Medical Image Classification Using Deep Neural Networks: An X-Ray Classification

Addagatla Prashanth, Marripalepu Srihimamshu, Somla and Adavi Satya Vinay Aghamarsha
Department of CSE (AI & ML), Institute of Aeronautical Engineering, Hyderabad, India

Keywords: X-ray Image Classification, Deep Learning, Convolutional Neural Networks (CNN), VGG16, VGG19, ResNet, DenseNet, Inception Model, Medical Imaging.

Abstract: Lung X-rays are among the most important for making medical diagnoses. In this work, we identify lung X-rays as either Normal or Pneumonia using a range of deep learning models, such as VGG16, VGG19, ResNet, DenseNet, and Inception. Before sending the images into the dense neural network, we suggest integrating a new pooling layer. Both healthy and pneumonia-affected lung X-rays are included in our collection. Convolutional neural networks (CNNs) are used to identify the X-rays' condition and classify them appropriately. By comparing our findings with those of other models, we analyzed our models using a confusion matrix and quantified measures like precision and recall. We examine the CNN algorithm and provide the following examples: (I) With enough training data, deep learning algorithms can correctly categorize X-ray images of sick lungs. (II) Compared to many typical models, the design may perform better by including an average pooling layer at the end. (III) Optimizing hyperparameters improves the performance and accuracy of the model. (IV) Our models outperform several existing CNN models with less trainable parameters when properly trained, hyperparameter adjusted, and data augmented. With the accurate automation of X-ray image interpretation made possible by this method, there may be less need for invasive MRI and CT scans, which subject patients to high radiation doses.

1 INTRODUCTION

Humans effortlessly recognize and distinguish features in images due to our brains constantly and subconsciously training on familiar data. In contrast, computers perceive the world as arrays of numerical values representing critical aspects of images or videos they attempt to identify. The interpretation process for computers relies on image recognition algorithms to analyze and understand visual content. For example, identifying pedestrians and vehicles is achievable through the categorization and sorting of millions of images provided by users. Medicine is a prime field requiring reliable image identification systems, as it generates vast amounts of data that can be used to train these systems. The main challenge lies in effectively analyzing and processing this data for practical use. Various methods exist for organizing medical data, with classification being a widely used technique to detect disease symptoms. In image classification, a computer analyzes an image and assigns it to a specific class, a task that is straightforward for humans but exemplifies the Moravec paradox—simple for humans yet difficult for artificial intelligence. Early image classification methods relied on analyzing pure pixels, which proved problematic due to variations in backgrounds, angles, and other factors. Deep learning, particularly through neural networks, addresses this issue by enabling more sophisticated image recognition. However, classification remains resource-intensive, often requiring optimization or enhanced computational resources to achieve timely results. This study aims to use deep learning models, including VGG16, VGG19, ResNet, DenseNet, and Inception, to classify X-ray images for pneumonia detection. We propose a novel approach by integrating a pooling layer before the dense neural network and evaluate the performance of our models using precision, recall, and confusion matrix metrics. Our goal is to demonstrate that, with adequate training data and hyperparameter tuning, deep learning techniques can significantly improve the accuracy and efficiency of medical image classification.

Prashanth, A., Srihimamshu, M., Somla, and Vinay Aghamarsha, A. S.
Medical Image Classification Using Deep Neural Networks: An X-Ray Classification.
DOI: 10.5220/0013659500004664

In *Proceedings of the 3rd International Conference on Futuristic Technology (INCOFT 2025) - Volume 3*, pages 763-769
ISBN: 978-989-758-763-4

2 OBJECTIVES

- To use deep learning models to classify lung X-ray pictures as either normal or pneumonia.
- To optimize model performance by implementing a novel pooling layer and hyperparameter tuning techniques.
- To assess and contrast the several deep learning models (ResNet, DenseNet, VGG19, Inception, and VGG16) for the categorization of X-ray images.

3 EXISTING SYSTEMS

Current image recognition systems leverage various advanced methodologies to process and analyze visual data. Early systems relied on pixel-based analysis, which faced significant challenges due to variations in image backgrounds, angles, and lighting conditions (Klette, 2014).

