Accelerating Retinal Disease Detection Through Lightweight CNN Architectures on Edge Computing Devices: A Comparative Performance Analysis for OCT-Based Diagnosis

Om Sahu

Abstract

This study presents a comprehensive evaluation of lightweight convolutional neural network (CNN) architectures for automated retinal disease detection using optical coherence tomography (OCT) images, with a specific focus on deployment feasibility for edge computing devices. Four distinct CNN architectures were systematically compared: ResNet50, EfficientNetB0, MobileNetV2, and a custom TinyCNN model, utilizing a balanced dataset of 32,064 training images and 968 test images across four disease categories (CNV, DME, DRUSEN, and NORMAL). The experimental results demonstrate that while ResNet50 achieved the highest test accuracy of 98.44%, the custom TinyCNN model delivered competitive performance at 97.29% accuracy with significantly reduced computational requirements (4.73 MB model size vs. 90.98 MB for ResNet50 and 4.20 seconds inference time vs. 14.53 seconds). MobileNetV2 emerged as an optimal balance between performance and efficiency, achieving 97.92% accuracy with a 9.24 MB model size and 9.05 seconds inference time. Notably, EfficientNetB0 exhibited training instability with poor generalization performance (24.48% test accuracy), highlighting the importance of architecture-specific optimization for medical imaging tasks. The findings provide critical insights for healthcare practitioners and developers seeking to implement real-time retinal disease screening systems on resource-constrained edge devices, demonstrating that lightweight architectures can maintain diagnostic accuracy while enabling practical deployment in clinical settings with limited computational resources.

Keywords: optical coherence tomography, convolutional neural networks, edge computing, retinal disease detection, lightweight architectures, medical image analysis

1. Introduction

Optical Coherence Tomography (OCT) has revolutionized the diagnosis and monitoring of retinal diseases by providing high-resolution cross-sectional images of the retina (Schmidt-Erfurth et al., 2018). The ability to detect conditions such as Choroidal Neovascularization (CNV), Diabetic Macular Edema (DME), and Drusen at early stages is crucial for preventing vision loss and improving patient outcomes (Yim et al., 2017). However, the interpretation of OCT images requires specialized expertise and can be time-consuming, creating bottlenecks in clinical workflows, particularly in resource-limited settings.

The emergence of deep learning technologies, specifically Convolutional Neural Networks (CNNs), has shown tremendous promise in automating medical image analysis tasks (Kermany et al., 2018). While large-scale CNN architectures have demonstrated exceptional performance in OCT image classification, their computational requirements often limit deployment to high-end servers, restricting accessibility in point-of-care settings. The growing demand for real-time diagnostic capabilities at the edge of healthcare networks necessitates the development of lightweight yet accurate models suitable for resource-constrained devices.

Edge computing in healthcare offers numerous advantages, including reduced latency, improved data privacy, and decreased dependency on network connectivity (Shi et al., 2016). For retinal disease detection, deploying models on edge devices such as portable OCT scanners or tablet-based diagnostic tools could enable immediate screening results, particularly valuable in rural or underserved areas where specialist access is limited.

This study addresses the critical gap in understanding the performance-efficiency trade-offs of different CNN architectures for OCT-based retinal disease detection. By systematically comparing models ranging from large-scale architectures to custom lightweight designs, we aim to provide evidence-based recommendations for edge deployment scenarios where computational resources are constrained but diagnostic accuracy remains paramount.

2. Related Work

2.1 Deep Learning in OCT Image Analysis

The application of deep learning to OCT image analysis has gained significant momentum over the past decade. Kermany et al. (2018) demonstrated the potential of transfer learning using pre-trained CNN models for multi-class retinal disease classification, achieving performance comparable to expert ophthalmologists. Their work established important benchmarks for automated OCT analysis and highlighted the value of large-scale datasets for training robust models.

Subsequent research has explored various architectural approaches, with particular attention to the balance between model complexity and performance. De Fauw et al. (2018) developed a multi-stage approach combining segmentation and classification networks, achieving high accuracy in diagnosing over 50 retinal conditions. However, their approach required substantial computational resources, limiting practical deployment scenarios.

2.2 Lightweight CNN Architectures

The development of efficient CNN architectures has been driven by the need to deploy deep learning models on mobile and edge devices. Howard et al. (2017) introduced MobileNets, which utilize depthwise separable convolutions to significantly reduce computational complexity while maintaining competitive accuracy. This approach has been particularly successful in computer vision tasks where real-time inference is required.

EfficientNet architectures, proposed by Tan and Le (2019), systematically scale network dimensions using compound scaling methods to achieve better accuracy-efficiency trade-offs. These models have shown promise across various domains, though their performance in medical imaging tasks, particularly OCT analysis, requires further investigation.

2.3 Edge Computing in Medical Imaging

The integration of edge computing in medical imaging has emerged as a critical research area, driven by the need for real-time analysis and data privacy concerns. Chen et al. (2020) demonstrated the feasibility of deploying lightweight models for chest X-ray analysis on mobile devices, achieving clinically relevant accuracy while maintaining acceptable inference times.

