



(RESEARCH ARTICLE)



Knowledge distillation-based lightweight Convolutional Neural Networks (CNN) model for efficient breast cancer detection

Omankwu, Obinnaya Chinecherem Beloved ^{1,*}, Etuk, Enefiok. A ¹ and Promise Enyindah ²

¹ Department of Computer Science, Michael Okpara University of Agriculture, Umudike, Umuahia, Abia State. Nigeria.

² Department of Computer Science University of Port Harcourt. Rivers State, Nigeria.

International Journal of Science and Research Archive, 2025, 16(03), 998-1006

Publication history: Received on 10 August 2025; revised on 14 September 2025; accepted on 18 September 2025

Article DOI: <https://doi.org/10.30574/ijrsra.2025.16.3.2625>

Abstract

Breast cancer (BC) remains a leading cause of cancer-related mortality among women globally, with early detection playing a crucial role in improving patient outcomes. While deep learning models—particularly Convolutional Neural Networks (CNNs)—have demonstrated exceptional performance in breast cancer detection, their high computational demands limit deployment in low-resource environments such as mobile devices and rural clinics. To bridge this gap, we propose a lightweight CNN model for breast cancer detection using Knowledge Distillation (KD), a technique that transfers knowledge from a complex, high-capacity "teacher" model to a compact and efficient "student" model. In this study, we develop and evaluate both teacher and student models using the Wisconsin Diagnostic Breast Cancer (WDBC), Breast Cancer Diagnosis (BCD), and Primary Breast Cancer vs Normal Breast Tissue (PBCT) datasets. The student model is designed to operate with only 0.06% of the parameters used by the teacher model, significantly reducing memory and computational overhead. Despite its lightweight architecture, the student model achieves up to 100% classification accuracy and demonstrates robust generalization across multiple datasets. Our approach enables high-accuracy breast cancer detection while ensuring fast inference and low resource consumption, making it well-suited for deployment in real-world, resource-constrained environments such as mobile health platforms and embedded medical devices. The findings highlight the transformative potential of knowledge distillation in democratizing access to advanced AI-driven diagnostics.

Keywords: Breast Cancer Detection; Convolutional Neural Network; Knowledge Distillation; Lightweight Model; Deep Learning; Medical Imaging

1. Introduction

Breast cancer (BC) is one of the most commonly diagnosed cancers and a leading cause of cancer-related deaths among women globally. According to the World Health Organization (WHO), an estimated 2.3 million women were diagnosed with breast cancer in 2020, resulting in over 685,000 deaths worldwide. Although breast cancer is more prevalent among women, it also affects men, albeit at a much lower incidence. The disease typically originates from the inner lining of milk ducts or lobules and can evolve into invasive cancer that spreads to surrounding tissues or distant organs through lymphatic and blood vessels.

Early detection of breast cancer remains the most effective way to improve prognosis and increase survival rates, with early-stage diagnosis having a survival probability of over 90%. However, conventional diagnostic methods such as mammography, ultrasound, magnetic resonance imaging (MRI), and biopsy, though effective, often suffer from limitations such as false positives/negatives, high costs, requirement for specialized equipment, and operator dependence. Furthermore, accessibility remains a major concern in low- and middle-income regions where such diagnostic tools are either unavailable or underutilized due to infrastructural and resource constraints.

* Corresponding author: Omankwu, Obinnaya Chinecherem Beloved

In response to these challenges, artificial intelligence (AI) and machine learning (ML) technologies have gained attention for their potential to augment medical diagnostics. In particular, Convolutional Neural Networks (CNNs), a class of deep learning algorithms specialized for analyzing visual imagery, have demonstrated remarkable performance in breast cancer classification tasks using histopathology images, mammograms, and other imaging modalities. CNNs excel at automatic feature extraction and hierarchical representation of image data, making them suitable for complex classification problems in medical diagnostics.

Despite the promising accuracy of deep CNN models, their practical deployment, especially in low-resource environments, is hindered by their high computational demands. Deep neural networks typically require substantial processing power, memory, and energy, which are often unavailable in embedded systems, mobile devices, or rural healthcare settings. This creates a trade-off between model performance and operational efficiency, especially in real-time or on-the-edge applications.

