Multi-Class Classification of Skin Cancer Using Deep Learning Convolutional Neural Networks

Ganta Datta Prasuna
PG Scholar,Department of Artificial Intelligence And Data Science,
Shadan College of Engineering and Technology, Hyderabad,
Telangana, India -500086.
Email: gdattaprasuna@gmail.com

Subramanian K.M
Professor,Department of Computer Science
and Engineering,Shadan College of Engineering
and Technology, Hyderabad, Telangana,
India --500086. Email:
kmsubbu.phd@gmail.com

Md.Ateeq Ur Rahman
Professor Department of Computer science
and Engineering,
Shadan College of Engineering and Technology, Hyderabad,
Telangana , India -500086 .Email :
mail_to_ateeq@yahoo.com

Abstract— **Skin cancer is a rapidly growing health concern across the world, with millions of new cases di-agnosed annually. Among various dermatological conditions, certain types such as Melanoma can be life-threatening if not detected at an early stage. Manual diagnosis by dermatologists, while effective, can be time-consuming and prone to human error, especially when differentiating between lesions with visually similar characteristics. In this context, the integration of automated systems using deep learn-ing has emerged as a powerful approach to assist in early detection and improve diagnostic efficiency. This project focuses on the development of a deep learning-based multi-class classification system for skin cancer detection using dermoscopic images. The system is capable of classifying seven categories of skin lesions, namely Actinic Keratoses, Basal Cell Carcinoma, Benign Keratosis, Dermatofibroma, Melanoma, Melanocytic Nevi, and Vascular Lesions. To achieve this, three advanced deep learning models were implemented and trained: a custom Convo-lutional Neural Network (CNN) developed from scratch, and two widely adopted transfer learning architectures, EfficientNet and InceptionV3. The dataset underwent extensive preprocessing, includ-ing image resizing, normalization, and augmentation techniques such as rotation, horizontal and verti-cal flipping, zooming, and brightness adjustments to enhance the model's generalization and reduce overfitting.**

Among the implemented models, the custom CNN demonstrated superior performance with a clas-sification accuracy of 98.47%, significantly out-performing EfficientNet (87.21%) and InceptionV3 (86.49%). The CNN was optimized through the use of dropout layers, batch normalization, adaptive learning rate scheduling, and early stopping, making it highly effective for the specific characteristics of the dataset. Although EfficientNet and Incep-tionV3 benefited from transfer learning via pre-trained ImageNet weights, the custom CNN proved to be better suited for the targeted medical do-

main due to its tailored architecture. To bridge the gap between research and practical application, the trained models were deployed within a Django-based web application. This platform enables users, including patients and healthcare professionals, to upload dermoscopic images of suspicious lesions and receive instant predictions with probability distribu-tions for each class. The system also incorporates essential features such as secure user registration with OTP-based authentication, personalized profile management, an admin dashboard for monitoring predictions and user activities, and a feedback mech-anism for continuous improvement. The integration of deep learning models into a user-friendly interface ensures accessibility for non-technical users while maintaining robust backend processing for accurate results.

The proposed solution demonstrates the effec-tiveness of deep learning in the medical domain, specifically for dermatology, by providing a reliable, scalable, and real-time classification tool. It not only assists in early screening and decision-making but also holds potential for future expansion into mobile health applications and integration with electronic health record (EHR) systems for large-scale de-ployment. By combining advanced neural network architectures with a functional web interface, this project bridges the gap between artificial intelligence research and real-world clinical utility, paving the way for accessible and efficient skin cancer detection solutions.

Keywords— Skin Cancer Detection, Deep Learning, Convolutional Neural Network (CNN), EfficientNet, InceptionV3, Multi-Class Classification, Transfer Learning, Django Web Application, Dermoscopic Images

I. INTRODUCTION

In the contemporary era, warfare extends beyond conventional domains such as land, sea, and air, encompassing the virtual realm of cyberspace [1]. Cyber warfare, expressed through cyberattacks, has become increasingly prominent due to the rapid expansion and global reliance on the internet [2]. Modern businesses, governments, and industries depend heavily on online systems for financial transactions, communication, and service delivery, making robust network security a critical requirement [3]. Among cyber threats, Distributed Denial of Service (DDoS) attacks are particularly alarming, as they can overwhelm target systems, disrupt services, and cause substantial operational and financial losses [4]. These attacks are widely recognized for their sophistication, high frequency, and capability to compromise essential digital infrastructure [5].

