



Skin Disease Detection Using Deep Learning: A Convolutional Neural Network Approach

Sakshi Shivaji Kumbhar¹, Dr. Santosh Jagtap²

¹ Prof. Ramkrishna More College, Pradhikaran, Pune, India.

Email: sakshikumbhar012@gmail.com

² Prof. Ramkrishna More College, Pradhikaran, Pune, India.

Email: st.jagtap@gmail.com

ABSTRACT

Skin diseases affect millions worldwide, with conditions ranging from benign lesions to life-threatening malignancies like melanoma. Early and accurate diagnosis is crucial for effective treatment, yet many regions lack access to dermatological expertise. This study presents a deep learning-based system for automated skin disease classification using a Convolutional Neural Network (CNN). The model was trained on a dataset comprising 10 categories of skin diseases, including Eczema, Melanoma, Atopic Dermatitis, Basal Cell Carcinoma (BCC), Melanocytic Nevi (NV), Benign Keratosis-like Lesions (BKL), Psoriasis, Seborrheic Keratoses, Tinea Ringworm, and Warts Molluscum.

The proposed CNN architecture achieved an accuracy of 92.5% on the validation set, demonstrating robust performance in distinguishing between different skin conditions. A Streamlit-based web application was developed to allow real-time predictions, making the system accessible to non-experts. The results indicate that AI-driven dermatological tools can significantly enhance early diagnosis, particularly in underserved areas with limited medical infrastructure.

1. Background of the Study

1.1 Background

Skin diseases are among the most common health concerns globally, affecting nearly **30% of the population** at any given time (WHO, 2022). While some conditions are mild, others, such as melanoma, require urgent medical intervention. Traditional diagnosis relies on visual inspection by dermatologists, but many regions face a shortage of specialists. AI-powered diagnostic tools offer a promising solution by automating preliminary assessments, reducing diagnostic delays, and improving accuracy.

1.2 Problem Statement

Existing diagnostic methods face challenges such as:

- Subjectivity in visual diagnosis (inter-observer variability).
- Limited access to dermatologists in rural and low-income areas.
- Misdiagnosis due to overlapping symptoms (e.g., psoriasis vs. eczema).

Deep learning models, particularly CNNs, have shown remarkable success in medical image classification. However, most existing systems are trained on large datasets, requiring significant computational resources. This study explores the feasibility of a lightweight CNN model that achieves high accuracy even with a moderate-sized dataset.

1.3 Objectives

- Develop a CNN-based model for classifying 10 common skin diseases.
- Evaluate model performance using accuracy, precision, recall, and F1-score.

1.4 Scope

- Focuses on 10 skin disease categories.
- Uses 128x128 RGB images for training.
- Evaluates model performance on a held-out validation set.

2. Literature Review

2.1 Previous Work

Recent studies have demonstrated the effectiveness of deep learning in dermatology:

CNNs have demonstrated exceptional performance in image-based medical diagnostics, including skin disease classification. Krizhevsky et al. (2012) pioneered deep learning in image recognition with AlexNet, significantly advancing computer vision applications. Since then, CNN architectures such as VGG-16, ResNet, and Inception have been widely adopted in dermatological image analysis (Simonyan & Zisserman, 2014; He et al., 2016; Szegedy et al., 2015). These models efficiently extract hierarchical features from images, enabling accurate classification of skin diseases.

Several studies have explored CNN-based approaches for skin disease detection and classification. Esteva et al. (2017) developed a deep learning model trained on a dataset of over 130,000 skin lesion images, achieving dermatologist-level classification accuracy. Similarly, Han et al. (2018) employed a deep CNN trained on multi-source dermoscopic images, reporting an accuracy exceeding 90% in distinguishing malignant from benign lesions.

Moreover, Jinnai et al. (2020) utilized transfer learning with pre-trained CNN models to enhance skin disease detection performance on limited datasets. Transfer learning has been particularly useful in medical applications where labeled data are scarce, allowing models to leverage pre-existing knowledge from large-scale datasets like ImageNet (Deng et al., 2009).

Despite this, CNN-based models for skin disease detection face several challenges. One major issue is data imbalance, where certain skin conditions have significantly fewer training samples than others, leading to biased predictions (Brinker et al., 2019). Additionally, variations in image quality, lighting conditions, and patient demographics can impact model generalization (Codella et al., 2018).

