Abstract

Heart valve diseases (HVD) are the physiological conditions that prevent the normal functioning of heart valves. Detection of HVD is critical as untreated valve diseases often develop into life-threatening cardiac conditions. Recent development suggests that with signal processing and machine learning (ML)-based algorithms, the heart sound signal, also known as the phonocardiogram (PCG), can effectively predict the anomaly in the valvular activity. However, in a real-world environment, PCG signals are collected from subjects of varying demographics using different acquisition systems and under diverse recording conditions. This leads to substantial variability in the physical properties of the signal, even for the same pathological conditions. Each change in the parameter settings subtly alters the signal and downgrades the performance of an ML model.

This thesis first addresses the challenges of limited training data, computational resource constraints, and the imbalanced class distribution in the PCG classification task. It investigates the merits of transfer learning (TL) using a pre-trained convolutional neural network for automatic HVD detection. Instead of relying on computationally intensive cardiac cycle segmentation, this study introduces overlapping time-frequency fragment selection strategies. The same method later extends to solve class imbalance. A cross-domain and noise-based study has also been conducted to assess the generability of the model across diverse PCG databases.

Then, the thesis investigates the acquisition-induced variability in the PCG signal using six different PCG databases. The stethoscope, auscultation sites, environment, and data distribution have been hypothesized as the significant factors. This part of the study systematically explores a range of features in the time, frequency, joint time-frequency, and deep feature domains to comprehend the cause and pattern of the variation.

The thesis then presents domain-invariant preprocessing, domain-balanced variable hop fragment selection (DBVHFS), and a partial fine-tuning method to enhance

the robustness of the PCG classification under acquisition-induced variability. The domain-invariant preprocessing normalizes the PCG to reduce the stethoscope and environment-induced variations. The TL reduces the impact of data variability by generalizing feature representations. DBVHFS facilitates unbiased fine-tuning of the pre-trained model by balancing the training fragments across all domains and classes.

Furthermore, a loss condition domain adversarial finetuning method has been introduced to enhance the domain invariance during model training. Finally, the thesis incorporates all major propositions into a complete, dedicated system robust to acquisition-induced variability and accurately detects HVD from PCG signals. It implements an automatic disease detection algorithm, an interactive graphical user interface (GUI), and a low-cost plug-and-play hardware platform. The interactive GUI is developed for real-time data acquisition, digital auscultation, multi-domain feature visualization, and analysis to detect HVD under limited resources and at a low cost, targeting regular heart health checkups.

Overall, this thesis addresses the challenges related to acquisition-induced variability under limited resources for HVD detection from the PCG signal. Outcomes of this thesis advance cardiac health screening, offering clinicians a more accurate tool for detecting and managing various cardiac irregularities across diverse practical scenarios.

Keywords: Heart valve disease detection, Phonocardiogram signal, Transfer learning, Signal variability, Cross-domain analysis, Acquisition parameter.