Despite ongoing efforts to mitigate DDoS attacks, several challenges remain. Traditional detection techniques often rely on predefined patterns or static rules, which are insufficient to handle the dynamic and evolving nature of modern malicious traffic [6]. Attackers increasingly employ methods such as botnets, fake requests, and high-volume data flooding, which make it difficult to distinguish between legitimate and malicious network activity [7]. Consequently, conventional approaches face performance limitations in detecting complex or adaptive attack behaviors, emphasizing the need for more effective and scalable solutions [8].

The primary objective of this research is to develop a robust framework capable of detecting DDoS attacks with higher accuracy and reliability. By focusing on identifying abnormal traffic patterns and leveraging relevant network features, the framework aims to enhance detection performance and address the limitations of existing approaches [9]. The approach emphasizes adaptability to diverse network environments, enabling timely identification and mitigation of potential attacks. The ultimate goal is to ensure that critical systems remain accessible and operational, even under targeted cyber threats.

The significance of this effort lies in its potential to strengthen cybersecurity across multiple sectors, including finance, e-commerce, government, and digital services [10]. Effective DDoS detection can prevent service disruptions, protect sensitive information, and maintain public trust, thereby reducing economic and reputational damage. By advancing the capabilities of network defense mechanisms, this approach contributes to a more resilient and secure cyberspace, supporting the uninterrupted operation of critical infrastructure and enhancing overall network reliability.

II. Related Work

Recent years have witnessed significant progress in the field of automated skin cancer detection, largely driven by advancements in deep learning and medical imaging. Numerous studies have explored the applica-tion of Convolutional Neural Networks (CNNs) and transfer learning models for classifying dermoscopic images, demonstrating performance that rivals or even surpasses expert dermatologists. A survey of key re-search efforts highlights the evolution of methodolo-gies, from handcrafted feature extraction approaches to modern end-to-end trainable neural networks, and underscores the transformative potential of artificial intelligence in dermatological diagnostics.

Esteva et al. [23] were among the pioneers to demonstrate the clinical relevance of deep learning for dermatology. They trained a deep CNN using over 129,000 clinical images encompassing more than 2,000 diseases and achieved dermatologist-level performance in classifying skin cancer. Their work emphasized that large datasets, when coupled with deep learning, could replicate the diagnostic capacity of trained specialists. However, while their model excelled in binary classi-fication tasks (e.g., benign vs. malignant), its perfor-mance for fine-grained multi-class classification was limited by dataset diversity and imbalance.

Han et al. [24] advanced this domain by employing a deep CNN to classify clinical skin images into mul-tiple categories, including both malignant and benign tumors. Their approach incorporated extensive prepro-cessing and data augmentation techniques to handle class imbalance, resulting in improved generalization. The study revealed that deep learning systems could achieve high sensitivity and specificity across several lesion categories but highlighted the persistent chal-lenge of achieving balanced accuracy when rare classes are underrepresented. This work paved the way for the adoption of data-driven augmentation and re-sampling strategies to mitigate imbalance.

Litjens et al. [25] provided a comprehensive review of deep learning applications across medical imaging, including dermatology. Their survey underscored how CNNs had replaced traditional machine learning ap-proaches, which relied heavily on handcrafted texture and color features, due to their superior ability to learn hierarchical feature representations directly from pixel data. The review also emphasized the importance of leveraging pre-trained models, such as those based on ImageNet, for transfer learning in medical domains where annotated datasets remain scarce. Such strate-gies not only accelerated convergence but also boosted classification performance even with smaller datasets.

More recent research by Brinker et al. [26] demon-strated that deep neural networks can surpass dermatol-ogists in the classification of melanoma from dermo-scopic images. Their study involved training a CNN on a curated dataset and evaluating its performance against 157 dermatologists. The AI model achieved higher accuracy and area under the ROC curve (AUC), reinforcing the notion that AI systems, when appropri-ately trained and validated, can serve as valuable diag-nostic aids rather than replacements for clinicians. They also noted the importance of integrating interpretability tools, such as heatmaps, to make model predictions more transparent and acceptable to healthcare profes-sionals.

