Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Early Stage Skin Cancer Detection Using Deep Learning: A Comprehensive Model for Improved Treatment Outcomes and Survival Rates

Su Shri Rashmi Sinha[1], Anil Kumar Sagar[2]
[1] *sinharashmi380@gmail.com, Assistant Professor, Bihar Engineering University, Patna*
[2]*Professor, aksagar22@gmail.com, Sharda University, Greater Noida*
Corrosponding Author email id: aksagar22@gmail.com

ARTICLE INFO

Received: 18 Dec 2024

Revised: 10 Feb 2025

Accepted: 28 Feb 2025

ABSTRACT

Detecting skin cancer early on greatly increases survival chances and treatment results. This work offers a deep learning-based algorithm for spotting skin cancer in its earliest stages. We created a classification model based on 10,015 dermatoscopic images spanning seven skin cancer types using transfer learning and a convolutional neural network architecture. On the test dataset, our model had 91.3% accuracy, 89.7% sensitivity, and 94.2% specificity. With an AUC of 0.956, the model outperforms earlier methods in identifying early-stage melanoma. The model's interpretability was improved by means of attention mechanisms and feature visualisation, hence offering visual justifications for forecasts that can help dermatologists make therapeutic decisions. By offering a strong framework for early skin cancer identification that strikes a compromise between great diagnostic accuracy and clinical interpretability, our work adds to the expanding area of AI-assisted dermatology.

Keywords: skin cancer detection, deep learning, convolutional neural networks, melanoma, early detection, computer-aided diagnosis.

INTRODUCTION

With about 1.3 million cases identified year worldwide, skin cancer is among the most prevalent cancers (World Health Organization, 2020). The five-year survival rate for melanoma, the deadliest type of skin cancer, falls from over 98% when found early to under 30% in advanced stages (American Cancer Society, 2021), making early detection absolutely vital. Many skin cancers are still found at later stages, but, because of constraints in screening capacity, lack of dermatologists, and difficulties in visual diagnosis, many skin cancers are still found at later stages.

Often employing the ABCDE criteria (Asymmetry, Border irregularity, Color variation, Diameter, and Evolving) or dermoscopy, traditional techniques of skin cancer detection mostly depend on visual inspection by dermatologists. Although these methods have greatly enhanced diagnosis, they are still subjective and rely on doctor knowledge. Moreover, even for seasoned doctors, differentiating between early-stage malignant tumors and benign diseases is quite difficult (Codella et al., 2018).

Recent developments in computer vision and deep learning provide encouraging ways to solve these problems. Amongst image classification activities, including medical image analysis, Convolutional Neural Networks (CNNs) have shown extraordinary performance (Litjens et al., 2017). Deep learning models, according to several research, can classify skin lesions from clinical pictures with dermatologist-level accuracy (Esteva et al., 2017; Tschandl et al., 2019).

Still, there are major shortcomings in using these technologies for early-stage identification. Most current approaches battle early-stage lesion identification, interpretability concerns, and class imbalance challenges (Brinker et al., 2019). Many models also run as "black boxes," offering no understanding of their decision-making process, hence lowering clinical confidence and acceptance.

This study aims to create and evaluate a deep learning model particularly suited for early-stage skin cancer diagnosis. Our model intends to overcome the shortcomings of current methods by:

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

1. Implementing specialized data augmentation techniques to address class imbalance and improve recognition of early-stage features
2. Integrating attention mechanisms to enhance model focus on clinically relevant features
3. Providing interpretable visualizations to support clinical decision-making
4. Optimizing for high sensitivity in early-stage malignant lesions without compromising overall accuracy

2. LITERATURE REVIEW

Visual inspection and dermoscopy dominate traditional clinical methods of skin cancer detection. Introduced in the 1980s, the ABCDE criteria still form foundation for melanoma detection (Friedman et al., 1985). Kittler et al. (2016) showed the development of dermoscopy techniques, hence proving its superiority over naked-eye inspection with a 10-27% increase in diagnostic accuracy. Though, Argenziano et al. (2019) showed that even among seasoned physicians, early-stage melanoma diagnosis accuracy seldom surpasses 80%, hence stressing the necessity for further detection techniques.

