Indian Journal of Computer Science and Technology
Volume 5, Issue1 (January-April 2026), PP: 307-312.
www.indjcst.com **ISSN No: 2583-5300**

Dermanex: Next Generation Ai Skin Disease Detection Using Deep Learning & Machine Learning

Sudheer Kumar Kolahala[1], Venkat Sai Keesara[2], Tejaswini Avula[3], Dr. Srinivas Jagirdar[4]
[1,2,3] UG Scholar, Department of Computer Engineering, Matrusri Engineering College, Hyderabad, Telangana, India.
[4]Associate Professor & HOD, Department of Information Technology, Matrusri Engineering College, Hyderabad, Telangana, India.

To Cite this Article: *Sudheer Kumar Kolahala[1], Venkat Sai Keesara[2], Tejaswini Avula[3], Dr. Srinivas Jagirdar[4], "Dermanex: Next Generation Ai Skin Disease Detection Using Deep Learning & Machine Learning", Indian Journal of Computer Science and Technology, Volume 05, Issue 01 (January-April 2026), PP: 307-312.*

Abstract: *Skin diseases are among the most prevalent health concerns globally, affecting nearly 900 million people at any given time. In India, the burden is further increased due to climatic conditions, high population density, and limited access to dermatological specialists, especially in rural areas of the country. This study presents Derma Nex, a multi-model deep learning framework designed for skin cancer detection and common dermatological disease classification.The system integrates four specialized models based on Efficient Net architectures to perform binary cancer screening, multi-class lesion classification, melanoma-versus-nevus differentiation, and common skin disease recognition. A three-model ensemble consensus mechanism combines multiple probabilistic predictions to improve the robustness and reliability. The proposed approach achieves strong performance, with high cancer detection sensitivity, 96% AUC-ROC in screening tasks, and up to 85.8% accuracy in the classification of common diseases. To enhance interpretability, the system incorporates a multi-layer Grad-CAM for a visual explanation of the predictions. The models were trained on a diverse dataset of over 50,000 images collected from multiple sources, ensuring improved generalization. Furthermore, the system was deployed as a full-stack web application with risk-level analysis and user-friendly interaction, making it suitable for real-world applications. Overall, DermaNex provides an efficient, scalable, and interpretable solution for AI-assisted dermatological diagnoses.*

Key Words: *Skin Cancer Detection, Deep Learning, Efficient Net, Attention Mechanisms, Grad-CAM, Ensemble Consensus.*

I. INTRODUCTION

Skin disorders represent a significant, yet often underestimated, healthcare challenge, affecting hundreds of millions of individuals worldwide. In countries with high population densities and limited access to dermatology specialists, early diagnosis remains difficult, leading to delayed treatment and an increased risk of complications. Common conditions such as eczema, fungal infections, and acne frequently go untreated or are misdiagnosed, whereas more critical diseases such as melanoma require timely detection for effective intervention. Recent progress in deep learning has opened new possibilities for the automated analysis of medical images. Convolutional neural networks have demonstrated strong performance in recognizing visual patterns within dermoscopic images, and in some cases, they have approached expert-level accuracy under controlled conditions. Publicly available datasets, such as HAM10000 and ISIC, have played a major role in enabling this progress by providing large-scale labeled data for training and evaluation.

Despite these advancements, most existing systems are designed with a narrow focus, typically addressing only skin-cancer detection. Such approaches do not reflect real-world clinical scenarios in which a broad range of dermatological conditions must be considered. Additionally, single-model systems often struggle to distinguish between visually similar categories, particularly melanoma and benign lesions, owing to overlapping features. To address these limitations, this study proposes DermaNex, a multi-model framework that distributes diagnostic tasks across specialized architectures rather than relying on a single generalized model. The system integrates multiple prediction pathways, each optimized for a specific objective, and combines their outputs using a consensus-based strategy to improve the overall reliability.

Another key aspect of the proposed approach is its interpretability. Instead of producing only classification results, the system provides visual explanations that indicate which regions of an image influence the decision. This improves transparency and makes the system more suitable for practical applications. The primary contributions of this study include the design of a multi-model diagnostic pipeline, an ensemble-based decision mechanism that integrates multiple probabilistic perspectives, and the incorporation of explainable AI techniques within a deployable web application. Together, these components form a comprehensive solution aimed at supporting accessible and scalable dermatological analyses.

