

# Designing a Deep Learning-Driven Resource-Efficient Diagnostic System for Metastatic Breast Cancer: Reducing Long Delays of Clinical Diagnosis and Improving Patient Survival in Developing Countries

Cancer Informatics  
Volume 22: 1–8  
© The Author(s) 2023  
Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/11769351231214446



William Gao<sup>1</sup> , Dayong Wang<sup>2</sup> and Yi Huang<sup>1</sup>

<sup>1</sup>Department of Mathematics and Statistics, University of Maryland, Baltimore County, Baltimore, MD, USA. <sup>2</sup>Meta Platforms Inc., Boston, MA, USA.

**ABSTRACT:** Breast cancer is one of the leading causes of cancer mortality. Breast cancer patients in developing countries, especially sub-Saharan Africa, South Asia, and South America, suffer from the highest mortality rate in the world. One crucial factor contributing to the global disparity in mortality rate is long delay of diagnosis due to a severe shortage of trained pathologists, which consequently has led to a large proportion of late-stage presentation at diagnosis. To tackle this critical healthcare disparity, we have developed a deep learning-based diagnosis system for metastatic breast cancer that can achieve high diagnostic accuracy as well as computational efficiency and mobile readiness suitable for an under-resourced environment. We evaluated 4 Convolutional Neural Network (CNN) architectures: MobileNetV2, VGG16, ResNet50 and ResNet101. The MobileNetV2-based diagnostic model outperformed the more complex VGG16, ResNet50 and ResNet101 models in diagnostic accuracy, model generalization, and model training efficiency. The ROC AUC of MobileNetV2 (0.933, 95% CI: 0.930, 0.936) was higher than VGG16 (0.911, 95% CI: 0.908, 0.915), ResNet50 (0.869, 95% CI: 0.866, 0.873), and ResNet101 (0.873, 95% CI: 0.869, 0.876). The time per inference step for the MobileNetV2 model (15ms/step) was substantially lower than that of VGG16 (48ms/step), ResNet50 (37ms/step), and ResNet110 (56ms/step). The visual comparisons between the model prediction and ground truth have demonstrated that the MobileNetV2 diagnostic models can identify very small cancerous nodes embedded in a large area of normal cells which is challenging for manual image analysis. Equally important, the light weight MobileNetV2 models were computationally efficient and ready for mobile devices or devices of low computational power. These advances empower the development of a resource-efficient and high performing AI-based metastatic breast cancer diagnostic system that can adapt to under-resourced healthcare facilities in developing countries.

**KEYWORDS:** Breast cancer diagnosis, deep learning, resource efficiency

**RECEIVED:** May 2, 2023. **ACCEPTED:** October 27, 2023.

**TYPE:** Original Research

**FUNDING:** The author(s) received no financial support for the research, authorship, and/or publication of this article.

**DECLARATION OF CONFLICTING INTERESTS:** The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

**CORRESPONDING AUTHOR:** Dayong Wang, Meta Platforms Inc., 1601 Willow Road, Boston, MA 94025-1452, USA. Email: dayong.wangts@gmail.com

