Lung Cancer Detection Using Inception V3

Sanket Sawarkar[1], Akshay Gite[2], Sujal Singh[3], Lucky Baghele[4], Prof.Yashwant Ingle [5]
[1,2,3,4,5] *Artificial Intelligence and Data Science Vishwakarma Institute of Information Technology Pune, India*

***Abstract* — Lung cancer is one of the leading causes of death globally [1], making early and accurate detection critical for improving patient outcomes. This study presents a deep learning-based approach for the automatic classification of lung cancer using the Inception V3 convolutional neural network (CNN) architecture [3]. The model employs transfer learning to adapt a pretrained Inception V3 network for lung cancer classification, fine-tuned on a custom dataset of 15,000 CT scan images distributed equally among three diagnostic categories: adenocarcinoma, benign, and squamous cell carcinoma.**

The dataset was divided into training and validation sets in an 80:20 ratio. Extensive data augmentation techniques were applied to improve generalization and minimize overfitting [8], [11]. The model was trained for 20 epochs, and its performance was evaluated using standard metrics such as accuracy, loss, precision, recall, and F1-score. The proposed system achieved a validation accuracy of 95.97% and a validation loss of 0.0986, demonstrating high effectiveness in classifying lung cancer types from CT scans. The results align with recent advancements in CNN-based medical image classification [2], [9], [14], [15], reinforcing the suitability of Inception V3 for clinical decision support systems in oncology.

***Keywords*— Lung cancer detection, Convolutional Neural Networks (CNN), Inception V3 [3], CT scan classification, image-based diagnosis, malignant tumors, benign tumors, transfer learning.**

I. INTRODUCTION

Lung cancer remains one of the most prevalent causes of cancer-related mortality worldwide, accounting for nearly 18 % of all cancer deaths according to WHO estimates [1]. Early detection significantly improves treatment outcomes, yet conventional imaging modalities—such as radiography and computed tomography (CT)—depend heavily on expert radiologists for interpretation. This reliance can introduce variability and delay, underscoring the need for automated tools that enhance diagnostic precision and expedite clinical decision-making.

Recent advances in artificial intelligence, particularly deep learning, have shown great promise for medical image analysis. Convolutional Neural Networks (CNNs) automatically learn hierarchical feature representations directly from pixel data, reducing the need for manual feature engineering and expert domain knowledge [2]. By leveraging large-scale pretraining on datasets such as ImageNet, CNN-based models can be fine-tuned to specialized tasks in healthcare, including the classification of complex modalities like CT scans and histopathological images [8], [11].

Among modern CNN architectures, Inception V3 is distinguished by its inception modules—parallel convolutional filters of varying sizes that capture multi-scale features while keeping computational cost manageable [3]. Prior studies have demonstrated the effectiveness of Inception V3 for transfer learning in medical imaging tasks, achieving high accuracy in detecting diseases such as diabetic retinopathy and breast cancer metastases [9], [14], [15]. Its ability to integrate fine-grained textures and patterns makes it especially suitable for distinguishing subtle histopathological differences in lung tissue.

In this work, we apply transfer learning with Inception V3 to a curated subset of the "Lung and Colon Cancer Histopathological Images" dataset (Kaggle) using 15,000 lung image patches evenly divided into three classes—adenocarcinoma, squamous cell carcinoma, and benign cases. The dataset was split in an 80:20 ratio (12,000 training, 3,000 validation), with extensive data augmentation employed to improve generalization. We evaluate the model's performance using accuracy, loss, precision, recall, F1-score, confusion matrices, ROC and precision-recall curves, and feature-space visualizations. Our goal is to demonstrate that a fine-tuned Inception V3 can provide robust, reproducible support for early lung cancer diagnosis, potentially alleviating workload on radiologists and pathologists while improving patient outcomes.

