International Research Journal of Multidisciplinary Scope (IRJMS), 2024; 5(4):711-721 IRJMS

Original Article | ISSN (O): 2582-631X DOI: 10.47857/irjms.2024.v05i04.01289

Densenet-201 for Skin Melanoma Classification: A Comprehensive Performance Evaluation and Analysis

Sheetal Nana Patil [1]*, Hitendra D Patil[1], Krishnakant P Adhiya[2], Prashant G Patil[3]

[1]*Department of Computer Engineering, Shri Shivaji Vidya Prasarak Sanstha's Bapusaheb Shivajirao Deore College of Engineering, Dhule (M.S) India, [2]Department of Computer Engineering, Shram Sadhana Bombay Trust's College of Engineering and Technology, Bambhori, Jalgaon (M.S) India, [3]Department of Electronics and Telecommunication Engineering, R C Patel Institute of Technology Shirpur, Dhule (M.S) India. *Corresponding Author's Email: sheetalpatil2000@gmail.com

Abstract

Skin cancer stands out as the predominant malignancy, necessitating prompt detection and intervention due to its potentially fatal nature. Distinguishing between cancerous and benign skin lesions poses a formidable challenge to visual assessment, underscoring the intricacy of accurate cancer detection. The inherent similarity in the appearances of various lesions further compounds the precision required for effective skin cancer identification. The rapidly evolving technological landscape, notably in the fields of machine learning and deep learning, has seen greater interest in resolving the categorization issues inherent in skin lesions. In our proposed research work, we deploy a deep learning model incorporating a pre-trained DenseNet architecture. This strategic utilization of advanced computational methods aims to enhance the discriminatory capabilities in skin cancer identification. For our research work, we used Melanoma cancer dataset, which contain 10540 dermoscopic labeled images. The study involves extensive preprocessing, including dataset inspection, label encoding, and distribution analysis. The research focuses on training DenseNet-201 architecture, evaluating its performance, and interpreting the results through various metrics.

Keywords: Average Pooling, Classification, DenseNet, Precision, Recall, Skin Melanoma.

Introduction

The integumentary system is the biggest organ in the human body, serving as a key barrier to external elements such as heat, light, and infections. It also contributes significantly to thermoregulation and serves as a reservoir for adipose tissue and water. Consequently, the imperative nature of identifying skin diseases is underscored. Among these conditions, skin cancer ranks prominently as one of the most prevalent forms of malignancy (1). Benign tumors typically exhibit a localized impact without causing significant harm to adjacent tissues, and they generally do not pose life-threatening risks. In contrast, malignant tumors can have profound consequences on an individual's well-being, necessitating immediate treatment to prevent further tissue damage and potential metastasis to other bodily regions (2). Recognizing the pivotal role of early detection and appropriate therapeutic interventions in enhancing patient recovery and survival rates. There is a rising attention in the possible application of intelligent computer systems as diagnostic tools for aiding healthcare professionals in the initial cancer identification (3). Melanoma stands out as a cancer with inherent unpredictability, often defying conventional expectations by manifesting even in its early stages (4). Skin diseases encompass diverse lesion types, characterized by variations in size, color, symmetry, bulge, and lesion margins. Artificial intelligence offers various diagnostic approaches, with one notable method relying on databases, particularly images. This strategy employs the principles of pattern recognition and deep learning, wherein iterative algorithms enable the computer to discern specific sets of clinical signs or images (5). This research introduces DenseNet201, a convolutional neural network based on the concept of dense connections, in which each layer is connected to every other layer in a feed-forward fashion. This dense connectivity promotes better information flow and gradient propagation throughout the network, resulting in

(Received 11th June 2024; Accepted 17th October 2024; Published 30th October 2024)

higher performance. Classification problems in several domains, including picture segmentation, brain tumor detection, pneumonia classification, cotton aphid severity rating, and galaxy classification, are handled by the proposed Dense Net model. These models have shown improved accuracy and performance compared to other classical classification models and neural network architectures. DenseNet models have the ability to transmit learned features backwards, reducing model parameters and improving local feature learning (6). DenseNet enhanced with a feature channel attention block for pneumonia classification. Their experimental results demonstrated superior performance in terms of F1-score, precision, accuracy, and recall compared to the standard DenseNet architecture (7). A modified DenseNet neural network model, a convolutional neural network architecture known for its dense connectivity patterns, to effectively classify MRI images into two distinct categories. This approach capitalizes on the dense connections between layers, allowing for efficient feature reuse and gradient flow throughout the network (8). By leveraging this modified architecture, the researchers achieved a notable accuracy rate of 93%, underscoring the model's effectiveness in accurately distinguishing between the specified categories within MRI images. An enhanced DenseNet classification network that surpasses ResNet50, ShuffleNetv2, GhostNet, MobileNetv3, and traditional DenseNet in classifying cotton aphid severity. This highlights CA_DenseNet_BC_100's superior accuracy in assessing cotton aphid infestations compared to other advanced models (9). DenseNet121, achieving an impressive test-set accuracy of approximately 89%. They utilized the Galaxy10 DECals Dataset to investigate transfer learning from pre-trained convolutional neural network models for galaxy classification, achieving the high accuracy within just 30 minutes of training (10). For emotion identification and classification, Kousis I. (11) employed a dense net model that had been pre-trained on the ImageNet dataset, and they were more accurate than VGG19. By fine-tuning a Dense net model that has been pre-trained in the ImageNet dataset, a transfer learning technique was utilized to increase the accuracy of emotion detection. The network model's goal is to extract features from pictures in order to recognize and categorize seven different emotional expressions. Existing machine learning and Deep learning techniques have made significant advancements in the detection and classification of skin cancer images. However, there are still some gaps that need to be addressed: Limited availability of large-scale and diverse skin cancer datasets: One of the major gaps in skin cancer image detection and classification is the lack of comprehensive and well annotated datasets (12). This can hinder training and generalizability of ML models, as they may not have enough diverse examples to learn from. Restricted interpretability of DL models: Convolutional neural networks have shown remarkable performance in image classification, but their interpretability is lacking. This means that it can be challenging to understand how and why these models make their predictions, which can be crucial in the medical field. Integration of clinical data: The combination of clinical data with image analysis represents a key gap in the identification and categorization of skin cancer images. The accuracy and dependability of skin cancer classification systems may be increased by integrating patient history, risk factors, and diagnostic data into the image analysis process. In general, DenseNet models have demonstrated efficacy and efficiency in handling classification tasks across several domains, achieving cutting-edge results on multiple benchmark datasets.