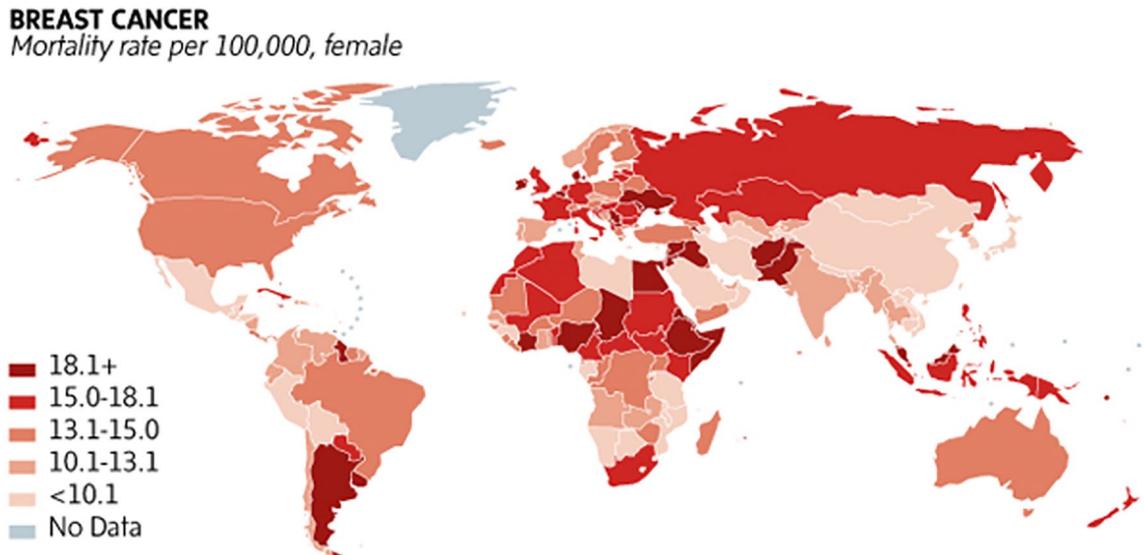
## Introduction

Breast cancer is the most diagnosed cancer among women and the fifth leading cause of cancer mortality, accounting for over 2.3 million new cases and 6.9% of all cancer deaths.<sup>1</sup> Breast cancer patients in developing countries, especially sub-Saharan Africa, South Asia, and South America, suffer from the highest mortality rates in the world (Figure 1). One crucial factor contributing to the global disparity in mortality rate is a long diagnosis delay. Research on Turn Around Time (TAT)—the time from biopsy till receipt of report by clinician—from Malawi showed that the median TAT for unpaid samples was 71 days. At 2 facilities in Rwanda, the delay between the initial development of symptoms and the ultimate receipt of a diagnosis could reach 15 months.<sup>2</sup> Consequently, the long diagnosis delay has led to a large proportion of late-stage presentation at diagnosis. An analysis of 83 studies across 17 sub-Saharan African countries reported that 77% of all staged cases were stage III/IV at diagnosis.<sup>3</sup> The long diagnosis delay is primarily attributed to a severe shortage of trained pathologists. The average number of pathologists per head of population is 1 to 1 000 000 in sub-Saharan regions, compared with the ratio of 1 pathologist to 15 000 to 20 000 in the US and UK.<sup>4</sup> Another

challenge is access to healthcare facilities. In sub-Saharan Africa, >170 million people are >2 hours from the nearest regional hospital, and about 40% of the population is >4 hours away from a national hospital.<sup>5</sup> Additionally, diagnosis quality may be impacted by a lack of resources and training. A retrospective assessment of reports of breast carcinoma in the university teaching hospital in Lagos, Nigeria showed a 46.9% discordance rate in the basic diagnosis when the cases were examined by pathologists in the UK.<sup>2</sup>

Artificial intelligence (AI) presents a new opportunity for advancing medicine and healthcare.<sup>6</sup> Currently, pathologists manually review histopathology slides to examine whether cancerous regions are present in tissue samples. However, manual review is time-consuming and subject to human error, especially in healthcare facilities in resource-constrained regions.<sup>2</sup> AI has a great potential to bring changes to cancer diagnosis by providing faster, more accurate, and more robust technological solutions. Machine learning-driven image analysis can assist pathologists by narrowing their search for high-risk areas on the whole slide of a biopsy sample and making their diagnosis more accurate and efficient. Progress has been made to develop machine learning-driven image analysis methods for breast





**Figure 1.** World map of breast cancer mortality rates (all ages), showing high mortality rates in sub-Saharan Africa, South Asia, and South America.<sup>1</sup>

cancer diagnosis.<sup>7-9</sup> However, real-world clinical application of these methods, especially in remote and under-resourced areas, remains a major challenge. The complex neural network architecture, computational cost, and lack of mobile capacity make the existing methods less feasible and adaptive to a geographically distributed and under-resourced environment.<sup>10</sup> An essential step toward better model computational efficiency is to reduce CNN architecture complexity while maintaining classification performance. To develop a deep learning-driven resource-efficient diagnostic system for metastatic breast cancer, we compared four CNN architectures, that is, MobileNetV2, VGG16, ResNet50 and ResNet101 to assess classification performance and computational efficiency. Our objective is to identify an appropriate CNN architecture that makes the diagnostic system suitable for remote and under-resourced health-care facilities by pursuing computational efficiency and mobile readiness while achieving high diagnostic accuracy.

## Methods

### *Framework of deep learning-driven diagnostic system for metastatic breast cancer*

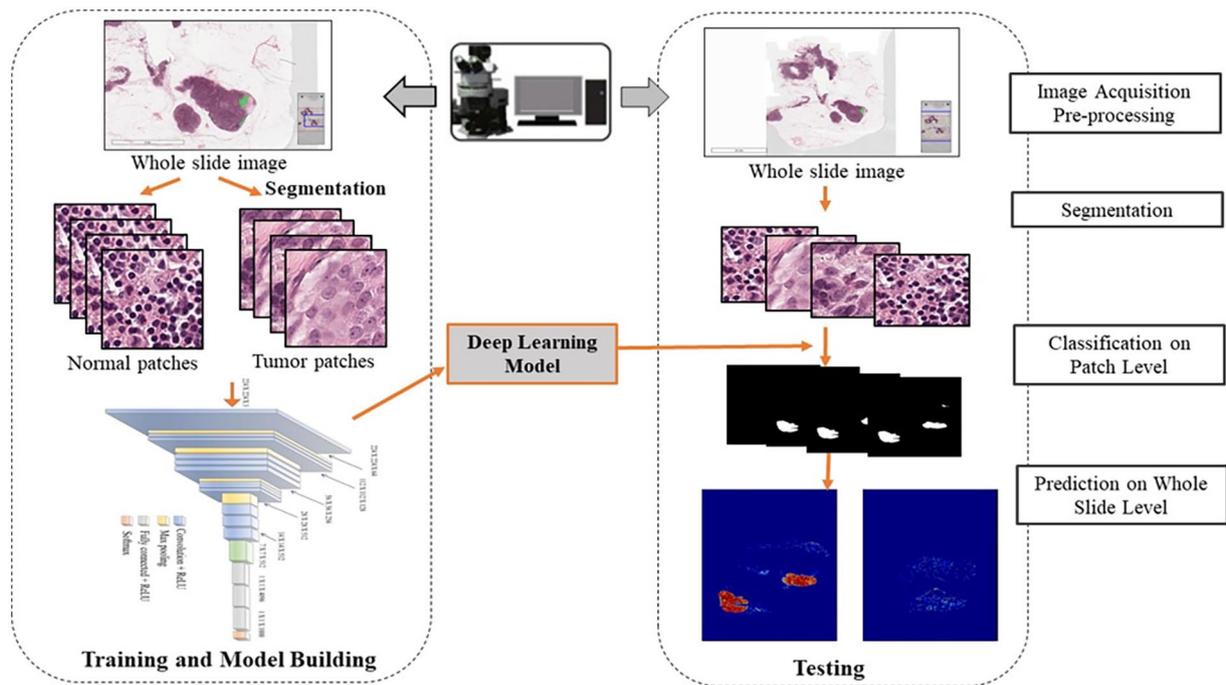
The deep learning-driven diagnostic system for metastatic breast cancer consists of 3 major components: histopathological image pre-processing, training to build deep learning-driven diagnostic models, and testing for prediction on unseen whole slide images (WSIs) (Figure 2).

