A Data-Driven Framework for Pediatric Cardiac Arrhythmia Detection

Hossein Babaei, Yanwan Dai, Ahmed Humayun,
Mario Paciuc, & McKell Stauffer

December 2019

Abstract

Arrhythmias can be lethal for children in the period following cardiac surgery, and junctional ectopic tachycardia (JET) is considered the most common type of tachycardia seen during early post-operative care. We present a novel classification algorithm that detects the JET onset based on electrocardiogram (ECG) waveforms. Our algorithm obtains an average cross-validation sensitivity of 87.07%, specificity of 87.12% and area under the receiver operating charactieristic curve (AUROC) of 90.71% on a dataset of 9 patients from Texas Children's Hospital. In addition, we present a "human in loop" development pipeline by creating a novel waveform visualization system. This pipeline will enable cardiac surgeons to better recognize and label arrhythmia onset.

1 Introduction

Congenital heart diseases are among the most common birth defects, affecting ∼1% of live births in the United States [1]. Of the post-operative pediatric cardiac patients, up to 48% develop post-operative arrhythmias [2]. Current arrhythmia detection algorithms are based on electrocardiogram (ECG) waveforms and result in a staggering number of false alarms (∼72-99% of clinical alarms are false [3]). Junctional ectopic tachycardia (JET) is considered the most common type of tachycardia seen during early post-operative care [2] and is very dangerous and difficult to treat in an infant heart [4].

In this paper, we present a novel JET detection algorithm. In this section, we provide a brief overview of cardiac physiology, cardiac monitoring waveforms, junctional rhythms, junctional ectopic tachycardia, and the current state of alarm fatigue. We will then conclude this section with our precise problem definition.

1.1 Cardiac Physiology

The movement of blood throughout the heart is dependant on muscle contractions governed by electrical signals. These signals start in a small region in the right atrium called the Sinoatrial (SA) node (Refer to Figure 1 for a depiction of heart anatomy). Firing of the SA node stimulates the right and left atria to contract, opening the passageway for blood to flow down into the ventricles. Then, the signal arrives at another small region between the atria and the ventricles called the Atrioventricular (AV) node. The electric signal travels down the heart via specialized cardiac cells and starts spreading all over the right and left ventricles. This causes the ventricles to contract, pushing the blood out of the heart. For a healthy heart, the natural cycle of electrical activity is called the sinus rhythm (which can be seen as an electrical waveform in Figure 2). An arrhythmia is a deviation from the sinus rhythm, causing an irregular heartbeat.

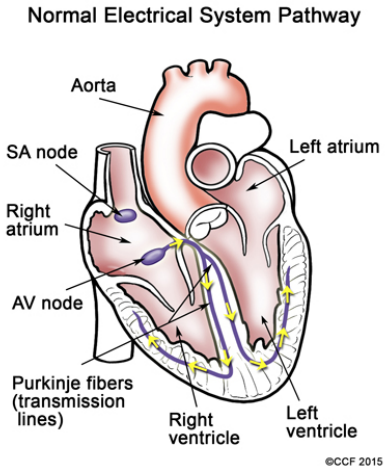

Figure 1: Basic heart anatomy [5].

1.2 Cardiac Monitoring Waveforms

Electrocardiogram The electrocardiogram (ECG) measures the electrical activity of the heart, and measurements are obtained by electrodes which are placed on a patient's chest. Every electrode is called a lead and records the voltage signal on the body surface. Each lead forms a separate channel in the ECG waveform. There are several standards for ECG recording, and each standard determines the number of leads and their placement. The 12-lead ECG is the most commonly used standard [6], while health gadgets and smartphones use less rigorous methods. The ECG waveform can be divided into modes called P, Q, R, S & T (Figure 2), each of which represent a different portion of the electrical activity of the heart. Due to physical movements and interference from other signals (such as the electromyogram), the ECG signal may contain noise artifacts.

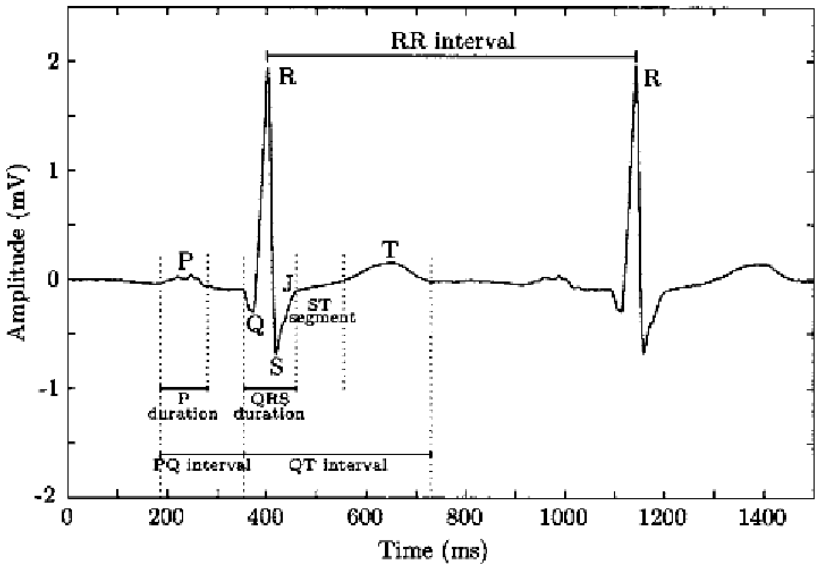

Figure 2: An example of an ECG waveform [7].

Central Venous Pressure Central venous pressure (CVP) is the average blood pressure measured in the superior vena cava, one of the large vessels that returns blood from the body to the right atrium [8]. It reflects the amount of the blood that is returned to the heart [9]. The CVP is measured using a catheter that is placed during heart surgery.

The CVP waveform morphology of a normal cardiac cycle is shown in Figure 3.

The CVP is a mechanical signal; therefore, it is generally invariant to artifacts caused by an electrical signal interference. It is transduced as pressure readings and also used for vitals estimation. However, the catheter is placed in body fluid, so CVP measurements can be affected by the intermittent body fluid infusion running through the lumen of the central line that is monitoring CVP. Pathological changes in CVP are usually noticed as blood pressure increases at the onset of JET. However, the morphology changes in CVP waveforms at JET onset can vary. One of the goals of our work is to understand the transition behaviour of the blood pressure with respect to JET.

