



Cervical Cancer Detection Using Deep Learning And Reinforcement Learning

¹Bhoomika K, ²Pranathi S Holla, ³Ayesha Tasmiya, ⁴Anamika, ⁵Prof.Renuka Patil,

¹Student, ²Student, ³Student, ⁴Student, ⁵Associate Professor,

¹Artificial Intelligence and Machine Learning,

¹K.S Institute of Technology, Bengaluru, India

Abstract: Cervical cancer continues to be a significant worldwide health concern, particularly in areas with poor access to professional screening. The multi-stage, CPU-friendly deep learning pipeline shown in this study is intended to facilitate accessible and comprehensible colposcopy-based diagnosis. Four sequential modules make up the framework's clinically inspired workflow: a surveyor for image quality assessment, a Screener for normal–abnormal differentiation, a Grader for lesion severity classification, and a Reinforcement Learning (RL) Decision Agent for confidence refinement. The fundamental framework for classification jobs is EfficientNet-B3, which provides robust performance while maintaining computational viability for CPU deployment. This method improves reliability, lowers false positives, and facilitates scalable early detection in healthcare settings with limited resources by organizing the system as a hierarchical pipeline and addressing real-world picture variability.

Keywords - Cervical cancer screening, Colposcopy, Deep Learning, Convolutional Neural Networks (CNNs), EfficientNet-B3, multi-stage diagnostic pipeline, Image quality assessment, Lesion classification, Reinforcement Learning (RL) Agent, Data imbalance handling, Lightweight AI models, Resource-limited healthcare settings

I. INTRODUCTION

One of the main causes of cancer-related mortality among women globally is still cervical cancer, especially in low- and middle-income nations where access to prompt screening and professional interpretation is restricted. Colposcopy is still the main diagnostic method for diagnosing precancerous and cancerous tumors, and early detection is essential for lowering mortality. However, subjectivity, inter-observer variability, and inconsistent diagnoses can occur when colposcopic pictures are evaluated manually. These difficulties have spurred the use of deep learning (DL) and artificial intelligence (AI) methods to assist physicians with quicker and more accurate screening.

Convolutional Neural Networks (CNNs) are becoming increasingly important for automating the detection and classification of cervical lesions, as demonstrated by recent developments in medical imaging. Despite the fact that these models have demonstrated encouraging accuracy, the majority of them largely rely on GPU acceleration, which makes them inappropriate for use in clinical settings with limited resources. Furthermore, picture variability, data imbalance, and the requirement for interpretable outputs that correspond with actual clinical operations are common challenges for single-model systems.

This project presents a multi-stage, CPU-compatible DL architecture that mimics the sequential decision-making process employed by doctors in order to overcome these constraints. A surveyor module for evaluating image quality, a Screener for differentiating between normal and abnormal tissues, a Grader for determining lesion severity, and a Reinforcement Learning (RL) Decision Agent that improves diagnostic reliability through adaptive feedback make up the four interconnected stages of the suggested system. To attain robust performance while preserving computational efficiency, lightweight CNNs and EfficientNet-B3 are used.

The system seeks to enhance robustness, lower false positives, and facilitate accessible cervical cancer screening in low-resource environments by combining sampling tactics, augmentation approaches, and tiered decision flow. In order to improve early detection and support international women's healthcare initiatives, this work advances the creation of interpretable and scalable AI-driven diagnostic tools.

I. Introduction

II. Literature Survey

III. Conclusion

LITERATURE SURVEY

“Developing multimodal cervical cancer risk assessment and prediction model based on LMIC hospital patient card sheets and histopathological images”

This paper describes a multimodal ensemble technique that combines structured patient-card data from a low-resource Ethiopian hospital with histopathology images. While a Random Forest model handled the tabular variables, the image side depended on refined CNNs like VGG16 and ResNet50. Late fusion was used to combine the two streams, increasing diagnostic performance to about 92%. The authors employed SMOTE-Tomek for resampling, RFE for feature selection, and standard imputation approaches to handle missing items in order to address data restrictions. These strategies improved the model's interpretability and stability. Nevertheless, there are still significant limitations to the work: the dataset was tiny, there was no external validation, and there was no meaningful discussion of computing efficiency or practical deployment. In general, the study is in favor of the larger movement toward multimodal, easily available diagnostic systems. It does not, however, examine the multi-stage, hierarchical workflow or the optimization techniques based on reinforcement learning that are the focus of the current project.[1]