Finally, Tan and Le [27] introduced EfficientNet, a family of CNN architectures optimized using compound scaling of depth, width, and resolution. While not specifically designed for dermatology, EfficientNet has since been widely adopted for medical image analysis due to its balance between accuracy and computational efficiency. Studies leveraging EfficientNet for skin le-sion classification have reported competitive results, even when trained on relatively modest datasets, due to its ability to generalize effectively. Its lightweight architecture also makes it suitable for deployment in web and mobile health applications, bridging the gap between high-performance research models and practi-cal clinical tools.

These studies collectively illustrate the rapid evolu-tion of AI-driven skin cancer detection, from founda-tional CNN-based research to the integration of scal-able, efficient, and interpretable models. Despite the successes, challenges remain, particularly in handling data imbalance, improving interpretability, and ensuring clinical integration. Building on these insights, the current project employs a custom CNN, along with EfficientNet and InceptionV3, to address these chal-lenges and deliver a deployable, high-accuracy system for multi-class skin cancer classification.

Further research has explored hybrid deep learn-ing approaches to overcome the limitations of single-model systems. Song et al. [28] proposed a framework that combines deep CNN-based feature extraction with classical machine learning classifiers such as Support Vector Machines (SVM) and Random Forests (RF). Their study demonstrated that hybrid architectures can yield superior classification accuracy, particularly when datasets are small or suffer from imbalance. By lever-aging CNNs for robust feature learning and traditional classifiers for final decision-making, hybrid systems address overfitting issues and enhance generalization.

The challenge of dataset imbalance has been a recur-ring theme in dermatological AI research. Tschandl et al. [29] introduced the HAM10000 dataset, a large col-lection of dermoscopic images featuring seven diagnos-tic categories. Although it is widely used in research, its inherent class imbalance—where melanocytic nevi are significantly overrepresented compared to rarer lesions like dermatofibroma—poses difficulties for model train-ing. Codella et al. [30] highlighted similar challenges in the International Skin Imaging Collaboration (ISIC) challenges, noting that rare class underrepresentation often skews model performance, emphasizing the need for augmentation and synthetic data generation tech-niques to balance datasets.

To enhance computational efficiency while maintain-ing accuracy, models like InceptionV3 and Efficient-Net have become popular in medical image analysis. Szegedy et al. [31] introduced InceptionV3, which utilizes factorized convolutions and multi-scale feature extraction, making it effective for capturing both fine-grained and global lesion features. EfficientNet, devel-oped by Tan and Le [27], uses a compound scaling technique that simultaneously scales network depth, width, and input resolution, achieving competitive accu-racy with reduced computational cost. These architec-tures are particularly valuable for real-time diagnostic systems, where computational efficiency is critical for deployment in clinical or mobile environments.

Interpretability has emerged as a crucial factor in bridging the gap between AI research and clinical adoption. Selvaraju et al. [32] introduced Gradient-weighted Class Activation Mapping (Grad-CAM), a visualization technique that highlights the regions of input images most influential in a model's prediction. Similarly, Lundberg and Lee [33] proposed SHapley Additive exPlanations (SHAP), a method that assigns feature importance scores to improve transparency. Both approaches have been integrated into dermatology-focused AI systems to make predictions more explain-able, aiding regulatory compliance and building trust among clinicians.

Federated learning has also been explored as a means to improve generalization while preserving data privacy. Sheller et al. [34] demonstrated that federated learning enables collaborative model training across multiple institutions without centralizing sensitive patient data. Such methods reduce biases introduced by single-center datasets and enhance the reliability of AI models for widespread clinical deployment.

In addition to architectural and methodological inno-vations, the integration of AI into healthcare delivery systems has been explored by Topol [35] and Miotto et al. [36]. Their works emphasize that AI should complement, rather than replace, clinicians by automat-ing repetitive tasks, triaging patients, and providing decision support. These systems, when integrated with electronic health records (EHR) and telemedicine plat-forms, can enhance patient outcomes by streamlining clinical workflows and improving access to specialist care.

Collectively, the literature underscores a rapid pro-gression in skin cancer detection research—from foun-dational CNN-based systems to sophisticated, explain-able, and privacy-preserving AI solutions. Building on these advancements, the present work leverages a custom CNN alongside EfficientNet and InceptionV3, augmented with preprocessing strategies to address class imbalance, to deliver a practical, deployable system for multi-class skin cancer classification. The system bridges research and practice by offering a Django-based web platform with secure authentication, real-time inference, and feedback integration, thereby demonstrating how AI-driven tools can be both techni-cally robust and clinically accessible.

