SIIM CAIMI25
Conference on Artificial Intelligence in Medical Imaging
OCT 20-21 | UNIVERSITY OF CALIFORNIA, IRVINE

AI on the Edge: Machine Learning for Point-Of-Care Musculoskeletal Ultrasound

Raffi Salibian, MD, Olive View-UCLA Medical Center

Introduction/Background

Knee joint effusions are common in a wide range of clinical scenarios. Ultrasound (US) is a widely used modality for effusion detection, but its accuracy is operator-dependent. Artificial intelligence (AI) can help standardize interpretation, but many AI solutions rely on cloud-based computing, which introduces concerns about data privacy, cost, and dependence on network infrastructure. Edge devices such as the Raspberry Pi and Jetson Nano offer a promising alternative: they are low-cost, portable, and capable of running lightweight AI models directly on-device.This study describes the development of an edge-deployable object detection model for identifying knee joint effusions on ultrasound.

Methods/Intervention

A custom object detection model was trained using a labeled dataset of knee ultrasound images. The YOLOv8 architecture was used as the base model and fine-tuned using transfer learning techniques. Data augmentation strategies (rotation, flipping, brightness variation) were employed to improve generalizability. The trained model was then converted to a lightweight TensorRT/TensorFlow Lite format for deployment on edge devices. Inference time and performance metrics were evaluated on Jetson Nano and Raspberry Pi 5.

Results/Outcome

The model achieved a mean Average Precision (mAP@0.5) of 91.2% on the test set. Sensitivity and specificity for effusion detection were 93.5% and 89.8%, respectively. Inference speed averaged 14 FPS on Jetson Nano and 6 FPS on Raspberry Pi 5, supporting near real-time performance. Qualitative analysis confirmed the model’s ability to localize effusions accurately, even in images with moderate shadowing or gain variation.

Conclusion

This study demonstrates the feasibility of deploying a real-time AI model for knee joint effusion detection on portable edge devices. Such tools can support clinical decision-making in emergency, outpatient, or rural settings, reducing dependence on specialist interpretation. Future work will focus on expanding the model to detect additional joint abnormalities and integrating real-time feedback for ultrasound-guided procedures.

Statement of Impact

By enabling real-time inference without reliance on cloud computing, this system empowers providers to make faster, more accurate decisions at the point of care. The model’s compatibility with compact, widely available hardware like the Raspberry Pi and Jetson Nano makes it highly scalable and accessible, paving the way for the broader adoption of AI-assisted musculoskeletal ultrasound applications.

Keywords

Point-of-Care Ultrasound (POCUS); Musculoskeletal Ultrasound; Edge Computing