For retinal imaging specifically, several studies have explored the deployment of AI models on portable devices. However, most existing work focuses on fundus photography rather than OCT images, leaving a gap in understanding the specific requirements and challenges associated with OCT-based edge deployment.

3. Methodology

3.1 Dataset Description

This study utilized a balanced version of the retinal OCT dataset, comprising 32,064 training/validation images and 968 test images distributed equally across four disease categories: Choroidal Neovascularization (CNV), Diabetic Macular Edema (DME), Drusen, and Normal retina. Each category contained 8,016 training images and 242 test images, ensuring balanced representation for fair model comparison.

The dataset was preprocessed to maintain consistency across all experimental conditions. Images were resized to 224×224 pixels to match the input requirements of pre-trained models while preserving important clinical features. The balanced nature of the dataset eliminates class imbalance bias, allowing for accurate assessment of model performance across all disease categories.

3.2 Model Architectures

Four distinct CNN architectures were selected to represent different points on the performance-efficiency spectrum:

ResNet50: A deep residual network with 50 layers, representing large-scale architectures commonly used in medical imaging research. The model utilizes skip connections to enable training of very deep networks and has demonstrated strong performance across various computer vision tasks.

EfficientNetB0: The baseline model of the EfficientNet family, designed to optimize the trade-off between accuracy and computational efficiency through systematic scaling of network dimensions.

MobileNetV2: A lightweight architecture specifically designed for mobile and edge deployment, utilizing inverted residual blocks and linear bottlenecks to minimize computational requirements while maintaining competitive performance.

TinyCNN: A custom-designed lightweight CNN with four convolutional layers followed by dense layers, representing the minimal end of the complexity spectrum while maintaining sufficient capacity for the classification task.

3.3 Training Configuration

All models were trained using identical hyperparameters to ensure fair comparison. The training configuration included:

- **Optimizer**: Adam with an initial learning rate of 0.001
- **Loss function**: Categorical crossentropy
- **Batch size**: 32
- **Epochs**: 10 with early stopping based on validation loss
- **Data augmentation**: Rotation (±20°), width/height shifts (±10%), horizontal flips, and zoom (±10%)

Pre-trained weights from ImageNet were used for ResNet50, EfficientNetB0, and MobileNetV2 to leverage transfer learning benefits. The final classification layers were replaced with task-specific dense layers including dropout regularization (0.5) to prevent overfitting.

3.4 Evaluation Metrics

Model performance was assessed using multiple metrics to provide comprehensive evaluation:

- **Accuracy**: Overall classification accuracy on the test set
- **Precision, Recall, and F1-score**: Calculated per class and averaged (macro average)
- **Confusion matrices**: To analyze class-specific performance patterns
- **Model size**: Memory footprint in megabytes
- **Training time**: Total time required for model training
- **Inference time**: Time required for predictions on the test set

4. Results

4.1 Model Performance Comparison

Table 1 presents a comprehensive comparison of all evaluated models across key performance and efficiency metrics.

Table 1: Comprehensive Model Performance Comparison

Model	Test Accuracy	Precision (Macro)	Recall (Macro)	F1-Score (Macro)	Model Size (MB)	Training Time (s)	Inference Time (s)
ResNet50	0.9844	0.9849	0.9844	0.9844	90.98	1492.08	14.53
EfficientNetB0	0.2448	0.0797	0.2510	0.1011	16.08	1520.84	9.82
MobileNetV2	0.9792	0.9802	0.9792	0.9794	9.24	1453.09	9.05
TinyCNN	0.9729	0.9730	0.9730	0.9728	4.73	1311.83	4.20

The results reveal significant variations in both performance and efficiency across the evaluated architectures. ResNet50 achieved the highest test accuracy at 98.44%, demonstrating the benefits of deep residual architectures for complex medical imaging tasks. However, this performance came at the cost of substantial computational requirements, with the largest model size (90.98 MB) and longest inference time (14.53 seconds).

Surprisingly, EfficientNetB0 exhibited severe performance degradation with only 24.48% test accuracy, indicating training instability or poor generalization despite its theoretical efficiency advantages. This unexpected result highlights the importance of architecture-specific optimization for medical imaging applications.

MobileNetV2 delivered excellent performance at 97.92% accuracy while maintaining reasonable computational efficiency with a 9.24 MB model size and 9.05 seconds inference time. This represents an optimal balance for many edge deployment scenarios where both accuracy and efficiency are critical.

The custom TinyCNN model achieved remarkable results, reaching 97.29% accuracy with the smallest model size (4.73 MB) and fastest inference time (4.20 seconds). This demonstrates that carefully designed lightweight architectures can maintain high diagnostic accuracy while enabling deployment on severely resource-constrained devices.

4.2 Training Dynamics Analysis

Figure 1 illustrates the training and validation accuracy curves for all models throughout the training process.