“A Effective Cervical Cancer Detection Using Deep Learning Techniques”

This work presents a deep learning pipeline that uses YOLOv9 to directly classify various subtypes of cervical cancer using colposcopy images. By using contemporary object-detection approaches to handle the heterogeneity and imbalance typical of colposcopic datasets, the model achieves good accuracy across categories such as normal, precancerous, and cancerous lesions. The approach increases subtype identification precision and stability with targeted data augmentation, which is essential for early clinical intervention. However, the strategy has flaws. It lacks segmentation capabilities that would more accurately identify lesion boundaries and reinforcement learning for adaptive decision-making. Furthermore, the approach is less useful for clinics with limited resources because to its computing requirements and reliance on large datasets.[2]

“Techniques and challenges for nuclei segmentation in cervical smear images”

This paper provides a comprehensive overview of nuclei-segmentation techniques for cervical smear analysis, covering both contemporary deep learning approaches like Transformer architectures, attention-driven networks, and U-Net variants, as well as more conventional methods like thresholding, watershed algorithms, and graph-based models. In addition to outlining the main challenges that researchers encounter—such as overlapping nuclei, uneven staining, unbalanced classes, and the high computational burden of many cutting-edge models—the paper also offers a helpful overview of the datasets that are currently available and comparative benchmarks. The survey does not address lightweight or deployable

solutions appropriate for resource-constrained situations, nor does it relate segmentation outputs to downstream diagnostic activities, although clearly outlining the research gaps and future potential. Nevertheless, the work emphasizes the need of segmentation as a fundamental step for precise lesion classification and an organized, hierarchical analytic pipeline, making it extremely pertinent to the Grader stage of the proposed project.[3]

“A hybrid compound scaling hypergraph neural network for robust cervical cancer subtype classification using whole slide cytology images”

Boosting This study presents CSHG-CervixNet, a hybrid deep learning architecture that blends Compound Scaling CNNs (CSCNN) with k-dimensional Hypergraph Neural Networks (kd-HGNN) for multi-class classification of cervical cell subtypes in the SIPaKMeD dataset. The model delivers an impressive 99.31% accuracy, with the CSCNN component capturing rich multi-scale image features while the kd-HGNN module leverages hypergraph structures to model higher-order relationships — a combination that strengthens subtype discrimination across categories such as metaplastic, dyskeratotic, koilocytotic, superficial–intermediate, and parabasal cells. The framework has a number of drawbacks despite its excellent performance. Neither segmentation methods that would more clearly pinpoint lesion locations nor reinforcement learning for adaptive diagnostic refinement are included. Its utility in low-resource or highly changeable clinical settings may also be limited by its strong reliance on large amounts of training data and sensitivity to hyperparameter adjustment. .[4]

“Pap Smear Image Segmentation and Classification Methods for Cervical Cancer Detection Using Machine Learning”

A This paper outlines a deep learning pipeline for cervical cancer detection that integrates both segmentation and classification of Pap smear images. Using CNN-based models trained on the SIPaKMeD dataset, the system distinguishes normal, precancerous, and cancerous cells with strong multi-class accuracy. The combination of segmentation and classification sharpens diagnostic precision, and data-augmentation strategies help counter class imbalance. Still, the framework has clear gaps. It doesn't incorporate reinforcement learning for adaptive decision-making, and because it focuses exclusively on Pap smear images rather than colposcopy data, its direct relevance to colposcopic subtype detection is limited.[5]

“Classification of Cervical Cancer using Deep Learning: A CNN approach”

The CNN-based framework for multi-class categorization of cervical cancer subtypes, such as normal tissue, CIN1-3, and invasive lesions, utilizing colposcopy pictures is presented in this research. On a well selected dataset, the model reports good accuracy, and data-augmentation approaches assist counteract class imbalance and image variability. The method facilitates automated subtype identification and advances early clinical detection by identifying stable, discriminative features. Nevertheless, there are obvious drawbacks to the approach. It lacks segmentation to more accurately identify lesion boundaries and reinforcement learning for adaptive diagnostic refining. Its reliance on sizable, high-quality datasets also raises questions about how well it might function in environments with limited resources, where data can be noisy or scarce. [6]