Methodology

In this study, we propose a methodology that integrates DenseNet-201 with a comprehensive dataset of dermoscopic images to develop an automated skin lesion detection system. The proposed approach involves several key steps, including data preprocessing, model training, and performance evaluation. Data preprocessing includes image augmentation techniques to increase the diversity of training samples and improve the model's generalization capabilities (13). The DenseNet-201 model is then trained on the preprocessed dataset using transfer learning, fine-tuning the pre-trained weights to adapt to the specific task of melanoma detection.

Dataset Used

The HAM10000 dataset has 10,015 dermatoscopic images distributed across seven classes. Below is the distribution of images for each class:

Melanocytic nevi (NV): 6,705 images (66.96%). Melanoma (MEL): 1,113 images (11.12%). Benign keratosis-like lesions (BKL): 1,099 images (10.98%). Basal cell carcinoma (BCC): 514 images (5.13%). Actinic keratoses (AKIEC): 327 images (3.27%). Vascular lesions (VASC): 142 images (1.42%). Dermatofibroma (DF): 115 images (1.15%)

Image Properties

The HAM10000 dataset consists of 10,015 high-resolution RGB dermatoscopic images in JPEG format, typically around 600x450 pixels. The images are annotated with seven skin lesion classes but do not include segmentation masks. Image resolutions and aspect ratios vary, with common preprocessing steps including resizing (e.g., to 224x224), normalization, and data augmentation to handle lighting variations and artifacts like hair. The dataset also provides metadata such as age, sex, and lesion location. It's widely used for classification tasks in skin lesion analysis. Figure 1 shows some sample images of benign and malignant lesions.

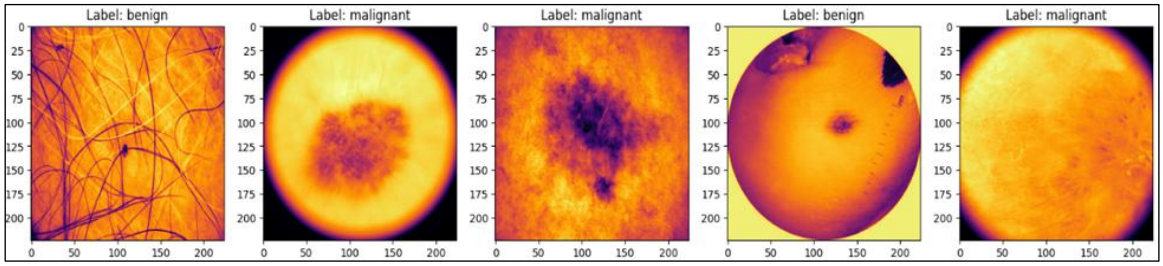

Figure 1: Benign and Malignant Images

Data Preprocessing

To facilitate effective model training, a robust preprocessing pipeline is established. This involves creating Pandas Data Frames for image paths and labels, label encoding, and shuffling the data to ensure unbiased model learning.

HAM-10000 Multiclass to Binary Conversion

Creating a binary dataset from the HAM10000 dataset involves classifying the images into two categories, typically based on the presence or absence of a specific condition, such as melanoma vs. non-melanoma. Following steps are used to convert multiclass to binary data. Step 1: Data Extraction: Load the HAM 10000 dataset and its annotations and Identify and separate images labeled as 'mel' (melanoma). Step 2: Label Assignment: Assign the label '1' to images of melanoma (mel). Assign the label '0' to images of all other categories (non-melanoma). Step 3: Data Balancing: Under-sampling: Randomly reducing the number of non-melanoma images. Over-sampling: Replicating melanoma images. Step 4: Preprocessing: Standardize image sizes, normalize pixel values and apply data augmentation techniques such as rotation, flipping and scaling—to enhance the diversity of the training data. Step 5: Splitting Data: Split the dataset into training, validation and test sets, ensuring that the distribution of melanoma and non-melanoma images is maintained in each subset. The binary dataset derived from HAM10000 would typically be used for a focused classification task, such as distinguishing between melanoma (mel) and non-melanoma lesions. This simplifies the multi-class classification problem into a binary classification problem, which in turn useful for machine learning or Deep Learning models.

