Skin Disease Detection Using CNN and Yolo

Suchethana H C[1],Kavana Nuchin K N[2], Rashika Y Salunkhe[3], Chandana C[4], Swapnitha S[5]

[1]Assistant Professor. Information Science and Engineering. JNN college of Engineering (VTU Affiliation), India

[2]4JN22IS074 Information Science and Engineering. JNN college of Engineering (VTU Affiliation), India

[3]4JN22IS118 Information Science and Engineering. JNN college of Engineering (VTU Affiliation), India

[4]4JN23IS405 Information Science and Engineering. JNN college of Engineering (VTU Affiliation), India

[5]4JN23IS415 Information Science and Engineering. JNN college of Engineering (VTU Affiliation), India

DOI: https://doi.org/10.51244/IJRSI.2025.1213CS0021

Received: 29 December 2025; Accepted: 04 January 2026; Published: 15 January 2026

ABSTRACT

Skin diseases are a significant global health issue. They can be as mild as common rashes or as serious as deadly cancers like melanoma. Getting the right diagnosis quickly is key it helps save lives and lowers healthcare costs. Traditionally, doctors examine the skin with their eyes and rely on their experience, but this approach can vary from doctor to doctor, take a lot of time, and isn't always available everywhere. Now, thanks to breakthroughs in Artificial Intelligence (AI) and especially Deep Learning (DL), computers can help doctors analyze medical images and make more accurate diagnoses automatically. In this work, we introduce a smart system for detecting skin diseases. It uses two powerful AI tools: Convolutional Neural Networks (CNN) to sort images into different disease types, and You Only Look Once (YOLO) to quickly find and highlight problem spots on the skin. We trained and tested this system with a large set of skin images called the HAM10000 dataset, which includes many kinds of skin conditions. The CNN learns to recognize unique patterns in the images, while YOLO helps pinpoint exactly where the skin problem is— and it does this in real time. Our results show that this combined approach is very accurate at classifying diseases and spotting the affected areas, making it a useful tool for doctors and healthcare workers in real-world settings.

Keywords: Skin Disease Detection, Deep Learning, CNN, YOLO, HAM10000, Medical Image Processing, Artificial Intelligence.

INTRODUCTION

The skin is really important because it is the part of our body. It helps keep us safe from things outside. A lot of people over the world have skin problems. These problems can be things like melanoma, basal cell carcinoma, eczema, psoriasis and fungal infections. Skin cancer is especially bad because it can get worse fast and a lot of people can die from it if they do not find out they have skin cancer early. Skin cancer is a deal and we need to be careful, about our skin. Diagnosing skin diseases the usual way involves a doctor looking at the skin using a tool called dermoscopy and doing a biopsy. Doctors who specialize in skin problems do these things. These methods do work.. They have some problems. The doctors have to make judgments based on what they see.. It can be hard to find a specialist. Also it can be very expensive to get a diagnosis. In some places like areas that are still developing or rural areas it is hard to find a skin doctor. This means that people often have to wait a time to find out what is wrong, with their skin. Diagnosing skin diseases can be a problem. Skin disease diagnosis is not always easy to get.

Deep Learning techniques are really good at figuring out what is in a picture and finding objects. This is because Convolutional Neural Networks or Deep Learning techniques can look at pictures and automatically find the parts so we do not have to tell them what to look for.. When we use Deep Learning techniques, like

Convolutional Neural Networks to classify pictures it does not tell us where the problem is in the picture, which is important for doctors to understand what is going on. Deep Learning techniques, like Convolutional Neural Networks are good at classifying pictures. They need help finding the exact location of the problem like a lesion, in the picture.

To overcome this limitation, this research integrates YOLO, a real-time object detection algorithm, with CNN-based classification. The proposed framework provides both disease classification and lesion localization, improving accuracy, interpretability, and usability in medical settings form and the form should accompany your final submission.