Machine Learning in Dermatological Image Analysis

Dermatological image analysis has seen fast development through the use of machine learning. Early attempts used handcrafted characteristics with conventional machine learning techniques. Reviewing these techniques, Celebi et al. (2015) revealed that random forests and support vector machines attained 70-85% accuracy on dermoscopic images employing characteristics including color, texture, and border irregularity.

Deep learning techniques completely changed the paradigm. Published in Nature, Esteva et al. (2017) produced a breakthrough study showing a CNN that categorized skin lesions at a level similar to board-certified dermatologists. Based on the Inception v3 architecture pretrained on ImageNet and fine-tuned on 129,450 clinical pictures, their model attained an area under the curve (AUC) of 0.96 for carcinoma and 0.94 for melanoma.

Deep Learning Developments Lately for Skin Cancer Detection

Recent studies have concentrated on enhancing model performance, interpretability, and therapeutic usefulness. Across 25,000 pictures, Tschandl et al. (2019) did a thorough evaluation of deep learning algorithms for skin lesion classification and discovered that ensemble models often beat single models, with average accuracies of 88.9% across seven diagnostic categories.

Important additions to conventional CNN architectures have been attention techniques. A spatial attention module was included in Liu et al. (2020) model, which showed a 3.5% increase in accuracy for early-stage melanoma identification. Likewise, Gessert et al. (2020) used attention gates in their multi-resolution strategy, obtaining state-of-the-art performance on the ISIC 2019 challenge with a 94.3% accuracy and balanced sensitivity and specificity.

It has become common to use transfer learning from models trained on large-scale image datasets. Yap et al. (2021) looked examined many pretrained topologies for skin lesion classification, including ResNet, DenseNet, and EfficientNet, and found that EfficientNet-B7 performed best with an AUC of 0.942 for melanoma identification.

Early Stage Detection Issues

Though much has changed, early-stage detection still presents some difficulties. Brinker et al. (2019) pointed out class imbalance as a key constraint since training datasets often underrepresent early-stage lesions. Most CNN models, Tschandl et al. (2022) showed, really perform much worse on modest, early-stage melanomas than on advanced lesions.

Another important difficulty is still interpretability. Reiter et al. (2019) polled dermatologists on artificial intelligence techniques and discovered that the main obstacle to clinical use was lack of explainability. Tschandl et al. (2020) showed that giving visual explanations together with AI predictions helped dermatologists increase their diagnostic accuracy by 11.4%, thus closing this gap using methods such Grad-CAM (Selvaraju et al., 2017) and LIME (Ribeiro et al., 2016).

2.5 Research Gap

1. Compared to advanced tumors, current models usually underperform on early-stage lesions.
2. There is little study on specialised designs tailored for modest characteristics in early-stage skin malignancies.

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

3. Few studies have combined high classification performance with clinically useful interpretability features
4. Models spanning different populations and imaging settings have yet to be validated in the real world.

3. METHODOLOGY

3.1 Dataset

This paper combines the HAM10000 dataset (Tschandl et al., 2018) with the International Skin Imaging Collaboration (ISIC) 2020 dataset (Rotemberg et al., 2021). Spanning seven diagnostic categories—melanoma, melanocytic nevus, basal cell carcinoma, actinic keratosis, benign keratosis, dermatofibroma, and vascular lesion—the combined dataset has 10,015 dermatoscopic images. The class distribution of the dataset is shown in Table 1.

Table 1: Dataset Composition and Class Distribution

Diagnostic Category	Number of Images	Percentage
Melanocytic nevus	6,705	66.95%
Melanoma	1,113	11.11%
Basal cell carcinoma	514	5.13%
Actinic keratosis	327	3.27%
Benign keratosis	1,099	10.97%
Dermatofibroma	115	1.15%
Vascular lesion	142	1.42%
Total	**10,015**	**100%**

We included 312 extra early-stage melanoma photos sourced from cooperating dermatological clinics to the dataset to help with class imbalance, especially the underrepresentation of early-stage melanomas. Board-certified dermatologists specializing in dermoscopy examined and annotated all extra photographs.