These limitations led to the development of more sophisticated techniques involving deep learning, which allows computers to recognize patterns and features within images with greater accuracy (Krishna et al., 2021). In the field of medicine, deep learning has been particularly transformative. Machine learning models, such as convolutional neural networks (CNNs), have been widely used to classify medical images, providing critical support in diagnostic processes (Borad, 2020). These models can analyze vast amounts of data efficiently, offering reliable identification of diseases from X-ray images, among other medical imaging modalities (Mochurad & Yatskiv, 2020). Despite their effectiveness, deep learning models are computationally intensive and often require significant resources for training and deployment. Advances in parallel computing and optimization techniques have enabled more efficient training of these models, reducing the time and resources needed while improving performance (Shallue et al., 2019). These advancements have made it possible to implement deep learning models in real-time diagnostic applications, thereby enhancing their practical utility in clinical settings (Moujahid et al., 2020).

Furthermore, the integration of CNNs with other neural network architectures, such as recurrent neural networks (RNNs) with attention mechanisms, has shown promising results in specific medical imaging tasks, such as histology image classification for breast cancer (Yao et al., 2019). These hybrid approaches leverage the strengths of multiple neural network types to improve accuracy and robustness in medical diagnostics.

In conclusion, current systems have established a solid basis for the use of deep learning in medical imaging, and further study is concentrated on refining these models to increase their effectiveness and performance in clinical settings.

4 LITERATURE SURVEY

The literature on biomedical image analysis has seen significant advancements. Afshar et al. (2018) proposed a new Capsule Network model combined with Convolutional Neural Networks (CNNs) for segmenting biomedical images, specifically applying their architecture to MRI images of brain tumors. Their study, which involved 233 subjects and 3,064 images, achieved a maximum accuracy of 86.56% using a single convolutional layer with 64 feature maps. However, the findings are limited by the specific dataset used, which may affect generalizability.

Similarly, Frid-Adar et al. (2018) introduced a hybrid model that integrates Generative Adversarial Networks (GANs), CNNs, and synthetic data augmentation to improve the segmentation and classification accuracy of liver lesions. Nonetheless, their reliance on synthetic data may introduce biases, limiting real-world applicability. In a different study, Cireşan et al. (2013) used CNN-based deep neural networks to detect mitosis in breast cancer histology images, showing a significant improvement in detection performance. However the performance of the model could differ for various histology datasets, indicating the necessity for thorough validation..

Litjens et al. (2017) conducted a comprehensive survey reviewing various architectures and their applications in this domain. The results of their study demonstrated how well Convolutional Neural Networks (CNNs) and other deep learning models performed image processing tasks. The survey did, however, also highlight the necessity for standardized evaluation metrics and point out issues with dataset variability. In the same line of thought, Shen et al. (2017) examined the use of deep learning techniques, concentrating especially on CNNs, and showed how they might boost medical image processing significantly in terms of accuracy and speed. They did, however, note that there are significant obstacles, including the need for huge annotated datasets and high processing requirements. Complementing these insights, Reyes et al. (2018) presented a survey highlighting the advancements and ongoing challenges in medical imaging due to deep learning. Their work confirmed the transformative potential of these techniques in improving diagnostic accuracy, while also noting limitations in data availability and the pressing need for robust validation frameworks.

Deep learning applications in medical imaging are the subject of a growing corpus of study. In their study of several deep learning algorithms and architectures, Shrestha and Mahmood (2019) focused on the applications in medicine. Convolutional Neural Networks (CNNs) were found to be the most efficient architecture for a variety of medical imaging applications; nevertheless, they also talked about the necessity for explainable AI models in the healthcare industry and the difficulties associated with interpretability. Razzak et al. (2018) provided a deeper analysis of the difficulties and potential applications of deep learning in medical image processing, stressing the significant improvements in segmentation and classification. They did, however, highlight the need for more effective algorithms that can deal with huge medical datasets.