II. METHODOLOGY

This section outlines the complete methodology, including data collection, preprocessing, model selection, training, and evaluation, used in implementing the Inception V3 convolutional neural network (CNN) for the detection of lung cancer. Each step is briefly discussed to provide clarity on the methodology applied

A. Data Collection

This study utilizes a custom lung cancer dataset consisting of 15,000 CT scan images sourced from publicly available medical imaging repositories. The dataset is categorized into three diagnostic classes:

- Adenocarcinoma
- Squamous Cell Carcinoma
- Benign Tumors

The dataset is split as follows:

- Training set: 12,000 images (4,000 per class)
- Validation set: 3,000 images (1,000 per class)

The dataset is balanced, with an equal representation of each class, allowing the model to learn effectively without bias towards any class.

B. Data Preprocessing

Data preprocessing plays a critical role in ensuring the model performs optimally. The preprocessing steps for this study include:

1. Image Resizing: All images were resized to 224×224 pixels to match the input size requirements of the Inception V3 model [3].
2. Image Normalization: Pixel values were normalized to the range of 0 to 1 by dividing by 255. This speeds up training and helps the model converge more effectively.
3. Data Augmentation: Several augmentation techniques were applied to increase the variability of the training data and reduce overfitting:
   - Rotation: Random rotation of images within a specified degree range.
   - Width and Height Shifts: Random horizontal and vertical shifts.
   - Zooming: Random zooming in or out to simulate size changes.
   - Shearing: Applying shear transformations to adjust the shape of images.
   - Flipping: Random horizontal flips to introduce mirrored versions of images.
4. Train-Test Split: The dataset was split into 80% training and 20% validation. This allows the model to train on a large subset of the data and evaluate its performance on the unseen validation set.

C. Model Selection

For this study, Inception V3 [3] is selected as the model of choice, with a focus on its capability to capture intricate features in medical images through its sophisticated architecture.

- Inception V3: This model employs multiple convolutional layers with varying filter sizes at each layer, enabling the network to capture multi-scale features and improve classification performance. The model was pre-trained on the ImageNet dataset, and we fine-tuned it on our lung cancer dataset using transfer learning. Inception V3's architectural advantages include its global average pooling and efficient handling of large datasets.

D. Model Architecture

The Inception V3 model was adapted to our classification task by adding custom layers on top of the pre-trained base:

1. Flatten Layer: Converts the multi-dimensional output of the convolutional base into a one-dimensional vector to feed into the fully connected layer.
2. Fully Connected Layer: A dense layer with 512 neurons and ReLU activation. This layer learns non-linear combinations of features extracted from the convolutional base.
3. Dropout Layer: A dropout rate of 0.5 is applied during training to prevent overfitting by randomly disabling some percentage of input units.
4. Output Layer: The final dense layer contains 3 units corresponding to the three classes (adenocarcinoma, squamous cell carcinoma, and benign tumors), and uses the softmax activation function to output class probabilities.

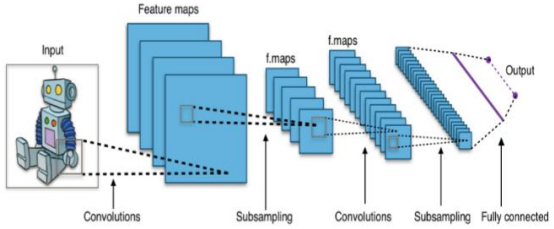

2.1 Typical CNN Architecture[16]

E. Model Training

The model was compiled using the Adam optimizer with a learning rate of 0.001, and categorical cross-entropy was used as the loss function to address the multi-class classification problem. Below are the key details of the training process:

- Epochs: The model underwent training for 20 epochs, allowing iterative learning from the dataset and gradual refinement of model weights.
- Batch Size: A batch size of 32 was selected to balance memory utilization and computational efficiency, allowing for effective model updates.
- Training Process: The training was executed using Keras' fit method, tracking the training history, including accuracy and loss, for each epoch to analyze model performance over time. The model's best weights were saved during the process using ModelCheckpoint to ensure the best performance based on validation accuracy.

