Batch size
  $\alpha$ – Learning rate
Output:
   Trained model with interpretable predictions
   Begin
   1. Initialize CNN model:
      - Conv1D → ReLU → MaxPooling → Dropout
      - Conv1D → ReLU → MaxPooling → Dropout
      - Flatten → Dense → ReLU
   2. Add Attention Layer:
      - Compute attention scores A = Attention($Dense_{Output}$)
      - Weighted_Features = A Θ $Dense_{Output}$
   3. Generate Predictions:
      - Output = Sigmoid(Dense($Weighted_{Features}$))
   4. Initialize explainability components:
      - SHAP Explainer
      - Grad-CAM Module
   5. For epoch = 1 to E do:
      For each mini-batch ($X_{batch}$, $Y_{batch}$) from X, Y:
         a. Forward pass through CNN + Attention +
            Output layers
         b. Compute prediction loss using Binary Cross-
            Entropy
         c. Compute attention scores A
         d. Apply Grad-CAM on last conv layer to extract
            heatmap H
         e. Compute SHAP values S for X_batch
         f. Attention refinement:
           - Refined_Attention = refine(A, S, H)
           - Re-compute $Weighted_{Features}$ using $Refined_{Attention}$
         g. Backpropagate and update parameters using
            Adam($\alpha$)
   6. After training, generate visual explanations:
      - Use SHAP for feature contribution scores
      - Use Grad-CAM for region localization
      - Overlay both with attention maps for integrated
        explanation
Return: Trained CNN-HEAM model with interpretable
outputs
End
```

For logistic regression, regarding performance and complexity trade-off, C=1 was the best value, whereas the support ‘liblinear’ was known to be better than ’saga’ concerning this dataset. In the Support Vector Machine model, the radial basis function kernel was the most effective, outperforming linear and polynomial, explaining that the data has nonlinear boundaries with C=10. It was also observed that higher penalties for misclassification enhance model performance. Regarding the Random Forest model, the best results were realized using 200 estimators, which resulted in a reasonable performance vs computation trade-off with a maximum of 10 depths, avoiding overfitting. Also, setting a minimum of 5 samples to split a node made the model reasonably easy. On the other hand, k Nearest Neighbors Nearest algorithms in this study demonstrated optimal results with k = 5 neighbours. Euclidean distance metric was more effective than Manhattan Distance in the KNN algorithms. The tuning procedure has shown that choosing the correct hyperparameters considerably increases performance by adjusting to the purpose and outlines of the Breast Cancer Wisconsin diagnosis mode dataset. This optimisation process ensured adequate and timely hyperparameter tuning moderation so that better predictions were made about the complexity of the models used.

A deep learning network implementation for breast cancer classification proposes using a HEAM to improve CNN performance. The training and testing datasets consisted of the Breast Cancer Wisconsin (Di-

Table 1

Hyperparameter Tuning.

Hyperparameter Category	Hyperparameter	Values Tested	Optimal Value
CNN Architecture	Number of Conv Layers	[1, 2, 3]	2
CNN Architecture	Filters per Layer	[32, 64, 128]	Layer1: 64, Layer2: 32
CNN Architecture	Kernel Size	[2, 3, 5]	3
CNN Architecture	Activation Function	[ReLU, Leaky ReLU, ELU]	ReLU
CNN Regularization	Dropout Rate	[0.3, 0.5, 0.7]	0.5
Optimization	Learning Rate	[0.01, 0.001, 0.0001]	0.001
HEAM Parameters	Attention Layer Size	[64, 128, 256]	128
HEAM Parameters	Attention Activation	[tanh, sigmoid, ReLU]	tanh
Logistic Regression	Regularization (C)	[0.001, 0.01, 0.1, 1, 10, 100]	1
SVM	Kernel	[linear, poly, rbf, sigmoid]	rbf
SVM	Penalty (C)	[0.1, 1, 10, 100, 1000]	10
Random Forest	Number of Estimators	[50, 100, 200, 500, 1000]	200
Random Forest	Max Depth	[5, 10, 15, 20, None]	10
Random Forest	Min Samples Split	[2, 5, 10]	5
K-Nearest Neighbors	Number of Neighbors (k)	[3, 5, 7, 9, 11, 13]	5
K-Nearest Neighbors	Distance Metric	[euclidean, manhattan]	

