***Scan, Click, Diagnose*: An AI-Powered System for Real-Time Skin Condition Detection**

Habiba Hye, Department of Computer Science, The College of Wooster

Introduction

Millions suffer from misdiagnosed or undiagnosed skin conditions—especially in underserved regions with limited access to dermatologists. This project introduces a real-time diagnostic tool that combines patient symptoms and images using machine learning to deliver fast, reliable skin condition predictions through a web app.

Research Questions

RQ1: How can the integration of an image-based CNN model with a metadata-driven machine learning model improve overall diagnostic accuracy?
RQ2: What impact do additional patient attributes, such as age, sex, and symptoms, have on model performance and prediction reliability?
RQ3: How can model architecture, training strategies, and ensemble learning techniques be optimized to ensure scalability and robustness across diverse datasets and populations?

Methodology - Datasets

Images

HAM10000: High-quality dermoscopic images

ISIC Archive: Verified, clinical-grade skin lesion dataset

SCIN Dataset: User-reported demographics

Metadata (from SCIN)

Age Group **Race** **Sex at birth**

Fitzpatrick skin type **Symptoms**

Methodology - Models

Image Model: EfficientNet (Convolutional Neural Network) trained on HAM10000 + ISIC

Metadata Model: Random Forest Classifier using structured fields from SCIN

Combined Model: Ensemble Meta-Learner that fuses predictions from image and metadata models. Final predictions generated using a weighted voting mechanism