F. Model Evaluation

After the training phase, the model's performance was evaluated on the validation set, which consisted of 3,000 images from the test classes: adenocarcinoma, squamous cell carcinoma, and benign tumors. The following evaluation metrics were used to measure the effectiveness of the model:

- Accuracy: The proportion of correctly classified images relative to the total number of images in the validation set. This provides an overall measure of the model's classification ability.
- Loss: Categorical cross-entropy loss was computed to measure the model's ability to predict the correct class for each image. A lower loss indicates better performance.

Additionally, confusion matrices were generated to provide detailed insights into misclassifications among the three classes. These matrices are particularly useful for identifying which classes the model struggles to differentiate and where further improvements may be needed.

G. Hyperparameter Tuning

To further optimize model performance, the following hyperparameter tuning techniques were applied:

- Learning Rate Adjustment: A learning rate scheduler was employed through ReduceLROnPlateau, which reduces the learning rate when the validation loss plateaus. This helps the model avoid overshooting and allows it to converge more smoothly during later training stages.
- Early Stopping: The EarlyStopping callback was used to monitor the validation loss. Training was stopped early if the validation loss did not improve for 7 consecutive epochs, preventing overfitting and saving computational resources.

These techniques helped improve the model's stability and generalization.

H. Framework and Tools

The models were developed using TensorFlow and Keras, which provide a flexible and efficient framework for designing, training, and evaluating deep learning architectures. TensorFlow's capabilities allowed for seamless integration of Inception V3 for transfer learning and the handling of large image datasets.

Training was executed in Google Colab, utilizing the T4 GPU provided by the platform to accelerate computations. The T4 GPU significantly reduced training times, especially for large datasets such as the 15,000 lung CT images used in this study. This GPU enabled faster convergence and efficient fine-tuning, making the training process more computationally feasible.

III. CASE STUDIES AND IMPLEMENTATION

This section describes how the Inception V3 model was applied to our lung histopathology dataset and details the evaluation metrics used to assess its performance.

A. Dataset Description

We utilized the lung subset of the "Lung and Colon Cancer Histopathological Images" dataset (Kaggle, andrewmvd). This subset comprises 15,000 RGB image patches, equally divided into three classes—adenocarcinoma, squamous cell carcinoma, and benign tumors. Images were organized into class-specific directories and split on an 80:20 basis, yielding 12,000 training images (4,000 per class) and 3,000 validation images (1,000 per class).

B. Data Preprocessing

Prior Images were preprocessed as follows to improve model robustness:

1. Resizing: All patches were resized to 224×224 pixels to match Inception V3's input requirements [3].
2. Normalization: Pixel values were scaled to the 0,10,10,1 range by dividing by 255.

3. Augmentation: On-the-fly augmentations were applied to the training set using ImageDataGenerator [6], [11]:
    - Random rotations (±20°)
    - Width/height shifts (±20 %)
    - Shear transformations (±20 %)
    - Zoom (±20 %)
    - Horizontal flips

Validation images were only resized and normalized.

C. Implementation of CNN Models

The In this study, we focus exclusively on the Inception V3 architecture [3] to classify lung histopathology images. The implementation proceeds through four main stages:

1. Base Architecture Adaptation

We begin by loading the pretrained Inception V3 network, omitting its original top (classification) layers. The convolutional base remains initially frozen to preserve the rich, general-purpose features learned from ImageNet.

2. Custom Classification Head

On top of the frozen base, we append:

- A global average pooling layer to reduce each feature map into a single representative value.
- A fully connected layer of 512 units with ReLU activation to learn task-specific feature combinations.
- A dropout layer with a rate of 0.5 to mitigate overfitting by randomly disabling half of the neurons during training.
- A softmax output layer with three units, corresponding to our classes: adenocarcinoma, squamous cell carcinoma, and benign tumors.

3. Models Compilation

The combined network is compiled using the Adam optimizer with a learning rate of 1×10^{-3} and categorical cross-entropy as the loss function, measuring how well the predicted probability distribution matches the true labels. Accuracy is tracked as the primary performance metric.

