Automated Detection and Classification of Multiple Diseases Using Deep Neural Networks

Ratna Priya Singh[1], Arnav Mazumdar[2], Arnav Agrawal[3], Sumit Kumar[4], Surbhi Vijh[5]
riyasingh21022003@gmail.com[1], forarnavm2004@gmail.com[2], arnavag31@gmail.com[3], sumitkumarbsr19@gmail.com[4], surbhivijh428@gmail.com[5]

[1]ASET, Amity University Noida, Uttar Pradesh, [2]ASET, Amity University Noida, Uttar Pradesh, [3]ASET, Amity University Noida, Uttar Pradesh, India
[4]ASET, Amity University Noida, Uttar Pradesh, [5]ASET, Amity University Noida, Uttar Pradesh, India

Abstract. Diseases are health problems which negatively impact the normal functioning of body and may involve one or more organs. This research targets three groups of diseases: pulmonary diseases from X-rays, neurological conditions like brain tumours from MRI scans, and skin diseases from dermoscopic images. Identifying these conditions accurately helps doctors choose the right treatment more effectively. We aim to create a smart model using deep learning methods to identify different diseases. Current methods rely on separate neural networks, where a dataset with multiple data types is given to different CNN architectures. Given a single CNN may not extract the right features, we propose a multi-modal approach using specific CNNs for each data type. Our plan involves three specialised CNN architectures: DenseNet-121 for chest X-ray analysis, ResNet-50 for brain MRI interpretation, and EfficientNet-B3 for skin lesion classification. Each CNN will be then combined for better accuracy in the disease detection process.

Keywords: Brain Disease, Skin Disease, Lung Disease

1. Introduction

The field of healthcare has always been among the first to take advantage of, and put in place, technologies that are cutting edge, and continually provides improvements to health that are necessary for human survival and improving quality of life, and an ability to provide a timely, accurate diagnosis and treatment. The medical field has made advances, but there still exist a number of diseases that have an early symptomatic phase and are challenging to correctly identify for a range of reasons, mainly due to the complexity of symptoms and symptomology, lack of trained professional personnel, and the limitations that traditional tools for diagnosing have by nature. The challenges become most pronounced in diseases that involve the brain, skin, and lungs; these diseases account for the most deaths and disabilities worldwide. Therefore, from an initial diagnostic perspective, early diagnosis is the most important, however, traditional medical diagnostic methodology or determining a diagnosis is fundamentally subjective, variable, and takes time [1], [2].

1.1 Brain Diseases:

IICTDS 2025, December 19-20, Chandigarh, India
Copyright © 2026 EAI
DOI 10.4108/eai.19-12-2025.2362548

Diagnosing brain tumors is one of the most difficult challenges in medicine. Tumors may be non-malignant (benign) or malignant (cancerous) but both types carry serious risks for neurologic injury. The most current and accepted method of diagnosis is Magnetic Resonance Imaging (MRI), which takes very detailed images of the brain tissues. However, reading and producing a report on MRI scans is a complex and time-consuming process even for highly trained radiologists, so even highly trained professionals can misinterpret a scan, which can lead to erroneous treatment plans and diminished survival rates. Alternatively, automated deep learning systems, especially those based on CNN architectures, have the potential to fairly accurately identify and delineate the boundaries of the tumor, which gives a better perspective to the pathologist, both in terms of accuracy and time [1], [5].

1.2 Skin Diseases:

Melanoma and skin cancer are two of the most common health problems. Despite being relatively uncommon, melanoma spreads quickly, making it the most aggressive type of skin cancer. Even the most experienced dermatologists still have difficulty distinguishing between benign lesions and malignant melanomas, although dermatoscopic imaging is one way to diagnose this condition. On the other hand, studies have shown that convolutional neural networks (CNN) trained on extensive dermatoscopic datasets are capable of achieving lesion-classification accuracy on par with dermatologists. These systems can be applied in care settings to reduce unnecessary biopsies and improve survival rates by identifying high-risk cases [6], [8].

1.3 Lung Diseases:

Among the leading causes of death worldwide are lung cancer, pneumonia, and tuberculosis. Typically, chest X-rays are considered the most common and effective ways to identify these diseases. But, in order to understand the pictures, one has to be specially trained and experienced. If the faint patterns like pulmonary nodules or opacities are wrongly classified, a delay in the treatments can ensue. It has been demonstrated that convolutional neural network (CNN) systems have made significant progress in accurately identifying lung abnormalities, which helps radiologists avoid needless work. Most notably, there are fewer false positives and false negatives, which increases the diagnosis's speed and accuracy [9], [10], and [11].