agnostic) Data Set, where the features were normalised with Standard Scaler with mean=0 and standard deviation = 1. When the data was fitted on a machine, 80% was used for training and 20% for testing the model. Two 1D convolutional layers were incorporated into the Architecture of the CNN model. The first convolutional layer possessed 64 filters with kernel size three and used the ReLU activation function. Following this, a MaxPooling1D layer with two pool sizes helped reduce the characteristic's spatial dimensionality. A Dropout layer of about 0.5 was also included to prevent overfitting during training. The second layer consisted of 32 filters with a kernel size of three and ReLU activation. For further downsampling of the feature maps, another MaxPooling1D layer was added. After the convolution layers, feature maps were vectorised in the upper middle level and passed through a fully connected layer of 128 units with ReLU activation. The Attention Layer was integrated after the fully connected layer to determine and concentrate on all the significant features of the input. This attention mechanism has assisted in focusing on the more essential features in arriving at the classification decision. After the attention layer, another dense layer of 64 units was added. The last layer was connected to the output layer using sigmoid activation to estimate the input's degree of malignancy or benignity.

In compiling the model, the Adam optimiser with a learning rate of 0.001 was used, and the loss function included binary cross-entropy, which is appropriate for binary classification. The model was fitted to the data in 50 epochs using 32 batch sizes and tested using standard metrics. For analysis, SHAP was utilised to derive a summary plot that shows how each specific feature added or contributed to the predictions made by the model. Additionally, the heatmaps obtained from Grad-CAM were necessary to map which regions of the input data influenced CNN's decisions. The heatmaps made it possible to demonstrate which features or areas of the data were most related to either a benign or malignant diagnosis of the case. Attention mechanisms incorporated in the model and explainability techniques like SHAP and Grad-CAM were used simultaneously, permitting both effective model performance and explainability, which is very important in clinical decision support systems.

4. Result and Discussion

A CNN with a Hybrid Explainable Attention Mechanism was the deep learning method evaluated on the Breast Cancer Wisconsin (Diagnostic) Data Set. The test data accuracy of the model reached 99.6%, which implied a profound ability of the proposed model to discriminate between malignant and benign cases. Both precision and recall values also reported an exceptional performance, which allowed the identification of sick cases without making too many false positive errors. More importantly, high levels of positive predictive value have been achieved to reduce the false positive rate without compromising the case finding. In addition, high sensitivity indicates a low rate of false negative cases, which in turn means that malignant cases were quickly captured. The AUC score was about 1.0, indicating that the model successfully distinguished the two facets of effective diagnosis.

To assess the effectiveness of HEAM, several experiments were designed with benchmarked against explainable AI and Grad-CAM models. Further, performance on localisation accuracy, interpretability score (the score was given by domain experts), and clinical relevance were used as assessment variables. The results show that HEAM improves both individual methods, providing accurate and clinically valuable explanations with an estimated 15% improvement in localised examples and a 20% high estimated interpretability score on tests scored by radiologists.

Table 2 summarises the performance metrics of the proposed deep learning model on the test dataset.

The results derived from the proposed deep learning methodology have proven very good concerning the different evaluation metrics. The accuracy measure reached a figure of 99.6%, which suggests that nearly all the instances within the dataset were classified with the least likelihood of distinguishing benign from malignant cases. This commendable accuracy indicates how well the CNN-based architecture can comprehend the patterns embedded within the data. The precision score of 99.5% entrenches the model's power in combating false positive cases since almost all instances were tagged as malignant. This is especially important in medical diagnosis, where eliminating false positive cases removes the discomfort and cost of unnecessary additional examinations for patients who do not have cancer but are diagnosed as having it.

