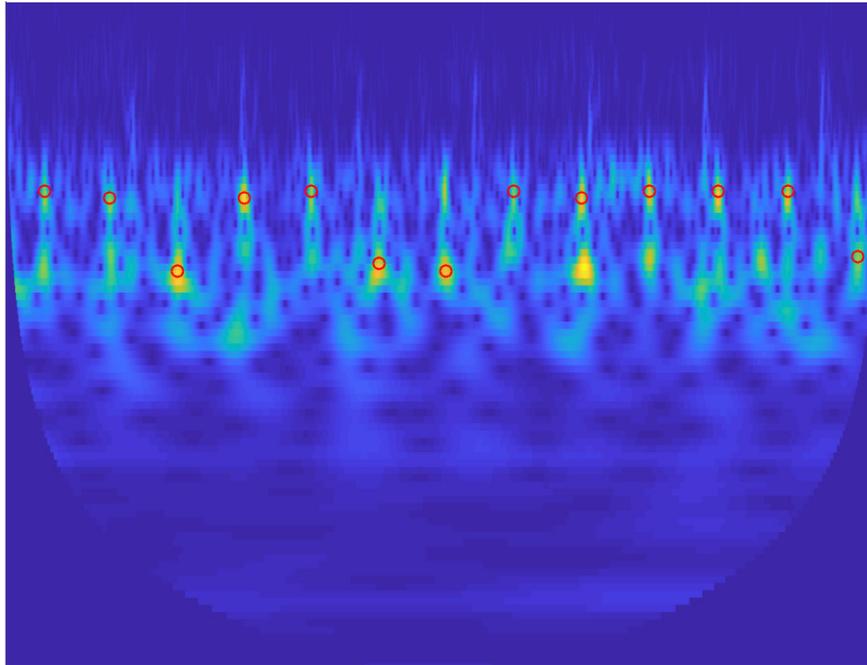




CHALMERS
UNIVERSITY OF TECHNOLOGY



Non-Invasive Fetal Monitoring using Non-Contact Electrodes and Recurrent Neural Networks

ALBIN ANNÉR
DAVID KASTÖ

MASTER'S THESIS 2019

**Non-Invasive Fetal Monitoring using
Non-Contact Electrodes and Recurrent
Neural Networks**

ALBIN ANNÉR
DAVID KASTÖ



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2019

Non-Invasive Fetal Monitoring using
Non-Contact Electrodes and Recurrent
Neural Networks
ALBIN ANNÉR
DAVID KASTÖ

© ALBIN ANNÉR, 2019. © DAVID KASTÖ, 2019.

Supervisor: Lars Hellberg, Department of Physics
Examiner: Peter Apell, Department of Physics

Master's Thesis 2019
Department of Physics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Continuous wavelet transform of an abdominal ECG of a pregnant woman, where the maternal ECG component has been removed. The circles denote detected fetal heart beats.

Typeset in L^AT_EX
Gothenburg, Sweden 2019

Abstract

The fetal well-being is routinely monitored using a cardiotocograph (CTG), a combination of a Doppler sensor used to measure the fetal heart beats and a pressure sensor to measure uterine muscle contractions. However, the equipment is expensive and there is a recurrent issue where the CTG confuses the maternal heart rate for the fetal heart rate, leading to ambiguities that have resulted in adverse fetal outcomes on multiple occasions. Non-invasive recordings of the fetal heart activity on the maternal abdomen could constitute a viable alternative to Doppler ultrasound recording. However, the potential sensed by an abdominal electrode is the combination of many sources, making it an onerous task to extract the components stemming from the fetal heart beats. The most difficult subsignal to circumvent is the contribution from the maternal heart.

This thesis explores the viability of using non-contact electrodes to reliably measure the fetal electrocardiogram, fECG, and an electrode is successfully designed for this purpose. This electrode could be implemented in an array embedded in a piece of clothing for a practical implementation of many uncorrelated electrodes, which normally facilitates the signal separation process and improves its accuracy.

Separately, using the Non-Invasive Fetal Electrocardiogram (NI-fECG) database, source separation and fetal heart rate extraction algorithms are evaluated, developed, and improved upon. It is shown that using said algorithms in combination with a novel fetal heart rate detector, fetal heart beats can be detected in a reliable manner on the evaluated data.

Acknowledgements

We would like to express our gratitude to our supervisor Lars Hellberg, Senior lecturer and Research engineer at the Department of Physics at Chalmers University of Technology, for always providing us with any equipment we might request in a swift and helpful manner. His constant availability, day and night, has many a time been crucial to the fast and smooth progression of the thesis work. We would also like to acknowledge our examiner Peter Apell, Professor of Living State Physics and Head of Department of Applied Physics at Chalmers University of Technology, who has been of great help with his unwavering support, tremendous experience, and vast network.

In addition, we would like to express our appreciation to Hans Odelius, Technician at the Department of Physics at Chalmers University of Technology, who helped us out on multiple occasions when we wanted to make in-house printings of our first couple of circuit generations.

Last but not least, we are highly indebted to the people at ReVibe Energy, who throughout the whole thesis period lent us their lab and various equipment of paramount importance to the project. Their hospitality and friendship not only played an important role in the positive outcome of the thesis work, but also made the endless hours we have spent on this project so much more pleasant.

Albin Annér and David Kastö
Gothenburg, June 2019

Contents

List of Figures	v
List of Tables	xii
I Context and problems to be solved	1
1 Introduction	2
1.1 Motivation and historical context	2
1.2 Outline of thesis	3
1.2.1 Part I - Context and problems to be solved	3
1.2.2 Part II - Hardware development and testing	3
1.2.3 Part III - Software development and testing	3
1.2.4 Part IV - Conclusion and future work	4
2 Electrophysiology of the heart	5
2.1 Cardiac muscle contraction	5
2.2 The cardiac cycle	6
2.3 Dipole moment of layers and closed surfaces	7
2.4 Electrocardiogram as a net dipole map	10
3 The fetal ECG	13
3.1 Similarities and differences to adult ECG signals	14
3.2 Prenatal development landmarks relevant to fECG extraction	15
3.3 Significant features of the fECG	15
3.4 Extraction of uterine contractions and fetal respiratory movements from the abdominal ECG	16
4 Issues with current technology and problems to be solved	18
4.1 Current technology and its drawbacks	18
4.2 The other subsignals and how to circumvent them or reduce their prominence	20
II Hardware development and testing	22
1 Basic theory/Fundamental principles	23

1.1	Capacitance	23
1.1.1	Displacement current and capacitive sensing	24
1.2	Faraday cage	25
1.3	Differential measurements and CMRR	26
2	Electrode design	27
2.1	Contact or non-contact electrode?	27
2.2	Desired properties for non-contact electrodes	28
2.3	Circuit analysis	29
2.3.1	Derivation of transfer function, maximal gain and cutoff frequency	29
2.3.2	Derivation of output noise	31
2.3.2.1	The effect of finite source resistance	32
2.3.2.2	Signal to noise ratio	33
2.3.3	Variations in body electrode coupling	35
2.4	Bootstrapping	36
2.5	Input offset voltage and discharge time	37
2.6	Driven guard and stray capacitance	38
2.7	Electrostatic shielding	39
2.8	Neutralization of input capacitance	40
2.9	Design procedure	40
3	Evaluation of electrode	45
3.1	ECG measurement setup	45
3.2	Transfer function measurement setup	46
3.3	Noise measurements setup	46
4	Results	48
4.1	Transfer function - gain and cutoff frequencies	48
4.2	Noise measurements	50
4.3	ECG measurements	52
5	Discussion	55
6	Conclusion	57
III	Software development and testing	58
1	Basic theory	59
1.1	Adaptive filtering	59
1.2	Blind source separation	60
1.3	Artificial neural networks	61
1.4	Recurrent neural networks	63
1.5	Genetic algorithms	63
1.6	Continuous wavelet transform	66
1.6.1	Scaling and shifting	67

1.6.2	CWT as a windowed transform	67
2	Database and figures of merit	69
2.1	Non-Invasive Fetal Electrocardiogram Database	69
2.2	Figures of merit	69
3	Review of non-invasive fECG extraction methods	71
3.1	Spatial methods	71
3.1.1	Principal component analysis	72
3.1.2	Independent component analysis	72
3.2	Temporal methods	72
3.2.1	Template subtraction	72
3.2.2	Adaptive filtering	73
3.2.2.1	Least mean square adaptive filter	74
3.2.2.2	Recursive least square adaptive filter	75
3.2.3	Kalman filtering	75
3.3	Performance	76
4	Experimentally evaluated methods - FECG extraction	79
4.1	Preprocessing	79
4.2	No maternal chest reference	80
4.2.1	Variants of ICA	80
4.3	Maternal chest reference	82
4.3.1	The long short-term memory network	82
4.3.2	Application and optimization of LSTM network	84
4.3.3	Echo state neural networks	85
4.3.4	Optimization of ESN performance using a genetic algorithm	88
5	Experimentally evaluated methods - fQRS extraction	90
5.1	Pan-Tompkins	90
5.2	QRS detection from a scalogram of a continuous wavelet transform	93
6	Results	95
6.1	Removal of the mECG from the AECG mixture	95
6.2	LSTM and ESN in combination with Pan-Tompkins	96
6.2.1	LSTM	96
6.2.2	ESN	97
6.3	LSTM and ESN in combination with a CWT based fQRS detection	98
6.3.1	LSTM	98
6.3.2	ESN	99
6.4	Performance with ± 5 BPM condition	100
7	Discussion	101
8	Conclusion	104

IV Conclusion and future work	105
Bibliography	108

List of Figures

I	Context and problems to be solved	
2.1	Figure showing the periodic action potentials generated in a pacemaker cell in the SA node. The orange line shows the slow depolarization due to sodium influx. At the threshold the cell rapidly depolarizes (red line) and then repolarizes (green line). In this manner the pacemaker cells initiate the propagating action potential, and therefore sets the rhythm of the heart. Reproduced from OpenStax College.	5
2.2	Anatomical illustration of a frontal plane cross section of the human heart. The illustration shows the propagation path of the action potential. The action potential originates from the SA node and propagates through both atria and also to the AV node. At the AV node, which is the only electrical connection between the atria and ventricles, the action potential is delayed before it propagates through specialized conductive fibers, Purkinje fibers, to both ventricles. Reproduced from OpenStax College.	6
2.3	Figure showing an electrostatic model of a partially depolarized heart. a) shows the cell as the action potential has depolarized part of the cell. The right side of the cell is yet to be depolarized by the propagating action potential. An equivalent model of a) is the superposition of b) and c) . Since there is no net electric field from the closed surface b) , the only contribution comes from c) . The model for the heart as a whole becomes equivalent as the action potential propagates from cell to cell. This means that the depolarization of the heart can be seen as a propagating charged surface that intersects the current position of the action potential wave.	8
2.4	Measurement made from two abdominal electrodes. The large loop is the contribution of the maternal QRS wave and the smaller loop is the contribution of the fetal QRS wave. Looking at the figure it is clear that the maternal and fetal net dipole moments traverse different paths in space. The data is taken from the Abdominal and Direct Fetal Electrocardiogram Database [1].	10

2.5	Illustration of an ECG waveform recorded on a lead on the skin surface. The P wave arises from atrial contraction, the PR segment is from the delay in the SA node and the QRS complex from ventricular depolarization. The ST is the isoelectric interval in between ventricular depolarization and repolarization. The T wave comes from ventricular repolarization. The U wave is not always visible and comes from repolarization of so-called purkinje fibers.	11
3.1	The most common fetal presentations and their concomitant prevalences in utero at the time of labor. The statistics are fetched from Symonds et al.[2].	13
3.2	a) Time-frequency spectrum of an SECG waveform. The yellow nodes are fetal heart beats. For each heart beat, most of the spectra is contained within the bandwidth 15-70 Hz. b) The time frequency spectrum of an abdominal electrode. Note that the fetal heart rate is higher than the maternal, the maternal heart beats have a higher amplitude than the fetal equivalents, and that the fetal wave packet contains higher frequencies than the maternal ditto. The data is taken from the Abdominal and Direct Fetal Electrocardiogram Database[1].	14
3.3	Timeline of prenatal development, including some key entries for fetal monitoring. Vernix caseosa is a thin, highly non-conductive layer normally surrounding the fetus between around the 28th-32nd week and the 37th-38th week that has a significant adverse affect on NI-fECG extraction. Based upon statistics given in Behar[3].	15
3.4	A uterine contraction, recorded by an intrauterine pressure transducer, and its effect on the signal quality of an ECG measured on the maternal abdomen. After 70 seconds there is a deterioration in the ECG quality that correlates with the onset of the contraction. The figure has been reproduced by permission of IOP Publishing.	16

II Hardware development and testing

1.1	Schematic of the charge separation in a parallel-plate capacitor and its rendered internal electric field. The dielectric material between the plates (orange ellipses) becomes polarized due to the charge displacement, which reduces the total internal field and increases the capacitance. Reproduced from November[4].	24
1.2	Illustration of the working principle of a Faraday cage. In a) , an external field has just been applied and thus the electrons in the conductive cage are starting to move towards the positive pole of the field. This redistribution of charges creates a polarity between the left and right edge of the cage, as seen in b) . This polarity, in turn, renders its own electric field in the opposite direction of the external field as displayed in c) . These two opposing fields cancel out each other to a great extent, in theory (for a perfect conductor) completely. This is illustrated in d) . Adapted from Skowron[5]. . . .	25

2.1	Circuit model of the proposed circuit, shown without external feedback circuitry such as bootstrapping and input neutralization. The amplifier is an instrumentation amplifier and hence has no gain feedback. e_{R_B} is the Johnson-Nyquist noise arising from the bias resistor to ground, i_n is the input referred current noise and e_n the input referred current noise of the amplifier. C_{body} is the capacitive coupling between body and electrode and C_{lim} is a capacitor that limits variation in source capacitance. C_B is the bias capacitance, which is the contribution between parasitic coupling capacitances between other parts of the circuit and the input capacitance of the amplifier.	29
2.2	Noise sources in the proposed circuit, shown without external feedback circuits such as bootstrapping and input neutralization. The amplifier is an instrumentation amplifier and hence has no gain feedback. e_{R_B} is the Johnson-Nyquist noise arising from the bias resistor to ground, i_n is the input referred current noise and e_n the input referred current noise of the amplifier. C_{body} is the capacitive coupling between body and electrode and C_{lim} is a capacitor that limits variation in source capacitance. C_B is the bias capacitance, which is the contribution between parasitic coupling capacitances between other parts of the circuit and the input capacitance of the amplifier.	31
2.3	Amplification vs body-electrode coupling capacitance for different limiting capacitors C_{lim}	34
2.4	Circuit model showing input resistance bootstrapping network. The equivalent input impedance is $Z_{in} = R_{in} = R_1 + R_2 + \frac{R_1 R_2}{R_3}$	36
2.5	The top figure shows the decay time constant for a step response on the input of the amplifier as a function of resistance, see equation (2.5). The bias current $i_n = 100$ fA, the bias capacitance $C_B = 7$ pF and $C_S = 10$ pF. The time constant is directly proportional to $C_S + C_B$ and R_B . The bottom graph shows the equilibrium offset voltage on the input due to the bias current. The input offset is linear with respect to bias current and bias resistance. When the bias resistance is in the range of multiple T Ω , the response time becomes large and so does the input offset voltage.	37
2.6	Figure showing the principle of driven guard. b) shows the sensitive trace, 3, and some other trace, 1, on the circuit board, e.g. ground. The sensitive trace couples with 1 through C_p and R_B . a) shows the driven guard technique. The sensitive input, 3, couples with the guard, 2, through C_d and R_d . The guard couples with some other trace, 1, through C_p and R_p . As the guard is actively driven to the same potential as the input, there is no leakage current between these two traces. Since the guard is not sensitive, it is unaffected by the coupling with trace 1.	38

2.7	Figure showing the neutralization technique of input capacitance. The output is amplified by $1 + \delta$ fed back to the input through the coupling capacitance C_n . With the correct choice of C_n and δ ($C_n = C_B/\delta$) the current to the parasitic capacitance is drawn from C_n , hence canceling the effect of C_B . In this way the capacitive part of the input impedance Z_{in} is removed, or in practise, reduced.	39
2.8	Noise power spectral density for the desired configuration of the electrode. As can be seen the thermal noise and current noise is lower than the input referred noise on the interval 15-75 Hz. On this bandwidth the total noise is $0.9242 \mu V_{rms}$	43
2.9	Design of the electrodes in EAGLE PCB. a) shows the top copper layer of the electrode, b) shows the inner copper guard layer and c) shows the electrode surface with a surrounding guard ring.	44
3.1	The ECG measurement circuit. On the left side are circuit diagrams for the electrodes described in section 2.9. The right side shows the amplifier circuit. Both signals were passed through a high pass filter and then to INA128, a high CMMR instrumentation amplifier, before a differential signal was logged.	45
4.1	Figure showing the measured transfer function of the non-contact electrode at $C_S = 10$ pF. The upper figure shows the amplitude modulation and the lower figure shows the phase attenuation. The high pass cutoff frequency was measured to be 0.0268 Hz which is well below 0.05, the standard for ECG measurements [6]. From the analytic transfer function, equation (2.4), the bias capacitance C_B was calculated to 7.01 pF and the bias resistance R_B to 349 G Ω	48
4.2	a) shows measurements of the mean of input noise density measured at 10 Hz for different amplifier gains. The error bars show $\pm 1\sigma$. b) shows measurements of the mean of input noise density measured at 100 Hz for different amplifier gains. The solid line in c) shows the average noise over time for a measurement made at 10 Hz and gain 9.81. The dashed line shows the corresponding expected noise level, equation (2.10). The solid line in d) shows the average noise over time for a measurement made at 10 Hz and gain 8.92. The dashed line shows the corresponding expected noise level, equation (2.10). Note that there are large fluctuations in the noise at 100 Hz, b) , and that there is no obvious trend with respect to gain, b) . At this frequency, it is expected that the noise decreases with gain until the contributions from the thermal and current noise are larger than the input referred noise, see Figure (2.8). In the same way there is no trend in average noise with respect to gain. Finally the noise levels in d) are in the vicinity of the estimated noise and the noise seems to reach a noise floor at some time intervals.	50

4.3	Histogram of the average noise over time at 100Hz and gain 24.92 measured with the lock-in amplifier. The measurements were taken over 11 h over night. The average noise distribution is not Gaussian and has a large standard deviation.	51
4.4	Typical ECG of abdominal non-contact electrode measurement (blue) conducted with the setup developed in this thesis and abdominal contact electrode from the NI-FECG database, see Part III Section 2.1. The red markers denote fQRS locations. The noise and interference from the non-contact electrode is similar to the contact electrode used for measuring the fECG.	52
4.5	The top figure shows raw data from two non-contact electrodes placed on the chest near each shoulder. Dry ground and shielding were used to reduce PLI. The bottom figure shows the same signal with a 50 Hz notch filter as well as a 2 Hz high-pass filter and a 80 Hz low pass filter applied digitally.	53
4.6	Measurement of electrodes through a cotton t-shirt. One electrode was placed on the chest near the right shoulder and one was placed near the left the apex of the left ventricle. The repolarization of the ventricles (T-wave) is pronounced due to the placement of the electrodes. A 50 Hz notch filter as well as a 0.5 Hz high-pass filter and a 80 Hz low pass filter have been applied digitally.	53

III Software development and testing

1.1	A block diagram describing the adaptive noise cancelling process of the AECG $y(n)$ using a chest reference input signal $u(n)$ at time step n . The fECG is the signal of interest $s(n)$, $\eta(n)$ is the total noise, $\hat{\eta}(n)$ the estimated noise, $e(n)$ the estimation error signal, and $\hat{s}(n)$ the output signal. The aim of the filter is to recursively map the fECG free chest signal as closely as possible onto $y(n)$. This mapping should then correspond to the primary noise source in the AECG, the mECG, and can then be subtracted from the abdominal signal. . . .	60
1.2	The principle and structure of a layered feed forward neural network (a) with McCulloch-Pitts neurons (b). In such a network, an input signal is propagating through forward-pointing (left to right) connections, where a numerical operation is conducted in each neuron (apart from in the input layer) before the signal is passed on to the neurons of the next layer. In a McCulloch-Pitts neuron (circumscribed by the dash-dotted circle), the output signal of the neuron is obtained by first calculating a weighted sum of all the n inputs x_j (which correspond to the outputs from the previous layer), then adding a bias term specified for each neuron and finally using the sum as the argument to some activation function σ . a) is reproduced from Cburnett[7]. . . .	61

1.3	To the left: the network structure of a one-unit recurrent neural network. A single input state x is propagated to a hidden state h with the connection weight U , before getting passed on with the weight W to the output state o . The output of the hidden layer is also passed on to itself the subsequent time step with the weight V , where it will be combined with the succeeding (weighted) input state to form the next hidden layer output. To the right: an unfolded version of the same network, where each of the three “columns” correspond to one time step. Note how the previous output of the hidden layer is combined with the current input state to obtain the new output of the hidden layer. Reproduced from Deloche[8].	63
1.4	A chromosome that encodes two variables with 5-bit accuracy.	64
1.5	The typical crossover procedure in genetic algorithms. Each square is a gene and the dashed, red line is the crossover point which is picked randomly. The first part (i.e. before the crossover point) of the first chromosome is combined with the second part of the second chromosome and vice versa, forming two new chromosomes/individuals.	65
1.6	An illustration of the qualitative difference in time and frequency resolution of a short-time Fourier transform and a wavelet transform. The resolution ratio of the two quantities is constant for the STFT, whereas it can be frequency dependent for a wavelet transform.	67
4.1	A typical result of applying the preprocessing steps on an AECG signal.	79
4.2	Typical independent components from abdominal electrodes. The figures to the left shows three independent components and the figures to the right shows three abdominal leads. The yellow stars are maternal QRS annotations and the red stars are fetal QRS annotations. Looking at the independent components, the ICA algorithm does not separate the fECG and mECG well.	81
4.3	Figure showing the standard LSTM cell. c_t , x_t and h_t are vectors with cell state, input and output. The yellow boxes are neural network layers with sigmoid or tanh activation functions. Red circles are point-wise operations. The two sigmoid units to the left control the information flow to the cell state and are called the forget gate and external input gate. Data from the current input x_t and previous output x_t are presented to the input gate through a tanh activation layer. The sigmoid unit to the right is called the output gate and controls to what extent the current cell state is presented to the output.	82
4.4	Figure showing the MSE between the LSTM prediction and an abdominal lead as a function of the number of units in the LSTM network. The number of units govern how much information that is remembered between each time step. As seen, the networks reaches a small MSE at only 10 units.	84

4.5	The mean and maximum F1 values of the population for each generation. Both measures had seemingly converged properly after the 30 generations that the GA was run for. The fact that the mean F1 value did not converge immediately indicates that the chosen values of the crossover and tournament operators did not favor stronger individuals excessively, which could have got the whole population stuck in a local maximum almost instantly.	88
5.1	A block diagram describing the main features of the Pan-Tompkins QRS detection algorithm, as applied to an AECG where the maternal contribution has been removed.	90
5.2	Flowchart of the algorithm that makes predictions of the QRS peak positions (the red circles) from a CWT of a signal.	91
5.3	Flowchart of the algorithm that combines predictions of the QRS peak positions of multiple channels into a final set of predictions for the given time period.	92
5.4	a) to c) show an example of QRS peak predictions (red circles), based upon the procedure described in Figure 5.2, made on CWTs of the signals from the three abdominal channels for a chosen five second period with an erroneous prediction (the leftmost in b)). The combined, final predictions based upon the predictions for each channel and the algorithm described in Figure 5.3 are shown in d) . The dotted vertical lines denote the corresponding annotations that are used as reference. Note how the erroneous prediction from the second electrode has been omitted in the combined version because the algorithm deemed it as unreliable.	93
6.1	Figure showing the potential from an abdominal lead (blue) and the LSTM mapping from thoracic to abdominal electrode (red). The red circles show annotations for the fQRS complexes. The maternal R peaks (i.e. the high peaks) are more or less completely mirrored by the mapping (the tip of the blue underlying line can be seen slightly for the second peak). Meanwhile, as desired, the fetal QRS complexes are not included in the mapping.	95

List of Tables

II Hardware development and testing

2.1	Input referred noise in INA116 at different amplifier gain and frequencies. The noise levels are given in nV/\sqrt{Hz}	41
-----	--	----

III Software development and testing

3.1	List of methods and their respective performance on a commercial database. The results are reproduced from Behar[3].	77
3.2	List of methods and their respective performance in the PhysioNet Computing in Cardiology Challenge 2013. The results are reproduced from Behar[3].	78
4.1	List of used parameter values in the GA.	89
4.2	List of preprocessing and ESN parameters that were optimized by the GA and their respective search ranges and obtained optimized values.	89
6.1	Positive predictive value for the case of LSTM in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9847.	96
6.2	Sensitivity for the case of LSTM in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9723.	97
6.3	Positive predictive value for the case of ESN in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9847.	97
6.4	Sensitivity for the case of ESN in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9723.	98
6.5	Table showing the positive predictive value, sensitivity and F1 measure for LSTM as the fECG extraction technique and the CWT method for the fQRS detection.	99
6.6	Table showing the positive predictive value, sensitivity and F1 measure for ESN as the fECG extraction technique and the CWT method for the fQRS detection. Files marked with * were part of the test set, and thus not included in the GA optimization procedure. Despite being previously unseen by the network, flawless performance was obtained for these files.	100

6.7	The positive predictive value, sensitivity and F1 measure for the different fECG and fQRS extraction method combinations when the rule of a maximum heart rate difference of ± 5 BPM is applied as tolerance during the grading procedure. The evaluation was performed on the full database.	100
-----	---	-----

Glossary and abbreviations

AECG - Abdominal ECG, the signal mixture one obtains upon measurements on the maternal abdomen.	vation of oxygen.
BPM - Beats per minute	KF - Kalman filter
CTG - Cardiotocograph	LSTM - Long short term memory neural network
CWT - Continuous wavelet transform	mHR - Maternal heart rate
ECG - Electrocardiography	mQRS - Maternal QRS wave
EKF - Extended Kalman filter	MSE - Mean square error
ESN - Echo state neural network	NI-fECG - Non-invasive fetal electrocardiogram
fECG - Fetal electrocardiogram	PCB - Printed circuit board
FFNN - Feed forward neural network	PLI - Power line interference
fHR - Fetal heart rate	PPV - Positive predictive value
fQRS - Fetal QRS wave	PT - The Pan-Tompkins QRS detection algorithm
GA - Genetic algorithm	SNR - Signal-to-noise ratio
Hypoxia Condition characterized by a deprivation of oxygen.	STFT - Short-time Fourier transform

Part I

Context and problems to be solved

Chapter 1

Introduction

1.1 Motivation and historical context

In 2017, the global neonatal mortality rate (death within 28 days after live birth) was 18 in 1000. In the native country of the authors, Sweden, that figure is 2 in 1000. While that is still undesirably high, the situation is substantially worse in e.g. sub-Saharan countries whose neonatal mortality rate is more than 13 times as high (27 per 1000)[9].