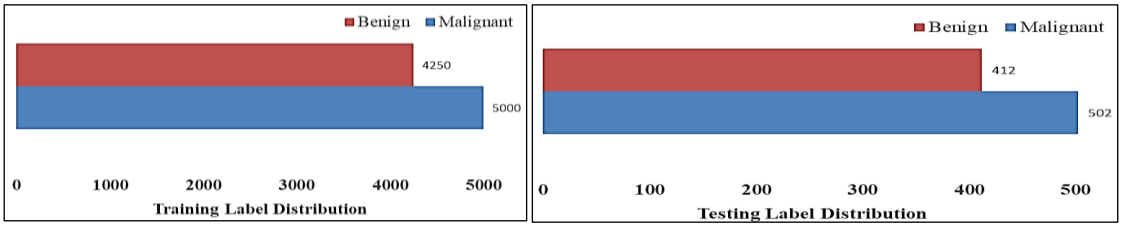

Figure 2: Original Train and Test Label Distribution

Data Augmentation

A complex data augmentation layer is provided to address over fitting problems and improve the model's generalizability. Horizontal and vertical flips, together with random zooming, help to diversity the training dataset, giving the model a more sophisticated grasp of melanoma features. From the original dataset, 85% of the photos are utilized for training, while the remaining 15% are used for testing. Figure 2 depicts the original train and test label distribution, whereas Figure 3 depicts the effect of augmentation on the original training and testing labels (14).

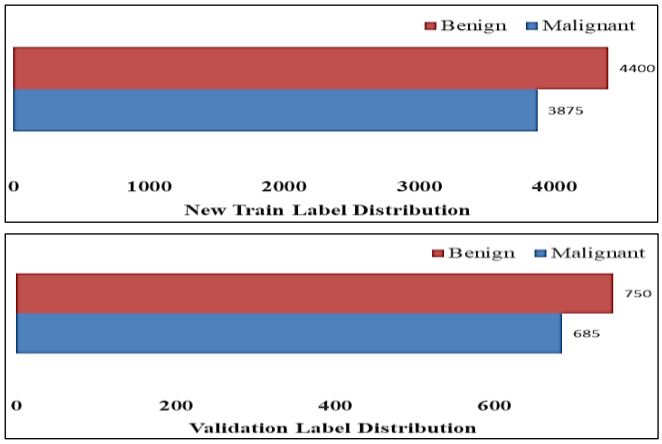

Figure 3: Train and Test Label Distribution After Augmentation

Dense Net201 Model Architecture

The DenseNet-201 architecture is a deeper version of DenseNet with 201 layers. In the provided code, the DenseNet-201 model is applied for the ImageNet dataset and then tuned for the melanoma classification. Dense Convolutional Networks, or DenseNet, connect each layer in a feed-forward fashion. This novel design handles various deep learning issues, including the vanishing-gradient problem. DenseNet supports strong feature propagation, facilitates effective feature reuse, and significantly decreases the overall amount of parameters by linking one layer to the next. Dense Net's key premise is that convolutional networks may achieve increased depth, accuracy, and training efficiency by establishing shorter connections between layers near the input and those near the output. This design concept promotes efficient information flow across the network. Dense Blocks: DenseNet introduces the notion of dense blocks, in which each layer takes input from every previous layer. This dense connection network helps to alleviate the vanishing gradient problem and allows for improved feature reuse. In a Dense net block, each layer gets feature maps from the layers before it and passes them on to the layers after it. This creates dense connectivity between layers. Mathematically, H_l denote the feature maps produced by the lth layer. Where $l = 0$ represents the input layer. The output of the lth layer is computed by the following equation.

$$H_l = H_{l-1} \oplus H_l . F_l \quad [1]$$

Bottleneck Layers: DenseNet utilizes bottleneck layers, consisting of 1x1 convolutional layers, to reduce the number of input feature maps before passing them to 3x3 convolutions as shown in below equation. This aids in reducing computational complexity while preserving representational capacity.

$$F_l = H_{l+1}[H_1\ (Conv\ (H_l)] \quad [2]$$

Transition Layers: Transition layers are used between dense blocks to limit the expansion of feature maps and minimize spatial dimensions. These transition layers are often composed of 1x1 convolution, batch normalization, and pooling procedures. The transition layer's output may be expressed mathematically as

$$H_l = H_1[Avarge\ pooling(F_l)] \quad [3]$$

Global Average Pooling (GAP): Instead of using fully connected layers at the end of the network, DenseNet employs global average pooling. This operation calculates the average value of each feature map, resulting in a single value for each channel. GAP reduces the number of parameters and helps mitigate over fitting. Given a 3D tensor of dimensions (H,W,C)where H is the height, W is

the width, and C is the number of channels (feature maps), Global Average Pooling computes the average value for each feature map as follows: For each channel c (where c=1, 2….. C):

$$GAP_C = \frac{1}{H.W}\sum_{i=1}^{H}\ \sum_{j=1}^{W}\ X_{i,j,C} \quad [4]$$

Where $X_{i,j,c}$ is the value at position (i, j) in the C^{th} feature map, GAP_c is the result of the Global Average Pooling for the ccc-th feature map. Growth Rate: The growth rate is a hyper parameter that defines how many new feature maps each layer in a dense block should produce. It determines the width of the network and affects the model's capacity. Figure 4 shows the detailed DenseNet201 architecture with its input and output layers. Let the growth rate be denoted by k. Consider the l^{th} layer in a dense block. The number of input feature maps to this layer is determined by the original number of feature maps plus the feature maps produced by all the previous layers in that block.

