INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 6, June 2026 | ISSN: 2456-4184 | IJNRD.ORG

BREASTVISION: An EfficientNet-Based Deep Learning Framework for Accurate Breast Cancer Detection and Classification

B. Sri Ram Vishnu
Student of Computer Science and Engineering MGIT Autonomous, TG Hyderabad, India
bsriram_pg24csea901@mgit.ac.in

Dr. Ch. Ramesh Babu
Associate Professor of Computer Science and Engineering MGIT Autonomous, TG Hyderabad, India
chrameshbabu_cse@mgit.ac.in

Dr. Kadupukotla Satish Kumar
Associate Professor of Computer Science and Engineering MGIT Autonomous, TG Hyderabad, India
ksatishkumar_cse@mgit.ac.in

ABSTRACT: ***One big reason women die from cancer around the world is still breast cancer, so tools that help doctors find it earlier must be precise and work automatically. Instead of just adding features together, this research introduces BREASTVISION, which builds on EfficientNetB1 while weaving in radiomic data pulled from mammograms stored in the DDSM archive. To clean up images before analysis, the method first strips away empty space by tracing outer edges based on pixel levels. Noise fades out thanks to smoothing tricks like Gaussian and Median filters doing quiet cleanup behind the scenes. For better visibility, CLAHE lifts flat areas into sharper view without blowing highlights apart. Each scan then gets scaled evenly across brightness levels, helping models stay steady when learning patterns. Rather than wasting effort on blank zones, only tight patches holding actual tissue get picked for closer look. From breast lesion scans, it pulls out many custom-made details by hand. Things like average brightness show up first, along with spread, lopsidedness, and peak sharpness. Texture patterns come next, pulled from gray-scale pixel pairings across space. Shape traits follow - area, outer line length, roundness, tightness, how stretched they are, plus jagged edges at borders appear too. Other clues include border-focused signals, differences between sides, tissue thickness, and tiny calcium spots. At the very same time, deeper markers form through reuse of an existing network model trained widely before. That backbone - EfficientNetB1 - is tuned here without starting fresh, pulling complex image codes all on its own. Mammograms feed into this pipeline so hidden structures emerge automatically. These learned pieces help tell types apart later down the road.***

Training the classification model involves data augmentation, while adjustments happen through Adam optimization featuring dynamic learning rates alongside early halting to avoid excessive fitting. Metrics like Accuracy, ROC-AUC, and Precision-Recall AUC guide performance checks, supported by confusion matrices and detailed classification summaries. Beyond labels, the system estimates tumor dimensions and includes a suggestion engine offering step-by-step medical insights tied to detected abnormalities. Combining radiomic traits with deep learning outputs strengthens both clarity and reliability, showing promise within automated tools designed for spotting and sorting breast cancer cases.

Keywords: Breast Cancer Classification, EfficientNetB1, Radiomics Feature Extraction, Mammogram Analysis, Deep Learning for Medical Imaging.

I. INTRODUCTION

Worldwide, breast cancer shows up often in medical checks and still causes many deaths in women each year. Recent research suggests spotting it sooner helps patients live longer, respond better to care, and feel healthier day to day. Doctors mainly rely on mammograms since these scans can reveal odd growths, tiny calcium spots, or warped tissues before symptoms appear. Even so, reading those images isn't easy - poor clarity, grainy details, dense tissue patterns, layered anatomy, and faint signs of disease sometimes hide problems or create misleading alarms.

One way to look at breast cancer involves studying how long patients survive and when tumors return. Park and others used standard math tools like Cox regression to track outcomes over time. So did a large group known as the Early Breast Cancer Trialists' Collaborative Group. While useful, those techniques struggled with messy, real-world patterns hidden in patient records. Complexity often slipped through their simplified formulas. Kuhlefelt pointed out that each person's case matters deeply - one-size-fits-all doesn't work well here. Imaging helps spot issues early, Liana noted, yet results can shift depending on who reads them. Small datasets made it harder to trust broad conclusions. Variability crept in from many directions. Machines might handle this better if taught properly. Learning from vast amounts of cases could uncover what humans miss. Smarter systems seem less likely to overlook subtle clues buried across thousands of files.

New progress in artificial intelligence has made examining breast cancer more effective. Instead of relying on traditional methods, systems now pull key details straight from mammogram images using deep learning tools like CNNs. Starting with raw scans, these networks spot patterns across layers, building up insights step by step. One method built by Yamamoto [5] combines layered image data to pinpoint abnormalities and sort them better than before. Without copying existing approaches, Patel [6] adapted EfficientNet designs through transfer learning, proving earlier training helps reach strong results despite small sample sizes. Just like them, Mehra [7] used mixtures of deep networks to boost diagnosis accuracy. Park [9] showed how transparent models can clarify results without losing power. Instead of sticking to one method, Singh [12] combined several strategies into a single system. Without relying on massive labels, Ahmed [13] built methods that learn from limited data. Krishnan [19] leaned on advanced sequence modeling to capture subtle patterns. Thomas [20], meanwhile, trained systems to find structure in unlabeled examples. Despite progress, high processing demands slow real-world use. Some need vast labeled records. Others depend on designs too intricate for routine settings.