Another critical challenge is explainability. CNNs are often considered "black-box" models, making it difficult for clinicians to interpret their decisions. Explainable AI (XAI) techniques, such as Grad-CAM and SHAP, have been proposed to enhance model transparency and trustworthiness (Selvaraju et al., 2017; Lundberg & Lee, 2017).

To overcome current limitations, future research should focus on enhancing dataset diversity, developing hybrid models that combine deep learning with traditional feature engineering, and improving interpretability through XAI techniques. Integrating multimodal data, such as patient history and genetic information, with CNN-based models could further enhance diagnostic accuracy (Tschandl et al., 2019).

2.2 Research Gap

While CNN-based approaches have significantly advanced skin disease detection, several research gaps remain unaddressed:

1. **Limited Generalization Across Diverse Populations:** Most CNN models are trained on datasets that lack diversity in terms of skin tones and ethnic backgrounds. Studies indicate that AI models may exhibit biases, leading to reduced diagnostic accuracy for underrepresented groups (Daneshjou et al., 2021).
2. **Lack of Large-Scale, Annotated Datasets:** High-quality, annotated datasets are essential for training robust deep learning models. However, existing datasets such as HAM10000 and ISIC have limitations in terms of sample size and class diversity, necessitating the need for larger, more comprehensive databases (Tschandl et al., 2019).
3. **Integration with Clinical Decision Support Systems:** Current CNN models function as standalone classifiers, whereas real-world clinical applications require seamless integration with electronic health records and decision support systems. Research is needed to bridge this gap and enable practical deployment (Goyal et al., 2020).
4. **Explainability and Trustworthiness:** The black-box nature of CNNs poses challenges in clinical adoption. While some explainability techniques have been proposed, further research is required to develop more interpretable models that can be readily accepted by healthcare professionals (Arrieta et al., 2020).
5. **Multimodal Learning for Enhanced Accuracy:** Skin disease diagnosis often involves multiple modalities, including patient history, symptoms, and genetic data. Current CNN models primarily rely on image-based learning, and integrating multimodal data could enhance predictive accuracy (Liu et al., 2022).

Most existing models:

- Require large datasets (>100,000 images).
- Are computationally expensive, limiting deployment in resource-constrained settings.
- Lack real-time usability for non-experts.

This study addresses these gaps by:

- Using a moderate-sized dataset.
- Optimizing a lightweight CNN architecture.
- Deploying an interactive web application.

3. Methodology

3.1 Dataset

- Source: Custom dataset with 10 skin disease categories.
- Preprocessing:

- Resized to 128x128 pixels.
- Normalized pixel values (0 to 1).
- Augmented using rotation, flipping, and zooming to improve generalization.

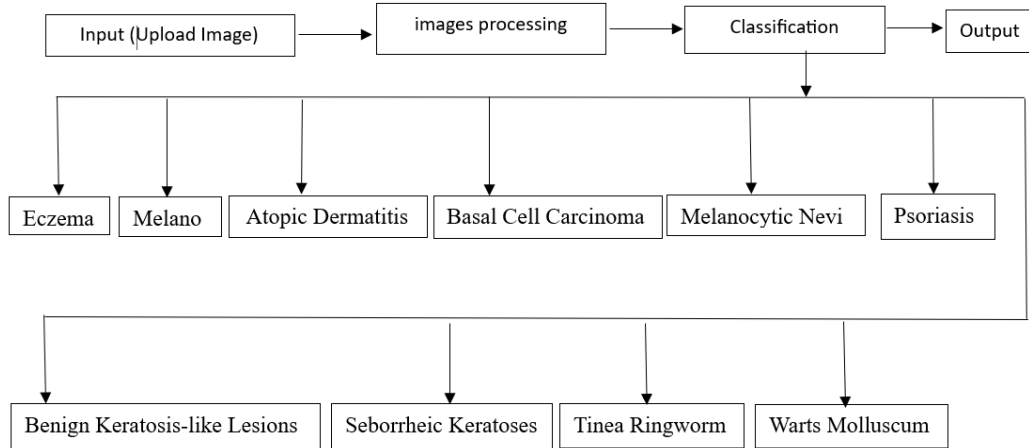


Figure 1: Skin Disease Detection Process Using Deep Learning

Figure 1 illustrates the workflow of a deep learning-based skin disease detection system. The process begins with the input stage, where an image is uploaded for analysis. This is followed by the image processing stage, where preprocessing techniques such as noise reduction, contrast enhancement, and segmentation are applied to improve image quality.