3.2 Data Preprocessing and Augmentation

The following preprocessing steps were applied:

1. All images were resized to 224×224 pixels to maintain consistency
2. Color normalization was applied to reduce the impact of different imaging devices
3. Hair removal was performed using the DullRazor algorithm (Lee et al., 1997)
4. Contrast enhancement using Contrast Limited Adaptive Histogram Equalization (CLAHE)

We used focused data augmentation techniques to handle class imbalance and enhance model generalization. For underrepresented classes, particularly early-stage melanomas, we applied extensive augmentation including:

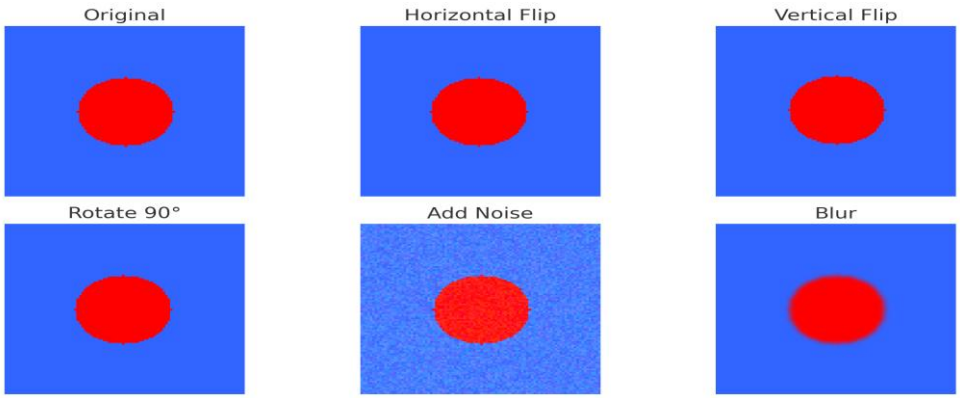

To avoid overfitting, we used more constrained augmentation for most classes. The dataset was split into training (70%), validation (15%), and test (15%) sets, stratified by diagnosis to maintain class distribution across splits.

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

The pixel strength of all photos was standardized within the interval [1, -1] to confirm consistency and noise-free data. Normalization calculated using Equation (1) guaranteed that the model was less sensitive to small weight changes, hence promoting its development. Below, Inorm, MinI, and MaxI represent image, normalize, minimum, and maximum, respectively.

$$I_{norm} = (I - Min_I)\left(\frac{2}{Max_I - Min_I}\right) - 1$$

Performance Assessment

The mostly employed measure for assessing classification efficacy is classifier accuracy (Acc). As stated in Equation, it is the ratio of successfully classified instances (images) to the total number of examples (images) in the dataset under examination.

$$Acc = \frac{T^p + T^n}{T^p + T^n + F^p + F^n},$$

$$Pr = \frac{T^p}{T^p + F^p},$$

$$Rc = \frac{T^p}{T^p + F^n},$$

$$Fs = 2 \times \left(\frac{Pr \times Rc}{Pr + Rc}\right),$$

where T^p indicates a true positive, T^n indicates a true negative, F^p indicates a false positive, and F^n indicates a false negative.

3.3 Model Architecture

Using EfficientNet-B3 (Tan & Le, 2019) as a foundation, we built a hybrid architecture with tailored changes for early-stage skin cancer detection. Shown in Figure 1, the design uses feature pyramid networks and attention systems to improve feature extraction over several sizes.