II. MATERIALS AND METHODS

2.1 System Overview

The proposed DermaNex system is a multimodal deep learning framework designed for comprehensive skin disease analysis. The architecture consists of four specialized models organized into two pathways. The first pathway processes

dermoscopic images using three models for cancer detection, whereas the second pathway analyzes clinical images using a separate model for common skin disease classification. Models 1–3 work together in a consensus-based approach to perform cancer screening, multi-class classification, and melanoma-versus-nevus differentiation. Model 4 independently classifies the common dermatological conditions. This dual-path design enables the system to handle both critical and routine skin diseases within a single workflow.

Table 1 summarizes the architecture of the proposed models.

Model	Backbone	Task	Input Size	Output
Model 1 – Cancer Screener	EfficientNet-B4	Binary	288×288	P(cancer)
Model 2 – Disease Classifier	EfficientNet-B4 + SE + SA	7-class	288×288	7 probs
Model 3 – MEL vs NV	EfficientNet-B5 + BAM	Binary	384×384	P(melanoma)
Model 4 – Common Diseases	EfficientNet-B1	7-class	224×224	7 probs

Table 1: DermaNex Model Architecture Summary

2.2 Model Architecture

The DermaNex system consists of four EfficientNet-based models, each designed for a specific diagnostic task.

Model 1 performs binary cancer detection using multi-scale feature extraction and integrates clinical metadata, such as age, sex, and lesion location, to improve prediction accuracy.

Model 2 extends EfficientNet-B4 by incorporating attention mechanisms, including squeeze-and-excitation and spatial attention, to enhance feature learning and improve the seven-class classification performance.

Model 3 focuses on melanoma versus nevus differentiation using EfficientNet-B5 with a Border Attention Module. This module emphasizes irregular lesion boundaries, which are critical indicators of melanoma, thereby improving the classification sensitivity.

Model 4 uses EfficientNet-B1 to classify common skin conditions, such as acne, eczema, fungal infections, psoriasis, and warts. It is optimized for clinical images and balances the performance with computational efficiency.

2.3 Ensemble Consensus Mechanism

To improve the prediction reliability, the system employs an ensemble consensus mechanism that combines the outputs of Models 1, 2, and 3. The ensemble is based on three probabilistic perspectives. The first view adjusts the class probabilities using the cancer likelihood predicted by Model 1. The second view uses the direct multiclass output from Model 2. The third view refines the predictions using melanoma-specific probabilities from Model 3.These outputs were averaged to generate a final prediction, and temperature scaling was applied to improve the confidence calibration. This approach enhances the overall robustness and reduces individual model errors.

2.4 Dataset Description

DermaNex was trained on a combined dataset of over 50,000 unique images assembled from five major dermatological image sources. Each model utilizes a specific combination of datasets that are optimized for its classification task.

Dataset	Images	Type	Primary Use
HAM10000 [11]	10,015	Dermoscopic	Models 1, 2, 3
ISIC 2019	~25,000	Dermoscopic	Model 3
ISIC 2020 [12]	~33,000	Dermoscopic	Models 1, 2, 3
PAD-UFES-20	2,298	Smartphone	Models 1, 2, 3
DermNet	~17,000	Clinical Photos	Model 4

Table 2: Dataset Sources

2.4.1 Model 1 Dataset

Model 1 combined the HAM10000, ISIC 2020, and PAD-UFES-20 datasets for binary cancer classification. The ISIC 2020 melanoma cases included 8,000 capped benign samples. Total: ~20,000 images after benign capping. Patient-aware stratified splitting ensures that no data leakage occurs.

2.4.2 Model 2 Dataset

Model 2 used HAM10000 as the primary source with supplemental melanoma cases from ISIC 2020 and cross-mapped PAD-UFES-20 entries. Total: ~12,000 images.