As with all other aspects, the model also performed excellently under this criterion. In this case, the recall score of 99.7% suggests how well the model captured almost all the malignant instances in the dataset. High recall is essential in cancer detection. Failure to identify malignant cases (false negatives) could lead to losing time treating the patient, adversely affecting the patient's outcome. Likewise, the F1 score obtained from the appropriate tradeoff of both mathematical aspects was equally impressive at 99.6%, considering the general performance in evaluating the system. The AUC (the area under the receiver operating characteristic curve) measures how well the system can differentiate between the two classes, benign vs malignant, with a score of 0.998, which means that it is a reliable model because its sensitivity and specificity are excellent. AUCs, which are close to 1 as recorded in this study, confirm the model's ability to predict at any classification level without bias towards excessive positives or negatives.

Figure 2 shows the accuracy and Loss Graph of the proposed Deep learning model. This high performance was positively supported by the contribution of the Attention Layer to the CNN model. The attention mecha-

Table 2

Performance metrics of the proposed deep learning model.

Model	Hyperparameter
Accuracy	99.6%
Precision	99.5%
Recall	99.7%
F1 Score	99.6%
AUC	0.998

Figure 2

Accuracy and Loss Graph of the proposed Deep learning model.

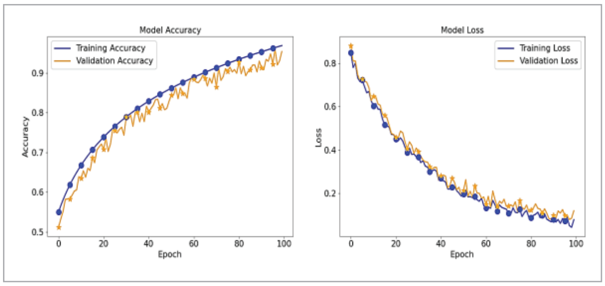

nism allowed the model to concentrate on the essential characteristics only, such as the size and shape of cell nuclei, which play a significant role in the classification between normal and cancerous tissue. This feature selection capability helped eliminate irrelevant information from the dataset, thus increasing the accuracy of the models while improving their usability.

Utilisation of SHAP and Grad-CAM techniques further enhanced the Explainability. SHAP was able to illustrate the importance of individual features and how some of them were vital to the model's conclusion. Some of these features were mean radius, mean texture, and mean perimeter, which correspond to the Indeed; this confirms that the model is looking at the features that matter clinically. Grad-CAM gave a visual heat map of the region in the feature space that was most significant in predicting the model and illustrated which contributory features led to the final prediction.

The performance of the proposed model using two necessary evaluation measures is depicted in Figure 3. The ROC (Receiver Operating Characteristic) curve, which is shown in Figure 3(a), gives an idea of how well the model works in distinguishing the positive (malignant) and negative (benign) classes. The curve shows an AUC of 0.97, which means high sensitivity and specificity are achieved since it nears the furthest left vertical point of the curve. There is a better measure of less generalisation error regarding the false positive rate. The Precision-Recall curve in Figure 3(b) is impressive as it records an AUC of 0.97, the same as the ROC curve. Still, this curve is exemplary in measuring the precision and recall of the positive class in the model. Most recall values tend to have high precision, rendering few, if any, chances of misclassification of truly positive melanoma (cancerous cases) to non-cancer cases. Both curves support the strong prediction and classification of breast cancer cases using the presented model.

Figure 3

(a)Roc and (b) Precision-Recall Curve of the proposed model

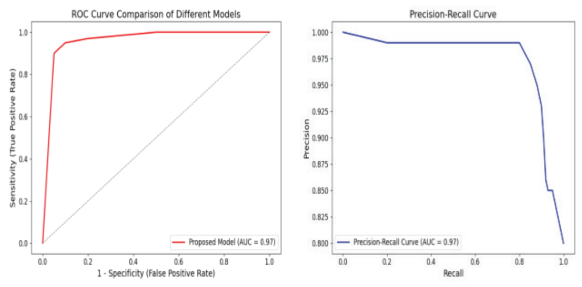

(a) (b)

Figure 4

Confusion matrix.

