ASEC 2024 Conference
The 5th International Electronic Conference on Applied Sciences
04–06 December 2024 | Online
MDPI

Evaluating TPUs and GPUs in a Two-Views EfficientNet-based architecture for cancer classification on mammograms: performance and speed analysis

Sergio A. G. Fuentes, Daniel G. P. Petrini, Hae Y. Kim
Escola Politécnica, Universidade de São Paulo, Brazil

INTRODUCTION & AIM

Breast cancer is the most prevalent cancer among women worldwide. Mammography is the primary exam used to detect this disease in its early stages. Currently, radiologists interpret these radiological images, but CAD (Computer-Aided Detection and Diagnosis) systems have been developed to assist in this process.

Traditionally, GPUs (Graphics Processing Units) have been used to train these systems due to their high computational power. However, newer hardware, such as TPUs (Tensor Processing Units), offers significant advantages. TPUs are optimized for machine learning tasks and provide larger memory capacity than GPUs, enabling the processing of higher-resolution mammograms with less resizing needed. This capability is crucial, as resizing can result in the loss of critical details, such as microcalcifications, a type of câncer that often appear as subtle textures in mammographic images.

This study explores the use of TPUs for classifying mammograms, aiming to determine whether this hardware can enhance the accuracy of breast cancer detection by preserving the full resolution of input images.

METHOD

The first step involved converting the convolutional neural network (CNN) developed by Petrini et al. [1] for detecting breast cancer in mammograms from the PyTorch framework to TensorFlow/Keras. This change was driven by the better integration of TensorFlow with TPUs, since both products belong to Google. The CNN model selected is based on EfficientNet-B0 [2] and has three components: a patch classifier, which processes small regions of interest potentially indicating cancer; a one-view image classifier; and a two-view classifier that integrates craniocaudal and mediolateral-oblique views. The components are illustrated in Figure 1.

The experiments utilized the CBIS-DDSM dataset (Curated Breast Imaging Subset of Digital Database for Screening Mammography) [3], which is a public dataset for mammographic research. The original training and test set divisions were used. The images were initially resized to 1152x896 pixels, as in Petrini's work, and additional versions with doubled resolution (2304x1792 pixels) were generated to evaluate the impact of preserving higher-resolution details.

This study compared the performance of the two-view classifier proposed by Petrini et al., alongside its one-view and patch classifiers, on CBIS-DDSM. The experiments leveraged the specialized architecture and memory capacity of TPUs over GPUs by using higher-resolution mammograms.