This research is about a way of doing things that uses basic ideas to help with diagnosis using artificial intelligence. The main thing it does is collect lots of information, which it gets from pictures. These pictures are from types of sicknesses. For example, the research uses MRI pictures of brain tumors, special eye pictures to detect problems, skin pictures to detect skin cancers and chest X-ray pictures to detect lung diseases [1] [3] [6] [10]. The research collects all these pictures, like the MRI pictures of brain tumors and the special eye pictures and the skin pictures and the chest X-ray pictures to help with diagnosis using artificial intelligence. Preprocessing steps including resizing, normalization and noise reduction are also an important step as they help ensure robustness and generalisation across an array of image types.

The model will use some architectures like VGG16, VGG19 ResNet and DenseNet. VGG networks are really good for transfer learning. Getting features from data when we do not have

a lot of it. This is because VGG networks can use the things they learned from datasets. ResNet is also good because it fixes the problem of vanishing gradients with its connections. This means ResNet can learn things accurately across all its layers. DenseNet is similar to ResNet. It also reuses the features it learned across layers. This makes DenseNet more efficient. It needs fewer parameters. The VGG networks and these other models like ResNet and DenseNet are all useful, for the model. Combining all of these structures serves as a strong framework for tackling different types of complex medical imaging tasks [14].

Training is supervised learning, usually with transfer learning from ImageNet pre-trained models. Optimizers (e.g., Adam) and loss functions (e.g., cross-entropy loss) are used to make the training process stable. Overfitting is controlled by frequently validating the data, and hyperparameters are adjusted to get better performance. Metrics measuring performance include not just accuracy, but also precision, recall, F1-score, confusion matrices, and ROC-AUC, allowing performance to align with true clinical practice when there are risks of false negatives [13].

An important addition to the methodology is the application of an explanation tool such as Gradient-weighted Class Activation Mapping (Grad-CAM) [12]. These techniques permit visualization of the portions of an image that contributed the most to decision-making of the algorithm, allowing clinicians to see inside the "black box" to make the clinical reasoning explicit. Providing real-time visual feedback on the prediction process supports clinician trust and allows for the clinical uptake of CNNs from "black box" to explainable decision support tools.

2. Literature Review

The rapid developments in the field of artificial intelligence, and especially deep learning, have profoundly influenced medical diagnosis since their application allows the analysis of multiscale and multimodal data in a highly automatic way. In the last decade, researchers have developed various AI methods for disease detection and classification by focusing on several medical imaging domains including MRI, CT scans, X-ray images, and dermoscopic images[15]. These studies vary significantly in disease condition, model structure, data types and assessing measures. Although numerous works claim encouraging results, due to the variety of methods, it is difficult to reach unified conclusions over the effective design choices. Thus, there is a need for some review or synthesis of existing material in structured manner to analyze the research conducted so far systematically, leading to common trends and limitations. In this work we aim to investigate the theme and classify previous works according to such methodological aspects, so one can have a better understanding of how deep learning is being applied in medical disease detection.