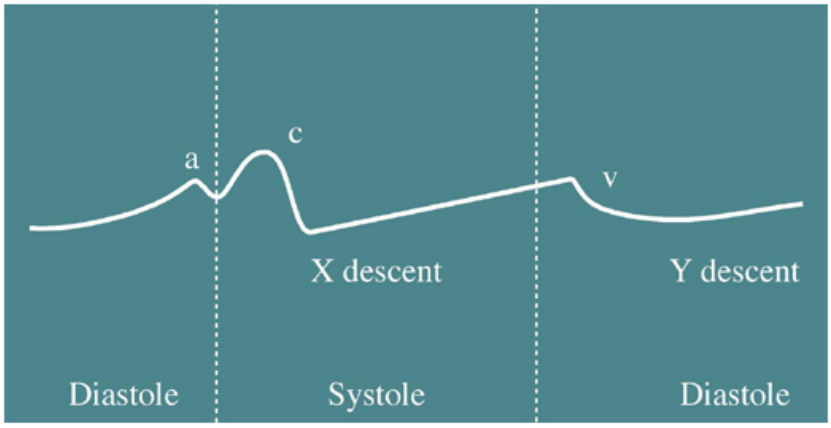

Figure 3: An example of an CVP waveform [10].

1.3 Junctional Rhythms

The electrical activity of the heart is measured through the ECG signal (Section 1.2). This signal is periodic in nature and contains waves that correspond to specific electrical activities in the heart (Figure 2). The P wave is the part of the waveform that corresponds to the firing of the SA node (the first electrical firing of a normal heart beat). A junctional rhythm is characterized by the absence of the P wave in the ECG signal. During a junctional rhythm, the AV node starts the electrical firing while the SA node

fires later, causing the P wave to become absorbed into the R wave (Figure 4).

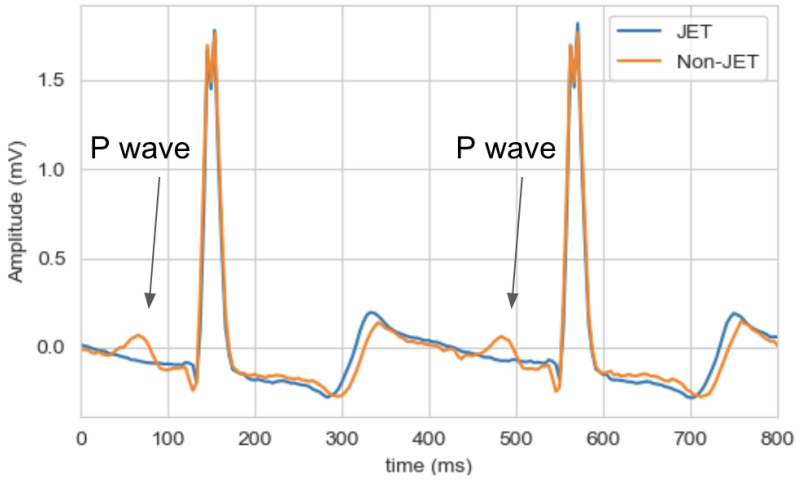

Figure 4: In junctional rhythms, the P wave gets absorbed into the R wave.

1.4 Junctional Ectopic Tachycardia

Any change in a regular heartbeat can be considered an arrhythmia (Section 1.1). Arrhythmias can be fatal in post-operative children; thus, accurate detection of such arrhythmias is absolutely imperative [11]. Tachycardia is an arrhytmia characterized by a fast heart beat [12]. Tachycardias are especially dangerous for post-operative children, as their hearts are already under stress following the surgery [4]. The most common post-operative tachycardia, seen in ∼2-11.2% of post-operative pediatric cardiac patients, is called Junctional Ectopic Tachycardia (JET) [2]. It is characterized by both a junctional rhythm (Section 1.3) and tachycardia.

JET is caused in post-operative patients by inflammation and swelling in the area around the AV node [4]. A source near the AV node begins firing electricity rapidly, causing a dramatic increase in heart rate (in infants, the heart rate may become as fast as 250-350 beats per minute). This puts the infants at severe risk for congestive heart failure. Thus, immediate medical intervention is needed to decrease the heart rate.

1.5 Alarm Fatigue

In most hospitals, there are alarm systems that monitor patients' cardiac waveforms and notify doctors when heart abnormalities are detected. However, these alarm systems are not completely precise and often set off many false alarms. The joint commission determined that 85-99% of alarms created by these systems are actually false alarms [3]. In fact, doctors and patients hear so many alarms that they are often turned off, a practice that is incredibly dangerous for the patients. Better detection methods would eliminate this "alarm fatigue", thus giving patients better care.

1.6 Problem Definition

Failure to intervene within three minutes of JET onset can prove fatal for post-operative infants. Current detection systems are imprecise and result in too many false alarms. Currently, in order to detect JET, a pediatric cardiologist must recognize it from a noisy ECG signal while monitoring many patients' signals at once. To aid pediatric cardiologists, we aim to build a model that will detect the onset of JET within a 30 second window. Our model will take into account the presence (or lack thereof) of P waves in the ECG morphology and the CVP waveform. To the best of our knowledge, we are the first to incorporate signals other than ECG into a JET detection system. We believe that multimodal JET detection will be more robust compared to current systems which utilize only ECG.

2 Literature Review

Arrhythmia detection from ECG signals has been a topic of research for many years. However, past works have not incorporated other signals, so our review focuses on methods that use only ECG signals.

Older algorithms used feature extraction, selection, reduction, signal processing and hand-engineered heuristics. Because these methods showed a substantial amount of misdiagnosis, recent research has focused on using more complex machine learning algorithms. In 2016, Luz et al. concluded that the most popular algorithms for cardiac arrhythmia detection were support vector machines (SVMs), artificial neural networks (ANNs), linear discriminants (LD) and Reservoir Computing With Logistic Regression (RC).

A more in-depth review of these approaches can be found in [13]. To complement their work, we review models created since 2016.

Rajpurkar et al. [14] trained a convolutional neural network (CNN) on ECG readings of about 54,000 patients. Their goal was to classify arrhythmias into 12 rhythm classes. Their model outperformed a team of trained cardiologists on both precision and recall. Similarly, Isin et al. [15] achieved a 98.51% accuracy on three cardiac conditions using only ECG waveforms. For their model, they used AlexNet as a feature extractor and then cascaded these features into a conventional CNN. Rajput et al. [16] filtered, segmented, and used Morlet wavelet and short-time Fourier transforms to preprocess the ECG signals before they were passed to a CNN. Their model outperformed cardiologists on 9/10 arrhythmias using the F1-score metric. [17] used no preprocessing and combined a CNN and LSTM to predict atrial fibrillation with a 83% F-score. [18] used wavelets and a LSTM to get near 99% accuracy on a dataset containing 6 different kinds of arrhythmias, outperforming previous work.

