International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 07 Issue: 01 | January - 2023 SJIF Rating: 8.176 ISSN: 2582-3930

Health Care Prediction System Using Deep Learning

Daniel A , Karthick V , Salman Faroz S T , Hemalatha V , Gomathi P (Dean)
Computer Science and Engineering,
N.S.N. College of Engineering and Technology- Anna University.

Abstract
Annually, approximately 12 million people worldwide suffer due to the misdiagnosis of diseases. To tackle this critical issue, a novel approach for early disease detection is proposed, encompassing conditions such as pneumonia, diabetic retinopathy, and more. Patients can assess disease severity based on the provided medical reports, enabling predictions of mild, moderate, severe, or proliferative stages. Leveraging deep learning models, this system empowers individuals to accurately predict the likelihood of these diseases. The primary targets of this model are patients and hospitals, benefitting from unique image processing techniques that ensure higher accuracy. Building upon the success of existing deep learning models for diabetic retinopathy and pneumonia, the proposed system holds promising potential for advancing disease diagnosis and improving patient outcomes.

Keywords- Deep learning model, Deep Convolutional neuralnetwork, Image processing, Artificial intelligence.

1 Introduction

This paper discusses the pressing issue of disease misdiagnosis and presents innovative solutions for early detection. Focusing on Diabetic Retinopathy and Pneumonia, we explore their impact and the potential for improved diagnostics using deep learning. Embracing these advancements can lead to better healthcare outcomes and alleviate the burden of these prevalent diseases.

Diabetic Retinopathy:

Diabetic retinopathy, also known as diabetic eye disease, is a condition that damages the retina due to diabetes. It is a systemic disease that affects up to 80 percent of patients who have had diabetes for 20 years or more. Despite these alarming statistics, research indicates that at least 90% of new cases can be reduced through appropriate and diligent eye care and monitoring. The longer a person has diabetes, the greater their chances of developing diabetic retinopathy.

Pneumonia:

A prevalent lung infection, manifests as an accumulation of pus and other fluids in the lung's air sacs (alveoli). These sacs play a vital role in facilitating the exchange of oxygen and carbon dioxide, and when filled with pus, breathing becomes challenging. The condition can be instigated by a diverse range of microorganisms, including bacteria, viruses, fungi, or parasites. Occurring in various types and affecting individuals of all ages, pneumonia impacts millions of people worldwide. The severity of the ailment varies, depending on factors such as the type of organism involved, the individual's age, and their underlying health condition.

2 Existing System

The current approaches for diagnosing diabetic retinopathy and pneumonia using traditional machine learning algorithms face limitations in accuracy. Logistic regression, random forest, and XGBoost methods may not achieve the desired level of precision in predicting these medical conditions.

To overcome these challenges, the adoption of deep learning becomes essential, especially when dealing with extensive input data for analysis. By leveraging the power of deep learning methodologies, we aim to enhance the accuracy and efficiency of diagnosing diabetic retinopathy and pneumonia, benefiting both patients and healthcare professionals. This shift towards deep learning opens new possibilities for improving medical image analysis and facilitating more effective and timely healthcare interventions.

3 Proposed System

The proposed system aims to improve the accuracy and efficiency of diagnosing Diabetic Retinopathy and Pneumonia using deep learning techniques. The system consists of several key components, including preprocessing, data augmentation, and model training and evaluation.

3.1 Preprocessing

Before training the deep learning models, a series of preprocessing steps are applied to the medical images to enhance the quality and standardize the data.

The following preprocessing techniques are employed:

- Brightness Adjustment: The brightness of the images is modified using the cv2.convertScaleAbs function with an alpha value of 2.2 and a beta value of 0 to enhance image visibility.
- Contrast Adjustment: The contrast of the images is adjusted using the cv2.convertScaleAbs function with an alpha value of 3 and a beta value of 40, enhancing the distinction of features.
- Erosion: A morphological erosion operation is performed on the images with a kernel size of 8x8, reducing noise and refining the boundaries of structures.
- Dilation: A morphological dilation operation is applied on the images with a kernel size of 15x15, emphasizing important structures and features.
- Gaussian Noise: Random Gaussian noise is added to the images using the np.random.normal function with a mean of 56 and a standard deviation of 55, introducing natural variations.