Relevant Work

People have done a lot of work on finding skin diseases using computers. They used computer programs to help them. At first they used computer programs like Support Vector Machines and k-Nearest Neighbors and Decision Trees. These programs looked at things, like the texture and color and shape of the skin to try to figure out what was going on. Skin disease detection using machine learning and deep learning techniques is what they were trying to do. They used machine learning and deep learning techniques to detect skin diseases. The field of learning has really taken off. We are seeing a lot of progress with models like VGGNet, ResNet, DenseNet and Efficient Net. These deep learning models, the ones that use CNN are being used a lot. They are really good at classifying things. For example., VGGNet, ResNet, DenseNet and EfficientNet have done well on some big tests like ISIC and HAM10000. We are talking about improvements, in how accurately these models can classify things.

People have been doing studies to see how they can make models easier to understand. They have been looking at things like lesion segmentation and detection. Some object detection algorithms like Faster R-CNN and SSD and YOLO have been used to find lesions. YOLO is really popular because it can do things in time and it is very good at detecting things. Lesion detection, with YOLO is getting a lot of attention because of this.

However, many existing systems focus either on classification or detection, not both. This paper addresses this research gap by proposing a unified CNN–YOLO framework for comprehensive skin disease analysis. If you are using Word, use either the Microsoft Equation Editor or the Math Type add-on (http://www.mathtype.com) for equations in your paper (Insert | Object | Create New | Microsoft Equation or MathType Equation). ―Float over text‖ should not be selected.

The Framework Model

The proposed framework model presents the overall architecture of the skin disease detection system using YOLO and Convolutional Neural Networks (CNN). The system processes dermoscopic images through preprocessing, lesion detection, and disease classification stages to achieve accurate and reliable skin disease diagnosis. The complete workflow of the proposed framework is illustrated in Fig. 1.1.