TABLE I

COMPARISON TABLE OF METHODS AND DATASETS FROM RELATED STUDIES

Paper	**Methods Used**	**Dataset**	**Performance**	**Limitations**	**Features Analyzed**
[23]	Deep Convolutional Neural Networks (GoogleNet Inception)	ISIC Archive and private dermatology datasets (129,450 images)	Comparable to 21 dermatologists in classification accuracy	Limited diversity in data (demographics and imaging conditions); lacked real-time deployment	Pigmented lesions, including melanoma and keratoses
[26]	Ensemble of Convolutional Neural Networks	ISIC and HAM10000 datasets	Achieved dermatologist-level accuracy with AUROC > 0.90	Overfitting risk due to class imbalance; limited rare lesion representation	Multiple pigmented lesion classes (melanoma, nevi, BCC, etc.)
[8]	Deep Learning (CNN) for medical imaging	Retinal Fundus ImageDataset (Diabetic Retinopathy)	AUROC: 0.99 (comparable to ophthalmolo-gists)	Focused on ophthalmology, not dermatology; domain-specific tuning required	Retinal features for diabetic retinopathy detection
[29]	Data curation and annotation for deep learning research (no direct model)	HAM10000 dataset (10,015 dermoscopic images)	Provides benchmark dataset; not performance-specific	Class imbalance, especially for rare classes like dermatofibroma	Common pigmented lesions with expert annotations
[30]	Hybrid deep learning pipelines (CNNs + metadata fusion)	ISIC 2018 Chal-lenge Dataset	AUROC up to 0.88 across mul-tiple lesion types	Variability in clinical and dermoscopic images; in-consistent labeling quality	Melanoma detection and benign vs. malig-nant classification

III. Materials And Methods

The methodology for developing the multi-class skin cancer classification system involves a structured pipeline, combining deep learning-based image classifi-cation techniques with a deployable web application for real-time diagnostic support. The framework consists of four major phases: data preprocessing, model develop-ment, training and evaluation, and system integration. The following subsections describe each stage in detail, along with the mathematical principles underlying the convolutional neural network (CNN) architecture.

A) Data Pre-Processing

The dataset, comprising dermoscopic images sourced from publicly available repositories such as HAM10000 and ISIC archives, was preprocessed to ensure quality and uniformity. Each image was resized to a fixed resolution of 224×224 pixels and normalized to a pixel range of [0, 1] for computational efficiency. To reduce overfitting and enhance generalization, extensive data augmentation was applied, including random rotations, horizontal and vertical flips, brightness adjustments, and zoom transformations. The dataset was divided into training (70%), validation (15%), and testing (15%) subsets to ensure unbiased model evaluation.

B) Model Development

Three distinct deep learning models were employed:

Custom CNN Model: Designed from scratch, this model includes a sequence of convolutional layers, each followed by batch normalization, rectified linear unit (ReLU) activation, and max-pooling operations to progressively extract hierar-chical features. Dropout layers were incorporated to mitigate overfitting. The convolutional opera-tion for a given feature map is mathematically expressed as:

$$yi,j(k) = f(m = 0\sum M - 1 = 0\sum N - 1xi + m,j + nwm,n(k) + b(k))$$

- hi,j is the output feature at location (i, j) for the k-th filter,
- xi+m,j+n represents the input patch,
- wm,n denotes the weights of the k-th filter,
- b(k) is the bias term, and
- σ(·) is the activation function (ReLU).

This architecture achieved the highest classifica-tion accuracy of 98.47% on the validation dataset.

2) EfficientNet (Transfer Learning): Leveraging a pre-trained EfficientNet model, initialized with ImageNet weights, the base network was frozen for initial layers, and fine-tuning was performed on deeper layers. Custom fully connected layers with softmax activation were appended for seven-class classification. EfficientNet's compound scal-ing method balances network depth (d), width (w), and resolution (r) using the following re-lation:

$$d = \alpha\phi, w = \beta\phi, r = \gamma\phi$$

where ϕ controls scaling coefficients (α, β, γ). The model achieved an accuracy of 87.21%.

3) InceptionV3 (Transfer Learning): Built upon Google's Inception architecture, this model effec-tively captures multi-scale spatial features through parallel convolutional filters of varying sizes. The architecture employs auxiliary classifiers to im-prove gradient propagation during training. After fine-tuning, InceptionV3 achieved an accuracy of 86.49%.