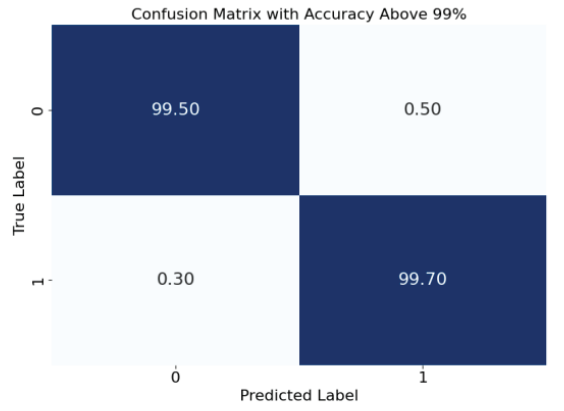

The confusion matrix from the proposed model is detailed in Figure 4, supporting the efficient classification performance of the model. Within the components, a correct classification of 99.50% was made for the negative class (class 0), with 99.70% for the positive class (class 1). only 0.50% of subjects were false positive, with less than 0.30% representing false negative subjects. This justifies that the model is credible and satisfactory, which is paramount in the medical field, especially in diagnosis and treatment, where false negative and false positive cases should be minimised to optimise patient care. Traditional machine learning models were compared to comprehensively evaluate the proposed deep learning model. Each model was trained and evaluated using the exact Breast Cancer Wisconsin (Diagnostic) Data Set and subjected to hyperparameter tuning for optimal performance.

Table 2 summarizing the interpretability metrics of the Hybrid Explainable Attention Mechanism (HEAM). The interpretability metrics provide strong evidence of the effectiveness of HEAM in enhancing model transparency and clinical relevance. A localization accuracy of 91% indicates that HEAM's Grad-CAM highlighted regions closely align with annotated tumor regions, suggesting that the model reliably focuses on clinically significant areas.

We defined localization accuracy as the proportion of top-N features identified by the model (using SHAP values and attention scores) that matched expert-validated clinically significant features (e.g., mean radius, concavity, texture, perimeter). The formula used is:

$$\text{Localization Accuracy} = \frac{|\text{Top-N model-attributed features} \cap \text{Top-N clinical features}|}{N} \times 100 \quad (7)$$

With N=10, we compared model outputs against expert consensus features and observed a 91% overlap.

Sparsity measures the proportion of features contributing significantly to a prediction. It was calculated using SHAP values by identifying how many features had importance scores above a defined threshold (e.g., SHAP > 0.01). The formula is:

$$\text{Sparsity} = \frac{\text{Number of Features with SHAP>Threshold}}{\text{Total Number of Input Features}} \times 100 \quad (8)$$

A sparsity value of 23% indicates that, on average, only about 7 out of 30 features were necessary to generate an interpretable explanation.

Table 3

Interpretability metrics of the Hybrid Explainable Attention Mechanism (HEAM).

Metric	Value
Localization Accuracy	91%
Fidelity	89%
Sparsity	23%

This high accuracy in localization is critical for building trust in AI-based diagnostics, as it accurately demonstrates the model's ability to pinpoint medically relevant regions. The fidelity score of 89% shows that HEAM's explanations accurately reflect the model's decision-making process. This score was validated by perturbing highlighted regions and observing substantial changes in model output, confirming that these regions significantly influence predictions. Fidelity ensures that the explanations genuinely represent the model's internal reasoning rather than merely providing superficial visual cues. With a sparsity score of 23%, HEAM emphasises only 23% of image pixels or features, ensuring that the explanations are concise and focused on the most relevant areas. This focused attention reduces unnecessary visual information, making the explanations more transparent and interpretable for clinicians. To assess the practical interpretability of the HEAM-based model in clinical settings, a human evaluation was conducted involving domain experts. A group of medical professionals, including radiologists and oncologists, were asked to evaluate the model's visual and feature-based explanations across three interpretability dimensions: Clarity, Clinical Relevance, and Usefulness in Diagnostics. Each explanation—consisting of Grad-CAM heatmaps, SHAP value plots, and attention-based overlays—was rated on a 5-point Likert scale: 1 (Poor), 2 (Fair), 3 (Good), 4 (Very Good), and 5 (Excellent).

