Machine Learning Project: Fetal ECG Extraction from Mixed Signal Input

Jeroen Muller (s2590182)
Wester Coenraads (s2938928)
Mario Cavero (s4113152)
Diarmuid Kelly (s4091221)

***Abstract*—In this paper, we present a system for isolating a child's heartbeat from an ECG signal that has the mother's heartbeat mixed in. Using several signal processing algorithms, the data is filtered and modified for the linear regression. A simple linear regression predicts the mother's heartbeat, allowing the system to predict the timing of the child's heartbeat and reconstruct the exact heartbeat.**

I. Introduction

In healthcare, monitoring of heart rate signals can provide insights to underlying conditions [6]. These signals are measured using Electrocardiogram Machines, or ECG's. Monitoring the highly sensitive intrapartum stage of human fetal development can indicate fatal fetal abnormalities and give physicians the potential to perform corrective procedures [3].

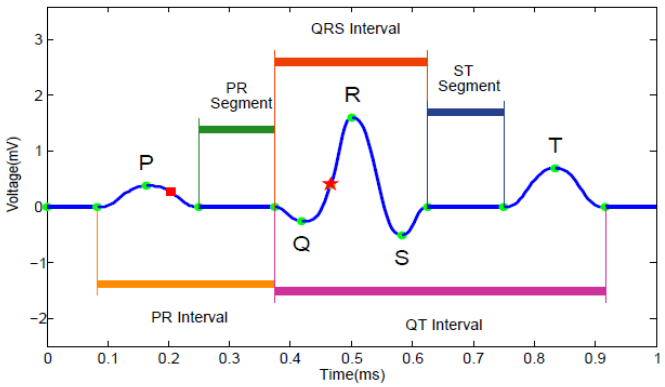

Fig. 1: Main Morphological Features in an ECG signal [8]

Non Invasive-fetal Electrocardiogram (NI-fECG) systems take measurements from the maternal abdomen [5]. These measurements are composed of mixtures of the fetal ECG (fECG) signal, maternal ECG (mECG) signal and noise (caused by e.g. uterine contractions). The mECG signal amplitude is predominantly stronger than the fetus' and both mECG and fECG share the same frequency and time domains [2]. Separating these signals is the focus of this paper.

Advances in diagnosing fetal abnormalities such as hypoxia are made possible by Morphological feature analysis, seen in Figure 1, specifically focusing on ST segment analysis. By measuring the ratio of T wave to QRS amplitude (T/QRS) it is possible to recognise the physiological response caused by hypoxia as an increase in the ST segment [1].

In order to differentiate between the fECG and mECG signals to perform analysis on fECG signals we investigate applying a linear regression to a dataset containing 5 measurements taken from a pregnant mother. The measurements are taken from the abdomen and thorax. The main goal here is to isolate the desired data – the child's heartbeat – out of the abdomen signal.

Our expected result is to get the average child heartbeat. As there are no quantitative success metrics to check whether the result is correct, we expect to see a distinguishable QRS structure for the fetus' heartbeat.

A. Problem

The problem addressed in this paper is the weakness of the fetal heartrate (fHR) with respect to the maternal heartrate (mHR) in utero ECG measurement. In order to perform Morphological Feature analysis on the continuous fetal heart rate, we need to distinguish the fHR from the mixed abdominal signal. In this case we treat the mHR from the thorax reading as noise when evaluating the abdominal readings.

II. Methods

A. Data Description

The data provided for this assignment was used. It contains five signals, each with 20,000 samples: three thorax signals and two abdomen signals (see Figure 2). As is clearly visible, the abdomen signals are much weaker than the thorax signals. The signals have a sampling frequency of 1000Hz, meaning each signal spans 20s.

