Enhanced Pneumonia Detection in Chest X-Rays Based on Integrated Denoising Autoencoders and Convolutional Neural Networks

Yufeng Xia[a]
School of Informatics, University of Edinburgh, Edinburgh, U.K.

Keywords: Computer Vison, Denoising, Classification, Pneumonia Detection.

Abstract: This research presents a new hybrid model that improves pneumonia detection from chest X-ray images by combining denoising autoencoders (DAEs) with convolutional neural networks (CNNs). The model concurrently performs image denoising and disease classification, leveraging both processes to enhance diagnostic accuracy. Preprocessing steps for the Chest X-Ray Images (Pneumonia) dataset included resizing to 150x150 pixels, image augmentation, and normalization to facilitate effective training. The integrated model architecture uses CNNs for feature extraction and classification, paired with DAEs for image denoising, all implemented using TensorFlow and optimized with the Adam optimizer on an NVIDIA RTX 4080 GPU. This setup allows dynamic adjustments of the learning rate, improving performance metrics. The model achieved a peak validation accuracy of 98.4% and demonstrated a substantial reduction in image noise, evidenced by a low Mean Squared Error (MSE) of 0.0049. These results highlight the model's capability to deliver precise classifications and superior image quality, thus enabling more reliable diagnoses. This study points to the potential for applying such integrated models more broadly in medical imaging, enhancing both interpretability and reliability of automated medical diagnostics. Future efforts will aim to extend this model's application to additional medical conditions and enhance its robustness and generalizability.

1 INTRODUCTION

Pneumonia, known for causing inflammation in the lung air sacs, poses a significant global health risk. Traditional ways to diagnose it rely heavily on radiologists reading chest X-rays, which can lead to errors and are quite expensive. (El-shafai et al., 2022) The rise of Machine Learning (ML) in medical imaging, especially in spotting and sorting diseases like pneumonia from chest X-ray pictures, brings a major shift in diagnosis due to their excellent performance in many domains (Li, 2024; Liu, 2023; Zhao, 2023). This change introduces a new healthcare era where AI-augmented diagnostics promise to overcome these longstanding hurdles.

While past research has explored various ML designs, like Convolutional Neural Networks (CNNs), for processing and classifying medical images (Liu, 2024; Lambert, 2024; Qiu, 2022), there's still a gap in making these models ready for real-world clinical use. A key challenge that's not fully tackled yet is improving image quality to boost model accuracy. Early work by El-shafai et al. (2022) and Thomas et al. (Thomas, 2022) highlights how denoising autoencoders could play a role in medical diagnostics. These studies point out the urgent need for further efforts to make models more reliable and adaptable to different datasets and conditions.

There's a pressing demand for new models that can do both noise reduction and accurate medical condition classification together. Previous methods mostly focused on one or the other. Merging these tasks could vastly improve diagnostics, cutting out the need for separate noise reduction and disease sorting steps. This would not only make the diagnostic process smoother but could also lower the computational resources needed by combining the tasks into one efficient model.

This study introduces a cutting-edge hybrid model that blends denoising and classification into one unified system. This innovative approach aims to make machine analysis easier to understand by producing clean images with precise diagnostic labels. By moving past the old division between focusing on

[a] https://orcid.org/0009-0008-5968-8774

Xia, Y.
Enhanced Pneumonia Detection in Chest X-Rays Based on Integrated Denoising Autoencoders and Convolutional Neural Networks.
DOI: 10.5220/0012973600004508

In *Proceedings of the 1st International Conference on Engineering Management, Information Technology and Intelligence (EMITI 2024)*, pages 799-803
ISBN: 978-989-758-713-9

noise reduction or classification, this combined method provides a deeper insight into the disease, helping medical professionals not just to rely on Artificial Intelligence (AI)'s diagnosis but also to review the high-quality images behind the AI's conclusions. The outcome is a model that predicts accurately and shares its results clearly, boosting trust and clarity in AI-supported medical diagnostics.

2 METHODS

2.1 Dataset Preparation

The study used the Chest X-Ray Images (Pneumonia) dataset from Kaggle (Mooney, 2018), consisting of grayscale images. For efficiency, images were resized to 150x150 pixels. This dataset is vital for developing machine learning models to automate pneumonia detection, featuring images labeled as 'Pneumonia' and 'Normal' for binary classification. Preprocessing, including augmentation (like adding noise) and normalization, was done to boost model strength. The dataset was divided into training and testing sets to ensure each category was well represented.