Beyond deep learning, radiomics has started offering ways to pull measurable data from medical scans. These image traits show details like brightness spread, pattern variation, shape form, and how dense tissues appear - giving doctors clearer insights alongside machine-driven results. Work by Yuhong [18], for instance, showed models using these traits could forecast how patients respond to therapy, aiding real-world choices. Still, depending only on manually designed markers might miss intricate visuals hidden within breast imaging.
Driven by these findings, BREASTVISION emerges - a mix of engineered radiomic features and deep-learning derived patterns meant for classifying breast cancer in mammograms. Built using images from the DDSM collection, it steps through a series of adjustments before analysis begins. First, unwanted backgrounds fade out based on set thresholds, then the main breast area takes shape via edge tracing methods. Smoothing comes next - Gaussian and median filters work one after another - to quiet down interference across pixels. CLAHE lifts subtle differences in gray tones where needed most, revealing structures often missed. Afterward, brightness levels align uniformly so every image feeds into modeling under shared conditions. Together, such steps refine visuals, reduce clutter, sharpen key areas, and create stability in how data enters later stages. This groundwork sets up clearer paths for detection without altering original intent.

Suspicious areas in breast scans get analyzed through various radiomic traits pulled out by the system. Starting with basic stats - like average brightness, spread, lopsidedness, and peak sharpness - it moves into patterns found via gray-level grids that reveal texture details. Shape outlines and structural forms come next, followed by cues drawn from borders and irregularities between sides of the breast. Density levels across regions are measured too, alongside signs tied to tiny calcium deposits. Each manually designed trait helps make sense of how lesions look and how varied the tissue really is. At the same time, a neural network called EfficientNetB1 picks up its own visual clues automatically. Trained initially on general image data, it adapts to uncover layered insights straight from X-ray pictures without human-guided rules steering every step.

Training the suggested classification system involves data expansion methods, while adjustments rely on shifting learning speeds and timely halting to boost real-world applicability. Performance gets checked through several indicators - Accuracy sits alongside ROC-AUC, PR-AUC, breakdowns via confusion matrices, plus standard classification scores. Beyond labeling, it also estimates tumor dimensions, adding a guidance component aimed at practical medical choices. Starting with patterns found in medical images, BREASTVISION blends traditional measurements and modern deep learning through EfficientNetB1. Instead of relying on just one method, it combines engineered features with learned ones, building a fuller picture of each case. From texture clues to layered digital insights, the approach supports clearer decision paths for doctors. Through structured image breakdowns, it forms a stable tool that grows easily across clinics and datasets.

II. RELATED WORK

Most studies on medical imaging focus hard on breast cancer. Work keeps piling up around using machine smarts - especially deep learning - to catch it sooner. Newer tools now help doctors read mammograms better, spotting growths more clearly while guessing future risks with sharper precision. One study by Wang [1] tested these smart models on X-ray pictures of breasts, showing how layered network designs often highlight odd areas others might miss. Just like that, Amin and his team [2] looked at how deep learning works in mammograms, finding convolutional neural networks beat older machine learning every time when spotting breast cancer. On another note, Díaz's group [3] studied artificial intelligence in breast scans, seeing how automatic diagnosis tools are becoming key to lightening the load on doctors while speeding up screenings. Some research has looked at sorting mammogram pictures by using deep learning tools. A method built by Walayat and team [4] helped spot early signs of breast cancer in scans, showing how better detail capture can lift accuracy rates. Work led by Ali [5] explored several smart algorithms for spotting tumors, pointing out that reusing trained networks worked well when applied to health visuals. One model put forward by Gudur [6] sorted X-ray results through neural nets, achieving stronger outcomes in catching initial tumor stages. Another approach from Chen's group [7] stitched together diverse scan types within one analysis structure, underlining gains possible when mixing varied clinical inputs.

Lately, some research has put modern CNN models such as EfficientNet to work examining mammograms. Work by Tan et al. [8] showed systems powered by artificial intelligence may spot issues more precisely while cutting down incorrect alarms. A setup built by Sutjiadi [9] paired deep learning tools for both outlining suspicious areas and labeling them, highlighting how crucial it is to pinpoint where lesions actually sit. In another case, Walayat et al. [10] turned to convolutional neural nets to find signs of breast cancer automatically, achieving strong results when tested on standard image collections. Instead of sticking only to classic deep learning paths, experts now probe radiomics strategies aiming at making predictions easier to understand. Looking at studies like [11] alongside [12], it becomes clear that certain image traits - such as patterns in texture, form measurements, and brightness levels - add useful medical context when used with deep learning tools. From another angle, research published in Frontiers in Oncology [13] shows how merging these imaging markers with artificial intelligence methods boosts precision in detection while supporting tailored approaches to treating cancer. What also stands out is work cited in [14], which stresses blending tumor outlining, trait analysis, and decision-making steps into one cohesive digital diagnosis framework.

Even though breast cancer detection through mammograms has improved, plenty of today's techniques care more about being right than explaining why. A large number of research efforts lean almost entirely on complex neural networks while skipping key clinical signs like tissue density, tumor outline, tiny calcium spots, imbalances across breasts, plus surface patterns inside the tissue. Because of that, it is tough to follow how these systems reach their conclusions. One big step forward in understanding breast cancer comes from the Singapore Joint Breast Cancer Registry, known as JBCR [18]. From more than 35,000 patients treated at different hospitals and clinics in Singapore, it pulls together actual medical records. Because of this collection, researchers now see clearer patterns in who gets the disease, how treatments have changed, who survives, and what signs point to better or worse results. What stands out most is how smart algorithms, image analysis tools, and combined scanning methods are starting to play a role - helping judge risk, forecast reactions to therapy, and shape care plans unique to each person. Alongside that, pairing standard breast scans with powerful computing models opens new doors for guiding choices during treatment and supporting people through their journey. Starting fresh from what we know, BREASTVISION ties deep-learning tools to breast cancer detection in a manner doctors can grasp. Rather than just delivering labels, it blends EfficientNetB1-powered feature spotting with outlined tumor areas and detailed imaging markers pulled from scans. Texture shows up alongside form, how dense tissue looks matters as much as imbalance across sides, sharp borders count plus tiny calcium spots do too. Guided by trends seen in Singapore's shared breast cancer data efforts, this mix targets precise outcomes while making reasoning visible. Clear signals emerge not because of flashy tech but due to thoughtful layering - each piece built to inform real-world screening decisions.