Class	Code	Images	Type
Melanoma	MEL	1,113	Cancer
Basal Cell Carcinoma	BCC	514	Cancer
Actinic Keratosis	AKIEC	327	Cancer
Melanocytic Nevus	NV	6,705	Benign
Benign Keratosis	BKL	1,099	Benign
Dermatofibroma	DF	115	Benign
Vascular Lesion	VASC	142	Benign
Total		10,015	

Table 3: HAM10000 Class Distribution

2.4.3 Model 3 Dataset

Model 3 used the largest combined dataset: HAM10000 (MEL: 1,113, NV: 6,705), ISIC 2019 (MEL: ~4,522, NV: ~12,875), ISIC 2020 (~584 melanoma), and PAD-UFES-20 (MEL: 52, NEV: 244). Total after deduplication: ~24,000 images.

2.4.4 Model 4 Dataset

Assembled from DermNet and supplementary Kaggle datasets. Classes with fewer than 1,500 training images received 2–6× the oversampling weights.

Class	Train	Test	Sampling Weight
Psoriasis	4,018	723	1.0
Eczema	3,234	582	1.0
Warts	2,499	450	1.0
Fungal Infection	2,402	433	1.0
Nail Fungus	939	169	2.0
Acne	832	150	2.5
Scabies	389	70	6.0
Total	14,313	2,577	—

Table 4: Common Skin Disease Dataset Distribution

2.5 Training Strategy

All models were trained using transfer learning with pre-trained EfficientNet weights. Training was performed in two phases, starting with training the classification layers while freezing the backbone, followed by full fine-tuning with a lower learning rate.

The AdamW optimizer was used along with cosine annealing learning rate scheduling. Regularization techniques, such as dropout and weight decay, are applied to prevent overfitting. For the common disease model, class imbalances were handled using weighted sampling.Test-time augmentation is applied during inference to improve the prediction stability by averaging the outputs from multiple augmented versions of the input image.

2.6 Explainability

The system integrates Grad-CAM to provide visual explanations of the model predictions. A multilayer fusion approach is used to combine attention maps from different network layers, improving the localization of important regions.

These visualizations highlight critical features, such as lesion borders and patterns, enhancing interpretability. Additionally, an AI-based assistant provides contextual explanations while maintaining safety constraints, thereby ensuring responsible use in healthcare applications.

III. RESULTS

3.1 Model 1 Results — Cancer Screener

Binary cancer screening achieved 82.48 an overall accuracy of 95.00% malignant recall at a threshold of 0.4599. The deliberately low threshold trades specificity for higher cancer sensitivity, consistent with clinical screening priorities, where missed cancers carry substantially higher costs than false positives.

Metric	Value	Target	Status
AUC-ROC	96.15%	≥ 95%	PASS
Cancer Recall	95.00%	≥ 95%	PASS
Negative Predictive Value (NPV)	98.51%	≥ 97%	PASS
Overall Accuracy	82.48%	—	—
Missed Cancer Cases	33	—	—