Normalization involves scaling image pixel values to the range [0,1] by dividing them by 255, a standard practice to aid model training convergence. Training images also had noise added, with a noise factor of 0.09, to mimic real-world imperfect images and possibly increase model robustness.

2.2 Proposed Model

2.2.1 Convolutional Neural Network

The proposed model uses Convolutional Neural Networks (CNNs), renowned for their ability to recognize, classify, and analyze images by extracting spatial features. The model architecture uses convolutional and pooling layers together for feature extraction and dimensionality reduction, adding batch normalization and dropout to improve and stabilize learning. Then, it splits into two paths: one for classifying pneumonia using dense layers, and another for denoising images to improve diagnostic accuracy.

This dual-pathway approach not only demonstrates the versatility of CNNs but also aligns with the goal of enhancing diagnostic precision by providing denoised images with reliable disease classification. It aims to advance pneumonia detection, showcasing the potential of CNNs in medical image analysis.

2.3 Denoising Autoencoder Model

Denoising Autoencoders (DAEs) are a variant of the autoencoder (Qiu, 2020), which is a type of artificial neural network used for unsupervised learning of efficient coding. The key feature of DAEs is their ability to recover clean data from data corrupted by noise. This is achieved through a process where the DAE learns to encode the input data into a latent-space representation and then decode it back to the original input's clean version. By training on noisy data, DAEs learn to ignore the noise and reconstruct the significant underlying patterns of the input data. DAEs are particularly useful in preprocessing steps for enhancing the quality of data before further analysis. The structure of the proposed model is shown in Figure 1.

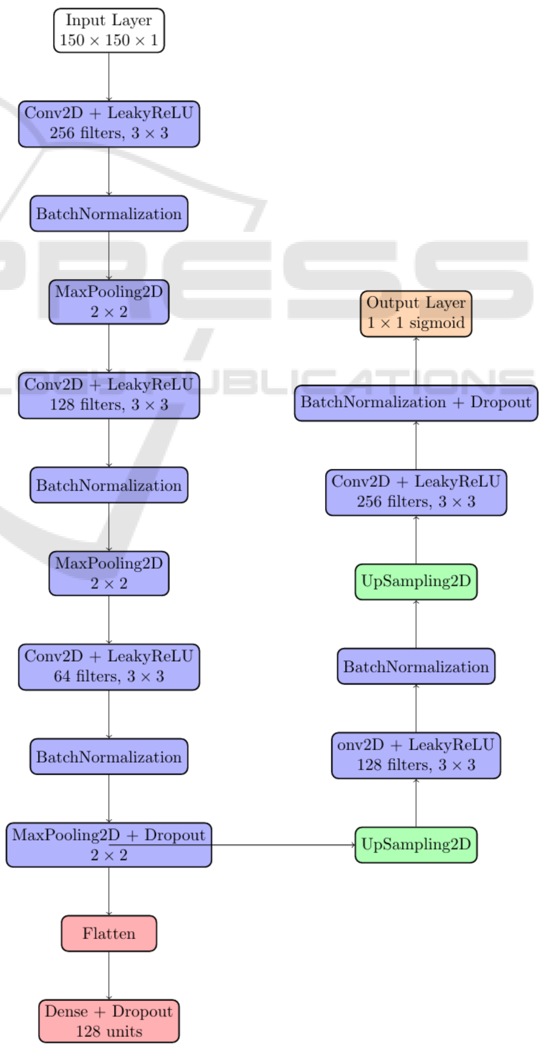

Figure 1: The structure of denoising autoencoder (Picture credit: Original).

2.4 Implementation Detail

In this project, TensorFlow was used for its efficiency and flexibility in deep learning projects, specifically for classifying chest X-ray images as Pneumonia or Normal. Using TensorFlow's high-level Keras API made it easier to build, train, and evaluate the neural network model. An NVIDIA RTX 4080 GPU accelerated the computation, significantly speeding up training times, vital for model iterations and experimentation.

The Adam optimizer, known for its adaptive learning rate feature, was chosen with an initial learning rate of 0.001. This choice is backed by the optimizer's wide success in various deep learning projects. A learning rate scheduler was also used to adjust the learning rate based on the validation set's performance, systematically lowering it to enhance optimization when the classification accuracy plateaued. The scheduler reduces the learning rate by 0.3 after every 2 epochs without accuracy improvements, to a minimum level.