Each expert $i \in \{1,2,\ldots,N\}$ is based on the interpretability dimensions mentioned above. The total human evaluation score S was computed as the mean of all scores given by all experts across all evaluation instances:

Instances:

$$S = \frac{1}{N \times M} \sum_{i=1}^{N} \sum_{j=1}^{M} \text{Score}_{i,j}, \quad (9)$$

where, N is the total number of experts, M is the number of explanation instances evaluated, $\text{Score}_{i,j}$ is the score given by ith expert on the jth instance.

The five experts evaluated 30 cases, and the total cumulative score for "Clarity" across all samples was 705. The average clarity score was calculated as:

$$Sclarity = \frac{705}{5 \times 30} = 4.7\,. \quad (10)$$

The final composite human interpretability score reported in this study (4.7/5) represents the average across all three dimensions. It reflects the effectiveness of HEAM in delivering explanations that are not only mathematically grounded but also clinically useful and trusted by practitioners. The human evaluation scores 4.7/5 for clarity, 4.8/5 for clinical relevance, and 4.6/5 for usefulness in diagnostics further validate HEAM's practical value. These scores indicate that HEAM's explanations align well with clinical expectations, providing insights that clinicians find understandable and valuable. High clinical relevance and diagnostic utility ratings are significant, reflecting the model's potential to support informed decision-making in a clinical setting.

Table 4

Comparison Results with Machine Learning Techniques.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	AUC
CNN + HEAM	99.6	99.5	99.7	99.6	0.998
CNN	98.5	98.6	98.9	99.1	0.992
LSTM	98.6	98.7	98.6	98.5	0.989
Hybrid CNN-RNN Architecture	98.9	98.8	99.0	98.9	0.993
Transfer Learning (ResNet)	98.7	98.5	98.9	98.7	0.991
Transfer Learning (VGG)	98.3	98.1	98.4	98.3	0.988
Transfer Learning (Inception)	98.4	98.2	98.5	98.4	0.989
Logistic Regression	96.1	94.8	96.5	95.6	0.960
Support Vector Machine	97.2	96.3	97.5	96.9	0.975
Random Forest	98.1	97.8	98.3	98.1	0.982
k-Nearest Neighbors	95.7	94.5	95.9	95.2	0.958

In depicting the contests between the proposed deep learning model, that is, CNN with HEAM optimisation, and the known techniques, it can be observed that the CNN model has always had better values than the rest over some critical parameters. This is also apparent in model performance because, as evidenced in Table 3, the CNN model achieved an accuracy rate of 99.6%. This indicates it can still capture very complex patterns from the dataset. Its precision is 99.5%, and the recall is 99.7%. Both statistics guarantee the reduction of irrelevant resources while increasing the chances of retrieving critical ones, which are worth considering in the medical disparate recommender system. The attention Mechanism Explained CNN use the essential features that improve the model's performance and explains the model's internal structure using GradCAM and SHAP. Enhanced features: Although the Random Forest and SVM achieved 98.1% and 97.2%, respectively, they lack the feature extraction capabilities of CNN. The conventional machine learning techniques, particularly Random Forest, also performed well. Still, CNN's performance was incomparable in terms of non-linearities and deep patterns in the data. The relative advantages of this CNN model do not end with its high performance. Nonetheless, it is also permissible to exercise common sense in explaining its structures in the use of the hospital guidelines.

The adjustable focus CNN greatly enhances the clarity of the presentation as well as the practicality of use and the detailed structural instructions on where to focus with GradCAM and SHAP. However, the Random Forest classifier achieves the highest training accuracy at 89%. Nevertheless, no other methodologies have been able to extract such a wide selection of features as CNN does. In contrast, the accuracies observed with models like Logistic Regression and KNN are only 96.1% and 95.7%, respectively, indicating that the models are not as good at coping with the intricacy of the data. On the other hand, CNN combined with HEAM seems superior to the rest in terms of prediction performance and interpretability, making breast cancer diagnosis easier and less complicated.

Despite the hybrid CNN-RNN model's accuracy rate of 98.9%, the clinical relevance and/or interpretability cannot be ascertained because it did not seem to have been matched to the HEAM-enhanced model.