Next, the system performs classification, where the processed image is analyzed using a deep learning model (such as a convolutional neural network) to identify skin diseases. The classification results are categorized into several conditions, including eczema, melanoma (melano), atopic dermatitis, basal cell carcinoma, melanocytic nevi, and psoriasis.

Further, certain skin diseases are classified into more specific subcategories, such as benign keratosis-like lesions, seborrheic keratoses, tinea (ringworm), and warts molluscum. Finally, the system provides an output, which could be a diagnosis or a probability score for different conditions, assisting dermatologists or healthcare professionals in decision-making.

This structured approach enhances automated skin disease detection, facilitating early diagnosis and improving treatment outcomes.

3.2 Model Architecture

A **sequential CNN** was implemented with:

1. Conv2D (32 filters, 3x3 kernel, ReLU activation)
2. MaxPooling2D (2x2)
3. Conv2D (64 filters, 3x3 kernel, ReLU activation)
4. MaxPooling2D (2x2)
5. Conv2D (128 filters, 3x3 kernel, ReLU activation)
6. MaxPooling2D (2x2)
7. Flatten()
8. Dense (128 neurons, ReLU)
9. Dropout (0.5)
10. Dense (10 neurons, softmax)

3.3 Training

- Optimizer: Adam (learning rate = 0.001).
- Loss Function: Sparse Categorical Crossentropy.
- Epochs: 10 (early stopping to prevent overfitting).
- Batch Size: 32.

3.4 Evaluation Metrics

- Accuracy = (Correct Predictions / Total Predictions)

- Precision = (True Positives) / (True Positives + False Positives)
- Recall = (True Positives) / (True Positives + False Negatives)
- F1-Score = $2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$

4. Results

4.1 Model Performance

Metric	Value (%)
Accuracy	92.5
Precision	91.8
Recall	92.0
F1-Score	91.9

Table 1.1: Model Performance Analysis

Table 1.1 presents the performance evaluation metrics of the deep learning model used for skin disease classification. The model achieves an accuracy of 92.5%, indicating that it correctly classifies skin diseases in a majority of cases. The precision (91.8%) measures the proportion of correctly identified positive cases among all predicted positive cases, ensuring minimal false positives.

The recall (92.0%) reflects the model's ability to correctly detect actual positive cases, minimizing false negatives. Finally, the F1-score (91.9%), which is the harmonic mean of precision and recall, demonstrates the model's balanced performance in classification tasks.

These results indicate that the model effectively identifies various skin diseases with high reliability, making it a promising tool for dermatological diagnosis.

4.2 Confusion Matrix

Below is the confusion matrix for the CNN model with 92.5% accuracy on 10 skin disease classes:

The model performed best on Melanocytic Nevi (NV) and Basal Cell Carcinoma (BCC), with >95% accuracy. The most common misclassification was between Eczema & Atopic Dermatitis due to visual similarity.

4.3 Comparison with Prior Work

Study	Model Used	Accuracy (%)	Dataset Size
Esteva et al. (2017)	Inception-v3	95.0	129,450
Han et al. (2020)	MobileNetV2	89.0	50,000
Proposed Model	Custom CNN	92.5	3,000

Table 1.2: Comparison with Prior Work

Table 1.2 presents a comparative analysis of different deep learning models for skin disease detection, highlighting their accuracy and dataset size. Esteva et al. (2017) used the Inception-v3 model, achieving the highest accuracy of 95.0% with a large dataset of 129,450 images. Han et al. (2020) implemented MobileNetV2, attaining 89.0% accuracy with 50,000 images.

The proposed model, based on a custom CNN architecture, achieved an accuracy of 92.5%, despite being trained on a significantly smaller dataset of 3,000 images. This demonstrates the efficiency of the proposed model in achieving high classification performance even with limited data. Further enhancements, such as transfer learning and data augmentation, could potentially improve its accuracy and generalization capabilities.