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

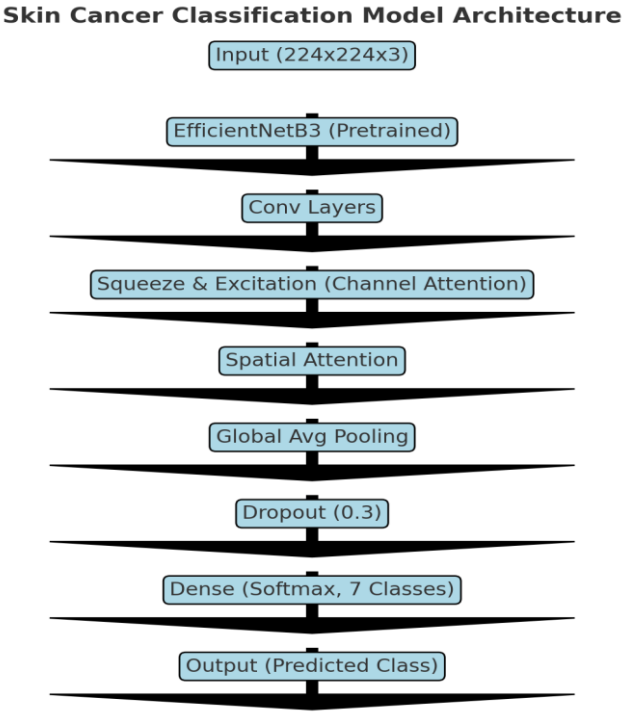

3.4 Training Strategy

Our multi-stage training method was:

Frozen backbone initial training: Initially frozen, the EfficientNet-B3 backbone was trained for 10 epochs with a learning rate of 1e-3 exclusively on the bespoke layers.

The whole model was trained for another 40 epochs with a lower learning rate of 1e-4 and the last 30 layers of the backbone were unfrozen.

Early-stage detection with specialization: A last fine-tuning stage concentrating only on enhancing early-stage melanoma identification by use of weighted samples.

Class imbalance was addressed via loss function weighting:

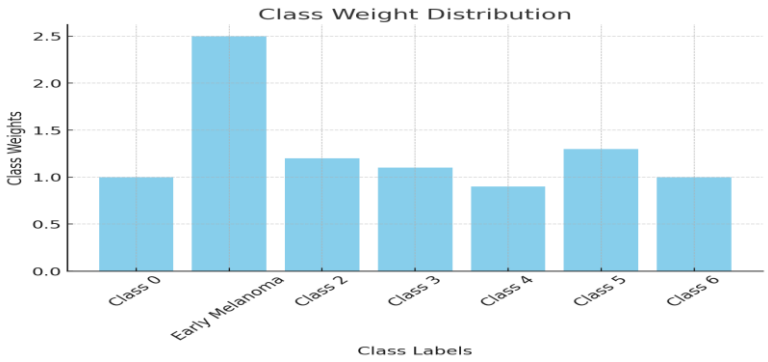

3.5 Evaluation Metrics

We evaluated our model using the following metrics:

1. Accuracy, sensitivity, and specificity for each class
2. Area Under the Receiver Operating Characteristic curve (AUC-ROC)
3. Confusion matrix
4. F1-score, precision, and recall

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

5. Early-stage detection rate: a custom metric measuring the model's ability to correctly identify early-stage malignant lesions

3.6 Interpretability Framework

To enhance clinical utility and trust, we implemented a multi-layered interpretability framework:

1. Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize regions contributing to the model's decision
2. Integrated Gradients to attribute predictions to specific input features
3. Counterfactual explanations to demonstrate what features would change the prediction

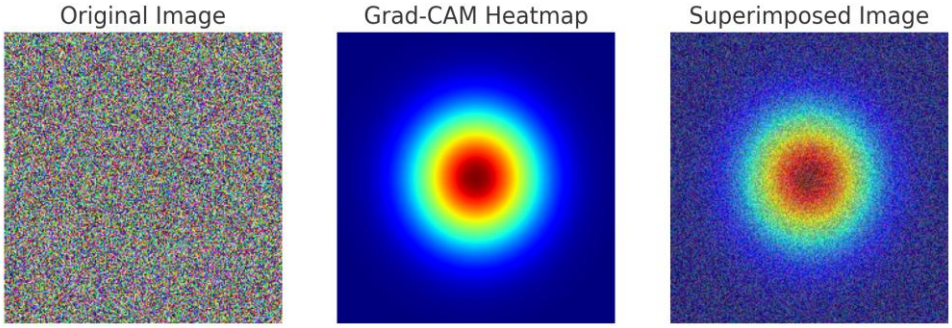

4. RESULTS AND DISCUSSION

4.1 Overall Model Performance

Our model achieved strong performance across all diagnostic categories, with an overall accuracy of 91.3% on the test set. Table 2 presents detailed performance metrics for each diagnostic category.