The deep CNN model training for metastatic breast cancer was based on more than two hundred thousand (200K) training patches that were segmented from the training set of histopathological whole slide images (WSIs). For testing, the trained model made patch-level predictions to classify tumor patches versus normal patches. Then the patch-level predictions were aggregated to generate tumor probability heatmaps for the

whole slide-level classification. In addition, we used transfer learning based on ImageNet pre-training.<sup>11</sup> The parameters in the pre-trained deep CNN models have a good initiation for image data, therefore reducing the requirement of training histopathological dataset size for metastatic breast cancer. In model training, we used a fixed learning rate of 0.01 and a dropout value of 0.2 for regularization. Each model ran through 10 iterations or “epochs” with data batch sizes of 32. Model optimization was run based on the stochastic gradient descent (SGD) optimizer. We evaluated 4 CNN architectures with varying structures, including the number of layers, depth, filter size, etc.

*Histopathological image pre-processing.* The diagnostic models were trained based on the ground truth data, that is, whole slide images (WSI) of sentinel lymph nodes with a pathologist’s delineation of regions of metastatic cancer. The image normalization was conducted with the WSI Color Standardization procedure<sup>12</sup> to minimize potential variations in the color and intensity of Hematoxylin and Eosin (H&E) staining. Tissue areas within the normalized WSIs were identified and extracted using a threshold-based segmentation method.<sup>13</sup> The mask images were generated for model training from the HSV representation transformed from the original RGB images.<sup>14</sup>

*Patch-based diagnostic model building.* Model training used WSIs and the ground truth image annotation indicating the delineation of cancerous regions as input data. The WSIs were randomly extracted into a large number of small patches. Using the small patches, the deep learning model was trained to recognize cancerous cells on a small scale, improving data volume and model training efficiency. These patches extracted from each WSI were categorized as a positive tumor, negative tumor, and negative normal. A positive tumor patch was extracted



**Figure 2.** The workflow of deep learning-driven diagnostic system for metastatic breast cancer.

from a tumor slide, containing cancerous regions; a negative tumor patch was extracted from a tumor slide but did not contain a cancerous region; a negative normal patch was extracted from a noncancerous/normal slide. During patch extraction, masks were created using a lower and upper bound of pixel color. After reducing mask noise with the “opening” and “closing” morphology processes, the pixel values of each mask were used to categorize the patch.

Four Convolutional Neural Network architectures, that is, MobileNetV2, VGG16, ResNet50 and ResNet101, were used for building diagnostic models. The 4 CNN architectures were assessed for diagnosis performance and computational efficiency. The basic characteristics of the CNN architectures are summarized in section 2.2. The diagnostic models were trained to discriminate between cancerous and noncancerous patches using a large number of small positive and negative patches randomly extracted from the set of training WSIs. Five-fold cross-validation and independent (unseen) image data were used to evaluate diagnostic performance.

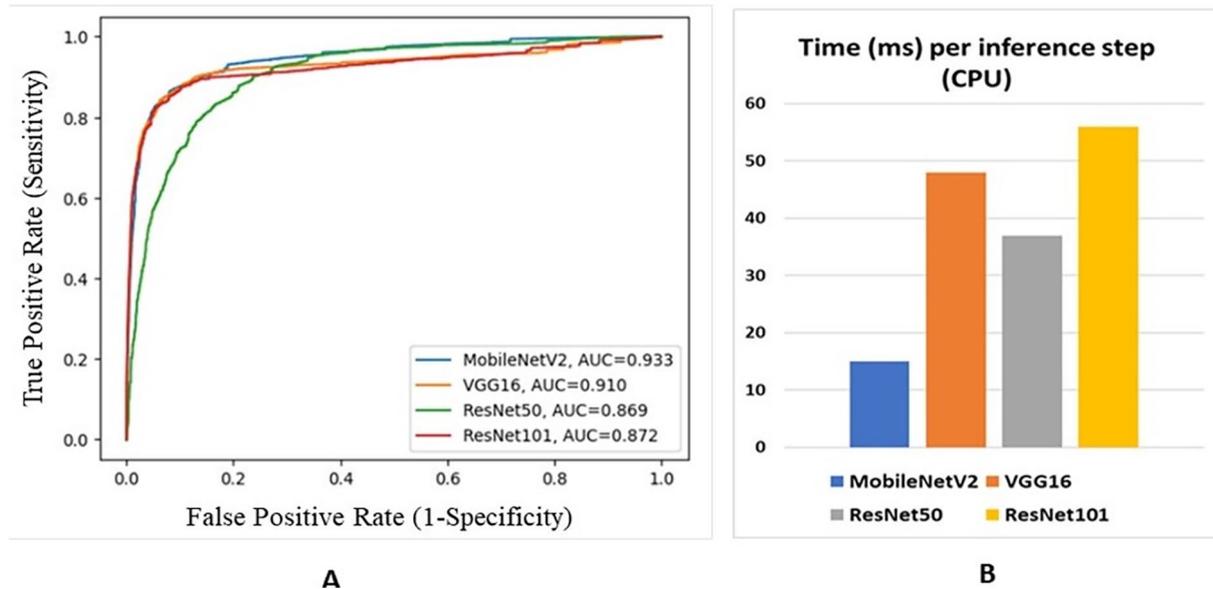
*Whole slide-based diagnosis.* As shown in Figure 2, unseen WSIs were pre-processed and segmented following the procedures described above. The diagnostic models were used to discriminate the cancerous versus non-cancerous patches. With the patch-level classification, a tumor probability heatmap was generated for each WSI to indicate the probability of each pixel being cancerous, with each pixel containing a value between 0 and 1, indicating the probability that the pixel contains tumor.