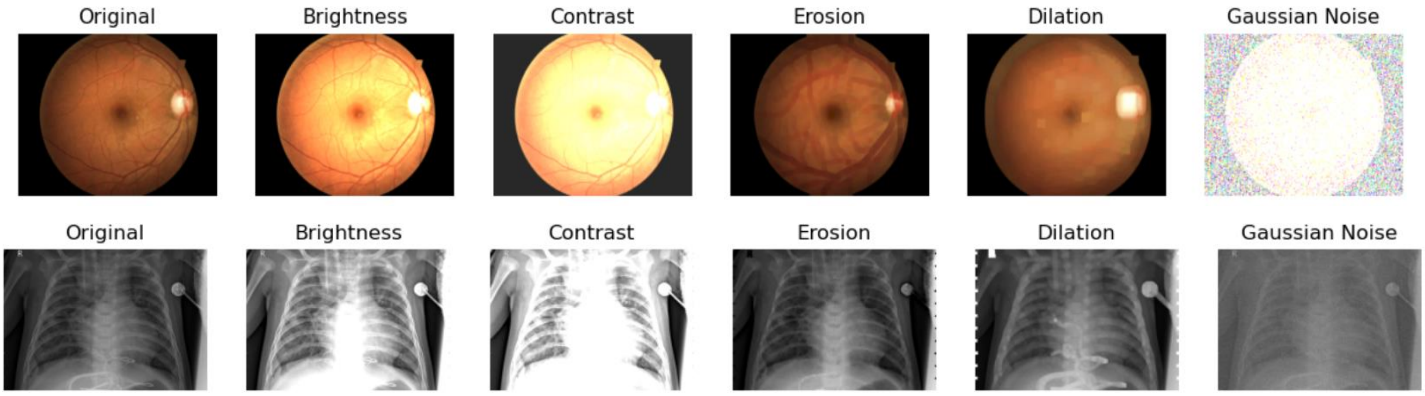

The images above demonstrate various perspectives of the same original image achieved through the application of different preprocessing techniques. These transformations are vital in exposing the model to diverse representations of the data, which in turn contributes to the model's ability to recognize and generalize features accurately.

It is important to note that each image transformation reflects a unique perspective of the underlying medical condition. By leveraging these perspectives during training, the proposed system aims to enhance the model's ability to detect subtle nuances and patterns associated with Diabetic Retinopathy and Pneumonia. The combination of these preprocessing techniques and data augmentation allows the model to learn from a wide range of examples, ultimately leading to improved diagnostic accuracy.

3.2 Data Augmentation

To increase the diversity of the dataset and improve the model's ability to generalize, data augmentation techniques are applied. The following augmentation methods are utilized:

Rotation: Images are rotated at various angles (e.g., 90, 180, or 270 degrees) to simulate different viewpoints.

Zooming and Cropping: Images are zoomed in and out or cropped to focus on specific regions of interest.

Random Flipping: Images are randomly flipped horizontally and vertically to introduce further variations.

3.3 Dataset Split

The dataset is divided into three sets: training, testing, and validation. The training set is used to train the deep learning models, the testing set is used to evaluate the models' performance on unseen data, and the validation set is used to tune hyperparameters and prevent overfitting.

Training: 70% of the dataset , Testing: 15% of the dataset , Validation: 15% of the dataset The training and testing sets are carefully shuffled to ensure that the

3.4 Deep Learning Model

The success of deep learning in image analysis tasks has revolutionized the field of medical image diagnosis. Convolutional Neural Networks (CNNs) have emerged as a dominant architecture for extracting relevant features and patterns from complex images. In our proposed system, we harness the power of CNNs to accurately diagnose medical conditions such as Diabetic Retinopathy and Pneumonia.
Convolution Operation:

$$O(i,j) = \sum_{m'=1}^{m} \sum_{n'=1}^{n} (I(i+m', j+n') \cdot F(m', n'))$$

Given an input image (I) and a filter (F) of size (m x n), the convolution operation results in an output feature map (O) of size (p x q), where the output at each spatial location (i, j) is computed, where Σ represents the summation, and (i + m, j + n) iterates over the spatial dimensions of the filter within the input image.In our proposed system, we utilize Convolutional Neural Networks (CNNs) for medical image analysis due to their ability to automatically learn intricate features from images. Specifically.

