REMINDER

REMINDER:

pRivacy-prEserving Machine LearnIng through secure managemeNt of Data's lifecyclE in distRibuted systems

Deliverable number: D5.1

Retrospective clinical data collection

chist-era

FWF Austrian Science Fund

UKRI Engineering and Physical Sciences Research Council

uefiscdi Unitatea Executivă pentru Finanțarea Învățământului Superior, a Cercetării, Dezvoltării și Inovării

AGENCIA ESTATAL DE INVESTIGACIÓN

Project Acronym:	REMINDER
Project Full Title:	pRivacy-prEserving Machine LearnIng through secure manageMeNt of Data's lifecyclE in distRibuted systems
Call:	Security and Privacy in Decentralised and Distributed Systems (SPiDDS). 2022
Grant Number:	PCI2023-145989-2
Project URL:	https://ants.inf.um.es/en/reminder
Editor:	SIE
Deliverable nature:	Report
Dissemination level:	Public
Delivery Date:	31/08/2024
Authors:	SIE

Table 1: Project details.

Abstract

Recent technological advancements have revolutionized medical practices, driving a transformative shift in the healthcare industry. The rapid growth in computational power and data storage has paved the way for innovative, data-driven healthcare solutions. However, these solutions often rely on access to sensitive patient data, which is strictly regulated by privacy laws. Federated Learning (FL) offers a promising approach by enabling decentralized model training, where patient data remains local and only model updates are shared. Despite this, privacy and security risks persist, as malicious actors may attempt to manipulate the training process or infer sensitive information from the exchanged model updates.

The REMINDER project introduces a novel FL-based framework to tackle these challenges, enhancing distributed systems with multiple security layers. This decentralized approach allows collaborative Machine Learning (ML) model training while safeguarding data privacy. REMINDER proactively addresses security threats by implementing protocols to authenticate legitimate participants and protect against adversarial attacks.

This document presents two clinical applications — aortic coarctation prediction and heart arrhythmia detection from ECG data — and their associated datasets. These use cases, encompassing both time series and tabular data, provide a solid foundation for testing and validating the REMINDER framework within the context of eHealth innovations, supporting the development of secure and privacy-preserving AI solutions for healthcare.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

WPs	Work Packages
ECG	Electrocardiogram
CDS	Clinical Decision Support
AF	Atrial Fibrillation
ML	Machine Learning
CFD	Computational Fluid Dynamics
SBP	Systolic blood pressure
DBP	Diastolic blood pressure
ESBP	End systolic blood pressure
HR	Heart Rate
AAo	Ascending Aorta
TAA	Transverse aortic arch
CoA	Aortic Coarctation

1 Introduction

1.1 Purpose and Scope of the Document

This document presents the retrospective clinical data collection conducted as part of the ongoing efforts to support the integration, evaluation, and validation of privacy-preserving techniques within the healthcare use cases outlined in the project. Specifically, this deliverable focuses on the generation / curation of data relevant to the eHealth use case, which encompasses two critical cardiovascular disease applications: (i) aortic coarctation prediction and (ii) the assessment of heart arrhythmia from Electrocardiogram (ECG) data. The two use cases were chosen in a way that allows the project to leverage different types of data: time series for atrial fibrillation and tabular data for aortic coarctation.

The clinical data generation or collection is a crucial step in ensuring the availability of high-quality, well-organized datasets for use in developing, integrating and testing privacy-preserving methods within decentralized training settings. These datasets will serve as the foundation for developing and evaluating Clinical Decision Support (CDS) systems designed for the targeted cardiovascular conditions. The data will be either synthetically generated or collected and securely stored in a shared location, with access and infrastructure provided by Siemens, leveraging tools and resources developed in previous projects.

The objective of this data collection process is not only to gather or generate relevant clinical information but also to ensure that the data is prepared in a manner that facilitates seamless integration with the privacy-preserving techniques developed across other Work Packages (WPs). This will enable comprehensive testing and validation of these techniques within realistic healthcare environments, ensuring their applicability and robustness.

This document outlines the methodologies employed for the data collection or generation, the types of data, the storage mechanisms, and the tools utilized. It also highlights how this deliverable aligns with the broader project goals, ultimately contributing to the secure and effective use of sensitive healthcare data in clinical decision-making processes.