Results

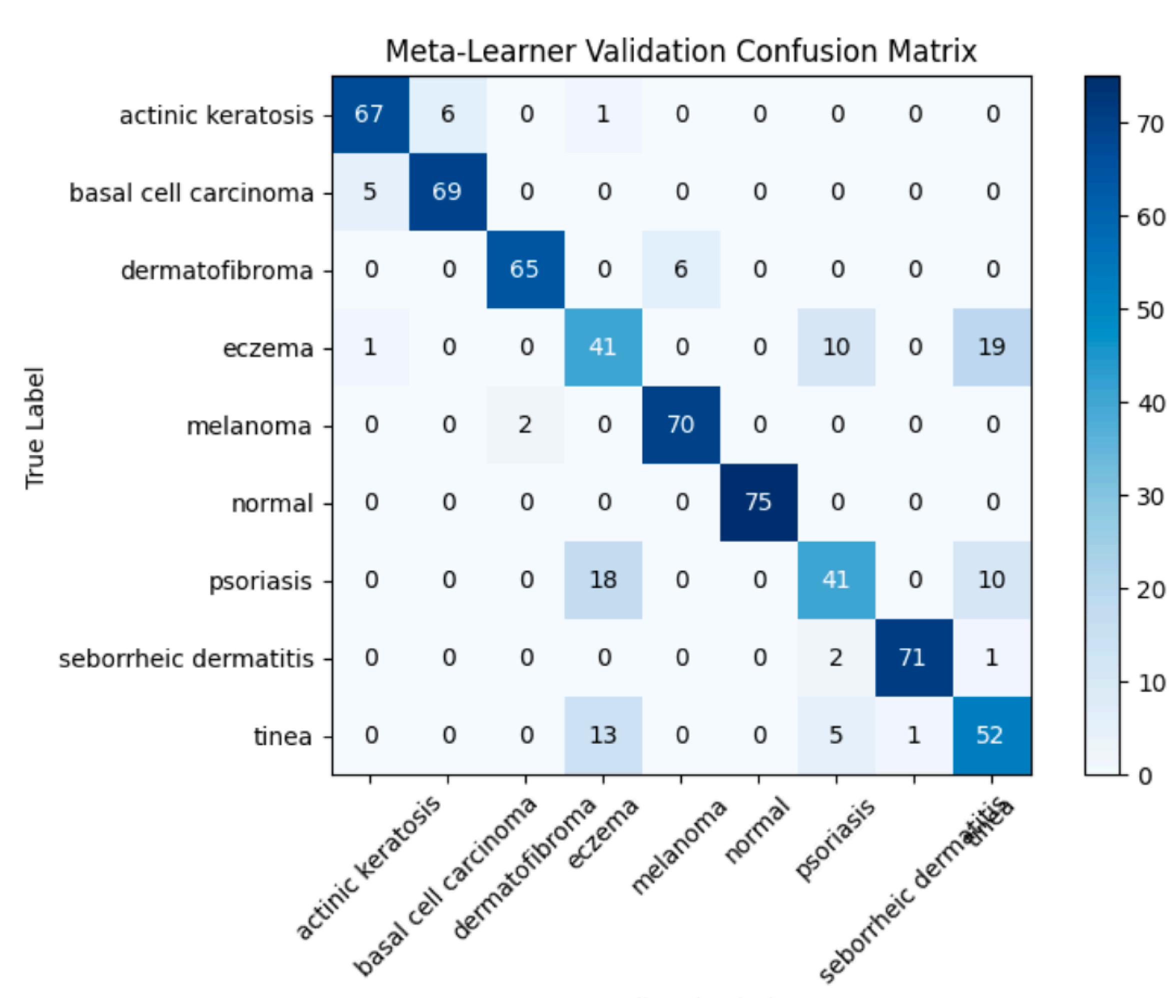

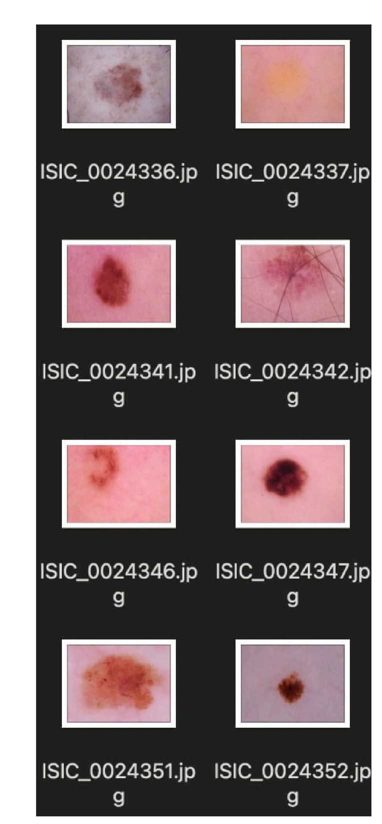

Figure: Confusion Matrix of Ensemble Model, Examples of Medical Graded Images from HAM10000

The final ensemble model, which combines predictions from both the CNN image classifier and a metadata-based model using, achieved an accuracy of 0.91—outperforming the image-only model (0.84) and metadata-only model (0.88). The confusion matrix illustrates the ensemble model's strong classification performance across nine skin conditions, with high accuracy.

Findings

- **Data cleaning was critical** — preprocessing steps like rotation, cropping, and balancing of skin condition cases significantly improved model accuracy.
- **Medical-grade images made a major difference** — boosting image model accuracy from 0.22 to 0.84.
- **Selective metadata use** — focusing on the most impactful demographics (like symptoms, age group, and skin type) led to stronger model performance.
- **Model fusion boosted accuracy** — combining the metadata model (0.88) with the image model (0.84) resulted in a final ensemble accuracy of 0.91.

Limitations

- **Limited dataset size** — many skin conditions had too few labeled images to train reliable models.
- **Capped condition count** — only 8 skin conditions (plus normal) were included due to class imbalance and data scarcity.
- **Lack of real diagnostic data** — the reliance on prediction-labeled data (not biopsy-verified) may have impacted true diagnostic accuracy.
- **Need for more real-world diversity** — future work should include larger, more representative datasets for improved generalization.

References

1. *Mudasir A Ganaie, Minghui Hu, Ashwani Kumar Malik, Muhammad Tanveer, and Ponnuthurai N Suganthan. "Ensemble deep learning: A review". In: Engineering Applications of Artificial Intelligence 115 (2022), p. 105151 (cit. on pp. 8, 12, 17, 22–23, 69, 72).*
2. *Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. "The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions". In: Scientific data 5.1 (2018), pp. 1–9 (cit. on pp. 2, 4, 10, 22, 26, 30–31, 35, 44, 68).*