Mobile Heart Sound Pathology Detection

Students: Tianqin Li[*1], Zhengyang Geng[*1] Mentors: Asim Smailagic[1], Alex Gaudio[1,2]

Carnegie Mellon University[1], John Hopkins University[2]

TL;DR: Mobile Heart Murmur Detector for enhanced diagnostic efficiency and accessibility in rural area.

Introduction

Figure 1: Left: Family Doctor Screening. Right: Weak Network Connection in Rural Area.

- Cardiovascular diseases (CVDs) cause the death of about **17.9 million** lives worldwide every year.
- The **early detection** of indications of cardiac abnormality is essential for preventing CVD development.
- Detection primarily relies on **expert diagnosis** of the heartbeat sound, which is subjective and limited to its scale.

There are heart murmur detectors/digital stethoscopes...

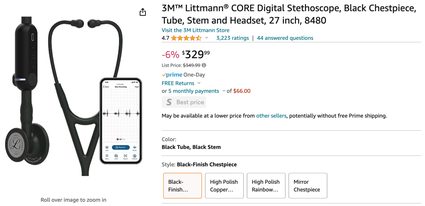

- **But**: The existing murmur detectors rely on cloud resources, take long time to collect data, consume lot of money.

Our work developed...

- **A Mobile Murmur Detector that dynamically allocates resources and hides the inference time into data collection.**
- On-device model inference and murmur diagnosis using a much lower budget and faster speed.
- Noninvasive, low cost, well suited for screening environments.

Mobile Murmur Detector

- **Remark 1**: Data Collection Time ≈ Inference Time. Model is waiting for data!
- **Idea 1**: Model pipelining amortizes the computation into the data collection time.

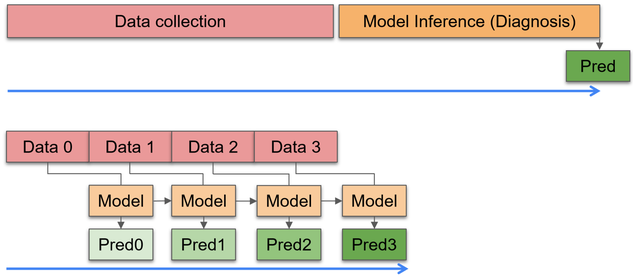

Figure 2: Use an adaptive model to hide the computation into data collection.

- **Remark 2**: Periodical data! The prediction converges as the model accesses more data.
- **Idea 2**: Use the convergence to decide when to stop data collection and computation.

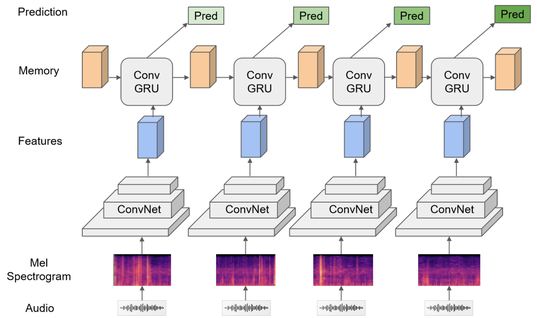

Figure 3: Our murmur detector employs dynamic computation upon convergent (fixed-point) predictions $\mathbf{z}^* = (\mathbf{h}^*, \mathbf{p}^*)$.

	Our Murmur Detector	Baseline / Prior Arts
Resource	On-Device	On-Device / Cloud
Computation	Pipelining & Adaptive	Sequential & Static
Accuracy	80.3%∼81.0%	81.2%
Speed	4.8s∼7.0s	13s

Table 1: The comparison between our mobile murmur detector and prior arts.

System & Results

- Our mobile deployment on Raspberry Pi.

Figure 4: System components of our murmur detector

- Our system acquires much less computation to match the performance of large models and runs directly on mobile devices.

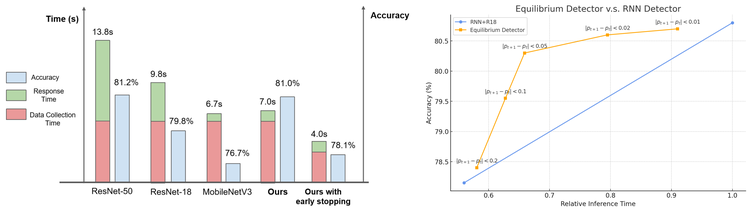

Figure 5: Left: Model pipelining enables much stronger equilibrium models to run on mobile devices. Right: Dynamic early stopping via model convergence entails a significantly smaller computing budget.

Conclusion

Our study presents a **Mobile Heart Murmur Detector**

- designed to be a low-cost, efficient, and accessible solution for cardiovascular health screening, particularly in resource-limited rural settings.
- reduces reliance on cloud resources, thereby enhancing privacy and decreasing operational costs.