1.2 Structure of the Document

The document is structured into four main sections. The first section provides an introduction to the deliverable, outlining the objectives, scope, and significance of the work presented. The second section delves into the heart arrhythmia data, discussing its relevance, characteristics, and how it contributes to the overall project. Section 3 focuses on the aortic coarctation use case, describing the specific application and methodology for analyzing this medical condition in the context of the project. Finally, the document concludes with a section summarizing the key insights and implications, offering a comprehensive wrap-up of the work and suggesting potential avenues for future work.

2 Atrial fibrilation dataset

Atrial Fibrillation (AF) is the most common type of arrhythmia, affecting millions of individuals worldwide. It is characterized by an irregular and often rapid heart rate due to abnormal electrical signals in the atria. This leads to inefficient blood flow, causing the heart to beat erratically rather than in a coordinated manner. The condition is often associated with an increased risk of stroke, heart failure, and other cardiovascular complications. Given its prevalence and potential to lead to serious health issues, atrial fibrillation remains a major focus of cardiovascular research.

2.1 Data collection

The ECG dataset provided for the Physionet Challenge [1, 2] serves as a foundation for the development of a privacy-preserving machine learning framework. This dataset, collected using AliveCor's single-channel ECG device, comprises a total of 8,528 recordings. These recordings are classified into several categories: "Normal," "Atrial Fibrillation," "Other Abnormal Rhythms," and "Noise." Specifically, 5,076 recordings are labeled as "Normal," 758 as "Atrial Fibrillation," 2,415 as "Other," and 279 as "Noise." Figure 1 provides a visual representation of the inherent class distribution in this dataset.

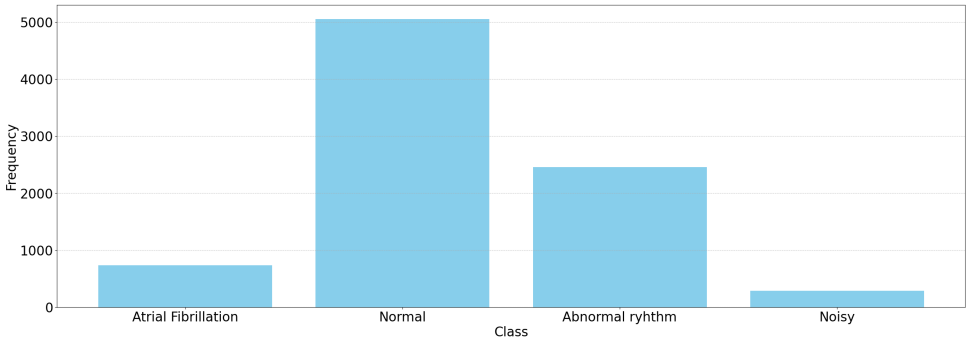

Figure 1: Class distribution of the PhysioNet data

2.2 Input data specifications

The dataset comprises single-lead electrocardiogram (ECG) recordings, which captures sequential data reflecting the heart's electrical activity. Each recording varies in length, spanning from a minimum of 9 seconds to a maximum of 60 seconds. The ECG signals were sampled at a frequency of 300 Hz, ensuring a sufficiently high resolution to capture the fine details of cardiac rhythms. To enhance the quality of the recordings, the AliveCor device applied band-pass filtering, effectively reducing noise and isolating the relevant frequency components associated with heart activity.

All recordings are provided in a MATLAB V4 WFDB-compliant format, a standard widely used for storing and processing physiological signals. Each ECG recording is accompanied by two files: a .mat file containing the raw ECG data and a corresponding .hea file that includes essential waveform metadata, such as sampling information and signal attributes. This structured format facilitates seamless integration with existing signal processing tools, enabling further analysis and model development.