In addition to general arrhythmia detection research, there has also been research focusing on P wave detection. [19] used an extended Kalman filter to estimate the state variables of the equations modelling the dynamics of the ECG signal. They were able to achieve 98.38% sensitivity and 96.75% precision on MIT-BIH Arrhythmia database from PhysioNet. On the same database, [20] used a method "based on two moving average filters, followed by a dynamic event-related threshold" and achieved 98.0.5% sensitivity and 97.11% precision. [21] used phase free stationary wavelets and achieved a -0.32± 12.41 ms error when compared to manual annotations on the QT database from PhysioNet. On the PTB database from PhysioNet, [22] achieved 97.78% sensitivity and 96.8% precision using slope detection in FPGAs. [23] achieved 96.47% sensitivity on the CSE database by removing the QRS complex (the combination of the Q, R, and S waves) from the signal, combining the signal from all leads, applying Teager Energy Operator and then thresholding. In all of the above mentioned research papers on P wave detection, the authors focused on sensitivity and precision instead of specificity as true positives are more valuable to doctors than true negatives. In contrast, we include specificity in our results as we aim to minimize false alarm rates in addition to maximizing true positives. Although all of these P wave detectors work well, our main goal is to build an arrhythmia detector, not a P wave detector. We include their research here as a proof of concept that creating a JET detection algorithm is feasible.

3 Data

We use a dataset that was gathered by Texas Children's Hospital at the Texas Medical Center in Houston, TX, using the Sickbay platform. It consists of about 1190 hours of signals from nine post-cardiac surgery children who developed JET. Every patient's recording lasts for approximately eight days and contains one instance of JET that starts about 72 hours in and lasts for 4-53 minutes.

Each recording is split into four hour time chunks. All chunks contain signals that have been synchronized, normalized, coded and de-indentified. Each signal corresponds to a channel as follows:

1. Four channels of ECG signals collected from four leads. All ECG signals are measured at 240 Hz;
2. A channel containing the CVP waveform, measured at 120 Hz;
3. A corresponding time channel.

Our dataset is missing a large amount of CVP data. The signal is completely missing for Patient 4 and in other patients is only missing for specific time chunks (Table 3). In our models that incorporate the CVP signal, we only train on the time chunks where it is present. We do not have any missing values for the ECG or time channels.

	1	2	3	4	5	6	7	8	9
CVP	100%	53%	74%	0%	91%	67%	25%	96%	44%

Table 1: This table contains the percentage of available CVP data for each patient (enumerated along the top row.

The labels for our dataset were generated by timestamps given to us by medical professionals. As 1190 hours of data is impossible to completely label, we only received one pair of starting and ending timestamps for both JET and a normal heart beat rhythm for each patient. This lead to a massively unlabelled dataset with only 1.6% of the time indices being labelled. Of this labelled data, 27.5% of it was labelled JET and the other 72.5% was labelled non-JET. To help mitigate this data labelling problem in the future, we created a data visualization and annotation tool that we present in Section 5.3.

4 Data Preprocessing

4.1 Corrupted Data Rejection

When leads are not attached to the patient, abnormal signal patterns arise. In particular, ECG measurements either remain constant at extremely low values or oscillate back and forth between these low values and the values we commonly see when the leads are attached. To identify these instances, we segmented the ECG data into 10-second windows and found all windows in which ECG remained constant. These constant values were classified as corrupted, and all windows containing these corrupted values were discarded (Figure 5).

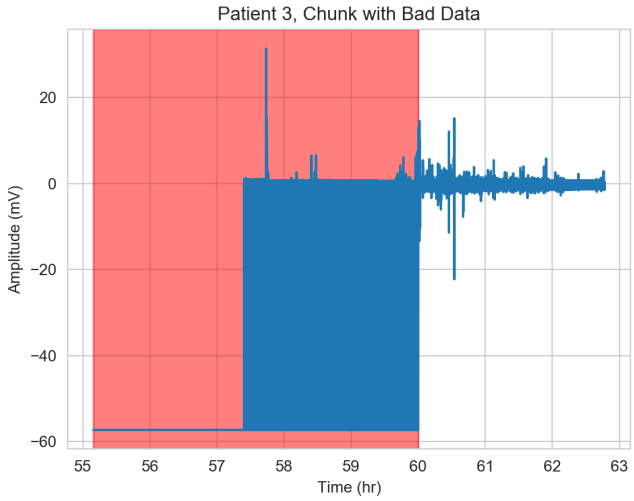

Figure 5: An example of corrupted data that was rejected in Patient 3. The red box is where our corrupt data rejection algorithm determined the corrupted data to be.

4.2 Frequency Filtering

The measuring of electrical cardiac rhythms can be confounded by other general body rhythms, such as the respiratory system. A general heart rate is 0.67–5 Hz and a general QRS complex is 10–50 Hz [24, 25]. As such, we use a 0.67 Hz high pass filter and a 50 Hz low pass filter to get rid of un-

wanted ECG frequencies (Figure 6). For the CVP, we use a 1.5 Hz high pass and a 50 Hz low pass filter. For reference, a high pass filter passes signals with frequencies higher than a specific threshold and attenuates signals lower than the threshold. In contrast, a low pass filter passes signals with frequencies lower than a specific threshold and attenuates signals higher than the threshold.

