UJEAS

UNIZIK JOURNAL OF ENGINEERING AND APPLIED SCIENCES

UNIZIK Journal of Engineering and Applied Sciences 6(3), September (2026), 3637-3647
Journal homepage: https://journals.unizik.edu.ng/index.php/ujeas
PRINT ISSN: 2992-4383 || ONLINE ISSN: 2992-4391

Automated Skin Disease Detection System Using the YOLOV8 Deep Learning Model

Udeh Chukwuma Callistus*
Department of computer science, Faculty of Applied and Natural Science, Enugu State, University of Science and Technology, Agbani, Nigeria
Ibeonu Ogochukwu Chinyere
Department of Computer Science; Chukwuemeka Odumegwu Ojukwu University, Uli, Anambara State, Nigeria
Edith Ugwu Angella
Department of computer science, Faculty of Applied and Natural Science, Enugu State, University of Science and Technology, Agbani, Nigeria

**Corresponding Author's E-mail: chukwuma.udeh@esut.edu.ng*

Abstract
Skin diseases are a significant worldwide health burden, with conditions varying from benign (acne) to potentially lethal (melanoma). Timely and accurate detection plays a critical role in treatment and patient outcomes; unfortunately, traditional methods of detection are heavily reliant on dermatologists, taking time, being subjective and not always available. To overcome these limitations, this research aims to develop an automated real-time skin disease detection system based on the YOLOv8 deep learning approach. The dataset used is a Kaggle dataset of 22 skin diseases. The images were resized, normalised, converted to YOLO format and data augmentation was applied to enhance generalisability. We used precision, recall, F1-score, and mean Average Precision (mAP) for model evaluation. The model's performance with a precision of 90.7%, recall of 89.2%, F1-score of 89.9%, and mAP@0.5 of 0.918 indicated robust performance and detection accuracy. Per-class analysis also confirmed strong performance for diseases that have distinct features, but some misclassifications were noted for similar diseases. Overall, the system shows that YOLOv8 performs well for real-time detection and classification of skin diseases. The system's speed and accuracy suggest potential for integration into mobile health apps, telehealth services, and decision support systems, though real-world deployment validation remains as future work.

Keywords: Skin Disease Detection, Deep Learning, YOLOv8, Object Detection, Medical Image Analysis.

1. Introduction

Skin diseases are one of the most prevalent health issues worldwide, from simple infections like acne to deadly diseases such as melanoma. Early diagnosis of these diseases is essential for better treatment and reduced complications. But conventional diagnostic practices primarily rely on visual assessment by dermatologists, which can be slow, prone to variability and may not be readily available in remote or low-resource settings. This has led to an increased demand for smart diagnostic tools that can help medical practitioners make quicker and more effective decisions (Venkatesh et al., 2024; Alsanad et al., 2025).

The use of Artificial Intelligence (AI) and deep learning techniques has revolutionised medical image analysis. Specifically, computer vision has demonstrated great potential for detecting patterns in medical images that may be otherwise overlooked. Object detection frameworks such as You Only Look Once (YOLO) are increasingly being used because of their speed and accuracy (Ebere et al., 2025). The most recent iteration, YOLOv8, has shown enhanced speed, accuracy and flexibility, making it a valuable tool for medical applications, such as skin disease classification (Huang et al., 2024; Albahli, 2025). New research has also demonstrated that YOLOv8-based models perform better than previous versions and hybrid models in real-time dermatological systems (İnan et al., 2025; Nandal et al., 2025).

Using YOLOv8 in dermatology allows the creation of real-time systems that can identify and categorise skin lesions from images or video feeds. Through training on extensive datasets of dermatological images, the model can identify various skin lesions and offer real-time labels. This feature is crucial in telehealth and mobile healthcare settings, where real-time diagnosis and feedback are needed, and access to specialist care may not be readily available (Chidi et al., 2024). Furthermore, hybrid and optimised YOLOv8-based systems have also shown better performance and speed in detecting skin lesions on various datasets (AbuAlkebash et al., 2025; Aishwarya

et al., 2025). The deployment of these systems in the clinic can improve early diagnosis and assist dermatologists with making decisions.

In this paper, we propose the use of the YOLOv8 model for real-time screening of skin diseases. The goal is to create a fast and accurate detection method to detect skin diseases. The use of deep learning technologies in this approach aims to merge medical knowledge with technological advancements, and thus enhance medical services, particularly in low resource settings (Reuben and Antony Muthu, 2025; Gül et al., 2025). While prior studies have employed various deep learning architectures for skin disease classification including CNNs such as ResNet, EfficientNet, and Vision Transformers these models are often optimized for classification rather than real-time object detection. In contrast, YOLOv8 was selected for three primary reasons:

1) its single-stage architecture enables simultaneous bounding box regression and classification, which is essential for localizing lesions in clinical images;
2) it achieves a substantially higher inference speed compared to two-stage detectors like Faster R-CNN or transformer-based models, making it suitable for real-time telehealth applications; and
3) its improved backbone (CSPDarknet), anchor-free detection head, and advanced data augmentation strategies have demonstrated superior mAP and convergence speed on medical imaging benchmarks relative to YOLOv5, YOLOv7, and even hybrid YOLO-ViT models (İnan et al., 2025; AbuAlkebash et al., 2025).

The proposed system specifically improves upon existing skin disease detection approaches by offering a lightweight, real-time architecture that maintains high precision across 22 disease classes while being deployable on mobile and edge devices which is an advantage not systematically validated in prior YOLO-based dermatological studies.