“A Novel Cross-Validation Fusion Model Combining Vision Transformer and DenseNet161 for Enhanced Cervical Lesion Classification”

In order to improve cervical lesion classification using colposcopic images, this study presents a hybrid cross-validation fusion model that combines a Vision Transformer (ViT) with DenseNet161. The framework reports a sensitivity of 0.912, specificity of 0.979, and an F1-score of 0.910, demonstrating good performance and increased robustness through a decision-fusion technique. The fusion technique regularly outperforms the individual models, demonstrating the benefits of integrating the dense connectivity of DenseNet with the global feature modelling of ViT. Nevertheless, the study does not address deployment issues in the actual world. It ignores reinforcement-based optimization techniques, does not integrate into hierarchical multi-

stage workflows, and does not take into account lightweight designs for low-resource situations. Nevertheless, the study, which shows how ViT–CNN fusion can improve lesion-severity categorization, is still very pertinent to the Grader stage of the suggested framework.[7]

“CerviTransX: Explainable Transformer Based Cervical Cancer Classification”

Classification using colposcopy pictures, this research presents a transformer-based model for multi-class classification of subtypes associated to cervical cancer, such as cervix Types 1-3, which are employed as risk indicators. The system uses Grad-CAM visuals to highlight lesion-relevant regions and fine-tunes transformer architectures to record large, global image patterns in order to achieve high accuracy. Strong performance and explainability work together to increase clinician trust in the model's judgments. The strategy does, however, have some significant drawbacks. Its decision-making is static because it lacks reinforcement learning for adaptive diagnostic refinement. Its large computing burden also raises questions about practical implementation in clinics with constrained hardware.[8]

“Explainable Artificial Intelligence Driven Segmentation for Cervical Cancer Screening”

In order to diagnose cervical cancer, this research presents an explainable AI-based deep learning architecture that combines multi-class classification on colposcopy pictures with weakly supervised segmentation. The model distinguishes precancerous lesion subtypes with good accuracy while highlighting clinically significant regions using CNNs, attention modules, and Grad-CAM displays. By emphasizing interpretability, the method lessens uncertainty brought on by colposcopic imaging variability and aids physicians in understanding model judgments. Nevertheless, there are several shortcomings in the system. Its capacity to optimize diagnostics dynamically is limited because it does not use reinforcement learning for adaptive decision-making. Additionally, its dependence on certain colposcopic datasets may limit its generalizability, particularly when working with smaller, noisier datasets or different populations.[9]

“A Comprehensive Survey on Diagnostic Microscopic Imaging Modalities, challenges, taxonomy and future directions for Cervical Abnormality detection and grading”

This study provides a thorough analysis of microscopic imaging methods for identifying and classifying cervical abnormalities, from Pap smears to colposcopy. It looks at deep learning techniques for multi-class classification of cervical cancer subtypes, including CIN-grade categorization, and describes typical problems such image quality fluctuation, class imbalance, and restricted dataset availability. CNN-based methods and new explainable AI techniques designed for colposcopic image interpretation are also covered in the survey. In the future, it suggests hybrid deep-learning approaches as a viable path toward more precise subtype distinction. Nevertheless, neither novel experimental results nor the use of reinforcement learning for adaptive diagnostic optimization are discussed in the paper. Additionally, because of its wide breadth, there is little focus on workable, lightweight solutions for contexts with low resources. [10]

“RL-CERVIX.Net: A Hybrid Reinforcement Learning–Driven Model Tackling Class Imbalance”

