Abstract

Physiological sensor data, such as those obtained from ECG, PCG, PPG, and VGRF modalities, medical records and metadata contain a rich collection of biomedical indicators that reflect a wide array of underlying medical conditions. Over the past decades, researchers have extensively employed machine learning (ML) and deep learning (DL) techniques to model the relationship between such sensor-derived data and disease-specific biomarkers.

Conventional shallow learning approaches typically involve the extraction of a limited number of handcrafted features, guided by domain expertise. These features may include time-domain statistics, frequency-domain coefficients, or combinations thereof. Although effective to a certain extent, such methods are often constrained by their reliance on feature engineering and limited representational capacity. In contrast, deep learning models automatically learn hierarchical feature representations from raw data. These models are capable of generating millions of parameters, enabling powerful non-linear mappings and classification decisions. However, these features often remain uninterpreted or disconnected, hindering their transparency and clinical interpretability.

To address this gap, the present thesis introduces a graph-based framework to capture the interrelationships between features in a structured and interpretable manner. Specifically, for the task of coronary artery disease (CAD) classification, we construct subject-specific graphs using shallow learning features. Unsupervised testing on photoplethysmogram (PPG) and metadata of 32 participants achieved 88% accuracy, demonstrating superior performance over a baseline and two state-of-the-art methods. Building upon this, we propose a hypergraph-based model that further enhances interpretability by representing individuals as hypergraphs—where hyperedges encapsulate higher-order interactions among multiple features. Validated using multimodal PPG, phonocardiogram (PCG) and metadata, the proposed method demonstrates 92% classification accuracy, driven by 98% sensitivity and 82% specificity.

A hybrid pipeline for Parkinson's Disease (PD) classification merges a novel explainable weighted hypergraph for feature extraction (Shallow Learning) with a ResNet architecture (Deep Learning) using multimodal Vertical Ground Reaction Force (VGRF) signals, gait speed and metadata. This dual-pronged strategy combines the interpretability of the graph-based models with the superior predictive performance of deep networks. The proposed pipeline achieves a superior 0.979 AUC on the PhysioNet Parkinson Dataset compared to isolated shallow (0.878) and deep learning (0.852) methods. The framework improves explainability through enhanced feature importance and demonstrates, via case studies, how the hybrid model rectifies individual component errors.

Furthermore, to improve model performance for highly imbalanced data, such as that found in sepsis prediction, we propose a novel loss function—Focal Weighted Binary Cross Entropy Loss. This approach not only addresses class imbalance but also assigns differential importance to more informative samples. Our sepsis prediction model leverages multimodal sensory data, medication records, and categorical metadata to ensure comprehensive and timely detection. The proposed method achieves a utility score of 0.4305 and 0.393 on two sepsis test sets which is at par with the best performance provided by the state of the art methods. Our SHAP analysis demonstrates that covering a wide range of medical specialties—a feature lacking in competitors—is our most impactful trait. This confirms that our loss function effectively addresses sepsis diagnosis by addressing multiple clinical domains.