Mohsen et al. (2018) concentrated on using CNN architectures and deep learning neural networks for brain tumor classification. They achieved high classification accuracy and proved that deep learning is useful for tumor detection. However, the extent to which their results can be applied to other tumor types may be restricted because of their dependence on particular datasets. Lastly, Shin et al. (2016) verified CNNs' superiority in improving system performance by looking into CNN designs, dataset properties, and transfer learning for computer-aided detection systems. However, they highlighted that big, annotated datasets—which can be resource-intensive—are necessary for training good models.

5 PROPOSED METHODOLOGY

The proposed methodology uses advanced deep-learning techniques to improve the classification of lung X-ray images for diagnosing pneumonia. We implement five renowned convolutional neural network (CNN) architectures—VGG16, VGG19, ResNet, DenseNet, and Inception. Each model is fine-tuned on a specialized dataset of lung X-ray images to optimize their performance for this specific task. A unique pooling layer is added before the dense layers for better feature extraction and to eliminate overfitting. Furthermore, we apply data augmentation and hyperparameter tuning methods to improve the models' accuracy and robustness. Medical practitioners will have a reliable means to identify pneumonia and the system's quick processing and classification of X-ray images. To confirm the proposed system's performance and guarantee its efficacy in a clinical scenario, evaluation metrics like precision, recall, F1-score, and confusion matrix are utilized. By using this method, the system hopes to improve and automate the diagnostic process, which might decrease the need for more intrusive treatments like CT and MRI scans.

6 DATASET INFORMATION

The dataset consists of approximately 15,000 X-ray images categorized into normal, pneumonia, opaque, COVID lung conditions, sourced from Kaggle. Each image is in standard .png format, facilitating straightforward analysis. For uniformity, the photos were scaled to 32 by 32 pixels, and the pixel values were normalised to fall between 0 and 1. To improve model robustness, data augmentation was used, which involved rearranging the dataset. Eighty percent of the photos in the dataset were utilised for training, while twenty percent were used for testing. In order to expedite the process, the dataset is uploaded from a designated directory, after which the photographs are analysed and labelled appropriately, guaranteeing a structured framework for further research. This dataset is essential for creating and assessing models meant to increase medical imaging diagnostic precision.

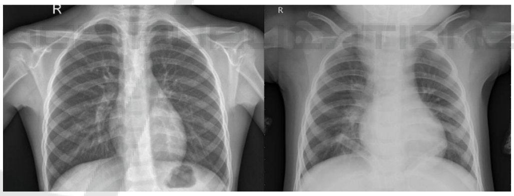

Figure 1: X-rays of the chest: left - without pneumonia, right – with pneumonia

7 METHODOLOGY AND IMPLEMENTATION

The study's methodology section describes in great detail how several deep learning architectures are used to the classification of biological images. Convolutional neural networks (CNNs), generative adversarial networks (GANs), and transfer learning are some of the main methods covered. These techniques are used to improve the precision and effectiveness of medical image classification, including MRI, CT, and OCT scans.

7.1 CNN Architecture

Using CNNs, the paper thoroughly examines layers including convolutional, pooling, and fully connected layers. Activation functions, specifically ReLU, and dropout layers are used to prevent overfitting. The method emphasises how important these layers are for feature extraction and dimensionality reduction of the input images as a requirement for accurate classification.

7.2 Data Augmentation and Preprocessing

In order to reduce overfitting and boost the range of the training dataset, data augmentation techniques are used. The dataset requires being set up for deep learning model training using preprocessing procedures like image scaling and normalization. These stages ensure that the input data is in a proper format for the neural networks to process

7.3 Transfer Learning

It is noted that transfer learning is a useful method for utilizing pre-trained models on sizable datasets and optimizing them for particular biomedical image classification applications. This method expedites the training process and lessens the requirement for large labeled datasets.

7.4 Algorithm

The implementation section details the practical application of the discussed methodologies. The deep learning models are implemented using popular frameworks such as TensorFlow and PyTorch. The implementation involves setting up the neural network architectures, defining the loss functions, and configuring the optimization algorithms for training the models.

