Abstract

Early and accurate detection of cervical cancer (CC) is critical for reducing mortality and enabling timely clinical intervention. Conventional screening methods, particularly biopsy-based diagnosis, provide high accuracy but are invasive, costly, and time-consuming. These limitations disproportionately affect women in low-resource and underserved regions, motivating the development of non-invasive, rapid, privacy-preserving, and explainable diagnostic solutions deployable on resource-constrained devices. This PhD thesis presents a unified suite of lightweight personalized federated learning (PFL) frameworks, namely CerviSpectraDiag, CerviImagingDiag, and GeoFed-Cervix, for cervical precancer detection using unimodal and multimodal data.

CerviSpectraDiag focuses on spectral-domain analysis for early-stage CC detection by integrating a lightweight personalized federated CerviSpectraYOLO classifier with CerviSpectraLangChain for explainable diagnostic report generation. The framework is deployed across decentralized local network nodes, preserving patient data privacy while enabling collaborative learning. CerviSpectraDiag achieves 81.83% precision, 81.22% recall, 94.90% specificity, and 84.73% top-1 accuracy, outperforming lightweight personalized federated YOLO baselines, including YOLOv5, YOLOv6, and newer YOLO variants, as well as compact deep learning (DL) models such as MobileNetV3, SqueezeNet, and EfficientNet. The system distinguishes normal tissue from Grade I–III precancerous conditions while generating patient-friendly explanations. The associated web application (app) delivers real-time diagnosis within seconds and supports bilingual explanations in English and Hindi through dedicated clinician and patient dashboards.

CerviImagingDiag extends the framework to image-based diagnosis using differential interference contrast (DIC) cervical images. It employs a lightweight personalized federated CerviImagingYOLO model combined with CerviImagingLangChain to generate interpretable diagnostic reports and treatment guidance. Designed for edge and low-resource devices, CerviImagingDiag achieves 80.41% precision, 81.93% recall, 93.82% specificity, and 81.11% top-1 accuracy, outperforming state-of-the-art lightweight personalized federated YOLO models and compact architectures while producing clear, patient-oriented explanations. The accompanying web app provides rapid diagnostic outputs with bilingual clinician and patient interfaces.

Finally, GeoFed-Cervix introduces a geometry-aware multimodal diagnostic framework that integrates spectral and imaging features within a Riemannian geometry–based PFL paradigm. GeoFed-CervixYOLO incorporates differential geometric priors to address feature heterogeneity and non-Euclidean data distributions, while GeoFed-CervixLangChain leverages large language models (LLMs) aligned with Explainable AI (XAI) 2.0 principles. Implemented as a lightweight Android app for consumer-grade edge devices, GeoFed-Cervix achieves 98.27% precision, 98.28% recall, 99.57% specificity, and 98.27% top-1 accuracy, providing clinically meaningful, bilingual, and personalized explanations for both clinicians and patients.

Collectively, this thesis demonstrates that lightweight, privacy-preserving, explainable, and geometry-aware PFL frameworks can significantly advance early CC detection, enabling scalable, equitable, and patient-centric AI-driven healthcare in real-world, resource-limited settings.

Keywords: Cervical precancer detection; personalized federated learning; lightweight deep learning; geometric deep learning; explainable AI 2.0; multimodal data analysis; YOLO-based models; edge AI; privacy-preserving healthcare; large language models.