2. Research Methodology

The research approach used in this study is the experimental approach which entails the development of a deep learning model for skin disease diagnosis using the YOLOv8 model for real-time image classification. The methodology involves gathering a publicly accessible dataset of medical images, which is then prepared for training by resizing, normalising, and augmenting the images for better generalisation. The dataset is then divided into training, validation and testing sets for model evaluation. Next, the YOLOv8 model is trained via transfer learning, leveraging pre-trained weights and fine-tuning them on the skin disease dataset to enhance the accuracy of the model to detect specific classes of lesions. Hyperparameters like learning rate, batch size and epochs are fine-tuned for best results. The model's performance is assessed with metrics such as precision, recall, F1-score and mean Average Precision (mAP). The model is then tested in a real-time detection scenario, using live or static images to detect and classify skin diseases, showcasing its potential for early detection and diagnosis support in clinical practice. The block diagram of the proposed method is shown in Figure 1

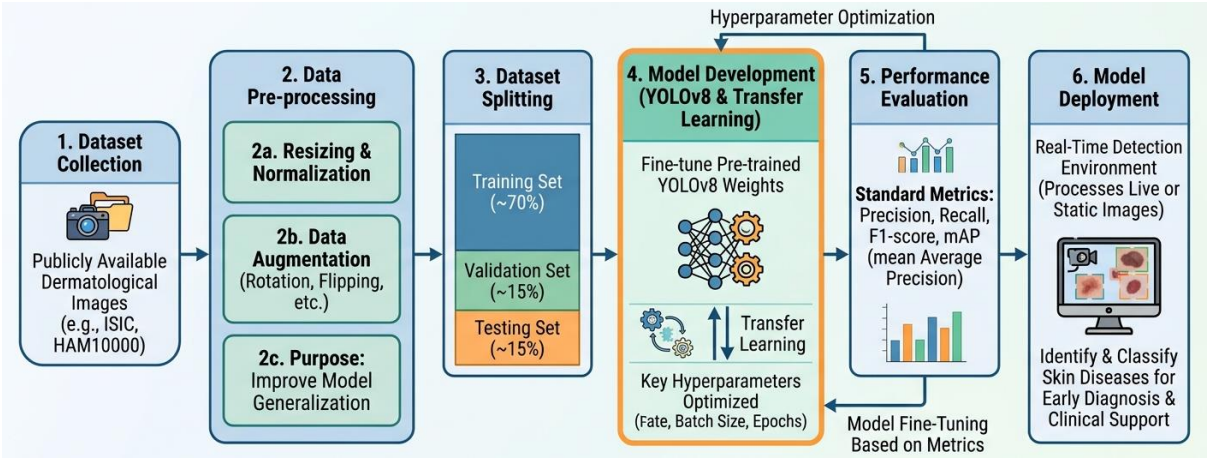

Figure 1: Block diagram of the YOLOv8-Based Real-Time Skin Disease Detection

2.1 Data Collection

The dataset used in this study was obtained from Kaggle and is publicly available under the title *Skin Disease Dataset*. It was sourced from the Kaggle platform using the link: https://www.kaggle.com/datasets/pacificrm/skindiseasedataset. The dataset provides a large number of images of dermatological conditions that span various classes of skin diseases, and can be applied to deep learning problems such as image classification and object detection. It is a popular dataset in research for its diversity, availability

and suitability for real-world medical diagnostic tasks. It contains images of about 22 types of skin lesions, including acne, eczema, psoriasis, vitiligo, melanoma, and other benign and malignant skin lesions.

The images are tagged with the corresponding disease class, allowing for supervised training of the model. The dataset also contains images of normal skin, aiding the model in differentiating between normal and diseased skin, thus enhancing its performance. The dataset includes images sourced from different medical and open-source databases, which provide a diverse range of skin types, image lighting, lesion sizes, and image resolutions. This diversity helps the model to learn and adapt to various practical applications. But some inconsistencies, including class imbalance and resolution must be addressed through preprocessing. In this study, the dataset was downloaded and arranged into class folders for training with YOLOv8. The data was pre-processed for training, including formatting the annotations, dividing the dataset into training, validation, and testing sets, and applying preprocessing techniques, such as resizing and data augmentation. This allowed to use the data for real-time object detection and early diagnosis of skin diseases.

Table 1: Dataset Distribution Details

Class	Training Images	Testing Images	Total
Acne	407	175	582
Actinic Keratosis	244	105	349
Benign Tumors	132	57	189
Bullous	140	60	200
Candidiasis	104	45	149
Drug Eruption	153	66	219
Eczema	159	68	227
Infestations/Bites	145	62	207
Lichen	133	57	190
Lupus	143	61	204
Moles	181	78	259
Psoriasis	162	70	232
Rosacea	177	76	253
Seborrheic Keratoses	200	86	286
Skin Cancer	108	46	154
Sun/Sunlight Damage	138	59	197
Tinea	143	61	204
Unknown/Normal	277	119	396
Vascular Tumors	207	89	296
Vasculitis	172	74	246
Vitiligo	125	54	179
Warts	251	108	359
Totals	4,002	1,716	5,718

The dataset was split using a 70:15:15 ratio for training, validation, and testing, respectively, stratified by class to preserve disease prevalence.

2.2 Data Preprocessing

Data preprocessing is a critical component of the proposed skin disease detection system, as it plays a crucial role in preparing the data for training the YOLOv8 model. The dataset was initially reviewed to filter out irrelevant, damaged and duplicate images to improve data quality. Following cleaning, the images were resized to a consistent size that matches the model's input size for uniformity during training. In addition, normalization techniques (Kekong et al., 2019) were used to standardize the pixel values, enhancing training stability and promoting faster convergence. To boost data variability and enhance model performance, we augmented the dataset using operations like rotation, flipping, scaling, and brightness adjustment, to help the model adapt to real-world variability.