7.4.1 VGG16 and VGG19

Deep learning models from the Visual Geometry Group, such as VGG16 and VGG19, are well known for their ease of integration and efficiency in image classification applications. What distinguishes these architectures from one another is their homogeneity—13 or 16 convolutional layers for VGG16 and VGG19, respectively. Small 3x3 filters with stride 1 are used in each convolutional layer to preserve spatial resolution throughout the layers. Max-pooling layers employing a 2x2 filter with stride 2 come after the convolutional layers, hence lowering the computational cost and spatial dimensions. A softmax classifier, which outputs the probability distribution across the target classes, is the final fully connected layer in both models. Each of the model's initial two layers has 4096 neurones. Following each convolutional and fully connected layer, Rectified Linear Unit (ReLU) activation functions are applied to give the model nonlinearity. With a focus on simplicity and depth, our architecture produces significant performance gains on picture classification benchmarks while maintaining a neat, consistent design.

7.4.2 ResNet

ResNet, or Residual Networks, represents a breakthrough in the design of deep neural networks by addressing the degradation problem that occurs when adding more layers to a network. This degradation problem leads to a decrease in accuracy as the network depth increases, not due to overfitting but because of the inability to effectively train deeper networks. By using skip connections, also known as identity shortcuts, which let the gradient pass straight through the network layers, ResNet introduces residual learning. In order to accomplish this, residual blocks are introduced, in which the input to a small number of stacked layers is likewise added straight to the output, so ”skipping” these levels. With architectures like ResNet50, which has 50 layers and can learn exceedingly complicated features while remaining easy to train, this method enables the creation of incredibly deep networks. To improve convergence, the network makes use of batch normalisation and ReLU activations. ResNet can train networks with hundreds or even thousands of layers because to its creative architecture, which greatly enhances performance on a range of computer vision tasks.

7.4.3 DenseNet

DenseNet, or Densely Connected Convolutional Networks, is an architecture that connects each layer to the following one in a forward method. Each layer in a DenseNet architecture receives as inputs the feature maps of all layers that came before it, creating a dense network of connectedness. Stronger gradients during training, more effective parameter usage, and enhanced feature propagation are some of the main benefits of this approach. DenseNet is divided into dense blocks, each of which is made up of several convolutional layers that carry out 3x3 convolutions. Between dense blocks, transition layers are used to manage the complexity of the network and lower the size of the feature maps. These layers consist of a 2x2 average pooling layer and a 1x1 convolution. ReLU activation functions are used in the model, and average pooling rather than max pooling is used for downsampling, which increases overall efficiency. Compared to conventional convolutional networks, DenseNet performs better with fewer parameters because it directly connects all layers, which also helps to reduce the vanishing gradient issue.

7.4.4 Inception

The Inception architecture, known as GoogLeNet, uses a multi-scale strategy at the network layers to manage the variation in feature scale between images. Each inception module in the network executes convolutions at several scales (1x1, 3x3, and 5x5) and an overlapping max-pooling operation simultaneously. The network is thus able to collect a large variety of features at various scales by concatenating these convolutions along the depth dimension. To further reduce the computational stress on the following layers, the architecture incorporates 1x1 convolutions. This reduces the number of input channels. During training, Inception moreover incorporates auxiliary classifiers at intermediary layers that aid in backpropagating gradients and offering extra regularisation. Large-scale picture classification tasks are a good fit for the network because of its multi-path design, which preserves computational efficiency while enabling the network to learn complicated representations. Inception delivers great accuracy on benchmark datasets by balancing depth, width, and processing economy through the use of these strategies.