Figure 1: Model Architecture

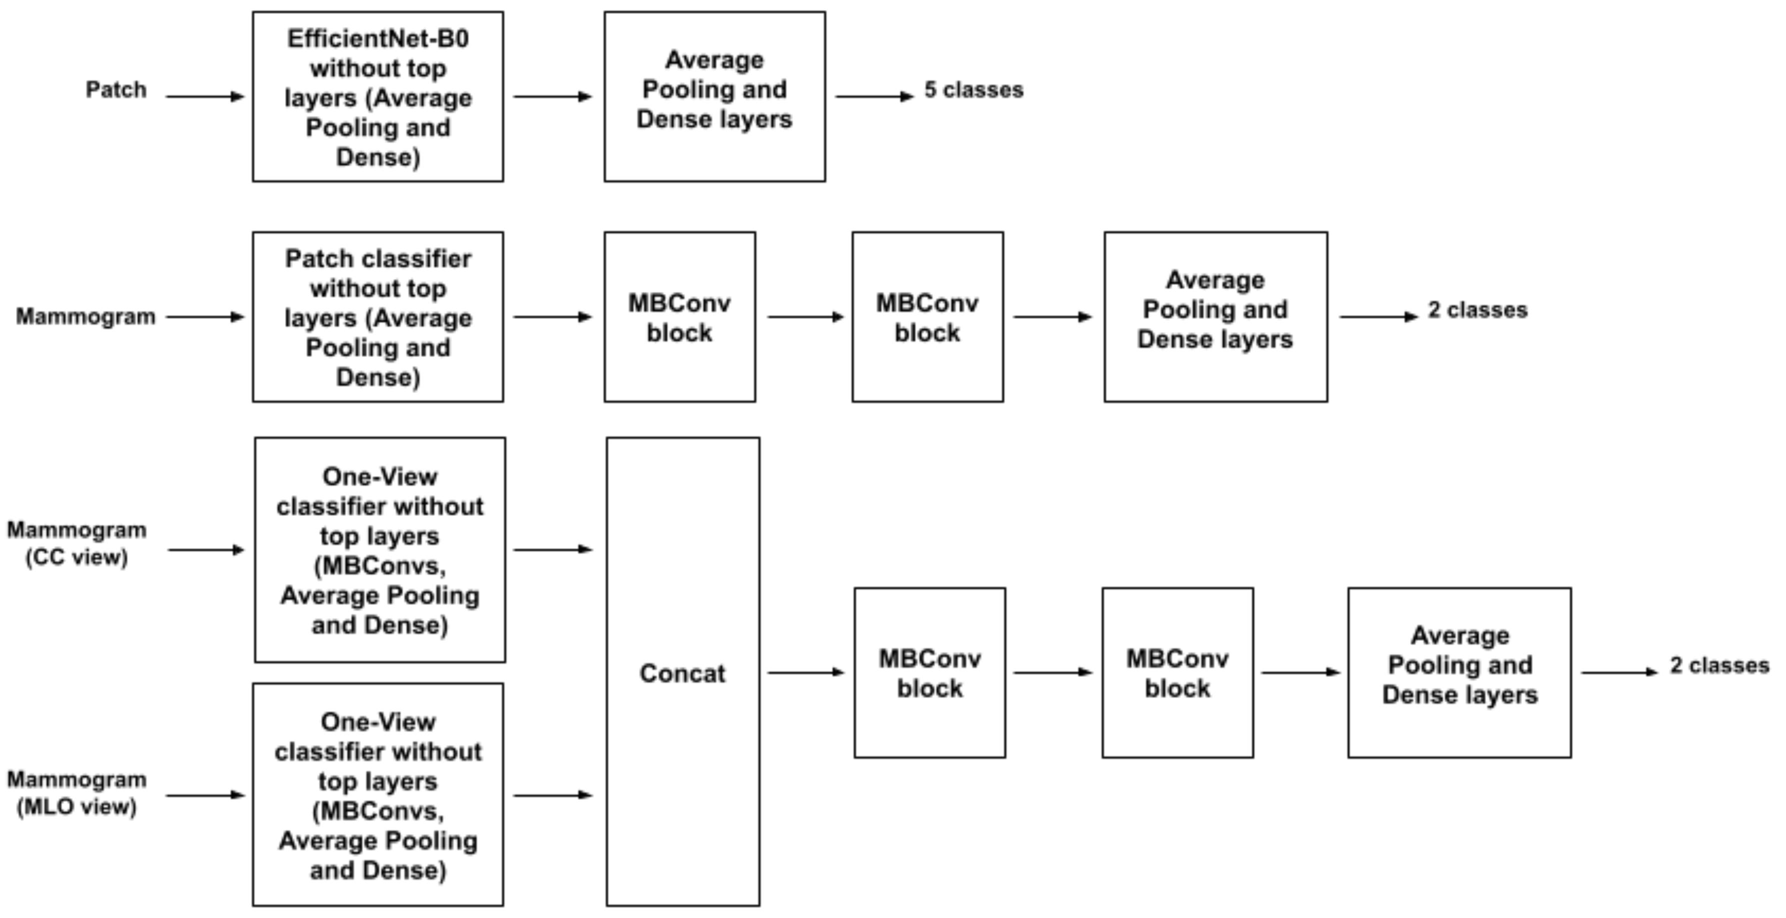

Source: Author

RESULTS & DISCUSSION

Tables 1 and 2 show the results of each classifier regarding their specific classification metric and the training time per epoch in seconds, respectively.

Table 1: Classifier performances

Classifier (Metric) / Device (Resolution)	Patch Classifier (Accuracy)	One View Classifier (AUC-ROC)	Two View Classifier (AUC-ROC)
Original (Petrini et al.)	75,54%	0,8033±0,0183	0,8418±0,0258
GPU (1152x896)	74,04%	0,8143±0,0179	0,8498±0,0227
TPU (1152x896)	76,37%	0,8003±0,0184	0,8327±0,0264
TPU (2304x1792)	79,52%	0,8154±0,0178	0,8466±0,0264

Table 2: Training time per epoch in seconds

Device (Resolution)	Patch Classifier	One View Classifier	Two View Classifier
GPU (1152x896)	363	376	73
TPU (1152x896)	19	22	17
TPU (2304x1792)	50	71	28

Using higher resolution, there was a significant improvement only in the accuracy of the patch classifier. Higher resolution did not lead to an increase in the performance metrics of the other two classifiers (one-view and two-view). A likely cause for this result is that the CBIS-DDSM dataset consists of low-quality, blurry scanned analog mammograms, which makes it useless to use high-resolution images.

Regarding speed, training on TPUs was up to 18 times faster than on GPUs, a significant increase in training speed that could potentially lead to better models in future work. However, no conclusive evidence showed that using images with higher resolutions and TPUs improved model performance. Metrics (accuracy and ROC-AUC) were similar at 1152x896 (GPU) and 2304x1792 (TPU).

CONCLUSION

Although the classification performance did not improve when increasing exam resolution, the use of TPUs is justifiable due to the increase in training speed, opening up possibilities to train with more data and using more complex architectures, which could lead to better classification results. Furthermore, working at high resolution should likely improve the performance of classifiers using high-quality, fully digital mammograms. This is the subject of a possible future work.

FUTURE WORK / REFERENCES

1 D. G. P. Petrini et al, "Breast Cancer Diagnosis in Two-View Mammography Using End-to-End Trained EfficientNet-Based Convolutional Network”, in IEEE Access, vol. 10, 2022.
2 M. Tan, and Q. Le. "Efficientnet: Rethinking model scaling for convolutional neural networks.”, in International conference on machine learning. PMLR, 2019.
3 R.S. Lee et al. “A curated mammography data set for use in computer-aided detection and diagnosis research”. In Sci Data 4, 2017.

Contact: sergio.a.g.fuentes@gmail.com
https://sciforum.net/event/ASEC2024