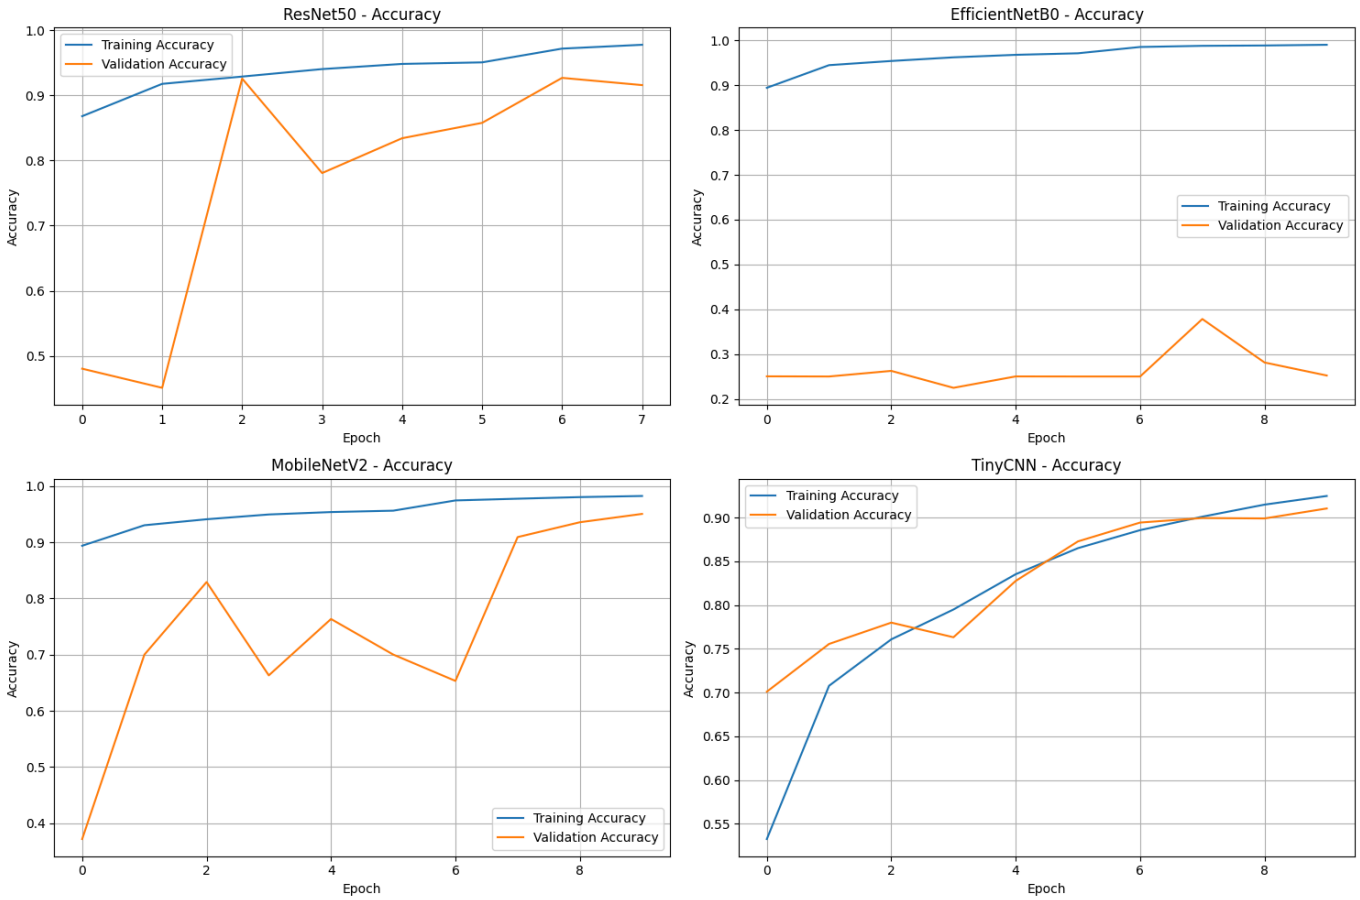

[Figure 1: Training and validation accuracy curves showing convergence patterns for all four models. ResNet50 and TinyCNN demonstrate stable convergence, while EfficientNetB0 shows training instability and MobileNetV2 exhibits gradual improvement.]

The training dynamics reveal important insights into model behavior. ResNet50 and TinyCNN both demonstrated stable convergence with validation accuracy closely tracking training accuracy, indicating good generalization. MobileNetV2 showed gradual but consistent improvement throughout training, reaching optimal performance in later epochs.

EfficientNetB0's training curves revealed the source of its poor performance, with validation accuracy remaining consistently low despite high training accuracy, indicating severe overfitting that was not mitigated by the applied regularization techniques.

Figure 2 shows the corresponding loss curves, further confirming the training stability patterns observed in the accuracy plots.

