SKIN CANCER DETECTION USING CONVOLUTIONAL NEURAL NETWORK (CNN)

Nishmitha R, Spoorthi G K, T Pavana, Bhuvan N,
Mrs Maseeha Banu (Assistant Professor)

Department of Computer Science and Engineering (Artificial Intelligence and Machine Learning), Malnad College of Engineering, Hassan, Karnataka, India

Corresponding author: Spoorthi G K (76761 33384)

Subject: Artificial Intelligence, Machine Learning and Medical Image Analysis

ABSTRACT

Skin cancer detection is an important application of artificial intelligence in medical image analysis. Identifying suspicious skin lesions at an early stage may help people seek timely medical attention. However, distinguishing between lesion categories can be difficult because some lesions share similar visual characteristics. This paper studies a Convolutional Neural Network (CNN)-based approach for classifying skin-lesion images. The workflow includes collecting labelled images, preprocessing them, learning visual features, training a classifier, and evaluating its predictions. Transfer learning using pretrained image-recognition architectures may improve learning efficiency when the available dataset is limited. Accuracy, precision, recall, and F1-score are suitable evaluation measures, with category-wise analysis particularly important when the dataset is imbalanced. Results quoted from earlier studies are presented as literature findings and are not claimed as experimental results from a newly implemented system. The proposed system is intended for preliminary image analysis only and cannot replace assessment by a qualified medical professional.

Keywords—Skin cancer, Convolutional Neural Network, deep learning, medical image analysis, image classification, transfer learning.

1. INTRODUCTION

Skin cancer is associated with the abnormal growth of skin cells. The visual appearance of a lesion can vary according to its type, location, and other characteristics. Some benign and malignant lesions may look similar, making reliable assessment challenging. Computer-aided analysis of skin images is therefore an active area of research. Artificial intelligence techniques learn patterns from examples and apply those patterns to new data. Deep learning methods are especially useful for image analysis because they can learn increasingly complex features from pixel data. A Convolutional Neural Network uses convolutional filters, activation functions, pooling operations, and classification layers to identify image patterns such as edges, textures, shapes, and higher-level structures. For skin-lesion classification, a CNN is trained on images with known labels and then predicts a category for an unseen image. Transfer learning adapts a model previously trained on a large image collection to a medical image task. The usefulness of such a model depends on dataset quality, the training procedure, and independent evaluation.

2. PROBLEM STATEMENT

Distinguishing skin-lesion categories from photographs is challenging because different lesions can have similar visual features and image quality may vary across devices and environments. This paper examines a CNN-based classification workflow that learns from labelled images and predicts a lesion category. The intended goal is preliminary image classification support, not autonomous medical diagnosis.

3. OBJECTIVES

- To study the application of deep learning to skin-lesion image analysis.
- To describe preprocessing steps required for labelled skin images.
- To outline a CNN-based classification workflow.
- To examine transfer learning with pretrained image-recognition architectures.
- To identify suitable classification metrics and evaluation considerations.
- To discuss the advantages, limitations, and possible future improvements of the approach.

4. LITERATURE SURVEY

4.1 CNN-Based Skin Cancer Classification

The source material supplied for this paper describes a study by Hasan et al. (2019) that investigated CNN-based classification of skin lesions into benign and malignant categories using images from the International Skin Imaging Collaboration (ISIC) archive. The supplied text reports test accuracy of 89.5% and training accuracy of 93.7%, along with evaluation using precision, recall, and F1-score. These values are specific to the experiment described by that study and should be checked against the original publication before formal submission.

4.2 Transfer Learning for Skin-Lesion Analysis

The supplied source also discusses work by Nahata and Singh on deep-learning approaches for skin cancer detection and diagnosis. Architectures listed include ResNet50, MobileNet, VGG16, InceptionV3, and InceptionResNet. The source reports accuracies of 85%, 85%, 87%, 90%, and 91%, respectively. In those reported experiments, InceptionResNet achieved the highest accuracy among the listed models. These figures are literature results, not results obtained by the present paper's author. Transfer learning can reduce the need to train a deep network from the beginning, since early layers of a pretrained model may already represent useful image features. However, comparisons are meaningful only when dataset splits, preprocessing, class distributions, and evaluation methods are considered.

5. DATASET DESCRIPTION

A skin-lesion classifier requires images with reliable category labels. The supplied source refers to the ISIC archive and describes a dataset assembled from ISIC 2018 and ISIC 2019 with seven shared categories. The categories listed are melanoma, melanocytic nevus, basal cell carcinoma, actinic keratosis, benign keratosis, dermatofibroma, and vascular lesion. The exact dataset version, image count, class distribution, licence, and access date should be recorded from the dataset actually used in an implementation. Before training, images and labels should be checked for missing entries, duplicate images, inconsistent formats, and incorrect labels. Data should be divided into training, validation, and test sets. The test set must remain independent of model training and model selection; where possible, duplicate or closely related images should not be distributed across different subsets.

6. PROPOSED METHODOLOGY

6.1 Image Collection

Collect labelled skin-lesion images from a suitable dataset such as the ISIC archive. Each image should be associated with a valid

target label. The collection should represent the categories that the system is intended to classify.

6.2 Image Preprocessing

Preprocessing prepares image data for consistent model input. Images are resized to the dimensions required by the selected architecture, pixel values are normalized as required by the model, and labels are encoded into a format supported by the training framework. Unreadable images and records with missing labels should be handled explicitly.

6.3 Data Augmentation

Controlled transformations such as small rotations, shifts, and zoom operations may be applied to training images to increase sample variety and reduce overfitting. Transformations should preserve medically relevant visual characteristics. Augmentation should be applied only to training data, not to the validation or test sets.