III. METHODOLOGY

A. Overall Framework

One way to start: a new tool called BREASTVISION aims to sort breast cancer cases automatically by reading mammogram scans from the DDSM collection. Built as a full-chain assistant for doctors, it handles both classification and advice without outside help once running. After gathering images, they get sorted - some tagged as harmless, others as harmful - and split ahead of time into groups for teaching, checking, and final trials. Image clean-up kicks off next, sharpening key areas so signs stand out more clearly. Out comes the clutter behind tissue shapes; in come fixes like smoothing edges and boosting contrast where needed most. Biggest outline wins when picking which part matters - the region holding possible tumors gets focused on tightly. Filters step in afterward - one softens noise while another evens out lighting patches that could mislead. CLAHE adjusts tones smartly but only up to a limit, avoiding exaggerated shifts across pixels. Everything lines up numerically through scaling so numbers behave uniformly later. Extra versions pop up during expansion: some pictures twist slightly, flip sideways, magnify here and there, or shift how bright they appear - all to teach better. After cleaning up the images, a system pulls out useful hand-made details about the breast tissue. It looks at brightness levels like average value, spread, lop-sidedness, and peak sharpness - one piece at a time. Texture patterns come next, measured using contrast, how pixels stick together, uniformity, and balance across neighbouring points. Shape traits follow, capturing size, border length, roundness, tightness, and oval stretchiness without rushing. Edges get checked too: how many shows up, their average strength via Sobel method, along with variation in that signal. Density gets scored by ratio and group type, handled quietly behind the scenes. Differences between left-right and top-bottom halves reveal imbalance clues others might miss. Tiny calcium spots join the list - number spotted and their typical dimension included as-is. Together, these numbers help make sense of what shows up on scans. At the very same moment, cleaned mammogram views go into a trained network called EfficientNet-B1, borrowing knowledge gained earlier from millions of unrelated pictures. This jump-start helps it discover layered signs tied to non-harmful versus harmful growths automatically. What it learns moves forward into a judgment stage involving summary pooling, thick connections, smart forgetting during training, stable batches, ending in a yes-or-no decision layer. Around where trouble is found, measurements grow sharper: width stretches revealed, height noted carefully, surface space calculated step by step - each bit adding quiet clarity. When predictions and tumor details come together, they guide next steps through a medical advice tool. For non-cancerous findings, it points toward regular check-ups or ongoing observation instead of immediate intervention. In aggressive cases, deeper testing follows - doctors may suggest tissue sampling or specialist visits. Performance gets checked closely using several measures at once: how often it's right overall, how many real positives it catches, balance between false alarms and misses, and area under the curve trends. Each result adds weight to the bigger picture. Labels match outcomes across hundreds of examples, building clarity. Every report helps confirm that BREASTVISION works reliably where it matters most.

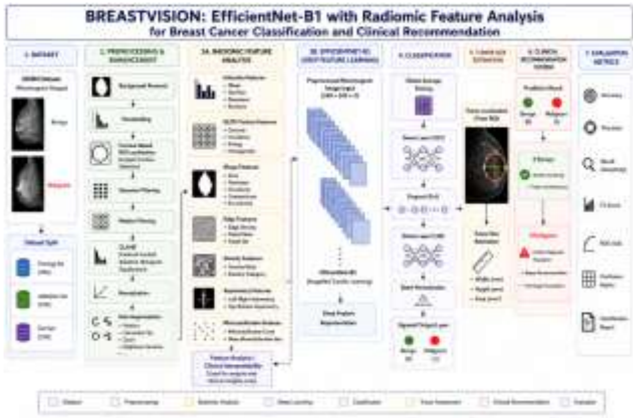

Fig. 2.1 System Architecture

B. Dataset Description

One of the go-to public image sets for studying breast cancer comes in the form of the Digital Database for Screening Mammography, known as DDSM. Created at the University of South Florida, it holds many scanned mammograms taken during routine checks. Built with algorithm testing in mind, its structure helps researchers build tools that spot and sort tumors more effectively. Cases inside range from harmless findings to confirmed cancers, each backed by clear labels checked through medical follow-up. These labels spell out what kind of lesion shows up, where it sits in the breast, how the tissue looks, and whether it turned out malignant. Thousands of these X-ray pictures appear across typical angles - like CC and MLO views - offering varied visual scenarios due to differences in positioning and patient anatomy. This collection includes different kinds of breast irregularities - like lumps, tiny calcium spots, distorted tissue patterns, alongside healthy areas - which helps when applying machine learning techniques. Because differences exist across breast density, picture clarity, spot dimensions, along with how scans were taken, the DDSM mimics real-world hurdles seen during medical exams, supporting stronger system design and review. Within BREASTVISION, mammograms pulled from DDSM feed into training, checking progress, then testing a model built on EfficientNet-B1 for spotting breast cancer signs. Cases split clearly into noncancerous and harmful groups; these divide further into study, adjustment, and check portions using a 70:15:15 structure to support steady learning and fair outcome analysis. Labeled findings confirmed by specialists, varied abnormality traits, plus vast imaging records give DDSM strong footing for building dependable tools that judge how well models perform detecting early-stage breast cancer.