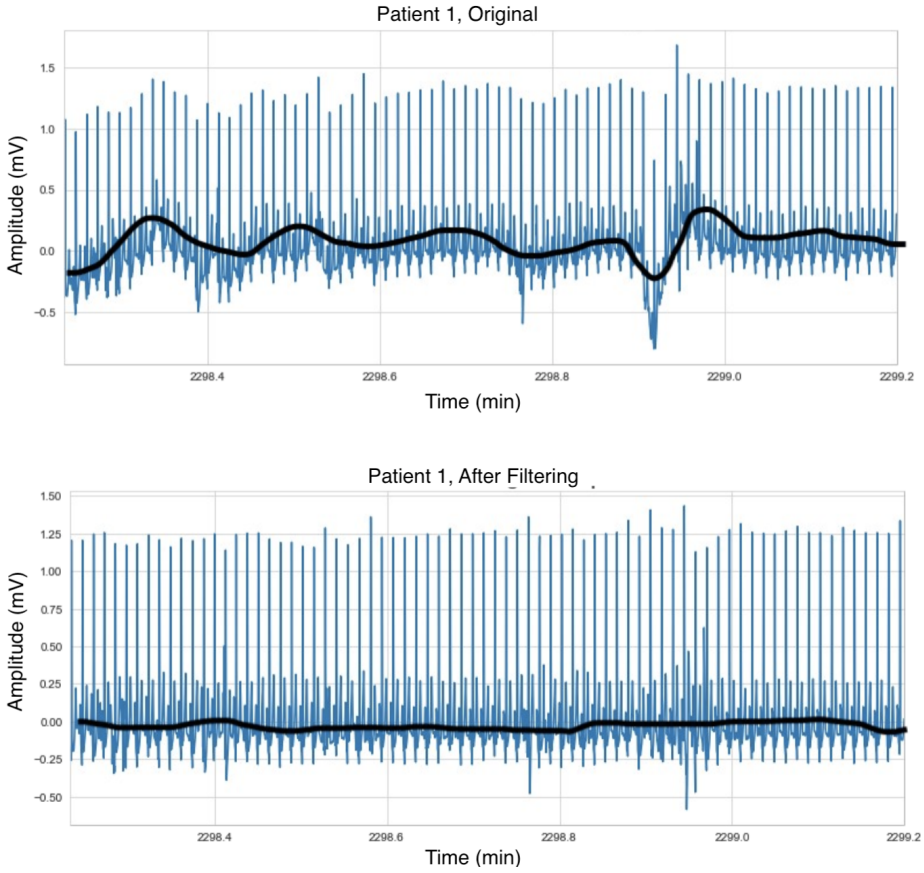

Figure 6: A subset of Patient 1's ECG 2 signal before and after the high and low pass filtering. Note that the filtering got rid of baseline shift. The black line is not part of the signal, but is included for illustrative purposes.

4.3 Abnormal Spike Removal

Part of the signals contains dramatic and noisy spikes, which can negatively impact our analyses. Therefore, we detected and removed noisy spikes in segmented 500 ms windows using a process introduced by [26]. This process iterates as follows:

1. Find the maximum absolute amplitude (MAA) in each window.
2. If at least one MAA is more than three times the median value of all the MAAs, continue. If not, the process is complete.
3. Set the values of the highest spike in the window containing the highest MAA to zero. The values of the spike are determined to be from the last zero-crossing point before the spike to the first zero-crossing point after the spike.
4. Iterate back to Step 2.

An example of the effect of spike removal is shown in Figure 7. In Section 5, we segment this signal into cardiac cycles (also known as R-R intervals) by detecting R peaks (R waves). Large spikes in the data will cause the R peak detection to fail. Although zeroing out these spikes still creates artifacts in the data, we can now perform the cardiac cycle segmentation. After we segment the data, we stack it and perform median filtering (each cardiac cycle is replaced by the median of its neighboring cycles, thereby smoothing the data) on the stack, thereby eliminating the artifacts.

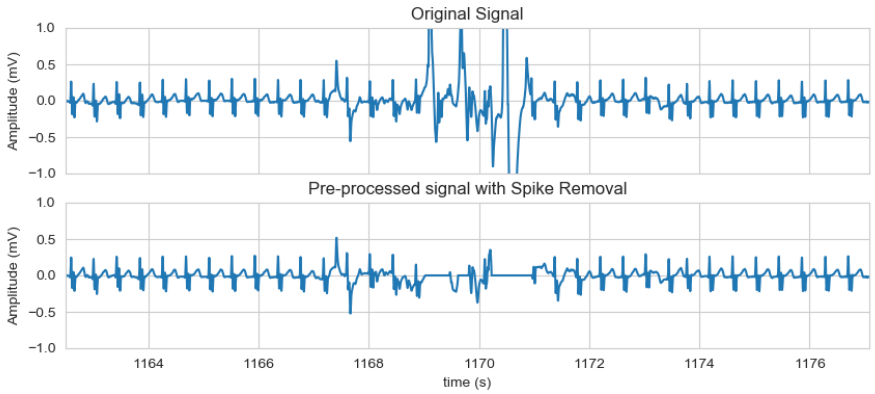

Figure 7: An illustration of noisy spikes that were removed.

5 Data Exploration

5.1 Cardiac Cycle Segmentation

As JET is canonically determined by the disappearance of the P wave in a cardiac cycle, we want to build our models on cardiac cycles. The most natural way to segment a ECG signal into cardiac cycles is to identify the R peaks, which are always present and clearly observable. We use the Python waveform-database (WFDB) package to detect the R peaks in the ECG waveform (Figure 8). We then apply a median filter and normalize each cardiac cycle to have the same length (100 data points). This results in cycles which are on the same time scale, allowing each sample within a cycle to represent roughly the same part of the wave. These 100 samples, or some combination of them, can then be used as features in a machine learning model.

We segment the preprocessed CVP signals based on the timestamps determined by the ECG R-R intervals. We then applied the same median filtering and normalization techniques used on the ECG cycles to the CVP cycles. Because the CVP is recorded at half the frequency of the ECG, each cardiac cycle was normalized to have 50 data points.

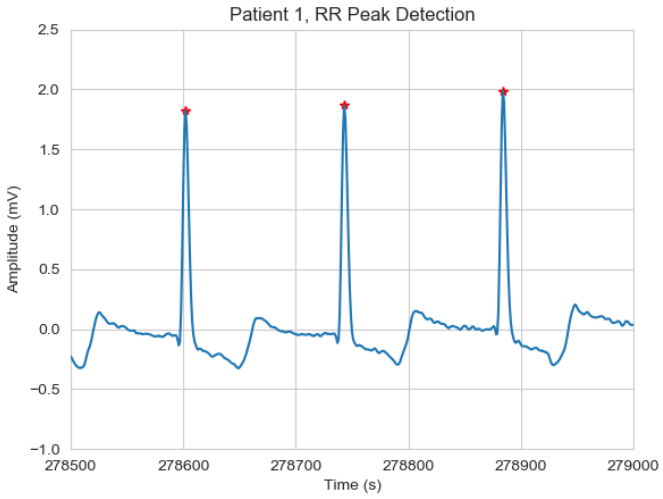

Figure 8: An illustration of R peak detection algorithm.

5.2 Cardiac Cycle Visualization

As each patient has a unique morphology, understanding how a patient's morphology has changed over time is essential to a clinician's ability to visually detect JET. The classical 2D visualization of a cardiac waveform can only display a few cardiac cycles at once; therefore, it is not a good representation of how the waveform changes over long periods of time. Consequently, we built a 3D visualization of the cardiac waveforms. Our visualization stacks many cardiac cycles together, showing hours of data at once, thereby allowing for easy detection of morphological changes (Figure 9).

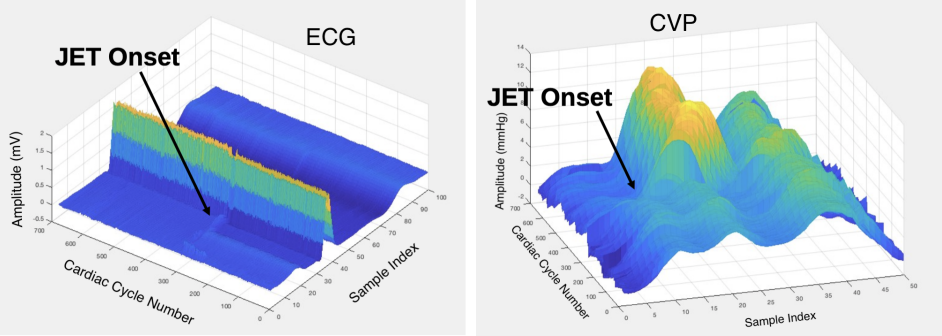

Figure 9: 3D Visualizations of JET onsets in a ECG waveform (left) and a CVP waveform (right). In the ECG waveform, one can easily see the P wave disappear and the height of the R peak increase in the ECG waveform. The P wave conjoining with the R wave is expected clinically for JET, as the electrical signal from the SA node is being masked by the signal from the AV node. The CVP waveform clearly shows the appearance of the cannon A wave in the CVP waveform following the onset of JET.