According to WHO, World Health Organization, there were 2.6 million stillbirths in the world in 2015, corresponding to almost 7200 deaths a day. 98% of these deaths occurred in low- and middle-income countries and about half during the intrapartum period, i.e. during childbirth. Constituting the most high-risk period, the intrapartum proportion of stillbirths still differ a lot worldwide, with 10% in some high-income regions to 59% in south Asia. A similar distribution of maternal deaths can be seen and the correlation with areas with substandard health care at birth is substantial[10].

Consequently, there is a dire need for a cheap, effective and accessible fetal health monitoring technique during pregnancy and labor. Furthermore, the main cause of birth defect-related deaths is heart defects[11]. Nowadays, many congenital cardiac defects are identified and diagnosed prepartum with routine ultrasound scans using the Doppler ultrasound sensor in cardiotocographs (CTG). The CTG apparatus consists of said Doppler sensor and pressure sensor to measure uterine muscle contractions. An electrode is sometimes placed directly on the fetal scalp if the fetal heart rate (fHR) measurements are inadequate during labour. Indicators are derived from the fHR and maternal contractions and used to monitor the fetal well being. However, some heart diseases cannot be detected using this modality and there is a recurrent issue where the CTG confuses the maternal heart rate for the fetal heart rate. In a study of approximately 10 000 deliveries, five examples of unexpected adverse fetal outcome attributed to this signal ambiguity were encountered, including still births, neonatal deaths, and severe brain damage[12].

The non-invasive fetal electrocardiogram (NI-fECG) could constitute a viable alter-

native to Doppler ultrasound recording. Using electrodes on or close to the maternal abdomen, both a more accurate estimation of the fetal heart rate (fHR) and additional information about the electrical activity of the fetal heart, obtained through the study of the FEKG morphology, could hopefully be achieved. Unfortunately, it is a highly complicated task to extract the fEKG from the abdominal signal mixture in a reliable manner. As a consequence, its usage has been very limited to this day.

1.2 Outline of thesis

The objective of this thesis was to investigate the feasibility of using NI-FEKG as a cheap, precise, and reliable modality for beat-to-beat fHR monitoring. To that end, a highly sensitive non-contact electrode was designed and a set of novel algorithms and tools for fHR extraction developed and evaluated. The thesis is split into three main parts: I) Context and problems to be solved, II) Hardware development and testing III) Software development and testing. These parts are then followed by a brief general summary of findings, conclusions and remaining future work.

1.2.1 Part I - Context and problems to be solved

The main intention of this part is to introduce the background of the field, address issues with the existing technology, and motivate the need and requirements for novel fHR extraction methods. The part includes succinct descriptions of the general electrophysiology of the heart and the physical origin of an electrocardiogram, as well as features and characteristics of the fEKG during different stages of gestation. Furthermore, a summary of the different subsignals comprising an EKG and their discerning characteristics is given.

1.2.2 Part II - Hardware development and testing

This part begins with an introduction to theoretical physical concepts crucial for the operation of the developed sensors. It then describes the various considerations and experimental evaluations leading to the final electrode prototype, including the choice of non-contact sensors and certain noise reduction features. Finally, a description of other vital components of an EKG measurement setup is given and some results from their implementation presented.

1.2.3 Part III - Software development and testing

Starting off with brief descriptions of relevant signal processing techniques, this part then reviews the pros and cons and the potential of the contemporary state-of-the-art fEKG extraction methods. After that, the data sets used for testing

and evaluation of the developed algorithms are presented and the score system for benchmarking defined. Lastly, the experimentally evaluated methods are described and their performances presented and discussed.

1.2.4 Part IV - Conclusion and future work

This part concisely summarizes the findings of this thesis and emphasizes promising directions for further future research.

Chapter 2

Electrophysiology of the heart

This chapter summarizes the basic physiology and key events of a human heart, which eventually leads to the potential difference on different parts of the skin that can be mapped in an electrocardiogram.

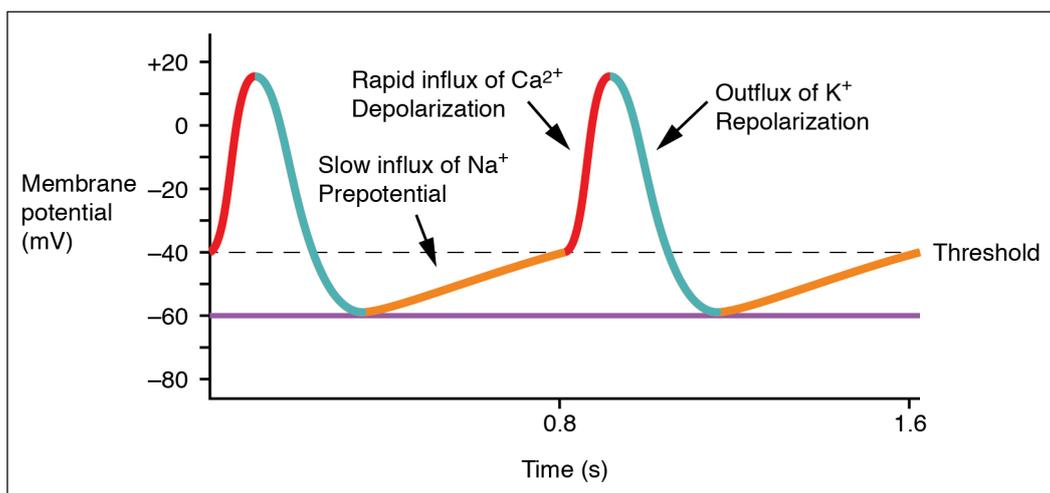


Figure 2.1: Figure showing the periodic action potentials generated in a pacemaker cell in the SA node. The orange line shows the slow depolarization due to sodium influx. At the threshold the cell rapidly depolarizes (red line) and then repolarizes (green line). In this manner the pacemaker cells initiate the propagating action potential, and therefore sets the rhythm of the heart. Reproduced from OpenStax College.

2.1 Cardiac muscle contraction

Virtually all cells have a potential across the cell membrane [13]. This potential is set by the interplay between ion pumps that actively transport specific ions across the membrane and ion channels that set the ion permeability of the membrane. The cardiac muscle cell (*myocardocyte*) has a membrane potential of approximately -90 mV at rest [14] measured from the inside of the cell. However, approximately one

percent of myocytes are never truly at rest and are known as pacemaker cells. Due to a net inflow of sodium ions the membrane potential slowly *depolarizes* over time[15]. At a certain threshold potential, approximately -40 mV[15], a series of rapid changes in ion permeability causes the cell to rapidly depolarize, and then repolarize back to the resting potential. This brief change in membrane potential is known as a *cardiac action potential* and triggers muscle contraction through an intracellular signal chain. A depiction of the membrane potential for a pacemaker cell is shown in Figure 2.1¹. The orange line is the slow depolarization due to sodium leakage, the red line is the rapid depolarization and the green line is the repolarization. The action potential propagates through the cell and to neighbouring cells through intercellular connections. Through a structured network of muscle cells and specialized conduction cells, the action potential propagates through the heart and causes the atrium and ventricles to contract in a coordinated fashion. The wave of depolarization can be detected on the skin surface, which is the basis of electrocardiography (ECG).

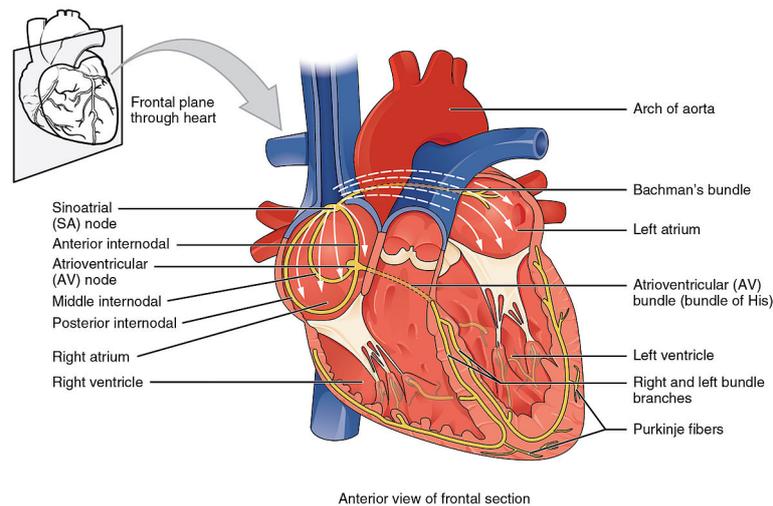


Figure 2.2: Anatomical illustration of a frontal plane cross section of the human heart. The illustration shows the propagation path of the action potential. The action potential originates from the SA node and propagates through both atria and also to the AV node. At the AV node, which is the only electrical connection between the atria and ventricles, the action potential is delayed before it propagates through specialized conductive fibers, Purkinje fibers, to both ventricles. Reproduced from OpenStax College.

2.2 The cardiac cycle

The pacemaker cells with the fastest rate of depolarization set the origin of the heart beat and the heart rate. A cluster of pacemaker cells in the sinoatrial (SA) node,

¹OpenStax College. *SA Node Tracing*. Wikimedia Commons. License: <https://creativecommons.org/licenses/by/3.0/deed.en>. Link: https://commons.wikimedia.org/wiki/File:2020_SA_Node_Tracing.jpg

just below the superior vena cava in the left atrium (Figure 2.2²), are in general the cells with the fastest rate of depolarization and therefore the natural pacemaker of the heart.

The rate of the pacemaker is influenced by sympathetic and parasympathetic nerve fibers. The action potential propagates to the left and right atrium and to a collection of pacemaker cells in the interatrial septa known as the atrioventricular (AV) node. The AV node is the only passage point for signals between the atria and ventricles. Here, the signals are delayed so that the atria have time to contract. After the delay, the action potential wave propagates to the ventricles through specialized conductive fibers, Purkinje fibers. The Purkinje fibers permeates both ventricles and the action potential proceeds through intercellular connections in the myocardium.

2.3 Dipole moment of layers and closed surfaces

As previously mentioned the resting potential of myocytes is approximately -90 mV and arises from concentrations of different ions across the membrane. A simple electrostatic depiction of the cell membrane would be two separated sheets with charge densities $+\sigma$ and $-\sigma$ separated at a distance a . At an observation point at a distance much greater than a , the field from a single dipole in vacuum is

$$V = \frac{1}{4\pi\epsilon_0} \frac{\mathbf{m} \cdot \mathbf{r}}{r^3} = \frac{1}{4\pi\epsilon_0} \left(\frac{m \cos \theta}{r^2} \right),$$

where \mathbf{m} is the dipole moment and r the distance between the observation point and the center point between the two charges. Consider the contribution to the potential at an observation point from a small part of the cell membrane,

$$dV = \frac{1}{4\pi\epsilon_0} \frac{\cos \theta}{r^2} dm = \{dm = \sigma a dS\} = \frac{\sigma a \cos \theta}{4\pi\epsilon_0 r^2} dS. \quad (2.1)$$

In spherical coordinates with θ defined as the angle from the horizontal plane to the z axis, the solid angle differential element is given by $d\Omega = \frac{\cos(\theta)}{r^2} dS$. Hence,

$$dV = \frac{\sigma a}{4\pi\epsilon_0} d\Omega \quad (2.2)$$

and the total potential at the observation point is therefore given by

$$V = \int_{\Omega_S} \frac{\sigma a}{4\pi\epsilon_0} d\Omega$$

For a closed surface of two sheets of separated charges, the potential will be zero as any differential solid angle will intersect two surfaces with opposite signs of the

²OpenStax College. *Conduction System of Heart*. Wikimedia Commons. License: <https://creativecommons.org/licenses/by/3.0/deed.en>. Link: <https://commons.wikimedia.org/w/index.php?curid=30148214>

charge density σ . Therefore a fully polarized or depolarized cell will have no net dipole moment. As this is true for all cells, the potential from a heart at rest is zero.

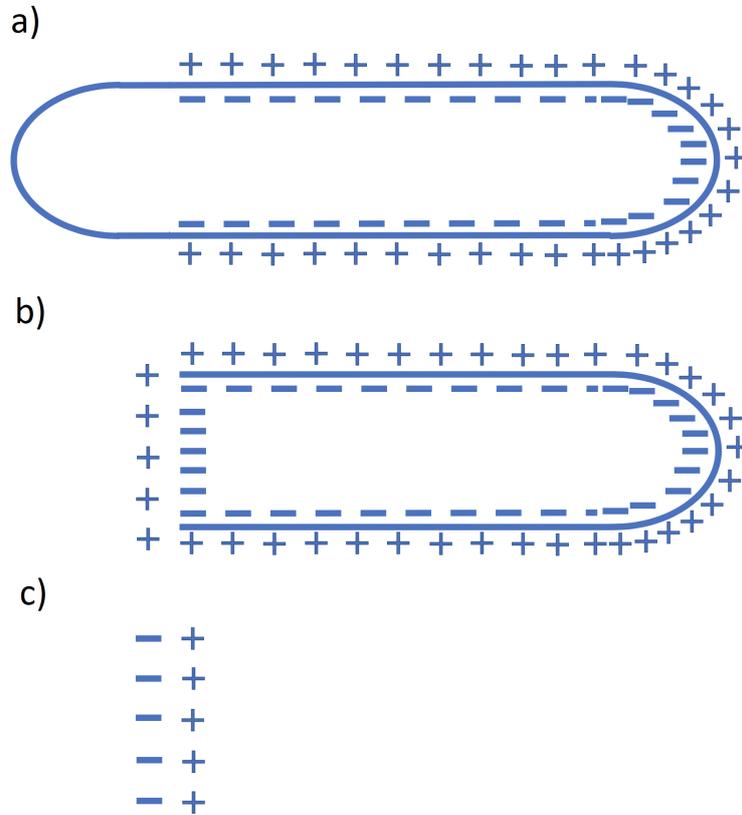


Figure 2.3: Figure showing an electrostatic model of a partially depolarized heart. **a)** shows the cell as the action potential has depolarized part of the cell. The right side of the cell is yet to be depolarized by the propagating action potential. An equivalent model of **a)** is the superposition of **b)** and **c)**. Since there is no net electric field from the closed surface **b)**, the only contribution comes from **c)**. The model for the heart as a whole becomes equivalent as the action potential propagates from cell to cell. This means that the depolarization of the heart can be seen as a propagating charged surface that intersects the current position of the action potential wave.

How can a heart cell which is partially depolarized be described within this picture? As the action potential propagates through the heart, parts of the myocytes are depolarized and parts are polarized. At any instant one can think of a cross section of the heart which divides the heart into a part that is fully depolarized and a part that is yet to be depolarized. As discussed above, both the depolarized and polarized part of the heart will have no net contribution to the potential at an observation point. However, the cells at the cross section do. A schematic figure of a partially polarized cell is shown in Figure 2.3. The cell **a)** can be formed through superposition of a closed sheet of a charge double layer, **b)**, and a residual surface **c)**. The potential from the closed surface is zero. Thus, only the residual

surface of separated charges in \mathbf{c}) will contribute to the potential at the observation point.

Integrating (2.1) over the residual surface of all partially depolarized cells in the cross section results in the total potential at an observation point P . Assuming that this point is far away from the heart, $\mathbf{r} = \mathbf{r}_P$ is approximately constant. With $r = r_P$ and $\theta = \theta_p$

$$V = \int_S \frac{\sigma a \cos \theta_P}{4\pi\epsilon_0 r_P^2} dS' = \frac{\sigma a \cos \theta_P}{4\pi\epsilon_0 r_P^2} \int_S dS' = \frac{\sigma a S \cos \theta_P}{4\pi\epsilon_0 r_P^2} = \frac{1}{4\pi\epsilon_0} \frac{\tilde{\mathbf{m}} \cdot \mathbf{r}_P}{r_P^3}$$

where S is the total cross section and $\tilde{\mathbf{m}}$ is the *net dipole moment*. If the observation point is close to the heart, the dipoles that are close to the heart will contribute more than dipoles far away. Either way the potential at the skin surface can be seen as the instantaneous projection of the net dipole moment on \mathbf{r}_P when P is far away from the heart, or as the sum of contributions of many small dipoles when P is close to the heart. The net dipole moment is often referred to as the "electrical force vector" or "activity" in medical literature[16].

One might question why the model was set out in vacuum when the body to an excellent degree is approximated as purely resistive[17]. One can however show that as the resistance inside and outside of the cell is the same, a conductive model would yield the same results as in a vacuum[16, 18].

Also, one might wonder if there will be any differences in an electrodynamic model compared to the electrostatic equivalent. Signals stemming from voltage sources will experience a negligible phase shift and attenuation for frequencies under 1000 Hz[17], which includes the contributions from both the fetal and the maternal heart. Therefore the associated time shift from voltage source to body surface is negligible and the propagation can be considered as quasi-static.

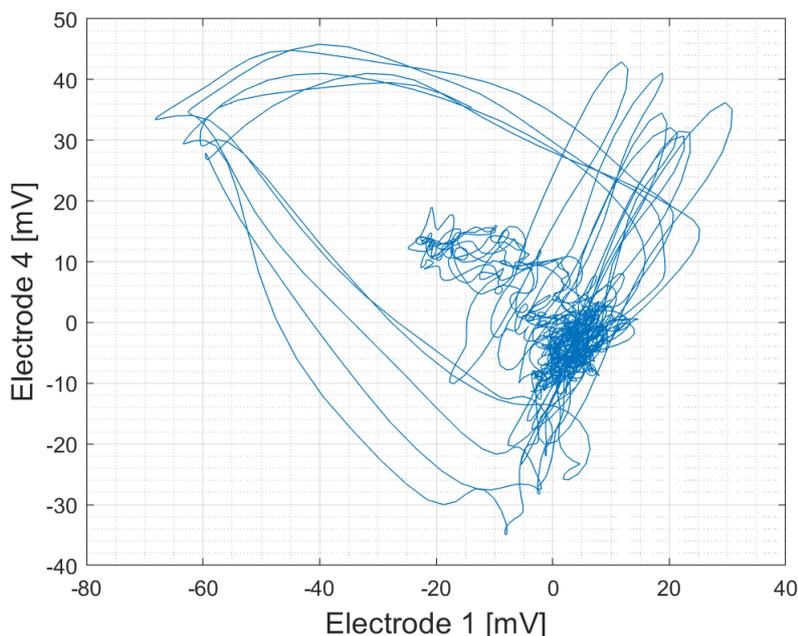


Figure 2.4: Measurement made from two abdominal electrodes. The large loop is the contribution of the maternal QRS wave and the smaller loop is the contribution of the fetal QRS wave. Looking at the figure it is clear that the maternal and fetal net dipole moments traverse different paths in space. The data is taken from the Abdominal and Direct Fetal Electrocardiogram Database [1].

2.4 Electrocardiogram as a net dipole map

The ECG is constructed by mapping the electrical activity from the heart on the skin surface. Measurements are taken as differences between the readings of two electrodes at the skin, also known as a lead. With electrode one situated at \mathbf{r}_1 and electrode two at \mathbf{r}_2 with the origin from the heart, according to (2.3),

$$V = \frac{1}{4\pi\epsilon_0 r^3} (\mathbf{r}_1 - \mathbf{r}_2) \cdot \tilde{\mathbf{m}} = \frac{1}{4\pi\epsilon_0 r^3} \mathbf{R} \cdot \tilde{\mathbf{m}}, \quad (2.3)$$

where \mathbf{R} is a vector pointing from electrode 1 to electrode 2. During the cardiac cycle the net dipole moment makes very distinct loops in space[16]. The ECG results from the projection of the dipole moment during the loops on to combinations of electrodes on the skin. A projection of the dipole moment from a maternal and fetal heart onto two abdominal electrodes is shown in Figure 2.4. Looking at the figure, it is clear that the maternal and fetal net dipole moments traverse different paths in space.

A conventional ECG contains measurements consists of at least 9 leads. Three electrodes are limb leads. They are placed one on each arm as well as on the left leg. As the body is well approximated as purely resistive[17], the limbs can be seen as cylindrical conductors. The potential on each limb lead is therefore proportional

to the point where the limbs meet abdomen and torso. The three limb leads form a geometrical reference system for other electrode measurements known as Einthoven's triangle. Three *augmented leads* are formed from combinations of the limb leads to show the projection of the dipole moment onto three new axes. These axes originate from the center of the triangle to each limb lead. aV_L to the left electrode, aV_R to the right electrode and aV_F to the electrode on the leg (F as in foot). Apart from limb leads, there are generally six precordial leads placed across the chest. Each of these measure the projection of the dipole moment from the center of Einthoven's triangle, i.e. the sum of the limb leads, to points around the heart. As these leads are placed close to the heart, the assumption that all dipoles are equidistant to the electrode breaks down. That is, each of these electrode are more heavily influenced by the myocardium tissue near to it. With strategic placement of the electrodes, each of the precordial leads accurately measure the contribution of the action potential from specific parts of the heart. This information can be useful within cardiac diagnostics.