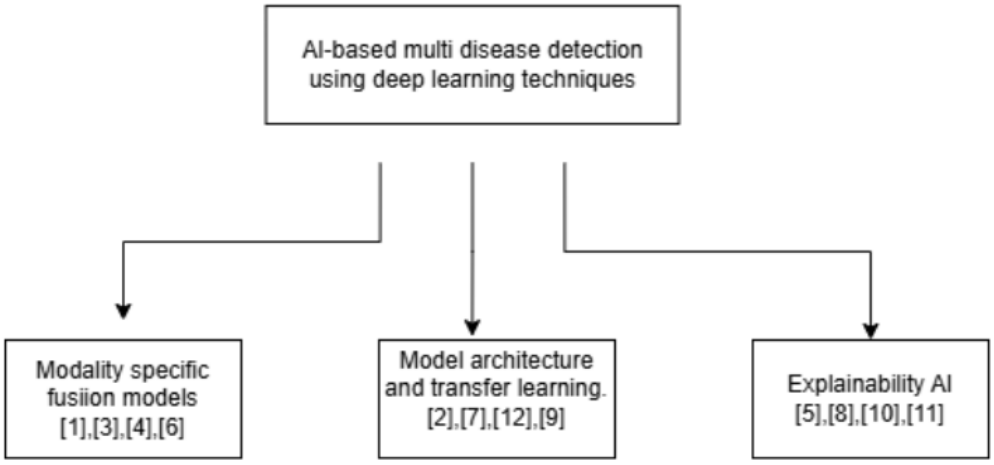


Fig. 1. Taxonomy Tree Of Literature Review

2.1 Disease-Specific Deep Learning Applications

Some early work concentrated on using deep learning models to identify certain diseases from medical images, validating the value of AI in healthcare diagnosis. Deepak and Ameer (2019) also worked on brain tumor classification and indicated that transfer learning combined with pretrained convolutional neural networks can detect useful patterns from MRI images while demonstrating the superior of the pretrained models in a context where medical data is scarce[1]. Abràmoff (2018) constructed a fully automated AI-based system for diagnosis of diabetic retinopathy and clinically validated it to prove that AI systems can reliably help with initial primary healthcare screening without human supervision[3]. Similarly, Gulshan (2016) trained deep CNNs with a large-scale retinal fundus image dataset and obtained high diagnostic accuracy, indicating the clinical applicability of using deep learning to identify DR[4]. Esteva (2017) introduced deep neural networks for dermoscopic image-based skin cancer classification and compared the diagnostic performance of their model with dermatologists, and demonstrated that a CNN approach can achieve expert-level diagnostic performance. Taken together, these studies demonstrate that deep learning achieves high performance in disease-specific medical image diagnosis for individual diseases[6].

2.2 Model Architectures and Transfer Learning Approaches

Another important trend at current research was the use of deep learning architectures and transfer learning in medical image classification. He (2016) proposed the ResNet architecture with short cut connections, which facilitate the training of much deeper networks and has been widely adopted in medical imaging applications[2]. Huang (2017) designed DenseNet, which can better reuse features and flow gradients by connecting each layer to all subsequent layers, thus is suitable for this situation with medical datasets having small-scaled samples[7]. Published by Simonyan and Zisserman, the VGG network is notably simple yet it was considered a common choice for transfer learning for early medical image analysis (Simonyan and Zisserman 2014)[12]. Kermany (2018) and they used deep learning models to detect a

variety of diseases with images and proved that pretrained CNN models are generalizable across different disease categories. These studies demonstrate that architectural innovations and transfer learning are crucial to improving diagnostic performance as well as reducing the complexity of training for healthcare applications[9].

2.3 Explainability, Preprocessing, and Evaluation in Medical AI

When it comes to using AI in medicine, it's not just about how well the model works. Explainability, how you prep the data, and how you actually test the model matter a lot too. Take Pratt and his team back in 2016—they used CNNs to spot diabetic retinopathy, but they ran into problems like class imbalance and stressed that you need solid pre-processing to trust the results[5]. Then there's Brinker (2019). They put deep learning head-to-head with dermatologists for diagnosing skin cancer and really dug into how these models hold up in real clinics, not just on paper[8]. Wang's group dropped the ChestX-ray8 dataset in 2017, which gave researchers a huge, standardized set of chest X-rays to really see how well models catch thoracic diseases, even when the data's only weakly labeled[10]. And Lyu (2020) took it a step further—they built a deep learning system with attention mechanisms so you can actually see which parts of a lung CT scan the model is looking at. That makes it easier for doctors to trust what the AI is doing. All this work shows one thing: if you want AI diagnostics that doctors will actually use, you need explainability, good datasets, and the right ways to measure performance[11].

3. Proposed Methodology

Given the data from the literature review, this work outlines the framework of deep learning methods for AI-driven multi-disease identification. The implemented approach essentially combines the major themes that were highlighted in the paper, such as modality-aware data processing, deep architectures based on transfer learning, and the usage of explainability and usual evaluation practice into the final solution.The resulting procedural model does not revolve around a particular ailment or dataset but rather seeks to furnish a single, adaptable framework that is capable of handling different medical imaging modalities and classes of diseases. The system's overall workflow, as proposed, is explained briefly after being depicted in the architecture diagram with the major elements.