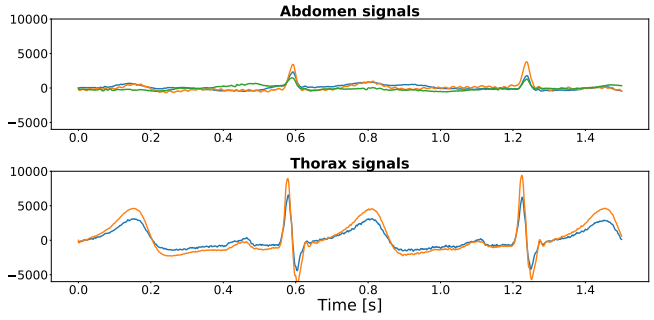

Fig. 2: The raw signals in the dataset.

B. Preprocessing

Before applying machine learning algorithms, the data needs to be preprocessed. The steps outlined in algorithm 1 are used to preprocess the data. Each step is further detailed in the following subsections.

Algorithm 1: Preprocessing steps applied to each signal.

```
input : An ECG signal of length 20000
output: A processed signal of length 20000
signal -= mean(signal);
signal /= std(signal);
filter(signal);
signal = analytic_signal(data);
```

1) Signal Normalisation: The first preprocessing step is to normalize the data. When measuring ECG signals, it is possible that these are in a different amplitude range depending on several factors. For example, the thorax signals have a much higher amplitude because those sensors are located much closer to the mother's heart. It's also possible that some signals have a DC component. For this reason, we subtract the mean from each signal, then divide them by their standard deviation. This ensures each channel is centered and has a similar range of values (see Figure 3).

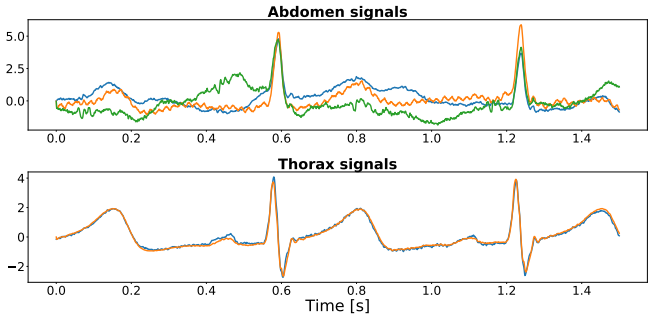

Fig. 3: The five signals after normalization.

2) Butterworth bandpass filter: The data is filtered using a digital Butterworth bandpass filter. The abdomen signals (and to a lesser extent the thorax signals) have a moving baseline, which throws off many machine learning algorithms. In addition, all the signals have high-frequency components that don't contribute meaningfully to the overall shape of the ECG. The bandpass filter reduces these unwanted high and low-frequency components.

The cutoff frequencies were experimentally determined at $20\ \mathrm{rad/s}$ for the low cutoff and $80\ \mathrm{rad/s}$ for the high cutoff. These frequency ranges were selected based on findings in [4], where the optimal ranges were determined between $\approx 3\mathrm{Hz} - 14\mathrm{Hz}\ (20 - 90\mathrm{rad/s})$. We then fine tuned our filter parameters to achieve the desired results. A 6th-order filter was used. Applying the filter to the inputs yields smoother signals without a moving baseline (see Figure 4).

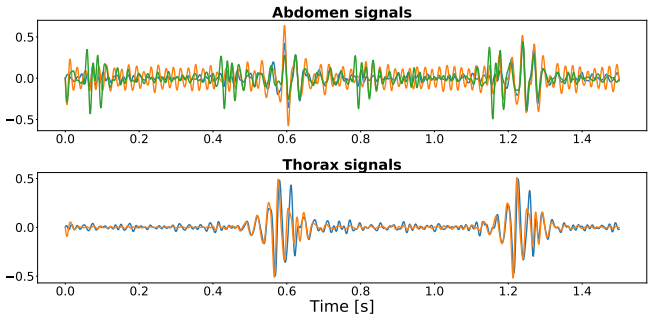

Fig. 4: The signals after applying the bandpass filter specified in Section II-B2.

3) Signal envelope: Next, the analytic signal is computed for each signal using the Hilbert transform (see Figure 5). For the case of periodic heartbeats, this essentially computes the envelope around each signal. This aids the peak detection in a later stage.