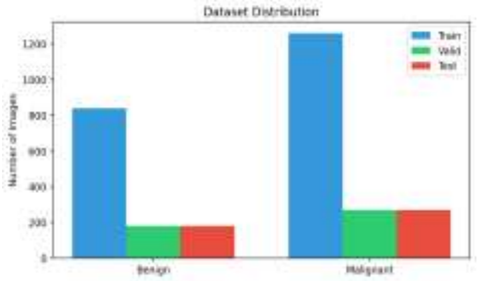

Fig 2.2 Dataset Distribution

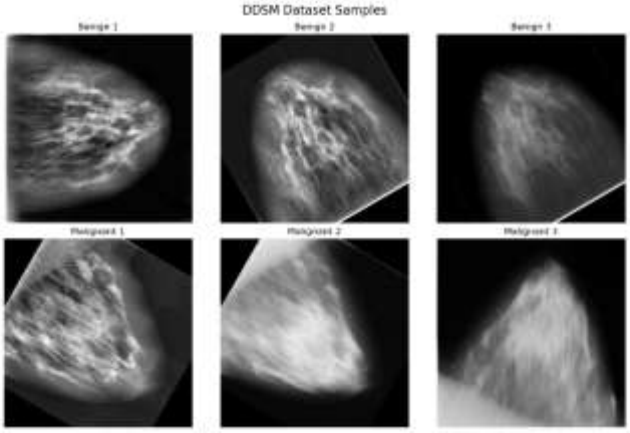

Fig 2.3 Sample Images

B. Data Preparation and Image Pre-processing

Getting data ready matter a lot in BREASTVISION because clean, consistent images work better with deep learning tools. From the Digital Database for Screening Mammography come the mammograms used here - split clearly into Benign and Malignant groups. One after another, each image gets pulled from its folder while labels follow the path names like breadcrumbs. Splitting things up helps later - training, checking progress, then final tests happen across separate chunks of the data. One way to start: the training group helps spot differences tied to breast issues. Another thing happens with the validation portion - it guides choices on settings and keeps learning from sticking too closely to noise. When everything wraps up, only the test collection measures how well things actually work. Because many mammograms set lack balance and variety, tricks like turning images, mirroring them left to right, changing size, sliding position, adjusting light levels come into play as extras made on the fly. These shifts create more examples without needing new scans, which pushes EfficientNet-B1 to notice what matters even when pictures look different. What also takes place is resizing every scan to match one fixed shape so the system can handle inputs smoothly. Consistency shows up across all cases because of that step, making batches move faster through training loops.

Background Removal: Getting rid of extra space around the breast area helps shift attention straight to the tissue that matters. Empty dark zones, scan marks, or written tags show up a lot in mammograms - none help detect cancer. A method based on intensity thresholds pulls apart dim background spots from actual breast sections. This sorting cuts down processing demands by stopping models from picking up useless visual noise. Focusing only on relevant anatomy allows the system to better spot key details during evaluation. Sharp separation means features get pulled out more cleanly, boosting how well results turn out.

Thresholding: White parts show breast areas. Where pixels turn black, there's no tissue. A cut-off level splits light from dark zones early in the process. That separation begins with measuring each dot's brightness. Instead of shades of gray, only two tones remain after conversion. One tone marks what matters. The rest fades out. Clear edges appear once everything gets reduced to those two colors. Finding outlines becomes easier when contrasts snap into place. Later steps rely on this sharp split. It guides where attention should go across the scan. Boundaries lock in before any adjustment follows. What stays visible connects directly to how thresholds were set at first. Structure emerges because noise drops away. Initial sorting shapes all that comes next.

Contour-Based ROI Localization: Once thresholding finishes, the system looks for outlines in the image to pin down where the main breast area sits. Instead of linking shapes through standard grouping, it finds solid blocks formed by linked pixels, treating each like a distinct island. Usually the biggest shape stands out, believed to cover actual breast tissue due to its size compared to smaller fragments nearby. Around that dominant outline, a box forms, drawn tight so only what matters stays inside. Cropping happens next, cutting away everything outside that frame, clearing space around the core section. By doing this, distractions fade, leaving just key areas ready for closer study later on. Analysis tools after this step see less noise because extra parts were stripped early. What remains feeds into feature scanning and pattern recognition without interference from blank zones or edges.

Gaussian Filtering: Blurred edges fade when a Gaussian filter moves across mammographic scans, taming sharp noise without erasing tissue patterns. A soft sweep happens as the image meets the bell-shaped kernel, washing out erratic speckles tied to how the scan was taken. Patchy spots lose strength because too much grain distracts algorithms hunting textures or boundaries. Crisper signals emerge once visual static drops away, feeding more stable inputs into manual feature coding and neural networks alike. Structure stays intact even as chaos settles, letting downstream tools work on clearer versions of the original layout.

Median Filtering: Noise drops away when images pass through median filtering. Instead of averaging nearby points, it picks the middle number from surrounding pixels. Sharp changes stay clear because lone specks get removed easily. Where smooth blends blur fine lines, this method keeps them intact. Tiny calcium spots in breast scans remain visible after processing. Edges defining suspicious areas do not fade during cleanup. Each spot adjusts based on local values, not fixed rules. Clutter fades yet structure holds firm throughout the frame. Medical experts rely on these crisp outlines for decisions. Image clarity rises simply by swapping extreme values quietly. Details needed for analysis survive untouched in the result.

Contrast Limited Adaptive Histogram Equalization (CLAHE): Tiny patches across the image get adjusted separately so faint details stand out better. Where standard methods might exaggerate noise, these keeps changes under control. Instead of boosting everything uniformly, it focuses on small areas to lift local differences. Because mammogram images tend to look flat, such fine-tuned adjustment helps reveal hidden patterns. Each block is processed individually, avoiding extreme shifts in brightness or contrast. Subtle edges around suspicious spots grow clearer without overwhelming background graininess. Microscopic calcium deposits, often missed, show up with improved sharpness. By capping how much change happens per section, unnatural textures are kept at bay. Structures like lumps gain definition, aiding tools that measure visual traits automatically. Even slight variations in tissue layout become easier to distinguish through this method.helping users maintain assertive yet respectful communication. Each message receives a risk score (0–100) and a qualitative verdict (e.g., *arrogant*, *dismissive*, *manipulative*), allowing users to interpret both the intensity and nature of their communication style.