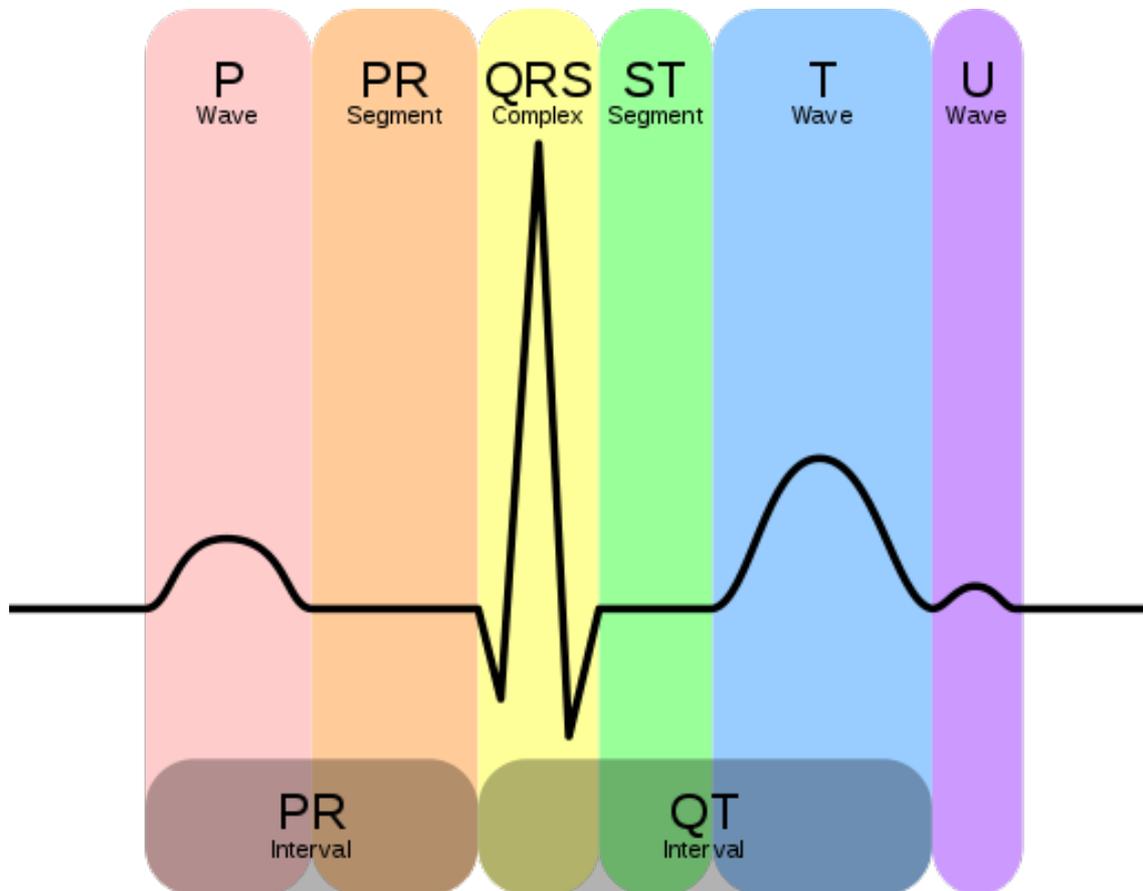


Figure 2.5: Illustration of an ECG waveform recorded on a lead on the skin surface. The P wave arises from atrial contraction, the PR segment is from the delay in the SA node and the QRS complex from ventricular depolarization. The ST is the isoelectric interval in between ventricular depolarization and repolarization. The T wave comes from ventricular repolarization. The U wave is not always visible and comes from repolarization of so-called purkinje fibers.

Each lead has three characteristic regions corresponding to cardiac events. First, the depolarization and contraction of the atrium are manifested as a small peak known as the P wave, see Figure 2.5. The dipole moment arising from ventricular depolarization is projected to the electrodes as a large wave known as the QRS complex. The repolarization of the ventricles manifests as a more slowly varying wave. The repolarization of the atrium is in general masked by the QRS complex. The sign and amplitude of the events at each electrode depend on electrode placement relative to the instantaneous net dipole moment.

Chapter 3

The fetal ECG

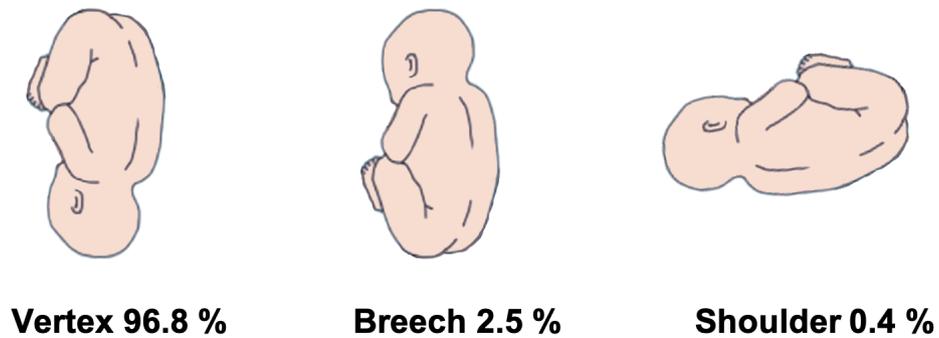


Figure 3.1: The most common fetal presentations and their concomitant prevalences in utero at the time of labor. The statistics are fetched from Symonds et al.[2].

The first functional organ to take shape in vertebrate embryos is the heart, which starts beating already at the end of the first month. Fetal movement is common during the first two trimesters of pregnancy, with an estimated frequency of once per four-five minutes between eight and 30 weeks of gestation[19]. At later stages the movements in the uterus are restricted by the fetus's size. Figure 3.1 illustrates the incidence of the most common fetal presentations during labor. The vertex position, i.e. head down and directed toward the birth channel, is by far the most common (96.8%).

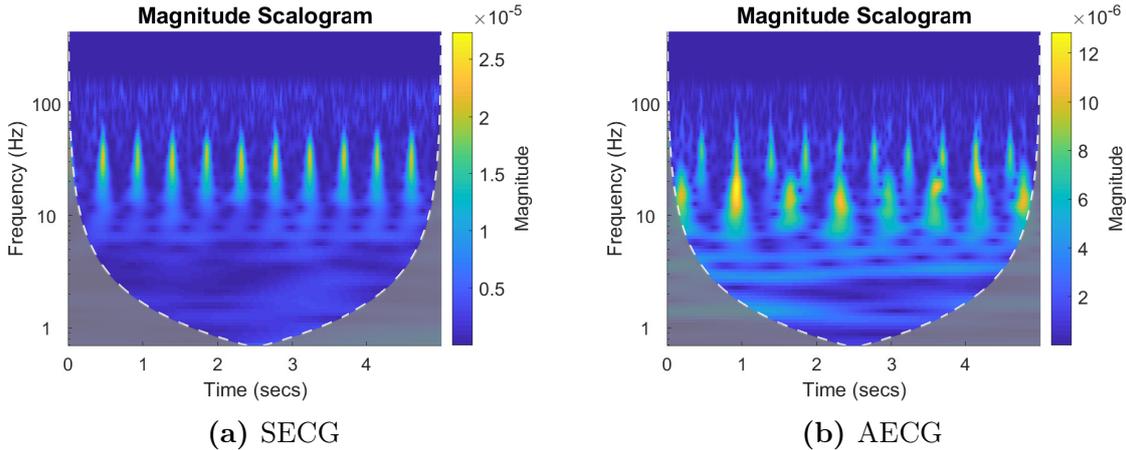


Figure 3.2: **a)** Time-frequency spectrum of an SECG waveform. The yellow nodes are fetal heart beats. For each heart beat, most of the spectra is contained within the bandwidth 15-70 Hz. **b)** The time frequency spectrum of an abdominal electrode. Note that the fetal heart rate is higher than the maternal, the maternal heart beats have a higher amplitude than the fetal equivalents, and that the fetal wave packet contains higher frequencies than the maternal ditto. The data is taken from the Abdominal and Direct Fetal Electrocardiogram Database[1].

3.1 Similarities and differences to adult ECG signals

The recorded body surface potential corresponds to the electrical wave that has propagated from the myocardium to the body surface. Usually the signal-to-noise ratio (SNR) is considerably worse for fECGs than for adult ECGs due to the additional obstructing media between the heart and the electrodes and simply because the fetal heart is smaller. Furthermore, the fetal heart’s development stage, fetal position and movements all affect the signal strength, orientation and non-stationary features of the fECG recorded from the maternal abdomen[3].

In general, however, there are many similarities with adult ECGs. The fetal electrocardiogram also contains the different characteristic waves as defined by Einthoven; P, QRS, and T, and the mECG and fECG usually overlap in both the time and frequency domain (although the fECG normally mainly comprises higher frequencies). Among the more pronounced dissimilarities, the fHR is normally higher than the adult equivalent (even though the normal range varies over the pregnancy)[20] and the variability of the heart rate is lower for the fetus[21]. Some of these characteristics are illustrated in Figure 3.2. The so-called magnitude scalograms show the time-frequency spectrum of an abdominal lead and a scalp electrode respectively. As seen, the fECG has a faster heart rate and each wave packet contains higher frequencies, although the spectra are overlapping.

Some of the smaller features of the fECG morphology are often more onerous to detect clearly than in adult ECGs. For instance, the T wave usually has a very low

amplitude for the fetus[22].

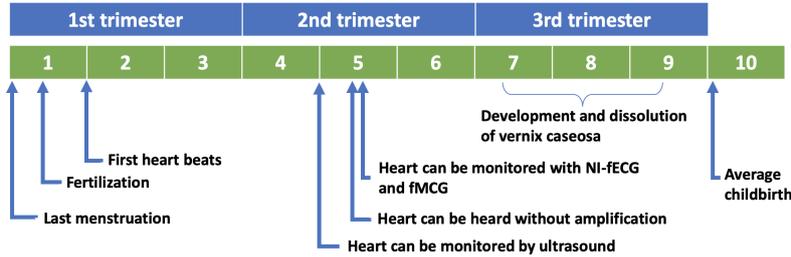


Figure 3.3: Timeline of prenatal development, including some key entries for fetal monitoring. Vernix caseosa is a thin, highly non-conductive layer normally surrounding the fetus between around the 28th-32nd week and the 37th-38th week that has a significant adverse affect on NI-fECG extraction. Based upon statistics given in Behar[3].

3.2 Prenatal development landmarks relevant to fECG extraction

Figure 3.3 displays important periods of the gestation, including parts that play important roles in the feasibility of fECG extraction. Starting around the 20th week, the fetal heart can be heard without amplification[21] and monitored through Doppler ultrasound, fECG and fetal magnetocardiography (fMCG)[23]. Somewhere around the 28th-32nd week of gestation, a thin layer called the vernix caseosa takes shape around the fetus and normally does not dissolve until the 37th-38th week[24]. Since this layer is highly non-conductive, it has been shown to aggravate the difficulty of NI-fECG extraction substantially[25, 26].

3.3 Significant features of the fECG

The most significant characteristics in fetal monitoring are related to rhythm (fHR) and morphology (e.g. changes in ST and QT segments). Fetal heart rate variability reflects the autoregulation by the autonomic nervous system, but its interpretation is challenging since it can be affected by factors such as fetal sleep state, stage of pregnancy, and drugs[27]. Antepartum, the fHR can be used for screening of intrauterine growth restricted fetuses[28] and intrapartum as a fetal distress indicator[29]. Morphological analyses of the fECG is outside of the scope of this thesis, but has the potential of many clinical applications such as intrapartum hypoxia detection[30]. However, as a consequence of the current inability to accomplish accurate extractions of the full fECG morphology (using either SECG or NI-fECG), a very limited number of studies have managed to show a substantial improvement regarding fetal outcomes or reductions in the proportion of cesarean deliveries[3].

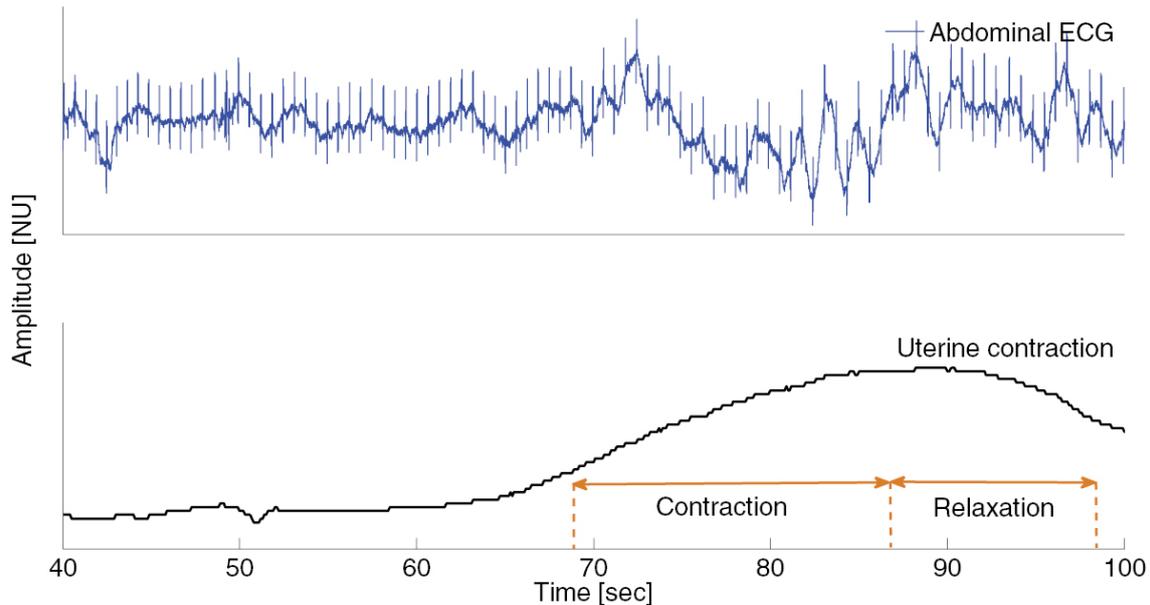


Figure 3.4: A uterine contraction, recorded by an intrauterine pressure transducer, and its effect on the signal quality of an ECG measured on the maternal abdomen. After 70 seconds there is a deterioration in the ECG quality that correlates with the onset of the contraction. The figure has been reproduced by permission of IOP Publishing.

3.4 Extraction of uterine contractions and fetal respiratory movements from the abdominal ECG

During childbirth, the timing of uterine contractions are normally tracked with an external tocodynamometer attached to the CTG or an intrauterine pressure catheter. It is however also possible to extract the contractions from the multi-sourced signal measured in abdominally recorded ECGs[31]. An example of the effects of a contraction on an abdominal ECG (AECG) can be seen in Figure 3.4¹. Even though it degrades the signal quality of the ECG and complicates NI-fECG extraction, it also indicates that contractions could be extracted through assessment of its concomitant noise in the AECG.

The first respiratory movements of a fetus were observed already in 1888[32]. Later on it was inferred that the breathing-like movements of the fetus were necessary as stimulation of lung development. Since oxygen is provided by the mother during the pregnancy, the fetal lungs ostensibly do not serve an oxygenation-related purpose antepartum[3]. The respiratory-like movements can be seen after the 10th pregnancy week[33] and healthy fetuses should have around 57 breaths per minute (brpm) 30-33 weeks into gestation and around 48 brpm at week 37-40[34]. Breathing rate extraction from adult ECGs is nowadays implemented in commercial applications[3],

¹Joachim Behar et al. *A practical guide to non-invasive foetal electrocardiogram extraction and analysis*. April 2016. *Physiological Measurement*. 37(5):R1–R35. DOI: 10.1088/0967-3334/37/5/r1. ©Institute of Physics and Engineering in Medicine. All rights reserved.

but the feasibility of the equivalent extraction from the fECG, although theoretically possible, has not yet been rigorously studied.

Chapter 4

Issues with current technology and problems to be solved

4.1 Current technology and its drawbacks

There are many different indicators and parameters which doctors can use to determine whether pregnancy and childbirth are developing as expected. Many crucial indicators are studied with CTG, cardiotocography, which measures uterine contractions via a pressure sensor and the fHR via an ultrasonic Doppler sensor or scalp electrode. Ever since the end of the 19th century, decelerations of the fHR have been associated with fetal distress. By the middle of the 20th century, sporadic observations of fetal heart sounds, auscultation, had become standard clinical practice and commercial fetal heart rate monitors became widespread during the mid-1970s. The expectations were high that continuous monitoring of the fHR drastically would reduce the incidence of undiagnosed fetal hypoxia (deprivation of oxygen)[11]. However, multiple studies showed that the outputs of fHR monitors were frequently unreliable and difficult to interpret, rendering a surge in the rates of painful, invasive, and expensive cesarean sections and subsequent postnatal depressions[35] and postoperative pain that negatively affected breastfeeding and infant care[36]. Conversely, little evidence supported that reductions in adverse outcomes could be attributed to the fHR monitor usage[11].

In addition to being expensive, the usage of Doppler ultrasound sensors comes with several challenges and drawbacks. Even though it is non-invasive, the measurement procedure is not passive and it has not been shown that long-term exposure to ultrasound radiation is harmless for the fetus[37]. Furthermore, it is beneficial for the mother to be mobile during labor and the sensor can often reduce her mobility significantly. Even if the sensors would be connected wirelessly, the transducer and cables are awkward and unwieldy and have substantial power requirements. Another disadvantage is that the technique is dependent upon a secondary phenomena (the mechanical activity of the heart). This, together with the utilization of the usual averaging procedures, makes the produced data ill-suited for the clinically significant beat-to-beat analysis of the fHR variability[38].

In particular, as mentioned in Section 1.1, Doppler ultrasound heart rate extraction occasionally confuses the fHR with the maternal heart rate, mHR, if the ultrasound transducer gets misoriented. In 10 000 monitored deliveries, five examples of unexpected adverse fetal outcome attributed to this signal ambiguity were encountered, including still births, neonatal deaths, and severe brain damage. Further cases of signal mix-ups were recorded during the study, but fortunately these did not have severe effects on the fetal outcome[12].

There are other electronic fetal monitoring techniques under investigation as well. Fetal magnetocardiography (fMCG) detects the magnetic field of the fetal heart through SQUIDS, superconducting quantum interference devices, positioned close to the maternal abdomen. While facilitating morphological analysis of the measured biopotential due to a higher signal-to-noise ratio compared to the fECG, it is expensive, demands skilled personnel, and as of now lacks the possibility of long term monitoring due to the size and cost of the apparatus[3]. Fetal pulse oximetry, FPO, is a fairly new method (the first device to get FDA approval received it in 2000[39]), but clinical studies showed that the usage of this technique in addition to the CTG was not associated with significant differences in neonatal outcomes. Its effects on cesarean frequency could not be determined[39, 40].

The second standard method in clinical practice today in addition to the CTG, the scalp fetal electrocardiogram (SECG), is measured by attaching an electrode to the scalp of the fetus. The method can only be used intrapartum and not at earlier stages of pregnancy and it also carries a slight risk of infection. Moreover, it is not unusual that the scalp electrode loses contact due to maternal and fetal movement[41]. Due to the invasiveness and the limitations in accessibility and consistency, SECG has not been adopted as a routine procedure in all pregnancies. Furthermore, since only one differential electrode is possible, the resultant cardiac activity will be projected on a specific lead axis. A mapping of the three dimensional field originating from the fetal heart is therefore not possible with this modality[3].

With the CTG only being able to give an estimate of the fHR, the SECG being invasive and mostly unavailable, and other methods being expensive, inaccessible and/or having dubious efficacy, the NI-fECG received a lot of interest within the scientific community during the past decades. It could constitute an alternative monitoring method that combines the advantages of the current clinical methods, being both non-invasive and accurately determining the fHR and other useful information about the fetal heart's electrical activity. However, unsurprisingly, the NI-fECG comes with its own inherent challenges. The rapidly developing nature of the fetal heart requires the signal processing methods to be able to manage large extents of intra- and inter-subject variability[3]. Moreover, the embedding of the weak fECG field strength within the recorded abdominal mixture makes an accurate reconstruction of the full fECG waveform a highly complicated task.

4.2 The other subsignals and how to circumvent them or reduce their prominence

The fECG signal itself usually manifests frequencies from around 0.05 to 100 Hz. For composite abdominal signals the maximal amplitude of the maternal QRS complex is normally in the range of 100-150 μV , whereas the fetal equivalent is merely up to 60 μV at best. Moreover, the fECG is often embedded in electrical noise as well as other biopotentials and the shape of the fECG signal is dependent on the electrode placement (there is no well-defined standard positioning for optimal fECG acquisition[42]), position of the fetus and the gestational age. All these additional signals aggravate the fECG extraction substantially. Consequently, understanding their characteristics and the properties that distinguish them from the fetal signal is of paramount importance for a successful extraction of the fECG. The following categorization of relevant subsignals can be made[6]:

- mECG signal: the maternal heart signal resembles and interferes with the fECG the most out of all the subsignals in the composite abdominal signal. It's frequency range partly overlaps with the fetal signal and thus will not regular filtering alone suffice to separate the two signals.
- Power line interference (PLI): PLI comprises 50 Hz and its harmonics and may, thanks to the well-defined frequency content, be modeled as (a combination of) sinusoids. Normally it should not vary noticeably during a measurement in a specific environment and keeping the electrode within a fairly electrically shielded environment such as a Faraday cage together with post-processing should reduce this noise significantly.
- Maternal muscle noise (EMG): this signal stems from maternal movement, mainly from the abdominal and leg muscles and could e.g. be registered through a reference electrode on the thighs. This electromyographic noise has a wide frequency span and can sometimes be difficult to identify and separate from the rest of the abdominal signal, but usually it should not have the same level of periodicity as the respective heart rates.
- Electrode contact noise: this kind of transient interference originates from loss of contact between a contact electrode and the skin, which disconnects the modality from the subject. The ensuing noise can be modeled as a random baseline transition, occurring once or several times in close succession, that decays exponentially to the baseline value and has a superimposed 60 Hz component. Naturally, this noise could be avoided through usage of non-contact electrodes.
- Motion artifacts: predominantly originating from two sources, electrode-skin interface and the electrode cables, this noise is present for both contact and non-contact electrode configurations. Proper design of the electrode circuitry and full setup as well as a filter that crops low frequencies can suppress these artifacts.

- Inherent noise in electronic components: intrinsic noise is present in all electronics and cannot be eliminated. This adds a restriction on the available selection of components, which have to be below the (highly application dependent) allowed noise levels.
- Ambient noise: electromagnetic radiation is omnipresent and it is practically impossible to completely avoid exposing the skin of the subject to some level of ambient noise.
- Baseline drift and ECG amplitude modulation with respiration: the respiration-induced baseline drift can be modeled as a sinusoid with the frequency of the respiration cycle superimposed on the ECG signal. Unfortunately, the amplitude of the ECG signal itself may vary by about 15% with respiration.

Part II

Hardware development and testing

Chapter 1

Basic theory/Fundamental principles

1.1 Capacitance

A capacitor comprises two conductors separated by a non-conductive region. The non-conductive region can either consist of complete vacuum or a dielectric material (an electrical insulator that can be polarized by an applied electric field) such as glass, air and ceramics. As stated by Coulomb's law, a charge carrier in one conductor will exert a distance- and charge-dependent force on the charge carriers of the other conductor. Since opposite polarities will get attracted and equal polarities will get repelled, an opposite surface charge will be induced on the other conductor. Consequently, the two conductors will hold opposite charges of the same magnitude on the surfaces that face each other which will render an electric field between them[43].

An ideal capacitor is characterized by a constant capacitance C , which has the unit farads (F) in the SI system and is defined as the ratio of the charge magnitude Q on each conductor to the voltage V between them:

$$C = \frac{Q}{V}.$$

Thus, a capacitance of one F corresponds to that one coulomb of charge on each conductor would result in a voltage of 1 V across the capacitor.