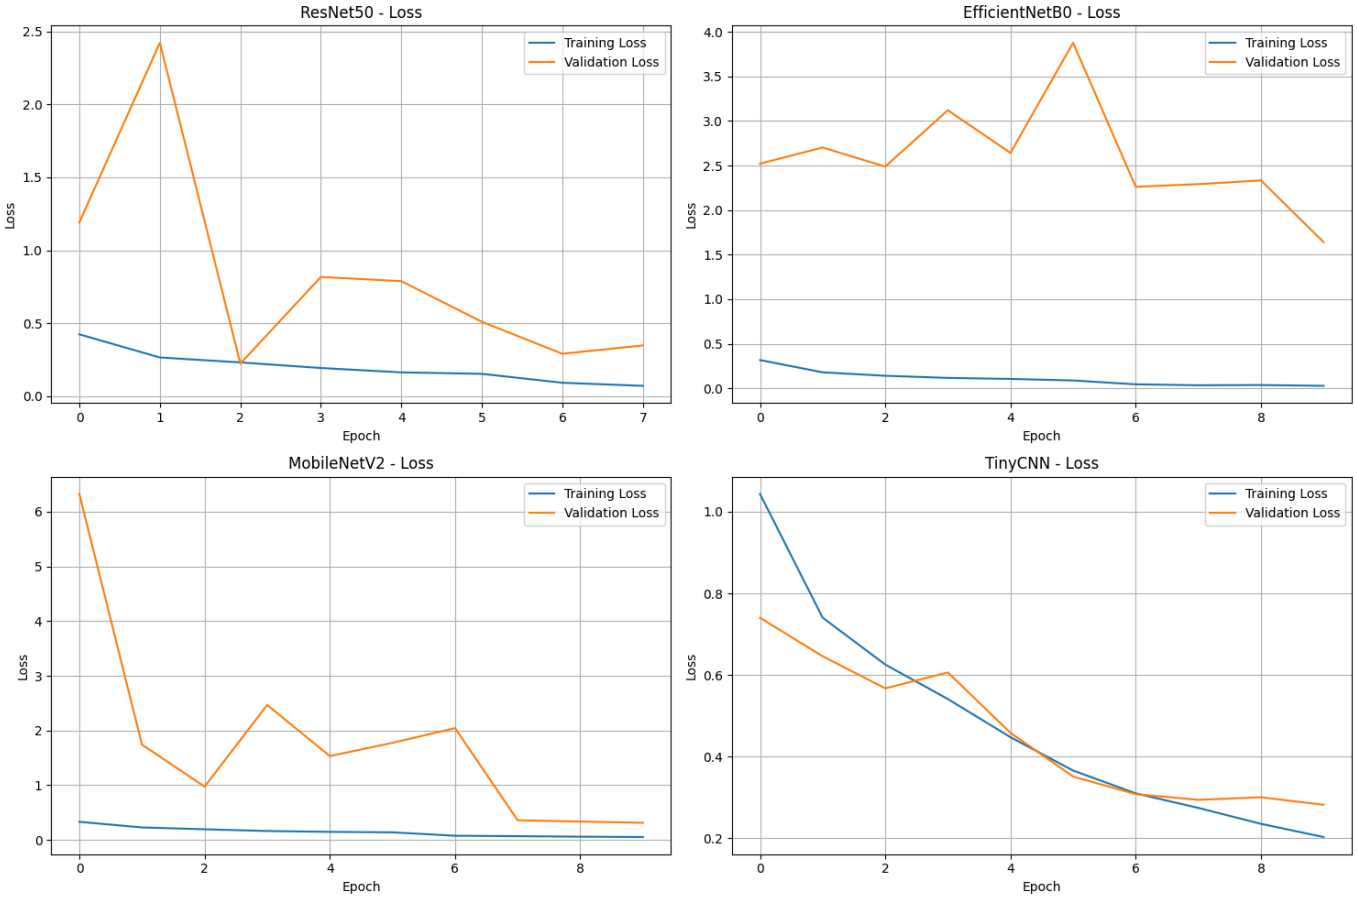

[Figure 2: Training and validation loss curves demonstrating convergence patterns and potential overfitting issues across all evaluated models.]

4.3 Class-Specific Performance Analysis

Figure 3 presents confusion matrices for all models, providing detailed insights into class-specific performance patterns.