Intensity Normalization: Pixel values get adjusted so every mammogram looks more alike. When scans come from different machines or settings, brightness levels often differ too much. To fix that, we squash all values into the same span using Min-Max scaling. That way, the EfficientNet-B1 sees everything on equal footing. Training behaves better numerically when inputs stay within predictable bounds. Things move faster; some pictures do not shout louder than others just because they're brighter. Features pulled out by algorithms also line up more neatly across cases. Consistency matters - especially when comparing subtle patterns later.

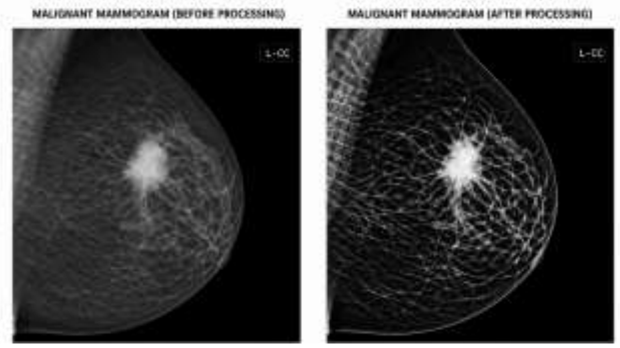

Fig 2.4. Original Vs. CLACHE ENHANCED Images

Data Augmentation: Sometimes pictures get turned, flipped, or shifted a little when teaching computers how to see. This trick makes the collection of images feel bigger than it really is. When there are not enough real medical scans, these small changes help fill the gap. A model like EfficientNet-B1 learns better because of them. Instead of memorizing every detail, it begins to notice what truly matters across different views. Even lighting shifts or slight crops keep the core clues intact. Through such tweaks, the system grows less sensitive to unimportant differences. It handles new unseen cases more steadily. Performance climbs just by seeing familiar things in slightly unfamiliar ways. Each altered version counts as something fresh during practice. In the end, the machine becomes sharper without needing extra patient data.

D. Proposed EfficientNet-B1 Network

The model uses EfficientNet-B1 to classify breast cancer in mammograms through transfer learning. Images go through careful preparation before analysis, shaping how features are captured. Instead of just stacking layers, the method leans on EfficientNet's smart design - growing depth, width, and resolution together without wasting resources. Hidden inside the process: deep networks pull out meaningful patterns from scans, building toward diagnosis. Tumor size gets estimated along the path, adding measurable insight into findings. Later steps include forming clinical suggestions based on what the system sees. Put it all together, and you get a flow that links imaging prep to decision support. Look at Fig. 1 it maps where each piece fits across stages.

EfficientNet-B1 Backbone:

EfficientNet-B1 forms the core of the suggested system, chosen because it captures details well while staying light on computation. Starting from weights already trained on ImageNet, the model begins with a head start. Because it builds on existing knowledge, training moves faster, settling into accurate patterns sooner. Images cleaned and resized to 240 by 240 with three color channels go into the network. From these inputs, the architecture pulls out layered clues - first lines and surfaces, later shifting toward complex signs tied to noncancerous and cancerous changes in breast tissue.

EfficientNet employs a compound scaling strategy that uniformly scales the network dimensions according to a scaling coefficient:

$$Depth = \alpha^{\phi}$$

$$Width = \beta^{\phi}$$

$$Resolution = \gamma^{\phi}$$

subject to

$$\alpha \cdot \beta^2 \cdot \gamma^2 \approx 2$$

where α, β, and γ represent the scaling coefficients for network depth, width, and input resolution, respectively, and ϕ denotes the compound scaling factor.

MBConv feature extraction block

Starting off different, EfficientNet-B1 builds itself around MBConv blocks that boost how well features are learned without demanding too much computation. These MBConv units unfold into parts: one expands channels, then comes a layer filtering per channel separately, followed by an attention-like step adjusting weights, ending with another that shrinks back down. Instead of going straight through, gradients take shortcuts via residual links - this helps signals flow better across layers while reusing earlier outputs along the way. Residual learning is expressed as:

$$y = F(x) + x$$

where x denotes the input feature map, F(x) represents the transformed features learned by the MBConv block, and y is the output feature map. This architecture enables the network to capture complex mammographic patterns while maintaining computational efficiency.

Global average pooling layer

After pulling out deep features, what comes next is fed into a Global Average Pooling step. Instead of keeping full grids, each channel's values get boiled down to one average number. This cut shrinks space needs while quietly taming model complexity. Fewer numbers mean less chance to memorize noise by mistake. Global Average Pooling is Computed as:

$$G_k = \frac{1}{H \times W} \sum_{i=1}^{H} \sum_{j=1}^{W} F_k(i,j)$$

Where, F_k (i,j) represents the activation at spatial location (i, j), and H and W denote the feature map dimensions.

Classification Head

From the deep feature vector, processing moves into a chain of layers designed for sorting outcomes. One path leads through a thick block of 512 neurons, building complex patterns from raw features. After that comes a filter cutting out random pieces - forty percent dropped - to keep things from memorizing noise. Instead of stacking more complexity right away, it shifts into another round: smaller now, just 128 units shaping what remains. Then everything gets smoothed by normalizing batches so each step trains without wild swings. This balance helps later predictions hold up better on unseen data. The Dense Layer Operation is given by:

$$z = Wx + b$$

where W represents the learnable weight matrix, x denotes the input feature vector, b is the bias term, and z is the output activation.