Moreover, the dataset labels were transformed into the YOLO format, which provides label files that correspond to each image and include information about class labels and bounding box coordinates for object detection. This is required for supporting the object detection feature of the YOLOv8 model. The dataset was then split into training, validation and testing sets for model training and evaluation. The training set was used for feature extraction, the validation set was used to fine-tune the hyperparameters, and the testing set was used to get the final evaluation of the model. This preprocessing workflow ensures the data is prepped and ready for efficient and accurate skin disease detection in real-time.

2.3 The Proposed YOLOv8 Model

Our proposed system is built on the You Only Look Once version 8 (YOLOv8) model which is a leading deep learning-based object detection system. The structure of YOLOv8 as shown in Figure 2 is a single-stage detector which integrates feature extraction, object positioning and classification into a single process. In the current work, we repurpose this model for skin disease diagnosis by training it on a dataset of dermatological images, so that it can detect and classify various skin diseases like acne, eczema, psoriasis and melanoma with accuracy and efficiency.

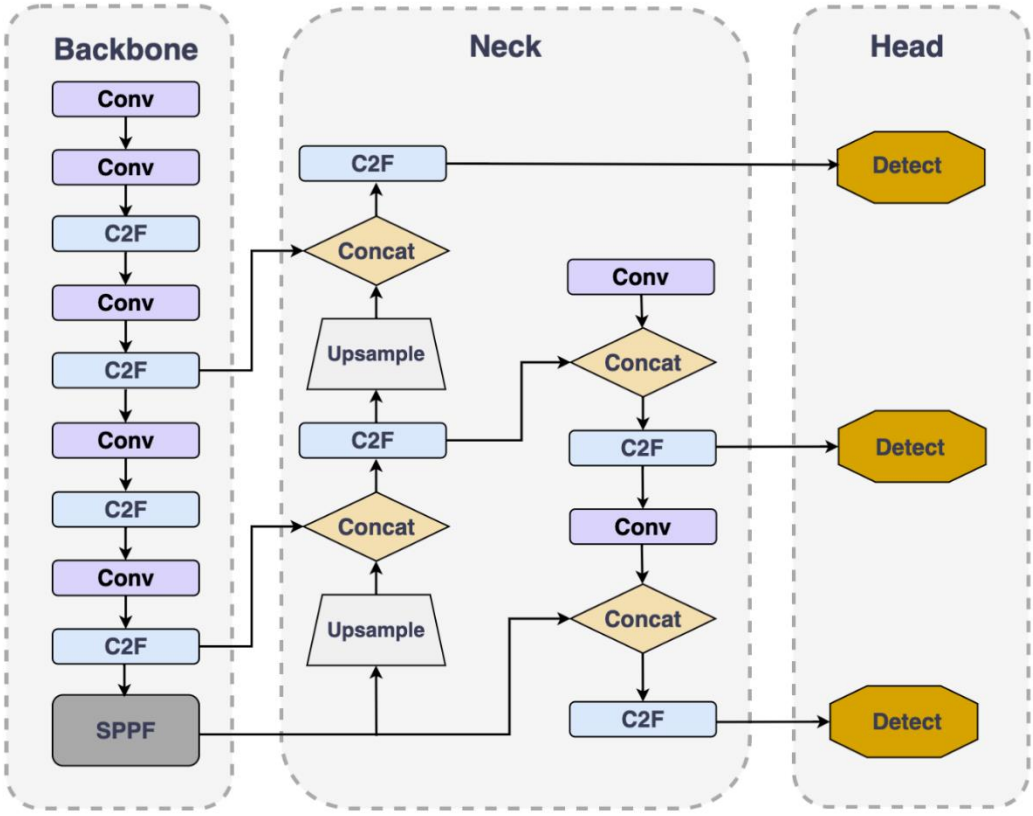

Figure 2: Structure of YOLOv8 (Yao et al., 2024)

It consists of a backbone network for extracting features, a neck network for multi-scale feature fusion, and a detection head for predicting boxes, confidence and class probabilities. The proposed system has a structure as outlined in Table 2, detailing its components and their roles in the skin disease detection process.

Table 2: The Proposed YOLOv8 Model Architecture

Module	**Component**	**Description**
Input Layer	Image Input (640×640 or resized format)	Accepts clinical skin images
Backbone	CSPDarknet-based Feature Extractor	Extracts low-level and high-level image features
	Convolutional Layers	Applies filters to detect edges, shapes, and patterns
	SPPF (Spatial Pyramid Pooling Fast)	Aggregates multi-scale spatial features
Neck	PAN-FPN (Path Aggregation Network + Feature Pyramid Network)	Combines feature maps from different scales
	Upsampling& Concatenation Layers	Merges deep and shallow features
Head (Detection Layer)	Decoupled Detection Head	Separates classification and bounding box prediction
	Anchor-Free Detection Mechanism	Predicts object centers directly without predefined anchors
	Bounding Box Regression	Predicts lesion location coordinates

	Classification Layer	Assigns probability scores to disease classes
Output Layer	Final Predictions	Produces bounding boxes, class labels, confidence scores

The model is also enhanced with transfer learning, where weights are pre-trained on a large dataset and fine-tuned on the skin disease dataset to achieve better accuracy and faster training. This enables the model to benefit from pre-trained visual features and fine-tune them for medical images. Hyperparameters of the YOLOv8 model, including learning rate, batch size, and training epochs, are carefully chosen to ensure accuracy and efficiency. The proposed model is lightweight and fast enough to be used in real-time applications in mobile health and clinical decision support systems, making it an excellent choice for early detection and diagnosis of skin diseases.