6.4 CNN Architecture

A typical CNN pipeline begins with a skin-lesion image, followed by convolutional layers that learn visual features. Activation functions introduce non-linearity, while pooling layers reduce the spatial dimensions of feature maps. The resulting representation is passed to a classification layer that generates scores for the target categories. The predicted category is selected from the model output.

Input image → Preprocessing → Convolution and activation → Pooling → Classification layer → Predicted category

Figure 1. General workflow of CNN-based skin-lesion classification.

6.5 Transfer Learning

A pretrained image-recognition network can be adapted by replacing or adding a final classification layer for the selected lesion categories. Some pretrained layers may initially be frozen, while the classification head is trained. Further fine-tuning can be considered if justified by validation performance and available computing resources. ResNet50, MobileNet, VGG16, InceptionV3, and InceptionResNet are candidate architectures discussed in the supplied literature.

6.6 Model Training

Training is performed in batches over multiple epochs. A loss function measures the difference between model predictions and the known labels, and an optimizer updates the model parameters. Validation results can guide model selection and help identify overfitting. The final selected model should be evaluated once on the held-out test set.

7. SYSTEM DESIGN

7.1 Context Model

The user submits a skin-lesion image to the application. The input module checks the file, the preprocessing module prepares the image, and the trained CNN produces class scores. The output module displays the predicted category and a clear limitation statement that the output is not a confirmed diagnosis.

7.2 System Modules

- Input module: accepts a skin-lesion image and checks file validity.
- Preprocessing module: resizes and normalizes the image.
- Classification module: runs inference using the trained CNN.
- Output module: displays the predicted category and relevant cautionary information.

7.3 Behavioural Model

The system begins when the user uploads an image. After validating the file, it preprocesses the image and passes it to the trained model. The model returns scores for the supported classes, and the interface presents the predicted category. If the image cannot be processed, the system should display an error message instead of an unreliable prediction.

8. PERFORMANCE EVALUATION

Accuracy, precision, recall, and F1-score can be used to evaluate classification performance. For binary classification, TP denotes true positives, TN true negatives, FP false positives, and FN false negatives.

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN)$$
$$\text{Precision} = TP / (TP + FP)$$
$$\text{Recall} = TP / (TP + FN)$$
$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

For multiclass datasets, the evaluation should explain whether precision, recall, and F1-score are macro-, micro-, or weighted-averaged. A confusion matrix and category-wise recall are also useful, particularly where missing a malignant lesion would be a serious concern. Metrics alone do not establish clinical safety.

9. RESULTS AND DISCUSSION

The following values are reproduced from the literature summary supplied for this paper. They are not results from a new experiment.

	Reported accuracy
	85%
	85%
	87%
	90%
	91%

Table 1. Accuracy values reported in the supplied literature summary.

Within the reported comparison, InceptionResNet has the highest accuracy, followed by InceptionV3. The difference does not prove that the same ranking will hold for another dataset or implementation. Differences in class imbalance, image quality, preprocessing, data splitting, and evaluation methods may change the results. A new implementation must report its own test-set metrics and must not reuse literature values as project results.

10. ADVANTAGES

- CNNs can learn image features automatically from labelled examples.
- A trained model can apply a consistent classification procedure to new images.
- Transfer learning can reduce training time and data requirements compared with training from scratch.
- Multiple metrics can be used to examine different aspects of model performance.
- A user interface can make preliminary image analysis easier to access.

11. LIMITATIONS

- Prediction quality depends on dataset quality, diversity, and label reliability.
- Imbalanced classes may lead to poor performance on less common lesion categories.
- Images from unfamiliar cameras, skin tones, or environments may produce unreliable predictions.
- High overall accuracy does not guarantee that every lesion is classified correctly.
- The model output cannot independently confirm or rule out a medical diagnosis.
- Clinical use would require appropriate independent validation and qualified medical oversight.

12. FUTURE SCOPE

Future work can expand dataset diversity, improve representation of underrepresented categories, and compare architectures under consistent experimental conditions. Explainability methods may help indicate which image regions influenced a prediction,

although such visual explanations do not prove that a model is correct. A web or mobile interface could support image submission and display outputs with clear guidance about limitations. Further work should include independent testing, evaluation across image sources and skin tones, and careful analysis of false-negative predictions.

Author note: Replace the author, affiliation, email, and correspondence placeholders. Verify all cited publication details and literature figures against the original sources before submitting the paper.

13. CONCLUSION

This paper presents a CNN-based workflow for skin-lesion image classification, covering labelled data collection, preprocessing, feature learning, model training, and evaluation. Transfer learning offers a practical option for adapting established image-recognition architectures to this task. The literature summary supplied with the paper reports promising results for several pretrained models, with InceptionResNet achieving the highest reported accuracy among the listed architectures. These figures must be verified against the original publications and should not be presented as results of a new implementation. CNN-based analysis may provide preliminary information, but it is not a substitute for assessment by a qualified medical professional.

REFERENCES

1. Hasan et al. (2019). “Skin Cancer Detection Using Convolutional Neural Network.” The supplied source does not contain the complete author list or publication details; verify the original paper before submission.

2. Nahata, H., and Singh, S. P. “Deep Learning Solutions for Skin Cancer Detection and Diagnosis.” The supplied source does not specify the complete publication year, venue, page range, or DOI; verify these details against the original publication.

3. International Skin Imaging Collaboration (ISIC). Skin Imaging Archive and associated challenge datasets. Record the exact dataset release, URL, licence, and access date used in the implementation.