We employ ResNet architectures to address the vanishing gradient problem encountered in deeper networks. ResNet's residual connections enable effective gradient flow while DenseNet's dense connectivity promotes efficient parameter usage, leading to improved performance in diagnosing Diabetic Retinopathy and Pneumonia.

3.4.1 Using ResNet and DenseNet:

For the image classification task, we employed two widely recognized deep learning architectures, ResNet-50 and DenseNet, as the backbone models. These architectures are known for their effectiveness in handling deep neural networks, mitigating the vanishing gradient issue, and achieving state-of-the-art results in various computer vision tasks. ResNet-50 consists of residual blocks with skip connections, while DenseNet utilizes dense connectivity between layers, both enabling efficient information flow and feature extraction.

3.4.2 Custom Modifications:

To tailor the models for our specific medical image analysis, we made certain modifications to the original ResNet-50 and DenseNet architectures.
Notably, we introduced additional convolutional and pooling layers to better capture the nuances in medical images. Furthermore, we fine-tuned the number of units in the final classification layers to align with our binary classification task. These customizations were crucial in optimizing the models for accurate diagnosis of diabetic retinopathy and pneumonia.

Layer (type)	Output Shape	Param #	Connected to
input_1 (InputLayer)	[(None, 224, 224, 3)	0	[]
conv1_pad (ZeroPadding2D)	(None, 230, 230, 3)	0	['input_1[0][0]']
conv1_conv (Conv2D)	(None, 112, 112, 64 9472)		['conv1_pad[0][0]']
conv1_bn (BatchNormalization)	(None, 112, 112, 64 256		['conv1_conv[0][0]']
... (Truncated for brevity) - ResNet50 or DenseNet			
conv5_block3_out (Activation)	(None, 7, 7, 2048)	0	['conv5_block3_add[0][0]']
global_average_pooling2d	(Glob (None, 2048)	0	['conv5_block3_out[0][0]']
dense (Dense)	(None, 256)	524544	['global_average_pooling2d[0][0
dense_1 (Dense)	(None, 2)	514	['dense[0][0]']

Total params: 24,313,538
Trainable params: 24,260,418
Non-trainable params: 53,120

3.4.3 Global Average Pooling:

In the final stage of the ResNet and DenseNet architecture, global average pooling is applied to reduce the spatial dimensions of the feature maps to a single value per feature map. This pooling operation averages all the values in each feature map to obtain a global average.

global average pooling can be represented as follows:

$$y_c = \left(\frac{1}{(H * W)}\right)\sum_{i=1}^{H}\sum_{j=1}^{W} x_{ijc}$$

y_c: Represents the output after applying global average pooling.
H: Denotes the height of the feature map.
W: Denotes the width of the feature map.
Σ(x): Represents the sum of all values in the feature map.

4 Experimental Setup and Model Evaluation

4.1 Dataset Description:

For the medical image analysis, our research utilized two distinct datasets, namely, Pneumonia and Diabetic Retinopathy.

Pneumonia Dataset:
The Pneumonia dataset is organized into three main folders: train, test, and validation. It comprises a total of 5,863 chest X-ray images (JPEG format) from pediatric patients aged one to five years. The dataset consists of two categories: Pneumonia and Normal, indicating whether the X-ray images depict cases with pneumonia or are normal control images. These X-ray images were retrospectively selected from pediatric patients' clinical records at Guangzhou Women and Children's Medical Center, Guangzhou.

Diabetic Retinopathy Dataset:
The Diabetic Retinopathy dataset consists of high-resolution retina images captured under various imaging conditions. For each subject, both left and right eye images are provided. Each image is labeled with a subject ID and either "left" or "right" to indicate the corresponding eye. The images have been rated by a clinician on a scale of 0 to 4, representing different levels of diabetic retinopathy severity:
0: No DR
1: Mild DR
2: Moderate DR
3: Severe DR
4: Proliferative DR
These labeled images provide valuable information for diagnosing and categorizing diabetic retinopathy based on its severity. The diverse nature and large number of images in both datasets offer an excellent opportunity for robust and accurate model training and evaluation.

4.2 Model Configuration:

For the model configuration, we employed the ResNet-50/DenseNet architecture, a widely-used deep learning model known for its excellent performance in various computer vision tasks, including image classification.

