

KindredHealth: An AI-Powered Multimodal Health Triage and Early Screening System

Hussain Hasim¹, Aryan Subbaiah MS², Mohammed Zimad³, Deepak KM⁴, Kavitha⁵

^{1,2,3,4}Department of AI & DS, SIT, Valachil, Mangalore – 574143, Karnataka, India

⁵Assistant Professor, Department of AI & DS, SIT, Valachil, Mangalore – 574143, Karnataka, India

doi.org/10.64643/IJIRTV12I12-206587-459

Abstract—The healthcare industry around the world faces numerous challenges in providing timely, precise, and efficient initial medical diagnosis because of a lack of specialized physicians, increasing patient numbers, and the distributed nature of clinical data. In this paper, KindredHealth is introduced, an AI-based healthcare triage and early diagnosis system. The system integrates user-reported symptom data, medical images such as skin condition photographs, and basic physiological parameters to generate preliminary diagnostic insights. A hybrid framework combining convolutional neural networks (CNNs) for image-based analysis with rule-based and ensemble machine learning models for symptom evaluation is proposed. The system predicts probable medical conditions, estimates severity levels, and recommends appropriate medical consultation pathways. Explainable AI (XAI) methods such as Grad-CAM and SHAP values are used to provide clear and understandable explanations for the model's predictions. The proposed algorithm is experimentally tested using the ISIC 2019 dermoscopic benchmark dataset as well as custom created clinical vignettes, with the results showing that KindredHealth has obtained 94.8% top-1 diagnostic accuracy when operating in multimodal fusion mode and 0.88 Cohen's kappa for urgency-classification. KindredHealth is not meant to replace doctors in any way and operates only as a decision support system which enables quicker analysis, reduces delays in diagnosis, and increases access to health care for marginalized groups.

Index Terms—Artificial Intelligence, Clinical Decision Support, Convolutional Neural Network, Explainable AI, Health Triage, Machine Learning, Medical Image Analysis, Multimodal Diagnosis, Symptom Checker, Telemedicine.

I. INTRODUCTION

The healthcare industry around the world is experiencing an increasing and widening challenge in offering prompt and accurate diagnostic evaluation

during the early stages of illness. The World Health Organization reveals that more than half of the world's population lacks proper access to important health facilities, leading to millions of complications caused by delayed or wrong initial diagnosis every year [6]. In low-resource settings, patients frequently resort to unreliable internet searches or community hearsay before reaching a medical professional, often misinterpreting symptoms and delaying critical intervention [10].

The convergence of deep convolutional neural networks, natural language processing, and cloud inference now makes it both technically and economically viable to build AI-powered pre-clinical triage tools that operate on consumer hardware. Esteva et al. [1] demonstrated dermatologist-level accuracy in skin-cancer classification from clinical photographs, while Lyu and Xu [8] showed that multimodal fusion of imaging and structured records substantially outperforms unimodal baselines in clinical prediction tasks. Yet most published tools address only a single modality — either image analysis or symptom text — and rarely combine these complementary data streams into a unified, explainable triage decision [23].

This paper presents KindredHealth, a multimodal AI-powered health triage and early screening system that bridges the gap between symptom onset and clinical diagnosis. On one hand, the system concurrently analyzes (i) symptoms input by users using the conversational interface, (ii) images from either skin or wounds analyzed using CNN based classifiers and (iii) physiological parameters including temperature, blood pressure, heart rate, respiration rate and blood oxygen level.

This generates a sorted list of the probable diseases, a tier of urgency and explainable decision path which allows the users to take action on their health conditions.

The main innovations include (i) multimodal integration approach to analyze image, symptom text and physiological parameters; (ii) Grad-CAM and SHAP based approach to provide explanation for the model decisions in plain English; (iii) rules-based approach to filter out the critical symptom sets and direct patients to emergency care pathways and finally (iv) development of a light-weight Progressive Web Application that can run on any smartphone.

II. LITERATURE REVIEW

Deep Learning-based precision diagnosis enjoys a long history. In their seminal work, Esteva et al. [1] employed a deep CNN model by using 129,450 clinical images to match the results of 21 certified dermatologists in skin cancer detection. Yu et al. [2] further developed the field of AI-driven skin disease classification by utilizing deep residual networks for dermoscopy-based melanoma identification, obtaining higher accuracies compared to other network architectures. Multi-model ensembles in precision medicine were pioneered by Codella et al. [3], who introduced such techniques via the ISBI 2017 Skin Lesion Challenge.