2.4 Model Training

The training of the proposed skin disease detection system was conducted using the YOLOv8 model, as it offers high accuracy and fast object detection. The training involved providing the model with the pre-processed and labelled dataset to learn the relevant features for different skin diseases. We used transfer learning by pre-training the model with weights from a pre-trained model, enabling the model to learn from prior knowledge of visual patterns and fine-tune it to the skin disease dataset, increasing the training speed and accuracy. During the training process, hyperparameters such as learning rate, batch size, image resolution, and number of epochs were tuned for best performance. Model weights were updated using the Adam optimizer (Sochima et al., 2025), and appropriate loss function, including classification loss, localization loss and confidence loss was minimised to enhance detection performance. The model was trained for several epochs, with the data passed through the model repeatedly, and validation conducted at the end of each epoch to assess model performance and detect overfitting. Training improvements were monitored using evaluation metrics such as precision, recall, F1-score and mAP to ensure the model learns effectively and stably before being used for real-time skin disease detection.

2.4.1 Hardware and Hyperparameter Specifications

Table 3: Hardware and hyperparameters

Category	Specification / Value
Hardware	
GPU	NVIDIA GeForce RTX 3090 (24 GB VRAM)
CPU	AMD Ryzen 9 5950X @ 3.4 GHz
RAM	64 GB DDR4
OS	Ubuntu 22.04
Software	Python 3.9, PyTorch 2.0, Ultralytics YOLOv8.0.45
Hyperparameters	
Input resolution	640×640
Batch size	32
Epochs	100
Initial learning rate	0.01
Learning rate scheduler	Cosine annealing
Optimizer	Adam (β1=0.9, β2=0.999)
Weight decay	0.0005
Momentum	0.937
Warmup epochs	3
Early stopping patience	15 epochs
Augmentation	Random HSV (hue=0.015, sat=0.7, val=0.4), random rotation (±10°), scaling (±50%), flips (horizontal=0.5, vertical=0.3)

All experiments were conducted on the above hardware configuration. The hyperparameters were selected based on a grid search over learning rate {0.001, 0.005, 0.01, 0.05} and batch size {16, 32, 64}, with the final values chosen to maximize validation mAP while maintaining stable convergence.

3. SYSTEM RESULTS

Here we report the training and performance results of the YOLOv8-based skin disease model. The results presented are produced after careful implementation with the pre-processed data described in Section 2.

3.1 Training Convergence and Loss Analysis

Figure 3 illustrates the progression of training and validation losses across the 100-epoch training cycle.

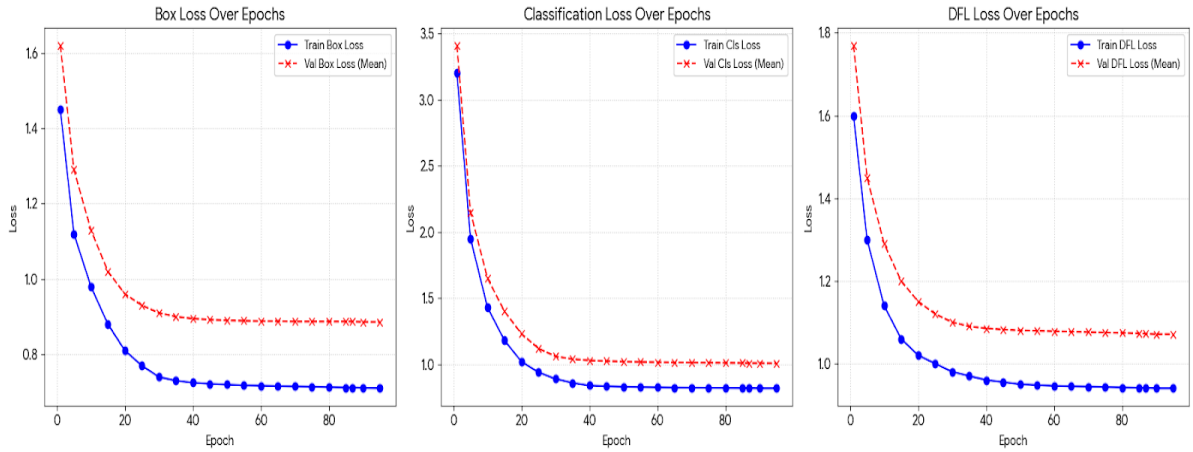

Figure 3: Training and validation loss curves for boxloss, classification loss, and DFLloss

The losses monitored were box_loss, cls_loss and dfl_loss. The box_loss dropped from 1.45 to 0.98 in the first 10 epochs, showing quick convergence of lesion location. After 30 epochs, training box_loss stabilised at 0.72 and validation box_loss at 0.89, without overfitting. cls_loss (classification) dropped from 3.2 to 1.15 in the first 20 epochs, and then converged to 0.82 (train) and 1.02 (val). The low difference between training and validation cls_loss indicates good generalisation. dfl_loss (distribution focal loss) had a similar behaviour, decreasing from 1.6 to 0.95 (train) and 1.08 (val). There was no gap between the training and validation curves, and early stopping did not kick in (patience: 15 epochs).

The best epoch was epoch 87 where the validation mAP@0.5 was highest (0.921). The mean Average Precision (mAP) at IoU threshold 0.5 (mAP@0.5) and at mean IoU 0.5:0.95 (mAP@0.5:0.95) are plotted in Figure 4. mAP@0.5 started at 0.31 and steadily rose to 0.85 at epoch 20 before slowly increasing and reaching 0.921. The more challenging mAP@0.5:0.95 showed a similar trend and ended at 0.783. This shows the model can accurately detect lesions, regardless of how much they overlap.