5.3 T-SNE Annotation Tool

As detailed in Section 3, it is impossible for a clinician to adequately label a large amount of cardiac waveform data. To help, we built a t-Distributed Stochastic Neighbor Embedding (t-SNE) annotation tool. T-SNE is a dimensionality reduction technique that works well for high dimensional visualization [27]. In brief, t-SNE creates a probability distribution using the Gaussian distribution to define relationships between points in a high-dimensional space. It then uses a Student t-distribution to recreate this relationships in a

low-dimensional space. It optimizes these relationship embeddings through gradient descent.

In Figure 10, we used t-SNE to cluster stacks of cardiac cycles. When one clicks on a data point in the cluster, it will bring up another window that one can use to label that cluster.

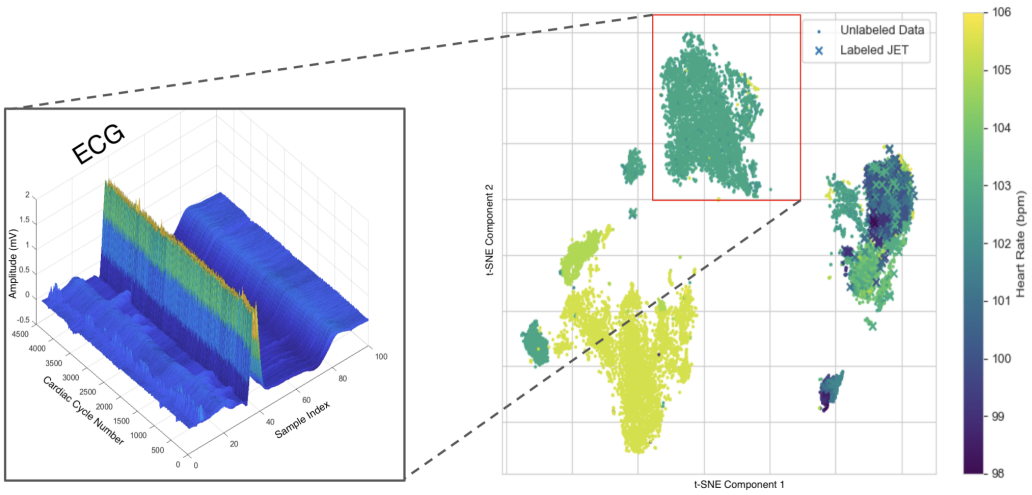

Figure 10: Example of a what would happen when one clicks on a point in a t-SNE clustering of ECG stacks of cardiac cycles. The coloring is based on heart rate, demonstrating that t-SNE seemed to cluster by heart rate in some instances. This is interesting as heart rate was not provided to the algorithm. Additionally, the cardiac cycles were normalized to the same length so the heart rate could not be computed by the algorithm.

6 Modeling

6.1 Supervised Classification Models

6.1.1 Convolutional Neural Network

As all current state-of-the-art arrhythmia detection models use CNNs on ECG signals (Section 2), we first experiment with a CNN on only ECG data. Following with the literature, we only train and test on the ECG 2 signal (the canonical signal used by doctors to detect arrhythmias as it shows the most

interpretable ECG morphology). We built a simple CNN with the following specifications:

1. The network has 4 convolutional layers followed by a fully-connected layer.

2. Each convolutional layer was followed by batch normalization and a ReLU activation, following best practices.

3. The fully-connected layer was followed by a softmax to prepare the outputs for our binary cross entropy loss (each output represents the probability that the dat point is either JET or non-JET).

4. We used a filter size of 2, a stride of 1 and no padding for simplicity.

5. The output dimensions of each respective layer were 128, 64, 32, 16, and 2. We chose layer sizes that were multiples of 2 for best practices and used a final layer size of 2 as we are doing binary classification.

6. We trained the network using a batch size of 128, the Adam Optimizer, and a learning rate of 0.001 for 50 epochs.

Although this is not a state-of-the-art CNN implementation, it is a simple one that follows best practices. It served as a proof of concept of whether or not CNNs could learn to predict the presence of JET. As will soon be discussed, the CNN was completely unable to generalize to new patients due to intrinsic limitations of CNNs. As such, we did not fine-tune our model to have the "best" implementation.

To fully utilize the convolutions in the CNN, we created a new dataset where each data point is a stack of 5 consecutive ECG 2 R-R intervals. We classified a stack as JET if at least one R-R interval in that stack was labelled JET. As the clinicians we were working with told us that JET only occurred once for each patient and as CNNs work best when given a large amount of data, we classified the 98.4% of our unlabelled data as non-JET. This created a highly unbalanced dataset. To help with this unbalanced nature of our dataset, we used a sliding window of 5 for each instance of JET and a non-sliding window of 5 for each instance of non-JET. Additionally, we built a data loader sampler so that each batch size was balanced by both class and patient, essentially randomly undersampling the majority class at each epoch.

Method	Sensitivity	Specificity	AUROC
CNN	96.37%	99.75%	98.06%

Table 2: The results of our CNN on a random 70/30 train/test split of all patients' ECG 2 data.

We tested the CNN on a random 70/30 train/test split of all of the entire dataset and got fairly good results (Table 2). But the CNN could not generalize to new patients. We trained the network on every combination of 8 patients and tested on the remaining patient and in each case, the network got 0% specificity. We believe this is due to the fact that morphologies vary greatly across patients (Figure 11), and we only trained on a few patient morphologies.