The model was meticulously designed to tackle both denoising and classification, requiring dual loss functions suitable for the binary classification task and for assessing the quality of denoised images against originals. Classification performance was evaluated using accuracy as the key metric, while denoising effectiveness was measured with Mean Squared Error (MSE).

Training ran for 12 epochs with a batch size of 32, balancing computational resources and update frequency. This setup often leads to solid results in various conditions. The training strategy aimed to boost both the denoising and classification capabilities of the model. It made changes to the learning rate to steadily improve performance based on important metrics.

3 RESULTS AND DISCUSSION

3.1 The Classification Performance

The performance graphs give an optimistic view about how good the deep learning model performs in detecting pneumonia from chest X-ray images over 10 epochs shown in Figure 2. Overall, the 'Model Classification Accuracy' graph displays a rapid increase in training accuracy, almost reaching perfection by the second epoch. This quick learning from the training data demonstrates a strong learning capacity. A minor dip in accuracy afterward indicates the model's adjustment to avoid overfitting, quickly regaining high accuracy levels.

Validation accuracy, however, varies more. It initially aligns with the training curve but peaks at 98.4%. The fluctuating accuracy in later epochs points to the model's ongoing development in applying its learning to new data, suggesting room for model improvements for more consistent validation performance.

The 'Model Loss' graph similarly indicates a fast decrease in training loss, highlighting quick progress. An initial increase in validation loss suggests early challenges in generalization, yet the model's swift adjustment implies an inherent adaptability.

These findings suggest a model that can accurately detecting pneumonia, with potential for further refinement. The variability in validation outcomes suggests areas for improvement, such as data augmentation, regularization, and hyperparameter tuning, to enhance its ability to generalize.

3.2 The Denoising Performance

The denoising results shown in Figure 3 presented in the images reflect the model's capability to clean up

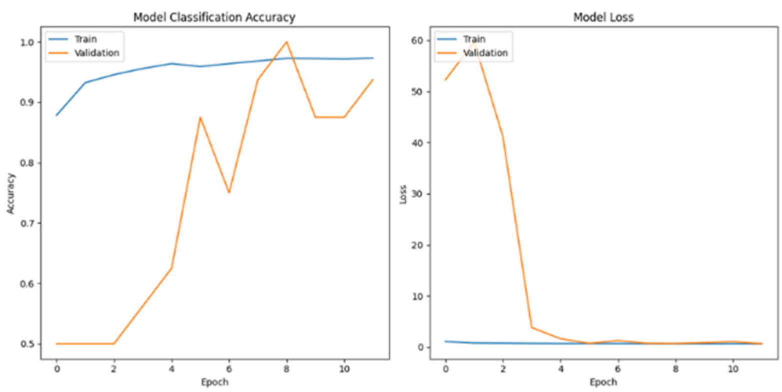

Figure 2: Model Classification Accuracy and Model Loss during the training process (Photo/Picture credit: Original)

noise from chest X-ray images effectively. On the left, the 'Noisy' images are visibly affected by granularity that could obscure diagnostic details. The 'Denoised' images on the right, processed by the model, show a marked reduction in noise, resulting in clearer images where anatomical structures appear more defined.

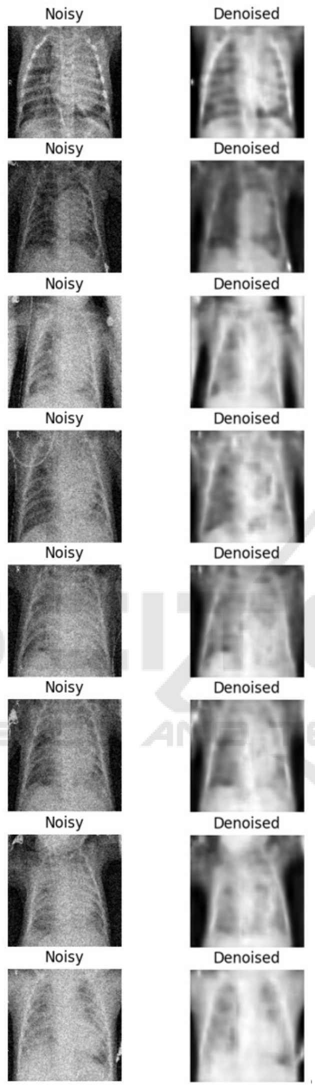

Figure 3: Denoising results (Photo/Picture credit: Original).