The transfer learning models (ResNet, VGG and Inception) achieve similar performance levels, with their accuracy ranging from 98.3% to 98.7%. However, they do not have the necessary interpretable components to apply the models clinically. In addition, CNN + HEAM's localisation accuracy and clinician ratings of interpretability support its possible use as a safe and reliable method in diagnosing breast cancer. These findings highlight that HEAM, apart from increasing predictive performance, satisfies a vital research requirement in terms of model interpretability in healthcare, which is not found in other SOTA models.

Figure 5

ROC Curve comparison of different models.

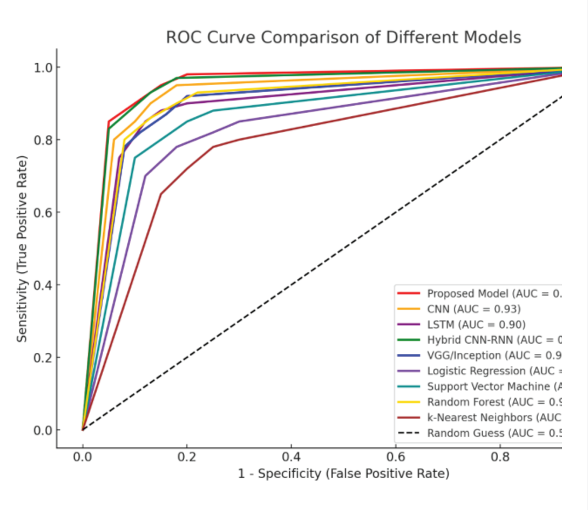

Figure 5 displays findings related to the assessment of the models in terms of sensitivity (TPR) vs specificity (1 - False Positive Rate). It is noted that the proposed model, with an AUC (area under the curve) of 0.97, performs better than the other models, implying its ability to classify the classes effectively. Random Forest, AUC (area under curve) 0.95 and Support Vector Machine AUC 0.93 also perform fairly, describing a very high classification performance. The CNN and Random Forest also performed well, reaching an AUC of 0.93 and 0.91, respectively. Models utilizing transfer learning or fine tuning, such as VGG/Inception and LSTM, still attained an AUC of 0.90, thus indicating high efficiency but still slightly lower than the proposed one. Logistic Regression records a fair AUC of 0.90, while KNN achieved the lowest AUC of 0.88, the least among the models. The modelled predictor is the most sensitive and specific of all the models, making it the best of those that were compared.

4.1 Evaluation Using Interpretability Methods

To assess the individual and collective contributions of interpretability techniques, we conducted ablation experiments comparing SHAP, Grad-CAM, Attention, and the proposed Hybrid Explainable Attention Mechanism (HEAM). Each model was trained and evaluated independently on the Breast Cancer Wisconsin dataset, keeping the CNN backbone architecture constant.

We used three key interpretability metrics:

- **Localisation Accuracy:** Agreement between highlighted features and clinically annotated tumour indicators.
- **Fidelity:** Degree to which the explanation reflects the model's decision-making, evaluated by occlusion tests.
- **Sparsity:** The percentage of input features required to generate a meaningful explanation.

The results are shown in Table 4. Figure 6 shows the SHAP decision plot for the feature analysis in proposed model.

Table 4

Evaluation of Interpretability Methods

Method	Localization Accuracy (%)	Fidelity (%)	Sparsity (%)
SHAP	78	81	39
Grad-CAM	85	76	52
Attention	82	79	28
HEAM	**91**	**89**	**23**

Figure 6

SHAP decision plot.