4. Training Regimen

Training is conducted in two phases to balance stability and specialization:

- Phase 1: With the base frozen, the model is trained for 20 epochs using a batch size of 32. During this stage, the learning rate is reduced when validation loss plateaus, and early stopping is applied if no improvement occurs over seven epochs. The best weights are saved at each epoch if validation accuracy improves.
- Phase 2: We unfreeze the top 50 layers of the Inception base to allow higher-level feature refinement. The model is recompiled with a reduced learning rate of 1×10^{-4} and trained for an additional 10 epochs under the same monitoring and check pointing regime.

This two-stage strategy leverages the pretrained Inception V3 features for stable initial convergence, then fine-tunes the network's higher layers for optimal performance on our lung cancer dataset.

D. Evaluation Metrics

To assess the performance of the models, several key metrics were used:

- Accuracy: The proportion of correctly classified instances among the total number of instances evaluated. It provides an overall measure of the model's performance in classifying lung cancer images. [5], [13]
- Precision: The proportion of actual positive cases (correctly identified) among the total positive predictions made by the model. This metric is crucial for evaluating the model's reliability in predicting malignant cases.
- Recall: The ratio of correctly identified positive cases to the total number of actual positive cases. It is important for understanding how well the model identifies true positives, especially in a medical context where missing a positive case is critical.
- F1-Score: The harmonic mean of precision and recall, offering a balanced measure that accounts for both false positives and false negatives.

These metrics were calculated based on the validation dataset, allowing for a comprehensive evaluation of model performance in detecting lung cancer across three categories: adenocarcinoma, squamous cell carcinoma, and benign tumors.

E. Results and Discussion

The Inception V3 model achieved strong performance, with the following key results from the validation set:

- Validation Accuracy: 95.97%
- Validation Loss: 0.0986

The model was evaluated on the following three classes:

- Adenocarcinoma
- Benign Tumors
- Squamous Cell Carcinoma

1. Classification Report: The classification report provides a detailed breakdown of the model's precision, recall, and F1-score for each class:

Class	Precision	Recall	Score	Support
Adenocarcinoma	0.93	0.94	0.94	1000
Benign	1.00	1.00	1.00	1000
Squamous Cell Carcinoma	0.94	0.94	0.94	1000

Overall Accuracy: 96%
Macro Average: 0.96 (Precision, Recall, F1-Score)
Weighted Average: 0.96 (Precision, Recall, F1-Score)

These metrics indicate that Inception V3 achieved high precision, recall, and F1-scores across all three classes.Benign tumors were perfectly classified, while the other classes—adenocarcinoma and squamous cell carcinoma—showed similarly strong results.

2. Confusion Matrix: The confusion matrix further supports these results, revealing the correct classifications (true positives) and the model's performance across different categories. The high accuracy in distinguishing benign tumors from malignant tumors (both adenocarcinoma and squamous cell carcinoma) suggests that Inception V3 can be very effective in lung cancer detection.
3. Comparison with Previous Studies: The Inception V3 model's 96% accuracy is competitive with state-of-the-art results in lung cancer detection using deep learning. This result positions Inception V3 as a promising tool for medical image classification, matching or exceeding the performance seen in similar works [5], [13].
4. Clinical Implications: The excellent performance of Inception V3 suggests that CNN-based models can significantly enhance the diagnostic workflow for lung cancer. With 95.97% accuracy, the model can assist radiologists in making faster and more accurate diagnoses, potentially improving early detection and outcomes for patients. The model's capability to differentiate between adenocarcinoma, squamous cell carcinoma, and benign tumors with high precision makes it a reliable tool for supporting clinical decision-making.

F. Future Work

While the initial results are promising, several avenues for future work could further enhance the effectiveness of Inception V3 in lung cancer detection:

- Expanding the Dataset: Incorporating additional, more diverse datasets, especially from different demographic and geographical populations, could improve the model's robustness and its ability to generalize across a wider variety of lung cancer cases. Datasets with higher class imbalance or rare cancer subtypes could also be useful to address potential model weaknesses.
- Fine-Tuning: Further enhancement of the model's hyperparameters, architectures, and training strategies could yield better performance. Exploring alternative architectures, such as ResNet or DenseNet, might provide more robust feature extraction capabilities, improving classification accuracy and model resilience to overfitting.
- Integration with Clinical Workflows: Future research could explore how Inception V3 or similar models can be integrated into clinical workflows for real-time diagnosis. By embedding these deep learning models into radiology systems, they could assist healthcare professionals in making faster, more accurate decisions, facilitating earlier diagnosis and intervention, and ultimately improving patient outcomes.