*Diagnostic performance evaluation.* Diagnostic performance was evaluated using accuracy and receiver operating characteristics

(ROC) curve. Accuracy is defined as the number of correct predictions divided by total number of predictions. ROC curve plots true positive rate (sensitivity) versus false positive rate (1-specificity); Area under ROC curve (ROC AUC) provides an aggregate measure of performance across all possible classification thresholds.

#### *Convolutional neural network architectures*

Four deep learning network architectures were evaluated for diagnostic performance for metastatic breast cancer and computational efficiency, including MobileNetV2, VGG16, ResNet50 and ResNet101. MobileNetV2 is a CNN architecture that aims to perform computer vision tasks efficiently on mobile devices. It is designed to achieve efficiency by incorporating the Inverted Residual Structure and the Depthwise Separable Convolution to significantly reduce the model size and complexity.<sup>15</sup> VGG16 is a CNN architecture that is designed for large-scale image recognition. VGG16 improves accuracy by increasing the depth to 16 weight layers using an architecture with very small ( $3 \times 3$ ) convolution filters, in combination with more non-linear activation layers for better discriminative decision function.<sup>16</sup> ResNet is a Residual Network, a category of CNN architectures that is intended to improve model accuracy by stacking layers to enrich the features of the model.<sup>17</sup> Resnet50 is a variant of ResNet whose neural network layers reach a depth of 50, with a bottleneck design to reduce the time taken to train the layers. ResNet101 is a large Residual Network with 101 neural network layers. In Keras Applications, MobileNetV2 uses 3.5 million (M) of parameters and has a size of 14MB; VGG16 uses 138.4M parameters and has a size



**Figure 3.** (A) Receiver Operating Characteristics (ROC) curves for independent set testing: comparison of ROC AUC among the CNN architectures MobileNetV2, VGG16, ResNet50 and ResNet101 and (B) comparison of time per inference step among the 4 CNN architectures.

of 528 MB; ResNet50 uses 25.6 M parameters and has a size of 98 MB; ResNet101 uses 44.7 M parameters and has a size of 171 MB.<sup>18</sup>

### Datasets

The data consists of 222 whole slide histopathological images from the 2016 Camelyon ISBM challenge.<sup>9</sup> The slides contain sentinel lymph node tissues extracted by the Radboud University Medical Center (Nijmegen, the Netherlands), as well as the University Medical Center Utrecht (Utrecht, the Netherlands).

## Results

### Evaluation of deep learning network architectures: Diagnostic accuracy and computational efficiency

To identify a CNN architecture that is suitable for a resource-constrained environment, 4 CNN architectures were evaluated for diagnostic performance and computational cost, including MobileNetV2, VGG16, ResNet50 and ResNet101. The diagnostic models were trained and assessed for accuracy using 5-fold cross-validation. For an individual fold, each model ran through 10 iterations or “epochs.” Each epoch allowed the model to reevaluate its weights to determine a more effective set of values. With a stochastic gradient descent (SGD) optimizer, the weights were retrained or optimized for the data with a 0.01 learning rate. Model accuracy was assessed for the training set and validation sets at each epoch. For training, the accuracy of MobileNetV2 (mean  $\pm$  standard deviation:  $0.963 \pm 0.005$ ) was higher than VGG16 ( $0.940 \pm 0.006$ ), ResNet50 ( $0.847 \pm 0.015$ ) and ResNet101 ( $0.851 \pm 0.015$ ). Similarly, the cross-validation results showed that the accuracy of MobileNetV2 ( $0.885 \pm 0.043$ ) was similar with VGG16