Upon analysing the results, it is evident that the model has successfully learned to filter out extraneous noise while retaining the essential features necessary for medical evaluation. The distinction between the original noisy images and the denoised outputs suggests that the model is not only distinguishing between signal and noise but is also enhancing the visibility of potentially critical diagnostic features.

Regarding denoising performance, the comparison between 'Noisy' and 'Denoised' images illustrates the model's efficiency in noise removal, making diagnostic details clearer. The model adeptly filters out irrelevant noise while keeping crucial features for medical assessment. The low Mean Squared Error (MSE) of 0.004951917566359043 for denoised images indicates a high pixel-wise similarity to original, clean images, underscoring the model's proficiency in preserving image quality while reducing noise.

This improvement in image clarity has significant implications for healthcare, as clear images are crucial for precise medical diagnosis. The model's ability to enhance images without compromising detail highlights its potential as a valuable tool for boosting diagnostic accuracy in clinical settings. The results confirm the model's noise-reduction capabilities and its practical value in healthcare. Future work could measure how the improved image quality affects diagnostic accuracy, comparing it against baseline models and traditional noise reduction methods for a fuller picture of the model's real-world medical benefits.

The harmonious optimization benefiting both functions. This convergence implies that features crucial for classification are maintained during denoising, focusing on details important for both clear diagnosis and disease identification. This synergy indicates compatible optimization paths for both tasks, enabling simultaneous improvements without conflicting outcomes. Such a balance is vital for multitasking in medical imaging, allowing the model to deliver clear diagnostic images and accurately detect pathological conditions.

4 CONCLUSION

This study proposes a medical denoising autoencoder for detecting pneumonia from the record of the chest x-ray, which combines image denoising and disease detection into a single model. By leveraging Convolutional Neural Networks, the goal was to address the shortcomings of traditional diagnostic methods and increase the clarity and reliability of automated medical diagnostics. The proposed model effectively created clear images and identified pneumonia. Testing showed the model's effectiveness, with performance measures significantly better than older methods. The potential for utilizing the model in healthcare is evident through its high accuracy in disease diagnosis and enhancement of image quality. Looking ahead, the goal is to broaden the model's use for more medical imaging tasks. The plan is to expand the dataset to include more diseases and use more advanced regularization methods to make the model more resistant to overfitting, thus improving its diagnostic accuracy.

REFERENCES

El-shafai, W., Abd El-Nabi, S., et al. 2022. Efficient Deep-Learning-Based Autoencoder Denoising Approach for Medical Image Diagnosis. Computers, Materials & Continua, 70(3).

Lambert, B., Forbes, F., Doyle, S., Dehaene, H., & Dojat, M. 2024. Trustworthy clinical AI solutions: a unified review of uncertainty quantification in deep learning models for medical image analysis. Artificial Intelligence in Medicine, 102830.

Li, S., Kou, P., Ma, M., Yang, H., Huang, S., & Yang, Z. 2024. Application of Semi-supervised Learning in Image Classification: Research on Fusion of Labeled and Unlabeled Data. IEEE Access.

Liu, Z., Zhang, Z., Lei, Z., Omura, M., Wang, R. L., & Gao, S. 2024. Dendritic deep learning for medical segmentation. IEEE/CAA Journal of Automatica Sinica, 11(3), 803-805.

Liu, Y., Yang, H., & Wu, C. 2023. Unveiling patterns: A study on semi-supervised classification of strip surface defects. IEEE Access, 11, 119933-119946.

Mooney, P. 2018. Chest X-Ray Images (Pneumonia) [Data set]. Kaggle. Available at: https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia

Qiu, Y., Wang, J., Jin, Z., Chen, H., Zhang, M., & Guo, L. 2022. Pose-guided matching based on deep learning for assessing quality of action on rehabilitation training. Biomedical Signal Processing and Control, 72, 103323.

Qiu, Y., Yang, Y., Lin, Z., Chen, P., Luo, Y., & Huang, W. 2020. Improved denoising autoencoder for maritime image denoising and semantic segmentation of USV. China Communications, 17(3), 46-57.

Thomas, J. M., & E, A. P. 2022. Bio-medical Image Denoising using Autoencoders. 2022 Second International Conference on Next Generation Intelligent Systems (ICNGIS).

Zhao, F., Yu, F., Trull, T., & Shang, Y. 2023. A new method using LLMs for keypoints generation in qualitative data analysis. In 2023 IEEE Conference on Artificial Intelligence (CAI) (pp. 333-334). IEEE.