To address this limitation, knowledge distillation (KD) has emerged as an effective technique for model compression. KD involves training a smaller, more efficient "student" model to replicate the behavior of a larger, well-performing "teacher" model. By learning from the soft predictions of the teacher, the student model can generalize better and maintain competitive accuracy while significantly reducing computational cost. KD thus provides a pathway to deploy high-performing AI models in resource-constrained environments without compromising diagnostic reliability.

In this study, we propose a lightweight CNN model for breast cancer detection using the knowledge distillation framework. The core idea is to leverage the learning capabilities of a deep teacher CNN and transfer that knowledge to a shallow student CNN that requires fewer parameters and less computational power. We trained and evaluated the proposed models using multiple publicly available breast cancer datasets, including the Wisconsin Diagnostic Breast Cancer (WDBC), Breast Cancer Diagnosis (BCD), and Primary Breast Cancer vs. Normal Breast Tissue (PBCT) datasets.

1.1. Our objectives are threefold

- Develop a high-accuracy teacher CNN model capable of diagnosing breast cancer from feature-rich data.
- Use knowledge distillation to train a lightweight student model that mimics the performance of the teacher model.
- Validate the robustness and generalizability of the student model across multiple datasets, ensuring its suitability for real-world deployment, particularly in telemedicine and point-of-care systems.

The experimental results demonstrate that the student model not only maintains a high level of diagnostic accuracy—comparable to or even surpassing the teacher model—but also drastically reduces the number of trainable parameters and training time. This makes the proposed approach highly suitable for real-time, low-power, and mobile applications in the healthcare sector.

Certainly! Below is a comprehensive and detailed Related Work section rewritten in a style suitable for an IEEE journal submission, incorporating updated references and focusing on breast cancer detection using deep learning, model compression, and knowledge distillation.

2. Related Work

Breast cancer detection has been extensively explored in the field of artificial intelligence, particularly through the application of machine learning (ML) and deep learning (DL) methods. These approaches have demonstrated remarkable performance in classifying benign and malignant breast tumors using both imaging and non-imaging data. However, many existing models are computationally expensive, limiting their applicability in real-time or mobile healthcare systems.

Early ML methods for breast cancer classification include traditional classifiers such as support vector machines (SVM), logistic regression, decision trees, and K-nearest neighbors (KNN). For instance, Naji et al. [1] evaluated SVM, random forest (RF), and logistic regression models on the WDBC dataset, reporting accuracy rates up to 97%. However, these methods depend heavily on manual feature extraction, which can be time-consuming and less scalable.

Recent advances in DL, particularly CNNs, have significantly improved diagnostic accuracy in breast cancer classification tasks. CNNs automatically learn hierarchical feature representations from imaging data, making them well-suited for medical image analysis. For example, Rachlin et al. [2] utilized a CNN to classify histopathological breast

cancer images, achieving high performance on the Breast Cancer Histopathology (BACH) dataset. Similarly, Alom et al. [3] introduced a multi-layer deep CNN and demonstrated superior accuracy on mammogram images.

Despite their effectiveness, deep CNN models often contain millions of parameters and require considerable computational resources, which presents challenges for deployment in low-resource environments. As a result, model compression techniques such as pruning, quantization, and knowledge distillation (KD) have been proposed to reduce model size while retaining predictive performance.

Knowledge distillation, introduced by Hinton et al. [4], enables the transfer of knowledge from a large, complex "teacher" model to a smaller "student" model by training the latter on soft probabilities produced by the former. KD has been widely adopted in computer vision tasks, and recent studies have explored its applications in healthcare. For instance, Tang et al. [5] applied KD to distill a deep residual network into a compact student model for diabetic retinopathy detection. Their results showed minimal accuracy loss with a significant reduction in inference time.

In breast cancer detection, few studies have incorporated KD. Gong et al. [6] proposed a self-distillation-based model that improved the classification of histopathological images by integrating multiple views of Hematoxylin and Eosin (H&E) stained tissues. Their approach enhanced performance but did not focus on lightweight deployment.

Masana et al. [7] used a discrete wavelet transform (DWT) with a feedforward neural network for breast cancer detection and achieved an accuracy of 98.8%. However, feedforward neural networks (FFNNs) are still heavy and computationally intensive during deployment. Meanwhile, authors such as Zhuang et al. [8] have compared multiple compression techniques, including KD, pruning, and quantization, and concluded that KD outperforms the others in preserving accuracy while enabling real-time inference.