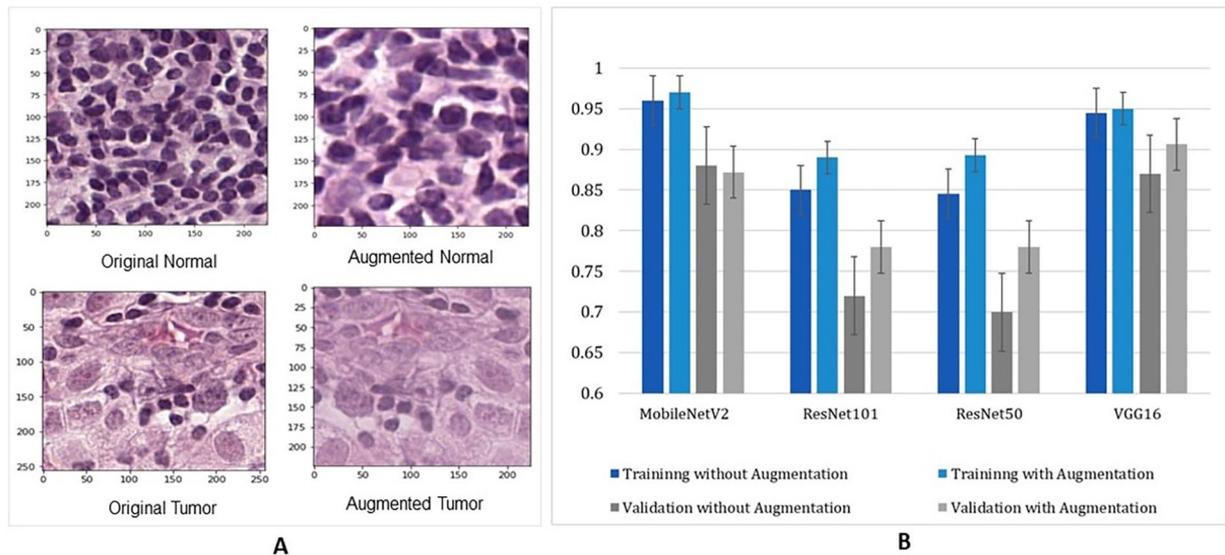
( $0.882 \pm 0.058$ ) but higher than ResNet50 ( $0.701 \pm 0.072$ ) and ResNet101 ( $0.707 \pm 0.076$ ).

The independent testing of the 4 diagnostic models with different CNN architectures was conducted using an independent testing set. The accuracies of MobileNetV2 (0.903) and VGG16 (0.898) were substantially higher than ResNet50 (0.721) and ResNet101 (0.857). Receiver operating characteristic (ROC) analysis was performed and ROC Area Under curve (AUC) was used to measure the diagnostic performance. As shown in Figure 3A, the ROC AUC of MobileNetV2 (0.933, 95% confidence interval (CI): 0.930-0.936) was higher than VGG16 (0.911, 95% CI: 0.908-0.915), Resnet50 (0.869, 95% CI: 0.866-0.873), and Resnet101 (0.873, 95% CI: 0.869-0.876).

The models were also evaluated for time per inference step, a measure of computational efficiency (Figure 3B). The time per inference step for the MobileNetV2 model was 15 ms/step, which was substantially lower than that of VGG16 (48 ms/step), ResNet50 (37 ms/step), and ResNet110 (56 ms/step). The result suggests that the MobileNetV2 model was more computationally efficient than the models with more complex CNN architectures.

### Data augmentation: Assessment of model generalization

Model generalization capacity is important for a model to properly adapt to previously unseen data and to achieve reasonable prediction performance. It is noticeable that the accuracy values of both the MobileNetV2-based and VGG16-based models exceeded 0.880 in the cross-validation while those of ResNet50 and ResNet101 dropped substantially in the cross-validation as compared with training. The result implies that



**Figure 4.** (A) Image transformation for data augmentation and (B) effects of data augmentation on the accuracy of the classification models with different network structures.

the MobileNetV2-based and VGG16-based models had a stronger ability of model generalization as compared with ResNet50 and ResNet101. We used data augmentation techniques to further evaluate model generalization of these CNN architectures. Data augmentation involves training models on data that has undergone various transformations.<sup>19</sup> This experiment was performed to evaluate whether the models can properly generalize to new instances. The image data transformation for data augmentation included 20-degree rotation, 20% zoom, horizontal flip, and vertical flip (Figure 4A).

As shown in Figure 4B, with or without data augmentation, the accuracies of MobileNetV2 and VGG16 models were substantially higher than those of ResNet50 and ResNet101 on both training and validation. Additionally, the models based on MobileNetV2 and VGG16 showed a substantially smaller gap in accuracy between training and validation as compared with other models, indicating a stronger generalization ability of the MobileNetV2 and VGG16 models. Data augmentation only caused slight changes to the accuracies of the MobileNetV2 and VGG16 models, while it drastically increased the accuracies of ResNet50 & ResNet101. These results suggest that the MobileNetV2 and VGG16 models were less reliant on bigger data volumes for achieving their peak performance.