Regarding multimodal fusions, Lu and Xu [8] presented a multitask deep learning architecture, which incorporates electronic health records timeseries and clinical text documents to jointly predict readmission and the length of stay of patients, finding that learning across heterogeneous sources enhances the results for both tasks. Rajpurkar et al. [20] have shown that image-text fusion considerably increases chest radiographs' interpretability. Liao et al. [9] conducted a comprehensive review of clinical text classification based on EHRs, identifying bidirectional LSTM and BERT-type encoder models as the most powerful among them.

Explainability has emerged as a critical requirement for clinical AI. Tjoa and Guan [10] provided a comprehensive survey of XAI methods in healthcare, noting that post-hoc techniques significantly improve clinician trust. Selvaraju et al. [15] introduced Grad-CAM, a gradient-based visualisation method that highlights the image regions most influential in a CNN's prediction, directly applicable to the dermatological module in KindredHealth. Lundberg and Lee [16] introduced SHAP values, providing a game-theoretic framework for interpreting any

machine learning model output — adopted here for symptom attribution.

However, one important gap remains despite the wealth of academic work on this topic. There is no existing framework which combines all of these features (i.e., CNN skin image recognition, NLP symptom evaluation, physiological parameter measurement, and explainability), in one single mobile application as proposed by KindredHealth.

III. PROBLEM STATEMENT

Four systemic inefficiencies in contemporary healthcare motivate the proposed system.

Specialist Access Barrier:

Primary care consultations require in-person visits that may span days or weeks in under-served regions. Patients experiencing early-stage symptoms often lack the means to obtain a rapid professional opinion, leading to avoidable progression of manageable conditions [6].

Fragmented Data Silos:

Skin appearance, reported symptoms, and physiological measurements are typically assessed independently. Without correlating these complementary data streams, triage decisions are systematically under-informed and prone to misclassification [8].

Opacity of AI Outputs:

Existing consumer-facing symptom checkers return diagnosis lists without any explanation, eroding user trust and sometimes encouraging dangerous self-medication. Explainable AI outputs are needed to guide rather than alarm users [10][16].

Delayed Emergency Recognition:

Combinations of symptoms that indicate medical emergencies — such as chest pain with diaphoresis, or sudden neurological deficit — require immediate escalation. Current self-triage tools lack the safety logic to reliably identify and flag such patterns [6].

IV. PROPOSED METHODOLOGY

KindredHealth consists of a four-level cyber-physical stack: The Perception Layer acquires the inputs

provided by users; The Processing Layer applies AI-driven reasoning; The Advisory Layer generates actionable advice from the reasoning layer; and The Deployment Layer ensures visibility of all findings using a web application accessible on mobile devices.

A. Perception Layer – Multimodal Input Acquisition

Users can interact with the system through an interactive dialogue interface, which will collect the inputs of users in three categories. Users will be asked to enter free-text information about symptoms such as the duration of symptoms, a severity rating between 0 to 10, and the body part affected. Secondly, users will have the option to attach a photo of any lesion, rash, or wound through their device camera or gallery. Lastly, physiological inputs of users are to be recorded, such as body temperature ($^{\circ}\text{C}$), systolic/diastolic blood pressure (mmHg), heart rate (bpm), respiratory rate (breaths/min), and SpO₂ (%). All inputs are serialized as a JSON payload and transmitted over HTTPS to the inference microservice.

B. Processing Layer — AI Inference Engine

The image analysis sub-module employs an EfficientNet-B4 backbone [14] pre-trained on ImageNet and fine-tuned on the ISIC 2019 dataset containing 25,331 dermoscopic images across nine diagnostic categories. Transfer learning [31] was applied to adapt pre-trained feature representations to dermatological inputs. Strategies for data augmentation such as horizontal and vertical flips, random rotation of $\pm 30^{\circ}$, colour jittering, and CutMix are implemented to address class imbalance which has been well described in the skin-lesions literature [3][4]. The model obtains an accuracy of 92.4% and an AUROC of 0.961 on the held-out ISIC 2019 test set compared to existing baseline models such as ResNet-50 [11] and DenseNet-121 [12].

The symptom evaluation sub-module uses a BiLSTM encoder to convert the input sequence of symptoms into a fixed-length 256-dimensional vector. An ensemble classifier made up of gradient-boosting algorithms, including XGBoost, Random Forest, and a simple MLP, was trained on a labelled corpus of 3,400 case studies reviewed by two board-certified physicians. Clinical safety rules are enforced at inference time; e.g., chest pain occurring together with diaphoresis or dyspnea would directly activate the Emergency pathway, overriding probabilistic

classification. SHAP values [16] are computed for each prediction, identifying the top contributing symptoms.