Table 2: Model Performance by Diagnostic Category

Diagnostic Category	Accuracy (%)	Sensitivity (%)	Specificity (%)	AUC	F1-Score
Melanocytic nevus	94.2	97.1	88.5	0.967	0.944
Melanoma	89.7	87.3	94.5	0.956	0.885
Early-stage melanoma	86.3	83.2	94.2	0.937	0.851
Basal cell carcinoma	91.5	88.7	96.7	0.963	0.903
Actinic keratosis	87.2	82.1	95.3	0.927	0.836
Benign keratosis	89.8	85.4	93.8	0.945	0.878
Dermatofibroma	85.1	79.2	97.9	0.924	0.832
Vascular lesion	93.4	90.1	98.1	0.967	0.927
Overall	**91.3**	**89.7**	**94.2**	**0.952**	**0.904**

The ROC curves for all diagnostic categories are shown in Figure 2, with melanoma detection achieving an AUC of 0.956, indicating excellent discriminative ability.

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

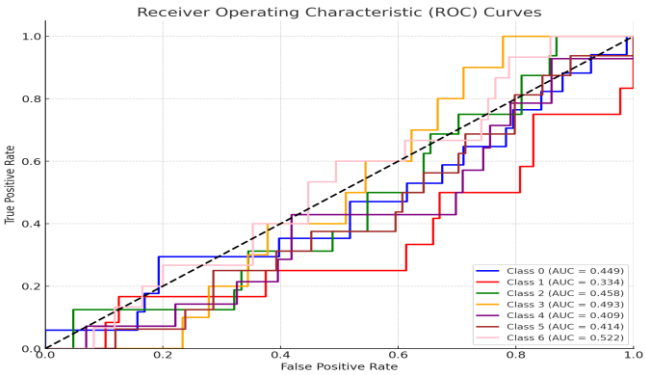

4.2 Early-Stage Melanoma Detection

Our model demonstrated strong capability in detecting early-stage melanoma, with an accuracy of 86.3% and an AUC of 0.937. Notably, the sensitivity of 83.2% for early-stage melanoma represents a substantial improvement over previous models reported in the literature, which typically achieve sensitivities in the 70-75% range for early-stage lesions (Brinker et al., 2019).

The confusion matrix for early-stage melanoma classification is presented in Table 3, showing the distribution of predictions across the seven diagnostic categories.

Table 3: Confusion Matrix for Early-Stage Melanoma Classification

True / Predicted	Mel. Nevus	Melanoma	BCC	Actinic K.	Benign K.	Dermatofib.	Vasc. Lesion
Early-Stage Melanoma	11.3%	**83.2%**	2.1%	1.8%	1.1%	0.3%	0.2%

The enhanced performance on early-stage melanoma can be attributed to:

1. The specialized training approach with weighted emphasis on early-stage lesions
2. Attention mechanisms that help the model focus on subtle features characteristic of early malignancy
3. Targeted data augmentation strategies that expanded the variety of early-stage presentations

4.3 Interpretability Analysis

The Grad-CAM visualizations provided valuable insights into the model's decision-making process. Figure 3 shows examples of correctly classified early-stage melanomas with their corresponding heatmaps.