2.3 Annotation specification

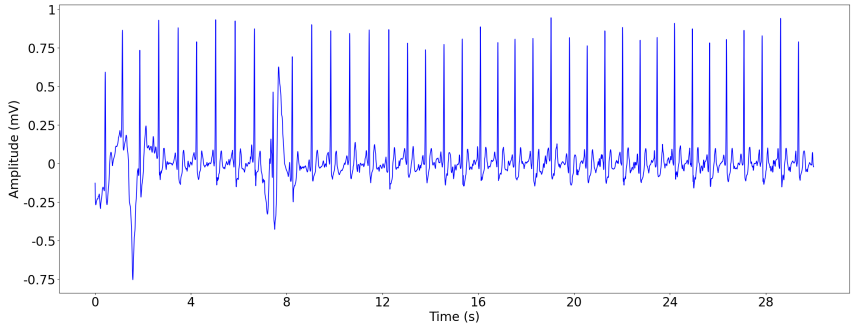

(a) Normal ECG

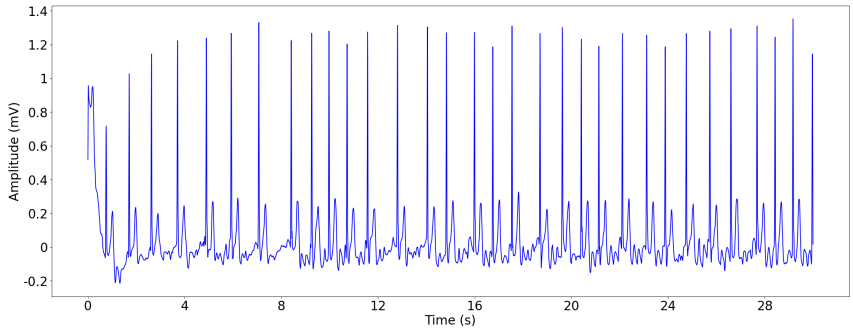

(b) Atrial Fibrillation ECG

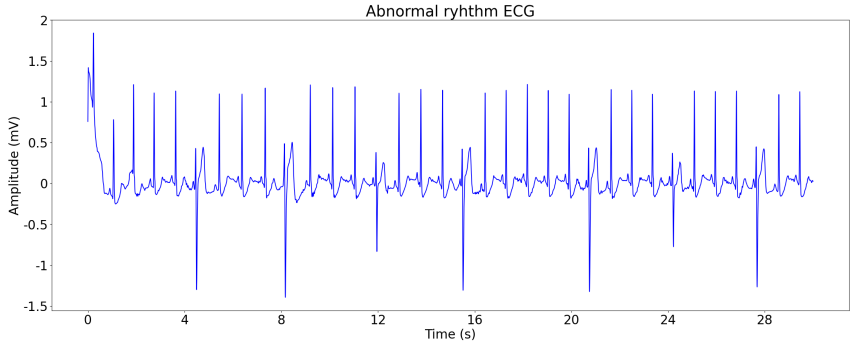

(c) Other abnormal rhythm

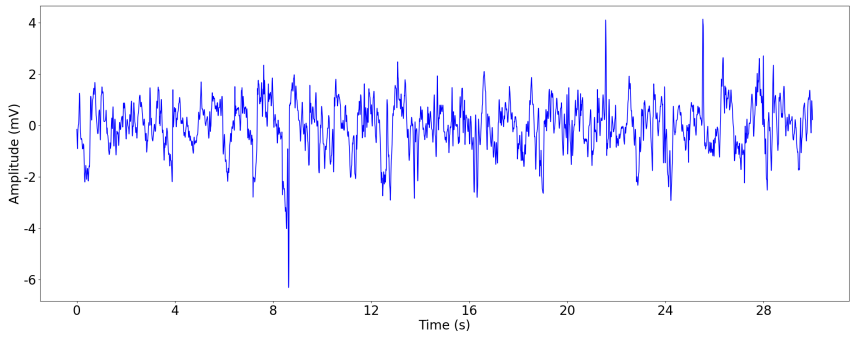

(d) Noisy ECG

Figure 2: Examples of different classes of ECG recordings

As depicted in figure 2, the dataset includes detailed annotations for each ECG recording, providing essential waveform information that serves as a label to categorize the heart's elec-

trical activity. Each recording is assigned to one of four distinct classes based on its rhythm characteristics, ensuring a clear and systematic classification of cardiac patterns.

The first class is Normal Sinus Rhythm (denoted by the symbol N), which represents a regular and consistent heart rhythm, characteristic of a healthy heart (figure 2a). This category indicates that the electrical signals are following the normal conduction pathway, resulting in a steady heartbeat without any irregularities.

The second class is Atrial Fibrillation (AF) (marked by the symbol A), where, as depicted in figure 2b, the ECG displays an irregular rhythm consistent with atrial fibrillation. This condition is identified by the absence of distinct P waves and the presence of an erratic R-R interval, reflecting the chaotic electrical activity in the atria.

The third class, ilustrated in figure 2c, is Alternative Rhythm (represented by the symbol O), which includes any heart rhythm that does not fit into the categories of normal sinus rhythm or atrial fibrillation. This class encompasses a variety of arrhythmias or other abnormal patterns, allowing for flexibility in capturing less common but clinically relevant cardiac irregularities.