#### *Predictive diagnosis on unseen whole slide images*

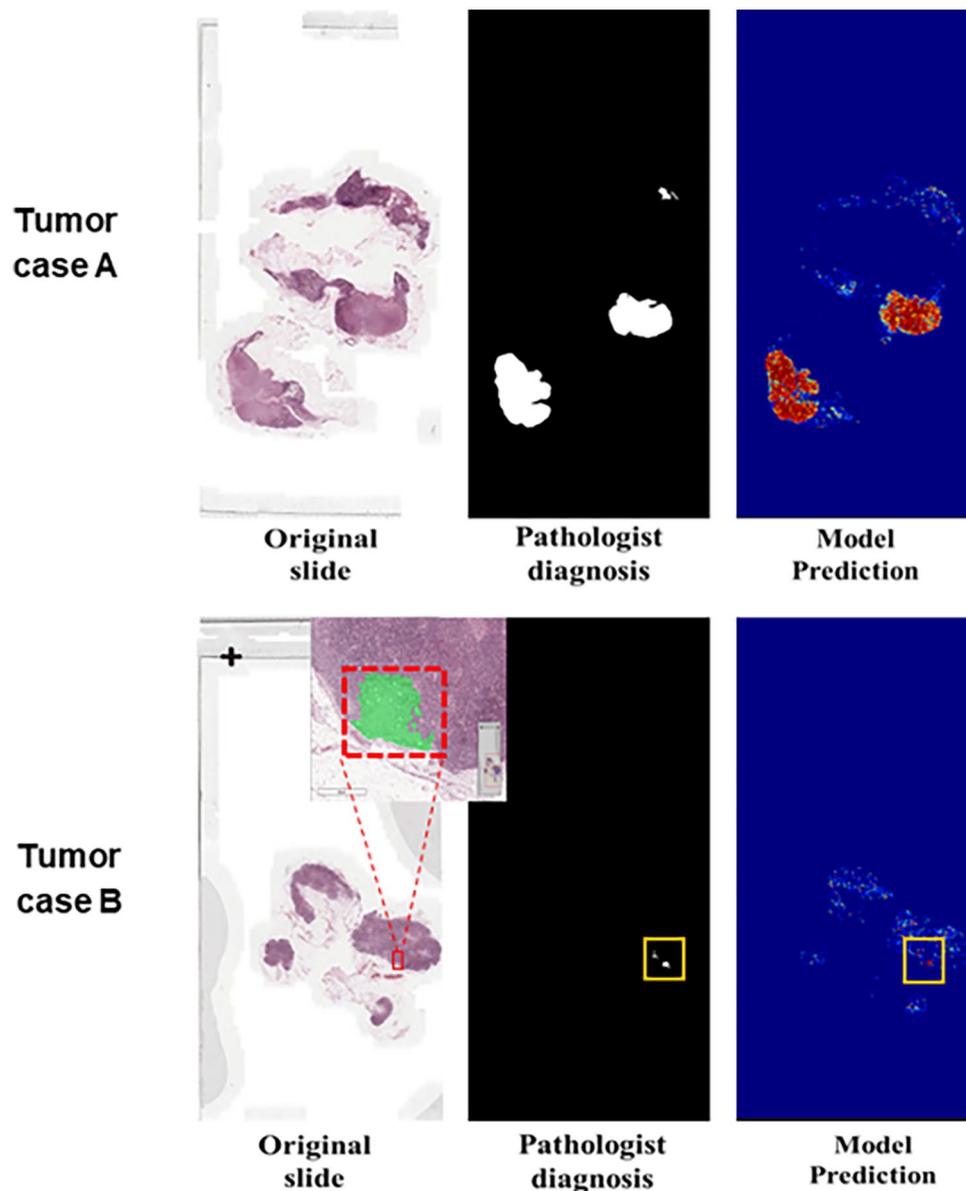
The MobileNetV2-based diagnostic model classified small image patches to cancerous versus non-cancerous with high accuracy. By iterating the model over each patch in the whole slide image, a probability rating was generated for each patch. After completion of the patch-based classification stage, the patch-level predictions were aggregated to create tumor probability heatmap for whole slide-based prediction. The likelihood of cancer presence on a whole slide image was shown on

the heatmap, where a pixel represents a patch, and a value represents the probability of the patch containing cancer with the range of 0 to 1. The threshold of 0.9 was used to control potential false positives.

Visual comparisons between the MobileNetV2 model prediction and the pathologist's diagnosis (ground truth) were shown in Figure 5. For tumor case A, the cancerous regions predicted by the model were consistent with those delineated by pathologists' diagnosis. Tumor case B was a more complicated case because the cancerous regions were small and embedded in a large area of normal cells. The MobileNetV2 model successfully identified the small cancerous node, and the prediction result was consistent with the pathologists' diagnosis. Identifying such a small cancerous region is challenging in the manual diagnosis process. The results demonstrated that the MobileNetV2-based diagnostic model's capacity in identifying the small regions of high risk and preventing false negative diagnoses.

#### **Discussion**

The recent research and development for computer-assisted cancer diagnosis has been focused on diagnostic accuracy.<sup>7-9</sup> Deep learning methods such as Convolutional Neural Networks have been used to develop automated diagnostic systems to improve the traditional manual assessment of histopathological images by pathologists. Research has shown that a deep learning-based diagnostic system can help a human pathologist to improve diagnostic accuracy.<sup>8</sup> Such a diagnostic system has a great potential to support under-resourced healthcare facilities, which often have severe shortages of trained pathologists, to expedite diagnosis. Despite the technical advances, the real-world application in developing countries is hindered by a lack of computational resources and general inadequacy of healthcare infrastructures. This