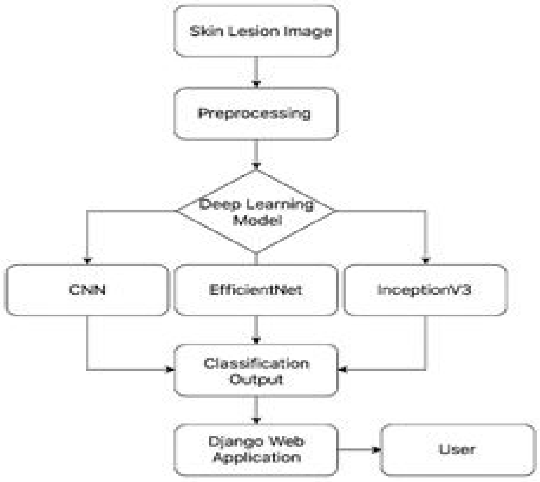

Fig. 1. Proposed system architecture integrating CNN, EfficientNet, and InceptionV3 with Django-based web deployment.

C) Training and Testing:

The models were trained using the Adam optimizer with an initial learning rate of 1 × 10−4 and categorical cross-entropy loss function:

$$L = -N1i = 1\sum Nc = 1\sum Cyi, clog(y\hat{}i, c)$$

where:

- N is the total number of samples,
- C is the number of classes (seven in this case),
- yi,c is the ground truth label (1 if sample i belongs to class c, else 0),
- yˆi,c is the predicted probability for class c.

Early stopping and learning rate reduction on plateau were used to optimize training convergence. Evaluation was performed using accuracy, precision, recall, and F1-score metrics for each class, revealing that the custom CNN consistently outperformed transfer learning mod-els.

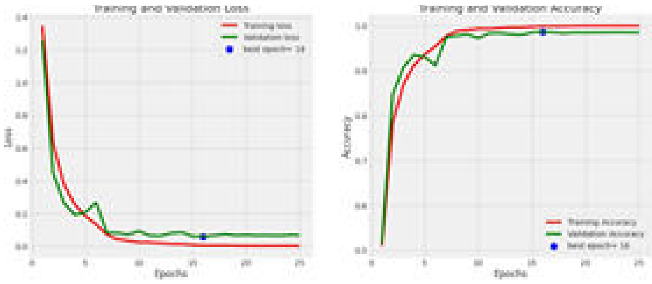

Fig. 2. Training and validation accuracy and loss curves for the custom CNN model.

D) System Integration

Upon achieving optimal model performance, the trained weights (.h5 files) were integrated into a Django-based web application. The platform allows users (patients or healthcare providers) to upload der-moscopic images, which undergo preprocessing before being passed through the CNN classifier. The system outputs a predicted class label along with class proba-bility scores. Additional features include:

- OTP-secured user registration and login,

- Admin dashboard for monitoring predictions and user activity,
- Profile management and feedback submission,
- Real-time prediction portal with visual results.

This integrated system not only demonstrates the technical feasibility of AI-assisted skin cancer detection but also ensures usability and accessibility for real-world applications.

IV. Experimental Results

The implementation of the multi-class skin cancer classification system integrates deep learning models with a Django-based web platform, enabling users to upload dermoscopic images for real-time diagnostic predictions. The system was designed with modularity in mind, where the deep learning models and web in- terface function cohesively while maintaining flexibility for updates or scaling.

A. Model Training and Export

The three models — Custom CNN, EfficientNet, and InceptionV3 — were developed and trained us-ing Keras with a TensorFlow backend. The training was performed on a system with GPU acceleration to ensure faster convergence. Each model was trained on preprocessed dermoscopic image datasets, where data augmentation techniques, such as rotation, scaling, and brightness adjustment, were applied to improve generalization.

Upon achieving satisfactory validation accuracy, the trained models were saved in the HDF5 (.h5) format, preserving both model weights and architectures. The custom CNN, which achieved the highest validation ac-curacy of 98.47%, was designated as the primary model for deployment, while EfficientNet and InceptionV3 were retained for comparative studies and optional use.