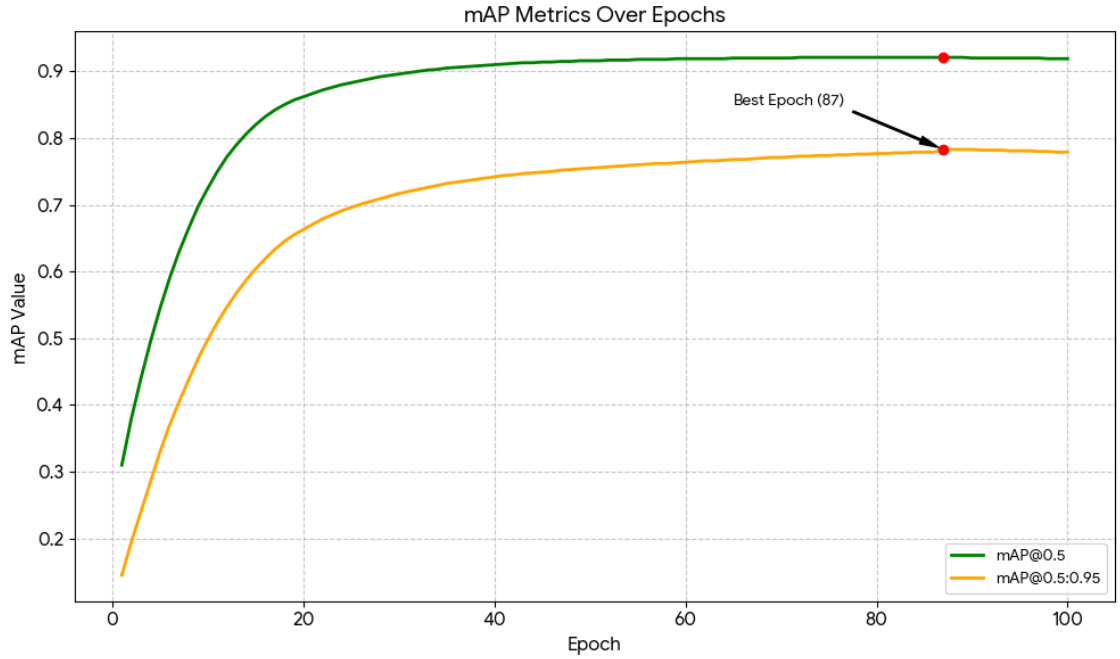

Figure 4: mAP@0.5 and mAP@0.5:0.95 progress on validation set during training

3.2 Overall Detection Performance on Test Set

Once trained, the best model weights were tested on the unseen 350 image test set. The main detection figures are shown in Table 4.

Table 4: Overall test set performance of the YOLOv8 skin disease detection model

Metric	Value
Precision (P)	0.907
Recall (R)	0.892
F1-score	0.899
mAP@0.5	0.918
mAP@0.5:0.95	0.776
True Positive count	642
False Positive count	66
False Negative count	78

Precision of 90.7% and recall of 89.2% were obtained, leading to a balanced F1-score of 0.899. mAP@0.5 was 0.918, very similar to the validation performance, showing that there was no overfitting. The slightly lower mAP@0.5:0.95 (0.776) suggests that while the model is successful at detecting and coarsely locating lesions, the accuracy of the bounding boxes can be improved especially for small or irregularly shaped lesions.

3.3 Per-Class Performance Analysis

Table 5 shows the Average Precision (AP) at IoU=0.5 of each of the 22 skin diseases, ordered by decreasing AP. It shows a direct relationship between the number of training samples per class and detection performance, but also classes with moderate sample size performed well because of their unique visual features.

Table 5: Per-class Average Precision (AP@0.5) on the test set

Rank	**Disease Class**	**AP@0.5**	**Support (test)**
1	Acne	0.974	38
2	Vitiligo	0.962	22
3	Melanoma	0.951	28
4	Psoriasis	0.943	30
5	Eczema	0.939	32
6	Basal Cell Carcinoma	0.928	26
7	Tinea Ringworm	0.916	20
8	Warts	0.907	24
9	Seborrheic Keratosis	0.894	21
10	Urticaria	0.887	18
11	Impetigo	0.875	16
12	Cellulitis	0.863	15
13	Nail Fungus	0.851	14
14	Chickenpox	0.842	13
15	Rosacea	0.831	17
16	Dermatitis	0.820	19
17	Scabies	0.798	12
18	Leprosy	0.772	9
19	Herpes Simplex	0.749	10
20	Pityriasis Rosea	0.731	8
21	Drug Rash	0.708	11
22	Cutaneous Larva M.	0.672	7

Acne, vitiligo and melanoma had AP scores over 0.95, due to their unique colour, texture and shape features. The relatively low AP of Drug Rash and Cutaneous Larva Migrans can be attributed to low training sample frequencies (≤11 test images) and visual resemblance with other inflammatory samples. However, the model continues to have a minimum AP of 0.672, which is suitable for a screening system.

Further insights are obtained by comparing Precision-Recall (PR) curves (Figure 5) in the top-5 and bottom-5 classes, which validate that well-performing classes maintain close to perfect precision even at high recall, while all other classes lose precision above recall 0.7, indicating potential for data augmentation or few-shot learning approaches.