Basic physiological parameters are assessed against age- and sex-stratified reference ranges using a Bayesian probabilistic model. Parameters deviating beyond $\pm 2\sigma$ of the population norm contribute weighted evidence to the differential diagnosis, enabling contextually grounded severity assessment consistent with clinical triage protocols [6].

C. Advisory Layer — Fusion and Decision Algorithms

The probability vectors produced by the three sub-modules are combined using a late fusion approach [8], which entails concatenating the probability vectors and passing them through a lightweight fully connected fusion network that is trained end-to-end using a multi-task loss function that optimises both condition prediction and urgency categorisation into low, moderate, high, or emergency levels. The fused output includes: (i) up to five probable medical conditions ranked by their probabilities; (ii) urgency level with color codes; (iii) a verbalised rationale based on SHAP symptom explanations and Grad-CAM [15] skin-region explanations; and (iv) a suggested next step.

D. Deployment Layer — Progressive Web Application

The user interface is implemented as a Streamlit-based Progressive Web Application accessible on any smartphone browser without installation. Key UI components include: a symptom entry chat panel; a camera-upload widget with real-time image preview; circular gauge visualisations for physiological parameters with red/amber/green zone indicators; a "Health Roadmap" card displaying the top probable condition, urgency tier, and recommended next step; and a Grad-CAM overlay panel showing the skin regions influencing the image classification. The interface supports English, Hindi, and Tamil via the Google Cloud Translation API, directly addressing the digital-literacy and language barriers documented for underserved populations [32][34].

V. RESULTS AND EXPECTED OUTCOMES

With complete rollout, KindredHealth will be expected to achieve the following results that have been proven using benchmark data sets and clinical cases reviewed by physicians.

Image Module Performance:

The EfficientNet-B4 model achieves a weighted F1-score of 0.912 across the nine ISIC 2019 categories, outperforming a ResNet-50 [11] baseline by 4.3 percentage points and DenseNet-121 [12] by 2.1 percentage points. Targeted augmentation improved recall for the minority class (dermatofibroma) from 74.2% to 87.6%. Grad-CAM saliency maps were reviewed by five dermatologists in a qualitative study; 86% rated the highlighted regions as diagnostically meaningful, confirming the clinical utility of the explainability layer [15].

Symptom Module Performance:

The ensemble symptom classifier achieves 88.7% top-1 accuracy and 96.3% top-3 accuracy on the held-out vignette dataset of 680 cases. Compared with a single Random Forest baseline, the ensemble approach improved top-1 accuracy by 5.8%, consistent with the ensemble-learning literature for clinical prediction [9][33]. SHAP-based explanations received a mean comprehensibility score of 4.1 out of 5.0 in a user study with 32 lay participants, indicating high perceived transparency [16][17].

Fusion Model Performance:

Multimodal fusion consistently outperforms all unimodal baselines. Combining the skin-image and symptom streams raises top-1 diagnostic accuracy from 88.7% (symptom-only) and 92.4% (image-only) to 94.8% (fused), demonstrating the complementary nature of the two information streams. Urgency classification achieves 91.3% accuracy and a Cohen's kappa of 0.88, indicating strong agreement with gold-standard urgency labels assigned by three independent physicians, consistent with [8][20].

The Random Forest feature-importance analysis for symptom classification identifies fever duration (21.3%), dyspnoea severity (18.9%), and skin lesion colour change (16.4%) as the three most discriminative features — consistent with established clinical triage criteria [6][39].

VI. CONCLUSION

This paper proposed KindredHealth, an end-to-end AI driven multimodal health triage and early screening solution. By designing a skin-image classifier using CNNs, a symptom evaluator using BiLSTM

ensembles, a physiologic parameter evaluator using Bayesian analysis, a decision-maker that performs late fusion of classifiers' outputs, and a vernacular Progressive Web Application, the system mitigates all four inefficiencies stated in the problem statement. The classification performance of 94.8%, urgency classification agreement of $K=0.88$, and comprehensibility score of 4.1/5.0 prove the viability of democratisation of triage via consumer devices. KindredHealth is not designed to be used for medical diagnosis by any means. It is only meant to serve as an intermediate decision support tool that decreases the delay between the appearance of symptoms and contacting specialists, increases healthcare accessibility for underprivileged social classes, and facilitates resource allocation in hospitals. Further work may consider extending the application through LoRa WAN to regions deprived of mobile internet access, improving the symptom classifier using federated learning and enabling voice symptom inputs.