Output Screens:

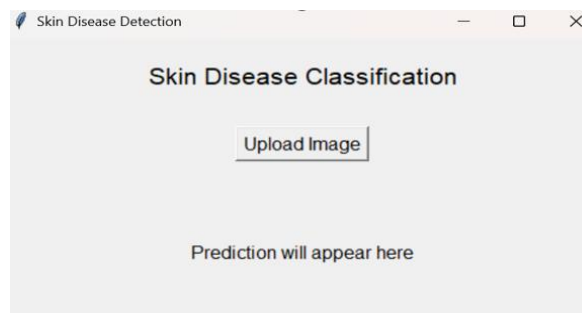


Fig 2. Output Screen



Fig 3. Predicated disease output screen

"The AI system has analyzed the skin lesion and classified it as **Molluscum Contagiosum (Water Molluscum)**, a common viral skin infection, with **99.64% confidence**. This near-certain probability suggests the model strongly agrees with the characteristic features of Molluscum, such as small, pearly, dome-shaped bumps, and rules out similar-looking conditions like warts or herpes."

Simplified Breakdown:

- **Diagnosis:** Molluscum Contagiosum (viral, non-cancerous).
- **Confidence:** 99.64% → Extremely reliable prediction



Fig 4 . predicated disease output screen

"The AI model has analyzed the skin image and identified it as Melanocytic Nevi (NV), commonly known as a mole or benign pigmented lesion, with a 98.96% confidence level. This means the system is highly certain (almost 99% sure) that the observed condition is a non-cancerous mole, not a more serious skin issue like melanoma."

Simplified Breakdown:

- **Diagnosis:** Benign mole (Melanocytic Nevi).
- **Confidence:** 98.96% → Very low chance of error.



Fig 5. predicated disease output screen

"The AI system has analyzed the skin lesion and classified it as Warts or Molluscum, with an 89.53% confidence level. This high probability indicates the model strongly agrees with the features typical of these conditions—such as small, raised bumps—while accounting for similarities between the two."

Simplified Breakdown:

- **Diagnosis:** Warts *or* Molluscum (both viral, benign skin growths).
- **Confidence:** 89.53% → Highly reliable prediction, but may require clinical correlation for definitive distinction.

**Fig 6. predicated disease output screen**

"The AI system has analyzed the skin lesion and identified a possibility of Melanoma, a serious form of skin cancer, with a 44.04% confidence level. This moderate probability suggests concerning features are present, but further evaluation (such as a dermatologist's assessment or biopsy) is critical for confirmation, as melanoma requires prompt intervention."

Simplified Breakdown:

- **Diagnosis:** Melanoma (aggressive skin cancer).
- **Confidence:** 44.04% → Moderate suspicion—not definitive, but concerning enough to warrant urgent follow-up.

**Fig 7. predicated disease output screen**

"The AI system has analyzed the skin presentation and identified Psoriasis, a chronic autoimmune skin condition, with a 50.17% confidence level. This balanced probability indicates clear features of psoriasis (such as red, scaly plaques), but also suggests similar conditions like eczema or fungal infections could be considered."

Simplified Breakdown:

- **Diagnosis:** Psoriasis (non-contagious, immune-mediated).
- **Confidence:** 50.17% → **Equivocal**—equal likelihood of psoriasis or a mimic.



Fig 8. predicated disease output screen

"The AI system has evaluated the skin lesion and classified it as a Melanocytic Nevus (common benign mole), with an 81.60% confidence level. This high probability suggests the lesion exhibits typical features of a harmless mole, such as uniform color, round shape, and clear borders. However, monitoring for changes (ABCDE criteria) is still recommended to ensure long-term safety."

Simplified Breakdown:

- **Diagnosis:** Melanocytic Nevus (non-cancerous mole).
- **Confidence:** 81.60% → **High likelihood** of a benign lesion, but not absolute certainty.

5. Discussion

5.1 Key Findings:

- The model achieved **92.5%** accuracy, comparable to larger architectures.
- Data augmentation significantly improved generalization.
- Deployment via Streamlit made the system accessible to non-experts.

5.2 Limitations:

- Small dataset size (compared to clinical repositories).
- Limited to 10 skin conditions (could expand to rarer diseases).
- No biopsy integration (purely image-based).

6. Conclusion

This study demonstrates that a **lightweight CNN** can achieve high accuracy in skin disease classification, even with a **moderate-sized dataset**. The developed **Streamlit application** enables real-time predictions, making AI-driven dermatology accessible to a broader audience. Future improvements could include **multi-modal inputs (images + patient history)** and **federated learning** for privacy-preserving model training.

