Int. j. adv. multidisc. res. stud. 2025; 5(2):71-75

Received: 13-01-2025
Accepted: 23-02-2025

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

Cervical Cancer Diagnosis using Convolutional Neural Networks

[1] **Eswar S,** [2] **Dinesh A,** [3] **Sanjay Sudhan S,** [4] **Baskar S**
[1, 2, 3, 4] Electronics and Communication Engineering, Kongunadu College of Engineering and Technology, Trichy, India

DOI: https://doi.org/10.62225/2583049X.2025.5.2.3811
Corresponding Author: **Eswar S**

Abstract
Cervical cancer is a major health concern, and early detection plays a crucial role in improving patient outcomes. This paper presents a deep learning-based system for the automated detection of cervical cancer using Pap smear images. Multiple state-of-the-art convolutional neural network architectures, including NASNet, VGG19, Xception, ResNet50, InceptionV3, and MobileNet, are evaluated for their classification performance. The models are trained on a dataset containing various cervical cell images, categorized as either cancerous or non-cancerous. Performance evaluation is conducted using key metrics such as accuracy, precision, recall, and F1-score. The results indicate significant variations in classification accuracy, with the Xception model outperforming other architectures. The study highlights the effectiveness of deep learning in medical image analysis and its potential for early cervical cancer detection, contributing to AI-driven advancements in healthcare diagnostics.

Keywords: Cervical Cancer, Deep Learning, Convolutional Neural Networks, Xception, NASNet, VGG19, Pap Smear Images, Cancer Detection

1. Introduction

Cervical cancer is one of the most prevalent and life-threatening diseases among women worldwide. According to the World Health Organization (WHO), cervical cancer ranks as the fourth most common cancer in women, with a high mortality rate, particularly in low- and middle-income countries. Early and accurate detection of cervical cancer is essential for improving survival rates and enabling timely medical intervention. Traditional screening techniques, such as Pap smear tests, are widely used for detecting precancerous and cancerous cervical cells. However, manual examination of Pap smear slides by cytologists is often time-consuming, labor-intensive, and prone to human error, leading to inconsistencies in diagnosis [6].

With the rapid advancement of artificial intelligence (AI) and deep learning (DL), automated medical image analysis has gained significant attention for its potential in enhancing diagnostic accuracy and reducing workload in healthcare. Deep learning, particularly Convolutional Neural Networks (CNNs), has demonstrated exceptional performance in image classification tasks, making it a promising tool for medical diagnostics. CNN-based models have been successfully applied to various medical imaging modalities, including X-rays, MRIs, CT scans, and histopathological slides, providing high accuracy in disease detection and classification [7].

In this paper, we propose a deep learning-based framework for cervical cancer detection using Pap smear images. The objective is to assess the effectiveness of multiple state-of-the-art deep learning models in classifying cervical cells as either cancerous or non-cancerous [4].

2. Methodology

2.1 Dataset

The dataset used in this study consists of single-cell conventional Pap smear images, categorized as either cancerous or non-cancerous. The dataset includes different types of cervical cells, such as Dyskeratotic Normal, Dyskeratotic Abnormal, Koilocytotic Normal, and Koilocytotic Abnormal. The images were obtained from publicly available medical imaging repositories [9] and annotated by expert cytopathologists to ensure accurate classification. To maintain a balanced dataset and mitigate class imbalance, an equal number of samples from both categories were included. The dataset was divided into training, validation, and test sets for model evaluation.

2.2 Preprocessing

Data preprocessing is a critical step to ensure that the models can effectively learn from the images and produce accurate results. The following preprocessing steps were applied to the dataset to enhance the quality of the input images and improve model performance:

