Title: Investigating Approximate Computing to design an energy-efficient deep learning architecture for anomaly detection from ECG signals

Cardiovascular diseases (CVDs), including coronary heart disease (CHD), stroke, and other circulatory diseases, contribute to around 30% of global mortality each year. CVD is a primary cause of premature death and the leading factor in morbidity among non-communicable diseases. The financial impact of CVD is significant, with an estimated cost of approximately €169 billion per year in the European Union. This cost includes 62% direct expenses in the healthcare system, as well as productivity loss and informal care. To address the costs and healthcare risks associated with CVD, continuous monitoring of physiological signals, such as electrocardiogram (ECG), using Internet of Things (IoT) enabled wearable devices is being explored as a potential solution. However, the practical implementation of continuous monitoring with medical-grade ECG has been hindered by the high power consumption associated with constant wireless transmission for processing in the cloud. Hence, it might be more beneficial to carry out inferences on the wearable device itself if the inference algorithms are energy and area-efficient.

Approximate computing is a technique that allows controlled errors to be introduced in order to improve computing efficiency such as power-/energy-efficiency. However, the introduction of errors may affect the precision of computing and compromise the quality of the output. The goal of approximate computing is to achieve the best efficiency design while maintaining an acceptable level of output quality. In this project, the focus is on implementing approximate multiply-accumulate (MAC) units to develop a deep-learning based algorithm for detecting heartbeat anomalies. By leveraging the benefits of approximate computing, it is expected that the algorithm can achieve efficient performance while providing sufficient quality in detecting abnormal heartbeats.

Tasks:

- Surveying and evaluating a suitable deep-learning-based anomaly detection algorithm using a suitable Arrhythmia ECG database (like the MIT-BIH arrhythmia ECG database).
- Implementing the model on FPGA to perform hardware cost analysis (chip-area, power consumption, latency)
- Researching approximate MAC techniques promising for deep-learning implementation.
- Optimizing the deep-learning MAC-based architecture for minimum power/energy consumption, which satisfies the output quality criterion of ECG anomaly detection.

Theory: 30%
Coding/implementing: 20%
Evaluation: 30%
Writing: 20%

Contact:

Arlene John, BSS research group, a.john@utwente.nl
Ghayoor Gillani, CAES research group, s.ghayoor.gillani@utwente.nl