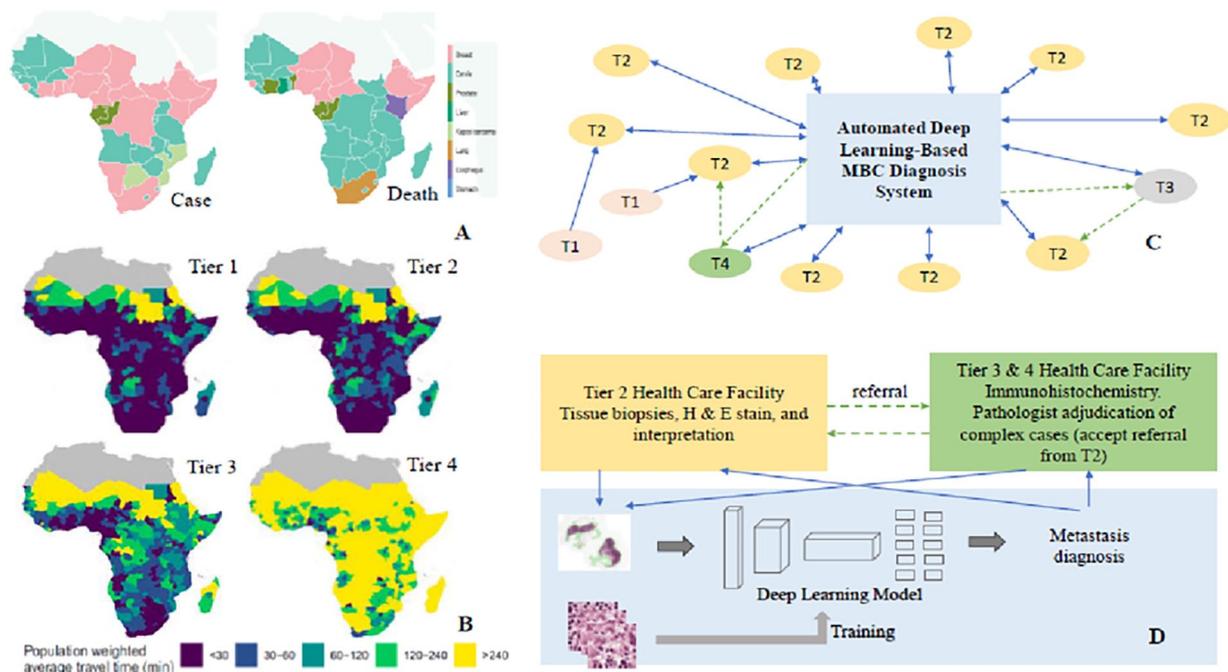


**Figure 5.** Visual comparisons between the pathologist's diagnosis (ground truth) and the model prediction on cancerous regions on Whole Slide Images.

research focuses on designing a diagnostic system with computational efficiency and mobile readiness while achieving high diagnostic accuracy.

MobileNetV2, a lightweight CNN architecture, is designed to achieve computational efficiency using the Inverted Residual Structure and the Depthwise Separable Convolution.<sup>15</sup> The model size and complexity of MobileNetV2 is substantially lower than VGG16, ResNet50, and ResNet101.<sup>18</sup> The high computational efficiency of MobileNetV2 makes it suitable for mobile devices or devices with low computational power. This research has demonstrated that the MobileNetV2-based diagnostic models outperformed its more complex counterparts VGG16, ResNet50 and ResNet101 in diagnostic performance and computational efficiency. The MobileNetV2 diagnostic models had a higher ROC AUC and a stronger model

generalization capacity as compared with VGG16, ResNet50 and ResNet101. The MobileNetV2 models reduced the time per inference step by 66.8%, 59.5%, and 73.2% as compared with VGG16, ResNet50, and ResNet101, respectively. The results justify the application of MobileNetV2 in an IT resource-constrained environment. Developing countries are experiencing severe shortages of trained pathologists, which causes long delays of metastatic breast cancer and high mortality rates. The average number of pathologists per head of population is 1 to 1 000 000 in sub-Saharan regions, compared with the ratio of 1 pathologist to 15 000 to 20 000 in the US and UK.<sup>4</sup> Meanwhile, diagnosis quality may be impacted by the limitations of resources and training. Therefore, the MobileNetV2-based diagnostic system provides an AI-driven, resource-efficient, automated, and standardized solution to



**Figure 6.** Prototypical schema for application of the deep learning-based automated diagnosis system in sub-Saharan Africa. (A) Sub-Saharan countries with high breast cancer cases (left) and high mortality rates (right) (both highlighted with pink). Breast cancer is the most common cancer in many sub-Saharan countries and has been the most common cause for cancer death in the Northern Sub-Saharan Africa.<sup>20</sup> (B) Accessibility to the nearest public health care facility (in travel hours), by Tier 1 to Tier 4 health care facility. Accessibility to public health care facilities is limited across the sub-Saharan region, especially in Northern sub-Saharan countries.<sup>5</sup> (C) The deep learning-based diagnosis system, highly efficient and mobile ready, can be deployed to all tiers of healthcare facilities. The collaborative diagnostic system can integrate knowledge from many institutions to support Tier 2 and possibly Tier 1 health care facilities for breast cancer diagnosis. (D) The system also provides connectivity between different tiers of health facility to expedite diagnosis of complex cases.

help relieve the severe shortages of trained pathologists and reduce the long diagnosis delays.