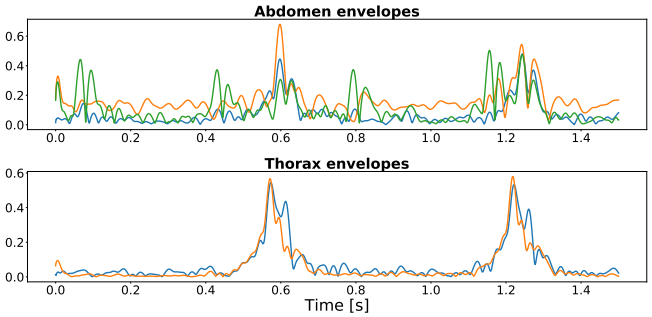

Fig. 5: The analytic signal for each signal.

C. Linear Regression

In the envelope of the filtered abdomen signal the child heartbeat is visible, which can be recognized as the double green peak at about 0.1s, 0.4s and 0.8s in Figure 5. However the amplitude is still smaller than that of the mother heartbeat. In order to isolate the child heartbeats we first need to remove the mother heartbeat from the abdomen signal as well as possible.

We know that the mother heartbeat originates in the heart and is transmitted to both the thorax electrode and the abdomen electrode. Both these transmissions are not instantaneous and may occur through multiple paths, so any impulse at the signal source becomes spread out in time in the electrode signals. This means that the information about the mother heartbeat contribution in the abdomen signal at time t is contained in some window about t in the thorax signal.

A comparison of these signals (see figure 5) shows that the mother heartbeat signal is very similar in both. Therefore it is reasonable to attempt to use a windowed linear regression to predict the mother heartbeat in the abdomen signal from the thorax signal. We found that this approach worked well enough for our purposes. The method is outlined below.

We wish to construct a signal $a'(t)$ as similar as possible to the abdomen signal $a(t)$ from the thorax signal $\theta(t)$. If the windows size is $2N+1$, let $\boldsymbol{\theta}_i \in \mathbb{R}^{2N+1}$ be a part of the thorax signal centered on sample i, such that.

$$\boldsymbol{\theta}_i = [\theta(i-N)\dots\theta(i)\dots\theta(N)]\prime \tag{1}$$

Then our estimation of the abdomen signal is determined by a weight vector $\boldsymbol{w} \in \mathbb{R}^{2N+1}$ according to $a'_i = \boldsymbol{w}\cdot\boldsymbol{\theta}_i$ [7]. We want to choose the weight vector that minimizes the squared error between the reconstructed abdomen signal and the real abdomen signal.

$$\boldsymbol{w} = \text{argmin}_{w^*} \sum_{i=1}^{t_{max}} \|a'_i - a_i\|^2 \tag{2}$$

This is a regression problem that we can solve with the least squares approach. If we let $X = [\boldsymbol{\theta}_1 \dots \boldsymbol{\theta}_{t_{max}}]'$ and $\boldsymbol{Y} = [a(1)\dots a(t_{max})]'$ the solution to this optimization problem is given by

$$\boldsymbol{w} = (X'X + \alpha^2 I^{(2N+1)\times(2N+1)})^{-1}X'\boldsymbol{y} \tag{3}$$

where the parameter α is used to regularize the solution by penalizing large values for the components of $\boldsymbol{w}$ [7]. One way to see why such regularization might improve the result is that it reduces the sensitivity of a' to noise in θ. For given i each component of $\boldsymbol{w}$ picks out a single sample of θ. This is problematic if there is noise in the thorax signal in addition to the heartbeat signal (which is true in reality). In that case large components in $\boldsymbol{w}$ tend to transmit the noise from the thorax to a', whereas more evenly distributed components tend to smooth the noise over a section over the window.

We used such a regression with $N = 250$ on the Hilbert envelopes of one of the thorax signals and one of the abdomen signal. The result shown in figure 6. The reconstructed signal $a'(t)$ closely follows the contribution of the mother heartbeats to $a(t)$ but not that of the child heartbeats. Therefore the difference $a(t) - a'(t)$ can be used to detect child heatbeats.