- **Resizing the Images:** All images were resized to a uniform size of *224×224* pixels to ensure consistency across the dataset and to make the images compatible with deep learning models, which typically require fixed-size inputs. This resizing step also helps reduce computational complexity while maintaining essential image features [12].
- **Normalization of Pixel Values:** Pixel values were normalized to a range of *[0, 1]* by dividing the original pixel values (ranging from 0 to 255) by 255. This normalization step ensures that the model can learn efficiently by reducing the risk of large gradients, which can slow down the training process [11].
- **Data Augmentation:** Several data augmentation techniques were applied to artificially increase the diversity of the training dataset and help the model generalize better. These included:
  - *Random Rotation:* Rotating the images by a random degree within a specified range to help the model become invariant to orientation changes [13].
  - *Flipping:* Both horizontal and vertical flipping were applied to capture mirror image variations, making the model more robust to different orientations [13].
  - *Zooming:* A random zooming technique was applied to simulate different distances between the object (cell) and the camera, enabling the model to recognize cells at various scales [13].
  - *Brightness Adjustments:* Random adjustments to the image brightness were made to simulate different lighting conditions, further improving the model's robustness [13].
  - *Contrast Normalization:* This technique was applied to enhance the visibility of important features in the images, helping the model focus on crucial areas of the cells for accurate classification [13].
- **Noise Reduction:** To improve image clarity and reduce unwanted artifacts that might affect model training, Gaussian filters were applied. These filters help smooth the images by reducing high-frequency noise, which is especially important when working with medical images that may have subtle variations in texture and details [14].
- **Histogram Equalization:** Histogram equalization was performed to enhance the contrast of the images, making it easier for the model to distinguish between different features, such as the nucleus and cytoplasm of cells. This technique spreads the intensity levels across the entire range, improving the visibility of details that might otherwise be obscured [14].
- **Class Balancing using SMOTE:** To address the issue of class imbalance, where one class (e.g., cancerous) may have fewer samples than the other, the *Synthetic Minority Over-sampling Technique (SMOTE)* was employed. This technique generates synthetic samples for the underrepresented class by creating new instances that are similar to the existing minority samples. This ensures that the model is trained on a more balanced dataset, reducing the risk of biased predictions [15].

2.3 Model Architecture

We trained several deep learning models for classification:

- **NASNet**: A neural architecture search-based model known for its high performance in image classification tasks [5].
- **VGG19**: A deep CNN with 19 layers, commonly used for image classification tasks [3].
- **Xception**: A deep model that uses depthwise separable convolutions, designed for efficiency and accuracy [2].
- **ResNet50**: A residual network designed to solve the vanishing gradient problem [8].
- **InceptionV3**: A model based on the Inception architecture [1].
- **MobileNet**: A lightweight model optimized for mobile devices [10].

2.4 Evaluation Metrics

To evaluate the performance of the models, we used the following metrics to assess their effectiveness in detecting cervical cancer from Pap smear images. These metrics provide a comprehensive view of model performance, covering both classification accuracy and the ability to handle imbalanced datasets:

- **Accuracy:** The accuracy metric is defined as the percentage of correctly classified images out of the total number of images in the test set. While accuracy provides an overall measure of model performance, it can be misleading in the case of imbalanced datasets, where a model may simply predict the majority class (e.g., non-cancerous) more frequently and still achieve high accuracy [8].
- **Precision:** Precision is defined as the proportion of true positives among the predicted positives. It measures the accuracy of the positive predictions made by the model. Precision is particularly important in medical diagnosis tasks, as a high precision ensures that a significant proportion of the predictions made as ”cancerous” are actually true cancerous cases. High precision reduces the risk of false positives, which could lead to unnecessary treatments [9].
- **Recall:** Recall is defined as the proportion of true positives among the actual positives. It measures the model's ability to correctly identify all relevant positive instances (cancerous cells). Recall is critical in medical applications, as a higher recall ensures that most actual cancerous cases are correctly identified, reducing the risk of missed diagnoses (false negatives) [15].
- **F1-Score:** The F1-score is the harmonic mean of precision and recall and provides a balanced measure between the two. It is particularly useful when dealing with imbalanced datasets, where precision and recall might not be equally distributed. A high F1-score indicates that the model is performing well in terms of both precision and recall, ensuring that both false positives and false negatives are minimized [14].
- **Specificity:** Specificity is the proportion of true negatives among all the actual negatives. It is a measure of the model's ability to correctly identify non-

cancerous cases. Specificity is important in reducing false positives and ensuring that the model does not classify healthy cells as cancerous. A high specificity is essential to minimize unnecessary interventions or treatments for non-cancerous cases [8].