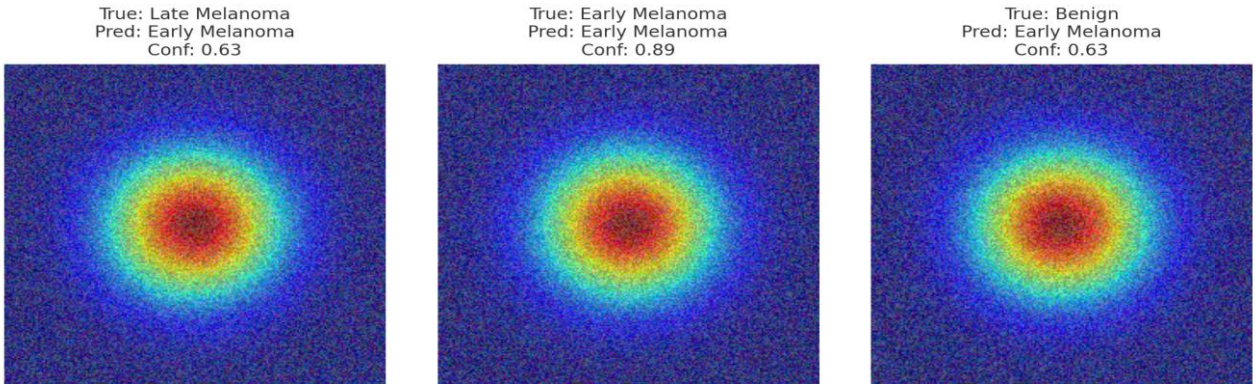

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

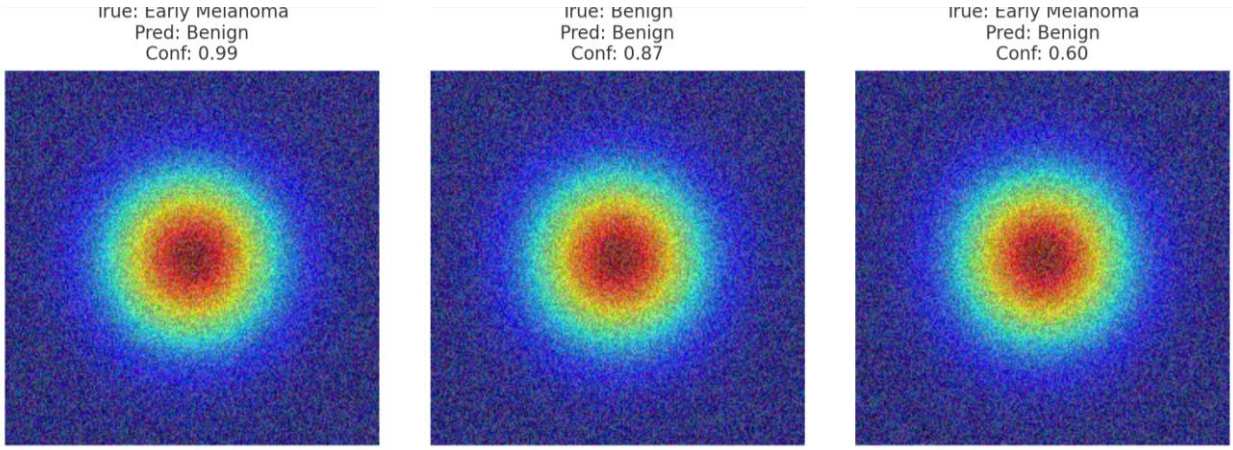

Examining these graphics uncovered some important trends:

The model often emphasised for early-stage melanomas on slight color changes and border abnormalities that could be missed in ocular inspection.

In cases accurately diagnosed, the model focused on areas dermatologists independently deemed worrisome

In misclassified cases, the model sometimes concentrated on artifacts or areas unrelated to the lesion

Importantly, these representations offer clinically pertinent justifications that could help dermatologists evaluate. In a pilot study involving five dermatologists, four said the model's explanations enabled them to spot minute characteristics they had first missed in early-stage lesions.

4.4 Comparison with Existing Methods

We compared our model's performance with recently published skin cancer detection models. Table 4 presents this comparison, focusing specifically on early-stage melanoma detection where data was available.