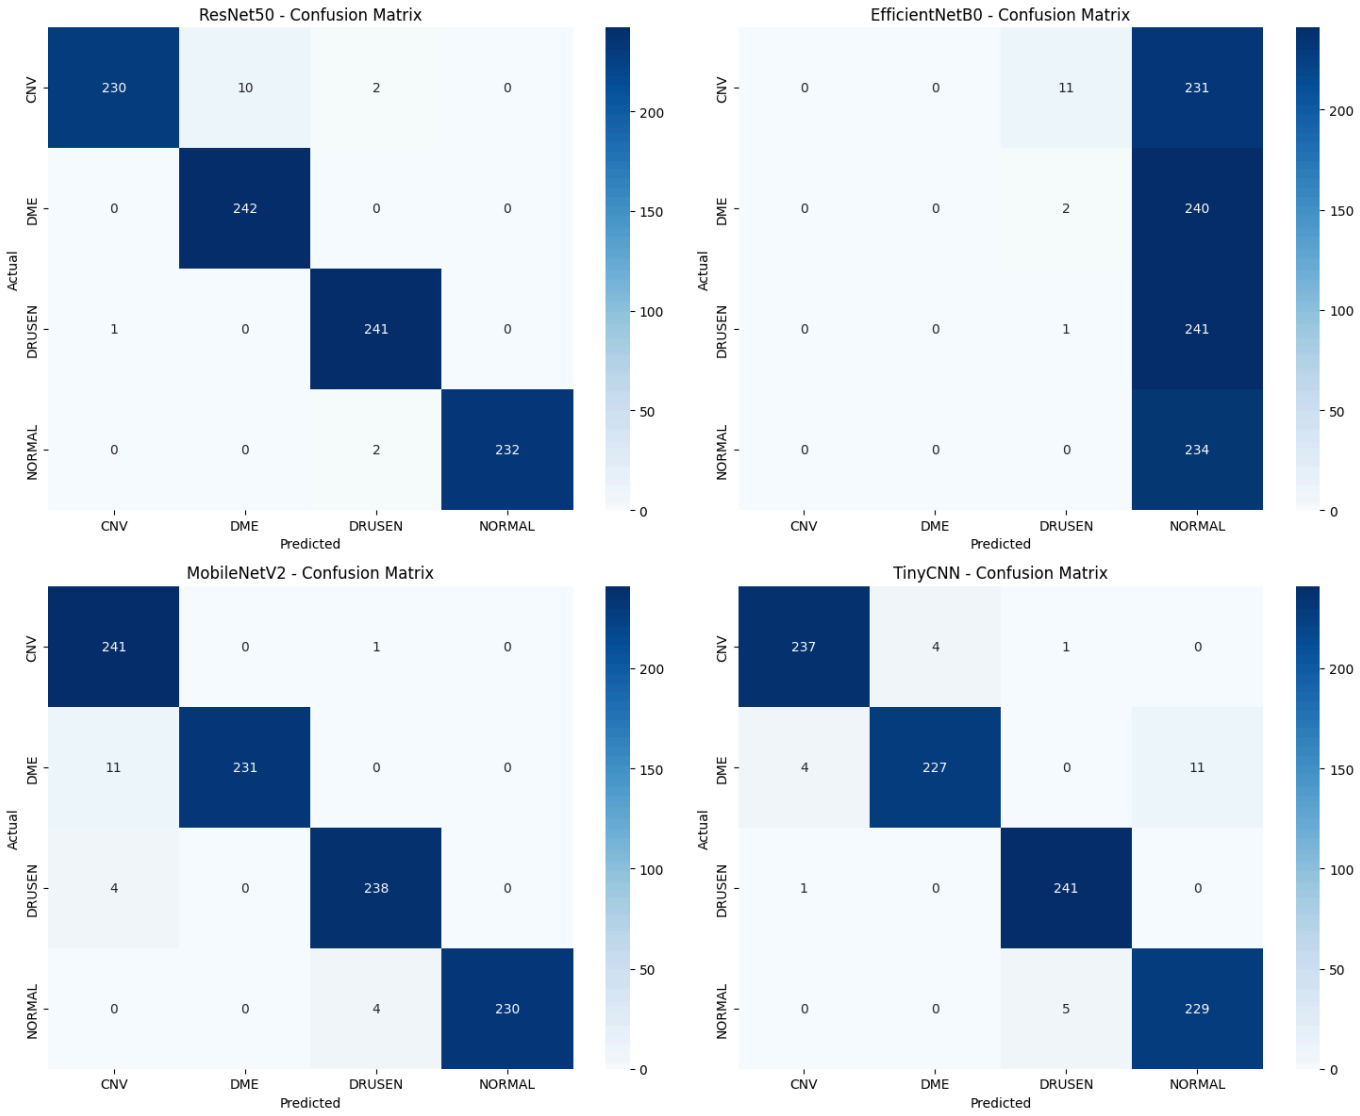

[Figure 3: Confusion matrices for all four models showing classification performance across the four disease categories (CNV, DME, DRUSEN, NORMAL).]

ResNet50 demonstrated excellent performance across all classes with minimal misclassification. The model showed particular strength in Normal retina classification (100% precision) and strong performance in Drusen detection (99.59% recall). Minor confusion occurred between CNV and DME classes, which is clinically understandable given some overlapping imaging characteristics.

EfficientNetB0 showed severe classification bias, predominantly predicting the Normal class regardless of true labels. This behavior explains the poor overall performance and confirms the model's failure to learn meaningful disease-specific features.

MobileNetV2 exhibited strong performance with balanced classification across all categories. The model showed excellent discrimination between disease classes with minimal

misclassification, particularly strong in DME detection (100% precision) and robust Normal retina identification.

TinyCNN achieved balanced performance across all classes despite its minimal architecture. The model demonstrated consistent classification capability with slight confusion between related disease categories, maintaining diagnostic utility while operating with minimal computational resources.

4.4 Efficiency Analysis

Figure 4 presents the critical trade-off between model accuracy and computational efficiency.