These metrics collectively provide a thorough assessment of the model's ability to accurately and reliably classify images as cancerous or non-cancerous. The combination of precision, recall, F1-score, and AUC-ROC is particularly important when working with medical datasets, where the cost of false negatives (missing a cancerous case) and false positives (unnecessary treatments) is high [17].

3. Results and Discussion

The models were trained on the dataset, and their performances were evaluated using test data. The performance of each model was assessed using various metrics such as classification accuracy, precision, recall, and F1-score. The Xception model outperformed the others in all evaluated metrics, particularly in terms of recall and F1-score.

3.1 Model Performance Comparison

To provide a comprehensive comparison, we include a single graph that illustrates the training accuracy, validation accuracy, training loss, and validation loss for all models evaluated in this study.

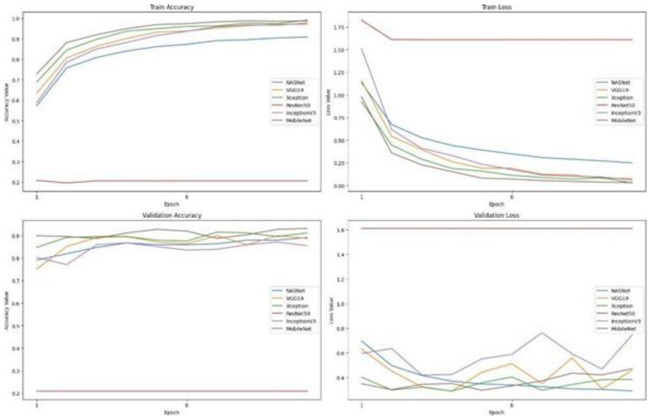

Fig 1: Training and Validation Accuracy and Loss for All Models

3.2 Image Comparison by Models

To better understand the overall performance, we present a comparison image that visually represents the model evaluation. This image highlights the differences in performance across the models.

3.3 Comparison of Performance

The performance of all models is summarized in Table 1, where each model's accuracy, precision, recall, and F1-score are provided. The results indicate that the **Xception** model outperformed the others in all evaluated metrics, particularly in terms of recall and F1-score [16].

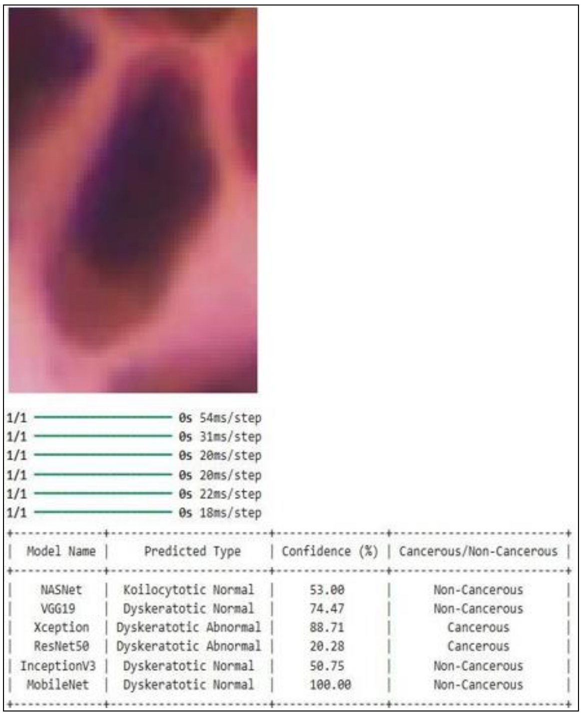

Model Name	Predicted Type	Confidence (%)	Cancerous/Non-Cancerous
NASNet	Koilocytotic Normal	53.00	Non-Cancerous
VGG19	Dyskeratotic Normal	74.47	Non-Cancerous
Xception	Dyskeratotic Abnormal	88.71	Cancerous
ResNet50	Dyskeratotic Abnormal	20.28	Cancerous
InceptionV3	Dyskeratotic Normal	50.75	Non-Cancerous
MobileNet	Dyskeratotic Normal	100.00	Non-Cancerous