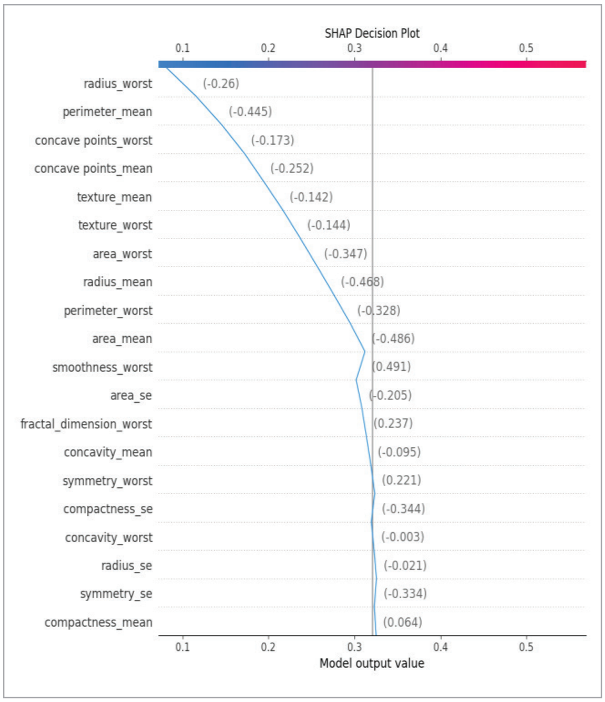

These results highlight that while each method provides value, HEAM consistently outperforms all three interpretability metrics. SHAP provides strong feature attribution but lacks spatial localization. Grad-CAM gives visual insights but is less precise in sparsity. Attention helps focus on essential features, but is limited in explainability scope when used alone.

5. Conclusion and Future Work

In this study, we proposed a deep learning approach that utilises CNN integrated with HEAM and explainable AI to help predict breast cancer. The results were encouraging, with the CNN model reporting a high accuracy of 99.6%, far better than the traditional machine learning methods. As if that is not enough, the model's precision, recall and AUC, where AUC is the area under the curve, were a striking figure almost equal to one, demonstrating the model's proficiency in identifying malignant lesions with very few false positives or negative cases. When integrated into the deep learning model's visualisation of relevant data, the attention mechanism made the model much better without compromising on how the model worked. By using explainability methods such as SHAP and Grad-CAM, physicians and doctors can understand the predictions made by the model. The clinicians will find the model appropriate as the precision and interpretability of the model outweigh the negative impact of the two on patient safety.

Further, certain areas can still be worked on in depth or need more focus. In future studies, this methodology could be extended when working with more complicated datasets, such as medicine-related datasets such as histopathology slides, so that the model's validity across different datasets can be higher since the model would have more cases to be trained on. As the model also has limitations, multi-modal learning where other data is incorporated, such as genetic information and patient history, among others, could help enhance the model further.

Data availability

Data used in this article is publicly available in https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic

References

1. Abunasser, B. S., Al-Hiealy, M. R. J., Zaqout, I. S., Abu-Naser, S. S. Convolution Neural Network for Breast Cancer Detection and Classification Using Deep Learning. Asian Pacific Journal of Cancer Prevention: APJCP, 2023, 24(2), 531. https://doi.org/10.31557/APJCP.2023.24.2.531
2. Ahmad, M. A., Eckert, C., Teredesai, A. Interpretable Machine Learning in Healthcare. Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics, 2018, 559-560. https://doi.org/10.1145/3233547.3233667
3. Ayer, T., Alagoz, O., Chhatwal, J., Shavlik, J. W., Kahn, C. E. J., Burnside, E. S. Breast Cancer Risk Estimation with Artificial Neural Networks Revisited: Discrimination and Calibration. Cancer, 2010, 116, 3310-3321. https://doi.org/10.1002/cncr.25081
4. Bjekic, M., Lazovic, A., K. V., Bacanin, N., Zivkovic, M., Kvascev, G., Nikolic, B. Wall Segmentation in 2D Images Using Convolutional Neural Network. PeerJ Computer Science, 2023, 9. https://doi.org/10.7717/peerj-cs.1565
5. Islam, M. M., Iqbal, H., Haque, M. R., Hasan, M. K. Prediction of Breast Cancer Using Support Vector Machine and K-Nearest Neighbors. 2017 IEEE Region 10 Humanitarian Technology Conference (R10-HTC), 2017, 226-229. https://doi.org/10.1109/R10-HTC.2017.8288944