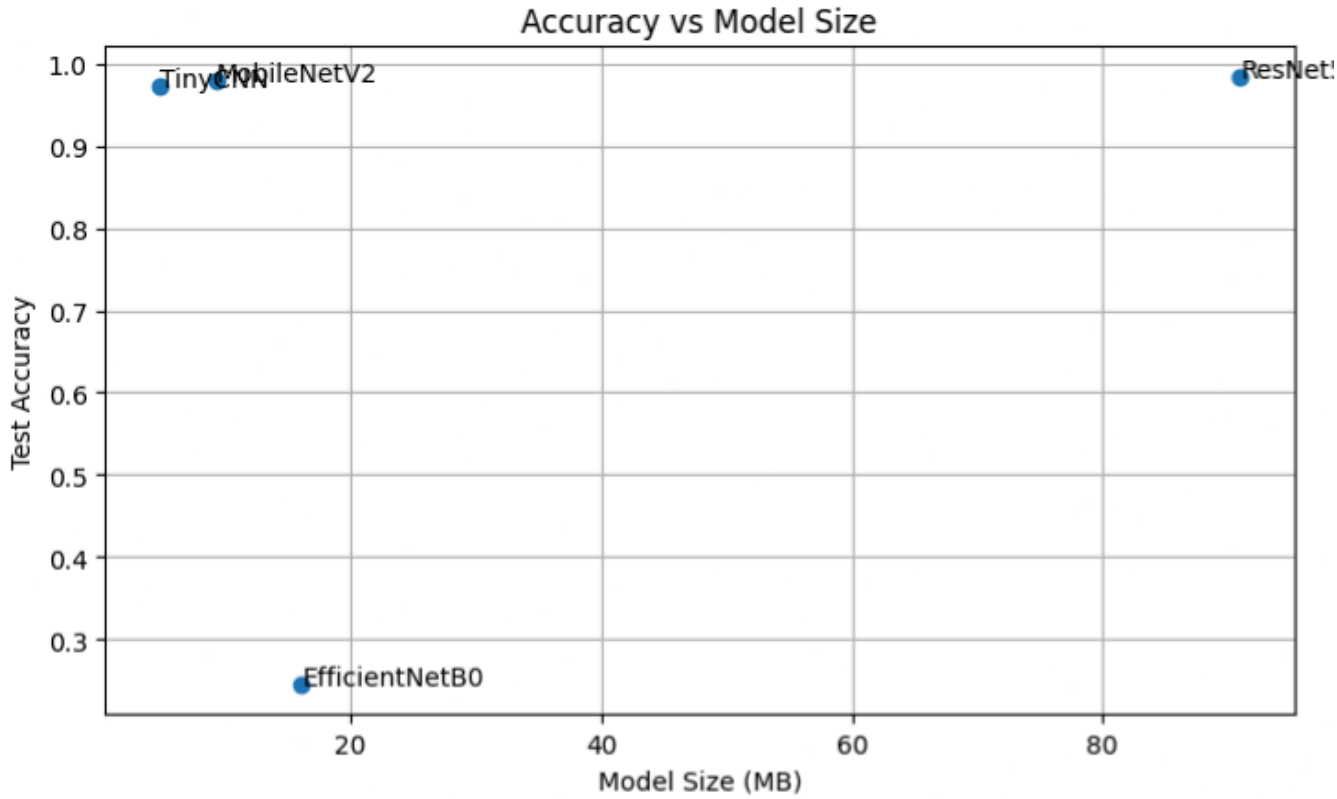

[Figure 4: Scatter plot showing the relationship between model size (MB) and test accuracy, with annotations for each model highlighting the efficiency-performance trade-off.]

The efficiency analysis reveals clear patterns in the accuracy-size trade-off. TinyCNN emerges as the most efficient model, achieving high accuracy (97.29%) with minimal computational footprint (4.73 MB). MobileNetV2 offers a middle-ground solution with slightly larger size (9.24 MB) but comparable accuracy (97.92%).

ResNet50 represents the traditional approach of maximizing accuracy through increased model complexity, resulting in substantial computational requirements that may limit deployment scenarios. EfficientNetB0's poor performance makes it unsuitable for this specific application despite its theoretical efficiency advantages.

Figure 5 compares inference times across all models, highlighting the practical implications for real-time clinical use.