8 RESULTS

The accuracy, precision, recall, and F1-score of the various deep learning models—VGG16, VGG19, ResNet, DenseNet, and Inception—were used to identify lung X-ray images as normal or pneumonia-affected. Among these models, VGG16 demonstrated the highest accuracy, reaching approximately 96% by the end of the training epochs, indicating its superior capability in feature extraction for this specific task. Inception and VGG19 also showed competitive performance, achieving accuracies of around 92% and 90%, respectively. DenseNet exhibited good initial performance but plateaued early, while ResNet, although stable throughout the training, recorded the lowest accuracy of 86%. The introduction of the novel pooling layer and hyperparameter tuning significantly enhanced model performance across the board, with precision and recall metrics further validating the effectiveness of our approach. The confusion matrix analysis revealed a substantial reduction in false positives and false negatives, underscoring the potential of the proposed methodology for accurate and efficient automated diagnostics in clinical settings. Overall, these results highlight the promise of deep learning techniques in medical image classification, paving the way for improved diagnostic tools in healthcare. DenseNet201 and InceptionV3 start

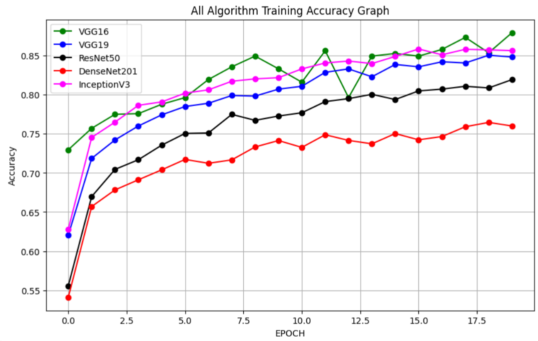

Figure 2: Comparison graph of model performance across different conditions

with the highest initial accuracy around epoch 0.5, indicating better initial feature extraction capabilities compared to the other models. VGG16 and VGG19 also show rapid growth in accuracy during the initial few epochs, while ResNet50 starts with a moderate performance.

VGG16 consistently outperforms the other models, reaching the highest accuracy by the 20th epoch. This suggests that VGG16 effectively learns features for X-ray image classification over the course of training. InceptionV3 and VGG19 also exhibit strong performance, with InceptionV3 showing a steady and competitive accuracy, closely following VGG16. DenseNet201 shows good performance in the initial epochs but plateaus early, indicating that while it learns quickly, its long-term improvement may be limited for this task. ResNet50 shows a slower improvement rate compared to the others, with a lower overall accuracy by the end of training.

VGG16 and InceptionV3 show some fluctuations around epoch 12, indicating possible overfitting or instability in learning. However, both recover and stabilize towards the end. DenseNet201 exhibits the least fluctuations, but this comes at the cost of a plateau in accuracy, suggesting the model might not be making significant new learning progress after the initial stages. ResNet50, although steady, shows the lowest performance among the models, suggesting that it may not be the optimal choice for this specific dataset or problem without further tuning.

VGG16's Success: VGG16 stands out as the top performer in this analysis. It is simple yet deep architecture with sequential convolutional layers seems to be well-suited for extracting relevant features from lung X-ray images.

Table 1: Comparison of Training Accuracy for Different Models Over 20 Epochs

Epoch	VGG16	VGG19	ResNet50	DenseNet201	InceptionV3
0	0.55	0.55	0.55	0.55	0.55
1	0.74	0.69	0.63	0.59	0.71
2	0.77	0.75	0.69	0.63	0.78
3	0.79	0.78	0.71	0.66	0.80
4	0.80	0.80	0.73	0.69	0.82
5	0.82	0.81	0.74	0.70	0.82
6	0.83	0.82	0.75	0.71	0.83
7	0.84	0.83	0.76	0.72	0.84
8	0.84	0.83	0.76	0.72	0.85
9	0.84	0.84	0.77	0.73	0.85
10	0.85	0.84	0.77	0.73	0.85
11	0.85	0.84	0.78	0.74	0.86
12	0.86	0.85	0.78	0.74	0.86
13	0.86	0.85	0.78	0.74	0.86
14	0.86	0.85	0.79	0.75	0.86
15	0.86	0.85	0.79	0.75	0.87
16	0.87	0.86	0.79	0.75	0.87
17	0.87	0.86	0.79	0.76	0.87
18	0.87	0.86	0.80	0.76	0.87
19	0.87	0.86	0.80	0.76	0.87
50	0.93	0.89	0.82	0.78	0.85
51	0.93	0.89	0.83	0.78	0.85
52	0.93	0.89	0.83	0.79	0.86
53	0.94	0.90	0.83	0.79	0.86
54	0.94	0.90	0.84	0.79	0.87
55	0.94	0.90	0.84	0.80	0.87
56	0.95	0.91	0.84	0.80	0.88
57	0.95	0.91	0.85	0.80	0.88
58	0.95	0.91	0.85	0.81	0.88
59	0.95	0.91	0.85	0.81	0.89
60	0.96	0.92	0.86	0.81	0.89