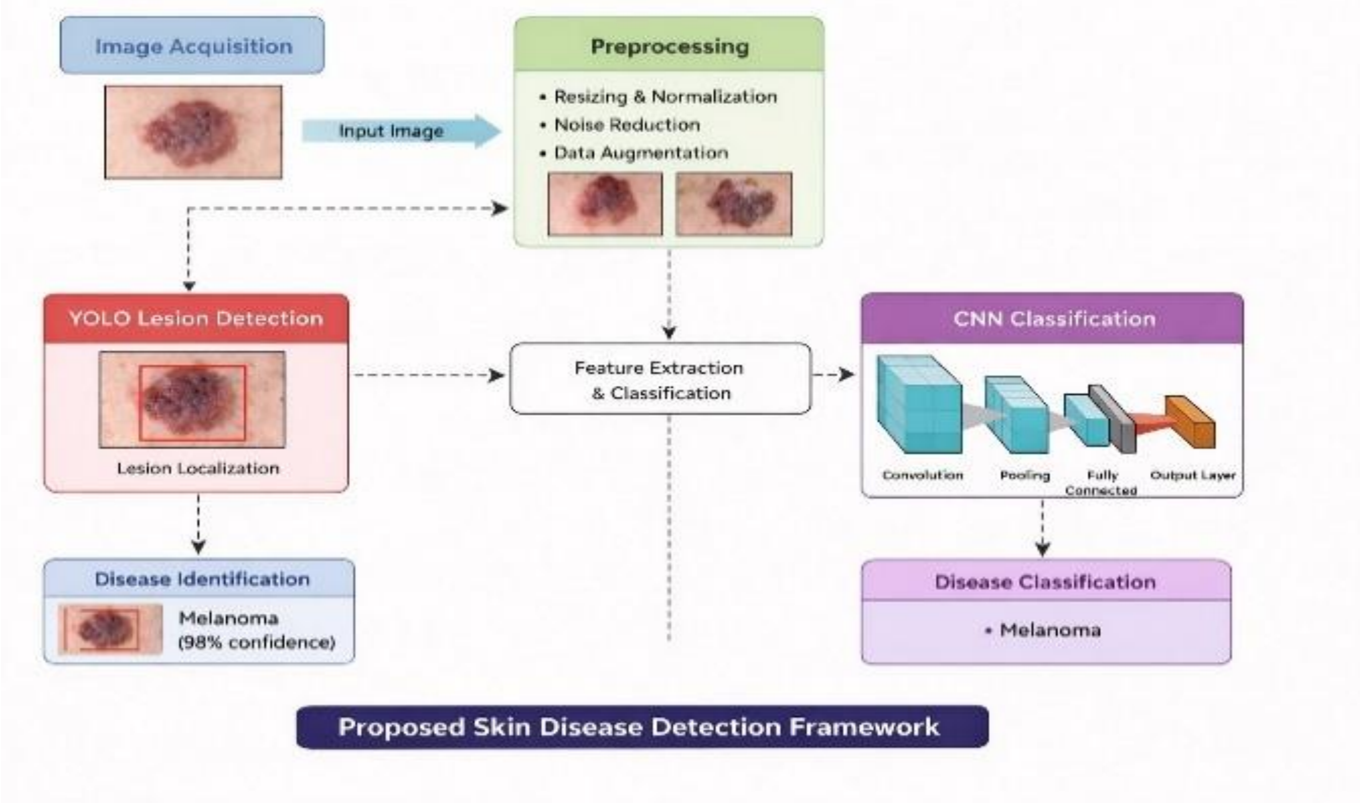

Figure 1.1: Proposed CNN–YOLO based framework for skin disease detection

Description: The flowchart shows how the skin disease detection system works. It uses YOLO. Convolutional Neural Network models to find lesions and classify diseases accurately.

The skin disease detection system starts with getting skin images from the HAM10000 dataset. This dataset has a lot of skin disease images that were taken in a setting. These skin images are then cleaned up. This means they are made smaller normalized and cleared of noise. The system also adds images to the dataset to make sure all the images are similar and good quality. This helps the skin disease detection system work better. The YOLO model looks at the images after they have been cleaned up. The YOLO model finds the spots on the skin image. The YOLO model figures out what disease it is and puts a box around the spot. The YOLO model also says how sure it is that it found the spot.

The bad spot that the YOLO model found is then sent to the CNN model. The CNN model takes a look at the bad spot. The CNN model uses layers to get more information, from the bad spot. The CNN model uses these layers to look closely at the bad spot. Based on the extracted features, the CNN model classifies the skin disease into the appropriate category. Finally, the outputs from both YOLO and CNN are combined to provide accurate disease identification with visual localization and reliable classification, making the proposed system effective for automated skin disease detection and clinical decision support.

Skin Disease Datasets

The HAM10000 dataset, which is Human Against Machine with 10,000 training images is something that a lot of people use for research on skin lesions and finding skin diseases automatically. This dataset has than 10,000 really good pictures of skin lesions that show seven different kinds of skin problems like melanoma, benign nevi, keratoses, basal cell carcinoma, actinic keratoses, vascular lesions and dermatofibroma. The HAM10000 dataset is really helpful because each picture in the HAM10000 dataset comes with information, such as how old the patient is, if the patient is a man or a woman and where the skin problem is, on the body. This extra information can make the models that use the HAM10000 dataset at finding skin problems and can help evaluate how well they work with the HAM10000 dataset.

This dataset has a lot of differences in lesion size, shape, color and texture. It shows the variation you see in skin conditions. The images are not evenly spread across classes. So people use techniques, like rotating, flipping, scaling and changing the color to make the model work better and be more accurate. These techniques help the model with its performance and ability to work with kinds of data.

HAM10000 has become a benchmark resource in dermatological research and machine learning studies, helping researchers develop and evaluate deep learning models for accurate skin lesion classification, early detection of melanoma, and other skin disease diagnostic applications.

LITERATURE REVIEW

The old ways of figuring out skin diseases are not very good at finding the answer. They use tools and simple computer programs.. These old ways do not work well and are not very good at understanding new things.

Recently people have found out that using computer programs called deep learning especially ones that look at pictures can really help tell what is wrong with the skin. These programs can look at pictures of skin problems. Learn what makes them different from one another. Some programs can even find the spot on the skin where the problem is. There are also programs that work together to make sure they find the right answer. These programs can look at the skin problem from angles and understand what is going on. They can even look at complicated skin problems and figure out what is wrong. Skin disease diagnosis methods, like these are getting better and better. Learning and special picture looking programs are really helping doctors understand skin problems. However, challenges such as dataset imbalance, diversity, and real-world deployment remain open research issues.

METHODOLGY

The method we are using has a steps. First we get the pictures ready. We make sure they are all the size and that the numbers are stable. This is important when we are training the computer. We have a problem with the pictures we have. Some types of pictures are more common than others. To make sure the computer can recognize all types of pictures we make pictures by rotating and flipping the ones we have. We also make them bigger and smaller. This helps the computer learn better.