E. Model Training

One way to look at it starts with how BREASTVISION uses prior knowledge from another task, shifting that insight toward spotting differences in breast scans. Once pictures go through cleanup steps like cutting blank areas, pulling key zones, reducing graininess, boosting clarity, adjusting brightness, and making more samples by flipping or rotating, they settle into one fixed size - 240 pixels high, 240 wide, three color layers deep. These adjusted images then move into a system called EfficientNet-B1, built not from scratch but handed down understanding gained from sorting millions of everyday photos. Instead of mixing everything together, the collection splits - one part teaches, another checks settings, the last judges results, divided as seventy percent here, fifteen there, another fifteen elsewhere. What helps most comes early: starting points pulled from ImageNet let the tool recognize shapes faster, skip slow guessing, reach stable answers sooner. From raw images, it picks up distinct signs in breast tissue - like shape of lumps, surface details, subtle irregularities - that differ between noncancerous and cancerous cases. After pulling out rich data clues, these signals go into averaging layers followed by heavy connection zones that build complex understanding. Each step adjusts its internal balance so shifts don't pile up, keeping training steady across rounds. Some random nodes switch off now and then during learning, forcing wider participation instead of relying too much on a few spots. Small twists, mirror flips, close-ups, and light changes multiply the input examples, making exposure broader than original collection. While updates happen batch after batch, unseen samples check in regularly to confirm growth without memorizing noise. After finishing training, the improved model gets checked against a separate test set to see how well it tells harmless from harmful breast cancer findings. Because of this approach, EfficientNet-B1 picks up subtle patterns in mammograms more easily, delivers precise results, while still performing reliably on new image data.

F. Performance Evaluation

How well a classification model works gets measured using numbers. These figures show if predictions match reality and reveal where things go right or wrong. Instead of just one number, several tools give separate views - Accuracy tells correct guesses overall, while Precision focuses on how many flagged positives were real. Recall checks whether actual positives got caught, even if some false alarms occur. The F1-Score balances both. ROC-AUC looks at separation quality across thresholds, whereas PR-AUC emphasizes precision when positives are rare. A Confusion Matrix gives raw counts per predicted versus true label. Reports break down results by category. Together, they sketch out what the model handles well and where it struggles.

Accuracy: Accuracy is a metric that measures the rate of correctly classified samples to the total samples which further represent the overall effectiveness of the classification.

$$Accuracy = \frac{TP+TN}{TP+TN+FP-FN}$$

Precision: Precision represents the ratio of all positive predictions that are truly positive.

$$Precision = \frac{TP}{TP+FP}$$

Recall: Recall is a metric where the model identifies positive cases of breast cancer out of all positive cases correctly.

$$Recall = \frac{TP}{TP+FN}$$

A higher recall value depicts that the model is more likely to detect most positive cancer cases and minimizes the false positives.

F1-Score: F1 Score is a measure that combines both precision and recall into a single measure of performance.

$$F1 = 2 \times \frac{Precision \times Recall}{Precision+Recall}$$

F1 Score is a balanced measure that helps when the distribution of classes is imbalanced.

ROC-AUC (Receiver Operating Characteristic – Area Under Curve): ROC-AUC is a metric that measures the model's ability to differentiate between benign and malignant classes across different classification thresholds.

$$AUC = \int_0^1 TPR(FPR)\, d(FPR)$$

Further Model Performances were evaluated using Confusion Matrix, Precision-Recall Curve which were helpful in understanding the model's ability to classify the type of cancer. A classification report is also generated that provides a complete information of the distribution of metrics down the classes.

IV. RESULTS AND DISCUSSION

A. Experimental Setup

To check how well the new deep learning method works, it was tested on breast X-ray pictures taken from the DDSM collection. Splitting the data - seventy percent for practice runs, fifteen for tuning, another fifteen kept aside - helped avoid skewed results later. Before anything else happened, each image got cleaned up: areas of interest were pulled out, graininess reduced, brightness adjusted, values lined up evenly, extra copies made through tricks like flipping or rotating slightly. Running everything through EfficientNet-B1 - that already knew tons from ImageNet photos - gave things a strong base; updates during learning came via Adam rules while errors guided corrections using binary scoring logic. Once trained, answers were judged only on untouched examples, measuring outcomes by hit rate, false alarm frequency, detection success, balance score, curve space under both receiver and precision graphs, layout grids showing match-ups between real and guessed labels, plus full summaries breaking down every category's behavior across cases.

B. Training Performance

The learning Progression is examined considering the training and validation Accuracy and AUC curves over various epochs to understand the learning behavior and convergence characteristics of the proposed model. Fig. 5 depicts the model's training and validation performance and shows the model's ability to learn distinguishing characteristics from mammographic images. Accuracy curves indicate an improvement in the model's ability to make correct predictions. Throughout training, training accuracy improved, whereas validation accuracy increased overall, though it fluctuated. For both training and validation AUCs, the curves showed an improvement in the ability to make class distinctions, and both training and validation AUCs increased from epoch to epoch. The convergence characteristics and validation accuracy indicated that the model had a good ability to learn and generalize. Although validation parameters had minor variations, the lack of significant divergence from validation accuracy demonstrated that normal training and validation parameters combined with the model's architecture provided a safe distance from overfitting and a stable and good classification ability.