The final class is Too Noisy (indicated by the symbol ~), used for recordings where the signal quality is too poor to allow for a reliable classification. An example is privded in figure 2d. Noise in these recordings may result from various factors such as motion artifacts, poor electrode contact, or external interference, making it difficult to accurately assess the heart's rhythm.

These annotations play a crucial role in training and evaluating machine learning models, as they provide the ground truth labels necessary for supervised learning tasks. The clear distinction between rhythm types ensures that models can learn to differentiate between normal and abnormal heart activity, ultimately contributing to the development of more accurate and reliable diagnostic tools.

2.4 Risks and mitigation strategies

In spite of the clear benefits stemming from leveraging AI algorithms in early disease detection, training Machine Learning (ML) algorithms by requiring data be sent to a server introduces several privacy risks, particularly due to the sensitive nature of the collected data, such as health metrics, location, etc. Privacy concerns stem from personal information leakage, that could expose sensitive details about the patient. Nonetheless, even in federated learning scenario where raw data is not externalized for training, patient re-identification risks exist through a series of attacks. Moreover, from a security standpoint, malicious actors could intefere with the training process by performing poisoning attacks, which could significantly degrade system's performance.

The REMINDER framework enhances its federated learning (FL) architecture by incorporating multiple layers of security to address key privacy and security concerns. To counter adversarial manipulation, it uses digital signatures to authenticate model updates, ensuring that only legitimate contributions are integrated into the global model. Additionally, REMINDER employs sophisticated aggregation techniques, such as Fast Fourier Transform (FFT)-based methods, to detect and filter out anomalous updates, effectively reducing the risk of model poisoning attacks.

In terms of data privacy, the framework implements differential privacy by introducing carefully calibrated noise to model gradients. This approach prevents the extraction of sensitive data

patterns while maintaining the integrity of the learning process. By combining these strategies, REMINDER not only strengthens privacy protections but also enhances the robustness and reliability of federated learning, ensuring that personal data remains secure without compromising model performance.

3 Aortic Coarctation dataset

Aortic coarctation is a congenital heart defect characterized by a narrowing of the aortic media, resulting in a reduced aortic lumen. It represents approximately 5% to 8% of all congenital heart disease cases. Diagnosis typically focuses on evaluating the pressure gradient across the coarctation, which can be estimated using non-invasive methods such as cuff-derived blood pressure differences between the upper and lower extremities, Doppler-based measurements applying the Bernoulli equation, or catheter-based assessments—the latter being the clinical gold standard.

In pursuit of a non-invasive alternative to catheterization, [3] proposed the use of personalized blood flow simulations through Computational Fluid Dynamics (CFD) to estimate the pressure drop. While innovative, this approach presents significant limitations: (1) reconstructing anatomical models from medical images involves semi-automated processes that still require manual adjustments, and (2) CFD simulations are computationally intensive and time-consuming, posing practical challenges for real-time or routine clinical use.

A potential solution involves utilizing a streamlined set of fully automated anatomical measurements and employing deep learning models to predict CFD outputs. However, training these models relies on patient-specific anatomical data, which is tightly regulated to protect patient privacy, posing a challenge for data access and usage.

3.1 Synthetic data generation

Due to the limited number of patient-specific cases, developing a robust training dataset for deep learning models presents a challenge. To support the current project's development, synthetic datasets have been generated by randomly sampling aortic measurements within predefined physiological ranges. These measurements are used to create anatomical models, which are then processed through CFD simulations [3] to produce ground truth values.

Since access to invasive, catheter-based measurements is currently unfeasible, the deep learning model will be trained to mimic CFD outputs, with its accuracy constrained by the precision of the CFD simulations. However, this synthetic data-driven approach serves as a foundational step. The results of this project will pave the way for decentralized training across hospitals, allowing future models to learn directly from real patient data without compromising privacy, ultimately enhancing their clinical relevance and performance.

Therefore, a synthetic training database comprising 6218 aortic coarctation cases has been generated. This dataset includes two key files: InputFeatures.csv and Output.csv. The InputFeatures.csv file contains the input features for each case, with one case represented per line, capturing the anatomical measurements sampled within predefined physiological ranges. The Output.csv file provides the corresponding output measures of interest derived from computational fluid dynamics simulations, spanning the entire heart cycle for each case.