We also attempted to apply the linear regression without first calculating the signal envelope, but this worked much less well.

D. Heartbeat detection

The reconstructed abdomen envelope for the child heartbeat $a(t) - a'(t)$ can be used to determine at which timesteps the child heart beat is present. This follows the steps outlined in algorithm 2. First a windowed max-filter is applied to the envelope, then a windowed min-filter is applied to the result. This yields a signal with clear peaks at the child heartbeat locations. Each of these peaks is interpreted as a single heartbeat (see Figure 7). This process detects child heartbeats at roughly twice the frequency of the mother's heartbeat, suggesting that it detects the heartbeats correctly.

This process is used because the reconstructed envelope has several possible peaks for each heartbeat (see Figure 6). By first filtering the signal, only one peak remains for each heartbeat. The window size is an important parameter here: if it is too large, the filter will combine several heartbeats into one. Conversely, it can yield too many peaks when the window size is too small. The window size was experimentally determined at 0.2s (or 200 samples).

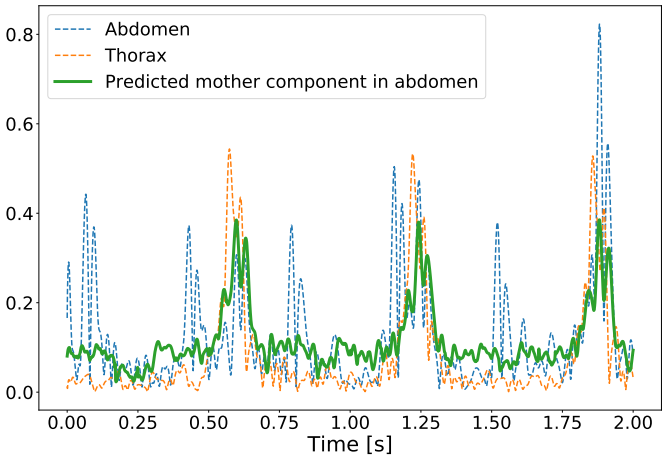

Fig. 6: The original abdomen and thorax signals and the predicted mother component in the abdomen signal.

Algorithm 2: The heartbeat detection algorithm.

input : The predicted child heartbeat envelope
output: The locations at which beats are detected
filtered = max_filter(envelope);
filtered = min_filter(filtered);
beat_locations = peak_indices(filtered);

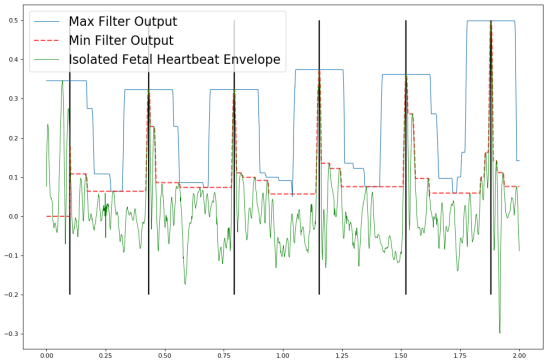

Fig. 7: The child heartbeat envelope with the output of the min- and max-filters. The vertical lines indicate detected beats.

E. Extracting the child ECG signal

Now that the locations of of the child heartbeats are known, the child ECG signal can be extracted using the algorithm described in Algorithm 3. The procedure outlined in Sections II-B and II-C is repeated, with exception of the step where the

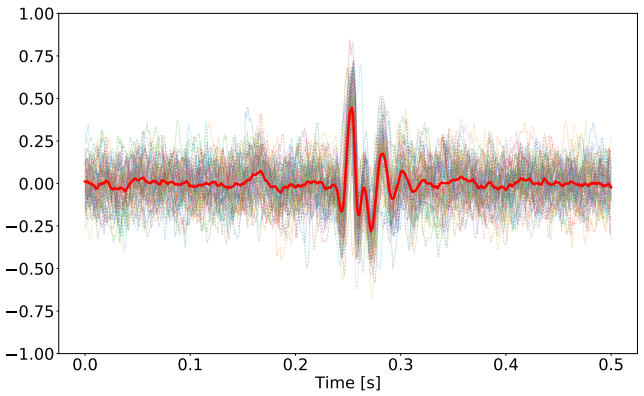

Fig. 8: The mean child heartbeat in red, with all predicted child heartbeats (thin dashed lines).

signals are changed into their envelope. This yields the predicted mother component of the abdomen ECG. We subtract this mother component from the abdomen ECG to obtain the child ECG.