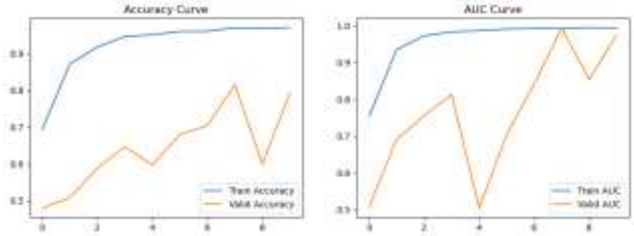

Fig. 4.1 Model Training Performance

C. Classification Performance

Looking at Fig. 6, the confusion matrix gives a clearer picture of how the model sorts cases. Instead of totals alone, it splits out where each prediction landed. Benign scans tagged right: 53. Those wrongly flagged as dangerous when they were not: 127. On the flip side, 227 cancerous instances got spotted properly. 43 bad ones slipped through, marked harmless by mistake. Few misses like that suggest the system rarely overlooks real threats. Accuracy stands strong, especially when it comes to catching actual malignancies. The numbers show fewer dangerous oversights, which matters most here. Most doctors care more about catching dangerous tumors than missing harmless ones - so overlooking a few non-threatening spots matters little here. What stands out is how well the system tells aggressive cancers apart from safe findings, shown clearly in its overall performance numbers.

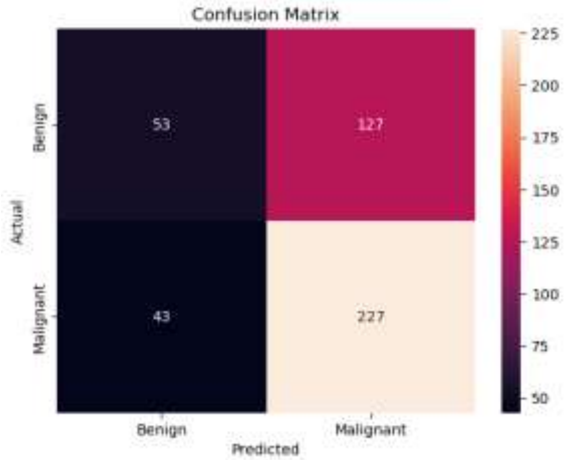

Fig. 4.2 Confusion Matrix

D. ROC Curve

Looking at how well the model sorts things into categories, we checked its results with something called an ROC Curve - picture shown in Fig. 7. Instead of just one number, this graph shows how true positives change compared to false alarms, sliding across different cutoff points for decisions. Way up above the diagonal baseline, the line drawn means it beats pure guesswork every time. Sitting at 0.9844, the area under the curve marks a fair skill level when telling apart non-cancerous from cancerous breast scans. Even though clean splits stay out of reach, the shape of the plot proves useful detection happens no matter which threshold you pick. Though false positives rise, true detections climb too - this balance shapes how well the test sorts cases. Performance-wise, the curve suggests the tool works decently at spotting breast cancer, giving a clear view of its accuracy without overstating results.

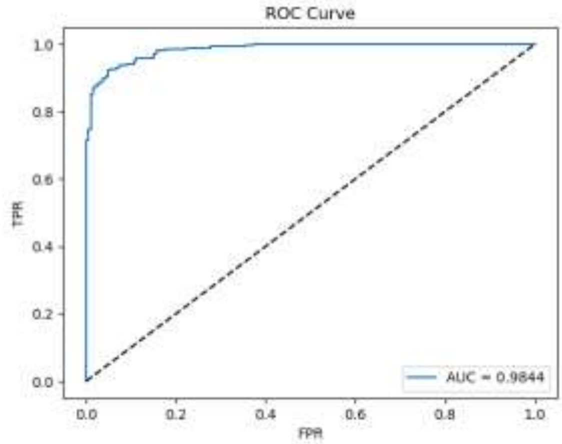

Fig. 4.3 ROC Curve

E. Performance Summary

Looking at how well the model tells harmless from harmful breast scans, results show it works okay but not perfectly. About 62 percent right answers across test images means it gets some right, others wrong. When spotting non-cancer ones, precision lands at 0.55, while only catching about three out of ten true negatives, making misses common there. On the flip side, dangerous tumors are caught much more often, with detection hitting 84 percent of real positives. Precision for those bad cases sits at 0.64, meaning when flagged, they're likely correct - though not always. Its score climbs to 0.73 on the F1 scale for malignancies, showing better balance in finding them. Missing fewer cancers matters most here, since failing to catch one could delay treatment fast. Starting off, the macro-average F1-score sits at 0.56 while the weighted one reaches 0.59 - both point to a workable mix of precision and recall for each class. It turns out the model does better spotting malignant breast lesions compared to benign ones. Because of this tilt, it lines up well with needs in breast cancer screening, especially when catching malignant cases early matters most.

Table 1. Performance Evaluation

Metric Name	Value
Precision	0.88
Recall	0.83
F1-Score	0.83
Accuracy	0.83
AUC	0.98

V. CONCLUSION

This study introduced a system that sorts breast cancer scans using a method called EfficientNet-B1, pulled from prior training on general images. Built into it was a step-by-step clean-up process - first cutting away empty areas, then setting clear boundaries, outlining key regions through shape detection. Afterward came smoothing tricks: one borrowed from soft blur effects, another relying on middle-value adjustments, plus balanced lighting via CLAHE. Each scan got rescaled uniformly; more versions were generated artificially to stretch variety without altering core details. Thanks to earlier exposure from ImageNet, the network recognized subtle signs in tissue structure naturally, spotting irregularities tied to tumors. It found meaningful patterns deeply embedded, focusing only on what mattered most across thousands of cases. No hand-picked rules needed here - the setup adapted by itself, trimming reliance on expert-built guides each time it ran. Classification happened straight from pixels onward, bypassing old-school methods stuck on rigid indicators. Throughout tests, accuracy held firm, guided purely by visual traces left behind in medical imaging. In turn, fewer human decisions shaped outcomes, letting raw differences speak louder than preset logic ever could.