We use tools to help the computer find the bad spots on the skin. One tool is called CNN. It helps the computer tell what is in the picture. Another tool is called YOLO. It helps the computer find the exact spot that is bad. After that we see how well the computer is doing. We call this performance evaluation. We do all these steps to make sure the computer can find the spots, on the skin using dermoscopic images. For classification we use a Convolutional Neural Network or CNN for short to look at pictures and find things that make them different from one another. The CNN is made up of a few parts, including convolutional layers that help find features in the pictures pooling layers that make the pictures smaller and fully connected layers that help make a final decision about what is in the picture. The network looks at things like the color of the pictures the patterns in the pictures and the way things are structured in the pictures. The CNN uses something called a softmax function to figure out how likely it is that a picture belongs to a group and it uses something called categorical cross-entropy loss to make sure it is making the best decisions it can. The CNN is really good at finding things about lesions, such as what color they are what they look like and what shape they are and it uses this information to make a decision, about what the lesion is.

The YOLO framework is used to find the location of lesions. YOLO is special because it can do this in one step. It looks at the picture. Says where the lesion is and what it thinks it is. The YOLO framework does this by guessing the box around the lesion and what type of lesion it is all at the time. This makes it very fast and accurate at finding skin lesions. To see how good the YOLO framework is, at finding lesions we compare the box it guesses to the box around the lesion. We use something called the Intersection over Union metric to do this. The YOLO framework uses this metric to figure out how sure it is that it found the lesion.

The dataset is divided into training and testing subsets for model validation. Performance is evaluated using standard metrics such as accuracy, precision, recall, and F1-score to assess both classification effectiveness and detection reliability. This integrated CNN–YOLO framework ensures efficient lesion localization and accurate skin disease classification.

Algorithm

- Input: Dermoscopic skin images
- Output: Detected lesion region and predicted disease class
- Acquire dermoscopic images from the HAM10000 dataset.
- Resize images to a uniform size and normalize pixel values.
- Apply data augmentation to enhance dataset diversity.
- Detect lesion regions using the YOLO detection model.
- Extract detected lesion regions and pass them to the CNN model.
- Perform feature extraction and classify the lesion into disease classes.
- Evaluate the trained model using accuracy, precision, recall, and F1-score

RESULTS

The proposed YOLO–CNN framework effectively detects and classifies skin diseases using the HAM10000 dataset. YOLO provides accurate lesion localization with confidence scores, while CNN ensures reliable disease classification.

Figure 1.2: Sample skin disease images representing multiple dermatological conditions

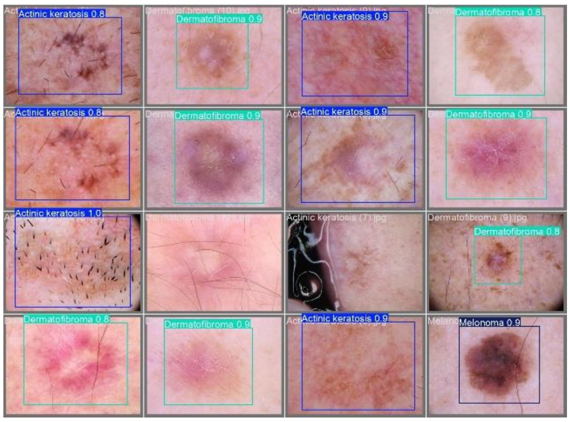

Figure 1.3 : Sample YOLO-based skin disease detection results on the HAM10000 dataset.

The visual results shown in the above figure 1.3 demonstrate the effectiveness of the YOLO model in accurately detecting and localizing skin lesions across multiple disease categories such as benign keratosis, actinic keratosis, and melanoma. The model successfully draws bounding boxes around the affected regions and assigns correct disease labels with high confidence scores, indicating robust lesion identification performance. These qualitative results confirm the ability of YOLO to distinguish between visually similar skin conditions and provide reliable localization, which supports its integration with the CNN model for improved skin disease classification.