Using the heartbeat locations detected, we can slice up the child ECG. For each heartbeat, we take a window of 500 samples (or 0.5s) around it. Each of these windows constitutes a single heartbeat. When plotted together, a clear pattern forms (see Figure 8). Taking the mean of each window yields a recognizable hearbeat shape (see Figure 8).

Algorithm 3: The heartbeat extraction algorithm.

```
input : The original data, the locations at which beats
        are detected
output: The child's ECG signal, sliced up per
        heartbeat
data = preprocess(data);
mother_heartbeat = linear_regression(data);
child_heartbeat = data - mother_heartbeat;
for each detected location do
    window = window of size 500 around location;
    slices.add(child_heartbeat[window]);
end
```

III. Discussion

As hypothesized in the introduction, we managed to get the predicted fetus' heartbeats and the average of it. As mentioned before, there is no way to check whether the predicted child heartbeat is exactly the same as the supposedly real heartbeat but the structure of itself looks decent.

In this project we not only have learnt how to do a ML project but also learnt about the healthcare sector. It has shown us that ML can be applied in so many fields, not only helping human tasks but also improving or speeding up processes. This may mean a faster detection of anomalies, or a better detection, or both.

One approach we tried early on was Independent Component Analysis. This method attempts to separate several sources from a mixed signal. After attempting to tune this method so it fits our problem, we discarded it in favor of the more straightforward linear regression.

Preprocessing the data in a proper way is key in order to have a good performance in any ML algorithm. We found that our results were very dependent on the preprocessing steps used, and not just the machine learning algorithm. In the end, the machine learning algorithm used was a simple ridge regression – with the pre- and postprocessing steps doing much of the work.

A. Improving the results

A desired approach for improving or validating our results would be from obtaining a fECG reading from one of the invasive fetal scalp electrodes. This could be used to verify our findings of what we believe to be the QRS complex in Figure 8. However, given the lack of detail and ground truth we cannot verify the presence of this. Another issue with our method is the inability to clearly determine the existence of the P or T peaks. An FSE measurement would provide our method with a quantitative measure for comparison and enable us to properly evaluate success.

References

[1] Isis Amer-Wåhlin et al. "Cardiotocography only versus cardiotocography plus ST analysis of fetal electrocardiogram for intrapartum fetal monitoring: a Swedish randomised controlled trial". In: *The Lancet* 358.9281 (2001), pp. 534–538.

[2] Gari D Clifford et al. "Non-invasive fetal ECG analysis". In: *Physiological measurement* 35.8 (2014), p. 1521.

[3] Matthias Egger, George Davey-Smith, and Douglas Altman. *Systematic reviews in health care: meta-analysis in context*. John Wiley & Sons, 2008.

[4] AA Fedotov. "Selection of parameters of bandpass filtering of the ECG signal for heart rhythm monitoring systems". In: *Biomedical Engineering* 50.2 (2016), pp. 114–118.

[5] EH Hon and O_W Hess. "The clinical value of fetal electrocardiography". In: *American Journal of Obstetrics & Gynecology* 79.5 (1960), pp. 1012–1023.

[6] EH Hon and EJ Quilligan. "The classification of fetal heart rate. II. A revised working classification." In: *Connecticut medicine* 31.11 (1967), pp. 779–784.

[7] Herbert Jaeger. "Machine Learning Lecture Notes". Nov. 2019.

[8] Rami Oweis and Basim Al-Tabbaa. "QRS Detection and Heart Rate Variability Analysis: A Survey". In: *Biomedical Science and Engineering*, 2 (Jan. 2014), pp. 13–34. DOI: 10.12691/bse-2-1-3.