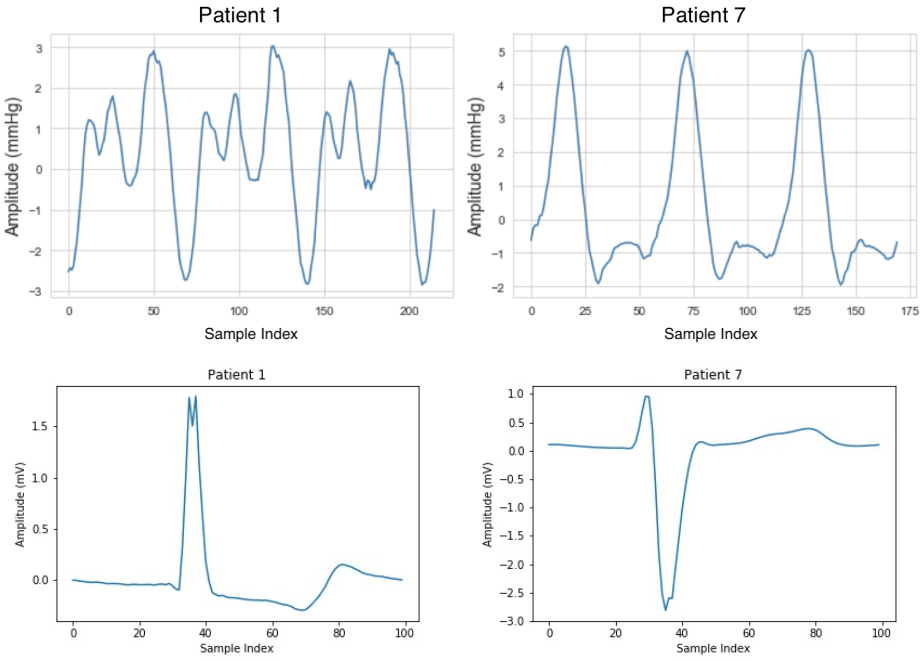

Figure 11: This figure shows the variations that can exist across patient morphologies. On top are the CVP morphologies for Patients 1 and 7. On bottom are the JET ECG morphologies for Patients 1 and 7.

6.1.2 Decision Trees

To assess the impact of CVP data on JET detection, two random forest models were trained: one that is trained on ECG data only and one that uses both ECG and CVP data. Random forest models are easier to interpret than CNNs, and they quantify the "importance" of each feature, making it easy to determine the effect of CVP. We balanced the classes by a method known as Synthetic Minority Over-Sampling Technique (SMOTE) that oversamples the minority class and undersamples the majority class [28]. We trained both models on patients 1, 7, and 9 (those for which the necessary CVP data was available) and used a 70/30 train/test split. We also trained two gradient boosting models using this same process. The results are displayed below (Table 3).

Method	Signals	Sensitivity	Specificity	AUROC
Random Forest	ECG	99.99%	99.92%	99.90%
Random Forest	ECG and CVP	99.99%	99.93%	99.91%
Gradient Boosting	ECG	99.19%	98.97%	98.10%
Gradient Boosting	ECG and CVP	99.50%	99.12%	98.31%

Table 3: The results of our supervised classification methods on a random 70/30 train/test split on patients with clear CVP signals. To make a fair comparison between ECG and ECG+CVP results, only Patients 1, 7, and 9 are used in this table.

The integration of CVP into the models helped improve their performance, but the model that uses ECG data performs quite well to begin with. However, same as the CNN, these models fail to generalize to different patients. When random forest and gradient boosting models are tested on patients that were not used in training, the models tend to classify all cycles into non-JET. In order to build a model which generalizes to new patients, more careful feature extraction is needed.

6.2 Semi-Supervised Classification Model

To better deal with variability among patient morphologies we would like to apply an unsupervised feature extraction method to the data before it is passed through a supervised algorithm. Using feature extraction to reduce dimensionality has been shown to be a successful approach in data science.

Feature extraction methods were applied to ECG signals to detect the position and/or the amplitude of the QRS complex, e.g. wavelet analysis and thresholding [29]. Such methods are usually deployed to parameterize ECG signals so that other algorithms can be used to process parameterized ECG instead of the signal itself. One of the problems with current parameterization methods is their interpretability. Interpretability is a crucial factor in medical applications as both doctors and patients want to know why the algorithm chose the answer that it chose. We propose Morphological Parameterization (MorPa), inspired by the way cardiologists interpret ECG signals, to address this issue.

6.2.1 MorPa Assumptions

In this method, we use the ECG 2 signal segmented into cardiac cycles of equal length. We use prior knowledge from cardiologists to model ECG signals:

$$x(t) = P(t) + R(t) + T(t) + b(t) + n(t)$$

where x is a cardiac cycle of ECG, $P(t)$ is the P waveform of ECG signal, $R(t)$ and $T(t)$, respectively, are the QRS complex and T waveforms in the cardiac cycle, $b(t)$ is a baseline that can be added to the ECG signal and $n(t)$ is the noise component. Figure 2 shows an example of this modeling.

6.2.2 Convex Baseline Removal

A good way to extract a low pass component of $x(t)$ which acts as the baseline, is solving this optimization problem:

$$y = \arg\min \ ||Dy||_2^2 + \lambda' ||x - y||_2^2$$

where D is the derivative matrix (i.e $Dx = x[n] - x[n-1]$). In order to solve this problem, we will write it as:

$$\min \quad (Dy)^T (Dy) + \lambda' (x - y)^T (x - y).$$

Setting the gradient with respect to y to zero, we can solve it and find the minimum without any iterative method.

$$(D^T D + \lambda' I) y = \lambda' x$$

$$\rightarrow y = \lambda'(D^T D + \lambda' I)^{-1} x := Hx$$

We must calculate H one time and then we can store it and use it for every cardiac cycle that we want to process. If we want to improve this baseline extraction algorithm, we can use Lasso for $y - x$ to impose sparsity on it (using the L1 norm for $y - x$ instead of the L2 norm), but doing so will make the algorithm slow and computationally expensive.

6.2.3 MorPa

MorPa aims to generate an ECG cardiac cycle that resembles a real cardiac cycle. Once we generate a signal similar to the given ECG signal, the parameters that led to generation of that signal are used as the extracted parameters of the given cardiac cycle.

We used the following model to generate cardiac cycles:

$$x(t) = C_P P^*(\frac{t - \theta_P}{\gamma_P}) + C_R R^*(\frac{t - \theta_R}{\gamma_R}) + C_T T^*(\frac{t - \theta_T}{\gamma_T})$$

where P^*, R^*, T^* are template waveforms for the P,Q, R, S and T waveforms, C is amplitude, θ is time positioning and γ is time scaling. We could also have a set of different templates for P, R, and T and could even modify these templates in order to have an adaptive algorithm.

We use singleMorPa to generate the most similar cardiac cycle to a given cardiac cycle.

6.2.4 SingleMorPa

Finding the Best Coefficients First we solve the problem:

$$[B, C] = \arg\min ||b + c\, M(t) - x(t)||_2^2$$

where $x(t)$ is a given signal with the length L and $M(t)$ is a template with the same length L.