Table 5: Model 1 — Cancer Screener Performance

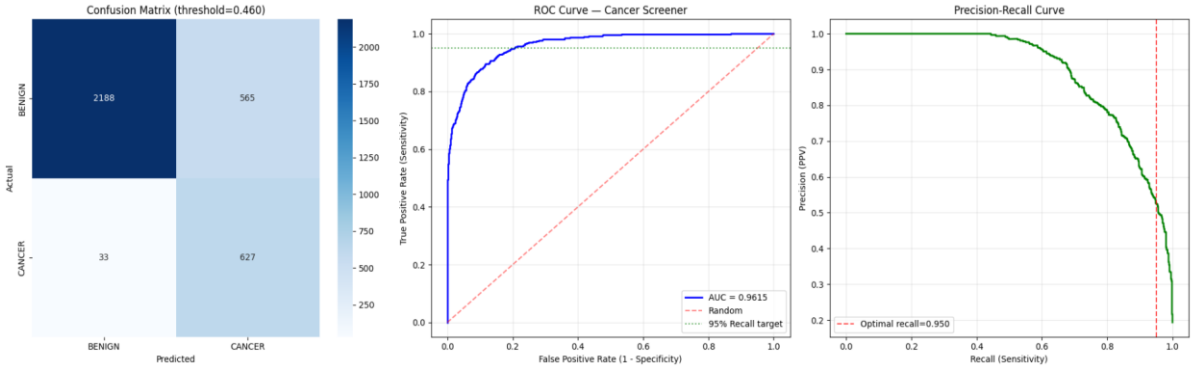

3.2 Model 2 Results — Skin Cancer Classifier

The seven-class classification achieved an overall accuracy of 77.25% with a Macro F1 of 0.7022. The cancer class safety check confirmed MEL and BCC recall at 87.23% (12.8% miss rate each). The primary challenge was NV with 71.86% recall owing to its dominance in the dataset (6,705 of 10,015 images).

Class	Recall	F1-Score	Status
MEL (Melanoma)	87.23%	0.67	PASS
BCC (Basal Cell Carcinoma)	87.23%	0.73	PASS
AKIEC (Actinic Keratosis)	72.73%	0.64	Close to Target
NV (Melanocytic Nevus)	71.86%	0.83	Close to Target
BKL (Benign Keratosis)	81.19%	0.66	PASS
DF (Dermatofibroma)	100.00%	0.45	PASS (Recall)
VASC (Vascular Lesion)	94.44%	0.92	PASS

Table 6: Model 2 — Per-Class Performance

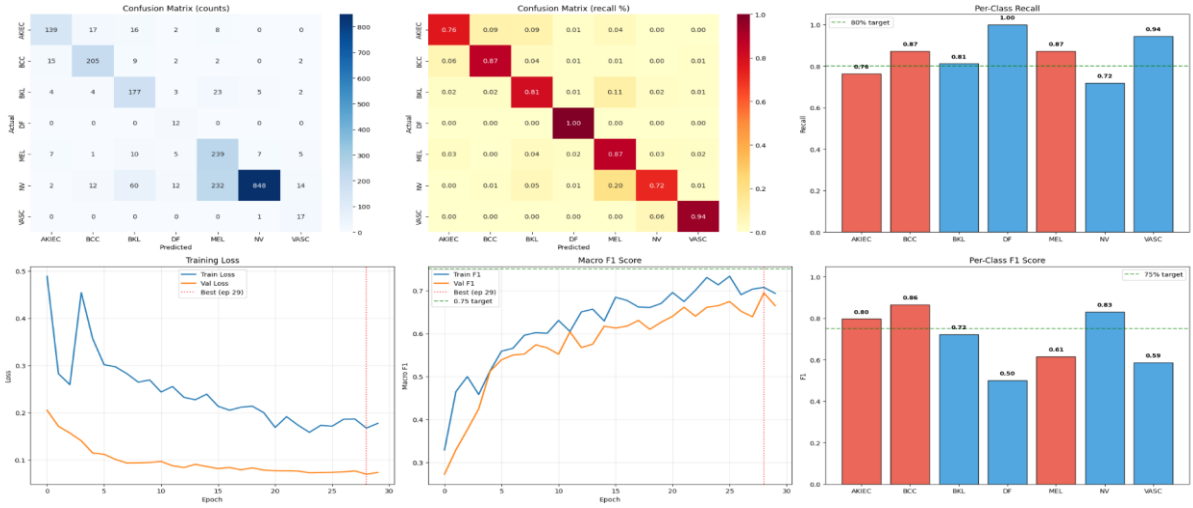

3.3 Model 3 Results — MEL vs NV Specialist

The specialized melanoma-nevus classifier achieved an accuracy of 84.72% and an AUC-ROC of 96.82 %. The Border Attention Module contributes to strong melanoma recall by focusing on irregular border patterns, the most discriminating visual feature between MEL and NV.