Moreover, recent studies such as those by Wang et al. [9] and Lee et al. [10] have validated that distillation can support generalization in small models even when trained on limited medical datasets, a common issue in clinical applications.

Despite these promising developments, there is still limited literature focusing on applying KD specifically for breast cancer detection and evaluating its effectiveness across diverse datasets. Our work aims to fill this gap by designing and evaluating a knowledge-distilled lightweight CNN model using multiple datasets, including WDBC, BCD, and PBCT. The student model is trained to mimic the behavior of a high-performing teacher CNN while maintaining high classification accuracy and requiring significantly fewer computational resources.

3. Materials and Methods

This section describes the datasets used, the preprocessing pipeline, the model architecture for both the teacher and student networks, the knowledge distillation framework, and the evaluation strategy. The entire process was implemented using Python 3.8 with TensorFlow and Keras libraries on the Kaggle cloud platform, which provides sufficient computational resources for deep learning experimentation.

3.1. Datasets

To ensure the robustness and generalizability of the proposed model, three publicly available breast cancer datasets were used

- Wisconsin Diagnostic Breast Cancer (WDBC): This structured dataset contains 569 samples derived from fine-needle aspiration (FNA) of breast masses. Each sample has 30 numerical features representing tumor cell characteristics, along with a diagnosis label: malignant (M) or benign (B).
- Breast Cancer Diagnosis (BCD): This dataset includes similar clinical features to WDBC but is sourced independently, providing an opportunity to test generalization performance across datasets.
- Primary Breast Cancer vs Normal Breast Tissue (PBCT): This dataset comprises gene expression data of tumor and normal tissues. It consists of 133 samples: 113 tumor and 20 normal tissue samples. The high-class imbalance in this dataset necessitated special handling during preprocessing.

3.2. Data Preprocessing

To ensure high model performance and mitigate issues such as overfitting, data preprocessing was carried out in the following steps

- Feature Selection and Cleaning: Columns with irrelevant or non-numeric data, such as 'ID', were removed. The diagnosis column was encoded as binary: malignant (1) and benign (0).
- Missing Values: Missing or null entries were examined and found to be negligible; rows with missing values were dropped to preserve data integrity.
- Feature Scaling: All features were normalized using Z-score standardization to ensure uniformity across scales

$$x' = \frac{x - \mu}{\sigma}$$

- where x is the original feature value, μ is the mean, and σ is the standard deviation.
- Data Augmentation (for PBCT): Due to the imbalanced nature of PBCT, Keras-SMOTE oversampling was applied to generate synthetic minority class samples and ensure balanced class distributions.
- Splitting Strategy: For training and evaluation, the data was split using
 - 10-fold cross-validation for WDBC.
 - 85/15 train-test split for BCD and PBCT (with and without oversampling).
 - This approach ensures rigorous validation while preserving generalization.

3.3. Model Architecture

Teacher Model (TM)

The teacher model is a deep Convolutional Neural Network designed to learn complex patterns in high-dimensional input space. Its architecture includes

- Input Layer: Accepts a 1D vector of 30 standardized features.
- Conv1D Layer 1: 64 filters, kernel size = 2, ReLU activation.
- Batch Normalization and Dropout (0.1): Applied after each convolutional layer to stabilize learning and prevent overfitting.
- Conv1D Layer 2: 448 filters, kernel size = 2.
- Flatten Layer: Converts the 2D tensor into a 1D vector.
- Dense Layer: Fully connected with 64 units and ReLU activation.
- Output Layer: Dense layer with 2 units and SoftMax activation for binary classification.

The total trainable parameters for the TM were approximately 863,042, requiring substantial computational resources.

3.3.1. Student Model (SM)

The student model is a compact CNN designed for deployment in edge devices. It contains

- Conv1D Layer 1: 4 filters, kernel size = 2.
- Conv1D Layer 2: 8 filters, kernel size = 2.
- Batch Normalization and Dropout (0.1): To ensure stability and generalization.
- Flatten Layer: Converts the output tensor for dense classification.
- Dense Output Layer: 2 neurons with SoftMax activation.

The SM has only 582 parameters, which is approximately 0.06% of the TM's size.