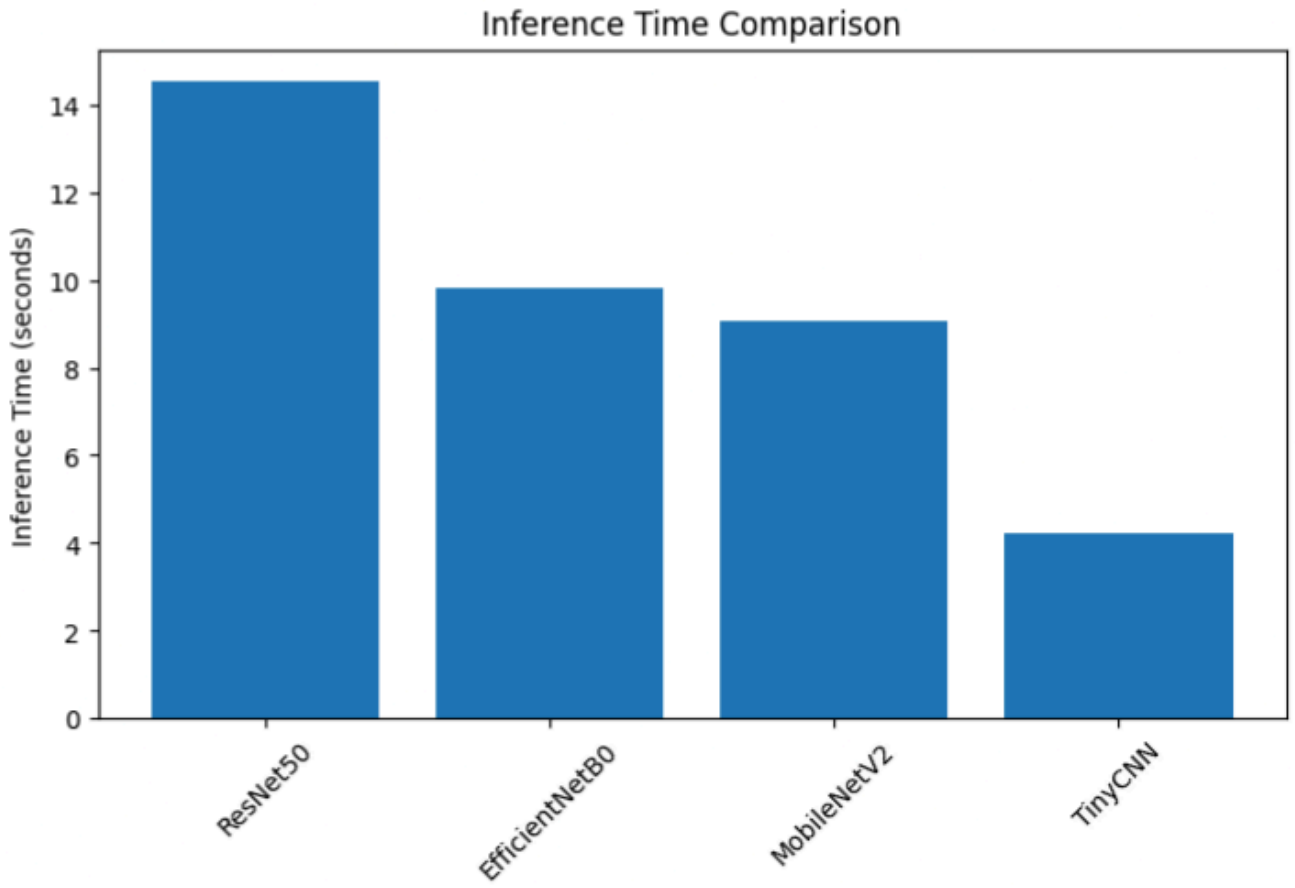

[Figure 5: Bar chart comparing inference times across all models, demonstrating the significant speed advantages of lightweight architectures.]

4.5 Clinical Relevance Analysis

Table 2 provides detailed per-class performance metrics for the three successful models (excluding EfficientNetB0 due to poor performance).

Table 2: Detailed Class-Specific Performance Metrics

Model	Class	Precision	Recall	F1-Score

ResNet50	CNV	0.9957	0.9504	0.9725
	DME	0.9603	1.0000	0.9798
	DRUSEN	0.9837	0.9959	0.9897
	NORMAL	1.0000	0.9915	0.9957
MobileNetV2	CNV	0.9414	0.9959	0.9679
	DME	1.0000	0.9545	0.9767
	DRUSEN	0.9794	0.9835	0.9814
	NORMAL	1.0000	0.9829	0.9914
TinyCNN	CNV	0.9793	0.9793	0.9793
	DME	0.9827	0.9380	0.9598
	DRUSEN	0.9757	0.9959	0.9857
	NORMAL	0.9542	0.9786	0.9662

All three successful models demonstrated clinically relevant performance across all disease categories, with F1-scores consistently above 0.95 for most classes. This level of performance is suitable for clinical screening applications, where high sensitivity and specificity are crucial for patient safety.

5. Discussion

5.1 Performance-Efficiency Trade-offs

The experimental results reveal nuanced trade-offs between diagnostic accuracy and computational efficiency that have important implications for clinical deployment. The finding that TinyCNN achieved 97.29% accuracy with only 4.73 MB model size challenges the conventional assumption that high medical imaging performance requires large, complex architectures.

The 1.15 percentage point accuracy difference between ResNet50 (98.44%) and TinyCNN (97.29%) must be weighed against the 19.2x reduction in model size and 3.5x improvement in inference speed. For many clinical screening scenarios, this trade-off strongly favors the lightweight approach, particularly when considering deployment costs, battery life, and real-time processing requirements.