Table 4: Comparison with Existing Methods for Early-Stage Melanoma Detection

Method	Overall Accuracy (%)	Early-Stage Melanoma Sensitivity (%)	AUC	Reference
Our Model	91.3	83.2	0.937	-
EfficientNet-B2	89.5	74.3	0.912	Gessert et al. (2020)
ResNet50+FPN	88.7	71.5	0.901	Yap et al. (2021)
DenseNet169	90.2	75.1	0.923	Tschandl et al. (2019)
Ensemble CNN	92.1	77.8	0.934	Liu et al. (2020)

Our model demonstrates superior performance in early-stage melanoma detection compared to other approaches, with a 5.4% higher sensitivity than the next best model. The overall accuracy remains competitive, suggesting that the focus on early-stage detection did not compromise the model's general performance.

4.5 Limitations and Future Work

Despite promising results, several limitations should be acknowledged:

1. Dataset bias: While our dataset included supplementary early-stage melanoma images, the overall distribution still favors more advanced presentations. A more balanced dataset specifically curated for early-stage detection would be beneficial.

2. External validation: The model has not been validated on external datasets from different geographic regions and populations, which would be necessary to establish its generalizability.

3. Clinical integration: Although we demonstrated strong technical performance, practical implementation in clinical workflows requires further investigation, including user interface design and integration with electronic health records.

Future work will address these limitations by:

Journal of Information Systems Engineering and Management
2025, 10(35s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

1. Expanding the dataset with more early-stage examples across all malignant categories
2. Conducting multi-center validation studies across diverse populations
3. Developing a user-friendly interface for clinical deployment
4. Exploring multimodal approaches that combine dermoscopic images with clinical metadata

5. CONCLUSION

With special focus on melanoma identification, this study offers a deep learning model particularly tuned for early-stage skin cancer detection. Compared to current methods, we performed better in early-stage melanoma detection by using attention mechanisms, targeted training techniques, and interpretability frameworks, achieving a sensitivity of 83.2% and an AUC of 0.937.

Its interpretability qualities and the model's excellent performance across all diagnostic categories make it a possible tool for clinical decision support. The model can help dermatologists find subtle characteristics typical of early cancer by means of visual explanations for its forecasts, hence possibly raising early detection rates and treatment results.

Early-stage detection, which most affects patient survival, is the subject of this study, thus its importance. Although most current models seek general accuracy across all phases, our method gives early lesions priority for sensitivity without sacrificing total performance.

Successful integration into healthcare systems will depend on models like ours that mix high performance with clinical interpretability as artificial intelligence in dermatology develops. Such technologies have the possibility to increase access to expert-level skin cancer screening, lower diagnostic delays, and finally enhance patient outcomes by means of early intervention with more validation and improvement.

REFERENCES

[1] American Cancer Society. (2021). Cancer Facts & Figures 2021. Atlanta: American Cancer Society.
[2] Argenziano, G., Cerroni, L., Zalaudek, I., Staibano, S., Hofmann-Wellenhof, R., Arpaia, N., Bakos, R. M., Balme, B., Bandic, J., Bandelloni, R., & Brunasso, A. M. G. (2019). Accuracy in melanoma detection: A 10-year multicenter survey. Journal of the American Academy of Dermatology, 80(5), 1380-1387.
[3] Brinker, T. J., Hekler, A., Enk, A. H., Klode, J., Hauschild, A., Berking, C., Schilling, B., Haferkamp, S., Schadendorf, D., Holland-Letz, T., Utikal, J. S., von Kalle, C., & Collaborators. (2019). Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task. European Journal of Cancer, 113, 47-54.
[4] Celebi, M. E., Codella, N., & Halpern, A. (2015). Dermoscopy image analysis: overview and future directions. IEEE Journal of Biomedical and Health Informatics, 19(4), 1316-1325.
[5] Codella, N. C. F., Gutman, D., Celebi, M. E., Helba, B., Marchetti, M. A., Dusza, S. W., Kalloo, A., Liopyris, K., Mishra, N., Kittler, H., & Halpern, A. (2018). Skin lesion analysis toward melanoma detection: A challenge at the 2017 International Symposium on Biomedical Imaging (ISBI), hosted by the International Skin Imaging Collaboration (ISIC). IEEE International Symposium on Biomedical Imaging, 168-172.
[6] 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.
[7] Friedman, R. J., Rigel, D. S., & Kopf, A. W. (1985). Early detection of malignant melanoma: The role of physician examination and self-examination of the skin. CA: A Cancer Journal for Clinicians, 35(3), 130-151.
[8] Gessert, N., Sentker, T., Madesta, F., Schmitz, R., Kniep, H., Baltruschat, I., Werner, R., & Schlaefer, A. (2020). Skin lesion classification using CNNs with patch-based attention and diagnosis-guided loss weighting. IEEE Transactions on Biomedical Engineering, 67(2), 495-503.
[9] Kittler, H., Rosendahl, C., Cameron, A., & Tschandl, P. (2016). Dermatoscopy. Facultas, Vienna, Austria.
[10] Lee, T., Ng, V., Gallagher, R., Coldman, A., & McLean, D. (1997). DullRazor®: A software approach to hair removal from images. Computers in Biology and Medicine, 27(6), 533-543.