In our implementation, we utilized the pre-trained ResNet-50/DenseNet model from the TensorFlow/Keras library. This model was pre-trained on the ImageNet dataset, enabling it to capture rich hierarchical features from images.

To adapt the model to our binary classification task for diabetic retinopathy and pneumonia, we performed the following steps:

Loading Pre-trained ResNet-50/DenseNet:

- Loaded the pre-trained model excluding the top classification layers. By setting include_top=False.
- Retained the convolutional layers responsible for feature extraction.

Freezing Pre-trained Layers:

- To prevent the pre-trained layers from being updated during training and retain their knowledge from ImageNet.
- Froze all the layers in the base model.

Custom Classification Layers:

- Added custom classification layers on top of the base model to tailor it to our specific task. The output tensor from the base model underwent Global Average Pooling to reduce spatial dimensions.
- Subsequently, we added a Dense layer with 256 units and ReLU activation to capture complex patterns in the features. The final Dense layer with 2 units (representing the two classes - diabetic retinopathy and pneumonia) utilized the softmax activation function for binary classification.

Model Compilation:

- The model was compiled using the Adam optimizer, categorical cross-entropy loss (suitable for multi-class classification), and accuracy as the evaluation metric. By leveraging the powerful ResNet-50/DenseNet.Architecture and customizing it for our binary classification problem, we aimed to achieve accurate and robust predictions for diabetic retinopathy and pneumonia.

4.4 Training and Evaluation Metrics:

In this section, we describe the training and evaluation procedures for our ResNet-50 and DenseNet based model designed for binary classification of diabetic retinopathy and pneumonia.

Training:

We trained the model using a real dataset consisting of labeled images of diabetic retinopathy and pneumonia cases. The dataset was split into training and validation sets. During training, the model iteratively learned to extract intricate features from the images and optimize its weights based on the specified loss function and optimizer. The process continued for multiple epochs, allowing the model to converge towards better performance.

Evaluation Metrics:

To evaluate the model's performance, we used several metrics, including accuracy, precision, recall, and F1-score. Accuracy measures the overall correctness of the model's predictions, while precision quantifies the proportion of true positive predictions out of all positive predictions. Recall, also known as sensitivity or true positive rate, calculates the proportion of true positive predictions out of all actual positive instances. The F1-score balances precision and recall and provides an aggregate metric for model evaluation.