$$Input\ to\ layer\ l = Initial\ feature\ maps + k \times (l-1) \quad [5]$$

DenseNet201 Feature Extractor (densenet201): The base of the model is DenseNet201, a convolutional neural network renowned for its dense connectivity. It operates as a feature extractor, transforming input images through multiple convolutional layers, dense blocks, and transition layers. The resulting feature maps have dimensions of (7, 7, 1920). Output Shape: (None, 7, 7, 1920), Parameters (Trainable): 18,321,984 (These parameters are adapted during model training to learn image representations.). Global Average pooling 2D: Following feature extraction, global average pooling is applied to spatially condense the feature maps. This operation computes the average value for each feature map, resulting in a reduced spatial dimensionality. Output Shape: (None, 1920) Parameters (Trainable): 0 (No trainable parameters in this layer; it performs a fixed operation.). Dense Layer for Intermediate Representation (dense_2): A dense layer with 128 units and Rectified Linear Unit (ReLU) activation is introduced to capture higher-level abstractions from the global average pooled features. Output Shape: (None, 128) Parameters (Trainable): 245,888 (Trainable parameters tuned to distill meaningful information.). Final Dense Layer for Binary Classification (dense_3): The final dense layer comprises 2 units, employing a sigmoid activation function for binary classification (distinguishing 'malignant' or 'benign'). Output Shape: (None, 2) Parameters (Trainable): 258 (Adjustable parameters enabling the model to discern between the two classes.)In summary, the architecture involves a robust feature extraction phase with DenseNet201, followed by dimensionality reduction through global average pooling. Subsequent dense layers refine the features, culminating in a final layer for binary classification. The model encompasses 18,569,130 trainable parameters, facilitating the learning of intricate image representations. Above figure 4 (A) shows the detailed DenseNet201 layered architecture with its input and output layers and Figure 4 (B) indicates DenseNet 201 configuration.

Model Training: The training phase involves the compilation of the model with appropriate loss functions (binary cross entropy loss) and optimizers like Adam. Noteworthy callbacks, such as early stopping and learning rate reduction, are integrated to optimize training efficiency. These callbacks are used to prevent over fitting and enhance training efficiency (15). The model undergoes training epochs on the designated training dataset, periodically validated on an independent dataset. Model Evaluation: The trained model is evaluated using various metrics such as accuracy, precision, recall, and F1-score. The confusion matrix and Region of convergence curves are employed to gain insights into the model's classification and performance.

Hyper Parameter Optimization

In the DenseNet201 melanoma classifier, binary cross-entropy loss and the Adam optimizer were used to ensure adaptive learning. A batch size of 32 balanced efficiency and performance with models trained for 15 epochs to prevent over fitting. DenseNet201 lacks dropout layers by default. Adding Dropout (e.g., with a rate of 0.2 to 0.5) after dense layers can prevent over fitting, especially for small or imbalanced datasets like HAM10000.

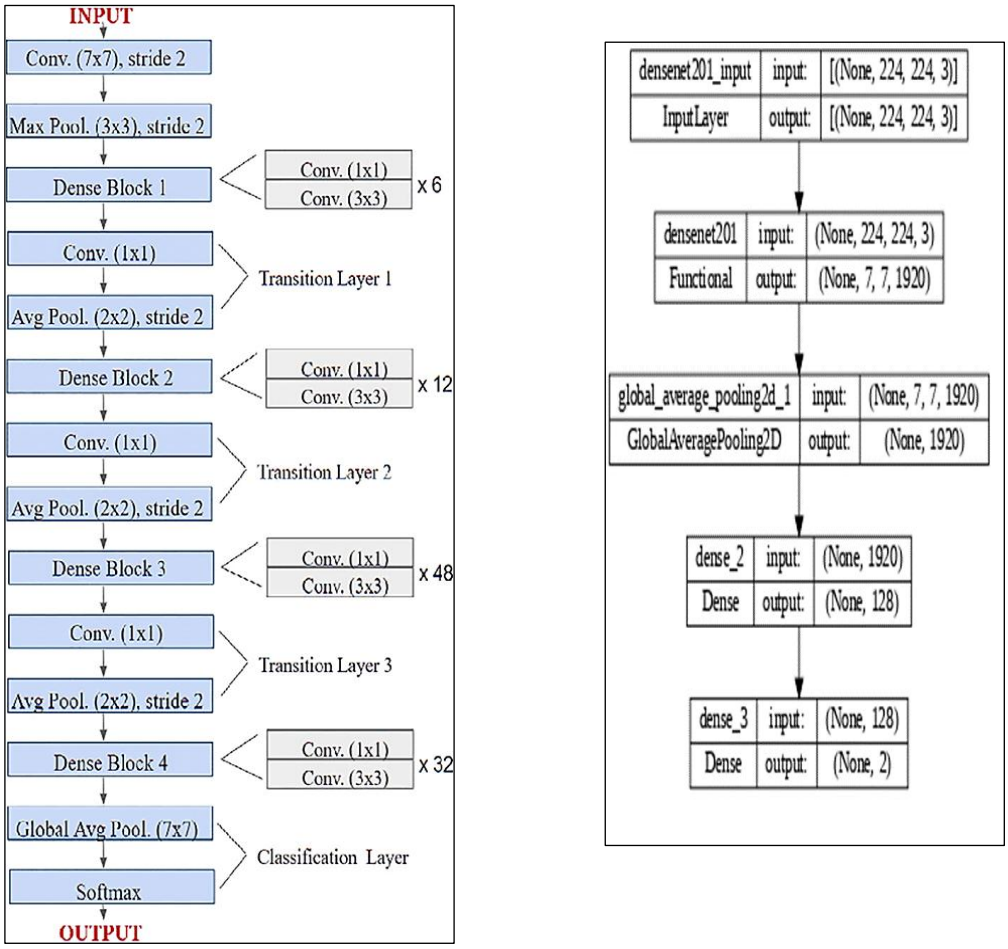

Figure 4 (A): DenseNet201 Layered Architecture **(B):** DenseNet201 Configuration

Results and Discussion

The DenseNet-201 model shows encouraging results in melanoma classification, with an accuracy of 92%. Metrics such as precision, recall, and F1-score indicate that malignant and benign classes function similarly. The study highlights the implications of these findings in terms of melanoma diagnosis. In Figure 5, when the number of epochs grows, training and validation loss decrease dramatically, while training accuracy rises to 94% and validation accuracy approaches 93%.