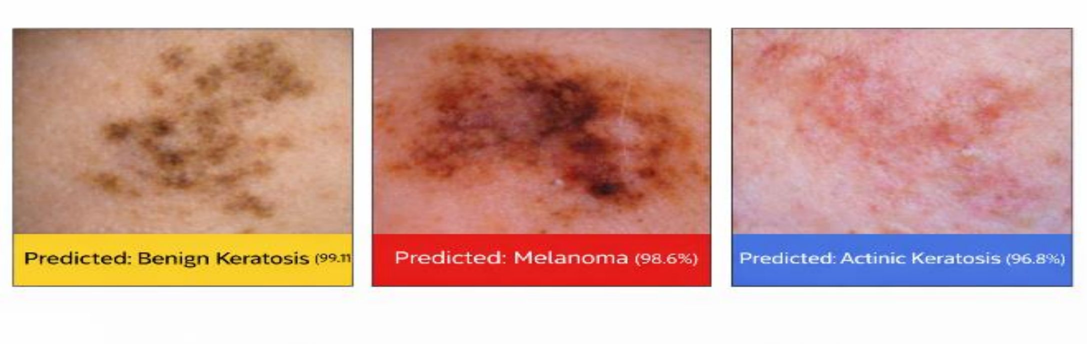

Figure 1.4: Sample prediction outputs of the CNN classification model.

The CNN prediction outputs shown in the figure 1.4 highlight the effectiveness of the proposed classification model in accurately identifying different skin disease categories. The model successfully classifies the detected lesion regions into their respective disease classes based on learned deep features such as texture, color distribution, and structural patterns. These qualitative results demonstrate the ability of the CNN to distinguish between visually similar skin conditions and validate its role in complementing the YOLO-based lesion detection process.

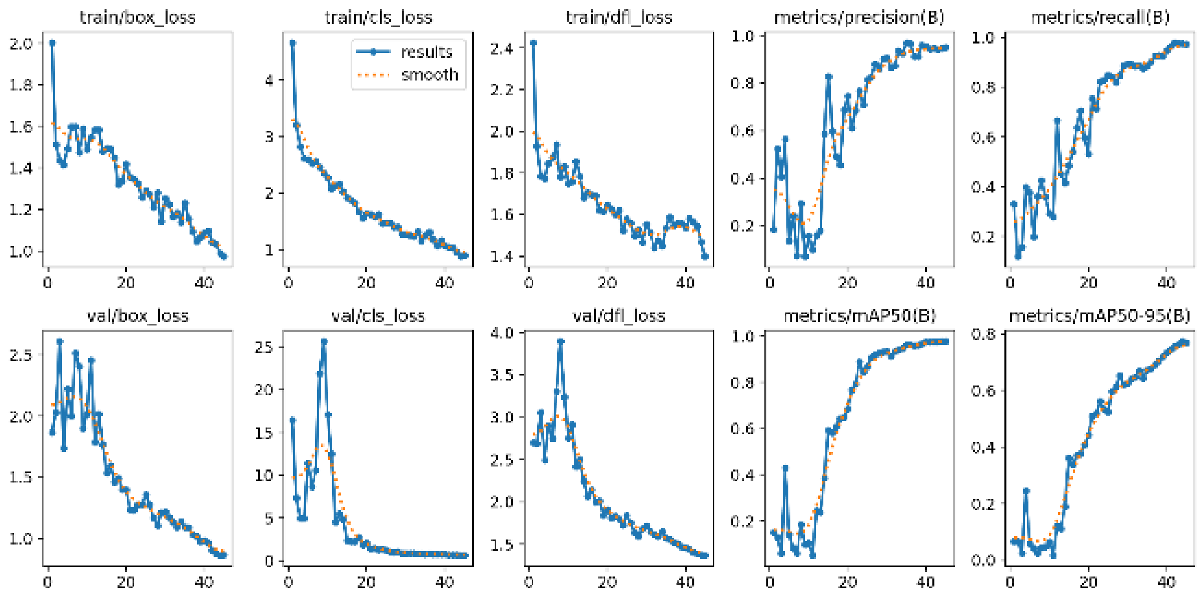

Figure 1.5: Training and validation accuracy of the CNN classification model

The above figure 1.5 illustratres the training and validation accuracy graph illustrates the learning behavior of the CNN model over multiple epochs. The steady increase in accuracy and the close alignment between training and validation curves indicate effective feature learning and good generalization, with minimal overfitting.