B. Backend Integration

The backend of the web application was developed using the Django framework, with Python as the pri-mary language. The workflow involves:

1) A user uploads a dermoscopic image through the web portal.
2) The uploaded image is preprocessed (resized, normalized, and augmented if necessary) using OpenCV and NumPy.
3) The preprocessed image is passed to the CNN classifier, which outputs a probability distribution across the seven lesion categories.
4) The predicted class label and associated confi-dence score are displayed to the user.

The backend also incorporates user management fea-tures, including OTP-secured registration and login, and an admin dashboard for monitoring predictions and managing user data. SQLite3 is used as the database for storing user information, prediction logs, and feedback data.

C. Frontend Development

The frontend interface was built using HTML, CSS, JavaScript, and Bootstrap to ensure a responsive and in-tuitive user experience. The user interface is structured into the following components:

- Prediction Portal: Allows users to upload dermo-scopic images and view classification results along with confidence percentages for each class.
- User Dashboard: Displays past predictions and provides tools for profile management.
- Feedback Section: Enables users to submit feed-back for system improvements, which is analyzed by the admin panel.
- Admin Panel: Allows administrators to manage users, view prediction statistics, and monitor feed-back analytics.

D. Prediction Workflow

Once the model predicts a class, the system provides the result in a user-friendly format, displaying:

- The predicted skin lesion type (e.g., melanoma, benign keratosis),
- The confidence score (as a percentage),
- Recommendations for consulting a dermatologist if the probability of a malignant lesion (e.g., melanoma or basal cell carcinoma) is significant.

This ensures that users can interpret the results clearly and take appropriate action based on the sys-tem's output.

E. System Deployment

The web application was deployed on a local de-velopment server for testing, ensuring that the predic-tion pipeline and user management systems operated seamlessly. The modular structure allows for deploy-ment on cloud platforms (e.g., AWS or Heroku) with minimal adjustments, enabling remote access for real-world users.

Through this integrated implementation, the project combines advanced AI-driven classification capabilities with a practical and accessible interface, making it suit-able for clinical assistance and patient-oriented usage.

V. RESULTS AND DISCUSSION

The multi-class skin cancer classification system was evaluated on a separate testing set after training the Custom CNN, EfficientNet, and InceptionV3 models. The primary objective of this phase was to assess the performance of each model, compare their results, and discuss the overall usability of the integrated web application.

A. Model Performance Analysis

The Custom CNN model demonstrated the highest accuracy of 98.47%, outperforming both EfficientNet (87.21%) and InceptionV3 (86.49%). The superior per-formance of the CNN can be attributed to its architec-ture, which was specifically designed and tuned for this dataset, capturing intricate texture and color features of dermoscopic images effectively.

The performance metrics, including precision, recall, and F1-score, were computed for each of the seven classes:

- Actinic Keratoses,
- Basal Cell Carcinoma,
- Benign Keratosis,
- Dermatofibroma,
- Melanoma,
- Melanocytic Nevi,
- Vascular Lesions.

The CNN model achieved consistently high scores across all metrics, particularly excelling in distinguish-ing between benign and malignant lesions. EfficientNet and InceptionV3 also performed reliably but demon-strated slightly reduced sensitivity for rare classes, such as Dermatofibroma and Vascular Lesions, due to their lower representation in the dataset.

B. Prediction Interface Evaluation

The Django-based web application successfully in-tegrated the trained CNN model, enabling real-time image-based predictions for users. The system's predic-tion portal allows users to upload dermoscopic images and instantly receive:

- The predicted lesion category,
- Confidence scores (as percentages) for all seven classes,
- Guidance to consult a dermatologist for lesions classified as malignant (e.g., Melanoma, Basal Cell Carcinoma).

The interface is designed to be user-friendly, en-suring that both patients and healthcare professionals can interact with the system without requiring technical expertise.

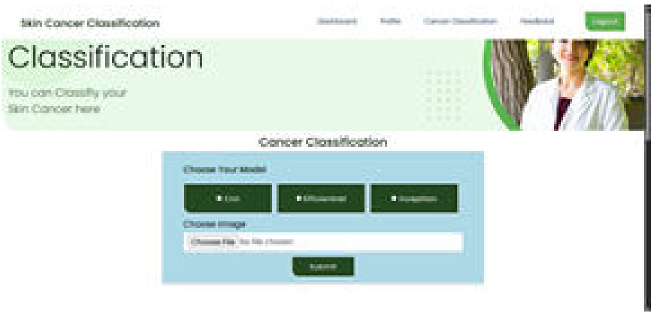

Fig. 3. Prediction portal interface where users can upload dermo-scopic images for classification.