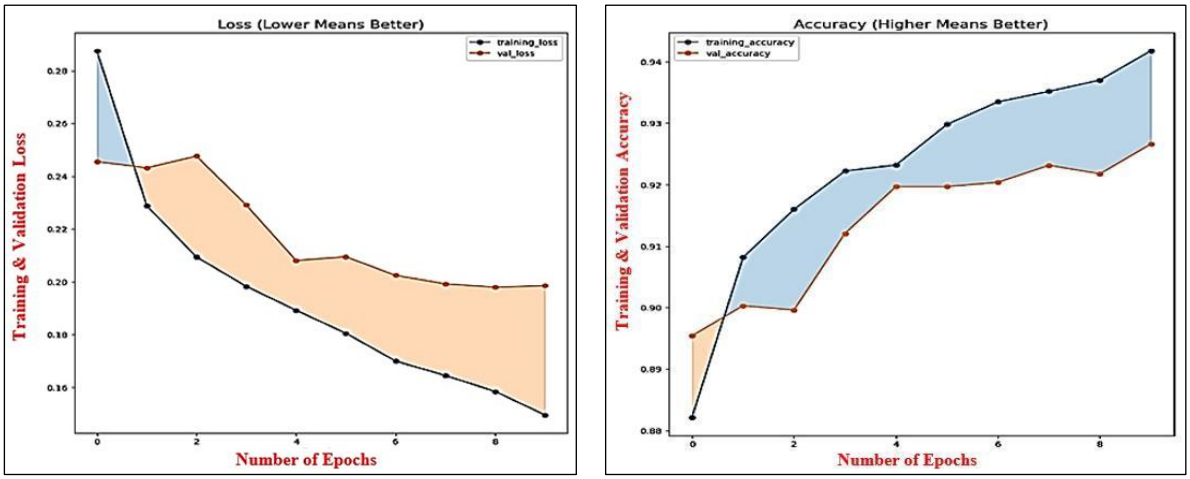

Figure 5: Densenet-201 Model Loss and Accuracy Curves for Training and Validation Samples

Table 1: Densenet201 Performance Comparison with Existing Models

Model	Precision	Recall	F1-score	Training Accuracy	Testing accuracy
ResNet (16)	84.94	97.50	90.79	92.78	85.20
VGG-16 (17)	88.27	95.18	91.59	88.83	86.09
Mobilenet (18)	84.62	94.57	89.32	92.93	82.62
DenseNet201	92.04	91.39	91.41	94.18	91.39

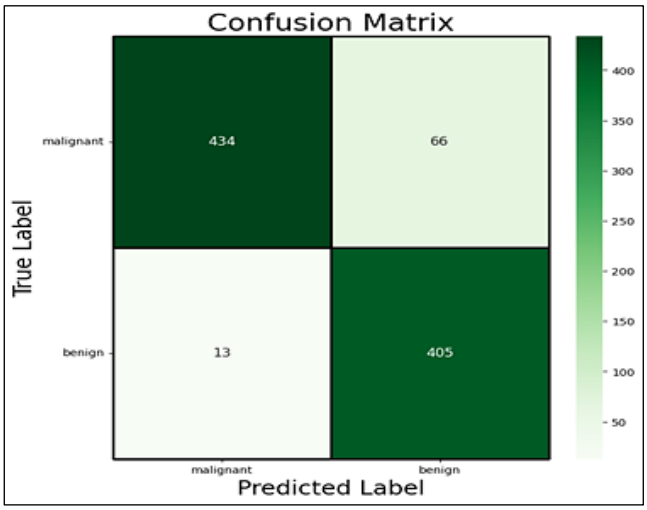

Figure 6: Confusion Matrix of the Proposed Densenet201 Model

In Table 1, DenseNet201 demonstrated high precision (92.04%) and recall (91.39%), resulting in an F1-score of 91.41%. This indicates a balanced performance in identifying both positive and negative instances. The model exhibited strong training accuracy of 94.18% and maintained a high testing accuracy of 91.39%, suggesting its robustness in generalization. Overall, DenseNet201 emerged as the top-performing model in terms of F1-score and testing accuracy, making it a promising choice for image classification tasks. However, the selection of the most suitable model should consider specific application requirements and computational constraints. In Figure 6, the Confusion matrix for DenseNet201 can be constructed to provide a more detailed understanding of its performance in classifying instances. A confusion matrix is a table that is often used to describe the performance of a classification model on a set of test data for which the true values are known. Each row of the matrix represents the instances in an actual class, while each column represents the instances in a predicted class.

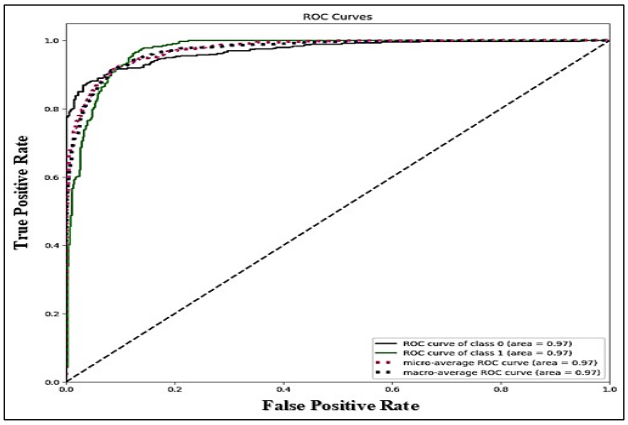

Figure 7: Receiver Operating Characteristics (ROC) Curve for Binary Classification