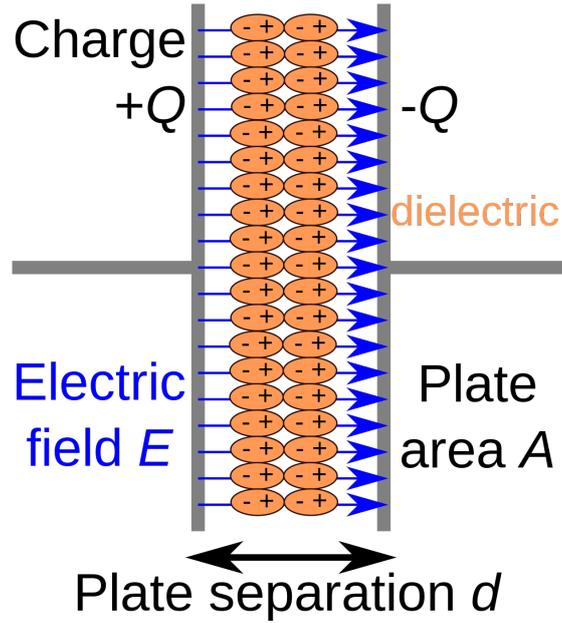


Figure 1.1: Schematic of the charge separation in a parallel-plate capacitor and its rendered internal electric field. The dielectric material between the plates (orange ellipses) becomes polarized due to the charge displacement, which reduces the total internal field and increases the capacitance. Reproduced from November[4].

1.1.1 Displacement current and capacitive sensing

Capacitive sensors are instruments that are able to detect electric signals through capacitive coupling. The principle can be modelled by a parallel-plate capacitor and is based on that the sensor makes up one of the two plates and the surface of the investigated object the other plate. The capacitance for a parallel-plate capacitor is given by

$$C = \varepsilon_0 \varepsilon_r \frac{A}{d}, \quad (1.1)$$

where ε_0 is the vacuum permittivity, ε_r the relative permittivity, A the area of each of the plates, and d the distance between them (i.e. the thickness of the dielectric medium). The charge distribution on the object surface at every instant of time gives rise to an electric field between the two “plates”. The redistribution of charge carriers on the sensor plate, as a response to the Coulomb forces exerted by the surface charges on the object, is called displacement current and acts to achieve temporary equilibrium. Since all quantities could be time-dependent and $I(t) = dQ/dt$, it may be written

$$I(t) = C(t) \frac{dV(t)}{dt}. \quad (1.2)$$

During the cycle of a heartbeat, the myocardium-generated potential will propagate to the skin, where a redistribution of free charges will occur so as to minimize the potential differences along all body dimensions at each instant of time. Thus a real

current, reflecting the shape and periodicity of the cardiac potential, will flow in the body and generate a time-varying electric field between the constantly redistributing skin surface charges and the sensor. By measuring the displacement current in the sensor plate, one should therefore be able to record an ECG.

In the case of capacitive sensing of the heartbeat induced skin current, with air as the dielectric medium between the sensor and the skin, the capacitance should according to (1.1) stay constant as long the distance is kept constant (since the areas normally should stay the same). This is desirable for ECG measurements, since there according to (1.2) otherwise would not be a linear relationship between the displacement current and the electric field. However, subject movements may lead to slight variations in distance between the body and the sensor. Thus, it is important to take precautionary measures to minimize these distance variations and the impact they will have on the total capacitance between the skin and the active components of the sensor. This will be discussed in Section 2.3.3.

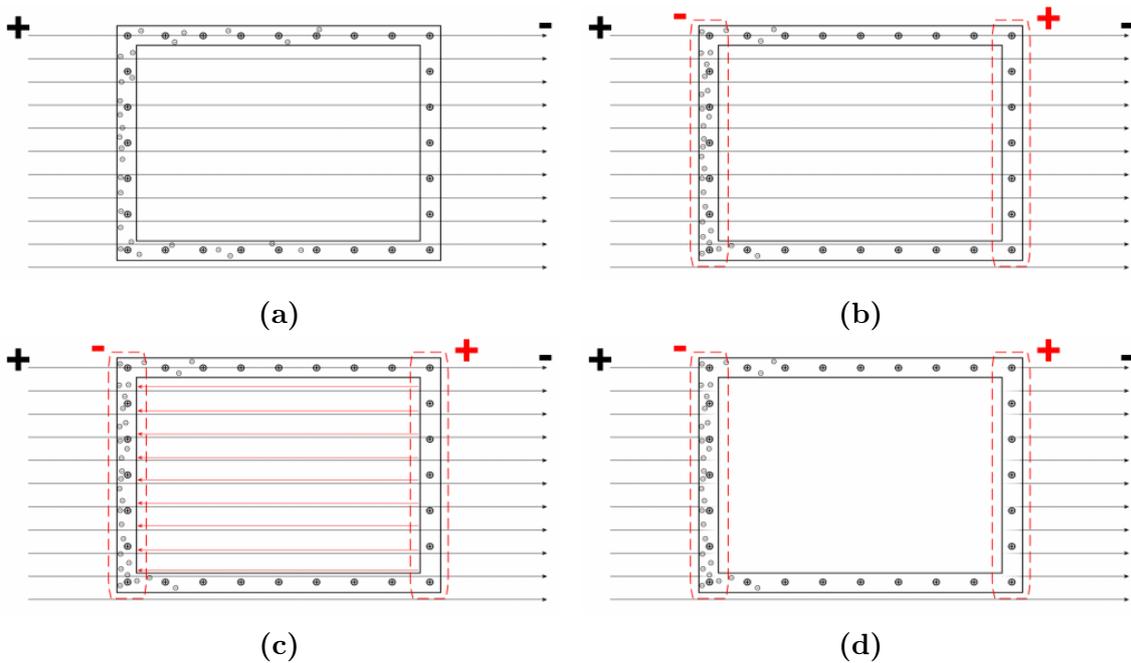


Figure 1.2: Illustration of the working principle of a Faraday cage. In **a**), an external field has just been applied and thus the electrons in the conductive cage are starting to move towards the positive pole of the field. This redistribution of charges creates a polarity between the left and right edge of the cage, as seen in **b**). This polarity, in turn, renders its own electric field in the opposite direction of the external field as displayed in **c**). These two opposing fields cancel out each other to a great extent, in theory (for a perfect conductor) completely. This is illustrated in **d**). Adapted from Skowron[5].

1.2 Faraday cage

To isolate electrical devices from their surroundings it is common to apply shielding, which aim to minimize disturbances from external electromagnetic fields by block-

ing the fields with enclosing conductive (or magnetic) barriers. Shielding can be realized with a Faraday cage, which is an enclosed mesh or continuous covering of a conductive material that to a very high degree cancels out external fields' within the interior of the cage. Mathematically, a Faraday cage may be described as a hollow perfect conductor[44]. Its operation is illustrated in Figure 1.2; an external field yields a response from the freely moving charge carriers in the conductor that, in turn, results in an electric field (in the opposite direction of the external field) between the edges of the cage that cancels out the external field within the cage.

Faraday cages are ubiquitous within contemporary technology and an invaluable tool for noise protection of weak analog signals. For instance, in addition to an insulating layer, the inner conducting wire of a coaxial cable is encapsulated inside a conducting layer to avoid external interference with the transmitted information.

1.3 Differential measurements and CMRR

An ECG corresponds to the potential difference between two points on the surface of the body. Since the two locations are fairly adjacent, they should be exposed to almost identical external electromagnetic fields (such as power line interference). By measuring the difference between the two electrodes, the external contributions should therefore to a great extent cancel each other out while the difference in the localized fields stemming from the relative positions to the heart should constitute the predominant output.

In order to further amplify the differences between the two signals and attenuate the similar elements, it is common to use a differential amplifier. A common measure of how well this accentuation of differences is accomplished is Common mode rejection ratio (CMRR). It can be defined as[45]

$$\text{CMRR} = 10 \log_{10} \left(\frac{A_{DM}}{A_{CM}} \right)^2 = 20 \log_{10} \left(\frac{A_{DM}}{|A_{CM}|} \right), \quad (1.3)$$

where A_{DM} is the amplification of the difference between the signals and A_{CM} the amplification of common elements in both signals. Evidently, a high value of A_{DM} and a small value of A_{CM} yields a high CMRR, i.e. a strong differential amplification. A CMRR of at least 100-120 dB is desirable for ECG measurements[45].

Chapter 2

Electrode design

This chapter describes the considerations that were taken into account when designing the non-contact electrode. Initially, the pros and cons with non-contact electrodes as compared to contact electrodes are discussed in section 2.1. A list of design priorities and desired properties are set in section 2.2. A simple non-contact electrode architecture based on previous work by [46] is described in section 2.3. The electrodes produced in this thesis are inspired by that design, with the same core component, the Burr-Brown INA116 instrumentation amplifier [47]. The transfer function and output noise profile is derived for this electrode. In 2.3.3 the issue of motion artifacts [48] and a solution is presented.

For non-contact probes, it is important to have a high bias impedance, high source capacitance and low bias capacitance as it reduces noise and sets the high-pass filter cutoff frequency, described in section 2.3.1 and 2.3.2. A previously used feedback circuit [49] to increase the input impedance, resistance bootstrapping, is described in section 2.4. A high input impedance can lead to large voltage offsets and long settling times. This phenomenon is treated in 2.5. Methods to reduce the bias capacitance, input neutralization and input guarding, is described in section 2.6 and 2.8. Shielding to reduce the effect of ambient noise and PLI (discussed in part I section 4.2) is described in 2.7.

All the findings from the initial, preparatory theory are eventually combined when the chapter is concluded with a summary of the electrode design procedure in section 2.9.

2.1 Contact or non-contact electrode?

The disposable Ag/AgCl contact electrode with hydro-gel interface is used universally for acquiring bio-potential signals and most of its properties are well understood and documented, e.g. [50] [51] [52]. It has for a long time proven to be a reliable and predictable interface for acquiring bio potential signals. However, there are some intrinsic issues with contact electrodes. Many electrodes in a near vicinity are prone

to cross-talk [53] and the gel dehydrates over time [50]. Also, there is a preparation procedure for each applied electrode. These issues are mitigated when using a non-contact electrode.

Also, the robustness with regards to motion artifacts and external electric fields have been shown to be the same or even better in non-contact electrodes than wet electrodes [50]. However, an issue with non-contact electrodes is that they tend to have an increased noise floor [54], especially at low body electrode coupling capacitance [55]. Since the fECG is often contaminated with electrical noise from amplifiers [6], it is important that noise in the electrode is minimized.

Non-contact electrodes offer reusability and a measurement device that easily could be integrated into a piece of clothing (that do not necessarily need to be in contact with the skin). This could have a positive effect on costs, environmental friendliness, and subject mobility during operation. In a long term perspective, it would also enable continuous health monitoring applications in home environments for both pregnant and non-pregnant individuals.

Taking all these aspects into account; for long term monitoring of minute signals and/or if creating a mesh of electrodes to improve signal quality after processing, non-contact electrodes would be a more realistic solution than wet electrodes.

2.2 Desired properties for non-contact electrodes

There are many requirements when designing electrodes, especially when trying to measure minute signals in a noisy and dynamic environment (see section 4.2). Most importantly, the electric components in the circuit should have a noise level that is much than the fECG, be resistant to power line interference and ambient noise, have a low-enough cutoff frequency and be resistant to motion artifacts. Furthermore, Non-contact electrodes must have a high input impedance to be able to detect low frequency signals arising from minute capacitive couplings. Previous coin sized electrodes have had coupling capacitances of around 10 pF through a shirt to multiple nF with contact through a good dielectric interface [48]. This resistance can lead to large offset voltages if the bias current of the front end amplifier is large (see section 2.5). In conclusion, the electrodes should be designed to fulfill the following requirements:

1. Be able measure the fECG with sufficient signal to noise ratio (SNR).
2. Have a desirable (low) offset voltage and be stable with respect to motion artifacts. This means that the electrode output should be resistant to variations in source impedance (i.e. in the skin-electrode coupling).
3. Be designed in a way such that it minimizes the influence from the environment, e.g. PLI and ambient noise.

2.3 Circuit analysis

A non-contact electrode schematic without feedback is shown in Figure 2.1. Note that the amplifier is an instrumentation amplifier and therefore does not necessarily need a feedback connection. The biopotential at the skin surface capacitively couples with the electrode surface through an insulating layer, i.e. dielectric coating and cloth, and is represented with an ideal capacitor C_{body} . The resistive element between electrode and skin is assumed to be very large compared to other circuit elements and is therefore neglected. Although, this approximation is discussed from a noise perspective in section This resistive element is treated from a noise perspective in section 2.3.2.1. Following the skin-electrode interface is a capacitor in series which limits the effect of variations in input capacitance. These two capacitors constitute the total source capacitance, $C_S = \frac{C_{body}C_{lim}}{C_{body}+C_{lim}}$.

Apart from the source capacitance, there are both a bias resistance and a bias capacitance to ground, set either explicitly or stemming from the nonideality of the amplifier, parasitic capacitance and noninfinite insulation resistances.

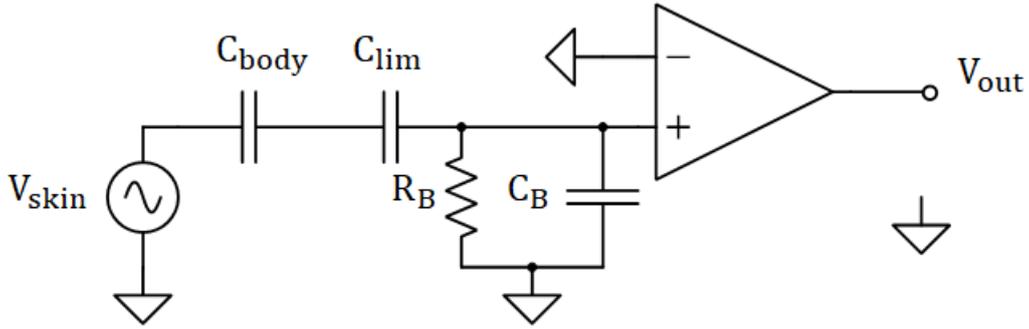


Figure 2.1: Circuit model of the proposed circuit, shown without external feedback circuitry such as bootstrapping and input neutralization. The amplifier is an instrumentation amplifier and hence has no gain feedback. e_{R_B} is the Johnson-Nyquist noise arising from the bias resistor to ground, i_n is the input referred current noise and e_n the input referred current noise of the amplifier. C_{body} is the capacitive coupling between body and electrode and C_{lim} is a capacitor that limits variation in source capacitance. C_B is the bias capacitance, which is the contribution between parasitic coupling capacitances between other parts of the circuit and the input capacitance of the amplifier.

2.3.1 Derivation of transfer function, maximal gain and cut-off frequency

A schematic illustration of the circuit is shown in Figure 2.2. Kirchoffs Voltage Law in the frequency domain on the circuit before the instrumentation amplifier

yields

$$V_{skin} = \frac{I}{j\omega C_S} + C_B || R_B I = I \left(\frac{1}{j\omega C_S} + \frac{R_B}{1 + j\omega C_B R_B} \right) \quad (2.1)$$

$$= I \left(\frac{1 + j\omega(C_S + C_B)R_B}{j\omega C_S(1 + j\omega C_B R_B)} \right), \quad (2.2)$$

where C_S is the combination of the limiting capacitance C_{lim} and the body electrode coupling C_{body} , $C_S = \frac{C_{bd}C_{lim}}{C_{bd} + C_{lim}}$. The current flowing through the bias resistor and capacitor is

$$I = \frac{V_{in}}{R_B || C_B} = V_{in} \frac{1 + j\omega C_B R_B}{R_B},$$

where V_{in} is the voltage on the input terminal. The output of the instrumentation amplifier is set by the the internal gain G . Hence, the output $V_{out} = GV_{in}$. This yields

$$I = \frac{V_{out}}{G} \frac{1 + j\omega C_B R_B}{R_B},$$

which in combination with (2.2) gives

$$V_{in} = V_{out} \frac{1 + j\omega C_S R_B}{G R_B} \frac{1 + 2j\omega C_S R_B}{j\omega C_S(1 + j\omega C_S R_B)} \Rightarrow \quad (2.3)$$

$$H = \frac{V_{out}}{V_{in}} = G \frac{j\omega C_S R_B}{1 + j\omega(C_S + C_B)R_B} \quad (2.4)$$

Maximal gain occurs when

$$\lim_{\omega \rightarrow \infty} |H| = \lim_{\omega \rightarrow \infty} \sqrt{\frac{(\omega C_S R_B)^2}{1 + (\omega(C_S + C_B)R_B)^2}} = G \frac{C_S}{C_S + C_B} \quad (2.5)$$

The cutoff frequency, f_c , is defined as the frequency where the square of the amplitude has decreased by a factor of $\frac{1}{2}$. Given the maximum amplitude in (2.5):

$$\begin{aligned} \frac{1}{\sqrt{2}} &= \left| \frac{j\omega C_S R_B}{1 + j\omega(C_S + C_B)R_B} \right| / \left(\frac{C_S}{C_S + C_B} \right) = \frac{j\omega(C_S + C_B)R_B}{\sqrt{1 + \omega^2(C_S + C_B)^2 R_B^2}} \Leftrightarrow \\ 2\omega^2(C_S + C_B)^2 R_B^2 &= 1 + \omega^2(C_S + C_B)^2 R_B^2 \Leftrightarrow \\ \omega &= \frac{1}{(C_S + C_B)R_B} \Rightarrow \\ f_c &= \frac{1}{2\pi(C_S + C_B)R_B}. \end{aligned} \quad (2.6)$$

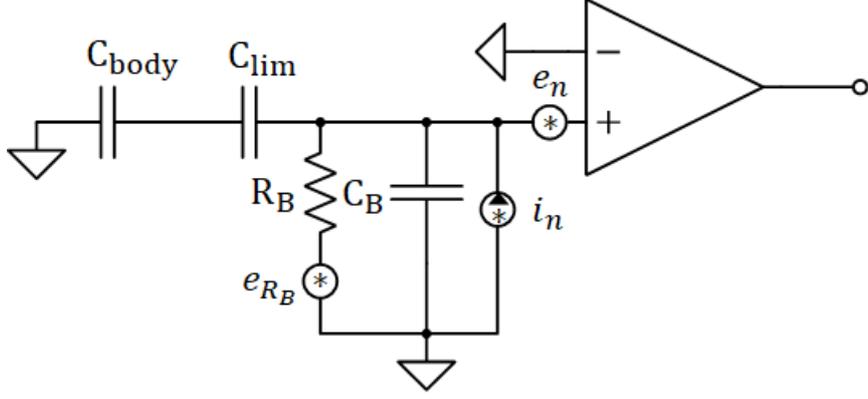


Figure 2.2: Noise sources in the proposed circuit, shown without external feedback circuits such as bootstrapping and input neutralization. The amplifier is an instrumentation amplifier and hence has no gain feedback. e_{R_B} is the Johnson-Nyquist noise arising from the bias resistor to ground, i_n is the input referred current noise and e_n the input referred current noise of the amplifier. C_{body} is the capacitive coupling between body and electrode and C_{lim} is a capacitor that limits variation in source capacitance. C_B is the bias capacitance, which is the contribution between parasitic coupling capacitances between other parts of the circuit and the input capacitance of the amplifier.

2.3.2 Derivation of output noise

The noise in an instrumentation amplifier is usually specified by three sources, voltage noise on the inputs e_{in} , current noise from the input to ground, i_{in} , and output noise e_o [56]. However, the INA116 noise equivalent does not include noise on the output [47]. These voltage sources are the square root of the power spectral density. On top of this, there's Johnson-Nyquist noise stemming from the bias resistance and given by $e_R^2 = 4\pi k_B T R_B$. An illustration of the noise sources in the proposed circuit is shown in Figure 2.2.

Beginning with the Johnson-Nyquist noise and setting all other noise sources to zero. C_B is in parallel with the source capacitance $C_S = C_{body} + C_{lim}$. Kirchoffs voltage law gives

$$e_{R_B} = IR_B + \frac{I}{j\omega(C_B + C_S)} = I \left(\frac{1 + j\omega(C_S + C_B)R_B}{j\omega(C_S + C_B)} \right) \Rightarrow$$

$$I = e_{R_B} \left(\frac{j\omega(C_S + C_B)}{1 + j\omega(C_S + C_B)R_B} \right).$$

The voltage on the input to ground, e_{in_v} , is

$$e_{in_v} = e_{R_B} - IR_B + e_n = e_{R_B} - e_{R_B} \left(\frac{j\omega(C_S + C_B)R_B}{1 + j\omega(C_S + C_B)R_B} \right) + e_n$$

$$= e_{R_B} \left(\frac{1}{1 + j\omega(C_S + C_B)R_B} \right) + e_n.$$

The same calculation for the current noise yields

$$e_{in_i} = i_n((C_S||C_B)||R_B) = i_n((C_S + C_B)||R_B) = \frac{i_n R_B}{1 + j\omega(C_S + C_B)R_B}. \quad (2.7)$$

The total voltage noise on the input is the sum of the two contributions:

$$e_{in} = e_{R_B} \left(\frac{1 + i_n R_B}{1 + j\omega(C_S + C_B)R_B} \right) + e_n.$$

Squaring the voltage sources and taking the time average, all cross terms cancel due to that the noise sources are uncorrelated:

$$\overline{e^2}(f) = \frac{4k_b T R_B + \overline{i_n^2} R_B^2}{1 + (C_S + C_B)^2 R_B^2 (2\pi f)^2} + \overline{e_n^2}. \quad (2.8)$$

From now to simplify notation, the mean of the power density of noise sources is defined through their root mean square values, $\hat{e} = \sqrt{\overline{e^2}}$, $\hat{i}_n = \sqrt{\overline{i_n^2}}$ and $\hat{e}_n = \sqrt{\overline{e_n^2}}$. The high-pass cutoff frequency is given by $1/(2\pi R_B(C_S + C_B))$, see (2.6), resulting in

$$\hat{e}^2(f) = \frac{4k_b T R_B + \hat{i}_n^2 R_B^2}{1 + \left(\frac{f}{f_c}\right)^2} + \hat{e}_n^2. \quad (2.9)$$

The main frequencies of interest for the fECG should be much higher than the desired cutoff frequency of the electrode. Hence, if $\frac{f}{f_c} \gg 1$ then

$$\begin{aligned} \hat{e}^2(f) &\approx \frac{4k_b T R_B + \hat{i}_n^2 R_B^2}{\left(\frac{f}{f_c}\right)^2} + \hat{e}_n^2 \\ &= \frac{4k_b T R_B + \hat{i}_n^2 R_B^2}{(C_S + C_B)^2 R_B^2 (2\pi f)^2} + \hat{e}_n^2 \\ &= \frac{1}{(C_S + C_B)^2 (2\pi f)^2} \left(\frac{4k_b T}{R_B} + \hat{i}_n^2 \right) + \hat{e}_n^2 \end{aligned} \quad (2.10)$$

2.3.2.1 The effect of finite source resistance

In previous calculations we have assumed that the source impedance is effectively infinite, however it is worth mentioning the effect if the source impedance is finite. To simplify the calculations the limiting capacitance is set to zero. The input noise density is then given by [57]:

$$v_{in}^2 = 4k_B T |Z_s||Z_B|^2 \left(\frac{1}{R_s} + \frac{1}{R_B} \right),$$

where R_s is the resistance between electrode and body, $Z_s = \frac{R_s}{1 + j\omega C_S R_s}$ is the corresponding source impedance, and $Z_B = \frac{R_B}{1 + j\omega C_B R_B}$ is the bias circuit impedance.

Rewriting the expression

$$\begin{aligned}
v_{in}^2 &= \left| \frac{R_s R_B}{(R_S + R_B) + j\omega R_s R_B (C_S + C_B)} \right|^2 \left(\frac{1}{R_s} + \frac{1}{R_B} \right) \\
&= \frac{R_s^2 R_B^2}{(R_S + R_B)^2 + \omega^2 R_s^2 R_B^2 (C_S + C_B)^2} \left(\frac{R_s + R_B}{R_s R_B} \right) \\
&= \left(\frac{R_s R_B}{R_S + R_B} \right) \frac{1}{1 + \omega^2 \left(\frac{R_s R_B}{R_s + R_B} \right)^2 (C_S + C_B)^2}
\end{aligned}$$

with $R_s \gg R_B$, $\frac{R_s R_B}{R_s + R_B} \Rightarrow R_B$ and the equation reduces to the first term in (2.9). If $R_s \ll R_B$, the source resistance instead dominates and the thermal noise is decided by this resistance. Previous authors have shown that this contribution can be dominating, e.g. the resistance through a cotton shirt can be around 300 M Ω [48] (depending on specific material, moisture level and thickness etc.). It is therefore important that the electrode surface is coated with a highly resistive dielectric.

