PakTech Today | Pakistan's Premier Technology Magazine
Vol. 01, Issue 01 • April '26 • paktech.today

AI & HEALTH TECH Research Communication / Medical AI / Explainable Intelligence

OncoLens AI: When a Machine Learns to Read a Biopsy and Explains What It Sees

Prostate cancer is the one of the major diseases that takes the lives of the 50 million men every year. This cancer resides in the tissues that takes too much time to be diagnosed by the doctors even approximately it takes 18 days to complete this process. Oncolens Ai is aiming to reduce this time by training the model that helps to recognize the prostate cancer at the early stages and help the professionals to diagnose this

Muhammad Muteeb Ramzan | University of Agriculture Faisalabad, Pakistan | Published: 14th April 2026
muteebramzan3@gmail.com | OncoLens AI Research Team

Abstract

Prostate cancer is the second most common cancer that is the diagnosed in the world but in the Pakistan its more dangerous due to the lack of the knowledge and the less access to the health care professionals. Oncolens Ai is the deep learning model that aims to address this gap present in the healthcare. The Oncolens ai currently applies the SAM-2 for the analysis of the biopsy histopathology images sourced from the Radboud Institute and the performing pixel-level tissue segmentation to identify the cancer regions. Initially the model is the trained on the 120 biopsy images out of the 5000. This segmentation is coupled with an Explainable Ai layer built on the GRAD-CAM which generates color-coded visual heatmaps that reveal the exactly which tissue features drove each prediction enables the pathologist to verify the system's output. Oncolens Ai is not designed to replace the clinical expertise. It is designed to help the healthcare professionals and the always available diagnostic assistant in the hands of every physician. This article presents the projects methodology, early results, active challenges, and the full scope of what AI will mean for the healthcare systems when in the production.

Keywords: Prostate Cancer Detection; Biopsy Image Analysis; Histopathology AI; SAM-3; Explainable AI; GRAD-CAM; Deep Learning; Pakistan Healthcare

Introduction: The Slide That Nobody Read in Time

There is the slide that is the present in the laboratory in the Pakistan at this time. A very small strip of the prostate tissue, prepared and the stained, and the ready for the examination by pathologist professional. The slide is present there for weeks. The pathologist who should read it have the burden of work. He has to examine the hundred of the slides and he has the constant grinding of pressure that the speed and the accuracy of that slide both are the non-negotiable when there is the risk of human life. On the other side, a man whose tissue is present on the slide, he does not know how much time he has to wait and this is the story of the most men effected with the prostate cancer when the diagnosis is completed the much time has been passed and the simple treatment is the no longer option of saving the life of effected patients. In Pakistan, the burden is increased to these reasons:

1. A Severe shortage of the specialist Pathologists and the weak infrastructure
2. A long diagonistc pipeline that moves far too slowly for a disease in which timing is the everything

Methodology

A. The Data Foundation: Radboud Institute Biopsy Dataset

A medical AI system is only trustworthy only due to the data which was used to train that system. In the medical AI, this principle has the particular importance. A deep learning model trained on the poorly or wrong annotated images will produce the results that are the looking fine but have dangerous impact on the human lives. Getting the good data is the foundation of the OncoLens AI. This deep learning models is the trained and the validated on the prostate biopsy images curated by the Radboud Institute one of the famous medical research universities. The Radboud institute dataset contains high-resolution while slide images of the prostate biopsies, annotated at the tissue level by experienced pathologists, with the Gleason grade assignments reflecting tumor aggressiveness and the clinical significance. Every slide is the created by the trained pathologist with the clinical care. The OncoLens Ai is learning from is as close to expert clinical reality as publicly available data can be.

Preparing these whole slide images for the deep learning training required a multi-stage preprocessing pipeline. High-resolution WSIs were tiled into overlapping patches at multiple levels of the magnification. Initially, the patches of the 120 sides created with the proper balancing of classes that are the Stroma, Healthy and cancerous patches. The final dataset split maintains patient-level separation between training, validation and the held-out tests' sets ensuring the model's performance metrics reflect genuine generalization.

B. The Segmentation Engine: SAM-2

Reading a prostate biopsy slide is not a binary task. A pathologist does not decide "cancer or "no cancer". They trace the boundaries of the suspicious tissue regions across a complex, multi-cellular image and the identification of the where glandular architecture ends and the where malignant patterns begin and the grading of the aggressiveness of the different zone. This is the task OncoLens AI must learn to perform: that is the pixel-level tissue segmentation

The segmentation architecture used for the Oncolens AI is the SAM-2 (segment anything model) in its second major iteration, developed by the Meta AI Research and the extensively adapted by the medical imaging research community for the histopathology applicants. SAM-2 is the transformer-based architecture gives it a critical advantage over earlier segmentation approaches: the SAM-2 processes the full spatial context of an image, rather than operating on isolate local patches. This global context awareness is particularly valuable in biopsy image analysis, where the morphological relationship between adjacent tissue structures often carries as much diagnostic information as the structures themselves.