We define the inner product of two signals as $< a(t), b(t) >= \sum_{t=0}^{L} a(t)b(t)$. We can then re-write the problem as:

$$\arg\min < b + c\, M(t) - x(t), b + c\, M(t) - x(t) >$$

if we set the gradient to zero, we will have:

$$\sum_{t=0}^{L} < b + c\,M(t) - x(t), 1 >= 0$$

$$\sum_{t=0}^{L} < b + c\,M(t) - x(t), M(t) >= 0.$$

Which is a system of linear equations:

$$\begin{pmatrix} < M, M > & < 1, M > \\ < M, 1 > & < 1, 1 > \end{pmatrix} \begin{pmatrix} c \\ b \end{pmatrix} = \begin{pmatrix} < x, M > \\ < x, 1 > \end{pmatrix}$$

If the determinant is non-zero, we can find b and c as:

$$\begin{pmatrix} c \\ b \end{pmatrix} = \begin{pmatrix} < M, M > & < 1, M > \\ < M, 1 > & < 1, 1 > \end{pmatrix}^{-1} \begin{pmatrix} < x, M > \\ < x, 1 > \end{pmatrix}$$

$$(c, b) \approx 1(x, M) \tag{1}$$

Finding the Best Positioning and Scaling The next step would be allowing the algorithm to shift the template $M(t)$ across time and find the optimal amount of time shifting:

$$(\theta, \gamma) = \arg\min ||b^*(\theta, \gamma) + c^*(\theta, \gamma)\, M(\frac{t-\theta}{\gamma}) - x(t)||_2^2$$

where $b^*(\theta, \gamma)$ and $c^*(\theta, \gamma)$ are answers to Equation 1 for the signal $x(t)$ and template $M(\frac{t-\theta}{\gamma})$. This means that for each amount of θ and γ we first marginally optimize the bias and the coefficient, and then we will find the best θ and γ by greedy search and use $b^*(\theta, \gamma)$ and $c^*(\theta, \gamma)$.

There are some cases, however, where there are unwanted waveforms added to the given signal (e.g. there is sometimes a specific unwanted signal which is the result of the electrical activity of a pacemaker.) Since the template is non-zero in a small portion of the length L (lets say l), the algorithm is likely to find a region in the given signal that has more energy. In contrast, we want the signal to have a similar morphology in that region. To avoid over-attraction to regions with more energy, we define another cost function below:

$$\frac{||b^*(\theta,\gamma) + c^*(\theta,\gamma)\, M(\frac{t-\theta}{\gamma}) - x(t)||_2^2}{||x(t)||_{local}^2}$$

where $||x(t)||_{local}^2$ is energy of the signal in the regions where $M(\frac{t-\theta}{\gamma})$ is non-zero

$$||x(t)||_{local}^2 = \sum_{t=\theta}^{\theta+\frac{l}{\gamma}} x^2(t).$$

This time, the algorithm is vulnerable to attraction to regions that have a small noise that has similar morphology to the template. We are interested in minimizing the mean squared error (MSE) and the above cost function at the same time. So the final cost function we want to minimize is:

$$(\theta^*, \gamma^*) = \arg\min\, (1 + \frac{\lambda}{||x(t)||_{local}^2})\, ||b^*(\theta,\gamma) + c^*(\theta,\gamma)\, M(\frac{t-\theta}{\gamma}) - x(t)||_2^2$$

where λ is a hyper-parameter. The method we used to solve this minimization problem is a greedy search, as the cost function is highly non-convex. It has many local minima and saddle points, and, hence, the gradient descent family of optimization algorithms will not work. Future work would be to find an appropriate optimization algorithm.

6.2.5 MorPa Outputs

We can now use SingleMorPa as a tool to parameterize the R, P and T waveforms of each ECG cardiac cycle into 3 features each (C-the amplitude,

θ-the time scaling, and λ-the time positioning), as shown in Algorithm 1.

Algorithm 1: MorPa

1 Use SingleMorPa to extract features for the R wave.
$(C_R, \theta_R, \gamma_R) = \text{SingleMorPa}(x, \text{R-template})$
2 Remove the R wave.
$x(\theta_R < t < \theta_R + \frac{l}{\gamma_R}) = x(\theta_R)$
$x = x - H * x$
3 Use SingleMorPa to extract features for the P wave.
$(C_P, \theta_P, \gamma_P) = \text{SingleMorPa}(x, \text{P-template})$
4 Remove the P wave.
$x(\theta_P < t < \theta_P + \frac{l}{\gamma_P}) = x(\theta_P)$
5 Use SingleMorPa to extract features from the T wave.
$(C_T, \theta_T, \gamma_T) = \text{SingleMorPa}(x, \text{T-template})$

6.2.6 MorPa Model Results

We trained a shallow CNN on the extracted features from MorPa (i.e. the amplitudes, time scaling and time positioning of the P, R, and T waveforms). The average results from a 4-fold cross-validation are shown in Table 4. From these results we see that with MorPa involved, a CNN can now generalize to new patients, even if it had only been trained on a few patients (in this case, it was trained on 3 patients).

CNN with MorPa Features	Signals	Sensitivity	Specificity	AUROC
MorPa	ECG	87.07%	87.12%	87.50%

Table 4: The average of a 4-fold cross-validation results for MorPa. Note that this table presents results for generalizing to new patients, whereas the other tables presented results for a random split on all patients.

7 Conclusion

Even though arrhythmia detection algorithms have been studied for years, JET, despite being the most common post-operative arrhythmia, has not been particularly targeted yet. The work we presented in this project represents a significant step in building an accurate JET detection model.

With the pre-processed sample matrix as inputs, we built classification models that are quite accurate within patients, obtaining AUROC of over 98%. We also integrated CVP signals, which is a novel practice in the realm of arrhythmia detection research. Integrating the CVP signals could improve model performance according to our preliminary tests. Though these models failed when testing on new patients, we developed MorPa to extract a better set of features from the ECG waveform. The MorPa features helped the model generalize across patients, and our preliminary tests exhibited great performance.

Besides the models we developed, we also introduced a novel visualization method for cardiac waveforms that takes advantage of their periodical nature. This method will enable doctors to easily access and visualize data for a long period of time.

Finally, we also introduced an interface that enables simple labeling of the data. With this interface, doctors can label hours of data with one click. This interface could be combined with our algorithm to label instances that our model is confused about, further empowering better results.

7.1 Future Plans

7.1.1 Blood Pressure Channels

The random forest and gradient boosting models showed that CVP data has the potential to improve the detection of JET. A logical next step, then, would be to perform feature extraction on the CVP signals and use them, along with the features extracted by MorPa, as input into a CNN. While MorPa itself is specific to ECG waveforms, a similar algorithm could be developed to extract CVP features. In addition to CVP, the hospital records arterial blood pressure (ABP). Changes in ABP can also be indicative of the presence of an arrhythmia, so further analysis and incorporation of the ABP data might further improve the model.