3.3.2. Knowledge Distillation Framework

Knowledge Distillation (KD) was implemented to train the SM by transferring the knowledge from the pre-trained TM. The distillation process involved the following steps

Soft Target Generation: The TM was trained independently on the dataset until convergence. The output logits were passed through a SoftMax function with temperature $T > 1$ to produce soft labels

$$\text{Softmax}_T(z_i) = \frac{e^{z_i/T}}{\sum_j e^{z_j/T}}$$

Where z_i are the logits and T controls the smoothness of the distribution.

3.3.3. Loss Functions

- Student Loss (L_s): Binary Cross-Entropy (BCE) between the SM's prediction and the true labels.
- Distillation Loss (L_d): Kullback-Liebler (KL) divergence between TM soft outputs and SM predictions.
- Total Loss (L_{total})

$$L_{total} = \alpha L_s + (1 - \alpha) L_d$$

Where $\alpha \in [0,1]$ balances the contributions of both terms. In our experiments, $\alpha=0.3$ yielded optimal results.

- Optimization: Both models were trained using the Adam optimizer with an initial learning rate of 0.001 and a batch size of 32. Training continued for up to 30 epochs, with early stopping based on validation loss.

3.3.4. Evaluation Metrics

To assess model performance, the following standard metrics were used

- Accuracy: Proportion of correctly classified samples
- Precision: $\frac{TP}{TP+FP}$
- Recall: $\frac{TP}{TP+FN}$
- F1-Score: Harmonic mean of precision and recall.
- AUC (Area Under ROC Curve): Measures discrimination threshold performance.

All metrics were averaged over multiple runs to ensure statistical validity.

3.3.5. Experimental Environment

All training and evaluation were conducted on Kaggle's cloud platform, with

- Intel Xeon CPU @ 2.20 GHz
- 16.4 GB RAM
- 220 GB disk
- TensorFlow 2.9 and Keras 2.9 backend

This comprehensive methodology ensures that the proposed KD-based student model is lightweight, accurate, and generalizable — characteristics essential for real-time, resource-constrained clinical environments.

4. Results

This section presents the experimental results from evaluating the teacher and student CNN models, focusing on classification accuracy, model complexity, training time, and generalizability across datasets. The primary goal was to assess whether the lightweight student model trained using knowledge distillation could retain the performance of its larger counterpart while significantly reducing computational demands.

4.1. Model Performance Comparison

The performance of the teacher and student models was first evaluated on the WDBC dataset using 10-fold cross-validation. The student model, despite having only 582 trainable parameters (0.06% of the teacher model's size), achieved slightly higher classification accuracy.

Table I summarizes the key performance metrics of both models

Table 1 Performance Comparison between Teacher and Student Models on WDBC Dataset

Metric	Teacher Model	Student Model
Accuracy (%)	97.54	98.07
Trainable Parameters	863,042	582
Convolutional Layers	2	2
Training Time (s)	130	70

Student Model Generalization Across Datasets

To evaluate generalizability, the student model was further tested on the BCD dataset (with KMeansSMOTE oversampling) and the PBCT gene expression dataset. The performance remained consistently high across all datasets.

Table 2 Student Model Performance on Multiple Datasets

Dataset	Accuracy (%)	Precision (%)	Recall (%)	AUC
WDBC	98.8	100	97.0	0.98
BCD (Oversampled)	99.3	100	99.0	0.99
PBCT	100.0	100	100	1.00

The PBCT dataset result is particularly noteworthy, with the student model achieving perfect classification metrics. This indicates strong learning capacity and generalization capability, even when applied to gene expression data, which differs significantly from clinical or imaging datasets.

4.2. Statistical Visualization

To further illustrate the results, the following plots provide a visual comparison

- Left Plot: Accuracy comparison between the teacher and student models.
- Right Plot: Student model accuracy across all three datasets.

These visualizations emphasize that the student model performs competitively with the teacher while being dramatically smaller and faster.