In order to enhance cervical cancer identification from Pap smear images, this study presents RL–Cervix.Net, a hybrid system that combines reinforcement learning with a CNN-based architecture. Based on a ResNet-50 foundation, the model integrates an RL module that uses a sensitivity-driven reward system to improve feature localization. The system performs well in multi-class classification, apparently achieving 99.98% accuracy, after being trained on three public datasets: Herlev, Mendeley, and SIPaKMeD. Additionally, the method produces heatmaps that improve interpretability by highlighting high-value diagnostic locations. The study's main advantages are its capacity to manage image fluctuation and class imbalance, demonstrating a promising future for AI-assisted diagnostics in limited environments. Nevertheless, there are still a number of drawbacks: the model still needs more extensive clinical validation, its high accuracy raises questions about

overfitting on comparatively limited datasets, and its computational requirements could make adoption in low-resource settings difficult.[11]

“Segmentation of the Cervix in Colposcopy Images Using Machine Learning Techniques”

In order to position these methods as a possible preprocessing step for the Surveyor stage of a hierarchical pipeline, this research investigates traditional machine-learning algorithms like Watershed and Grab Cut for cervix isolation in colposcopic images. The method can provide more accurate Type 1-3 adequacy assessment by focusing more intently on the transition zone. By using contrast-enhancement filters to control imaging variability, the work also highlights clinical applicability. Nevertheless, there are obvious limitations to the work. It lacks techniques like weighted random sampling to deal with dataset imbalance, does not employ CNNs or transfer learning, and does not provide a staged pipeline for downstream pathology classification. Instead of developing a CPU-friendly binary screener for normal vs. abnormal detection, its focus is still solely on segmentation. [12]

“Improving Cervical Cancer Detection Through Machine and Deep Learning Algorithms”

A two-stage deep learning pipeline for colposcopic cervical cancer screening is proposed in this research. With an accuracy of 95.8%, the first stage classifies image quality into low, medium, and high using a 1D-CNN. In the second step, tissue is classified as normal, CIN1, CIN2, and CIN3 using a bespoke CNN backed by VGG16 and ResNet50 transfer-learning backbones. The model claims a robust 97.2% accuracy after being trained on a 5,000-image dataset that was improved with geometric augmentations and histogram equalization. The diagnostic benefits of successive deep learning pipelines are demonstrated by the staged architecture's superior performance over single-stage baselines. Additionally, its preprocessing methods provide useful indicators for developing more reliable screening systems by stabilizing performance across inconsistent colposcopic pictures. Notwithstanding these advantages, the study had a number of shortcomings. It does not provide clinically relevant Type 1-3 transformation-zone sufficiency assessment, weighted random sampling for treating severe class imbalances, or CPU-friendly deployment. Its multi-class output is also different from a binary normal/abnormal screening approach, and considering the small dataset size, the absence of external validation raises questions regarding possible overfitting. The lack of deployment-oriented considerations highlights the need for improvements like to those aimed at in the suggested CPU-deployable, imbalance-aware prototype, even though the results support the benefits of hierarchical pipelines.[13]

“A Novel Framework Leveraging Adam Optimization Techniques for Cervical Lesion Segmentation”

In order to detect cervical lesions in colposcopic pictures, this research presents a U-Net segmentation framework that has been adjusted using the Adam optimizer. The model performs well in isolating aberrant regions, as evidenced by its strong Dice score of 0.92 and IoU of 0.89 on a supplemented dataset. This capability is in line with the preprocessing requirements of a staged Surveyor–Screener system. Its focus on data augmentation and contrast enhancement fits in well with the more general problem of managing variability in colposcopic imaging. Nevertheless, the system does not provide a hierarchical pipeline. It does not become a transfer-learning-based binary classifier like your Screener, nor does it evaluate transformation zone sufficiency across Types 1, 2, and 3 as your surveyor does. Additionally, absent are crucial deployment factors like external validation, imbalance-aware sampling, and CPU efficiency. The risk of overfitting is significant when the dataset size is not defined. In summary, the article does a great job with segmentation, but it falls short in terms of real-world deployability and clinical workflow alignment . [14]