7.1.2 Expanding the Dataset

The small number of patients included in the dataset proved to be a substantial challenge in creating models that are generalizable to new patients. Morphologies vary widely from one patient to another, so models that are trained on only a few patients struggle to detect JET in new patients. While

the development of MorPa helped overcome this challenge, adding more patients to the dataset is likely to improve the performance of our models. We have requested more data from the hospital and are hoping to obtain it shortly.

7.1.3 Human-in-the-loop Model Developments

As mentioned earlier, only a small proportion of our dataset was labeled, but we have built a tool that enables efficient labeling of large amounts of data. We intend to use this tool to obtain more accurate labels and thereby increase the amount of data that the models are trained on. This results in a "human-in-the-loop" model, in which unsupervised machine learning helps a human label the data, which in turn helps train a supervised machine learning model.

7.1.4 Detecting Various Arrhythmias

While JET is the most common type of postoperative arrhythmia in children, a system that can detect multiple types of arrhythmias and distinguish among them would be even more useful to clinicians. A future project could incorporate data from various types of arrhythmias and aim to detect each one.

References

[1] Data and statistics on congenital heart defects — cdc, Nov 2018. 1

[2] Mohamed Salim Kabbani, Hayan Al Taweel, Nasib Kabbani, and Saleh Al Ghamdi. Critical arrythmia in postoperative cardiac children: Recognition and management. *Avicenna journal of medicine*, 7(3):88–95, Jul 2017. 1, 5

[3] https://www.jointcommission.org/assets/1/18/sea_50_alarms_4_5_13_final1.pdf, 04 2013. 1, 6

[4] Junctional tachycardia. 1, 5

[5] Arrhythmias in children. https://my.clevelandclinic.org/health/diseases/14788-arrhythmias-in-children. 2

[6] Alan H. Kadish, Alfred E. Buxton, Harold L Kennedy, Bradley P. Knight, Jay W. Mason, Claudio D. Schuger, Cynthia M. Tracy, Alan H. Boone, Michael Elnicki, John W. Hirshfeld, Beverly H. Lorell, George P. Rodgers, Cynthia M. Tracy, and Howard H. Weitz. Acc/aha clinical competence statement on electrocardiography and ambulatory electrocardiography - a report of the acc/aha/acp-asim task force on clinical competence. *Journal of the American College of Cardiology*, 38(7):2092, 2001. 3

[7] Ahmad Khalil Jaffery, Zainul and Afroz. Ecg baseline drift elimination using wavelet packets transform. *Australian Journal of Basic and Applied Sciences*, 8:459–466, 03 2014. 3

[8] Richard K Klabunde. Central venous pressure, Apr 2014. 3

[9] Central venous pressure - evaluation, interpretation, monitoring, clinical implications. *Bratisl Lek Listy*, 109(4):185–187. 3

[10] The University of Hong Kong. https://www.anaesthesia.hku.hk/LearNet/interpretation.htm. 4

[11] NA Haas, K Plumpton, R Justo, H Jalali, and P Pohlner. Postoperative junctional ectopic tachycardia (jet). *Zeitschrift für Kardiologie*, 93(5):371–380, 2004. 5

[12] Christian Nordqvist. Tachycardia: Causes, symptoms, and treatments. https://www.medicalnewstoday.com/articles/175241.php, Nov 2017. 5

[13] Eduardo José Da S. Luz, William Robson Schwartz, Guillermo Cámara-Chávez, and David Menotti. Ecg-based heartbeat classification for arrhythmia detection: A survey. *Computer Methods and Programs in Biomedicine*, 127:144–164, 2016. 7

[14] Pranav Rajpurkar, Awni Y. Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y. Ng. Cardiologist-level arrhythmia detection with convolutional neural networks. *CoRR*, abs/1707.01836, 2017. 7

[15] Ali Isin and Selen Ozdalili. Cardiac arrhythmia detection using deep learning. https://www.sciencedirect.com/science/article/pii/S187705091732450X, Dec 2017. 7

[16] K S Rajput, S Wibowo, C Hao, and M Majmudar. On arrhythmia detection by deep learning and multidimensional representation. *CoRR*, abs/1904.00138, Apr 2019. 7

[17] Philip Warrick and Masun Nabhan Homsi. Cardiac arrhythmia detection from ecg combining convolutional and long short-term memory networks. 09 2017. 7

[18] Saeed Saadatnejad, Mohammadhosein Oveisi, and Matin Hashemi. Lstm-based ecg classification for continuous monitoring on personal wearable devices. *IEEE Journal of Biomedical and Health Informatics*, 12 2018. 7

[19] M Rahimpour and B Mohammadzadeh Asl. P wave detection in ecg signals using an extended kalman filter: an evaluation in different arrhythmia contexts. *Physiological Measurement*, 37(7):1089–1104, 2016. 7

[20] Mohamed Elgendi, Marianna Meo, and Derek Abbott. A proof-of-concept study: Simple and effective detection of p and t waves in arrhythmic ecg signals. *Bioengineering*, 3(4):26, 2016. 7

[21] Gustavo Lenis, Nicolas Pilia, Tobias Oesterlein, Armin Luik, Claus Schmitt, and Olaf Dössel. P wave detection and delineation in the ecg based on the phase free stationary wavelet transform and using intracardiac atrial electrograms as reference. *Biomedical Engineering / Biomedizinische Technik*, 61(1):37–56, Jan 2016. 7

[22] H.k. Chatterjee, R. Gupta, and M. Mitra. Real time p and t wave detection from ecg using fpga. *Procedia Technology*, 4:840–844, 2012. 7

[23] Rene Gonzalez-Fernandez, Martha Rivero-Varona, and Gisela Montes de Oca-Colina. Detection of p wave in electrocardiogram. *Computing in Cardiology*, Sep 2013. 7

[24] Arnau-Vives M. Arnau-Vives A Buenda-Fuentes, F. High-bandpass filters in electrocardiography: Source of error in the interpretation of the st segment. *ISRN Cardiology*, 2012. 9

[25] Herbr-J. E. Herbrandson J. E. son Asirvatham S. J Venkatachalam, K. L. Signals and signal processing for the electrophysiologist: part i:

electrogram acquisition circulation. *Arrhythmia And Electrophysiology*, 2011. 9

[26] Segmentation of heart sound recordings by a duration-dependent hidden markov model. *Physiological Measurement*, 31(4):513–529, 2011. 11

[27] L. van der Matten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008. 13

[28] N. V. Chawla, Hall-L. O. Bowyer, K. W., and W. P. Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artifical Intelligence Research*, 16:321–357, 2002. 17

[29] Juie D. Peshave and Rajveer Shastri. Feature extraction of ecg signal. *Conference 2014 International Conference on Communication and Signal Processing*, 2014. 18