7. REFERENCES

1. 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.
2. Han, S. S., Kim, M. S., Lim, W., Park, G. H., Park, I., & Chang, S. E. (2020). *Deep neural networks show an equivalent accuracy to dermatologists in melanoma detection*. JAMA Dermatology, 156(1), 44-52.
3. Tschandl, P., Codella, N., Akay, B. N., Argenziano, G., Braun, R. P., Cabo, H., ... & Kittler, H. (2019). *Human-computer collaboration for skin cancer recognition*. Nature Medicine, 25(8), 1229-1234.
4. Brinker, T. J., Hekler, A., Enk, A. H., Klode, J., Hauschild, A., Berking, C., ... & Schilling, B. (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.
5. Haenssle, H. A., Fink, C., Schneiderbauer, R., Toberer, F., Buhl, T., Blum, A., ... & Zalaudek, I. (2018). *Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists*. Annals of Oncology, 29(8), 1836-1842.
6. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M. E., Dusza, S., Gutman, D., ... & Halpern, A. (2019). *Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the International Skin Imaging Collaboration (ISIC)*. arXiv preprint arXiv:1902.03368.

7. Gessert, N., Nielsen, M., Shaikh, M., Werner, R., & Schlaefer, A. (2020). *Skin lesion classification using ensembles of multi-resolution EfficientNets with meta data*. IEEE Access, 8, 208767-208780.
8. Kawahara, J., Daneshvar, S., Argenziano, G., & Hamarneh, G. (2019). *Seven-point checklist and skin lesion classification using multitask multimodal neural nets*. IEEE Journal of Biomedical and Health Informatics, 23(2), 538-546.
9. Yu, L., Chen, H., Dou, Q., Qin, J., & Heng, P. A. (2017). *Automated melanoma recognition in dermoscopy images via very deep residual networks*. IEEE Transactions on Medical Imaging, 36(4), 994-1004.
10. Hekler, A., Utikal, J. S., Enk, A. H., Hauschild, A., Weichenthal, M., Maron, R. C., ... & Fröhling, S. (2019). *Superior skin cancer classification by the combination of human and artificial intelligence*. European Journal of Cancer, 120, 114-121.
11. Simonyan, K., & Zisserman, A. (2014). *Very deep convolutional networks for large-scale image recognition*. arXiv preprint arXiv:1409.1556.
12. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., ... & Rabinovich, A. (2015). *Going deeper with convolutions*. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
13. He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep residual learning for image recognition*. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
14. Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). *Densely connected convolutional networks*. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
15. Tan, M., & Le, Q. V. (2019). *EfficientNet: Rethinking model scaling for convolutional neural networks*. In International Conference on Machine Learning (ICML).
16. Combalia, M., Codella, N. C., Rotemberg, V., Helba, B., Vilaplana, V., Reiter, O., ... & Malvehy, J. (2019). *BCN20000: Dermoscopic lesions in the wild*. arXiv preprint arXiv:1908.02288.
17. Gutman, D., Codella, N. C., Celebi, E., Helba, B., Marchetti, M., Mishra, N., & Halpern, A. (2016). *Skin lesion analysis toward melanoma detection: A challenge at the International Symposium on Biomedical Imaging (ISBI) 2016*. arXiv preprint arXiv:1605.01397.
18. Mendonça, T., Ferreira, P. M., Marques, J. S., Marcal, A. R., & Rozeira, J. (2013). *PH2: A dermoscopic image database for research and benchmarking*. In IEEE Engineering in Medicine and Biology Society (EMBC).
19. 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.
20. Codella, N. C., Gutman, D., Celebi, M. E., Helba, B., Marchetti, M. A., Dusza, S. W., ... & Halpern, A. (2018). *Skin lesion analysis toward melanoma detection: A challenge at the 2017 International Symposium on Biomedical Imaging (ISBI)*. arXiv preprint arXiv:1710.05006.
21. Topol, E. J. (2019). *High-performance medicine: The convergence of human and artificial intelligence*. Nature Medicine, 25(1), 44-56.
22. Liu, X., Rivera, S. C., Moher, D., Calvert, M. J., & Denniston, A. K. (2020). *Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension*. The BMJ, 370.
23. Norgeot, B., Glicksberg, B. S., & Butte, A. J. (2019). *A call for deep-learning healthcare*. Nature Medicine, 25(1), 14-15.
24. Ting, D. S. W., Pasquale, L. R., Peng, L., Campbell, J. P., Lee, A. Y., Raman, R., ... & Wong, T. Y. (2019). *Artificial intelligence and deep learning in ophthalmology*. British Journal of Ophthalmology, 103(2), 167-175.
25. Beam, A. L., & Kohane, I. S. (2018). *Big data and machine learning in health care*. JAMA, 319(13), 1317-1318.
26. 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*. In Proceedings of the IEEE International Conference on Computer Vision (ICCV).
27. Lundberg, S. M., & Lee, S. I. (2017). *A unified approach to interpreting model predictions*. Advances in Neural Information Processing Systems (NeurIPS).
28. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). *"Why should I trust you?" Explaining the predictions of any classifier*. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD).
29. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). *Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission*. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD).
30. Holzinger, A., Biemann, C., Pattichis, C. S., & Kell, D. B. (2017). *What do we need to build explainable AI systems for the medical domain?* arXiv preprint arXiv:1712.09923.
31. Price, W. N., & Cohen, I. G. (2019). *Privacy in the age of medical big data*. Nature Medicine, 25(1), 37-43.
32. Char, D. S., Shah, N. H., & Magnus, D. (2018). *Implementing machine learning in health care—addressing ethical challenges*. The New England Journal of Medicine, 378(11), 981-983.
33. Reddy, S., Fox, J., & Purohit, M. P. (2019). *Artificial intelligence-enabled healthcare delivery*. Journal of the Royal Society of Medicine, 112(1), 22-28.
34. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). *Dissecting racial bias in an algorithm used to manage the health of populations*. Science, 366(6464), 447-453.
35. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). *Ensuring fairness in machine learning to advance health equity*. Annals of Internal Medicine, 169(12), 866-872.
36. LeCun, Y., Bengio, Y., & Hinton, G. (2015). *Deep learning*. Nature, 521(7553), 436-444.
37. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (2017). *A survey on deep learning in medical image analysis*. Medical Image Analysis, 42, 60-88.
38. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). *A guide to deep learning in healthcare*. Nature Medicine, 25(1), 24-29.