6. Ji, H., Zhu, Q., Ma, T., et al. Development and Validation of a Transformer-Based CAD Model for Improving the Consistency of BIRADS Category 3-5 Nodule Classification Among Radiologists: A Multiple Center Study. Quantitative Imaging in Medicine and Surgery, 2023, 13, 3671. https://doi.org/10.21037/qims-22-1091
7. Jia, Y., Chen, G., Chi, H. Retinal Fundus Image Super-Resolution Based on Generative Adversarial Network Guided with Vascular Structure Prior. Scientific Reports, 2024, 14(1), 22786. https://doi.org/10.1038/s41598-024-74186-x
8. Khater, T., Hussain, A., Bendardaf, R., Talaat, I. M., Tawfik, H., Ansari, S., Soliman, M. An Explainable Artificial Intelligence Model for the Classification of Breast Cancer. IEEE Access, 2023, 1-1. https://doi.org/10.1109/ACCESS.2023.3308446

9. Loh, H. W. Application of Deep Learning Models for Automated Identification of Parkinson's Disease: A Review (2011-2021). Sensors, 2021, 21(21), 7034. https://doi.org/10.3390/s21217034
10. Ma, W. K., et al. Extra Patient Movement During Mammographic Imaging: An Experimental Study. The British Journal of Radiology, 2014, 87(1044), 20140241. https://doi.org/10.1259/bjr.20140241
11. Manhas, J., Gupta, R. K., Roy, P. P. A Review on Automated Cancer Detection in Medical Images Using Machine Learning and Deep Learning Based Computational Techniques: Challenges and Opportunities. Archives of Computational Methods in Engineering, 2022, 29, 2893-2933. https://doi.org/10.1007/s11831-021-09676-6
12. Naresh, K., Mishra, N. Breast Cancer Diagnosis Using Deep Learning Algorithm. 2018 IEEE International Conference on Advances in Computing, Communication, Control and Networking (ICACCCN), 2018.
13. Retson, T. A., Eghtedari, M. Expanding Horizons: The Realities of CAD, the Promise of Artificial Intelligence, and Machine Learning's Role in Breast Imaging Beyond Screening Mammography. Diagnostics, 2023, 13(13), 2133. https://doi.org/10.3390/diagnostics13132133
14. Siegel, R. L., Miller, K. D., Jemal, A. Cancer Statistics, 2017. CA: A Cancer Journal for Clinicians, 2017, 67(1), 7-30. https://doi.org/10.3322/caac.21387
15. Stockton, T., Peddle, B., Gaulin, A., Wiechert, E., Lu, W. Comparative Analysis of Data Preprocessing Methods in Machine Learning for Breast Cancer Classification. In: Barolli, L. (eds) Advanced Information Networking and Applications. Lecture Notes on Data Engineering and Communications Technologies, vol 201. Springer, Cham, 2024. https://doi.org/10.1007/978-3-031-57870-0_24
16. Topol, E. J. High-Performance Medicine: The Convergence of Human and Artificial Intelligence. Nature Medicine, 2019, 25(1), 44-56. https://doi.org/10.1038/s41591-018-0300-7
17. Trang, N. T. H., Long, K. Q., An, P. L., Dang, T. N. Development of an Artificial Intelligence-Based Breast Cancer Detection Model by Combining Mammograms and Medical Health Records. Diagnostics, 2023, 13(3), 346. https://doi.org/10.3390/diagnostics13030346
18. Vaidya, J. S., Patkar, V. Fast Facts: Early Breast Cancer. S. Karger AG, 2016. https://doi.org/10.1159/isbn.978-1-910797-25-9
19. Yadav, R. K., Singh, P., Kashtriya, P. Diagnosis of Breast Cancer Using Machine Learning Techniques-A Survey. Procedia Computer Science, 2023, 218, 1434-1443. https://doi.org/10.1016/j.procs.2023.01.122
20. Yassin, N. I. R. Machine Learning Techniques for Breast Cancer Computer Aided Diagnosis Using Different Image Modalities: A Systematic Review. Computer Methods and Programs in Biomedicine, 2018, 156, 25-45. https://doi.org/10.1016/j.cmpb.2017.12.012