2.3.2.2 Signal to noise ratio

The signal to noise ratio is given by the power of a signal divided by the power of the surrounding noise

$$SNR = \frac{P_{signal}}{P_{noise}}. \quad (2.11)$$

The signal frequency amplitude spectrum on the input of the amplifier is the transfer function H multiplied by the source frequency spectrum, V . The signal power is proportional to the amplitude squared. The signal power at the input is therefore

$$\tilde{V}_{in}^2 = \left| \tilde{V}_{skin} H \right|^2.$$

The transfer function to the input of the amplifier is simply the the transfer function to the output, equation (2.4), divided by the amplifier gain. Assuming that the frequency band is far away from the cutoff frequency, the transfer function is estimated to be at its maximal value, i.e. equation (2.5) divided by signal gain.

$$\tilde{V}_{in}^2 \approx \tilde{V}_{skin}^2 \left(\frac{C_S}{C_S + C_B} \right)^2.$$

The total signal power within a frequency bandwidth Δf is the integral of the power spectrum. Assuming that the latter is flat within the bandwidth, the total signal power becomes:

$$P_{signal} = \int_{\Delta f} \tilde{V}_{skin}^2 \left(\frac{C_S}{C_S + C_B} \right)^2 df = V_{skin}^2 \left(\frac{C_S}{C_S + C_B} \right) \Delta f.$$

The total noise power on the corresponding bandwidth is the integral of equation (2.10). Analytically, it is in this case more convenient to observe the noise to signal

ratio $NSR = \frac{1}{SNR}$ since each noise contribution is additive. The noise to signal ratio

$$NSR = \frac{1}{V_{skin}^2 \Delta f} \left(\frac{C_S + C_B}{C_S} \right)^2 \int_{\Delta f} \hat{e}_n^2 + \frac{1}{(C_S + C_B)^2 (2\pi f)^2} \left(\frac{4k_b T}{R_B} + \hat{i}_n^2 \right) df \quad (2.12)$$

$$= \frac{1}{V_{skin}^2 \Delta f} \int_{\Delta f} \left(\frac{C_S + C_B}{C_S} \right)^2 \hat{e}_n^2 + \frac{1}{C_S^2 (2\pi f)^2} \left(\frac{4k_b T}{R_B} + \hat{i}_n^2 \right) df. \quad (2.13)$$

For the noise to signal ratio to be minimized, each of the terms in the integral should be as small as possible. The contributions from the thermal noise and current noise reduces with C_S^2 , while the contribution from the input referred noise depends on the relationship between C_S and C_B , i.e. the gain to the input. Also, the current noise contribution reduces with R_B . An electrode with a high signal to noise ratio and therefore a low noise to signal ratio have a high source capacitance, a low bias capacitance and a high bias resistance. It is however very important to know which noise contribution that dominates. This is found by inserting tabulated/measured values for the three noise sources with actual values for C_S , C_B and R_S .

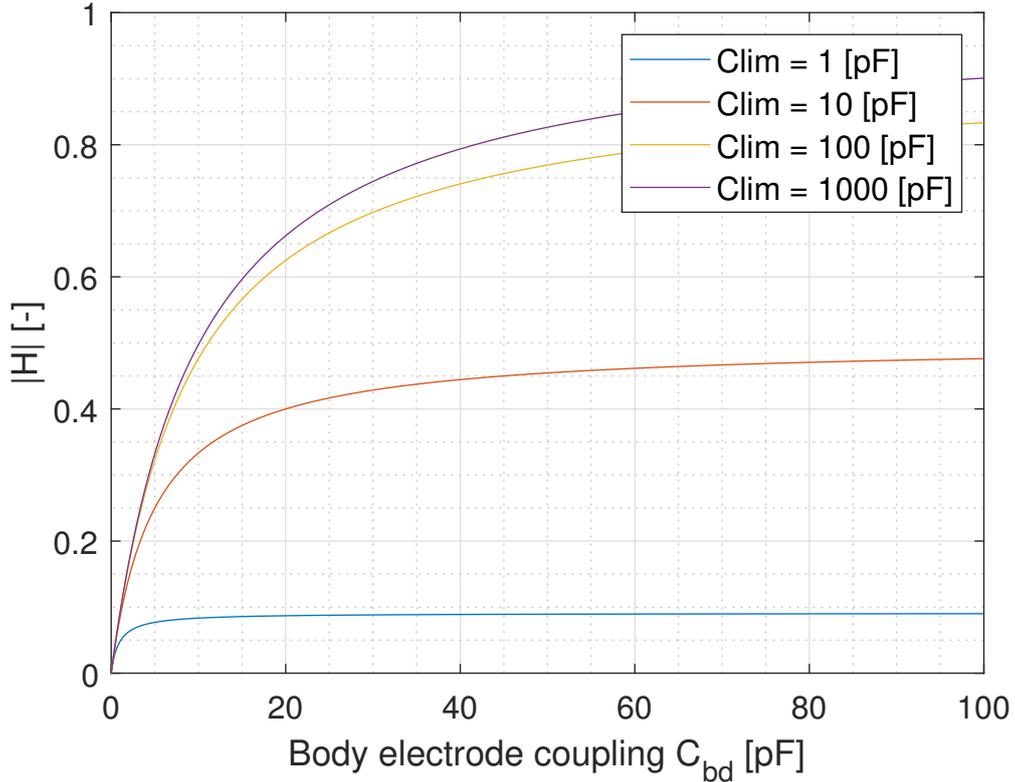


Figure 2.3: Amplification vs body-electrode coupling capacitance for different limiting capacitors C_{lim} .

2.3.3 Variations in body electrode coupling

An issue with previous non-contact electrodes has been the sensitivity to motion artifacts[48]. This is due to that even a modest change in skin-electrode distance may induce a significant change in the coupling capacitance between the electrode and the skin surface. To limit the influence of variations in this capacitance, a small capacitance C_{lim} was introduced. This capacitor sets an upper limit for the input capacitance C_S as $C_S = \frac{C_{lim}C_{body}}{C_{lim}+C_{body}}$. The amplitude modulation of the measured signal with respect to changes in electrode-body coupling is seen in the derivative of the maximal transfer function $|H|_{\omega=\infty} = C_S/(C_S + C_B)$ between the skin and the input of the instrumentation amplifier:

$$\frac{d|H|_{\omega=\infty}}{dC_{body}} = \frac{d|H|_{\omega=\infty}}{dC_S} \frac{dC_S}{dC_{body}} = \frac{C_B}{(C_S + C_B)^2} \left(\frac{C_{lim}}{C_{lim} + C_{body}} \right)^2 \quad (2.14)$$

If there was no C_{lim} , then $C_S = C_{body}$ and

$$\frac{d|H'|_{\omega=\infty}}{dC_{body}} = \frac{C_B}{(C_{body} + C_B)^2} \quad (2.15)$$

The ratio at which the derivative is reduced is

$$\frac{d|H'|_{\omega=\infty}}{dC_{body}} \bigg/ \frac{d|H|_{\omega=\infty}}{dC_{body}} = \quad (2.16)$$

$$\left(\frac{C_{lim} + C_{body}}{C_{lim}} \right)^2 = \left(1 + \frac{C_{body}}{C_{lim}} \right)^2 \quad (2.17)$$

In regions where $C_{body} > C_{lim}$ the derivative of the transfer function with respect to body electrode coupling is reduced and in regions where $C_{body} < C_{lim}$ it is increased. Hence, to minimize variations in source capacitance C_{lim} should be set as low as possible. However, this should be done with caution as it reduces C_S which in turn reduces the term signal to noise ratio, or equivalently, increases the noise to signal ratio, equation (2.13). Figure 2.3 displays the transfer function as a function of the body-electrode capacitance for various values of C_{lim} .

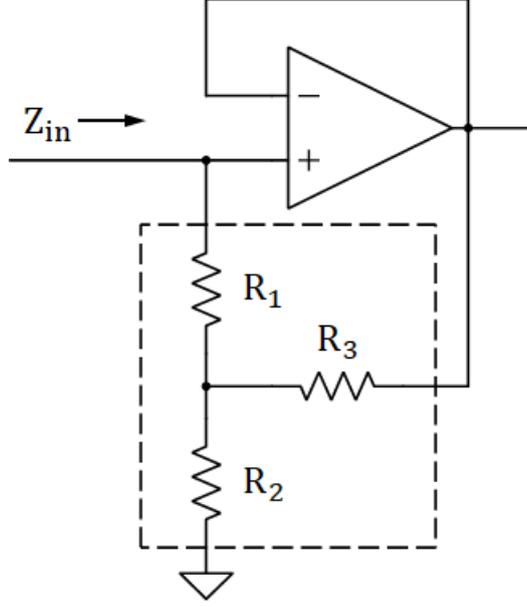


Figure 2.4: Circuit model showing input resistance bootstrapping network. The equivalent input impedance is $Z_{in} = R_{in} = R_1 + R_2 + \frac{R_1 R_2}{R_3}$

2.4 Bootstrapping

Bootstrapping is a method for increasing the resistance of the bias resistor, which ideally should be as high as possible for a low noise electrode, see (2.10). A bootstrapped bias resistance was tested and implemented by [49]. The bootstrapped circuitry is shown in Figure 2.4. The R_1 , R_2 and R_3 form an equivalent input resistance to ground $R_{eq} = R_1 + R_2 + \frac{R_1 R_2}{R_3}$. The analytic noise profile is similar to (2.10), with the only difference being that the bias resistance R_B is replaced with R_1 [49]

$$\begin{aligned} \hat{e}^2(f) &= \frac{4k_b T}{(C_S + C_B)^2 R_1 (2\pi f)^2} + \frac{\hat{i}_n^2}{(C_S + C_B)^2 (2\pi f)^2} + \hat{e}_{in}^2 \\ &= \frac{1}{(C_S + C_B)^2 (2\pi f)^2} \left(\frac{4k_b T}{R_1} + \hat{i}_n^2 \right) + \hat{e}_{in}^2 \end{aligned} \quad (2.18)$$

The R_1 resistance is much smaller than that of a non-bootstrapped circuit with a large bias resistor R_B with the same value as R_{eq} . This is the purpose of this type of bootstrapping, smaller resistors can be combined to create a huge equivalent resistance. However, since R_1 and not R_{eq} goes into (2.18), the noise contribution from the bias circuitry becomes much larger for a bootstrapped setup than for a bias resistor with the same resistance as R_{eq} . In conclusion, bootstrapping is a simple way to increase the effective resistance to ground, but will introduce more noise from the bias resistors. As long as large enough resistors are available, they should be preferable to an equivalent bootstrapping setup. That being said, whether or not the first term in (2.18) is dominating over the other terms is dependent on the

actual values involved.

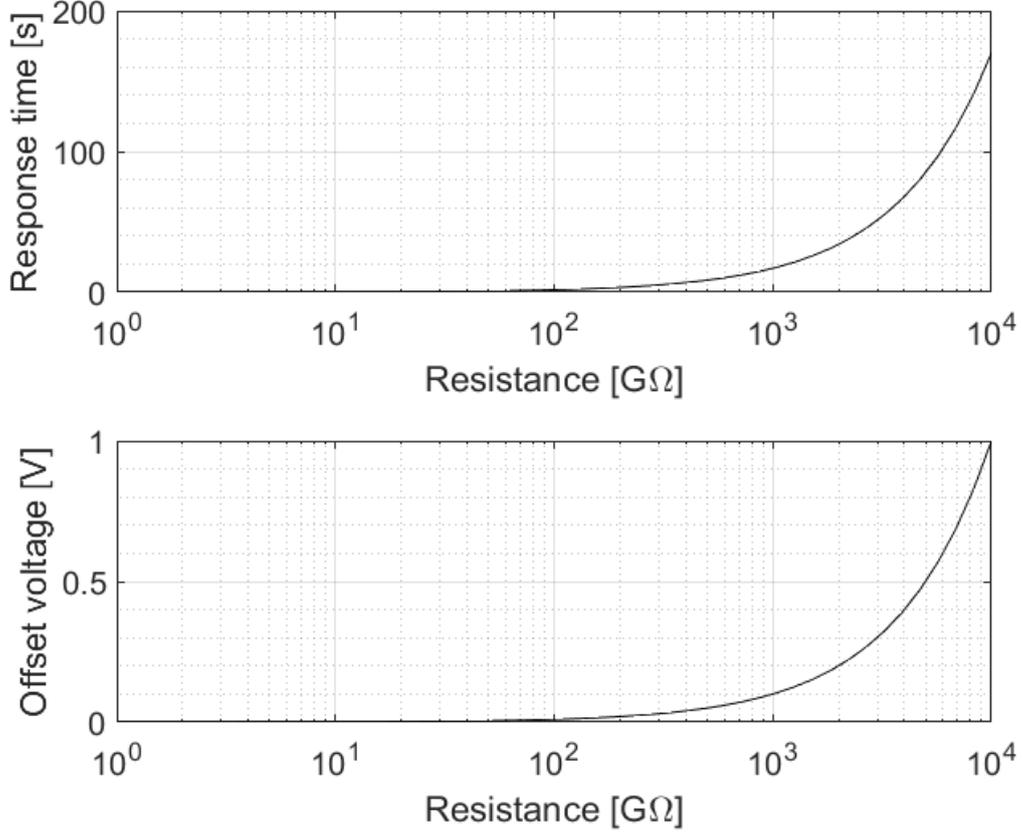


Figure 2.5: The top figure shows the decay time constant for a step response on the input of the amplifier as a function of resistance, see equation (2.5). The bias current $i_n = 100$ fA, the bias capacitance $C_B = 7$ pF and $C_S = 10$ pF. The time constant is directly proportional to $C_S + C_B$ and R_B . The bottom graph shows the equilibrium offset voltage on the input due to the bias current. The input offset is linear with respect to bias current and bias resistance. When the bias resistance is in the range of multiple TΩ, the response time becomes large and so does the input offset voltage.

2.5 Input offset voltage and discharge time

If there were no bias current path on the input terminal, charges would inevitably build up on the electrode input due to bias currents. One of the main features of INA116 is the low bias current. In the datasheet[47], the typical bias current is specified to 3 fA and the maximum bias current 100 fA. As in the previous calculations, assume that the source capacitance C_S is around 10 pF. The voltage over the input is the same expression as for the current source in the noise calculations:

$$e_{in} = \frac{i_B R_B}{1 + j\omega(C_S + C_B)R_B} = e_{i_B} \frac{1}{1 + j\omega(C_S + C_B)R_B}.$$

At $\omega = 0$ there will be a constant offset voltage $e_{i_B} = i_B R_B$. Such one-pole systems as the one in (2.5) are well-studied[58] and it can be shown that the corresponding transient response is

$$e_{in}(t) = e_{i_B}(0) \left(1 - e^{\frac{t}{(C_S + C_B)R_B}} \right),$$

where $\tau = (C_S + C_B)R_B$ is the response time. For $C_S = 10$ pF, $C_B = 7$ pF and $i_B = 100$ fA, the response time and offset voltage with respect to resistance is shown in Figure 2.5.

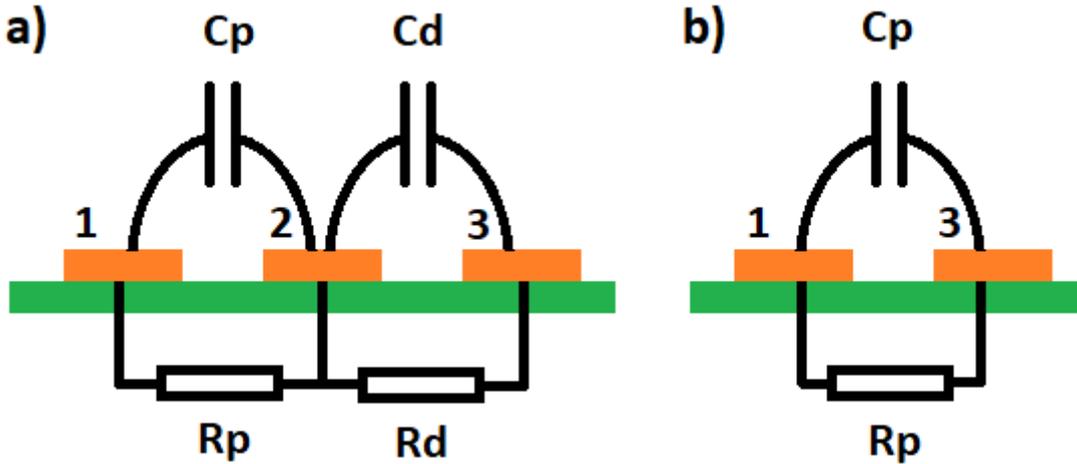


Figure 2.6: Figure showing the principle of driven guard. b) shows the sensitive trace, 3, and some other trace, 1, on the circuit board, e.g. ground. The sensitive trace couples with 1 through C_p and R_B . a) shows the driven guard technique. The sensitive input, 3, couples with the guard, 2, through C_d and R_d . The guard couples with some other trace, 1, through C_p and R_p . As the guard is actively driven to the same potential as the input, there is no leakage current between these two traces. Since the guard is not sensitive, it is unaffected by the coupling with trace 1.

2.6 Driven guard and stray capacitance

Low current circuits can sometimes be sensitive to leakage currents in the PCB, printed circuit board, material. Driven guard is a technique used to minimize these leakage currents. A unit gain output of the amplifier is directed in a loop around the sensitive input, the principle is illustrated in Figure 2.6. In this way there will be no potential difference between the input and the surrounding guard loop, and if designed correctly, a very high barrier to other parts of the board. The guard ring itself is not sensitive to leakage currents as the output of the amplifier can draw large currents. A well used technique [59] for guarding the input of a high impedance electrode is to use a three layer board with a guard ring around the input on the top layer, a middle layer completely filled with a copper guard plate and a guard ring around the electrode surface. By using the middle plate and the bottom plate, the large electrode surface does not couple with any other circuitry on the board

nor other electronics.

However, the small input pin on the top of the electrode is not isolated from the environment. The influence of the environment on the input is reduced through electrostatic shielding.

2.7 Electrostatic shielding

Shielding is a technique used to reduce external electromagnetic field in an enclosed space. Electrostatic shielding builds on the same principle as the Faraday cage, see Part II Section 1.2. The sensitive circuitry and signal carrying wires are surrounded with a metal sheet that neutralizes external electrostatic fields. Furthermore, shielded cables are used as well as a shielded enclosure around the backside of the electrode.

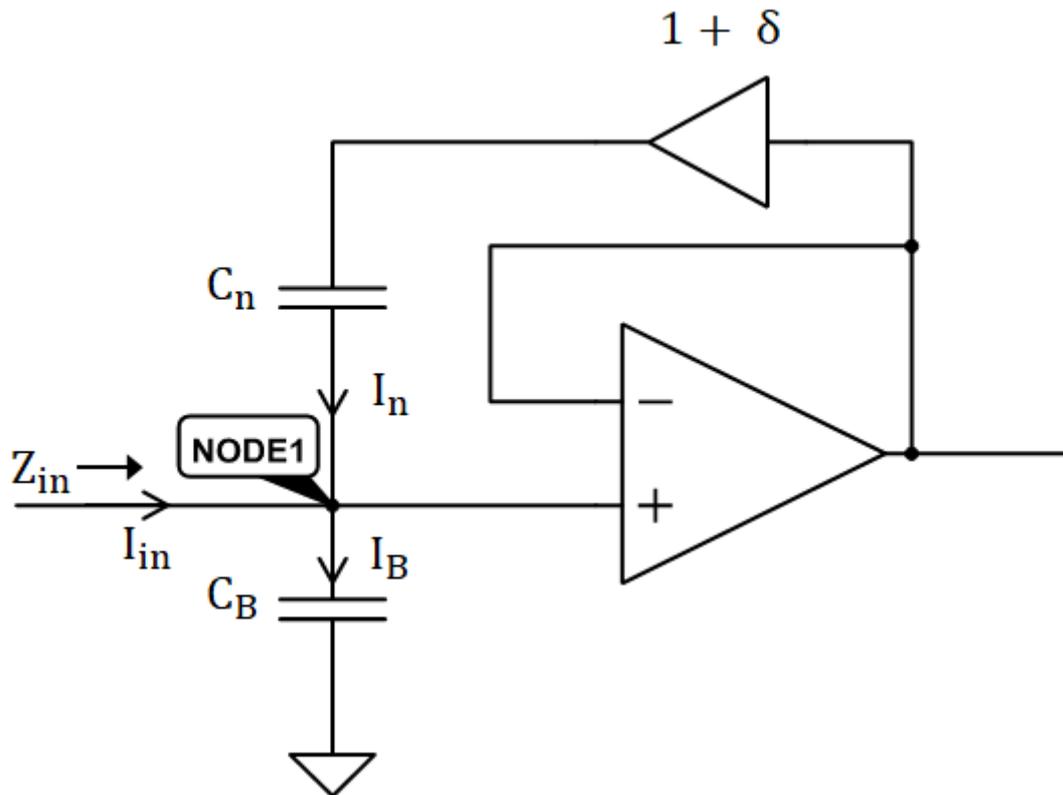


Figure 2.7: Figure showing the neutralization technique of input capacitance. The output is amplified by $1 + \delta$ fed back to the input through the coupling capacitance C_n . With the correct choice of C_n and δ ($C_n = C_B/\delta$) the current to the parasitic capacitance is drawn from C_n , hence canceling the effect of C_B . In this way the capacitive part of the input impedance Z_{in} is removed, or in practise, reduced.

2.8 Neutralization of input capacitance

Guarding can reduce much of the parasitic capacitance. However, there will always be a capacitance associated with the input of the an amplifier. For example, INA116 has an input capacitance of 7 pF[47]. A technique used to reduce the effect of the input capacitance is neutralization. A small portion of the output is amplified by a gain $1 + \delta$ and fed to the input to a capacitor C_n . By setting both C_n and δ appropriately, the current to the parasitic capacitance is drawn from C_n , hence canceling the effect of C_B .

Denote the output V_o . For a the configuration in figure 2.7 $V_i \approx V_o$ (This is a so called voltage follower configuration). The current over C_n is $I_n = \frac{((1+\delta)V_i - V_i)}{j\omega C_n} = \frac{\delta V_i}{j\omega C_n}$. The current over C_B is $I_B = \frac{V_i}{j\omega C_B}$. Looking at the input node, node 1 in Figure 2.7, I_B is effectively removed if

$$I_{in} = I_B - I_n = 0 \Rightarrow C_n = \frac{C_B}{\delta}. \quad (2.19)$$

This requires the input impedance observed from the input to be

$$Z_{in} = \frac{V_i}{I_{in}} = j\omega(C_B - \delta C_n). \quad (2.20)$$

Nevertheless, there are drawbacks with this technique. The feedback loop introduces noise to the circuit. Denote the voltage noise on the input of the voltage follower amplifier e_{1n} and the corresponding noise on the feedback amplifier e_{2n} . The output noise e_o of the amplifier is given by[55]

$$e_o^2 = e_{1n}^2 \left(1 + \left(1 + \frac{1}{\delta}\right) \frac{C_{in}}{C_S}\right)^2 + e_{2n}^2 \left(\left(1 + \frac{1}{\delta}\right) \frac{C_{in}}{C_S}\right)^2 \quad (2.21)$$

A small δ leads to undesired noise through feedback in both amplifiers. On the contrary, an excessively large δ demands very small neutralization capacitors.

2.9 Design procedure

To make design decisions based on the noise analysis, it is important to know which type of noise that is dominating. Equation (2.10) shows that R_B reduces the thermal noise on the input, and C_S and C_B reduces both thermal and current noise, however C_B reduces the signal to noise ratio. For INA116, the current noise is exceptionally low, a typical value of $0.1 \text{ fA}/\sqrt{\text{Hz}}$ at 1000 Hz[47]. Even though this current is small, it will be comparable to the input referred noise because of the small source and bias capacitors. From a noise perspective, the design proceeds as following:

1. Choose a reasonably large gain in the instrumentation amplifier. Gain reduces the input referred noise [47].
2. Choose a reasonable bandwidth to minimize the noise on, i.e. the same as the frequency spectrum of the fECG.