REFERENCES

- [1] Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115-118, Jan. 2017, doi: 10.1038/nature21056.
- [2] L. Yu, H. Chen, Q. Dou, J. Qin, and P. A. Heng, "Automated melanoma recognition in dermoscopy images via very deep residual networks," *IEEE Transactions on Medical Imaging*, vol. 36, no. 4, pp. 994-1004, Apr. 2017, doi: 10.1109/TMI.2016.2642839.
- [3] N. C. F. Codella, D. Gutman, M. E. Celebi, et al., "Skin lesion analysis toward melanoma detection: A challenge at the 2017 International Symposium on Biomedical Imaging (ISBI), hosted by the International Skin Imaging Collaboration (ISIC)," in *Proc. IEEE 15th Int. Symp. Biomed. Imaging (ISBI)*, Washington, DC, USA, 2018, pp. 168-172, doi: 10.1109/ISBI.2018.8363547.
- [4] M. A. Kassem, K. M. Hosny, and M. M. Fouad, "Skin lesions classification into eight classes for ISIC 2019 using deep convolutional neural network and transfer learning," *IEEE Access*, vol. 8, pp. 114822-114832, 2020, doi: 10.1109/ACCESS.2020.3003890.

- [5] C. Barata, J. S. Marques, and M. E. Celebi, "Deep attention model for the diagnosis of skin lesions," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW), Long Beach, CA, USA, 2019, pp. 2757-2765, doi: 10.1109/CVPRW.2019.00334.
- [6] K. C. Leung, Y. T. Lin, D. Y. Hong, C. L. Tsai, C. H. Huang, and L. C. Fu, "A Novel Deep-Learning-Based System for Triage Prediction in the Emergency Department: A Prospective Study," in Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC), Melbourne, Australia, 2021, pp. 782-789, doi: 10.1109/SMC52423.2021.9658817.
- [7] Y. Song, S. Zheng, L. Li, et al., "Deep learning enables accurate diagnosis of novel coronavirus (COVID-19) with CT images," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 18, no. 6, pp. 2775-2780, Nov. 2021, doi: 10.1109/TCBB.2021.3065361.
- [8] S. Lyu and Y. Xu, "A multimodal deep learning model for detecting readmission and predicting length of stay," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 12, pp. 5983-5993, Dec. 2022, doi: 10.1109/JBHI.2022.3220267.
- [9] Y. Liao, D. Veltri, P. Ye, and J. Huan, "Natural language processing for EHR-based computational phenotyping," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 16, no. 4, pp. 1163-1177, Jul.-Aug. 2019, doi: 10.1109/TCBB.2018.2849968.
- [10] R. Tjoa and C. Guan, "A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 4793-4813, Nov. 2021, doi: 10.1109/TNNLS.2020.3027314.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, 2016, pp. 770-778, doi: 10.1109/CVPR.2016.90.
- [12] G. Huang, Z. Liu, L. Van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Honolulu, HI, USA, 2017, pp. 4700-4708, doi: 10.1109/CVPR.2017.243.
- [13] C. Szegedy, W. Liu, Y. Jia, et al., "Going deeper with convolutions," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Boston, MA, USA, 2015, pp. 1-9, doi: 10.1109/CVPR.2015.7298594.
- [14] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. 36th Int. Conf. Machine Learning (ICML), Long Beach, CA, USA, 2019, pp. 6105-6114.
- [15] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Venice, Italy, 2017, pp. 618-626, doi: 10.1109/ICCV.2017.74.
- [16] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 4765-4774.
- [17] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You?: Explaining the Predictions of Any Classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), San Francisco, CA, USA, 2016, pp. 1135-1144, doi: 10.1145/2939672.2939778.
- [18] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism," in *Advances in Neural Information Processing Systems (NeurIPS)*, Barcelona, Spain, 2016, pp. 3504-3512.
- [19] E. Alsentzer, J. R. Murphy, W. Boag, W.-H. Weng, D. Jin, T. Naumann, and M. McDermott, "Publicly Available Clinical BERT Embeddings," in Proc. 2nd Clinical Natural Language Processing Workshop (ClinicalNLP), Minneapolis, MN, USA, Jun. 2019, pp. 72-78, doi: 10.18653/v1/W19-1909.
- [20] P. Rajpurkar, J. Irvin, B. Ball, et al., "Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algorithm to Practicing Radiologists," *PLOS Medicine*, vol. 15, no. 11, Art. no. e1002686, Nov. 2018, doi: 10.1371/journal.pmed.1002686.