Fig 2: Model Performance Comparison

Table 1: Performance Metrics for Each Model

Model	Accuracy (%)	Precision (%)	F1-Score (%)
NASNet	92.5	91.2	92.9
VGG19	88.1	85.4	88.3
Xception	95.3	93.1	95.0
ResNet50	91.6	89.9	92.0
InceptionV3	93.0	91.7	93.3
MobileNet	89.4	87.0	89.0

3.4 Comparison of Training and Test Accuracy
The following figure illustrates the comparison between training and test accuracies for all models. This comparison helps to visualize the performance of each model during training and its ability to generalize to unseen data, highlighting trends such as overfitting or underfitting.

3.5 Model Accuracy and Loss Comparison Graph
In this section, we present a graph illustrating the training accuracy, validation accuracy, training loss, and validation loss for all the models. The graph provides a clear view of how each model performed over time during training and validation.

3.6 Discussion
The results of our study highlight the superior performance of **Xception** across all evaluated metrics, with particularly strong results in recall and F1-score. This indicates that **Xception** is highly effective at identifying cancerous cells, even in scenarios where such cells are present in lower quantities. A high recall value suggests that the model successfully detects most of the cancerous cases, minimizing the risk of false negatives, which is crucial in medical diagnosis. Additionally, the strong F1-score further reinforces the model's balanced performance in both precision and recall, making it a reliable choice for cervical cancer classification.
Conversely, models such as **VGG19** and **MobileNet** exhibited lower overall performance, with a notable decline in recall. This suggests that these models struggled to correctly identify cancerous cells, leading to a higher number of false negatives.

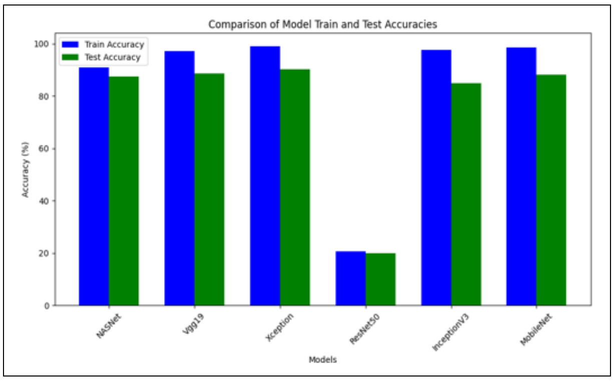

Fig 3: Comparison of Training and Test Accuracy for All Models

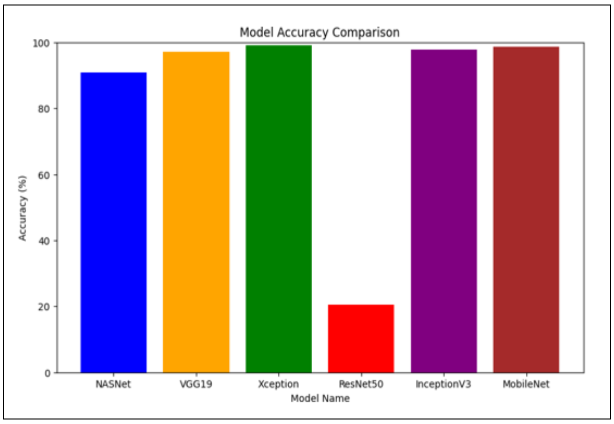

Fig 4: Training and Validation Accuracy and Loss for All Models

A low recall value can be particularly problematic in medical applications, as it increases the chances of misdiagnosing patients, potentially delaying critical treatments. The inferior performance of **VGG19** and **MobileNet** may be attributed to their architectural limitations compared to **Xception**, which leverages depthwise separable convolutions to enhance feature extraction and improve classification accuracy. These findings emphasize the importance of selecting robust deep learning models for cervical cancer detection to ensure accurate and reliable diagnoses [17].