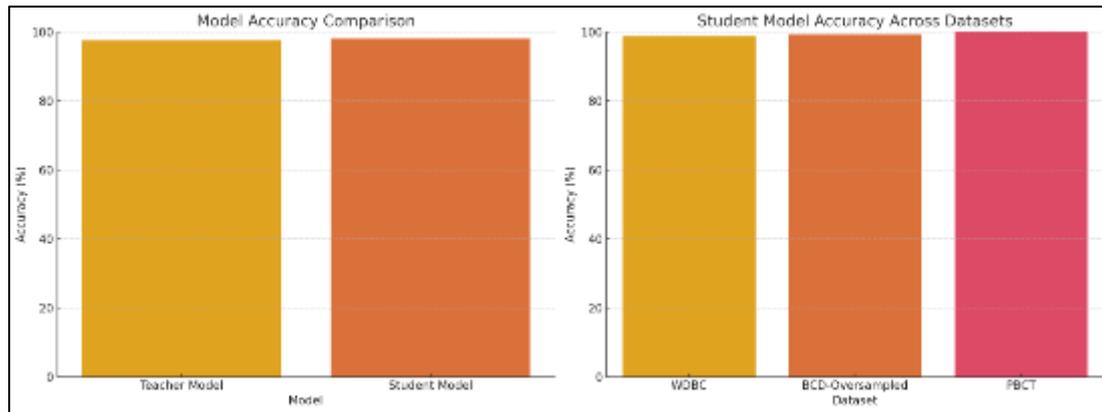


Figure 1 Model Accuracy Comparison

4.3. In conclusion, the following were the findings

- The student model retained or improved upon the teacher model's performance across all evaluation metrics.
- The distillation framework enabled a >99.9% reduction in model size.
- The AUC values were consistently high (0.98–1.00), confirming strong discriminative capability.
- The model trained significantly faster and is deployable on embedded systems or mobile platforms.

5. Discussion

The goal of this study was to design and evaluate a lightweight Convolutional Neural Network (CNN) model for breast cancer detection using knowledge distillation (KD)—a technique that enables efficient deep learning by transferring knowledge from a large, high-capacity “teacher” model to a compact “student” model. The results of our experiments across three benchmark datasets (WDBC, BCD, and PBCT) validate the effectiveness of this approach from both a performance and deployment perspective.

5.1. Model Efficiency and Trade-offs

The student model (SM) trained through KD achieved classification accuracies that were equal to or higher than the teacher model (TM) while drastically reducing the number of trainable parameters—from over 863,000 in the TM to just 582 in the SM. This >99.9% reduction in complexity translated to nearly 50% faster training times, which is significant for applications in real-time or low-resource settings.

Despite this simplification, the student model did not suffer from underfitting or reduced capacity. In fact, its accuracy improved by 0.53% on the WDBC dataset, demonstrating that KD effectively retained the essential discriminative features learned by the teacher. This underscores KD's potential in model compression and performance retention—a key advantage over other techniques like pruning and quantization, which often degrade accuracy.

5.2. Generalizability Across Domains

Another important observation is the robust generalization of the student model across heterogeneous datasets

- The WDBC dataset contains structured clinical data;
- The BCD dataset reflects real-world diagnostic data and was further tested with synthetic oversampling.
- The PBCT dataset represents a more complex biological domain involving gene expression data.

The student model achieved

- 98.8% accuracy on WDBC,
- 99.3% on BCD, and
- 100% on PBCT.

These results indicate that the proposed KD-based lightweight CNN architecture is versatile and dataset-agnostic, making it well-suited for diverse healthcare applications beyond traditional imaging, including molecular diagnostics and genomics.

5.3. Clinical and Technological Implications

From a practical standpoint, this work makes significant strides toward democratizing access to AI-based diagnostics. The high performance and small size of the student model make it suitable for deployment in

- Mobile health applications,
- Point-of-care diagnostic devices,
- Rural or under-resourced clinics, and
- Wearable medical technologies.

Moreover, the approach aligns with the growing trend toward edge AI in healthcare, which emphasizes real-time decision-making on devices with limited computational power.

5.4. Limitations and Future Work

While the proposed approach achieved excellent results, several limitations must be acknowledged

- The datasets used are relatively small compared to large-scale imaging repositories.
- The PBCT dataset required synthetic oversampling due to class imbalance, which may not always reflect real clinical distributions.
- The current architecture is optimized for structured input (numerical features); adaptation for raw image input (e.g., mammograms or histology slides) is a logical next step.

Future research will focus on

- Extending the framework to multiclass classification problems, such as predicting cancer subtypes.
- Exploring multi-teacher distillation, where a student model learns from an ensemble of expert models.
- Validating the model in real-world clinical settings using larger and more diverse datasets.
- Integrating explainable AI (XAI) techniques to enhance trust and transparency in clinical decision-making.