The Receiver Operating Characteristics (ROC) curve shown in Figure 7 is a graphical representation used to evaluate the performance of a binary classification model. It plots the True

Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold settings. The curve rises sharply towards the top-left corner of the plot, indicating a high True Positive Rate and a low False Positive Rate. This suggests that the model is effective in identifying true positives while minimizing false positives. The ROC value for the model is 0.95, which indicates excellent discriminatory power. An AUC value close to 1.0 signifies that the model performs well in differentiating between melanoma and non-melanoma cases. Figure 8 represents simulated results for Actual vs. predicted classification labels using Densenet201 algorithm. The majority of the model's predictions align with the actual labels, often with very high probability scores (e.g., 99.99% for malignant and benign cases). This indicates that DenseNet-201 is highly confident in its correct classifications. The images reveal some instances of misclassification, for example, one image labeled "Actual Label: malignant" is predicted as benign with a probability of 67.61%. This indicates that the model occasionally fails to detect certain melanoma cases. There are also benign cases misclassified as malignant, such as an image with an actual benign label but predicted as malignant with a probability of 85.05%. This suggests that some benign lesions exhibit features that the model mistakenly identifies as melanoma. The high probability scores in correctly identified malignant cases (e.g., 100.00% probability) demonstrate the model's strong sensitivity to melanoma features. This is crucial for minimizing missed diagnoses of melanoma, which can have serious clinical implications. In this study, we evaluated the performance of four popular CNN architectures: ResNet, VGG-16, Mobilenet, and DenseNet201, for the task of image classification. Each model was trained and tested on a standardized dataset to assess its precision, recall, F1-score, training accuracy, and testing accuracy. ResNet achieved a high recall of 97.50%, indicating its effectiveness in correctly identifying positive instances from the dataset. However, its precision of 84.94% suggests a moderate level of false positives. The F1-score, which combines precision and recall, was calculated at 90.79%. The model demonstrated robust performance during training with an accuracy of 92.78%, which slightly dropped to 85.20% on unseen test data. VGG-16 exhibited slightly higher precision compared to ResNet, with a value of 88.27%, while maintaining a recall of 95.18%. Consequently, the F1-score for VGG-16 was computed at 91.59%, indicating a balanced performance in terms of precision and recall. The training accuracy was recorded at 88.83%, with a marginal decrease to 86.09% on the testing dataset. The DenseNet201 model, while powerful for skin melanoma classification, has several limitations. It requires high computational resources due to its deep architecture and large number of parameters, leading to longer training times and high memory consumption, which can be challenging for high-resolution dermatoscopic images. Additionally, DenseNet201 is prone to over fitting, especially on smaller datasets like HAM10000, and struggles with class imbalance, often favoring majority classes over rare ones like melanoma. Extensive hyper parameter tuning is often necessary to optimize performance, and fine-tuning for domain-specific tasks can be complex in transfer learning scenarios. Furthermore, DenseNet201, like other deep learning models, lacks interpretability, which is crucial in medical applications where explaining ability is important for decision-making. The model's performance is also sensitive to image quality, with noise, artifacts, and lighting variations affecting predictions. Despite its dense connections, training such a deep network still poses risks like the vanishing gradient problem, particularly when fine-tuning for specialized tasks like melanoma classification.