3.2 Input data specifications

Synthetic datasets were created by randomly sampling aortic measurements within specified physiological ranges, as shown in Figure 3. The deep learning algorithm should consider 20 anatomical and physiological measurements, outlined in Table 2, to accurately predict the pressure drops across the coarctation.

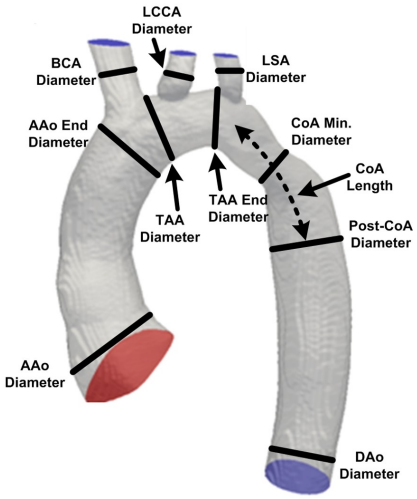

Figure 3: Reduced set of anatomical features that can be leveraged by a DL algorithm.

To ensure a coherent anatomical and physiological measurements, the following constraints were applied in the sampling strategy:

- Systolic blood pressure (SBP): integer between 80 and 180 mmHg
- Diastolic blood pressure (DBP): integer between 30 and $(SBP - 30)$ mmHg
- End systolic blood pressure (ESBP): integer between $(DBP + 0.2 * (SBP - DBP))$ and $(DBP + 0.8 * (SBP - DBP))$ mmHg
- Heart Rate (HR): integer between 30 and 180
- Systolic duration: integer between 30 and 70
- Ascending Aorta (AAo) diameter: float between 1 and 3.5
- AAoMax diameter: $(p/100) * AAoDiameter + AAoDiameter$, where p is a float between 5 and 25
- AAoEndDiameter: float between $0.6 * AAoDiameter$ and $1.1 * AAoDiameter$
- Transverse aortic arch (TAA) diameter: float between $0.9 * AAoEndDiameter$ and $AAoEndDiameter$
- TAAEnd diameter: float between $0.7 * TAADiameter$ and $1.1 * TAADiameter$
- Aortic Coarctation (CoA) Min Diameter: float between $0.2 * TAAEndDiameter$ and $0.8 * TAAEndDiameter$
- CoA Length: float between 0.5 and 6
- PostCoA Diameter: float between $0.85 * TAAEndDiameter$ and $1.1 * TAAEndDiameter$
- DAo Diameter: float between $0.85 * postCoADiameter$ and $1.15 * postCoADiameter$
- BCA Diameter: float between $0.4 * DAoDiameter$ and $0.8 * DAoDiameter$
- LCCA Diameter: float between $0.4 * DAoDiameter$ and $0.8 * DAoDiameter$
- LSA Diameter: float between $0.4 * DAoDiameter$ and $0.8 * DAoDiameter$
- Cardiac Output: float between 1 and 8

- Stiffness - $(3/4) * \beta$, where $\beta = (PSys - PDia)/(1 - \sqrt{ADia/ASys})$, $PSys = SBP * 1333.22$, $PDia = SBP * 1333.22$, $Asys = AAoMaxDim^2/4$ and $ADia = AAoDim^2/4$
- Percentage flow: $100*(DAoDiam/2)^2/[(DAoDiam/2)^2+(BCADiam/2)^2+(LCCADiam/2)^2+(LSADiam/2)^2]$

Table 2: Input features for DL.

Name	Measurement type	Unit	Description
SBP	Physiological	mmHg	systolic blood pressure (measured at the left arm)
DBP	Physiological	mmHg	diastolic blood pressure (measured at the left arm)
ESBP	Physiological	mmHg	end systolic blood pressure (measured at the left arm)
HR	Physiological	bpm	heart rate
Systolic duration	Physiological	[%]	-
Cardiac Output	Physiological	l/min	-
Percentage flow descending aorta	Physiological	[%]	percentage of flow in through the descending aorta from the total cardiac output
AAoDiameter	Anatomical	cm	diameter at the inlet of the ascending aorta
AAoEndDiameter	Anatomical	cm	diameter at the outlet of the ascending aorta
TAADiameter	Anatomical	cm	diameter at the inlet of the transverse aortic arch
TAAEndDiameter	Anatomical	cm	diameter at the outlet of the transverse aortic arch
CoAMinDiameter	Anatomical	cm	minimum diameter of the aortic coarctation
CoALength	Anatomical	cm	length of the aortic coarctation
postCoaDiameter	Anatomical	cm	diameter at the outlet of the aortic coarctation
DAoDiameter	Anatomical	cm	diameter at the outlet of the descending aorta
BCADiameter	Anatomical	cm	diameter of the brachiocephalic artery
LCCADiameter	Anatomical	cm	diameter of the left common carotid artery
LSADiameter	Anatomical	cm	diameter of the left subclavian artery
AAoMaxDiameter	Anatomical	cm	maximum diameter at the inlet of the ascending aorta (at peak systole)
Stiffness	Other	$g/(cm \cdot s^2)$	stiffness of the aortic wall