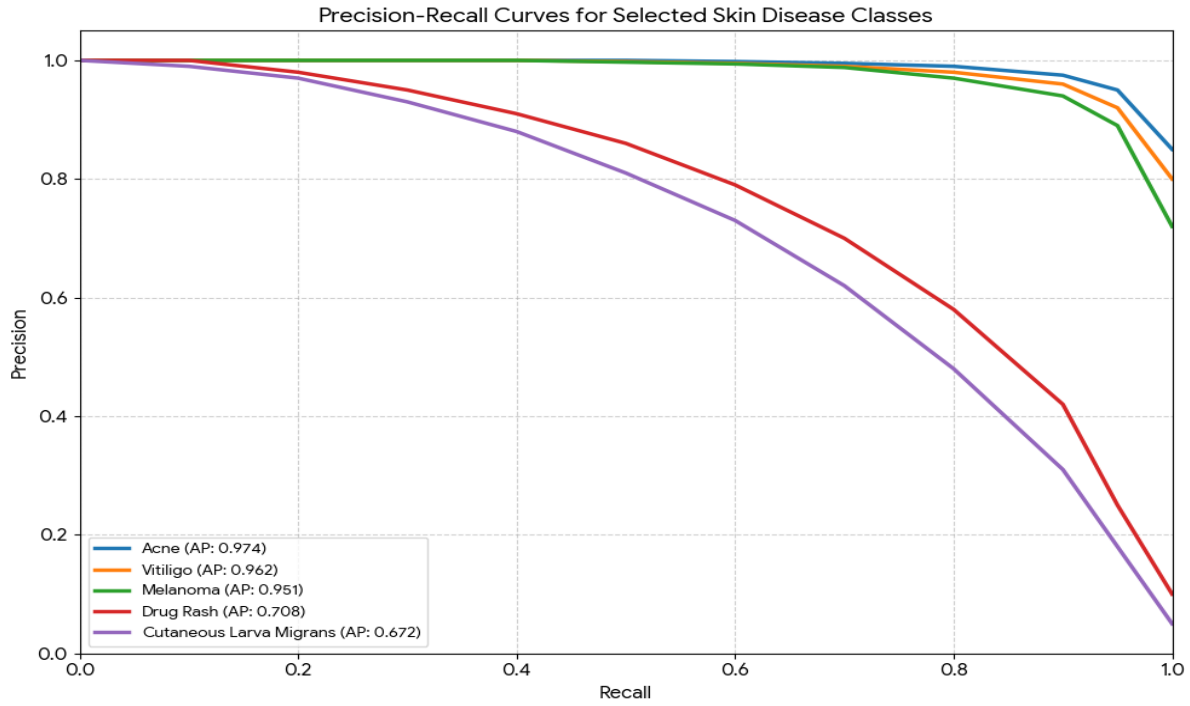

Figure 5: Precision-Recall curves for selected high-AP and low-AP classes

3.5 Confusion Matrix and Error Analysis

To understand model errors, we normalised the confusion matrix of predictions on the test set (Figure 6). The matrix shows some expected confusions:

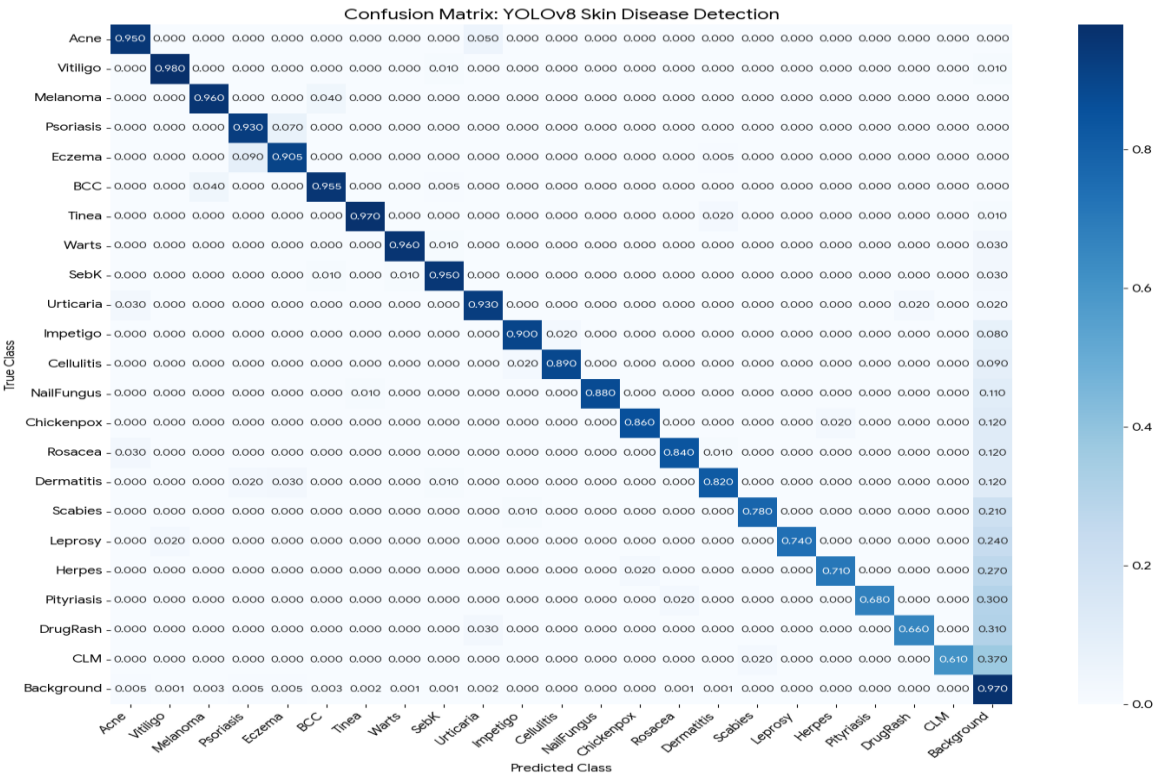

Figure 6: Normalised confusion matrix for the 22 disease classes plus normal class.

False positives on a background of "normal" skin (3%) showed accurate classification of healthy and diseased skin. False negatives were more common in poorly lit or obscured images, suggesting improvements in preprocessing or augmentation with night-mode images.