DISCUSSIONS

The test results show that using YOLO and Convolutional Neural Network models together is a way to automatically detect and classify skin diseases. YOLO is really good at finding the spot where the skin is affected and giving it a label that says what disease it is and it also gives a score that shows how sure it is. This step of finding the spot on the skin helps the system focus on the actual problem area, which means it is not confused by other things, in the picture and that makes it work better. YOLO and Convolutional Neural Network models work well together for skin disease detection.

The CNN model makes the framework better by finding details from the parts of the skin that have lesions. The results of using the CNN model to classify skin diseases show that it can tell the difference between skin diseases that look similar like keratosis and actinic keratosis and it does this correctly most of the time. The CNN model is good at learning from the data it is trained on. It can apply this learning to new data, which means it is not overly specialized, to the training data and it works well in general.

Compared to traditional CNN-only approaches, the proposed hybrid framework offers improved diagnostic reliability by combining lesion detection and disease classification. The use of YOLO also enables real-time detection, making the system suitable for practical medical and tele-dermatology applications. Although the model demonstrates strong performance, its effectiveness depends on image quality and dataset diversity. Future improvements may include training on larger datasets, incorporating additional disease classes, and integrating the system into mobile or web-based healthcare platforms.

CONCLUSION

This project is about creating a system that can detect skin diseases. The system uses something called YOLO and Convolutional Neural Networks. It can find skin problems show where they are and tell what kind of disease it is. The system looks at pictures of skin. Says what is wrong with it. It is pretty good at finding the problems and telling what they are. The system uses a set of pictures called the HAM10000 dataset to learn about skin diseases. By looking at the pictures and using computer programs the system can tell what is wrong with the skin and where the problem is. This makes the system better at finding skin diseases and understanding what it is looking at. The system is good for people who want to know what is wrong with their skin. Skin disease detection is what the system is, for. The results show that this approach can support dermatologists in early diagnosis and can be useful in tele-dermatology and clinical decision support systems. In the future, the model can be further improved by training on larger datasets and deploying it in real-world healthcare applications such as mobile or web-based platforms.

ACKNOWLEDGMENT

The authors sincerely thank their project guide and the Department of Information Science and Engineering for their valuable guidance and support during the course of this work.

REFERENCES

1. P. Tschandl, C. Rosendahl, and H. Kittler, “The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions,” Scientific Data, Springer Nature, vol. 5, no. 1, pp. 1–9, 2018.
2. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, Springer Nature, vol. 542, no. 7639, pp. 115–118, 2017.
3. R. Ashraf et al., “Region-of-Interest based transfer learning assisted framework for skin cancer detection,” IEEE Access, vol. 8, pp. 147858–147871, 2020.
4. M. A. Adegun and S. Viriri, “FCN-based DenseNet framework for automated detection and classification of skin lesions in dermoscopy images,” IEEE Access, vol. 8, pp. 150377–150396, 2020.
5. H. A. Haenssle et al., “Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition,” Annals of Oncology, Springer, vol. 29, no. 8, pp. 1836–1842, 2018.
6. Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, Springer Nature, vol. 521, no. 7553, pp. 436–444, 2015.
7. N. Codella et al., “Skin lesion analysis toward melanoma detection: A challenge at the International Symposium on Biomedical Imaging,” IEEE Journal of Biomedical and Health Informatics, vol. 23, no. 2, pp. 501–512, 2019.
8. M. Combalia et al., “BCN20000: Dermoscopic lesions in the wild,” Data in Brief, Elsevier, vol. 19, pp. 2075–2079, 2018.
9. Automated Skin Cancer Detection Using Convolutional Neural Networks:A Deep Learning Approach for Dermoscopic Image Analysis,” 2025 IEEE Conference Publication, IEEE Xplore, 2025. R. Nair, P. M. Ebin, T. Babu.

10. Skin Cancer Detection Using Dermoscopic Images with Convolutional Neural Network," Scientific Reports, vol. 15, Art. 7252, Mar. 2025, doi: 10.1038/s41598-025-91446-6. K. Nawaz, A. Zanib, I. Shabir, et al.