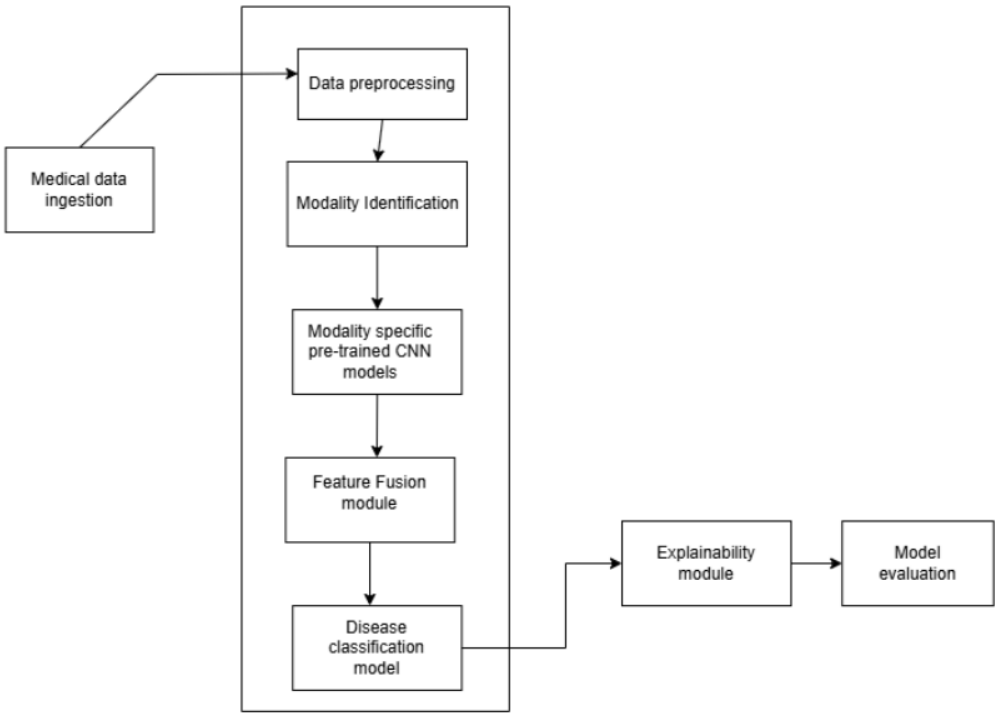


Fig. 2. Workflow of the proposed system.

3.1 Medical Data Input: This unit is a representation of getting medical pictures from various means such as MRI, CT scans, X-rays, and dermoscopic images, which are generally used in disease-specific diagnostic studies. Having the ability to work with different means allows the system to take care of a large number of medical problems.

3.2 Data Preprocessing: Preprocessing is the stage that guarantees uniformity and high quality of medical pictures through the use of several methods such as resizing, normalizing, noise elimination, and data augmentation. There are studies which stakeholders in the field say that highly effective preprocessing is money in the bank when given the issues of data variability, class imbalance, and imaging artifacts in medical datasets.

3.3 Modality-Specific Feature Extraction: This level processes each image source through a separate feature extraction path to unveil source-specific and disease-related visual features. The method of handling different modalities separately to extract relevant features has been referred to in the talked of the same kind of the deep learning research in the disease area previously .

3.4 Transfer Learning-Based Deep Backbone: The deep learning structures that have already been trained are utilized as one of the main components for feature extraction to get hold of the idea that has been grasped by the large-scale image datasets. Transfer learning imparts a medical imaging task with better output with less effort and only a few training instances particularly when there is a scarcity of labeled data .

3.5 Feature Fusion Module: The feature fusion component is instrumental in weaving together the features pulled out of different layers or sourced from different places into one representation. Fusion-based techniques, which are used to integrate complementary information, have been employed in various deep learning–based diagnostic systems to improve disease characterization .

3.6 Disease Classification Module: The integration of the features is just the first step. The resultant features are then handed over to a classification unit that theoretically can deal with multi-disease prediction. This kind of classification tactics has been utilized in the previous deep learning different types of disease studies to detect the presence of a certain disease or its category from medical images .

3.7 Explainability Module: An explainability element is put in place so as to facilitate understanding the underlying reasons behind the model predictions. The current work in this field underlines that the visual explanations and attention mechanisms play a major role in gaining the clinician's trust and providing the medical decision-making support.

3.8 Evaluation Module: The evaluating unit measures the diagnostic effectiveness with the application of a set of clinical metrics like accuracy, sensitivity, specificity, precision, and recall. The research that led to these conclusions emphasize that only with careful evaluation and use of metrics that have real clinical value can AI-based healthcare systems be employed in the real world .