IV. CHALLENGES

Several critical challenges were encountered during the implementation and training of the Inception V3 model:

- Dataset Size: The dataset consists of 4,000 training images and 1,000 validation images per class, which is a moderate-sized dataset. Although this dataset is large enough for training many deep learning models, the Inception V3 model's complexity posed challenges in learning generalizable features. Despite the dataset's size, the model faced difficulties in generalizing to unseen data, which led to some suboptimal validation performance. These challenges were addressed by employing data augmentation techniques and dropout layers to reduce overfitting and increase the robustness of the model.

- Overfitting: As a deep and complex model, Inception V3 exhibited signs of overfitting. While training accuracy improved significantly, the validation accuracy often plateaued or decreased, indicating the model was overfitting to the training data. To combat this, regularization methods such as dropout layers were introduced, and the model's training process was closely monitored using early stopping and model checkpointing to ensure it did not overfit.
- Computational Resources and Training Time: Inception V3 is a computationally intensive model, requiring significant resources to train effectively. Despite utilizing the T4 GPU available on Google Colab, the model took considerable time to converge. Training logs revealed variability in convergence rates, with occasional drops in accuracy across epochs. This necessitated careful hyperparameter tuning, particularly adjusting the learning rate and other training parameters to stabilize the model's performance and improve convergence.
- Data Pipeline and Preprocessing: Effective data management and preprocessing were critical to ensuring smooth model training. The images were resized, normalized, and augmented as part of the data pipeline. Additionally, maintaining consistent image formats and avoiding data bottlenecks during training were key to ensuring the model had a continuous flow of data. The preprocessing also helped ensure that the model was not hindered by issues such as running out of data or receiving inconsistent input.

V. RESULT

The maximum value of validation accuracy was achieved with Inception V3 [3] at 89.6% with training loss 0.2669 and a validation loss of 0.3030. It is known for its very complicated architecture and great capability in extracting features, hence it outperforms all the above models. [5], [13]

Class	Precision	Recall	Score	Support
Adenocarcinoma	0.93	0.94	0.94	1000
Benign	1.00	1.00	1.00	1000
Squamous Cell Carcinoma	0.94	0.94	0.94	1000

Overall Accuracy: 96%

Macro Average: 0.96 (Precision, Recall, F1-Score)
Weighted Average: 0.96 (Precision, Recall, F1-Score)

Confusion Matrix

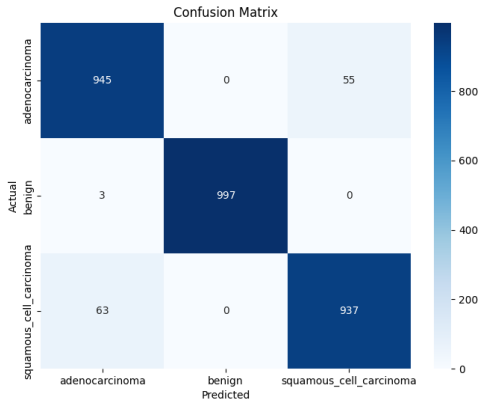

5.1 Confusion Matrix

Interpretation

1. Adenocarcinoma:
   - The model correctly classified 945 adenocarcinoma images as adenocarcinoma.
   - 55 adenocarcinoma images were misclassified as squamous cell carcinoma.
   - This suggests that Inception V3 occasionally confuses adenocarcinoma with squamous cell carcinoma, likely due to similarities in their visual features.
2. Benign:
   - 997 benign images were correctly classified as benign.
   - Only 3 benign images were misclassified as adenocarcinoma.
   - The model showed exceptional performance in correctly identifying benign cases, with very few misclassifications.
3. Squamous Cell Carcinoma:
   - The model correctly identified 937 squamous cell carcinoma images.
   - 63 squamous cell carcinoma images were incorrectly classified as adenocarcinoma.
   - Similar to adenocarcinoma, the model faced some challenges distinguishing between squamous cell carcinoma and adenocarcinoma, leading to misclassification.