3. Pick a bias resistance such that the thermal noise is in the same range as the current noise on the desired bandwidth.
4. Find the source capacitance that makes the current and thermal noise in the same range as the input referred noise on the desired bandwidth.

The input noise reduces drastically with gain for small amplifier gain. A table of input referred noise at 1, 10 and 100 Hz for gain 1, 10 and 1000 for INA116 is shown below. As can be seen, the noise reduces sharply from gain 1 to gain

Gain [-]\Frequency [Hz]	1	10	100
1	2000	450	195
10	500	140	56
1000	330	90	38

Table 2.1: Input referred noise in INA116 at different amplifier gain and frequencies. The noise levels are given in $\text{nV}/\sqrt{\text{Hz}}$

10, but not much from gain 10 to gain 1000. The gain should be in the range of 10-1000. Exactly which gain that is used depends on experimental validation of the robustness and offset voltage of the electrode.

The complete fECG contains frequencies in the range from DC to around 300 Hz[20], however most of the spectrum is contained within 15 and 70 Hz as can be seen in Figure 3.2.

Looking at (2.10), the noise contribution from the current source exceeds the thermal noise contribution when $\frac{4k_bT}{R_B} < i_n^2$. Using the specified typical current noise at 1000 Hz, $i_n = 0.1 \text{ fA}/\sqrt{\text{Hz}}$, the bias resistor that would be needed for the current noise to be larger than the thermal noise is approximately 1.6 T Ω . For FET type input amplifiers current noise is generally white for low frequencies and has a linear slope at a certain cutoff frequency [60]. Hence, the result is valid for frequencies lower than 1000 Hz.

To attain this bias resistance a bootstrapping feedback network is not feasible as the purpose of the network is to reduce the need for large resistors, and the corresponding resistance in the noise calculations (R_1 in equation (2.18)) would increase the thermal noise to a level several times greater than the current noise. It is noteworthy that attempts have been made to create extremely large input impedances by using bootstrapped reverse diodes [61] or not setting the bias resistance explicitly at all [62], leaving the bias current path through the PCB, skin-electrode interface and amplifier. However, for the purposes of this thesis, this only serves to reduce the overall reliability of the electrode as the settling time and offset voltage would become large, thereby indirectly reducing noise performance through a lower amplifier gain. With this in consideration, a bias resistor with a resistance of 500 G Ω was implemented without bootstrapping. With this resistance, the input voltage offset would be around 0.05 V (Figure 2.5), allowing for maximum input gain in the range of ~ 200 if the INA116 is operated with the specified typical voltage supply, ± 15 V.

The source capacitance needed to reduce the current and voltage noise to levels of the input referred noise depends on whether a neutralization network is implemented or not. As described in section 2.8, the coupling capacitance must be relatively small not to introduce excessive noise. This is in general not desired as the coupling capacitance between PCB traces may be in the range of the neutralization capacitance. However, either a strict manufacturing process and well designed board layout, or manually tuning the feedback gain can resolve this issue. A neutralization capacitor in a non-contact electrode has successfully been implemented by Spinelli[49], although the authors used a bootstrap technique that increased the voltage noise, see Part II Section 2.4). It has also been suggested that Harland and Prance used a neutralization technique in their moderately secretive publication in 2002[62], [48]. Input neutralization is a must when dealing with very small coupling capacitances as the gain drastically reduces as the bias capacitance approaches the source capacitance, see equation (2.5).

A neutralization technique was not implemented in the current electrode due to two reasons. Firstly, the body-electrode capacitance is most of the time large with respect to the input capacitance of the INA116, 10 pF to several nF compared to 7 pF. Secondly, setting the feedback gain of the feedback amplifier while adjusting the gain of the INA116 for optimal performance would result in more experimental error sources, which was deemed highly undesirable during the early evaluation stages discussed in this report. However, the technique can be implemented in a next generation electrode to marginally increase the electrode performance with respect to stability to variations in coupling capacitance and input referred voltage noise levels.

That being said, assuming that the driven guard is effective, the bias capacitance is 7 pF from the amplifier. The source capacitance that would reduce the current and thermal noise contributions to the same level as the input referred noise is found through equation (2.10):

$$\frac{1}{(C_S + C_B)^2(2\pi f)^2} \left(\frac{4k_b T}{R_B} + i_n^2 \right) = \hat{e}_n^2 \Leftrightarrow \quad (2.22)$$

$$(C_S + C_B) = \frac{1}{\hat{e}_n(2\pi f)} \sqrt{\frac{4k_b T}{R_B} + i_n^2} \Leftrightarrow \quad (2.23)$$

$$C_S = \frac{1}{\hat{e}_n(2\pi f)} \sqrt{\frac{4k_b T}{R_B} + i_n^2} - C_B. \quad (2.24)$$

The thermal noise and current noise contributions decrease faster with frequency than the input referred noise [47]. Hence, the thermal and current noise should preferably be lower than the input referred noise at 15 Hz. Using gain 10, the corresponding source capacitance is given by inserting the value into (2.24), which results in $C_S = 8.6425$ pF. Knowing this, the limiting capacitance C_{lim} was set to 10 pF as it should be as low as possible to reduce the effect of variations in input capacitance. However, as described in 2.3.2.2, a reduced ratio between source and bias capacitance increases the term in the noise to signal arising from the input referred noise, equation (2.11). An upper bound of the increase in NSR can be set

by assuming that the other noise sources are negligible. The NSR then becomes:

$$NSR = \frac{1}{V_{skin}^2 \Delta f} \int_{\Delta f} \left(\frac{C_S + C_B}{C_S} \right)^2 \hat{e}_n^2 = \frac{1}{V_{skin}^2} \left(\frac{C_S + C_B}{C_S} \right)^2 \hat{e}_n^2. \quad (2.25)$$

If the body-electrode coupling is 50 pF and the bias capacitance is 7 pF, the upper limit is equation 2.25 evaluated at $C_S = \frac{C_{body}C_{lim}}{C_{body}+C_{lim}}$ divided by the same equation evaluated at $C_S = C_{body} = 50$ pF. At the specified parameters the limiting capacitor leads to a in an increase NSR by a factor 1.61 or equivalently, a reduction of the SNR by a factor 0.619. This was deemed acceptable with regards to the increase in robustness with respect to variations in source capacitance.

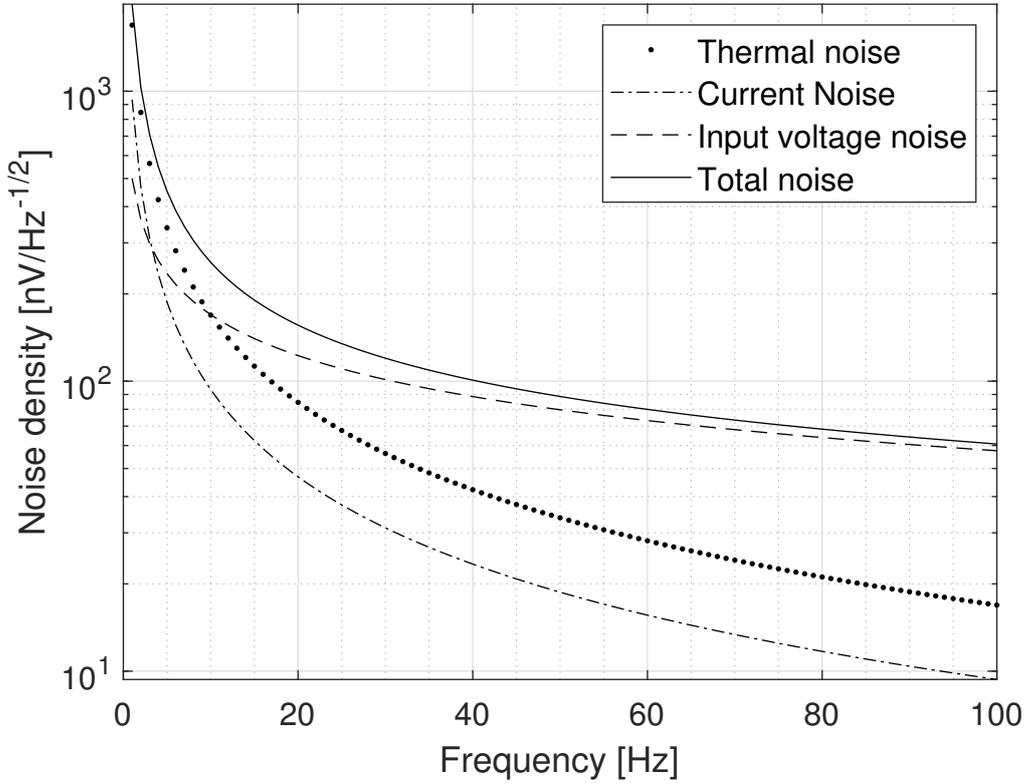


Figure 2.8: Noise power spectral density for the desired configuration of the electrode. As can be seen the thermal noise and current noise is lower than the input referred noise on the interval 15-75 Hz. On this bandwidth the total noise is $0.9242 \mu V_{rms}$.

The noise power density spectrum with the set parameters and $C_S = 10$ pF is shown in Figure 2.8. The total noise in the bandwidth 15-75 Hz is $0.9242 \mu V_{rms}$.

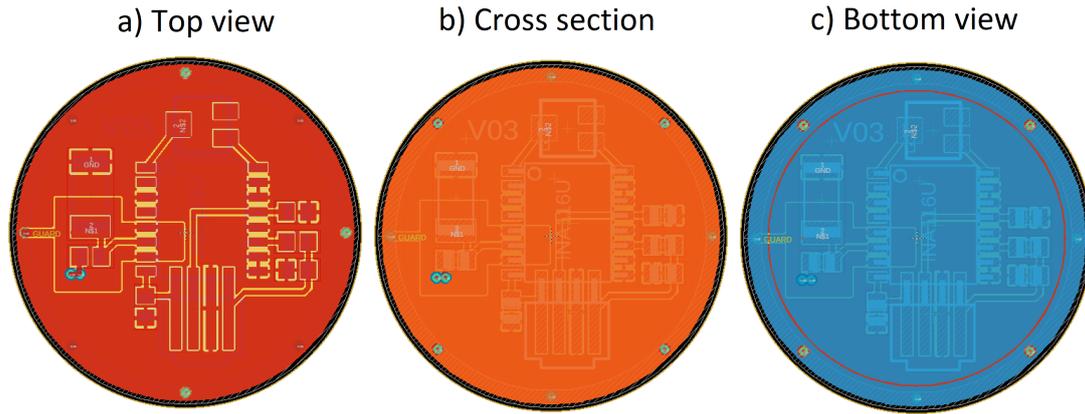


Figure 2.9: Design of the electrodes in EAGLE PCB. a) shows the top copper layer of the electrode, b) shows the inner copper guard layer and c) shows the electrode surface with a surrounding guard ring.

The three layer guarding technique was implemented as described in Section 2.6. The three layers of the board is shown in Figure 2.9. The guard ring on the top layer can be seen on the left side of the top view. The cross section shows the entire guard layer in the middle of the electrode. The bottom view shows the electrode surface as well as the surrounding guard ring. The four vertical stripes by the lower edge on the top view are pads for output, ground, and power supply external connections going to and from the electrode. Shielded cables were used and a shielded conductive casing was placed around the electrode in accordance with Section 2.7.

Chapter 3

Evaluation of electrode

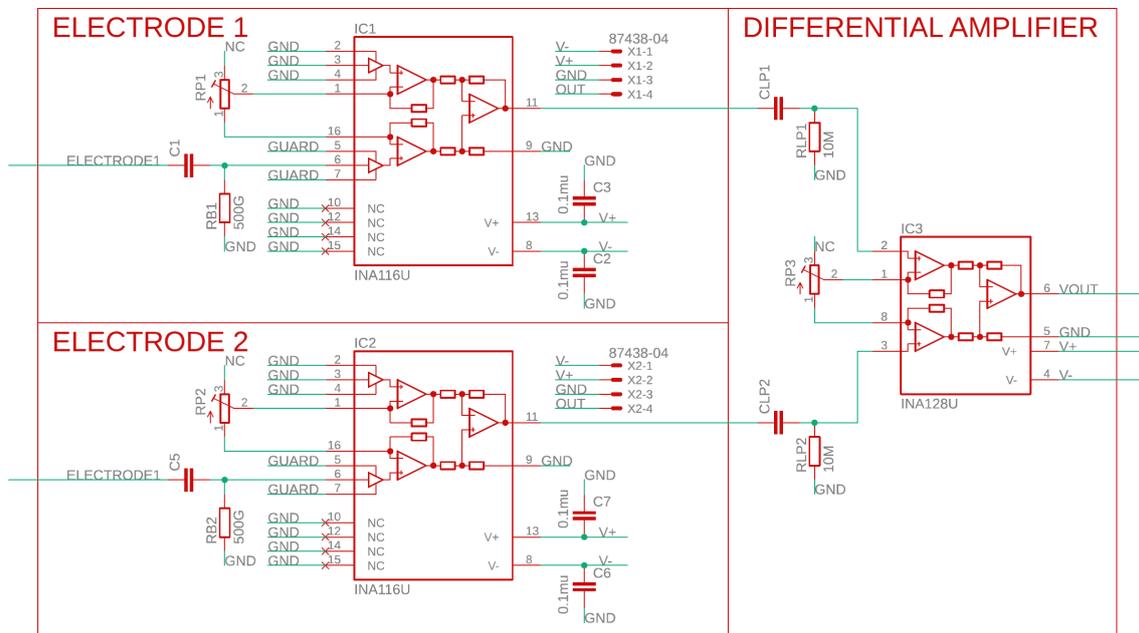


Figure 3.1: The ECG measurement circuit. On the left side are circuit diagrams for the electrodes described in section 2.9. The right side shows the amplifier circuit. Both signals were passed through a high pass filter and then to INA128, a high CMMR instrumentation amplifier, before a differential signal was logged.

3.1 ECG measurement setup

The ECG measurement circuit is shown in Figure 3.1. Both of the outputs from the electrodes were propagated through a high pass filter that removed the DC components. The signals were then amplified by an INA128 [63], a high CMRR instrumentation amplifier. Gain of both electrodes and differential amplifier were varied to find optimal performance. No notch filter for the PLI was implemented as this distorts the ECG waveform. The signal from INA128 was acquired using the data acquisition device DT9837A connected to a PC.

3.2 Transfer function measurement setup

The electrode was mounted inside a closed metal box. An 820 nF ceramic capacitor was connected to the input of the electrode. The electrode input was supplied with sinusoidal signals from the digital to analog converter DT9837A through BNC connectors in the metal box. The frequency of the driving signal was increased in logarithmic steps from 0.05 to 1000 Hz. Both driving signal and output signal were measured using the DT9837A. Exactly ten periods of the driving signals was used to estimate the Fourier transform using MatlabTM. The maximum amplitude in the frequency domain of the driving signal and the measured signal respectively were used to calculate the amplitude and phase of the transfer function at each frequency step. The amplitude of the transfer function at 1000 Hz was used to estimate the bias capacitance C_B through equation (2.5), which in combination with equation (2.6) and measurement of the cutoff frequency was used to estimate the bias resistance. The bias capacitance was compared to the capacitance measured on the board with no components on to verify that the driven guard was working as expected.

Also, a measurement was conducted with the electrode surface (i.e. the insulating solder mask made of photosensitive ink) connected to ground by conductive copper tape. The amplitude of the transfer function at the cutoff frequency was compared with and without copper tape to verify that the body electrode resistance was much larger than than the bias resistance.

3.3 Noise measurements setup

The electrodes were connected in a shielded box with BNC connectors. Supply voltage for the electrodes was supplied through the BNC connections in the box. As the sensor is very sensitive to movement and vibrations, measurements varied a lot depending on where the shielded box was situated and on what type of equipment that surrounded it. The shielded box was placed on a stack of thick, vibration attenuating foam cuboids, to reduce the influence of external vibrations. The output of the electrode was connected to a BNC connector and a BNC cable was connected to the input on the box and to a lock-in amplifier, SR510. The lock-in amplifier was supplied with a reference signal from a signal generator, Agilent 33220A. Noise measurement was selected on the amplifier, the dynamic range was set to low, the bandwidth was set to 1 Hz and the resolution was set depending on the magnitude of the noise levels. The outgoing signal from the shielded box was pre-amplified ten times before the lock-in amplifier measured the noise signal. This to enable accurate detection of noise signals of smaller magnitude. All amplifications were accounted for in the noise measurements. The measurements show the noise levels at the input of the INA116.

The analog output of the amplifier was connected to a 16 bit analog-to-digital converter, NI PCI-6014. Signals were recorded for on average 15 minutes. Longer measurements were made (up to 11h over night) to make sure that the signal did

not deviate over time. Before each measurement a delay of one or a few minutes was introduced so that the averaging method in the lock-in amplifier had settled. Measurements were made with various amplifier (INA116) gain and at various frequencies.

Also, the lock-in amplifier was tested on various resistors and short circuiting to make sure that the amplifier was working as expected.

Chapter 4

Results

4.1 Transfer function - gain and cutoff frequencies

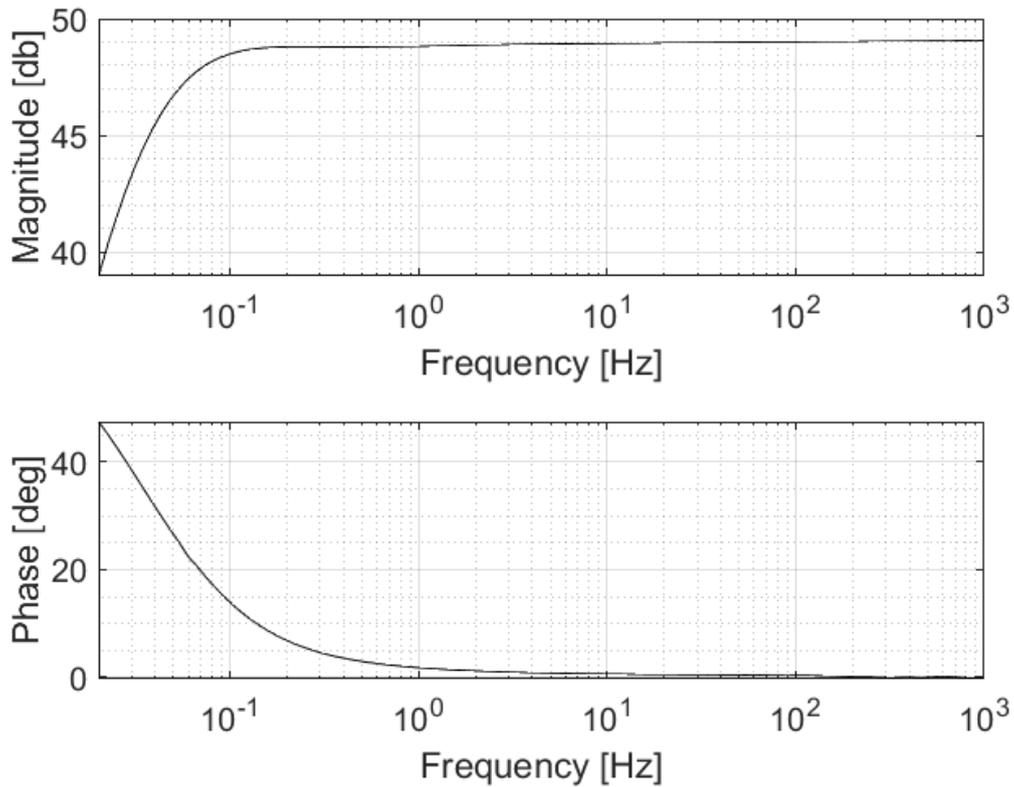


Figure 4.1: Figure showing the measured transfer function of the non-contact electrode at $C_S = 10$ pF. The upper figure shows the amplitude modulation and the lower figure shows the phase attenuation. The high pass cutoff frequency was measured to be 0.0268 Hz which is well below 0.05, the standard for ECG measurements [6]. From the analytic transfer function, equation (2.4), the bias capacitance C_B was calculated to 7.01 pF and the bias resistance R_B to 349 G Ω .

The measured transfer function for the electrode is shown in Figure (4.1). The gain resistor was set to 2663Ω which corresponds to gain 19.78. The maximum value of the transfer function, i.e. at 1000 Hz, was $|H|_{max} = 11.626$. This is well above the cutoff frequency, which implies that the amplification is almost at a maximum, equation (2.5). Using equation (2.5) with $C_{body} = 820 \text{ nF}$ and $C_{lim} = 10 \text{ pF}$, the bias capacitance was determined to $C_B = 7.01 \text{ pF}$. Without any components on the board, the capacitance was 5 pF from input to ground and 35 pF from the electrode surface to ground.

The measured cutoff frequency was 0.0268 Hz , well below standards for ECG measurements. This is slightly above the calculated cutoff frequency, 0.0187 Hz . The bias resistance was not measured due to difficulty finding a valid reference resistor. Using equation (2.6) with $C_{body} = 820 \text{ nF}$, $C_{lim} = 10 \text{ pF}$ and $C_B = 7.01 \text{ pF}$, the bias resistance was calculated to $349 \text{ G}\Omega$. This is less than the specified resistance of the bias resistor, $500 \pm 100 \text{ G}\Omega$. At 0.0268 Hz , the amplitude of the transfer function was 8.67 with no connection between electrode surface and ground (air) and 8.65 with a direct conducting path to ground.

4.2 Noise measurements

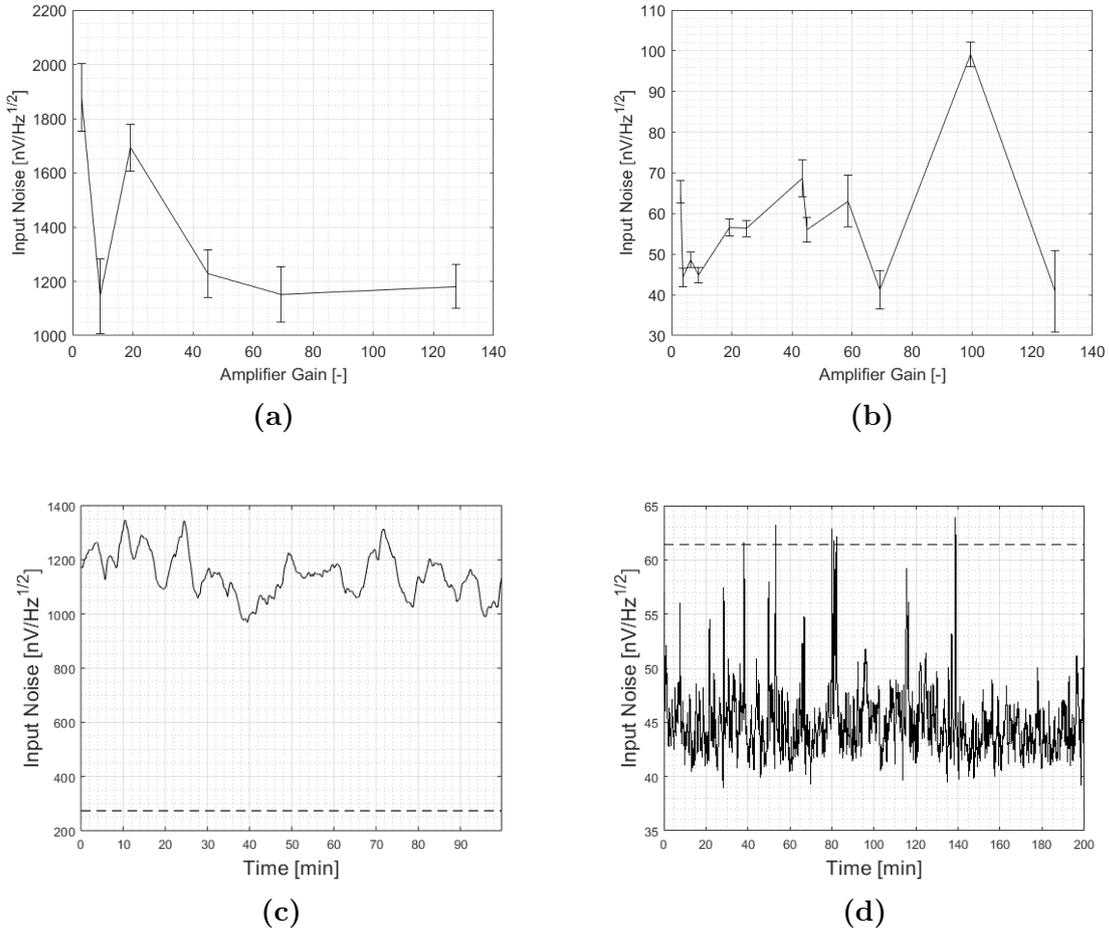


Figure 4.2: **a)** shows measurements of the mean of input noise density measured at 10 Hz for different amplifier gains. The error bars show $\pm 1\sigma$. **b)** shows measurements of the mean of input noise density measured at 100 Hz for different amplifier gains. The solid line in **c)** shows the average noise over time for a measurement made at 10 Hz and gain 9.81. The dashed line shows the corresponding expected noise level, equation (2.10). The solid line in **d)** shows the average noise over time for a measurement made at 10 Hz and gain 8.92. The dashed line shows the corresponding expected noise level, equation (2.10). Note that there are large fluctuations in the noise at 100 Hz, **b)**, and that there is no obvious trend with respect to gain, **b)**. At this frequency, it is expected that the noise decreases with gain until the contributions from the thermal and current noise are larger than the input referred noise, see Figure (2.8). In the same way there is no trend in average noise with respect to gain. Finally the noise levels in **d)** are in the vicinity of the estimated noise and the noise seems to reach a noise floor at some time intervals.