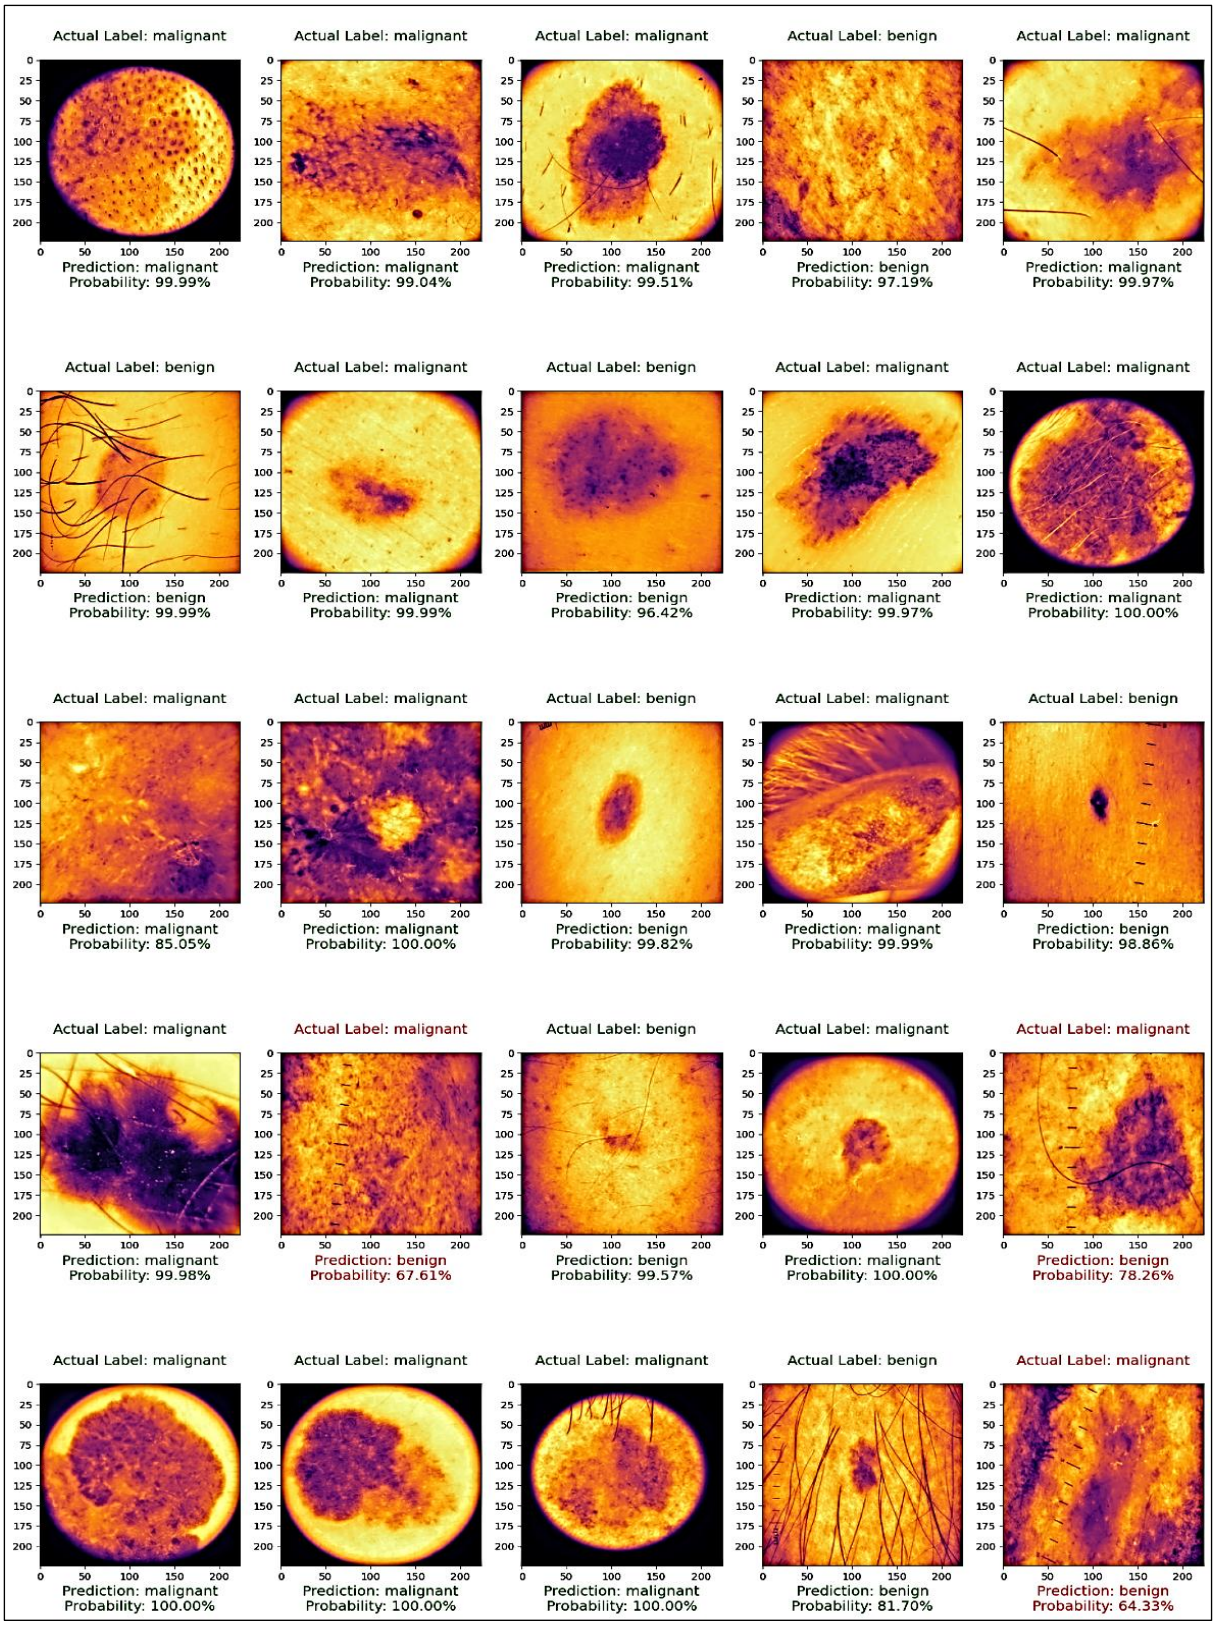

Figure 8: Actual vs. Predicted Classification Results with Probability

Conclusion

This research successfully explores the application of deep learning models in melanoma classification. The comprehensive preprocessing pipeline, data augmentation techniques, and the adoption of a powerful CNN architecture contribute to the model's robust performance. The study concludes with insights into potential future

improvements and applications of the developed model in clinical settings. In conclusion, research studies demonstrate that DenseNet architectures offer significant advantages in skin melanoma detection tasks, including improved accuracy, efficient parameter usage, effective feature extraction, and interpretability of learned features. These findings highlight the potential of DenseNet-based models as valuable tools for assisting dermatologists in early and accurate diagnosis of melanoma.

Due to its deep architecture, DenseNet-201 is prone to over fitting, especially when trained on small or imbalanced datasets. This can result in poor generalization to new, unseen data, limiting its practical applicability. DenseNet-201, like many deep learning models, functions as a black box, making it challenging to interpret its decision-making process. This lack of interpretability can be a barrier to clinical acceptance, as medical professionals often require understanding of how decisions are made. DenseNet-201 has demonstrated impressive performance in various image classification tasks, including melanoma detection; several limitations and challenges need to be addressed for its effective deployment in clinical settings. DenseNet-201 is computationally intensive, requiring significant processing power and memory. This can be a challenge for deployment on standard medical devices or in settings with limited computational resources.