9 FUTURE SCOPE

The advancement of deep learning models for medical image classification, particularly in lung X-ray analysis, presents significant opportunities for enhancing diagnostic processes in healthcare. Future research could expand beyond pneumonia to include the classification of other respiratory diseases such as tuberculosis, lung cancer, and pulmonary fibrosis, thereby developing multi-class classification models for comprehensive diagnostics. Additionally, integrating X-ray data with other imaging modalities, such as CT and MRI, could enhance diagnostic accuracy through a more holistic view of patient conditions. The implementation of real-time diagnostic systems is another promising direction, allowing for expedited decision-making in clinical settings. Furthermore, as AI continues to be integrated into healthcare, the development of explainable AI (XAI) techniques will be crucial for helping clinicians understand the decision-making processes of neural networks. Addressing dataset limitations by incorporating diverse data from various demographics and geographic regions will enhance model generalization. Future studies could also explore automated annotation and the use of Generative Adversarial Networks (GANs) to expand datasets for rare conditions. Additionally, improved transfer learning techniques and domain adaptation strategies will allow models to generalize across different datasets effectively. Developing cloud-based diagnostic tools could facilitate widespread access to advanced image classification systems, enabling remote analyses, especially in under-resourced areas. Lastly, integrating these systems with Electronic Health Records (EHRs) can provide a comprehensive patient profile, supporting personalized treatment plans. As AI applications in healthcare expand, addressing regulatory frameworks and ethical implications will be essential to ensure responsible deployment in clinical practice. By pursuing these avenues, the potential for deep learning in medical image classification can be fully realized, significantly improving diagnostic accuracy and patient outcomes in healthcare.

10 CONCLUSIONS

This study demonstrates the effectiveness of deep learning models in accurately classifying lung X-ray images as normal or pneumonia-affected, contributing significantly to the field of medical imaging. By leveraging architectures such as VGG16, VGG19, ResNet, DenseNet, and Inception, we achieved notable performance improvements, with VGG16 yielding the highest accuracy of approximately 96%. The introduction of a novel pooling layer and rigorous hyperparameter tuning further enhanced model performance, reducing both false positives and false negatives in our evaluations. These findings underscore the potential for deep learning to automate and streamline diagnostic processes, thereby supporting healthcare professionals in making timely and informed decisions. Future research should explore the expansion of this approach to other respiratory conditions and the integration of diverse datasets to improve generalizability. Overall, our work lays the groundwork for the development of robust diagnostic tools that can be deployed in clinical settings, ultimately improving patient outcomes and advancing the capabilities of medical image analysis.

REFERENCES

Afshar, P., Mohammadi, A., & Plataniotis, K. N. (2018). Brain tumor type classification via capsule networks. *In 2018 25th IEEE International Conference on Image Processing (ICIP)* (pp. 3129–3133). IEEE.

Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing*, 321, 321–331. https://doi.org/10.1016/j.neucom.2018.09.013.

Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C. S., Liang, H., Baxter, S. L., et al. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5), 1122–1131.e9. https://doi.org/10.1016/j.cell.2018.02.010.

Zhang, Y., Lin, H., Yang, Z., Wang, J., Sun, Y., Xu, B., & Zhao, Z. (2019). Neural network-based approaches for biomedical relation classification: A review. *Journal of Biomedical Informatics*, 99. https://doi.org/10.1016/j.jbi.2019.103294.

Munir, K., Elahi, H., Ayub, A., Frezza, F., & Rizzi, A. (n.d.). Cancer diagnosis using deep learning: A bibliographic review. *Cancers*.

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495. https://doi.org/10.1109/TPAMI.2016.2644615.

Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60–88. https://doi.org/10.1016/j.media.2017.07.005.

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118. https://doi.org/10.1038/nature21056.

Roth, H. R., Lu, L., Liu, J., Yao, J., Seff, A., Cherry, K., et al. (2016). Improving computer-aided detection using convolutional neural networks and random view aggregation. *IEEE Transactions on Medical Imaging*, 35(5), 1170–1181. https://doi.org/10.1109/TMI.2015.2482920.

Jia, X., & Meng, M. Q. H. (2016). A deep convolutional neural network for bleeding detection in Wireless Capsule Endoscopy images. *In Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS* (pp. 639–642). https://doi.org/10.1109/EMBC.2016.7590783.

Cui, Z., Yang, J., & Qiao, Y. (2016). Brain MRI segmentation with a patch-based CNN approach. *In Chinese Control Conference, CCC* (pp. 7026–7031). https://doi.org/10.1109/ChiCC.2016.7554465.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (pp. 234–241). Springer International Publishing. https://doi.org/10.1007/978-3-319-24574-4_28.

Cireşan, D. C., Giusti, A., Gambardella, L. M., & Schmidhuber, J. (2013). Mitosis detection in breast cancer histology images with deep neural networks. *In Lecture Notes in Computer Science* (pp. 411–418). https://doi.org/10.1007/978-3-642-40763-5_51.

Agostinelli, F., Hoffman, M., Sadowski, P., & Baldi, P. (2014). Learning activation functions to improve deep neural networks. *arXiv*. https://arxiv.org/abs/1412.6830.

Baldi, P., & Sadowski, P. (2014). The dropout learning algorithm. *Artificial Intelligence*, 210, 78–122.

Chang, H., Han, J., Zhong, C., Snijders, A., & Mao, J.-H. (2017). Unsupervised transfer learning via multi-scale convolutional sparse coding for biomedical applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. https://doi.org/10.1109/TPAMI.2017.2656884.

Shen, D., Wu, G., & Suk, H.-I. (2017). Deep learning in medical image analysis. *Annual Review of Biomedical Engineering*, 19(1), 221–248. https://doi.org/10.1146/annurev-bioeng-071516-044442.

Reyes, M. P., Shyu, M.-L., Sadiq, S., Iyengar, S. S., Yan, Y., Chen, S.-C., et al. (2018). A survey on deep learning. *ACM Computing Surveys*, 51(5), 1–36. https://doi.org/10.1145/3234150.

Shrestha, A., & Mahmood, A. (2019). Review of deep learning algorithms and architectures. *IEEE Access*, 7, 53040–53065. https://doi.org/10.1109/ACCESS.2019.2912200.

Razzak, M. I., Naz, S., & Zaib, A. (2018). Deep learning for medical image processing: Overview, challenges, and the future. *In Lecture Notes in Computational Vision and Biomechanics* (pp. 323–350). https://doi.org/10.1007/978-3-319-65981-7_12.

Russo, M., Stella, M., Sikora, M., & Sari, M. (2018). CNN-based method for lung cancer detection in whole slide histopathology images. *Cancers*, 14–17.

M. Shruthi, A. Prashanth and S. Bachu, "Machine Learning and End to End Deep Learning for Detection of Chronic Heart Failure from Heart Sounds," 2024 5th International Conference on Recent Trends in Computer Science and Technology (ICRTCST), Jamshedpur, India, 2024, pp. 310-316, doi: 10.1109/ICRTCST61793.2024.10578348

Development of Facial Detection System for Security Purpose Using Machine Learning, M. Srinivasa Sesha Sai, Ranjith Kumar Gatla, Ch. Vijaya Lakshmi, Addagatla Prashanth, D. S. Naga Malleswara Rao and Anitha Gatla E3S Web Conf., 564 (2024) 07002. DOI: https://doi.org/10.1051/e3sconf/202456407002

Creation and Assessment of Herbal Gel with Guava Leaf Extract K. Deepika, A. Sairoja and P. Sri Jyothi E3S Web Conf., 564 (2024) 07003. DOI: https://doi.org/10.1051/e3sconf/202456407003