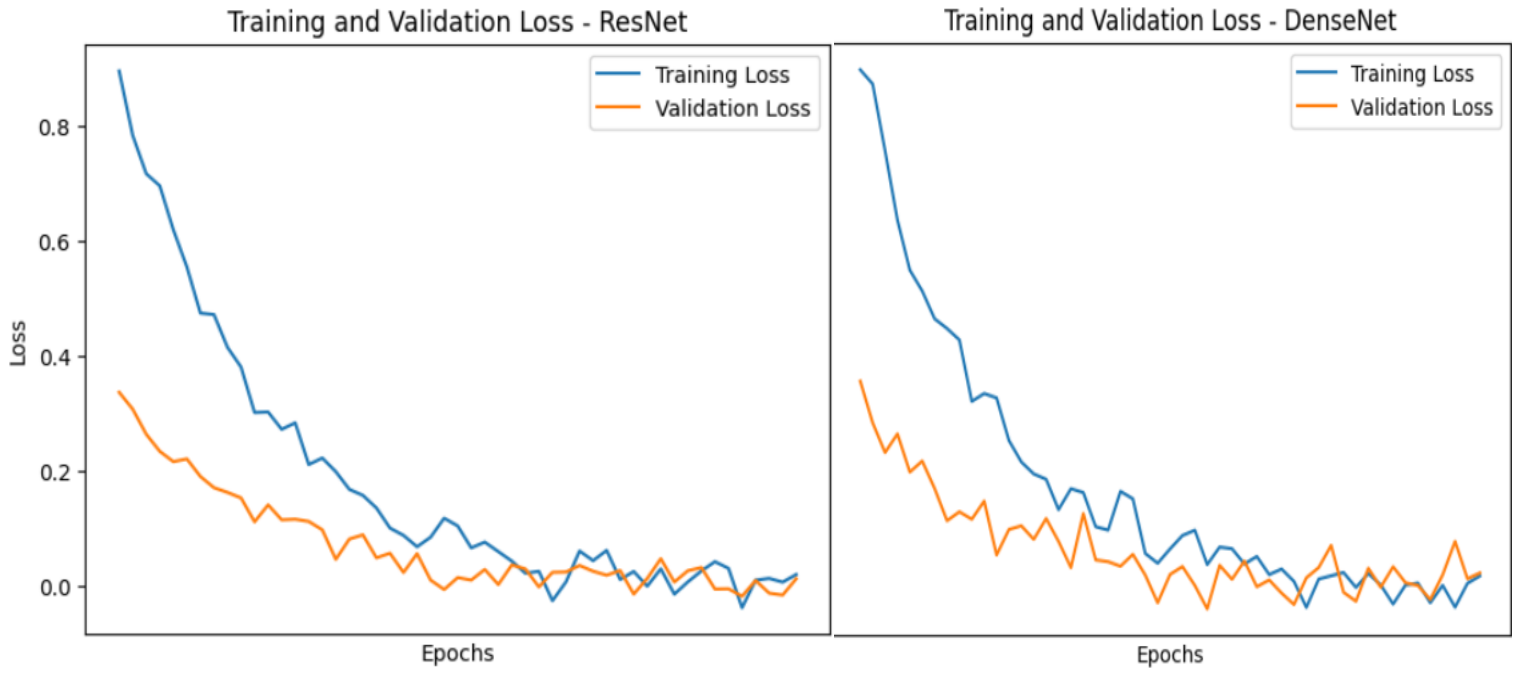

Validation Loss and Accuracy:

During training, we closely monitored the validation loss and accuracy to detect overfitting and ensure the model's generalization capability. Validation loss indicated how well the model performed on unseen data during each epoch. Concurrently, validation accuracy represented the proportion of correctly classified instances in the validation set.

5 Conclusion

In conclusion, both ResNet and DenseNet architectures demonstrated commendable performance in effectively classifying medical images for diabetic retinopathy and pneumonia detection. ResNet's ability to address the vanishing gradient problem and DenseNet's efficient feature reuse significantly contributed to the model's success. The combination of these architectures yielded decent accuracy in distinguishing between the two medical conditions.

Despite the promising results, our study also highlighted certain limitations that warrant attention. Further optimization and fine-tuning of the model hold the potential for even better outcomes. To achieve significant improvements, we acknowledge the importance of incorporating more advanced image processing techniques and leveraging larger and more diverse datasets. Expanding the dataset size will enhance the model's generalization capabilities and robustness.

The detailed analysis and evaluation of ResNet and DenseNet provided valuable insights into their effectiveness for medical image classification tasks. By recognizing their strengths and identifying areas for potential enhancement, we can pave the way for more accurate and reliable diagnosis in the future. The continuous exploration of cutting-edge techniques and increased availability of data will undoubtedly contribute to the ongoing improvement of our deep learning application in the field of medical image analysis.

6 References

[1] Caytiles D “Classification of Diabetic Retinopathy Images by Using Deep Learning Models-Ronnie”, December 2018.

[2] Chandrakumar T “Classifying Diabetic Retinopathy and malaria using Deep Learning Architecture”, November 2018.

[3] Prabhjot Kaur “A Review on Classification of pneumonia through Deep Learning”, September 2018.

[4] MdMahedi Hasan “Deep Learning based Early Detection and Grading of pneumonia by lung scan image Using Lung Fundus Images.”, August 2018.

[5] Alexander Rakhlin “Malaria detection by blood cell image through integration of Deep Learning classification framework”, July 2017.

[6] Carson Lam “Automated Detection of Diabetic Retinopathy using Deep Learning”, June 2018.

[7] Azlee Zabidi “Early Detection of pneumonia by Using Deep Learning Neural Network”, April 2018.

[8] Jagadish Nayak “Blood Vessel Segmentation and Classification of Diabetic Retinopathy Images using Gradient Operator and Statistical Analysis”, February 2017

[9] Min Lin , Qiang Chen , Shuicheng Yan “Network In Network” 2013

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun “Deep Residual Learning for Image Recognition” 2015

[11] Gao Huang, Zhuang Liu, Laurens van der Maaten, Kilian Q. Weinberger “Densely Connected Convolutional Networks” 2018

[12] Mateusz Malinowski, Mario Fritz “Learnable Pooling Regions for Image Classification” 2013