The model was initializing from the Segmentation Anything Model 2 foundation checkpoint and the finetuned on the patches created by us from the Radboud institute dataset. Training used a combined loss function of Binary Cross-Entropy weighted by inverse class frequency and the Dice Loss to handle class imbalance and improve segmentation accuracy. A learning-rate schedule with warm up and the cosine decay was applied to stabilize the fine-tuning process. During early training the model identified some wrong tissue regions and the abnormal Gleason patterns. As the training progressed on the Radboud dataset, Predictions become more better.

C. The Explainability Layer: Opening the Black Box

The history of the Medica Ai is the full of the systems that were accurate an unused. Systems that achieved impressive benchmark scores, publish their results in the reputable journals, and then sat dormant because the physicians could not understand how to use and the and its difficult for them to take a shift from their regular methods. This problem name is the "Black Box". A model that produces the outputs without any explanation just like a student says I know this answer but has no answer to the question "Why?". The explainability of the model is the innovation in the OncoLens Ai, the explainability layer uses the GRAD-CAM which works by computing the gradient of the model's prediction with respect to the feature maps of its final convolutional layers. These gradients are use to create the spatial heatmaps. Where the heatmap glows red and orange, the model's attention was concentrated. Where it fades to cool blue, the tissue was considered less significant.

Is it responding to genuine pathological change or to a staining artefact? Is there a region it has weighted highly that I should look at more carefully? That capacity for interrogation and override is not a limitation of the system. It is the feature that will determine whether the OncoLens AI is actually used and whether, when it is used, it makes outcomes better

The system is currently facing challenges regarding this but we are working to make it better and the useful for the physicians

Results

OncoLens AI is an active research project. And the recent results reflect the current early-stage evaluation of the system on unseen biopsy data; at this stage the model have hit promising capability in the identifying cancerous tissue regions and the achieving moderation segmentation performance and the cancer detection **AUC-ROC of 0.72** with an average precision of the 0.79. this shows that the system can identify the cancer form non-cancer patters better than chance while still leaving a major room for the improvement

Qualitative inspection of predictions suggests that the model is still learning meaningful pathological signals rather than the simple visual shortcuts. The strongest performance is the observed in the stromal tissue segmentation while the cancer regions are detected with the moderation. The most notable limitation in this current fine-tuned model is the difficulty in distinguishing health glandular tissue and the realty cancer patterns

These findings are the expected for the model because the very less patches were taken at this stage and the high leigh the next research priorities improve the sensitivity for the cancer patterns and the enhancing the class separation between health and the malignant glandular structures Some images of the current training are