Metric	Value	Target	Status
AUC-ROC	96.82%	≥ 95%	PASS
MEL Recall	95.07%	≥ 95%	PASS
Negative Predictive Value (NPV)	97.76%	≥ 97%	PASS
Overall Accuracy	84.72%	—	—
Missed MEL Cases	41	—	—

Table 7: Model 3 — MEL vs NV Specialist Performance

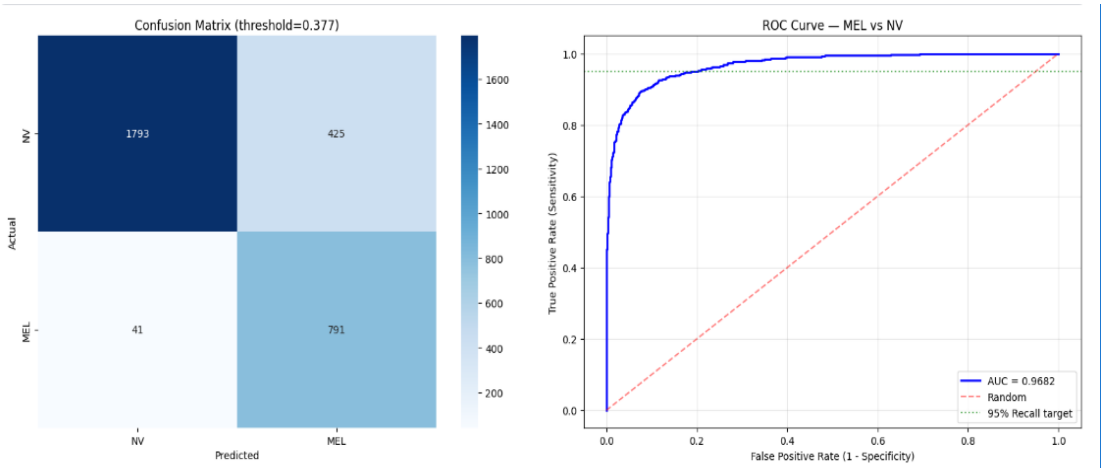

3.4 Model 4 Results — Common Skin Disease Classifier

Model 4 achieved an overall accuracy of 85.8% with a Macro F1 of 0.831 and Weighted F1 of 0.857 using test-time augmentation (TTA). Six augmentation variants (original, horizontal flip, vertical flip, rotation ±15°, and CLAHE) were applied with averaged softmax probabilities. TTA improved the accuracy from 84.9% to 85.8% (+0.9%).

Class	Recall	F1-Score	Status
Fungal Infection	93.5%	0.90	Target Met
Psoriasis	91.7%	0.91	Target Met
Nail Fungus	84.6%	0.82	Close to Target
Warts	84.2%	0.84	Close to Target
Acne	80.0%	0.78	Close to Target
Eczema	79.0%	0.80	Close to Target
Scabies	60.0%	0.59	Below Target

Table 8: Model 4 — Per-Class Performance with TTA

IV. DISCUSSION

The DermaNex framework demonstrates the effectiveness of integrating multiple specialized deep learning models within a unified pipeline for comprehensive dermatological analysis. Unlike conventional single-model approaches, the proposed architecture can address both cancerous and non-cancerous skin conditions, thereby expanding its clinical applicability. By combining multiple diagnostic pathways, the system provides a structured and scalable solution that can handle diverse dermatological cases in a single workflow. This multi-model design not only improves coverage but also ensures that different diagnostic tasks are handled by models that are optimized for specific objectives.

A key strength of the system lies in its ensemble-based decision mechanism, which significantly enhances the prediction reliability. The integration of binary cancer screening, multi-class classification, and melanoma-specific analysis allowed the system to capture complementary information from different models. This layered approach reduces the likelihood of misclassification, particularly in challenging scenarios, such as distinguishing melanoma from melanocytic nevi. In addition, incorporating attention mechanisms improves feature representation by enabling models to focus on diagnostically relevant

regions, including lesion borders, textures, and color variations. The proposed Border Attention Module further refines this process by emphasizing irregular boundary patterns that are critical indicators of malignant lesions.

Interpretability and usability are important aspects of the DermaNex system. The integration of multi-layer fusion Grad-CAM provides more precise and informative visual explanations than traditional single-layer methods. By combining both structural and semantic feature maps, the system can highlight key regions that influence predictions, thereby improving transparency and user trust. Furthermore, the deployment of the system as a full-stack web application enhances its practicality. Features such as confidence-based risk assessment, analysis history tracking, and user-friendly interaction make the system accessible to a wider audience, particularly in regions where access to dermatological expertise is limited.