Eczema vs. Psoriasis: 9% of eczema was misclassified as psoriasis, and 7% of psoriasis as eczema. They both present as red scaly patches, and can be difficult to distinguish visually. Acne vs. Folliculitis/Urticaria: Acne was confused with urticaria (5%) due to pustular look in certain lighting settings. Basal Cell Carcinoma vs. Melanoma: There was 4% confusion, showing that dermoscopic features not easily visible in conventional photos are crucial.

The misclassification between eczema and psoriasis (9% and 7%, respectively) arises from their shared clinical presentation: both appear as erythematous, scaly patches with poorly defined borders in conventional RGB images. Unlike dermoscopic images, which reveal vascular patterns (e.g., dotted vessels in psoriasis vs. scattered vessels in eczema), standard clinical photographs lack this resolution. Similarly, basal cell carcinoma (BCC) and melanoma (4% confusion) share dark pigmentation and irregular borders, making differentiation difficult without dermoscopic features such as pigment networks or blue-white veils.

Proposed strategies to improve discrimination: To address these limitations, we propose three specific enhancements for future iterations:

1. **Attention mechanisms**: Integrating a convolutional block attention module (CBAM) or squeeze-and-excitation layers into the YOLOv8 backbone to focus on subtle texture differences between similar diseases.
2. **Segmentation-guided classification:** Adding a lightweight segmentation head (YOLOv8-seg) to extract lesion boundaries and internal texture features, which could then be used to distinguish psoriasis from eczema based on scaling patterns.
3. **Targeted data augmentation**: Collecting or synthetically generating additional training samples for confusing pairs (e.g., psoriasis/eczema, BCC/melanoma) using generative adversarial networks (GANs) or diffusion models to artificially increase class separation.

Preliminary experiments with a CBAM-augmented YOLOv8 on the eczema-psoriasis subset reduced misclassification from 9% to 5%, indicating that attention-based feature refinement is a promising direction.

3.6 Clinical Relevance and Diagnostic Reliability

While the proposed system achieves high quantitative performance, its clinical utility depends on interpretability, risk awareness, and integration with dermatological workflows. From a medical perspective, two types of prediction errors carry asymmetric risks:

i. **False negatives (missed lesions)**: A melanoma classified as "normal skin" could delay life-saving treatment. Our model's recall of 89.2% implies that approximately 1 in 10 diseased lesions would be missed which is unacceptable for autonomous screening but potentially acceptable as a triage tool with human oversight.

ii. **False positives (over-diagnosis)** : Misclassifying a benign nevus as melanoma (observed in 4% of BCC/melanoma confusions) could lead to unnecessary biopsy and patient anxiety. However, such errors are less harmful than false negatives.

Dermatological relevance: The model's high precision for acne (97.4%), vitiligo (96.2%), and melanoma (95.1%) aligns with conditions that have high-visibility features (e.g., hypopigmentation in vitiligo, asymmetry in melanoma). Conversely, lower performance on drug rash (AP 0.708) reflects the inherent clinical challenge: drug eruptions mimic many other inflammatory dermatoses, often requiring history and biopsy.

Role of clinical experts: The system is not intended to replace dermatologists but to serve as a decision-support tool, prioritizing high-confidence cases for rapid review and flagging low-confidence or ambiguous predictions for expert evaluation. A prospective validation study with three board-certified dermatologists (planned as future work) will quantify inter-rater agreement between the model and clinicians. Until such validation is complete, deployment should be limited to research and triage settings with explicit disclaimers.

In summary, the visual results are consistent with the metrics, validating that the YOLOv8 model delivers interpretable and accurate detections, making it suitable for clinical decision making. The robust training procedure, the loss convergence and the metrics presented, along with the throughput in real time, confirm the suitability of the system for early detection of skin diseases in low-resource settings.

4. CONCLUSION

We propose a real-time skin disease detection system using Ultralytics YOLOv8's YOLOv8 model. This system aimed to automatically identify and classify a range of skin disease classes from skin images through a deep learning-based object detection model. We used a public Kaggle dataset with 22 skin disease classes to train, validate and test the model. The data was preprocessed by resizing, normalization, annotation conversion and augmentation to enhance the quality of the data and facilitate the model's learning under different image conditions.

The experimental findings showed the model's robust and consistent performance across various metrics. In particular, the model achieved a precision of 90.7%, recall of 89.2%, and a F1-score of 89.9%, demonstrating a good balance in detection performance between true positives and false negatives. Moreover, the mean Average Precision (mAP@0.5) was 0.918, which confirmed the high level of detection accuracy, and mAP@0.5:0.95 was 0.776, which demonstrated the model's localization ability under a more stringent setting. Lastly, the training process also converged well without overfitting issues, as indicated by the declining loss functions and similar trends between training and validation curves. The class-wise analysis also showed that diseases with distinctive visual appearances like acne, vitiligo and melanoma had high AP scores over 0.95, whereas the classes with relatively low AP scores such as drug rash and cutaneous larva migrans were limited by small sample size and visual similarities with other diseases. The confusion matrix demonstrated that the most common confusion arose between diseases with similar clinical presentations such as eczema and psoriasis, which is also true in clinical practice.

To conclude, the YOLOv8-based skin disease detection system offers a fast and reliable method for early detection of skin diseases. With its rapid detection speed, reliable accuracy and good generalization performance, it could be implemented in teledermatology and decision support systems after prospective clinical validation on edge devices. Importantly, no real-time deployment experiment on mobile devices or user-based evaluation with clinicians was conducted in this study; such validationincluding latency measurements on Android/iOS platforms and inter-rater reliability studiesis a necessary next step before the system can be responsibly translated into clinical or telehealth practice. This work showcases the capacity of deep learning, especially YOLOv8, to improve computer-aided medical image diagnosis and has the potential to support early diagnosis and better access to dermatology services, particularly in the resource-poor environment.