4. Findings

Deep learning models depicted in the reviewed studies to be strongly capable of identifying individual diseases such as brain tumors, diabetic retinopathy, and skin cancer through medical images. In particular, pretrained CNNs turn out to be very potent in disease-specific scenarios, notably when there is a scarcity of labeled data. As shown in **Figure 3**, most of the studies have their focus on single diseases and hence are not able to generalize multiple conditions.

Studies reveal that the utilization of deep learning architectures of advanced level and transfer learning plays a very significant role in improving the performance of medical image classification (see **Figure 4**) . The models like residual and densely connected networks help to enhance feature learning as well as training stability. But still, the majority of the works are on general-purpose vision architectures, which are not specifically designed for healthcare.

The literature suggesting that for preprocessing, explainability, and evaluation clinical adoption that is reliable depends on these factors is very strong. The studies emphasize that among challenges are data imbalance, weak supervision, and inconsistent evaluation practices. Even if explainability techniques contribute to transparency, they are usually not at a level where they are completely merged into the system design.

The scores represented in the figures are qualitative and based on our literature review, so they are not experimental measurements.

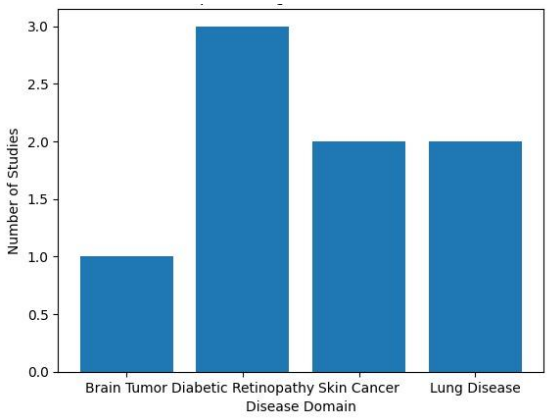


Fig. 3. Distribution of reviewed deep learning studies across disease domains

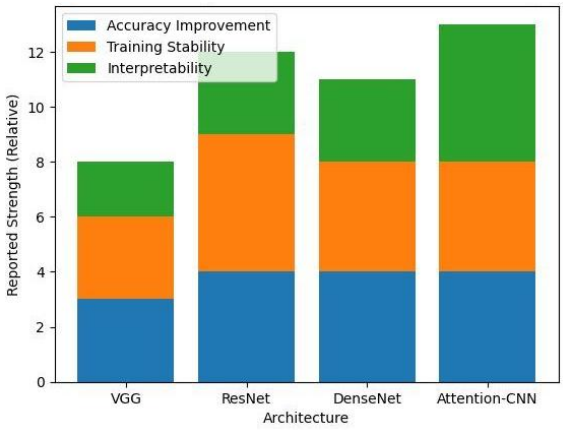


Fig. 4. Reported advantages of deep learning architectures in medical imaging literature.

Table 1. Evidence of performance and contribution across reviewed studies.

Study	**Domain**	**Model / Method**	**Key Performance Evidence Reported**
Deepak & Ameer [1]	Brain MRI	Transfer learning CNN	Showed improved accuracy over traditional ML for brain tumor classification

Gulshan et al [4]	Eye (DR)	Deep CNN	Achieved high diagnostic accuracy comparable to clinicians
Abràmoff et al. [3]	Eye (DR)	Autonomous DL system	Clinically validated in real-world screening
Esteva et al. [6]	Skin	Deep CNN	Achieved dermatologist-level classification performance
Wang et al. [10]	Lung X-ray	CNN on ChestX-ray8	Demonstrated large-scale disease classification feasibility
Brinker et al. [8]	Skin	Deep learning vs dermatologists	Model outperformed majority of clinicians
Lyu et al. [11]	Lung CT	Attention-based DL	Improved interpretability competitive performance

Table 1 shows a list of disease-specific deep learning applications across several medical domains, highlighting different approaches and their respective outcomes from a clinical perspective. It contains seven key studies along with their target domain, employed methodology, and reported performance outcome. Note that the distribution spans three primary medical imaging modalities: brain MRI analysis for tumor detection, dermoscopic imaging for skin cancer detection, and radiological chest imaging for respiratory disease identification. The evidence column demonstrates a consistent pattern of deep learning models achieving clinically significant performance levels, with several of them performing at par or better than those of trained medical specialists. These studies indicate a strong fundamental foundation confirming that deep learning technologies have matured sufficiently to support automated disease detection across multiple clinical domains. The diversity in evaluation approaches, however, highlights the need for standardized assessment methodologies for medical applications.