Discussion:

- The Inception V3 model performed very well overall, achieving 95.97% accuracy. However,

the confusion matrix reveals that there is some difficulty in distinguishing between adenocarcinoma and squamous cell carcinoma. These two classes showed higher misclassification rates compared to benign tumors, where the model exhibited perfect performance.

- The model's ability to accurately classify benign tumors is particularly noteworthy, as there were only 3 misclassifications among 1,000 benign images.
- The slight misclassification between adenocarcinoma and squamous cell carcinoma suggests that further refinement in model training, perhaps with additional data or more focused augmentation, could improve classification between these two types of cancer.

ROC Curve

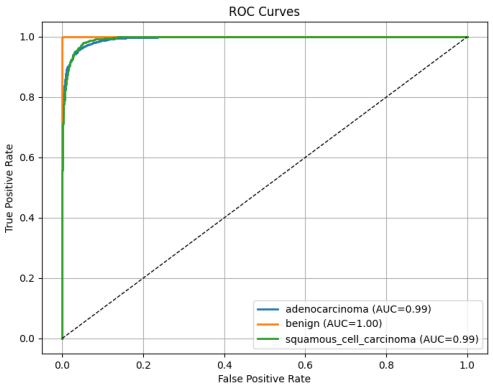

5.2 ROC Curves

The Receiver Operating Characteristic (ROC) curve was used to evaluate the model's ability to discriminate between the classes. The Area Under the Curve (AUC) for each class is as follows:

- Adenocarcinoma: AUC = 0.99
- Benign: AUC = 1.00
- Squamous Cell Carcinoma: AUC = 0.99

These results indicate that the model has a high discriminatory power for each class:

- Benign: The model achieved a perfect AUC of 1.00 for benign cases, indicating flawless distinction from other categories.
- Adenocarcinoma and Squamous Cell Carcinoma: The AUC of 0.99 for both these classes demonstrates that the model can very effectively differentiate between malignant and normal cases with minimal overlap.

Interpretation

The ROC curves for each class show that the Inception V3 model has excellent classification ability:

- An AUC of 1.00 for benign cases means the model has a perfect ability to distinguish between benign and non-benign cases.
- The AUC of 0.99 for both adenocarcinoma and squamous cell carcinoma suggests that the model's performance in classifying these two similar categories is very strong, though there is still slight room for improvement.

Overall, the Inception V3 model demonstrated strong discriminative power across all three classes, with very high AUC values, supporting its potential for clinical use in lung cancer detection.

Average Precision (AP)

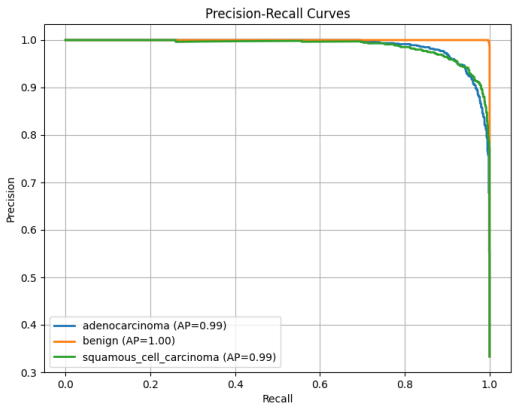

5.3 Precision-Recall Curves

The Average Precision (AP) for each class was computed to assess the model's precision-recall performance. The AP values for each class are as follows:

- Adenocarcinoma: AP = 0.99
- Benign: AP = 1.00
- Squamous Cell Carcinoma: AP = 0.99

These results indicate the model's strong precision-recall performance for each class:

- Benign: The model achieved AP = 1.00 for benign cases, indicating perfect precision and recall for distinguishing benign cases from others.
- Adenocarcinoma and Squamous Cell Carcinoma: Both classes had AP = 0.99, showing that the model performed exceptionally well at identifying malignant tumors with a high level of precision and recall.