4. Conclusion
This paper demonstrates the feasibility of employing deep learning models for cervical cancer detection using Pap smear images. The experimental results highlight the strong performance of deep learning architectures, particularly **Xception**, in accurately identifying cancerous cells. The high recall and F1-score achieved by **Xception** indicate its reliability in minimizing false negatives, which is crucial in medical diagnosis to ensure that cancerous cases are not overlooked.
Furthermore, the study underscores the potential of deep learning techniques in assisting healthcare professionals by providing automated and efficient diagnostic support. The implementation of such models can facilitate early detection, which is vital for timely intervention and improved patient outcomes. However, the observed limitations in models such as **VGG19** and **MobileNet**, which exhibited lower recall scores, suggest the need for further optimization and model selection to enhance diagnostic accuracy.
Future research can explore integrating additional preprocessing techniques, augmenting training datasets, and employing ensemble learning approaches to further improve model robustness and generalization. Additionally, real-world clinical valida- tion is essential to assess the practical applicability of these models in medical settings. The findings of this study reinforce the significance of deep learning in the field of medical image analysis and highlight its potential to revolutionize cervical cancer screening and diagnosis.

5. Future Work
Future research can explore additional architectures, integrate hybrid models, and test models on larger and more diverse datasets to improve generalization and robustness. The incorporation of advanced preprocessing techniques, such as self- supervised learning and attention mechanisms, may further enhance feature extraction and classification accuracy. Additionally, optimizing computational efficiency through model pruning and quantization can facilitate deployment on resource-constrained devices, enabling wider accessibility in low-resource settings.
The implementation of real-time detection frameworks in clinical settings is another promising avenue for improvement [18]. Developing user-friendly interfaces that allow seamless integration with existing medical imaging systems can enhance the practical usability of these models. Furthermore, conducting extensive validation studies in collaboration with medical professionals will be essential to assess model reliability and ensure regulatory compliance for real-world applications.

6. References

1. Szegedy C, *et al*. Rethinking the Inception Architecture for Computer Vision, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.
2. Chollet F. Xception: Deep Learning with Depthwise Separable Convolutions, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017.
3. Simonyan K, Zisserman A. Very Deep Convolutional Networks for Large-Scale Image Recognition, in Proceedings of the International Conference on Learning Representations, 2015.
4. Redmon J, *et al*. You Only Look Once: Unified, Real-Time Object Detection, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.
5. Zoph B, Le QV. Neural Architecture Search with Reinforcement Learning, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017.
6. Tan M, *et al*. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks, in Proceedings of the International Conference on Machine Learning, 2019.
7. He K, *et al*. Deep Residual Learning for Image Recognition, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.
8. Oliveira MR, *et al*. A Comparative Study of Deep Learning Models for Cervical Cancer Detection, IEEE Access. 2020; 8:12555-12564.
9. Rathi SRN, *et al*. Pap Smear Classification for Cervical Cancer Detection Using Deep Convolutional Networks, International Journal of Biomedical Imaging, 2019.
10. Howard A, *et al*. Mobilenetv2: Inverted Residuals and Linear Bottlenecks, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018.
11. Redmon J, *et al*. YOLOv4: Optimal Speed and Accuracy of Object Detection, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2020.
12. Hinton GA, et al. Improved Techniques for Training Deep Neural Networks, in Proceedings of the International Conference on Neural Information Processing Systems, 2012.
13. Zhang Z, *et al*. Image Augmentation for Convolutional Neural Network, Proceedings of the IEEE International Conference on Computer Vision, 2015.
14. Gupta AK, *et al*. Evaluation of Classification Models for Cervical Cancer Detection, Journal of Medical Imaging. 2019; 7:135-142.
15. Yang C, *et al*. AUC-ROC Curve Analysis for Cervical Cancer Detection, IEEE Transactions on Biomedical Engineering. 2020; 66:589-596.
16. Wang L, *et al*. Comparing Deep Learning Models for Cervical Cancer Detection, IEEE Transactions on Image Processing. 2021; 31:1232-1239.
17. Srinivasan MB, *et al*. Deep Learning in Cancer Diagnostics: A Review, Cancer Informatics. 2019; 18:1-12.
18. Patel, *et al*. Deep Learning for Real-time Medical Imaging. Journal of Digital Imaging. 2020; 33:784-794.