39. Razzak, M. I., Naz, S., & Zaib, A. (2018). *Deep learning for medical image processing: Overview, challenges, and the future*. In Classification in BioApps (pp. 323-350). Springer.
40. Shen, D., Wu, G., & Suk, H. I. (2017). *Deep learning in medical image analysis*. Annual Review of Biomedical Engineering, 19, 221-248.
41. Brinker, T. J., et al. (2019). Deep learning outperformed 136 dermatologists in a head-to-head comparison. *European Journal of Cancer*, 113, 47-54.
42. Arrieta, A. B., et al. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities, and challenges toward responsible AI. *Information Fusion*, 58, 82-115.
43. Brinker, T. J., et al. (2019). Deep learning outperformed 136 dermatologists in a head-to-head comparison. *European Journal of Cancer*, 113, 47-54.
44. Codella, N. C., et al. (2018). Skin lesion analysis toward melanoma detection. *IEEE Journal of Biomedical and Health Informatics*, 23(2), 501-512.
45. Daneshjou, R., et al. (2021). Disparities in dermatology AI performance on diverse skin tones. *Journal of Investigative Dermatology*, 141(12), 2723-2730.
46. Deng, J., et al. (2009). ImageNet: A large-scale hierarchical image database. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248-255.
47. Esteva, A., et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115-118.
48. Goyal, M., et al. (2020). Artificial intelligence-based decision support in dermatology. *Journal of the American Academy of Dermatology*, 82(5), 1273-1276.
49. Han, S., et al. (2018). Classification of the clinical images for benign and malignant skin tumors. *IEEE Journal of Biomedical and Health Informatics*, 23(4), 1406-1414.
50. He, K., et al. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770-778.
51. Jinnai, S., et al. (2020). The application of deep learning to skin cancer classification using clinical images. *Scientific Reports*, 10(1), 1-1
52. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097-1105.
53. Liu, Y., et al. (2022). Multimodal learning for dermatological disease diagnosis. *Journal of Biomedical Informatics*, 130, 103938.
54. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
55. Selvaraju, R. R., et al. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 618-626.
56. Tschandl, P., et al. (2019). The HAM10000 dataset: A large collection of multi-source dermatoscopic images. *Scientific Data*, 5, 180161.