3.3 Ground truth specification

The measurements enumerated in section 3.2 are used to create anatomical models, which are then processed through CFD simulations to produce ground truth values. Table 3 enumerates all values of interest along with their description.

Table 3: Ground truth values for DL.

Name	Output type	Unit	Description
$PP\Delta P_{AAo-DAo}$	Pressure drop	mmHg	peak-to-peak pressure drop between ascending aorta and descending aorta
$Avg\Delta P_{AAo-DAo}$	Pressure drop	mmHg	average pressure drop between ascending aorta and descending aorta
$PP\Delta P_{TAA-DAo}$	Pressure drop	mmHg	peak-to-peak pressure drop between transverse aortic arch and descending aorta
$Avg\Delta P_{TAA-DAo}$	Pressure drop	mmHg	average pressure drop between transverse aortic arch and descending aorta

3.4 Risks and mitigation strategies

Although we currently use synthetic dataset samples to support the development of this project, the resulting system will facilitate training on real patient data. Since the information required by the deep learning algorithm contains sensitive health information, it is subject to regulatory constraints. To address this, REMINDER will utilize the Federated Learning framework, enabling a decentralized training approach that does not require externalizing the data.

Even in decentralized training settings where the data itself is not externalized, there is a risk that sensitive health information can be indirectly leaked through model updates. In this case, attackers can try to reverse-engineer information about the data by examining the model updates received from each hospital. If an attacker gains access to the model updates or gradients, they might be able to reconstruct sensitive health data from the shared model parameters.

Nonetheless, malicious actors could disrupt the training process by executing poisoning attacks, intentionally degrading the system's performance. These poisoned updates may impair the model's accuracy or introduce biases that jeopardize patient privacy or result in unsafe medical recommendations.

As a result, REMINDER will incorporate multiple techniques to maintain a secure and privacy-preserving environment, including secure aggregation to protect against poisoning attacks, digital signatures to verify that only authorized users can participate in the learning process, and differential privacy to minimize the risk of sensitive information leakage.

4 Conclusions

The REMINDER project aims to enhance the Federated Learning framework by incorporating multiple security layers, ensuring a reliable, robust, privacy-preserving, and trustworthy training environment. Although healthcare data is abundant, it is subject to strict regulatory constraints due to the sensitive nature of patient information. Consequently, this deliverable explores two eHealth-related applications that will inform developments in other work packages to achieve the project's objectives. Specifically, time series ECG data, presented in Section 2, can be utilized to evaluate the effectiveness of the REMINDER framework in training a heart arrhythmia detection model in a secure, decentralized environment, while anatomical and physiological data stored in tabular format (presented in section 3) can be used for diagnosing aortic coarctation.

In future activities, the project will investigate methodologies to enhance the efficiency, scalability, and resilience of the FL framework. The upcoming work will concentrate on refining security mechanisms, implementing privacy-preserving techniques, and performing experimental validation in simulated eHealth environments using, but not limiting to, the datasets presented here.

References

[1] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C. K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals.," *Circulation*, vol. 101, pp. E215–20, 6 2000.

[2] G. D. Clifford, C. Liu, B. Moody, L.-W. H. Lehman, I. Silva, Q. Li, A. E. Johnson, and R. G. Mark, "Af classification from a short single lead ecg recording: the physionet/computing in cardiology challenge 2017.," *Computing in cardiology*, vol. 44, 9 2017.

[3] L. Itu, P. Sharma, K. Ralovich, V. Mihalef, R. Ionasec, A. Everett, R. Ringel, A. Kamen, and D. Comaniciu, "Non-invasive hemodynamic assessment of aortic coarctation: Validation with in vivo measurements," *Annals of biomedical engineering*, vol. 41, pp. 669–681, Apr. 2013.