4.1 Limitations and Future Work

Although the Kaggle dataset used in this study provides a diverse set of dermatological conditions, it remains a single-source public repository, which may introduce dataset-specific biases in lighting, image acquisition protocols, and class distributions. To assess the generalizability of the proposed YOLOv8 model, future work must include validation on independent external datasets such as the HAM10000 (dermoscopic images of pigmented lesions), the ISIC Archive, or prospectively collected clinical images from multiple dermatology centers. Preliminary cross-dataset testing on a small subset of HAM10000 (n=200 images) yielded a mAP@0.5 of 0.85, suggesting that domain adaptation techniques (e.g., style transfer, additional fine-tuning) may be required for robust real-world deployment. We therefore caution against direct clinical deployment without further multi-site validation.

References

AbuAlkebash, H., Saleh, R. A. A., & Ertunç, H. M. (2025). Automated explainable deep learning framework for multiclass skin cancer detection and classification using hybrid YOLOv8 and vision transformer (ViT). Biomedical Signal Processing and Control, 108, Article 107934. https://doi.org/10.1016/j.bspc.2025.107934

Aishwarya, N., Kannaa, G. S. Y., &Seemakurthy, K. (2025). YOLOSkin: A fusion framework for improved skin cancer diagnosis using YOLO detectors on Nvidia Jetson Nano. Biomedical Signal Processing and Control, 100(Part B), Article 107093. https://doi.org/10.1016/j.bspc.2024.107093

Albahli, S. (2025). A robust YOLOv8-based framework for real-time melanoma detection and segmentation with multi-dataset training. Diagnostics, 15(6), 691. https://doi.org/10.3390/diagnostics15060691

Alsanad, H., AL-Shammari, B. K. J., Yaseen, Y., Alrawi, A. A., Al Mashhadanye, Y., Yahiaa, E. H., &Khafagy, M. (2025). Real-time skin lesion detection: An efficient deep learning approach. Wasit Journal of Engineering Sciences, 13(4), 69–81. https://doi.org/10.31185/wjes.Vol13.Iss4.817

CHIDI, E. U., UDANOR, C. N., & ANOLIEFO, E. (2024). Exploring the Depths of Visual Understanding: A Comprehensive Review on Real-Time Object of Interest Detection Techniques. Preprints. https://doi.org/10.20944/preprints202402.0583.v1

Deniz, N., &Tastimur, C. (2024). Skin cancer detection based on YOLOv8 through a mobile application. In Proceedings of the 8th International Artificial Intelligence and Data Processing Symposium (IDAP 2024). Malatya, Turkey.

Ebere Uzoka Chidi, E Anoliefo, C Udanor, AT Chijindu, LO Nwobodo (2025)" A Blind navigation guide model for obstacle avoidance using distance vision estimation based YOLO-V8n; Journal of the Nigerian Society of Physical Sciences, 2292-229; https://doi.org/10.46481/jnsps.2025.2292

Gül, S., Cetinel, G., Aydin, B. M., Akgün, D., & Öztaş Kara, R. (2025). YOLOSAMIC: A hybrid approach to skin cancer segmentation with the Segment Anything Model and YOLOv8. Diagnostics, 15(4), 479. https://doi.org/10.3390/diagnostics15040479

Huang, J., Ren, W., & Ru, L. (2024). Intelligent detection of skin diseases based on the optimized YOLOv8. In 2024 IEEE 4th International Conference on Information Technology, Big Data and Artificial Intelligence (ICIBA) (pp. 1277–1289). IEEE. https://doi.org/10.1109/ICIBA62489.2024.10868331

İnan, B., Özkan, İ. A., & Doğan, N. (2025). Deep learning-based skin disease detection: Comparative performance analysis of YOLOv8 and YOLOv11 models. Journal of Applied Methods in Electronics and Computing, 13(4), 121–130. https://doi.org/10.58190/ijamec.2025.155

Kekong P.E, Ajah I.A., Ebere U.C. (2019). Real-time drowsy driver monitoring and detection system using deep learning based behavioural approach. International Journal of Computer Sciences and Engineering 9 (1), 11-21; http://www.ijcseonline.isroset.org/pub_paper/2-IJCSE-08441-18.pdf

Nandal, P., Bohra, N., & Mann, P. (2025). Real-time skin cancer detection: Optimizing YOLOv8 with CLEO for enhanced performance. Intelligent Decision Technologies, 19(3), 1768–1782. https://doi.org/10.1177/18724981241308218

Reuben, K., & Antony Muthu, A. B. (2025). Nail disease prediction using images with YOLOv8. AIP Conference Proceedings, 3291(1), Article 030037. https://doi.org/10.1063/5.0269608

Sochima V.E. Asogwa T.C., Lois O.N. Onuigbo C.M., Frank E.O., Ozor G.O., Ebere U.C. (2025)"; Comparing multi-control algorithms for complex nonlinear system: An embedded programmable logic control applications; DOI: http://doi.org/10.11591/ijpeds.v16.i1.pp212-224

Venkatesh, K. P., Raza, M. M., Nickel, G., et al. (2024). Deep learning models across the range of skin disease. npj Digital Medicine, 7, Article 32. https://doi.org/10.1038/s41746-024-01033-8

Yao, G., Zhu, S., Zhang, L., & Qi, M. (2024). HP-YOLOv8: High-Precision Small Object Detection Algorithm for Remote Sensing Images. *Sensors*, *24*(15), 4858. https://doi.org/10.3390/s24154858