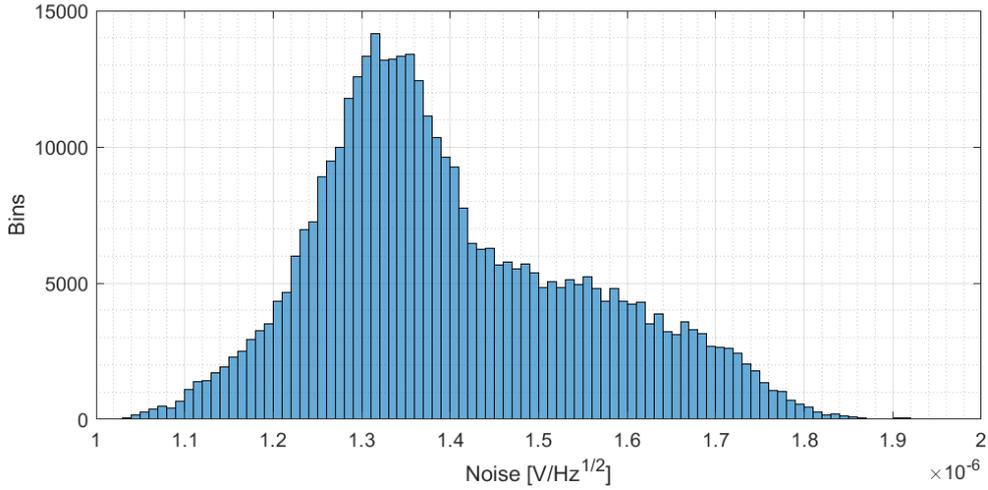


Figure 4.3: Histogram of the average noise over time at 100Hz and gain 24.92 measured with the lock-in amplifier. The measurements were taken over 11 h over night. The average noise distribution is not Gaussian and has a large standard deviation.

Noise measurements at 10Hz and 100Hz with different amplifier gain is shown in Figure 4.2a and 4.2b respectively. The error bars indicate $\pm 1\sigma$. According to the noise calculations in Section 2.9, input referred voltage and total noise density should decrease with gain, especially at 100 Hz, see Figure (2.8) and Table 2.1. At 10 Hz, the noise should decrease with gain until the thermal and current noise dominate over the input referred noise. As seen from Figure 4.2b, there is no clear trend for noise at 100 Hz, although it is in the expected range. At gain 10 and 100 Hz the input noise should be $61.432 \text{ nV}/\sqrt{\text{Hz}}$, equation (2.10). At gain 8.92 the noise is in the range 42-47 $\text{nV}/\sqrt{\text{Hz}}$. However, at 10 Hz the measured noise is markedly higher than what is expected and shows no clear trend. At this frequency, the noise is in the range 1000-1300 $\text{nV}/\sqrt{\text{Hz}}$ while it should be around 273.28 $\text{nV}/\sqrt{\text{Hz}}$. Looking at a measurement over time, Figure 4.2c, gain 9.81 at 10 Hz, it can be seen that the measurements drift over time. The dashed line indicate the expected noise level. From a long measurement at 100 Hz and gain 8.92, Figure 4.2d, it can be seen that drift still exists, but looking at the relative changes the measurement is more stable. A histogram of the measurement is shown in Figure 4.3. As can be seen, the distribution is not Gaussian.

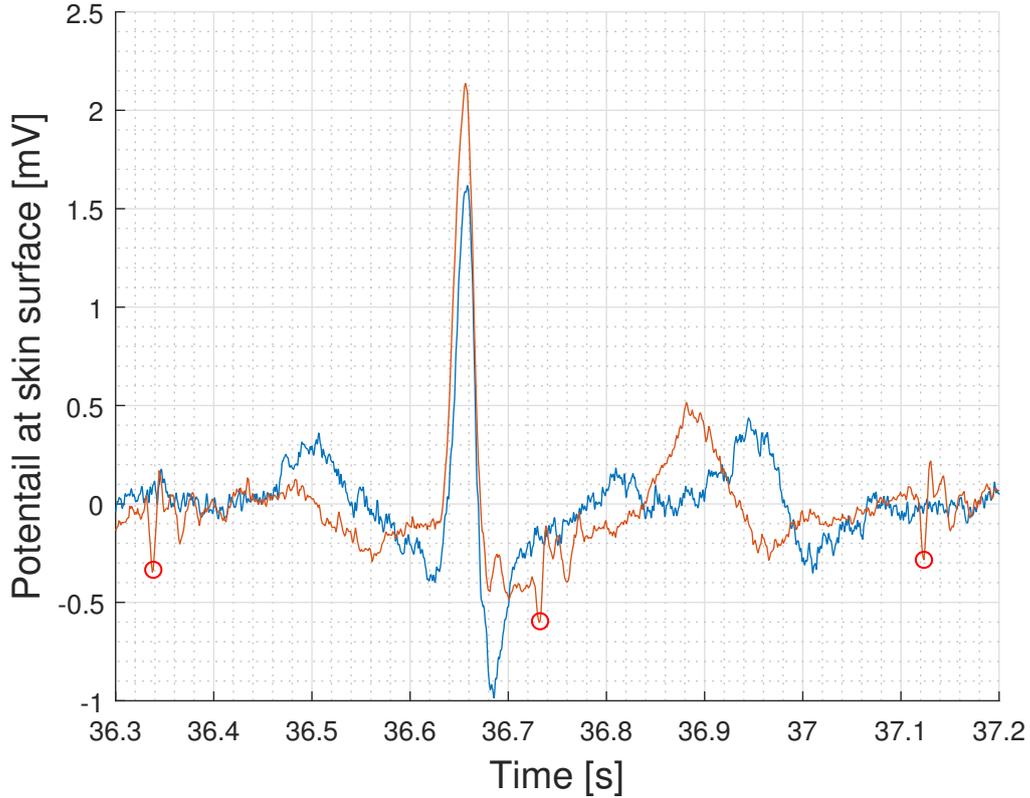


Figure 4.4: Typical ECG of abdominal non-contact electrode measurement (blue) conducted with the setup developed in this thesis and abdominal contact electrode from the NI-FECG database, see Part III Section 2.1. The red markers denote fQRS locations. The noise and interference from the non-contact electrode is similar to the contact electrode used for measuring the fECG.

4.3 ECG measurements

Figure 4.4 shows a typical abdominal non-contact electrode measurement conducted with the setup developed in this thesis and a typical contact electrode-based measurement taken from the NI-FECG database, see Part III Section 2.1. Both signals are compensated for gain, i.e. the graph shows the potential at the skin surface. The gain of INA116 was set to 20 and the gain of INA128 was set to 100. The red markers indicate fQRS locations. A digital low pass filter with 2 Hz cutoff and a digital high pass filter with 100 Hz cutoff have been applied to both signals. Note that the non-contact electrode signal is similar with respect to noise and interference to the contact electrode.

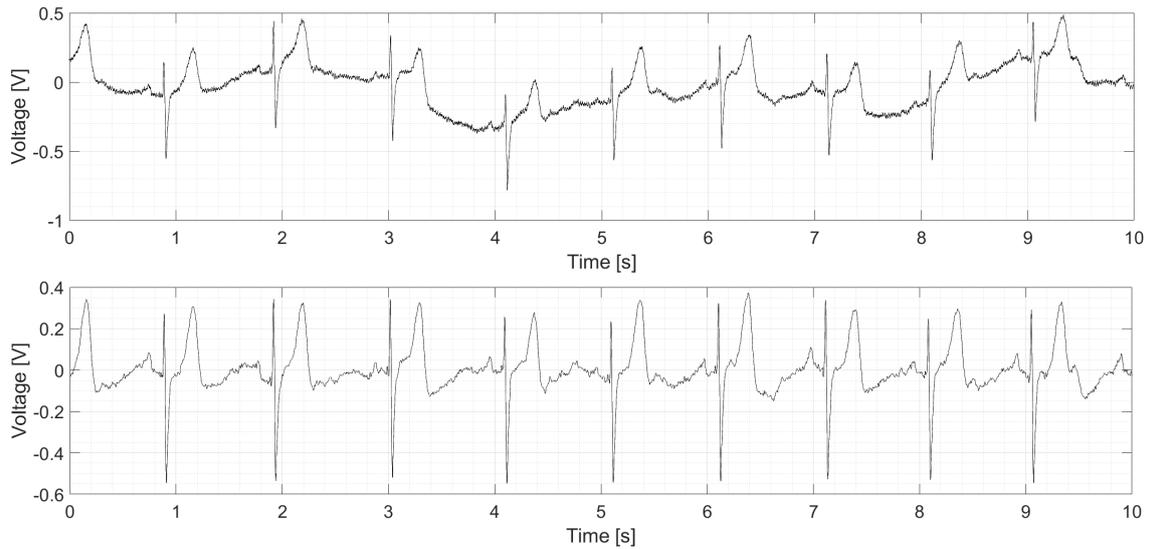


Figure 4.5: The top figure shows raw data from two non-contact electrodes placed on the chest near each shoulder. Dry ground and shielding were used to reduce PLI. The bottom figure shows the same signal with a 50 Hz notch filter as well as a 2 Hz high-pass filter and a 80 Hz low pass filter applied digitally.

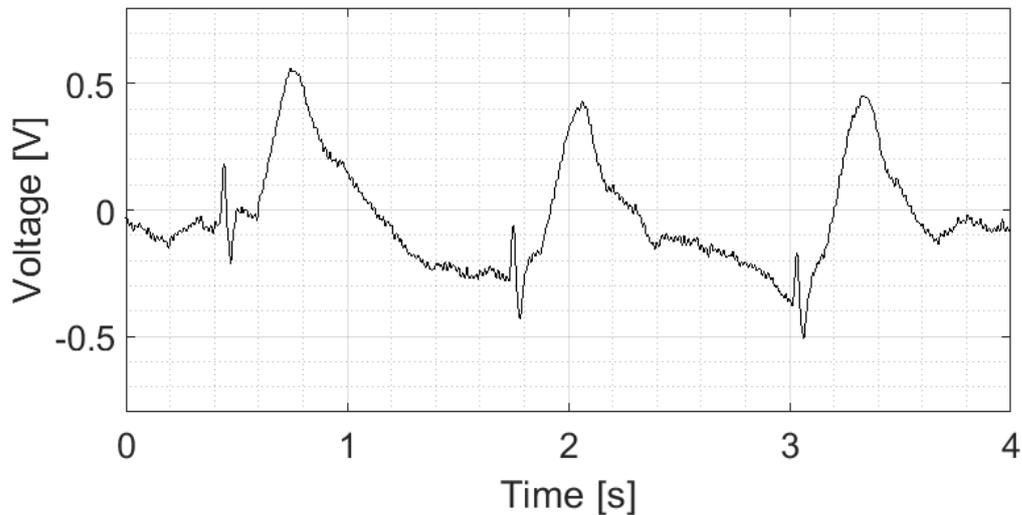


Figure 4.6: Measurement of electrodes through a cotton t-shirt. One electrode was placed on the chest near the right shoulder and one was placed near the left the apex of the left ventricle. The repolarization of the ventricles (T-wave) is pronounced due to the placement of the electrodes. A 50 Hz notch filter as well as a 0.5 Hz high-pass filter and a 80 Hz low pass filter have been applied digitally.

The top pane in Figure 4.5 shows an unfiltered measurement with the non-contact electrodes placed on the thorax near each shoulder. The bottom pane shows the same signal with a 50 Hz notch filter as well as a 2 Hz high-pass filter and a 80 Hz low pass filter applied digitally.

Figure 4.6 shows a measurement made through a thin cotton shirt. One electrode was placed on the chest near the right shoulder and one was placed near the apex of left ventricle. The repolarization of the ventricles (the normally quite inconspicuous T-wave) is pronounced due to the placement of the second electrode. A 50 Hz notch filter as well as a 0.5 Hz high-pass filter and a 80 Hz low pass filter have been applied digitally.

Chapter 5

Discussion

The noise measurements were particularly difficult as any type of vibration seemed to perturb the measurements, e.g. computer fans, walking in the vicinity of the setup etc. Even when the foam stand was introduced there were large variations in average noise levels, especially at 10 Hz. This noise could potentially come from electro-mechanical coupling between cables in the shielded box or from the voltage supply, although care was taken to filter out any noise from the DC source. From the unreliability of the measurements, it is believed that the measurements at 10 Hz are due to other noise sources than the circuitry. However, as no noise source can decrease the total noise, the absolute noise levels of the measurement at 100 Hz should be fairly accurate. However, the discrepancy between the predicted noise vs gain relationship and the measured ditto indicates that these measurements were affected by external interference as well. Hence, the measurements have to be made in an even more isolated environment to draw confident conclusions about the noise profile of the electrode. One method that could improve the setup is to directly solder a large ($\sim 10 \mu\text{F}$) capacitor to the input of C_S . In this way the source capacitance would be approximately the same, whereas basically any signal picked up by the electrode surface would be filtered through the external capacitor.

There were some issues with the settling time when applying the electrode, especially at higher gain. This issue can however easily be resolved. By directing a low passed version of the input signal (e.g. indirectly from the guard of the INA116) to the grounded input, the DC offset is removed on the output. Furthermore, this type of feedback makes higher gain on the amplifier more viable.

In the next version of the electrode, input neutralization should be implemented to some extent to reduce noise levels and make the electrode more stable against variations in source capacitance. Also, using this technique the electrodes can be even smaller while having the same signal gain. If using a reasonably small capacitor and fine tuning the feedback gain, the introduced noise should be almost negligible. Lastly, a lower input capacitance allows for a lower limiting capacitance which would make the electrodes more stable.

A technique called right leg drive is ubiquitous in medical ECG devices to cancel

common mode signals, especially 50/60 Hz power line interference. The common signal mode between two electrodes is inverted and amplified before it is returned to the body through an electrode on the right leg. In our measurements there were a lot of power line interference if not a dry ground and proper shielding were used as in Figure 4.5. A strategically placed right leg drive could improve the signal quality, however, the PLI in the measured signals is already small. Even a purely capacitive right leg drive can prove very effective[64], making all the electrodes entirely non-contact.

The amplification of the electrode was about the same at the cutoff frequency when there was no conducting path between the insulating electrode surface and when there was a direct connection to ground, Part II Section 4.1. There was small difference $\approx 0.2\%$. This minute difference can either come from a parallel circuit with a very large resistance and the $500\text{ G}\Omega$ or other changes in the measurement setup such as capacitive coupling to the input pin (which is simply added to the 10 pF source impedance). In either way the resistance in the solder mask is much larger than the bias resistance.

In the future other operational amplifiers should be considered. The extremely low bias current is not required if (as previously described) the output is high-pass filtered and connected to the other input. At $500\text{ G}\Omega$ and 100 fA bias current, the calculated input voltage offset was around 0.07 V , see Figure 2.5. If an amplifier with 5 pF bias current was used, the input voltage would be around 3.5 V and many amplifiers have supply voltages larger than this. A component search should be made for all amplifiers with a bias current smaller than a specific threshold, and then filtered with regards to the introduced noise, see Section 2.9. A possible candidate is Analog Devices LTC6240 with much lower voltage noise and almost the same current noise as INA116. It has a sufficiently low input bias current and is sold at a fraction of the price of INA116.

The electrodes have been designed around being able to measure the fECG, although to have the same functionality as the CTG it should also be able to measure contractions. As the electrohystrogram (mEHG) is normally much larger in amplitude and on a different frequency band (above the lower cutoff frequency of the developed electrode), around 0.1 to 1.1 Hz , [20] the electrodes should be able to reliably measure contractions. Furthermore, during experiments the electrodes had a visible shift in baseline voltage during breathing (where the amplitude naturally depended on the placement on the electrode). Also this frequency range is included in the bandwidth of the electrode.

Chapter 6

Conclusion

A high impedance non-contact electrode was developed and optimized with respect to noise, stability and resistance to artifacts. The electrode can reliably measure an ECG signal and should be able to measure the fECG with similar signal to noise ratios as state of the art contact electrodes, Figure (4.5). The electrodes are suitable for stationary implementation in stretchy fabric, which facilitates an easy application procedure, reusability and long term fECG monitoring.

Part III

Software development and testing

Chapter 1

Basic theory

This chapter briefly discusses concepts integral to the reviewed and developed fECG extraction methods, including adaptive filtering and different kinds of neural networks.

1.1 Adaptive filtering

In order to remove noise from a signal one usually filters it in some way. For a non-adaptive method the filter is fixed, meaning that it has a constant transfer function, whereas the opposite holds for an adaptive method. Within the field of fECG extraction, adaptive noise cancellation methods are usually employed to suppress noise within the AECG mixture[65]. In such methods, the AECG $y(n)$ is regarded as a sum of the desired subsignal, the fECG $s(n)$, and noise $\eta(n)$, that is $y(n) = s(n) + \eta(n)$. The noise $\eta(n)$ is the composite signal comprising the mECG, other biopotentials, and artifacts stemming from movement etc. It is adaptively subtracted by a finite impulse response filter (a filter whose impulse response, or response of any finite length input, has a finite duration) with coefficients/weights $\mathbf{w} = [w_1, \dots, w_N]$. N denotes the number of weights that are updated in a recursive manner in order to minimize an error signal $e(n)$ [3].

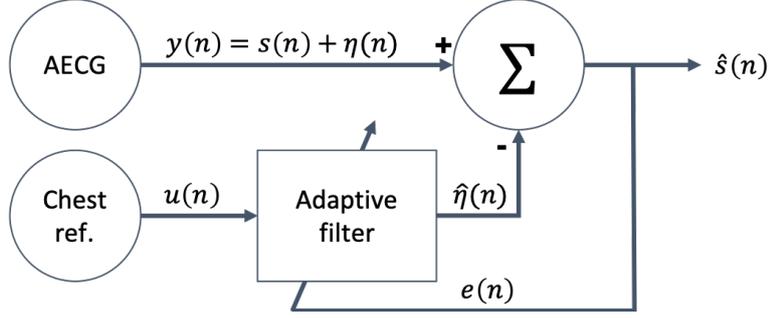


Figure 1.1: A block diagram describing the adaptive noise cancelling process of the AECG $y(n)$ using a chest reference input signal $u(n)$ at time step n . The fECG is the signal of interest $s(n)$, $\eta(n)$ is the total noise, $\hat{\eta}(n)$ the estimated noise, $e(n)$ the estimation error signal, and $\hat{s}(n)$ the output signal. The aim of the filter is to recursively map the fECG free chest signal as closely as possible onto $y(n)$. This mapping should then correspond to the primary noise source in the AECG, the mECG, and can then be subtracted from the abdominal signal.

A schematic of the adaptive filtering process is shown in Figure 1.1. Let an mECG measurement recorded on the chest be the input signal $u(n)$. Due to its placement, it should be virtually fECG free. The aim of the filtering methods is to find a model that maps $u(n)$ to an output $\hat{\eta}(n)$ such that $\hat{\eta}(n)$, in the least mean square error sense, resembles the target signal $y(n)$ as much as possible. By applying a filter that matches $u(n)$ with $y(n)$, the resulting signal $\hat{\eta}(n)$ will predominantly mimic the main source of noise in the AECG mixture, the mECG. Subtraction of $\hat{\eta}(n)$ from the AECG should therefore yield a significant suppression of the mECG component.

1.2 Blind source separation

In terms of NI-fECG extraction, the goal of blind source separation (BSS) methods is to split the underlying sources of the AECG into three different categories: mECG, fECG, and noise. For a standard BSS model, one considers p independent source signals $\mathbf{s} = [s_1^k, \dots, s_p^k] \in \mathfrak{R}^p$ and the observation of equally many mixtures $\mathbf{m} = [m_1^k, \dots, m_p^k] \in \mathfrak{R}^p$ for time index k . The mixtures are assumed to be momentary linear combinations of the sources, $m_i^k = \sum_{j=1}^p a_{i,j} s_j^k$ [66]. On matrix form, one gets the mixing equation $\mathbf{m} = \mathbf{A}\mathbf{s}$, where the observation vector \mathbf{m} at time k contains the AECG at that time instant and $\mathbf{A} \in \mathfrak{R}^{p \times p}$ is the mixing matrix. The general objective of BSS is to extract the random variable vector \mathbf{s} from the observed data in the random variable vector \mathbf{m} , which may be approached in multiple ways. Two spatial separation techniques within the BSS framework are discussed in the fECG extraction methods review in Part III Section 3.1, principal component analysis and independent component analysis.

1.3 Artificial neural networks

An artificial neural network (ANN) is a mathematical tool that function similarly to (simplified) biological neural systems. Oftentimes, it is designed to reproduce human brain functions such as information processing, pattern recognition, and memory. Originally introduced by neuroscientist Warren McCulloch and mathematician Walter Pitts in 1943, it was suggested early on that any computable function could be fully realized by their so-called McCulloch-Pitts artificial neural network, even though some functions would require impractically large networks[67]. Ever since, a myriad of ANN types have been introduced and it is nowadays widely used to solve a variety of problems within numerous research fields and commercial applications.

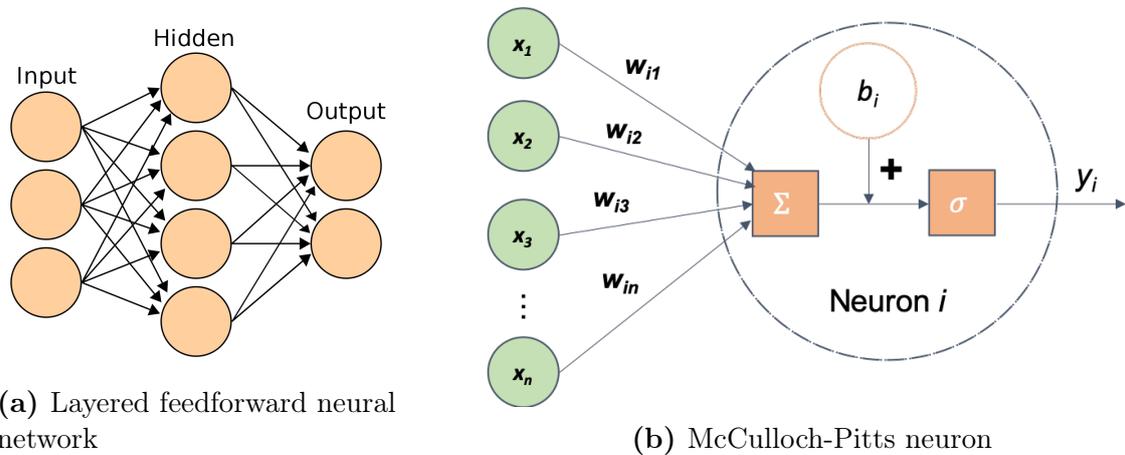


Figure 1.2: The principle and structure of a layered feed forward neural network (a)) with McCulloch-Pitts neurons (b)). In such a network, an input signal is propagating through forward-pointing (left to right) connections, where a numerical operation is conducted in each neuron (apart from in the input layer) before the signal is passed on to the neurons of the next layer. In a McCulloch-Pitts neuron (circumscribed by the dash-dotted circle), the output signal of the neuron is obtained by first calculating a weighted sum of all the n inputs x_j (which correspond to the outputs from the previous layer), then adding a bias term specified for each neuron and finally using the sum as the argument to some activation function σ . a) is reproduced from Cburnett[7].