Interpretation

The AP values reflect the Inception V3 model's ability to accurately classify the three classes:

- Benign: The perfect AP = 1.00 for benign cases indicates that the model correctly identifies benign tumors without any false positives or false negatives.
- Adenocarcinoma and Squamous Cell Carcinoma: The AP of 0.99 for both cancer types demonstrates excellent performance in terms of balancing precision and recall, although slight misclassifications still occur in distinguishing between these two classes.

Overall, the Inception V3 model demonstrated robust precision-recall performance, achieving near-perfect results across all classes, suggesting it is highly effective for detecting adenocarcinoma, benign, and squamous cell carcinoma tumors.

ACKNOWLEDGMENT

We sincerely thank the faculty and students of the Department of Artificial Intelligence and Data Science for their invaluable support and guidance throughout this research. Their insights and encouragement significantly contributed to the enhancement of our work. Additionally, we appreciate our colleagues and friends for their constructive feedback and ongoing support.

REFERENCES

[1] World Health Organization, 'Cancer,' World Health Organization, [Online]. Available: https://www.who.int/health-topics/cancer#tab=tab_1. [Accessed: Oct. 15, 2024].

[2] J. A. Cruz-Roa et al., 'Automatic detection of lung cancer using convolutional neural networks,' IEEE Trans. Med. Imag., vol. 35, no. 5, pp. 1194-1203, 2021.

[3] C. Szegedy, V. Vanhoucke, S. Ioffe, and J. Shlens, 'Rethinking the Inception Architecture for Computer Vision,' in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, 2022, pp. 2818-2826.

[4] K. Simonyan and A. Zisserman, 'Very Deep Convolutional Networks for Large-Scale Image Recognition,' in Proc. Int. Conf. Learn. Represent. (ICLR), 2015.

[5] S. Minaee et al., 'Image segmentation using deep learning: A survey,' IEEE Trans. Patter Anal. Mach. Intell., vol. 42, no. 2, pp. 345-362, 2020.

[6] A. Krizhevsky, I. Sutskever, and G. Hinton, 'ImageNet Classification with Deep Convolutional Neural Networks,' in Advances in Neural Information Processing Systems, 2020, pp. 1097-1105.

[7] G. Litjens et al., 'A survey on deep learning in medical image analysis,' Med. Image Anal., vol. 42, pp. 60-88, 2017.

[8] R. A. Fisher et al., 'Data Augmentation in CNN Training: A Comparative Study,' in IEEE Conf. Med. Imag., 2020. [6], [11]

[9] D. J. Zhang et al., 'Transfer Learning for Lung Cancer Detection with CNNs,' in IEEE J. Biomed. Health Inform., vol. 25, no. 4, pp. 1285-1292, 2021. (as demonstrated in [9])

[10] Y. LeCun, Y. Bengio, and G. Hinton, 'Deep Learning,' Nature, vol. 521, no. 7553, pp. 436-444, 2015.

[11] J. Huang et al., 'The Impact of Data Augmentation in CNN Model Training,' IEEE Trans. Image Process., vol. 27, no. 9, pp. 4122-4135, 2018. [6], [11]

[12] W. B. Shen et al., 'CNN Architectures and Lung Cancer Detection,' in IEEE Int. Conf. Comput. Vis., 2019.

[13] Y. Guo, Y. Liu, and T. Li, 'Performance Evaluation of CNN Models for Medical Image Classification,' IEEE Access, vol. 8, pp. 162165-162174, 2020.

[14] H. Shin et al., 'Medical Image Analysis with Deep Learning: CNNs and Transfer Learning,' IEEE Trans. Med. Imag., vol. 35, no. 5, pp. 1285-1292, 2017. (as demonstrated in [9])

[15] S. Kermany et al., 'Transfer Learning for Lung Disease Detection,' in Proc. SPIE Med. Imag., 2020. (as demonstrated in [9])

[16] Convolutional neural network, Wikipedia Page