To tackle the public health crisis of long diagnosis delays in developing countries, a major challenge is inadequate accessibility to healthcare facilities. As an example, in sub-Saharan Africa, >170 million people are >2 hours from the nearest regional hospital, and about 40% of the population is >4 hours away from a national hospital (Figure 6B).<sup>5</sup> While the accessibility to health care facilities is limited across sub-Saharan Africa, this is particularly challenging in Northern sub-Saharan countries, lacking access to all tiers of health care facilities, including basic health post (Tier 1), health center (Tier 2), regional hospital (Tier 3) and central hospital (Tier 4) (Figure 6B). The countries with high breast cancer mortality and those with severe health care resource shortage are geographically overlapped in Northern sub-Sahara (Figure 6A and B), suggesting a strong link between the elevated breast cancer mortality and limited health care resources. To address this challenge, the MobileNetV2-based diagnosis system can be integrated into the local healthcare network to facilitate histopathological image processing and assist pathologists' diagnosis. The system can be deployed to Tier 3 regional hospitals and Tier 4 central hospitals to assist pathology services (Figure 6C). Because the MobileNetV2 system is suitable for devices of low computational power, this system

can also be deployed to Tier 2 and possibly Tier 1 healthcare facilities with limited IT resources. This system also connects lower tiers and high tiers healthcare facilities to expedite referrals for diagnosis of complex cases. This will shorten the delay in diagnosis, reduce the late-stage presentation at diagnosis, and therefore improve patients' survival outcomes.

## Conclusion

To tackle the healthcare disparities in metastatic breast cancer diagnosis, this research has developed a deep learning-based diagnosis system for metastatic breast cancer that aims to achieve high diagnostic accuracy as well as computational efficiency suitable for an under-resourced environment. We evaluated 4 CNN architectures: MobileNetV2, VGG16, ResNet50 and ResNet101. The MobileNetV2-based diagnostic model outperformed the more complex VGG16, ResNet50 and ResNet101 models in diagnostic accuracy, model generalization, and model training efficiency. The visual comparison between the model prediction and ground truth has demonstrated that the MobileNetV2 diagnostic models can identify very small cancerous nodes embedded in a large area of normal cells which is challenging for manual image analysis. Equally Important, the light weighted MobileNetV2 models were computationally efficient and ready for mobile devices or devices of low computational power. These advances empower

the development of a resource-efficient and high performing AI-based metastatic breast cancer diagnostic system that can adapt to under-resourced healthcare facilities in developing countries. This research provides an innovative technological solution to address the long delays in metastatic breast cancer diagnosis and the consequent disparity in patient survival outcome in developing countries.

### Author Contributions

W.G. developed the methods, conducted the analyses, and wrote the manuscript. D.W. and Y.H. provided advice on machine learning and statistical analysis. All authors reviewed and approved the manuscript.

### ORCID iD

William Gao  <https://orcid.org/0009-0000-2090-2595>

### REFERENCES

1. Sung H, Ferlay J, Siegel RL, et al. Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer J Clin.* 2021;71:209-249.
2. Jedy-Agba E, McCormack V, Adebamowo C, Dos-Santos-Silva I. Stage at diagnosis of breast cancer in sub-Saharan Africa: a systematic review and meta-analysis. *Lancet Glob Health.* 2016;4:e923-e935. doi:doi
3. Pace LE, Shulman LN. Breast cancer in Sub-Saharan Africa: challenges and opportunities to reduce mortality. *Oncologist.* 2016;21:739-744.
4. Fleming K. Pathology and cancer in Africa. *Ecancermedicalscience.* 2019;13:945.
5. Falchetta G, Hammad AT, Shayegh S. Planning universal accessibility to public health care in sub-Saharan Africa. *Proc Natl Acad Sci.* 2020;117:31760-31769.
6. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. *Nat Med.* 2022;28:31-38.
7. Freeman K, Geppert J, Stinton C, et al. Use of artificial intelligence for image analysis in breast cancer screening programmes: systematic review of test accuracy. *BMJ.* 2021;374: n1872.
8. Wang D, Khosla A, Gargeya R, et al. 2016. Deep learning for identifying metastatic breast cancer. arXiv: 1606.05718v1.
9. Veta M, Heng YJ, Stathonikos N, et al. Predicting breast tumor proliferation from whole-slide images: the TUPAC16 Challenge. *Med Image Anal.* 2019;54:111-121.
10. Wang Y, Wang J, Zhang W, et al. A survey on deploying mobile deep learning applications: a systemic and technical perspective. *Digit Commun Netw.* 2022;8:1-17.
11. Morid MA, Borjali A, Del Fiol G. A scoping review of transfer learning research on medical image analysis using ImageNet. *Comput Biol Med.* 2021;128: 104115.
12. Bejnordi BE, Litjens G, Timofeeva N, et al. Stain specific standardization of whole-slide histopathological images. *IEEE Trans Med Imaging.* 2016; 35:404-415.
13. Otsu N. A threshold selection method from gray-level histograms. *Automatica.* 1975;11:23-27.
14. Sural S, Qian G, Pramanik S. 2002. Segmentation and histogram generation using the HSV color space for image retrieval. Proceedings. 2002 International Conference on Image Processing, vol. 2. IEEE; 2002:II-589.
15. Sandler M, Howard A, Zhu M, et al. 2019. MobileNetV2: Inverted residuals and linear bottlenecks. arXiv: 1801.04381v4.
16. Simonyan K, Zisserman A. 2015. Very deep convolutional networks for large-scale image recognition. arXiv: 1409.1556v6.
17. He K, Zhang X, Ren S, et al. 2015. Deep residual learning for image recognition. arXiv: 1512.03385v1.
18. Keras API Reference. <https://keras.io/api/applications/> (November 16, 2023)
19. Shorten C, Khoshgoftaar TM. A survey on image data augmentation for deep learning. *J Big Data.* 2019;6:60.
20. American Cancer Society. 2019. The Cancer Atlas.