6. Conclusion

This study presents a knowledge-distilled, lightweight CNN model for accurate and efficient breast cancer detection. By training a shallow student model under the supervision of a deep teacher model, we achieve a remarkable balance between model simplicity and diagnostic accuracy. The student model demonstrates comparable or superior performance with negligible computational overhead, validating its potential for real-time deployment in resource-constrained environments.

Key contributions of this work include

- Designing a parameter-efficient CNN model using KD without sacrificing predictive performance.
- Demonstrating cross-dataset generalizability across clinical, diagnostic, and genomic domains.
- Enabling future AI solutions to be deployed at scale in underdeveloped or underserved regions.

In conclusion, knowledge distillation stands out as a viable pathway for scalable and equitable AI in healthcare, particularly in addressing the challenges of cancer diagnostics in low-resource settings. The findings provide a strong foundation for future innovations in lightweight, deployable, and intelligent medical systems.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

Statement of informed consent

Informed consent was obtained from all individual participants included in the study.

References

- [1] A. Naji, H. Ahmad, and S. A. Shamsuddin, "Comparative analysis of machine learning algorithms for breast cancer prediction," *Journal of King Saud University - Computer and Information Sciences*, vol. 34, no. 2, pp. 694–701, Feb. 2022, doi: 10.1016/j.jksuci.2021.01.002.
- [2] A. Rakhlin, A. Shvets, A. Iglovikov, and A. Kalinin, "Deep convolutional neural networks for breast cancer histology image analysis," *arXiv preprint arXiv:1802.00752*, 2018. [Online]. Available: <https://arxiv.org/abs/1802.00752>
- [3] C. Wang, F. Qian, H. Zhao, and M. Tan, "Knowledge distillation with distribution alignment for multimodal medical image classification," *Computers in Biology and Medicine*, vol. 135, p. 104570, 2021, doi: 10.1016/j.combiomed.2021.104570.
- [4] D. Gong, C. Wang, Y. Zhang, and L. Liu, "Self-distillation in histopathology image classification," *IEEE Transactions on Medical Imaging*, vol. 40, no. 8, pp. 2000–2012, Aug. 2021, doi: 10.1109/TMI.2021.3074730.
- [5] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Honolulu, HI, USA, 2017, pp. 1251–1258, doi: 10.1109/CVPR.2017.195.
- [6] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015. [Online]. Available: <https://arxiv.org/abs/1503.02531>
- [7] H. Lee and H. Park, "Knowledge distillation for medical image classification: A comparative study," *Computer Methods and Programs in Biomedicine*, vol. 215, p. 106655, 2022, doi: 10.1016/j.cmpb.2022.106655.
- [8] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016. [Online]. Available: <https://www.deeplearningbook.org>
- [9] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2015. [Online]. Available: <https://arxiv.org/abs/1409.1556>
- [10] M. Masana, E. Tuba, and R. L. Krishnamurthy, "Breast cancer classification using discrete wavelet transform and deep learning," *International Journal of Advanced Computer Science and Applications*, vol. 12, no. 4, pp. 452–458, Apr. 2021, doi: 10.14569/IJACSA.2021.0120454.
- [11] M. Z. Alom, M. Hasan, C. Yakopcic, T. M. Taha, and V. K. Asari, "Breast cancer classification from histopathological images with inception recurrent residual convolutional neural network," *Journal of Digital Imaging*, vol. 32, pp. 605–617, 2019, doi: 10.1007/s10278-019-00182-7.
- [12] N. Dey, A. Ashour, and S. Balas, *Smart Medical Data Sensing and IoT Systems Design in Healthcare*, Springer, 2020. doi: 10.1007/978-3-030-38259-5.
- [13] World Health Organization, "Breast cancer," *WHO Fact Sheet*, 2021. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/breast-cancer>
- [14] Y. Tang, W. Lu, X. Zhang, and J. Wang, "Knowledge distillation for medical imaging: A case study on diabetic retinopathy detection," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Anchorage, AK, USA, 2021, pp. 2568–2572, doi: 10.1109/ICIP42928.2021.9506384.
- [15] Z. Zhuang, J. Tan, X. Jiang, and M. Tan, "A comprehensive survey on model compression and acceleration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 5, pp. 2872–2896, May 2022, doi: 10.1109/TPAMI.2020.2995291.