Despite its promising performance, the system has certain limitations that must be addressed. Reliance on publicly available datasets may introduce bias, particularly in terms of skin tone diversity and imaging conditions. Variations in lighting, resolution, and acquisition devices can affect the generalization of models in real-world environments. Additionally, although the system provides valuable decision support, it is not intended to replace professional medical diagnosis. Future work will focus on incorporating more diverse and region-specific datasets, improving robustness against real-world variations, and exploring advanced architectures, such as Vision Transformers, for enhanced feature learning. Conducting clinical validation studies with dermatologists is essential to evaluate the reliability and effectiveness of the system in practical healthcare settings. Overall, the DermaNex framework presents a robust, interpretable, and scalable solution that addresses the key challenges in AI-assisted dermatology.

V. CONCLUSION

This study presents DermaNex, a comprehensive multimodal deep learning framework designed for both skin cancer detection and common dermatological disease classification. By integrating four specialized models within a unified pipeline, the system effectively handles multiple diagnostic tasks, including binary cancer screening, multi-class lesion classification, melanoma-versus-nevus differentiation, and common skin disease recognition. The use of EfficientNet-based architectures, attention mechanisms, and an ensemble consensus strategy improved the prediction accuracy and robustness. The combination of multiple probabilistic views enables more reliable decision-making than conventional single-model approaches. Additionally, the integration of multi-layer Grad-CAM enhances the interpretability by providing clear and meaningful visual explanations of the model predictions.

Beyond the model performance, the deployment of DermaNex as a full-stack web application demonstrates its practical applicability. Features such as risk-level categorization, user-friendly interaction, and explainable outputs make the system suitable for real-world use, particularly in regions with limited access to dermatological expertise. Although the system shows strong performance, it is designed as a decision-support tool rather than a replacement for professional medical diagnosis.

Future work will focus on improving dataset diversity and enhancing model generalization across various skin types and imaging conditions. In addition, clinical validation studies and exploration of advanced architectures will be conducted to further improve real-world reliability and performance.

Overall, DermaNex provides a scalable, interpretable, and efficient solution for AI-assisted dermatology, addressing the key limitations of existing approaches and contributing to more accessible healthcare solutions.

REFERENCES

1. R. Hay et al., "The global burden of skin disease in 2010," Journal of Investigative Dermatology, vol. 134, no. 6, pp. 1527–1534, 2014.
2. American Cancer Society, "Cancer Facts & Figures 2023," Atlanta, GA, USA: American Cancer Society, 2023.
3. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.
4. N. C. F. Codella et al., "Skin lesion analysis toward melanoma detection: A challenge at the 2017 ISBI," in Proc. IEEE 15th Int. Symp. Biomed. Imaging (ISBI), 2018, pp. 168–172.
5. M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. Int. Conf. Mach. Learn. (ICML), 2019, pp. 6105–6114.
6. Y. Liu et al., "A deep learning system for differential diagnosis of skin diseases," Nature Medicine, vol. 26, no. 6, pp. 900–908, 2020.
7. J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 7132–7141.
8. S. Woo et al., "CBAM: Convolutional block attention module," in Proc. Eur. Conf. Comput. Vis. (ECCV), 2018, pp. 3–19.
9. R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), 2017, pp. 618–626.
10. H. A. Haenssle et al., "Man against machine: Diagnostic performance of a deep learning CNN for dermoscopic melanoma recognition," Annals of Oncology, vol. 29, no. 8, pp. 1836–1842, 2018.
11. P. Tschandl et al., "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," Scientific Data, vol. 5, no. 1, pp. 1–9, 2018.
12. V. Rotemberg et al., "A patient-centric dataset and challenge proposal for identifying melanomas using clinical context," Scientific Data, vol. 8, no. 1, pp. 1–8, 2021.
13. T.-Y. Lin et al., "Focal loss for dense object detection," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), 2017, pp. 2980–2988.
14. Y. Cui et al., "Class-balanced loss based on effective number of samples," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2019, pp. 9268–9277.
15. K. He et al., "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.
16. A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.