Federated Learning Approach for ECG Signal Classification

1D-CNN for Arrhythmia Detection: Federated Learning with Enhanced Privacy in a Multi-Hospital Scenario

Student

Dennis Shushack

Introduction: Machine learning (ML) and Deep Learning (DL) are shaping the landscape of modern medicine. These algorithms have reached a level of proficiency close to that of experienced medical professionals. Federated Learning (FL) supports this transition, promoting a decentralized learning approach without compromising sensitive data. One area that has received significant attention is Electrocardiogram (ECG) diagnostics for cardiac arrhythmias, such as Atrial Fibrillation (AFIB) and Atrial Flutter (AF). One of the challenges in using ECG data is its sensitive nature. This necessitates the use of solutions such as Differential Privacy (DP) to ensure the anonymity of patient data. Although several ECG classification methodologies have been established, a distinct research gap must be addressed. Research on combining one-dimensional convolutional neural networks (1D-CNN), Federated learning, and Differential Privacy for 12-lead ECG three-class classification remains sparse.

Definition of Task: To address the challenges of 12-lead ECG classification, this project designed and implemented a DL & FL-based framework. At its core is a 1D-CNN explicitly designed for the detection of arrhythmias. Tailored to ECG data, the framework utilizes domain-specific data cleaning, processing, and augmentation techniques. The framework operates in both a centralized and a federated setting. Authentic ECG data were obtained to replicate three virtual hospitals, emulating a real-world scenario with genuine ECG recordings. The framework further integrates a differential privacy optimiser, enriched with Rényi Differential Privacy, to enhance the privacy dimensions of the FL procedure. A multitude of experiments were performed to evaluate and compare the framework's efficiencies. These experiments encompassed centralized learning, federated learning, and a strategy incorporating differential private federated learning.

Conclusion: This comprehensive study on 12-lead ECG classification revealed several pivotal findings. The centralized learning experiment demonstrated outstanding results. An impressive 90% accuracy and F1-score were obtained on the aggregated non-augmented dataset. While ECG data augmentation showed potential to improve model performance, its efficacy largely depended on the specific hospital dataset used. Federated learning rivalled the centralized approach for distributed ECG classification, with a negligible performance gap of 1%. However, its integration with differential privacy presented challenges. Attempting to enhance privacy with DP by introducing slight noise dramatically hindered the effectiveness of the model. This underscores the imperative need to find a balance between ensuring privacy and sustaining model performance.

Architectural Overview of DL & FL Framework
Own presentment

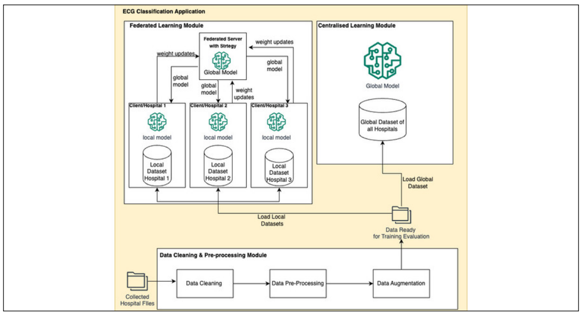

Example ECG Recording of a normal heartbeat (Sinus Rhythm)
P. Davey, ECG at a Glance. John Wiley & Sons, 2013.

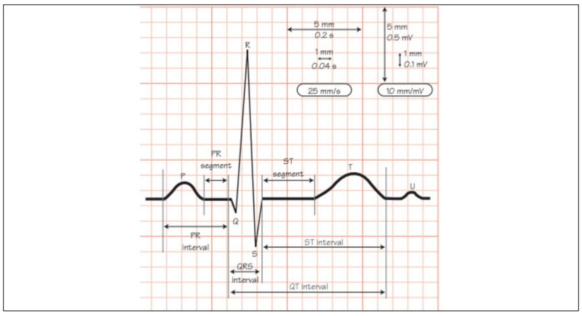

Federated Learning (FL) with ECG data from three hospitals
Own presentment

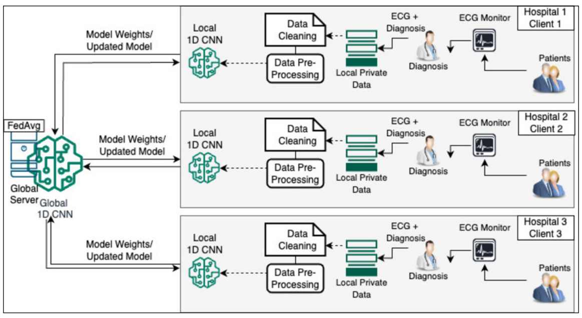

Advisor
Hannes Badertscher

Subject Area
Data Science, Computer Science, Medical Engineering