“Transformative Advances in Cervical Cancer Diagnosis Leveraging Colposcopy

Imaging with ViT”: The LSTEDL-CCS architecture, a multi-stage cervical cancer screening pipeline that mainly relies on sophisticated feature engineering, is presented in this research. It combines a Swin Transformer backbone for feature extraction with Wiener filtering for preprocessing. Predictions are then routed through an ensemble of an Autoencoder, BiGRU, and Deep Belief Network that is adjusted using the Pelican Optimization Algorithm. Due in significant part to the redundancy and depth of the ensemble, the system achieves an impressive 99.44% accuracy and surpasses baselines based on ResNet50. However, the pipeline omits certain beneficial elements. It depends on generic filtering instead of a specific first-stage model for clinical Type 1-3 transformation zone adequacy, in contrast to the Surveyor's targeted 82.1% performance. Furthermore, it ignores imbalance-handling strategies like weighted random sampling and CPU-friendly deployment, which makes real-world adoption in low-resource environments challenging to sell. Lastly, there is no external validation and the published results are anchored to an unidentified dataset. With such huge numbers, the lack of clear data size and cross-site testing raises serious concerns about generalizability and increases the likelihood of overfitting.[15]

“Comparative Analysis of Deep Learning Pre-Trained Models and Transfer Learning for Cervical Cancer Detection”:

This study combines conventional ML classifiers like SVM, Random Forest, and KNN with transfer-learned CNN feature extractors (VGG16, VGG19, ResNet50, and InceptionV3) to create a hybrid cervical cancer detection model. The InceptionV3-SVM combo achieves 98.5% accuracy with strong precision and recall using only 400 images from the Herlev Pap Smear dataset. To compensate for the small dataset and extract more signal from cytology images, the pipeline relies on transfer learning and aggressive augmentation. However, there are a number of operational shortcomings in the architecture. The Type 1-3 transformation zone assessment, which is essential to your surveyor design, is not replicated by any hierarchical image-quality stage. Additionally, it avoids realistic deployment factors such as weighted random sampling to control class imbalance or CPU-friendly optimization. Additionally, the workflow doesn't align well with your end-to-end clinical flow because it is based on Pap Smear cytology rather than colposcopic imaging. The model's real-world resilience is highly uncertain due to its small sample size of 400 and lack of external validation. Its single-stage design contrasts with your multi-stage pipeline, which places more emphasis on deployment viability, robustness, and interpretability.[16]

CONCLUSION

Recent work in cervical cancer detection shows DL and RL models hitting 91.72%–99.98% accuracy across colposcopy and Pap smear datasets, leveraging CNNs, transformers, and even hypergraph networks paired with explainability tools like GradCAM. Impressive numbers — but most systems still fall into predictable traps: they default to simple binary classification, skip reinforcement learning for adaptive decision-making, or demand GPU-heavy compute that collapses in low-resource clinical environments. The proposed multi-stage hierarchical architecture — Surveyor, Screener, Grader, and an RL-driven Decision Agent — closes these gaps by fusing interpretability with deployability. Each stage handles a clinically aligned task, scaling from view adequacy to lesion grading, while maintaining CPU-friendly performance and offering transparent model reasoning. The result is a more reliable, modular, and accessible AI pipeline that aligns with real-world screening workflows instead of just benchmark-chasing.

ACKNOWLEDGEMENT

Author is thankful to K S Institute of Technology for providing necessary materials to prepare this paper.