Table 2. Evidence supporting effectiveness of transfer learning in reviewed studies.

Study	**Used Transfer Learning**	**Evidence Reported**

Deepak & Ameer [1]	Yes	Pretrained CNN improved tumor classification accuracy
Kermany et al. [9]	Yes	Demonstrated strong cross-disease generalization
Esteva et al. [6]	Yes	Large pretrained CNN achieved expert-level results
Gulshan et al. [4]	Yes	Large-scale pretrained model improved DR detection
Simonyan & Zisserman [12]	Yes (VGG)	Established transfer learning suitability

Table 2 shows the performance of transfer learning across five key studies in medical imaging. The evidence demonstrates consistent benefits across different applications, including improved tumor classification, strong generalization potential, and expertise in diagnostic performance. This early adoption highlights transfer learning as an essential technique for addressing the limited availability of labeled medical data.

Table 3. Architectural advantages reported in reviewed literature.

Architecture	**Study**	**Reported Advantage**
VGG	Simonyan & Zisserman [12]	Simplicity and strong transfer learning capability
ResNet	He et al. [2]	Solves vanishing gradient, allows deep networks
DenseNet	Huang et al. [7]	Feature reuse, fewer parameters, better efficiency

Attention-based CNN	Lyu et al. [11]	Improves interpretability and performance
Hybrid DL systems	Kermany et al. [9]	Generalizes across disease categories

Table 3 shows the design advantages of deep learning models used in medical imaging applications. The table highlights five different architectures with their respective strengths. These architectural innovations together address critical challenges in medical image analysis, including training stability, computational efficiency, and clinical explainability.

5. Conclusion and Future Work

Deep learning has emerged as a powerful tool in medical image analysis, consistently outperforming traditional machine learning methods across multiple disease domains. CNN-based architectures, transfer learning, and attention mechanisms enable accurate feature extraction, improved interpretability, and strong performance even with limited data. Large-scale datasets have been the main drivers of advancements, thus, facilitating the development of scalable, generalizable multi-disease classification systems. However, the performance of automated diagnostic models that can sometimes be at par with or even better than that of experts, is still challenged by issues related to dataset diversity, interpretability, and clinical integration.

Future Work: In line with these discoveries, subsequent investigations could consider federated learning for safer joint training of models across different institutions, multimodal data utilization (linking imaging with clinical notes), and compact architectures that can be easily transported to and used in less developed healthcare sectors.

References

[1] S. Deepak and P. Ameer, “Brain tumor classification using deep CNN features via transfer learning,” *Computers in Biology and Medicine*, vol. 111, p. 103345, 2019.

[2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE CVPR*, pp. 770–778, 2016.

[3] M. Abràmoff, P. Lavin, M. Birch, N. Shah, and J. Folk, “Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices,” *NPJ Digital Medicine*, vol. 1, no. 1, pp. 1–8, 2018.

[4] V. Gulshan et al., “Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs,” *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016.

[5] A. Pratt, F. Coenen, and Y. Zheng, “Convolutional neural networks for diabetic retinopathy,” *Procedia Computer Science*, vol. 90, pp. 200–205, 2016.

[6] A. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.

[7] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proc. IEEE CVPR*, pp. 4700–4708, 2017.

[8] T. Brinker et al., “Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task,” *European Journal of Cancer*, vol. 113, pp. 47–54, 2019.

[9] D. Kermany et al., “Identifying medical diagnoses and treatable diseases by image-based deep learning,” *Cell*, vol. 172, no. 5, pp. 1122–1131.e9, 2018.

[10] X. Wang et al., “ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases,” in *Proc. IEEE CVPR*, pp. 2097–2106, 2017.

[11] Y. Lyu, F. Li, and H. Liang, “Attention-based deep learning for lung disease classification from CT images,” *IEEE Access*, vol. 8, pp. 119958–119967, 2020.

[12] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014

[13] Leibig, Christian, et al. "Leveraging uncertainty information from deep neural networks for disease detection." *Scientific reports* 7.1 (2017): 1-14.

[14] Mall, Shachi, and Jagendra Singh. "Heart diagnosis using deep neural network." *2023 International Conference on Computational Intelligence and Knowledge Economy (ICCIKE)*. IEEE, 2023.

[15] Al Reshan, Mana Saleh, et al. "A robust heart disease prediction system using hybrid deep neural networks." *IEEE Access* 11 (2023): 121574-121591.