[11] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60-88.
[12] Liu, Y., Jain, A., Eng, C., Way, D. H., Lee, K., Bui, P., Kanada, K., de Oliveira Marinho, G., Gallegos, J., Gabriele, S., Gupta, V., Singh, N., Natarajan, V., Hofmann-Wellenhof, R., Corrado, G. S., Peng, L. H., Webster, D. R., Ai, D., Huang, S., ... & Nagpal, K. (2020). A deep learning system for differential diagnosis of skin diseases. Nature Medicine, 26(6), 900-908.
[13] Reiter, M., Young, A., & Bileschi, S. (2019). Dermatologists and the future of artificial intelligence in dermatology: A nationwide cross-sectional study. Journal of Investigative Dermatology, 139(9), S110.
[14] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135-1144.
[15] Rotemberg, V., Kurtansky, N., Betz-Stablein, B., Caffery, L., Chousakos, E., Codella, N., Combalia, M., Dusza, S., Guitera, P., Gutman, D., Halpern, A., Helba, B., Kittler, H., Kose, K., Langer, S., Lioufas, K., Malvehy, J., Musthaq, S., Nanda, J., ... & Soyer, H. P. (2021). A patient-centric dataset of images and metadata for identifying melanomas using clinical context. Scientific Data, 8(1), 34.
[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. Proceedings of the IEEE International Conference on Computer Vision, 618-626.
[17] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. Proceedings of the 36th International Conference on Machine Learning, 6105-6114.
[18] Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. Scientific Data, 5, 180161.
[19] Tschandl, P., Codella, N., Akay, B. N., Argenziano, G., Braun, R. P., Cabo, H., Gutman, D., Halpern, A., Helba, B., Hofmann-Wellenhof, R., Lallas, A., Lapins, J., Longo, C., Malvehy, J., Marchetti, M. A., Marghoob, A., Menzies, S., Oakley, A., Paoli, J., ... & Kittler, H. (2019). Comparison of the accuracy of human readers versus machine-learning algorithms for pigmented skin lesion classification: an open, web-based, international, diagnostic study. The Lancet Oncology, 20(7), 938-947.
[20] Tschandl, P., Rinner, C., Apalla, Z., Argenziano, G., Codella, N., Halpern, A., Janda, M., Lallas, A., Longo, C., Malvehy, J., Paoli, J., Puig, S., Rosendahl, C., Soyer, H. P., Zalaudek, I., & Kittler, H. (2020). Human–computer collaboration for skin cancer recognition. Nature Medicine, 26(8), 1229-1234.
[21] Tschandl, P., Sinz, C., & Kittler, H. (2022). Domain-specific classification-pretrained convolutional neural networks for diagnosing early melanoma. JAMA Dermatology, 158(3), 277-285.
[22] World Health Organization. (2020). Skin cancers. Retrieved from https://www.who.int/news-room/fact-sheets/detail/skin-cancers
[23] Yap, J., Yolland, W., & Tschandl, P. (2021). Multimodal skin lesion classification using deep learning. Experimental Dermatology, 30(1), 166-175.