In conclusion, the DenseNet201 model demonstrates high potential for skin melanoma detection due to its deep architecture and efficient feature reuse. Its dense connectivity allows for better gradient flow, making it effective in extracting complex patterns from dermatoscopic images. However, the model's success is contingent on several factors. While DenseNet201 achieves impressive performance in terms of accuracy and classification metrics, it comes with significant computational demands, requiring substantial memory and long training times. Additionally, over fitting remains a challenge, especially when working with smaller or imbalanced datasets like HAM10000. Mitigating these issues through data augmentation, regularization, and hyper parameter tuning is essential to improve generalization. Despite its power, DenseNet201's lack of interpretability limits its direct applicability in clinical settings, where explain ability is crucial for medical decision-making. Furthermore, sensitivity to image quality and class imbalance can hinder model performance, necessitating careful preprocessing and weighting strategies. Overall, while DenseNet201 proves to be an effective model for skin melanoma detection, addressing these limitations is key to making it a more robust and clinically viable tool.

Abbreviations

ML: Machine learning, DL: Deep learning, CNN: Convolutional Neural Network, GAP: Global Average Pooling, ReLU: Rectified Linear Unit, ROC: Region of convergence, TPR: True Positive Rate, FPR: False Positive Rate.

Acknowledgement

We are grateful to the Research Center of SSVPS College of Engineering, Dhule, for granting permission to conduct our research on using the DenseNet-201 algorithm for automated melanoma detection.

Author Contributions

SP conceptualized the study and drafted the manuscript. HP collected and analyzed the data. SP, HP and KA finalized and critically reviewed the manuscript.

Conflicts of Interests

The authors do not have any conflicts of interest.

Ethics Approval

Approved by the Institutional Review Board.

Funding

No funding received.

References

1. Manne R, Kantheti S, Kantheti S. Classification of Skin cancer using deep learning, Convolutional Neural Networks-Opportunities and vulnerabilities-A systematic Review. International Journal for Modern Trends in Science and Technology. 2020; 6: 2455-3778.
2. Haenssle HA, Fink C, Schneiderbauer R, Toberer F, Buhl T, Blum A, Kalloo A, Hassen AB, Thomas L, Enk A, Uhlmann L. Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. Annals of oncology. 2018; 29(8):1836-42.
3. Adegun AA, Viriri S. FCN-based DenseNet framework for automated detection and classification of skin lesions in dermoscopy images. IEEE Access. 2020; 8:150377-96.
4. Bhatt H, Shah V, Shah K, Shah R, Shah M. State-of-the-art machine learning techniques for melanoma skin

cancer detection and classification: a comprehensive review. Intelligent Medicine. 2023; 3(03):180-190.
5. Malik P, Pathania M, Rathaur VK. Overview of artificial intelligence in medicine. Journal of family medicine and primary care. 2019; 8(7):2328-31.
6. Zhou Y, Chang H, Lu X, Lu Y. Dense U-Net: Improved image classification method using standard convolution and dense transposed convolution. Knowledge-Based Systems. 2022; 254:109658-78.
7. Wei L, Ding K, Hu H. Automatic skin cancer detection in dermoscopy images based on ensemble lightweight deep learning network. IEEE Access. 2020; 8:99633-647.
8. Girdhar N, Sinha A, Gupta S. DenseNet-II: An improved deep convolutional neural network for melanoma cancer detection. Journal of soft computing. 2023; 27(18):13285-304.
9. Fraiwan M, Faouri E. On the automatic detection and classification of skin cancer using deep transfer learning. Sensors. 2022; 22(13):4963.
10. Thurnhofer Hemsi K, Domínguez E. A convolutional neural network framework for accurate skin cancer detection. Neural Processing Letters. 2021; 53(5):3073-093.
11. Kousis I, Perikos I, Hatzilygeroudis I, Virvou M. Deep learning methods for accurate skin cancer recognition and mobile application. Electronics. 2022; 11(9):1294-99.
12. Patil PG, Jaware TH, Patil SP, Badgujar RD, Albu F, Mahariq I, Al-Sheikh B, Nayak C. Marathi speech intelligibility enhancement using I-AMS based neuro-fuzzy classifier approach for hearing aid users. IEEE Access. 2022 Nov 18;10:123028-42.
13. Patil SN, Patil HD. Modified CNNs for Automated Binary Classification of Skin Lesions: Benign vs. Malignant. In2023 3rd International Conference on Advancement in Electronics & Communication Engineering (AECE) 2023 Nov 23 (pp. 958-964). IEEE.
14. Patil SN, Patil HD. Performance Analysis of Image Denoising Techniques in Skin Cancer Detection. InInternational Conference on Data, Engineering and Applications. 2022 Dec 23:165-182. Singapore: Springer Nature Singapore.
15. Patil SN, Patil HD. Enhancing skin lesion segmentation with U-Net++: Design, analysis, and performance evaluation. Review of Computer Engineering Research. 2024;11(1):30-44.
16. Zhao C, Shuai R, Ma L, Liu W, Hu D, Wu M. Dermoscopy image classification based on Style GAN and DenseNet201. IEEE Access. 2021; 9:8659-79.
17. Qureshi AS, Roos T. Transfer learning with ensembles of deep neural networks for skin cancer detection in imbalanced data sets. Neural Processing Letters. 2023; 55(4):4461-79.
18. Thanka MR, Edwin EB, Ebenezer V, Sagayam KM, Reddy BJ, Genera H, Emadifar H. A hybrid approach for melanoma classification using ensemble machine learning techniques with deep transfer learning. Computer Methods and Programs in Biomedicine. 2023; 3:100103-127.