C. Result Visualization

Upon processing the uploaded image, the system presents the final prediction results clearly. The results page displays:

- The most probable skin lesion type,
- Confidence percentages for all classes,
- Recommended actions, such as consulting a med-ical professional when a malignant lesion is sus-pected.

The visual layout ensures the information is acces-sible and easy to interpret, supporting decision-making for both patients and clinicians.

Fig. 4. Result page displaying predicted skin lesion type and class probabilities for the uploaded image.

D. Discussion

The project demonstrates that a well-designed CNN, trained specifically for a curated dermoscopic dataset, can surpass pre-trained transfer learning models in performance. While EfficientNet and InceptionV3 offer advantages in terms of scalability and computational efficiency, the custom CNN achieved better accuracy and class balance by leveraging dataset-specific tuning and augmentation strategies.

The integration of the CNN into a Django-based web platform ensures practical usability. The OTP-secured login, user dashboard, and admin management features add an additional layer of reliability and security, mak-ing the system viable for clinical assistance and patient-facing applications.

Overall, the system proves to be a reliable diagnostic support tool, capable of assisting in early skin cancer detection. Future enhancements may include integrating ensemble techniques (combining all three models), ex-panding the dataset to improve rare-class detection, and deploying the platform on a cloud-based infrastructure for broader accessibility.

VI. CONCLUSION AND FUTURE WORK

A. Conclusion

The project successfully developed a multi-class skin cancer classification system using deep learning models, focusing on the detection of seven types of dermoscopic skin lesions. Among the models implemented — a custom-built Convolutional Neural Network (CNN), Ef-ficientNet, and InceptionV3 — the CNN demonstrated the highest validation accuracy of 98.47%, outperform-ing the transfer learning-based architectures. This result highlights the effectiveness of a carefully designed and dataset-specific network in capturing intricate patterns and visual cues within dermoscopic images.

The Django-based web application serves as a prac-tical interface, allowing users, including patients and healthcare providers, to upload images and receive real-time diagnostic predictions with class-wise probabil-ities. With additional features such as OTP-secured registration, profile management, and feedback ana-lytics, the system bridges the gap between research-driven AI models and real-world usability. The project demonstrates both technical robustness and practical ap-plicability, offering a scalable AI-assisted tool that may support dermatologists in early skin cancer screening and diagnosis.

B. Future Work

Although the proposed system achieved promising performance, several improvements can further enhance its accuracy, reliability, and real-world usability.

1) Dataset Expansion: Future research can in-corporate larger and more diverse dermoscopic datasets, including additional rare lesion cate-gories, to improve model generalization and re-duce class imbalance.

2) Ensemble Modeling: Combining predictions from CNN, EfficientNet, and InceptionV3 through ensemble techniques such as weighted averaging or stacking may further improve classification accuracy and robustness.

3) Explainable AI Integration: The integration of interpretability techniques such as Grad-CAM or SHAP will help visualize the regions influencing the model's predictions, increasing transparency and trust among clinicians.

4) Cloud Deployment: Deploying the system on scalable cloud platforms will allow remote access and enable integration with telemedicine plat-forms for broader accessibility.

5) Mobile Application Development: A mobile-friendly application can be developed to enable quick image capture and real-time analysis, es-pecially beneficial in rural or resource-limited regions.

6) Clinical Validation: Future studies may involve collaboration with dermatologists to conduct real-world clinical trials and evaluate the system's diagnostic reliability and practical usefulness.

Overall, the continued advancement of deep learning, combined with larger datasets and improved explain-ability, can transform this system into a reliable clinical decision support tool for early skin cancer detection and prevention.

REFERENCES

[1] Ramachandra, B., & Surekha, T. P. (2024). Development and evaluation of a network intrusion detection system for DDoS attack detection using machine learning. Bulletin of Electrical Engineering and Informatics, 13(6), 4207-4213.

[2] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

[3] Han, S. S., Kim, M. S., Lim, W., Park, G. H., Park, I., & Chang, S. E. (2018). Classification of clinical images for benign and malignant cutaneous tumors using a deep learning algorithm. *Journal of Investigative Dermatology*, 138(7), 1529–1538.