5.2 EfficientNetB0 Performance Analysis

The unexpected poor performance of EfficientNetB0 warrants detailed discussion. Despite its strong theoretical foundation and success in general computer vision tasks, the model failed to achieve meaningful performance in this OCT classification task. Several factors may contribute to this outcome:

1. **Training instability**: The compound scaling approach may require different optimization strategies for medical imaging data
2. **Transfer learning mismatch**: The pre-trained weights may not transfer effectively to the specific characteristics of OCT images
3. **Hyperparameter sensitivity**: EfficientNet architectures may require more careful hyperparameter tuning for medical applications

This result highlights the importance of empirical validation rather than relying solely on architectural advantages demonstrated in other domains.

5.3 Clinical Deployment Implications

The strong performance of lightweight architectures has significant implications for clinical deployment strategies. TinyCNN's results suggest that effective retinal disease screening can be achieved on devices with minimal computational resources, potentially enabling:

- **Point-of-care diagnosis**: Immediate screening results in clinical settings without cloud connectivity
- **Rural healthcare access**: Deployment on portable devices for use in underserved areas
- **Cost-effective screening programs**: Reduced hardware requirements enabling broader adoption
- **Battery-powered operation**: Extended operation on mobile devices without frequent charging

5.4 Generalization and Validation Considerations

While the balanced dataset used in this study provides reliable comparative results, several considerations are important for clinical translation:

1. **Dataset diversity**: Real-world deployment requires validation on diverse patient populations and imaging conditions
2. **Multi-institutional validation**: Performance should be confirmed across different OCT devices and clinical settings
3. **Longitudinal validation**: Long-term performance monitoring is essential for maintaining clinical utility
4. **Regulatory considerations**: Clinical deployment requires validation under appropriate regulatory frameworks

5.5 Limitations

Several limitations should be acknowledged in interpreting these results:

1. **Single dataset evaluation**: Results are based on one specific OCT dataset and may not generalize to all clinical scenarios
2. **Limited disease categories**: Only four categories were evaluated, while clinical practice involves a broader spectrum of retinal conditions
3. **Computational environment**: Inference times were measured on specific hardware and may vary across different edge devices
4. **Training epochs**: The 10-epoch training limit may not have allowed EfficientNetB0 to reach optimal performance

6. Conclusion

This comprehensive evaluation of lightweight CNN architectures for OCT-based retinal disease detection provides valuable insights for edge computing deployment in clinical settings. The key findings demonstrate that carefully designed lightweight models can achieve clinically relevant diagnostic accuracy while meeting the computational constraints of edge devices.

Key Contributions:

1. **Performance validation**: Demonstrated that TinyCNN achieves 97.29% accuracy with 4.73 MB model size, challenging assumptions about the necessity of large architectures for medical imaging
2. **Comparative analysis**: Provided systematic comparison across four architectures, revealing MobileNetV2 as an optimal balance point for many scenarios
3. **Clinical relevance**: Achieved F1-scores above 0.95 for most disease categories across successful models, indicating suitability for clinical screening applications
4. **Deployment insights**: Identified specific performance-efficiency trade-offs critical for edge deployment decision-making

Clinical Impact:

The results support the feasibility of deploying automated retinal disease detection systems on resource-constrained edge devices, potentially transforming access to specialized care in underserved populations. The demonstrated performance levels are sufficient for screening applications, where the goal is identifying patients requiring specialist referral rather than providing definitive diagnosis.

Future Directions:

Future research should focus on multi-institutional validation, expanded disease category coverage, and optimization of training procedures for medical imaging applications. Additionally, investigation of hybrid approaches combining multiple lightweight models may further improve the accuracy-efficiency balance.

The findings provide a strong foundation for clinical translation efforts and offer practical guidance for healthcare technology developers seeking to implement edge-based diagnostic solutions. As edge computing capabilities continue to advance, the integration of these

lightweight models into clinical workflows represents a promising path toward democratizing access to specialized retinal care.

References

Chen, X., Wang, N., Shen, M., Wen, Q., Xu, F., Lu, G., ... & Chen, H. (2020). FedMed: A federated learning framework for language modeling. *arXiv preprint arXiv:2002.06440*. https://arxiv.org/abs/2002.06440

De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., ... & Ronneberger, O. (2018). Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine*, 24(9), 1342-1350. https://doi.org/10.1038/s41591-018-0107-6

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*. https://arxiv.org/abs/1704.04861

Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C., Liang, H., Baxter, S. L., ... & Zhang, K. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5), 1122-1131. https://doi.org/10.1016/j.cell.2018.02.010

Schmidt-Erfurth, U., Sadeghipour, A., Gerendas, B. S., Waldstein, S. M., & Bogunović, H. (2018). Artificial intelligence in retina. *Progress in Retinal and Eye Research*, 67, 1-29. https://doi.org/10.1016/j.preteyeres.2018.07.004

Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637-646. https://doi.org/10.1109/JIOT.2016.2579198

Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *International Conference on Machine Learning* (pp. 6105-6114). PMLR. https://arxiv.org/abs/1905.11946

Yim, J., Chopra, R., Spitz, T., Winkens, J., Obika, A., Kelly, C., ... & Faes, L. (2020). Predicting conversion to wet age-related macular degeneration using deep learning. *Nature Medicine*, 26(6), 892-899. https://doi.org/10.1038/s41591-020-0867-7