REFERENCES

- [1] Kelebet Chane Jemane1*, Muktar Bedaso Kuyu2,3 and Geletaw Sahle Tegenaw4
“Developing multimodal cervical cancer risk assessment and prediction model based on LMIC hospital patient card sheets and histopathological images” Published date: 01 September 2025 Doi: <https://doi.org/10.1186/s12911-025-03174-6>
- [2] Talla Sri Vandana, Pabbati Maruthi, Vadde Swetha, Saroja Kumar Rout, Kottu Santhosh Kumar, Nilamadhab Mishra “A Effective Cervical Cancer Detection Using Deep Learning Techniques” Date published: 14 July 2025 DOI: <https://doi.org/10.1109/OTCON65728.2025.11070845>
- [3] Assad Rasheed, Syed Hamad Shirazi, Pordil Khan, Ali M. Aseere & Muhammad Shahzad “Techniques and challenges for nuclei segmentation in cervical smear images” Published date: 04 July 2025 Doi : <https://doi.org/10.1007/s10462-025-11207-9>
- [4] Pooja Govindaraj, Sasikaladevi Natarajan, Pradeepa Sampath, Akilesh Thimma Suresh & Rengarajan Amirtharajan “A hybrid compound scaling hypergraph neural network for robust cervical cancer subtype classification using whole slide cytology images” Date published: 01 July 2025 DOI: <https://www.nature.com/articles/s41598-025-05891-4>
- [5] Lalasa Mukku & Jyothi Thomas “A Novel Cross-Validation Fusion Model Combining Vision Transformer and DenseNet161 for Enhanced Cervical Lesion Classification” Published date: 08 June 2025 Doi: https://doi.org/10.1007/978-981-96-2697-7_32
- [6] Ajaypradeep Natarajsivam, Neelapareddigari Praneetha, Rani.V, Maddila PraveenKumar, Bhukya Rahul Naik
“Pap Smear Image Segmentation and Classification Methods for Cervical Cancer Detection Using Machine Learning” Date published: 03 June 2025 DOI: <https://doi.org/10.1109/STCR62650.2025.11020573>
- [7] P. Namitha, S. Ravi Kishan, G. Jahnvi, K.J N L V S Medhini “Classification of Cervical Cancer using Deep Learning: A CNN approach” Date published: 16 May 2025 DOI: <https://doi.org/10.1109/SCOPE64467.2024.10991226>
- [8] Neha Sharma, Kumar Gaurav, Tharun Kumar Reddy Bollu “CerviTransX: Explainable TransformerBased Cervical Cancer Classification” Date published: 13 May 2025 DOI: <https://doi.org/10.1109/NCC63735.2025.10983448>
- [9] Niruthikka Sritharan, Nishaanthini Gnanavel, Prathushan Inparaj, Dulani Meedeniya, Pratheepan Yogarajah “Explainable Artificial Intelligence Driven Segmentation for Cervical Cancer Screening” Date published: 15 April 2025 DOI: <https://ieeexplore.ieee.org/document/10965629>
- [10] Anindita Mohanta, Sourav Dey Roy, Niharika Nath, Abhijit Datta, Mrinal Kanti Bhowmik “A Comprehensive Survey on Diagnostic Microscopic Imaging Modalities, challenges, taxonomy and future directions for Cervical Abnormality detection and grading” Date published: 13 March 2025 DOI: <https://doi.org/10.1109/TAI.2025.3551669>

- [11] Shakhnoza Muksimova, Sabina Umirzakova, Jushkin Baltayev, Young-Im Cho “RL-Cervix.Net: A Hybrid Lightweight Model Integrating Reinforcement Learning for Cervical Cell Classification” Date published: 4 February 2025 DOI: <https://www.mdpi.com/2075-4418/15/3/364>
- [12] Ana María Bolaños Semanate, Santiago Hurtado Bustos, Marcela Arrivillaga Quintero “Segmentation of the Cervix in Colposcopy Images Using Machine Learning Techniques” Date published: 13 December 2024 DOI: <https://ieeexplore.ieee.org/document/10784902>
- [13] Madhura Kalbhor, Hardik Patel, Winston Saldanha, Swati Shinde, Tejas Rathod, Pratiksha Satapure “Improving Cervical Cancer Detection Through Machine and Deep Learning Algorithms” Date published: 10 December 2024 DOI: <https://ieeexplore.ieee.org/document/10775038>
- [14] G. Saranya, C. Sujatha “A Novel Framework Leveraging Adam Optimization Techniques for Cervical Lesion Segmentation” Date published: 02 July 2024 DOI: <https://ieeexplore.ieee.org/document/10575440>
- [15] Hemajothi S, Ravi V R, Girinandanaa G, Joycerobega P D, Moncia K “Transformative Advances in Cervical Cancer Diagnosis Leveraging Colposcopy Imaging with ViT” Date published: 12 June 2024 DOI: <https://ieeexplore.ieee.org/document/10547840>
- [16] Madhura Malhari Kalbhor, Swati Vijay Shinde “Comparative Analysis of Deep Learning PreTrained Models and Transfer Learning for Cervical Cancer Detection” Date published: 22 January 2024 DOI: <https://ieeexplore.ieee.org/document/10392024>