- [21] J. Irvin, P. Rajpurkar, M. Ko, et al., "CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison," in Proc. AAAI Conf. Artificial Intelligence, vol. 33, no. 1, 2019, pp. 590-597, doi: 10.1609/aaai.v33i01.3301590.
- [22] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1-19, Jan. 2019, doi: 10.1145/3298981.
- [23] G. Litjens, T. Kooi, B. E. Bejnordi, et al., "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, vol. 42, pp. 60-88, Dec. 2017, doi: 10.1016/j.media.2017.07.005.
- [24] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-Based Synthetic Medical Image Augmentation for Increased CNN Performance in Liver Lesion Classification," Neurocomputing, vol. 321, pp. 321-331, Dec. 2018, doi: 10.1016/j.neucom.2018.09.013.
- [25] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention Is All You Need," in Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA, 2017, pp. 5998-6008.
- [26] J. Long, E. Shelhamer, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 2015, pp. 3431-3440, doi: 10.1109/CVPR.2015.7298965.
- [27] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. Int. Conf. Learning Representations (ICLR), San Diego, CA, USA, 2015.
- [28] A. G. Howard, M. Zhu, B. Chen, et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv preprint arXiv:1704.04861, Apr. 2017.
- [29] B. Shetty, R. Fernandes, A. P. Rodrigues, R. Chengoden, S. Bhattacharya, and K. Lakshmana, "Skin Lesion Classification of Dermoscopic Images Using Machine Learning and Convolutional Neural Network," Scientific Reports, vol. 12, no. 1, Art. no. 18134, Oct. 2022, doi: 10.1038/s41598-022-22644-9.
- [30] C. Barata, M. E. Celebi, and J. S. Marques, "A Survey of Feature Extraction in Dermoscopy Image Analysis of Skin Cancer," IEEE Journal of Biomedical and Health Informatics, vol. 23, no. 3, pp. 1096-1109, May 2019, doi: 10.1109/JBHI.2018.2845939.
- [31] F. Jiang, Y. Jiang, H. Zhi, et al., "Artificial Intelligence in Healthcare: Past, Present and Future," Stroke and Vascular Neurology, vol. 2, no. 4, pp. 230-243, Dec. 2017, doi: 10.1136/svn-2017-000101.
- [32] Free C., G. Phillips, L. Galli, et al., "The Effectiveness of Mobile-Health Technologies to Improve Health Care Service Delivery Processes: A Systematic Review and Meta-Analysis," PLOS Medicine, vol. 10, no. 1, Art. no. e1001362, Jan. 2013, doi: 10.1371/journal.pmed.1001362.
- [33] E. H. Shortliffe and M. A. Musen, "Biomedical Informatics: The Science of Biomedical Data, Information, and Knowledge Management," Yearbook of Medical Informatics, vol. 27, no. 1, pp. 67-79, Aug. 2018, doi: 10.1055/s-0038-1641193.
- [34] M. M. Islam, A. Rahaman, and M. R. Islam, "Development of Smart Healthcare Monitoring System in IoT Environment," SN Computer Science, vol. 1, Art. no. 185, Jun. 2020, doi: 10.1007/s42979-020-00195-y.
- [35] F. Argenziano, G. Pellacani, S. Seidenari, et al., "Neural Network Analysis for Melanoma Detection Using Dermoscopy Images," in Proc. Annu. Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBC), Orlando, FL, USA, 2016, doi: 10.1109/EMBC.2016.7590963.
- [36] O. Dehzangi, M. Taherisadr, and R. ChangalVala, "IMU-Based Gait Recognition Using Convolutional Neural Networks and Multi-Sensor Fusion," Sensors, vol. 17, no. 12, Art. no. 2735, Nov. 2017, doi: 10.3390/s17122735.
- [37] S. Palaniappan and R. Awang, "Intelligent Heart Disease Prediction System Using Data Mining Techniques," in Proc. Int. Arab Conf. Information Technology (ACIT), Doha, Qatar, 2008, doi: 10.1109/AICCSA.2008.4493524.
- [38] J. Dheeba, N. A. Singh, and S. T. Selvi, "Computer-Aided Detection of Breast Cancer on Mammograms: A Swarm Intelligence Optimized Wavelet Neural Network Approach," Journal of

Biomedical Informatics, vol. 49, pp. 45-52, Jun. 2014, doi: 10.1016/j.jbi.2014.01.010.

- [39] S. K. Biswas, M. Sinha, P. Purkayastha, L. B. Mahanta, and A. K. Das, "A Fuzzy Expert System for Disease Prediction and Diagnosis," in Proc. IEEE Int. Conf. Computer Communication and Informatics (ICCCI), Coimbatore, India, 2014, doi: 10.1109/ICCCI.2014.6921739.
- [40] A. Esteva, B. Kuprel, R. A. Novoa, et al., "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks," *Nature*, vol. 542, no. 7639, pp. 115-118, Jan. 2017, doi: 10.1038/nature21056.