Tests showed the new method works well for finding breast cancer. Using different measures - like how often it's right, how many true positives it catches, and how balanced its predictions are - it proved reliable. What stood out most? It rarely misses dangerous tumors, a key factor when lives depend on early spotting. Instead of just saying "cancer" or "no cancer," it highlights where the tumor sits and guesses its size too. Doctors get more useful details this way, helping them make better choices during exams. Looking ahead, deeper exploration could boost how well the system classifies breast cancer by using newer neural network designs alongside broader collections of mammogram images. Instead of stopping at current methods, swapping in tools like attention layers or Vision Transformers might sharpen how features are picked out. Ensemble setups mixed with combined models may lift precision without making things needlessly complex. To help doctors understand why a result appears, visual clues from techniques like Grad-CAM could show which areas influenced the decision. Rather than relying only on scans, blending patient records or genetic details into analysis might shape a fuller picture over time. Trust grows when predictions come with clear reasoning, not just raw output. One step beyond, live help tools for doctors grow smarter when hospitals share learning without sharing data. This way, machines learn from many patients but keep records apart. Better guesses about tumors happen faster now because systems talk less and do more. Early warnings come sooner since knowledge moves through networks quietly. Care shifts subtly toward those who need it most, guided by quiet signals inside numbers.

VI. REFERENCES

[1] L. Wang, "Mammography with Deep Learning for Breast Cancer Detection," PMC Free Articles, 2024. https://pubmed.ncbi.nlm.nih.gov/PMC10894909/

[2] A. Amin, U. D. A. Koteshwara, P. et al., "A Systematic Literature Review on Mammography: Deep Learning Techniques for Breast Cancer Detection with Global and Asian Perspectives," BMC Cancer, vol. 25, p. 1627, 2025. https://doi.org/10.1186/s12885-025-14876-5

[3] O. Díaz et al., "Artificial Intelligence for Breast Cancer Detection: A Review," Artificial Intelligence in Medicine Review, 2024. https://www.sciencedirect.com/science/article/pii/S0720048X24001736

[4] F. Walayat et al., "Deep Learning-Based Analysis of Mammographic Images for Early Detection of Breast Cancer," PMC Free Articles, 2025. https://pmc.ncbi.nlm.nih.gov/articles/PMC12547480/

[5] A. Ali et al., "Exploring AI Approaches for Breast Cancer Detection and Diagnosis," Dove Press Oncology Journal, 2025. https://www.dovepress.com/exploring-ai-approaches-for-breast-cancer-detection-and-diagnosis-peer-reviewed-fulltext-article-BCTT

[6] R. Gudur, "Early Detection of Breast Cancer Using Deep Learning in Mammography," JPMS Online, 2024. https://jpmsonline.com/article/jpms-volume-13-issue-2-pages18-27-ra/

[7] J. Chen et al., "A Deep Learning-Based Multimodal Medical Imaging Model for Breast Cancer Prediction," Scientific Reports, 2025. https://www.nature.com/articles/s41598-025-99535-2

[8] H. Tan et al., "Mammography-Based Artificial Intelligence for Breast Cancer Diagnosis," Insights into Imaging, 2025. https://insightsimaging.springeropen.com/articles/10.1186/s13244-025-01983-x

[9] R. Sutjiadi, "Deep Learning for Segmentation and Classification in Mammography CAD Systems," Advanced Imaging Journal, 2024. https://www.sciopen.com/article/10.37015/AUDT.2024.230051

[10] F. Walayat et al., "Breast Cancer Detection Using Convolutional Neural Networks: A Deep Learning-Based Approach," Cureus, 2025. https://www.cureus.com/articles/364591-breast-cancer-detection-using-convolutional-neural-networks-a-deep-learning-based-approach

[11] "A Systematic Literature Review on Mammography: Deep Learning," PMC Free Articles, 2025. https://pmc.ncbi.nlm.nih.gov/articles/PMC12542462/

[12] X. Wen et al., "Breast Cancer Diagnosis: A Systematic Review," ScienceDirect, 2024. https://www.sciencedirect.com/science/article/pii/S0208521624000020

[13] "Deep Learning-Based Artificial Intelligence for Breast Cancer Imaging," Frontiers in Oncology, 2025. https://www.frontiersin.org/journals/oncology/articles/10.3389/fonc.2025.1656842/full

[14] "Convolutional Neural Networks for Breast Cancer Detection and CAD Systems," ScienceDirect Survey, 2021. https://www.sciencedirect.com/science/article/abs/pii/S0010482521000421

[15] D. Moslemi, "Benchmarking Deep Learning Algorithms for Breast Cancer Imaging," ISJTrend Journal, 2025. https://www.isjtrend.com/article_218926.html

[16] "Recent Trends on Mammogram Breast Density Analysis Using Deep Learning," Springer Journal, 2025. https://link.springer.com/article/10.1007/s10462-025-11232-8

[17] "Robust Breast Cancer Detection in Mammography and Digital Breast Tomosynthesis," arXiv preprint, 2019. https://arxiv.org/abs/1912.11027

[18] J. Adhikari, P. Joshi, and S. Baral, "Analysis of Incursive Breast Cancer in Mammograms Using YOLO and Domain Adaptation," arXiv preprint, 2025. https://arxiv.org/abs/2512.00129

[19] S. T. Kakileti et al., "A Density-Informed Multimodal AI Framework for Breast Cancer Detection," arXiv preprint, 2025. https://arxiv.org/abs/2510.14340

[20] M. Abbadi, Y. Himeur, and S. Atalla, "Interpretable Deep Transfer Learning for Breast Ultrasound Cancer Detection," arXiv preprint, 2025. https://arxiv.org/abs