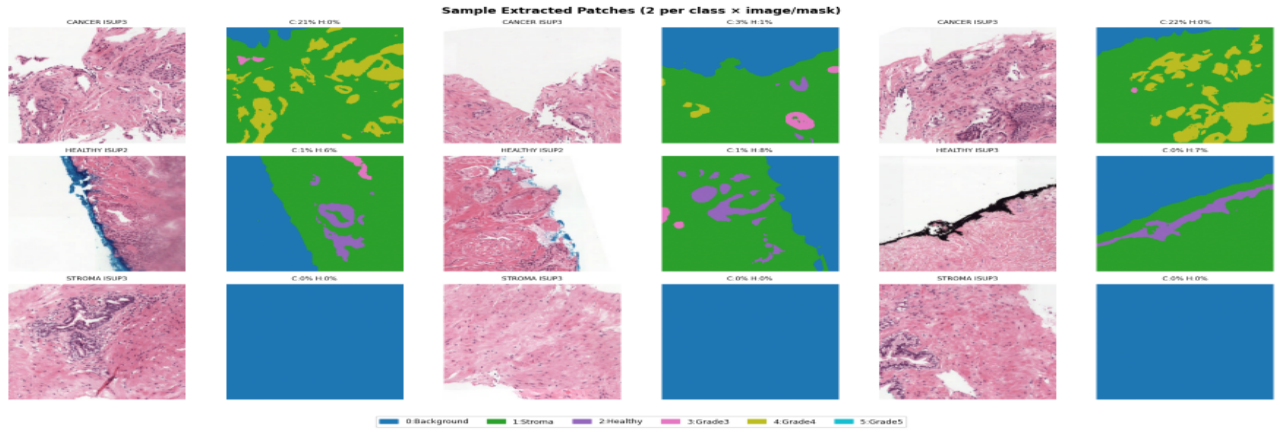

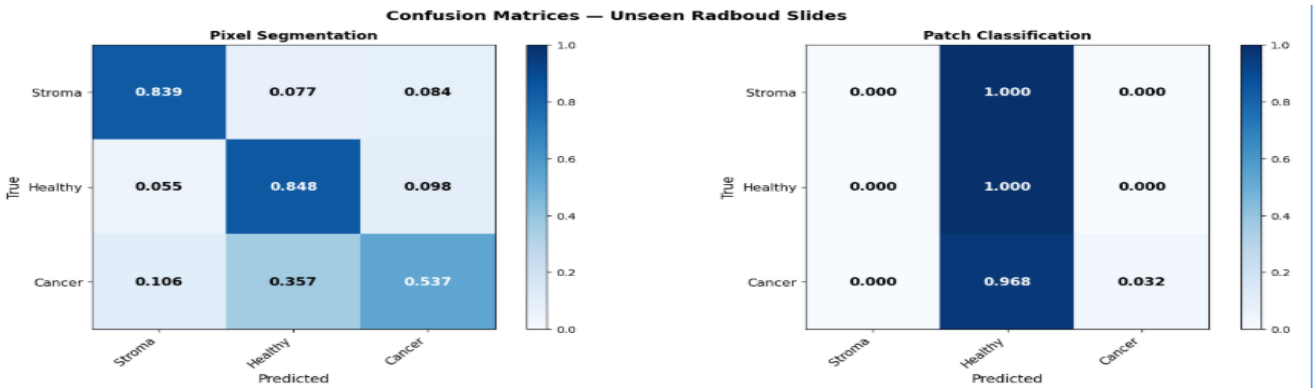

CONFUSTUON MATICES-UNSEEN RADBOUD SLIDES THAT WERE USE TO EVALUATE THE PERFORMANCE OF THE MODEL

Discussion: What OncoLens AI Will Change

Step away from the metrics and consider what it will mean, in practical human terms, when OncoLens AI is fully deployed and working. But we are assured that a general physician in a district hospital or the lab will be able to upload a prostate biopsy image to web application and receive a complete analysis of that image within seconds reducing the time to detect the cancer. The physician will not need the specialized training to interpret this output. They will need only the information that they have and the information OncoLens AI will provide. That combination will allow them to make the perfect decisions for the patient's care. Consider what changes when reliable, explainable biopsy analysis becomes available across Pakistan's district and rural hospitals not just in Lahore and Karachi, but in Dera Ghazi Khan, Zhob, and Chitral. Emergency referrals will become more precisely detected, because physician sending a patient to a specialist will be sending documented evidence, not a worried.

And the training data generated by OncoLens AI's deployment in Pakistani settings will, over time, be used to retrain and localize the model making it the progressively more accurate for the specific population it serves. OncoLens AI will also matter for the pathologists and physicians who use it. AI-assisted workflows in digital pathology have been shown to reduce the overload of the repetitive screening tasks, decrease inter observer variability, and free specialist time for the complex cases where expert judgement is irreplaceable. OncoLens AI is not coming to replace that expertise. It is coming to protect it from being consumed by volume.

About the Author

Muhammad Muteeb Ramzan is the lead developer of OncoLens AI, an ongoing research project developing deep learning solutions for explainable prostate cancer detection from biopsy histopathology images. The project focuses on making advanced medical AI accessible and trustworthy in underserved healthcare settings. Correspondence: muteebramzan3@gmail.com.

References

[1] H. Sung et al., "Global Cancer Statistics 2020: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers in 185 Countries," *CA: A Cancer Journal for Clinicians,* vol. 71, no. 3, pp. 209–249, 2021. doi: 10.3322/caac.21660

[2] World Health Organization, "Pakistan: Health Workforce Data," *WHO Global Health Observatory,* 2023. [Online]. Available: https://www.who.int/data/gho

[3] G. Litjens et al., "Evaluation of prostate segmentation algorithms for MRI: The PROMISE12 challenge," *Medical Image Analysis,* vol. 18, no. 2, pp. 359–373, 2014. doi: 10.1016/j.media.2013.12.002

[4] M. Bulten et al., "Automated deep-learning system for Gleason grading of prostate cancer using biopsies: a diagnostic study," *The Lancet Oncology,* vol. 21, no. 2, pp. 233–241, 2020. doi: 10.1016/S1470-2045(19)30739-9

[5] A. Kirillov et al., "Segment Anything," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV),* pp. 4015–4026, 2023. doi: 10.1109/ICCV51070.2023.00371

[6] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," *International Journal of Computer Vision,* vol. 128, pp. 336–359, 2020. doi: 10.1007/s11263-019-01228-7

[7] Q. Yao et al., "SAM-Med3D: Towards General-purpose Segmentation Models for Volumetric Medical Images," *arXiv preprint,* arXiv:2310.15161, 2023. [Online]. Available: https://arxiv.org/abs/2310.15161

[8] E. J. Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again.* New York: Basic Books, 2019. ISBN: 978-1541644649