[4] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A., Ciompi, F., Ghafoorian, M., ... & van der Laak, J. A. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88.

[5] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In Proceedings of the 36th International Conference on Machine Learning (ICML) (pp. 6105–6114).

[6] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2818–2826).

[7] Brinker, T. J., Hekler, A., Enk, A. H., Klode, J., Hauschild, A., Berking, C., ... & von Kalle, C. (2019). Deep neural networks are superior to dermatologists in melanoma image classification. European Journal of Cancer, 119, 11–17.

[8] Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., ... & Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of dia-betic retinopathy in retinal fundus photographs. JAMA, 316(22), 2402–2410.

[9] Ting, D. S., Cheung, C. Y., Lim, G., Tan, G. S., Quang, N. D., Gan, A., ... & Wong, T. Y. (2017). Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. JAMA, 318(22), 2211–2223.

[10] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10(1), 12598.

[11] Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. Nature Medicine, 28(1), 31–38.

[12] Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset: A large collection of multi-source dermato-scopic images of common pigmented skin lesions. Scientific Data, 5, 180161.

[13] Codella, N., Rotemberg, V., Tschandl, P., Celebi, M. E., Dusza, S., Gutman, D., ... & Halpern, A. (2019). Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the International Skin Imaging Collaboration (ISIC). arXiv preprint arXiv:1902.03368.

[14] Topol, E. J. (2019). High-performance medicine: the conver-gence of human and artificial intelligence. Nature Medicine, 25(1), 44–56.

[15] Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: review, opportunities and challenges. Briefings in Bioinformatics, 19(6), 1236–1246.

[16] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 618–626).

[17] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS) (pp. 4768–4777).

[18] Rogers, H. W., Weinstock, M. A., Feldman, S. R., & Coldiron, B. M. (2015). Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the US population, 2012. *JAMA Dermatology*, 151(10), 1081–1086.

[19] Dietterich, T. G. (2000). Ensemble methods in machine learning. In International Workshop on Multiple Classifier Systems (pp. 1–15). Springer.

[20] Song, Y., Zhang, Y. D., Yan, X., & Wang, S. (2020). Hybrid deep learning for skin lesion classification. Computers in Biol-ogy and Medicine, 127, 104067.

[21] Howard, A., Sandler, M., Chu, G., Chen, L. C., Chen, B., Tan, M., ... & Adam, H. (2019). Searching for MobileNetV3. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 1314–1324).

[22] Banbury, C., Reddi, V. J., Torelli, P., Holleman, J., Jeffries, N., Kiraly, C., ... & Mattina, M. (2021). MLPerf Tiny Bench-mark. In Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS) (pp. 254–263).

[23] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

[24] Han, S. S., Kim, M. S., Lim, W., Park, G. H., Park, I., & Chang, S. E. (2018). Classification of clinical images for benign and malignant cutaneous tumors using a deep learning algorithm. *Journal of Investigative Dermatology*, 138(7), 1529–1538.

[25] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A., Ciompi, F., Ghafoorian, M., ... & van der Laak, J. A. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88.

[26] Brinker, T. J., Hekler, A., Enk, A. H., Klode, J., Hauschild, A., Berking, C., ... & von Kalle, C. (2019). Deep neural networks are superior to dermatologists in melanoma image classification. European Journal of Cancer, 119, 11–17.

[27] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In Proceedings of the 36th International Conference on Machine Learning (ICML) (pp. 6105–6114).

[28] Song, Y., Zhang, Y. D., Yan, X., & Wang, S. (2020). Hybrid deep learning for skin lesion classification. Computers in Biol-ogy and Medicine, 127, 104067.

[29] Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset: A large collection of multi-source dermato-scopic images of common pigmented skin lesions. Scientific Data, 5, 180161.

[30] Codella, N., Rotemberg, V., Tschandl, P., Celebi, M. E., Dusza, S., Gutman, D., ... & Halpern, A. (2019). Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the International Skin Imaging Collaboration (ISIC). arXiv preprint arXiv:1902.03368.

[31] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2818–2826).

[32] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 618–626).

[33] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS) (pp. 4768–4777).

[34] Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10(1), 12598.

[35] Topol, E. J. (2019). High-performance medicine: the conver-gence of human and artificial intelligence. Nature Medicine, 25(1), 44–56.

[36] Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: review, opportunities and challenges. Briefings in Bioinformatics, 19(6), 1236–1246.