ANNs utilize intricate mathematical algorithms and computational functions in order to simulate biological neural networks. They comprise building blocks called “artificial neurons” that resemble the structure of real neurons and consist of three main parts: inputs, transformation functions, and output. This introduction will focus on one of the most common configurations, a layered feed forward neural network (FFNN) with McCulloch-Pitts (MCP) neurons. In such a network, shown in Figure 1.2, only forward-pointing connections (left to right) are used. Each of the connections are associated with a weight, w_{ij} , with a specific numerical value. An FFNN starts with a layer of input elements that simply transmits the incoming signal to the next layer without performing any numerical operations. Each input neuron normally corresponds to an element of an input vector.

Usually there is one or more so-called hidden layers between the input layer and the output layer. For each of the hidden layers and the output layer an elementary operation is carried out before the signal is passed on to the next layer. Once the output layer is reached, a final computation is performed to obtain the output of the network. An FFNN with multiple hidden layers is often denoted a “deep neural network”.

In an MCP neuron, the output signal y_i of neuron i in a layer is obtained by first calculating a weighted sum of all the n inputs x_j (which correspond to the outputs from the previous layer)

$$s_i = \sum_{j=1}^n w_{ij}x_j + b_i,$$

where b_i is a bias term. The output y_i is then finally generated by

$$y_i = \sigma(s_i),$$

where σ is an activation function that determines the output range. It could for instance be the logistic sigmoid function

$$\sigma(s_i) = \frac{1}{1 + e^{-cs_i}},$$

where c is a positive constant. The output is then limited within the range $[0, 1]$ and the bias term b_i thus regulates the propensity of getting an output close to 1 or 0. In the case of a completely binary activation function that takes the value 1 if s exceeds some threshold value and 0 otherwise (imitating the biological neuron property of all-or-none firing), the bias then adjusts the threshold value needed by the weighted sum to get a non-zero output from the neuron.

There are various ways to train neural networks, that is to set their weights and biases (and occasionally also structural parameters such as number of neurons in a certain layer or number of hidden layers) in order to solve a specific computational task. Usually, learning methods are divided into supervised and unsupervised methods. For supervised methods, there must be a desired output vector from the network for each input vector inserted into the ANN. The objective of the learning process then becomes to minimize the error (the difference) between the desired and actual output of the network. The weights and biases are then driven towards values that generate the desired outputs for each set of inputs. Conversely, unsupervised learning methods are used in situations where complete information about all target output values is not available.

The range of applications of feedforward neural networks is vast. Among other things, they can be used to make decisions or draw conclusions from incomplete or noisy data, they can recognize and classify patterns, conduct time series analysis and prediction, approximate functions, and perform signal processing operations[67].

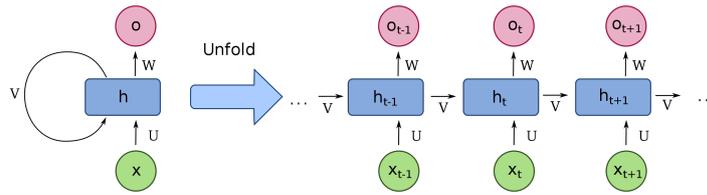


Figure 1.3: To the left: the network structure of a one-unit recurrent neural network. A single input state x is propagated to a hidden state h with the connection weight U , before getting passed on with the weight W to the output state o . The output of the hidden layer is also passed on to itself the subsequent time step with the weight V , where it will be combined with the succeeding (weighted) input state to form the next hidden layer output. To the right: an unfolded version of the same network, where each of the three “columns” correspond to one time step. Note how the previous output of the hidden layer is combined with the current input state to obtain the new output of the hidden layer. Reproduced from Deloche[8].

1.4 Recurrent neural networks

An FFNN can be designed and trained to perform an immense amount of input-output mappings, but one of its major disadvantages is that its output will always be the same for any given input signal. That means that FFNNs lack dynamic memory, which is crucial in various applications where previous input to some extent should affect the interpretation of the current input (such as for sentence handling in language translators). This could be solved by introducing feedback (recurrent) connections, yielding a recurrent neural network (RNN). In a fully recurrent neural network, any neuron may, in principle, receive input from any other neuron (including itself) in the network. For instance, the output of the network could be fed to a hidden layer or the output of a hidden layer could be added to the same hidden layer the subsequent network propagation sequence. It is common that each network sequence corresponds to a certain time step and that it is fed with the current input state for each time step[68]. Figure 1.3 illustrates one of the simplest topologies of an RNN and how it may be unfolded in time so as to become practically equivalent with a deep neural network, where each hidden layer corresponds to a timestep.

1.5 Genetic algorithms

A genetic algorithm, GA, is a search heuristic inspired by the theory of natural evolution. The algorithm reflects the process of natural selection, in which the fittest individuals are predominantly selected for reproduction and have a greater influence on the offspring of the next generation. GAs can be highly useful in e.g. optimization problems where the objective function cannot be explicitly expressed

as a mathematical function, but rather only obtained after, for instance, a lengthy simulation process[69]. There are different versions of GAs, but the one implemented in this thesis will be based upon the methodology, terminology, and definitions in Wahde[69].

Consider the optimization problem of finding the maximum of an n -dimensional function $f(x_1, \dots, x_n)$. The search space is defined as the set of allowed values for the variables $\mathbf{x} = (x_1, \dots, x_n)$. In order to solve this problem with a GA the variables are encoded in strings of digits, normally referred to as chromosomes. The digits comprising the chromosome are denoted genes and may encode the information in various ways. In the original genetic algorithms[70], introduced in the 1970s, a standard binary encoding scheme was suggested that is still common to this day. The individual genes can then take on the values 1 or 0 and a generic variable x is decoded from genes g_1, \dots, g_k through

$$x = -d + \frac{2d}{1 - 2^{-k}} (2^{-1}g_1 + \dots + 2^{-k}g_k),$$

which yields a value in the range $[-d, d]$. A drawback with this scheme, however, is that a small change in a variable might require many genes to switch values. For example, assume that the ideal chromosome that optimizes the objective function in a case with a single variable is 10000, but the best chromosome of a GA session has converged to 01111. The decoded variable will then have a value close to the optimal one, but to reach it all gene values need to be flipped (which is a highly unlikely event). To avoid that the algorithm gets stuck due to the encoding scheme, a so called Gray code representation[71] is used in this thesis. It may be written

$$x = -d + \frac{2d}{1 - 2^{-k}} 2^{-k} \gamma(g_1, \dots, g_k),$$

where $\gamma(g_1, \dots, g_k)$ is the corresponding integer in the range $[0, 2^k - 1]$ of the binary number constituted by the genes g_1, \dots, g_k . An increment of one in the value of this integer would only require a single (albeit specific) gene to change.

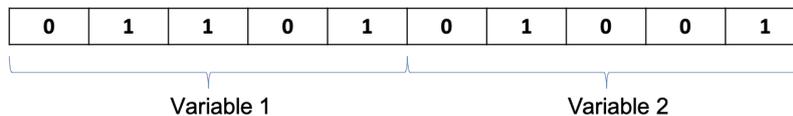


Figure 1.4: A chromosome that encodes two variables with 5-bit accuracy.

Upon initialization of the genetic algorithm, a set of chromosomes (denoted population) is randomly generated by assigning the genes of each chromosome with a 1 or 0 (with equal probability). These chromosomes make up the first generation. After they have been initialized, each of them is decoded to obtain their corresponding individual, i.e. their corresponding set of n variable values. Commonly, a chromosome of length m is divided into n parts of length m/n that constitute the binary

representation of the respective variable, see Figure 1.4. Evidently, the accuracy of a variable is governed by the number of genes it contains.

When the variables are decoded, the individual is evaluated. During this procedure, a problem-dependent fitness value is assigned to the individual depending on its variable based performance. Usually the goal is to maximize the fitness value, which is later used when individuals are selected and combined for reproduction. Obviously this do not exclude minimization problems, since minimization of a measure is equivalent to maximization of its inverse.

When all individuals of the population have been decoded and evaluated, it is time to shape the next generation. It is desirable to favor more fit individuals without making the selection procedure fully deterministic. This since an individual that is slightly superior to the others, but far from the global maximum, then probably would have dominated the population and therefore obstructed the genetic algorithm from finding the global optimum. Conversely, a worse performing individual may contain chromosome sequences that will excel when combined with certain genetic material from another individual. A common way to make the selection more stochastic is to pick individuals through a tournament selection. Two (or more, it is predefined by the tournament size parameter) individuals are then picked randomly from the population at a time (with or without replacement) and then the individual with the highest fitness is selected with a specified probability p_{tour} called tournament selection parameter. This implies that the weaker individual is selected with probability $1 - p_{\text{tour}}$. Normally, p_{tour} is set to around 0.7-0.8. In the case of more than two individuals the best individual is also selected with probability p_{tour} , but if it is not chosen the procedure is repeated for the remaining individuals in the tournament, i.e. the best of the remaining individuals is selected with probability p_{tour} .

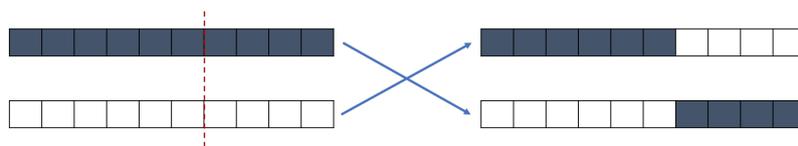


Figure 1.5: The typical crossover procedure in genetic algorithms. Each square is a gene and the dashed, red line is the crossover point which is picked randomly. The first part (i.e. before the crossover point) of the first chromosome is combined with the second part of the second chromosome and vice versa, forming two new chromosomes/individuals.

The selection process is used to pick pairs of individuals that get to shape two new individuals through reproduction, i.e. by combining their genetic material. In a procedure denoted crossover, the chromosomes of the two selected individuals are cut at a randomly chosen crossover point and then “cross-assembled” so that the first part of the first chromosome is attached to the second part of the second chromosome and vice versa, see Figure 1.5. Crossover based reproduction is a very efficient mechanism for spreading high-performing genetic material, so to find a balance between exploitation and exploration in the optimization algorithm and not converge to local optimas crossover is typically only applied with a certain

probability. In the other cases the selected pair of individuals are simply copied completely in the formation of the two new individuals.

After new individuals have been formed through the reproduction procedure, they are exposed to mutation in order to introduce new material for the GA to work with. Mutation is employed simply by changing the value of some randomly selected genes of the new individuals. A rule of thumb is that the mutation probability should be around $1/m$, where m is the chromosome length. Consequently, the value of one gene per individual will get flipped on average.

The final GA component employed in this thesis is elitism, which ensures that the best performing individual of each generation is preserved until the next generation. This is enforced by replacing one or more of the freshly formed individuals with exact copies of the best individual before replacing the old population with the new set of individuals and starting a new generation.

1.6 Continuous wavelet transform

Similarly to a Fourier transform, the continuous wavelet transform (CWT) utilizes inner products to quantify the similitude between a signal and an analyzing function. For a regular Fourier transform, the analyzing functions have the form of complex exponentials, $e^{j\omega t}$, and the transformed function will be a function of a single variable, the angular frequency ω . In the case of a short-time Fourier transform, STFT, the analyzing functions are time-dependent windowed complex exponentials $w(t)e^{j\omega t}$ that yield a transformed function of two variables. The coefficients of the transformed functions represent the similarity between the signal and a sinusoid with the angular frequency ω within a time region of a specific length centered at time τ [72].

In a CWT, the analyzing function is instead a wavelet, ψ and the signal is compared to dilated (i.e. compressed/stretched) and shifted versions of this wavelet. Through comparison of the 1D signal with the wavelet at different positions and scales (governed by the shifting and dilation respectively), the transform returns a function of two variables. If a complex-valued signal is transformed, the CWT will be a complex-valued function of position and scale. For a real-valued signal, the CWT will instead be a real-valued function of the same variables. The continuous wavelet transform of a function $f(t)$ can formally be written

$$C(a, b; f(t), \psi(t)) = \int_{-\infty}^{\infty} f(t) \frac{1}{a} \psi^* \left(\frac{t-b}{a} \right) dt, \quad (1.1)$$

where $a > 0$ is the scale parameter and b the position parameter, which are continuously varied to acquire the CWT coefficients $C(a, b; f(t), \psi(t))$. The complex conjugate is denoted by $*$. As seen in (1.1), in addition to the scale and position values, also the choice of wavelet all affect the obtained CWT coefficients values. In analogy with a regular Fourier transform, multiplication of the coefficients by their

respective shifted and scaled wavelet renders the constituent wavelets of the original signal[72].

A variety of wavelets can be used in CWTs and the choice is normally based upon what signal features one attempts to detect. For instance, one wavelet may be appropriate for signals with abrupt discontinuities, whereas another might be suitable for oscillations with smooth onsets and offsets[72]. Generalized Morse wavelets[73] are a family of exactly analytic wavelets. Analytic wavelets are complex-valued wavelets whose Fourier transforms are not supported on the negative real axis. Morse wavelets are widely used for signals with e.g. time-varying amplitude and frequency and/or localized discontinuities, and have been described as an ideal starting point for general purpose usage[74].

The resulting CWT can be plotted in a scalogram, which shows the absolute value of the CWT as a function of time and frequency. Frequency is plotted along the y-axis with a logarithmic scale. Normally, a so-called cone of influence is added to the plot, which delineates the boundary outside which edge effects become significant.

1.6.1 Scaling and shifting

As can be inferred from (1.1), longer scales (regulated by a) correspond to more stretched wavelets. A more stretched wavelet implicates that a longer portion of the original signal will be compared with it and thus that the wavelet coefficients will measure coarser features. Conversely, a small a yields a compressed wavelet and measures rapidly changing (i.e. high frequency) details.

To shift a wavelet means that its onset is delayed or advanced. In mathematical terms, a delay of a function $f(t)$ by k can be represented by $f(t - k)$.

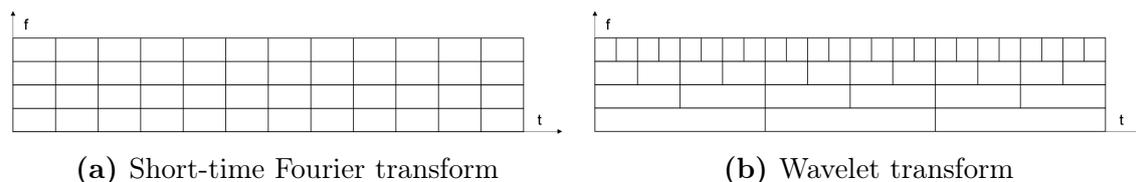


Figure 1.6: An illustration of the qualitative difference in time and frequency resolution of a short-time Fourier transform and a wavelet transform. The resolution ratio of the two quantities is constant for the STFT, whereas it can be frequency dependent for a wavelet transform.

1.6.2 CWT as a windowed transform

STFTs are often described as windowings of the signal that enables local frequency analysis. A significant limitation of the STFT is that the window size is constant, which imposes a trade off between good resolution of time or frequency. A long time window results in a good frequency resolution and a poor time resolution, since the Fourier transform within each window do not have any time resolution at all,

whereas a shorter time window has the opposite effect. Wavelet transforms, on the other hand, offer a windowing method with variable-sized regions. Through wavelet analysis, long time intervals can be used where more accurate low-frequency information is required, and shorter time spans can be employed when high-frequency information is desired[72]. An illustration of the resolution differences between the two transforms is displayed in Figure 1.6. Examples of a CWT of an SECG signal and an AECCG signal respectively were displayed in Figure 3.2 in Part I.

Chapter 2

Database and figures of merit

2.1 Non-Invasive Fetal Electrocardiogram Database

The NI-FECG database[1] comprises 55 multichannel abdominal NI-FECG recordings, measured with contact electrodes on a single subject between 21-40 weeks of gestation. The records have different length and were taken weekly apart from some weeks when two or more records were acquired. Each record contains two thoracic (chest) channels and three to four abdominal channels with variable placements (the electrode position was varied in order to improve SNR). The sampling rate of all signals was 1 kHz with 16-bit resolution and a bandpass filter (0.1 Hz-100 Hz) and a notch filter (50 Hz) were applied during the data acquisition.

14 of these records, where the fQRS complexes are fairly visible on at least one channel, have previously been manually annotated[3]. A minute-long signal, starting 30 seconds after the beginning of the record, was excerpted from three abdominal and two thoracic channels for each of the records. From these abdominal snippets, visually examining the channel with the most visible fECG for each record, 2148 fQRS complexes were annotated in total. Since the abdominal channels may be considered independently, $3 \times 14 = 42$ min of annotated data with $3 \times 2148 = 6444$ reference fQRS was available for fECG extraction efforts. The data was generally of quite good quality with varying SNRs for the mECG and fECG and some lesser artifacts that were not discarded.

2.2 Figures of merit

In compliance with the ANSI/AAMI guidelines[75], the positive predictive value (*PPV*) and the sensitivity (*Se*) should be defined as

$$Se = \frac{TP}{TP + FN}, \quad PPV = \frac{TP}{TP + FP}. \quad (2.1)$$

Here TP denotes true positive (a correctly identified fQRS), FN false negative (a missed existing fQRS), and FP false positive (an erroneously detected non-existing fQRS).

For optimization of algorithm parameters, the F_1 statistic was used as a performance indicator. For binary classification tasks, it is defined as the harmonic mean of PPV and Se

$$F_1 = 2 \cdot \frac{PPV \cdot Se}{PPV + Se} = \frac{2 \cdot TP}{2 \cdot TP + FN + FP}. \quad (2.2)$$

As can be seen in (2.2), flaws in FN and FP affect F_1 to the same degree. In contrast to the arithmetic mean, F_1 is suitable for situations when the average of rates is considered[76]. As an example of the higher suitability of F_1 over the arithmetic mean, one may consider the case where $Se = 1$ and $PPV \approx 0$ (which corresponds to every datapoint of the signal being annotated as an fQRS by the QRS detection algorithm). Naturally, such an algorithm is highly undesirable, which should be reflected in its score. Nevertheless, an arithmetic mean would yield around 0.5 (corresponding to 50% accuracy), while F_1 would be equal to roughly zero.

According to ANSI/AAMI guidelines[77], a tolerance of ± 5 beats per minute (BPM) is advised for heart rate monitoring. The fHR was derived from the time difference of two neighboring QRS detections made by the used algorithm. However, an alternative, slightly more lenient, acceptance criterion was primarily used for the methods developed in this thesis. This criterion accepted detections within 50 ms of the corresponding reference annotation. The reason for this tolerance criterion was uncertainties in the exactness of the manual annotations available for benchmarking of the extraction methods developed in this thesis. Furthermore, the main objective was to develop methods that reliably detects incidences of fetal heart beats (without making false detections), to optimize the precision of each correct detection was not top priority. This was partly due to the dubiousness of the reference annotation resolution and partly to address one of the main issues with current technology. Since 50 ms corresponds to merely one third of the refractory period of a fetal heart, the criterion should still be strict enough to demand that the right ECG cycle was detected for each annotation.

Chapter 3

Review of non-invasive fECG extraction methods

Quite some work has previously been made within the field of fECG extraction from electrodes on the maternal abdomen, although satisfactory performance is to the authors' knowledge yet to be reached. In particular, two distinct categories of algorithms are discernible: algorithms that utilize spatial characteristics of the different abdominal channels (electrodes), and algorithms that use temporal features in the signal of each channel. A few of the more promising methodologies within each category will be presented briefly below.

3.1 Spatial methods

Two of the more well-known BSS methods are principal component analysis (PCA) and individual component analysis (ICA). They have previously been applied also within the context of fECG extraction, using spatial information of a multi-electrode abdominal configuration to separate sources contributing to the AECG. A central assumption of classical PCA and ICA methods is that the mixing matrix between the sources is linear and stationary. The latter property is often not a good approximation over long recordings, since e.g. the relative positions of the sources and the electrodes might shift due to fetal or maternal motion. This issue could however be alleviated by regenerating the matrix frequently. Neither the assumption about linearity is unproblematic. In fact, fetal signals and other interferences and noises have been shown to not always be linearly separable[20].

The core idea of the techniques is to project the original signals into the “source domain”, i.e. the domain of the abdominal signals that have been transformed by a BSS method. The components representing the mECG and noise can then be removed before the signals are projected back, which finally should result in signals predominantly comprising fECG components.

3.1.1 Principal component analysis

PCA is an orthogonal transformation of the dataset in directions that maximize the variance. In this manner it is expected to decorrelate the dataset, that is map the possibly correlated variables (i.e. the sources) into a set of orthogonal variables named principal components.

3.1.2 Independent component analysis

PCA suffers from two limitations that severely impair its potential within fECG extraction, namely that the axes of the transformed principal components basis have to be orthogonal and that the algorithm attempts to decorrelate the data by removing second order dependencies. The fECG and the mECG are normally not orthogonal in the observation domain (the domain before BSS has been applied, consisting of the raw abdominal recordings), so to look for independence instead of decorrelation could potentially be a more effective criterion when separating the fECG from the AECG. ICA addresses these two limitations of PCA. Classical ICA methods (for example JADE[78] and FastICA[79]) assume that the source signals are statistically independent and non-Gaussian and apply higher order statistics than PCA to split the signal into its subcomponents[3]. There are also variants where the sub-components are assumed to be periodic rather than independent.

3.2 Temporal methods

A common approach for fECG extraction is to model the mECG contribution to the abdominal mixture in some manner and then subtract it from the latter. The only remaining signal, save for a small maternal residue depending on the accuracy of the model, in the AECG with “ECG characteristics” (see Part I, chapter 3) should then be the fECG. This modelling can be conducted in various ways. A few notable examples are described below.

3.2.1 Template subtraction

There have been various attempts and versions of template subtraction (TS) within the field of fECG extraction, see for instance [80, 81, 82, 83] for thorough descriptions. The algorithms have had varying degrees of adaptability, but they all share the same foundation; construction of an mECG template cycle to be subtracted from the subsequent mECG cycle within the AECG. The procedure is conducted for each abdominal channel.

In the simplest TS method, an mECG template cycle is centered on the maternal R-peaks with fixed durations for the P, QRS, and T waves respectively. Other adaptations include e.g. scaling of the whole template mECG cycle or the individual waves

with a constant. The scaling was based on minimization (in the mean square error sense) of the mismatch between an average of multiple selected cycles/individual waves (the template) and each individual incoming mECG cycle/wave. In the case of individual waves, the problem may be written[82]:

$$e^2 = \min (\|\mathbf{T}\mathbf{a} - \mathbf{m}\|^2), \quad \mathbf{T} = \begin{bmatrix} \mathbf{t}_P & 0 & 0 \\ 0 & \mathbf{t}_{QRS} & 0 \\ 0 & 0 & \mathbf{t}_T \end{bmatrix}$$

$$\implies \mathbf{a} = (\mathbf{T}^T\mathbf{T})^{-1} \mathbf{T}^T \mathbf{m}.$$

Here e^2 denotes the mean square error (MSE), \mathbf{a} the scaling vector, \mathbf{m} the individual mECG cycle, and \mathbf{t}_P , \mathbf{t}_{QRS} , and \mathbf{t}_T are the template parts adhering to the P, QRS, and T wave respectively.

In another implementation[84], the template mECG was instead assembled by weighting the previous cycles so the MSE was minimized (as opposed to the other versions where the weights were equal for all the involved cycles). Common to all the methods, however, is that the number of cycles used for the template construction is the most crucial parameter to optimize. Taking the non-stationary behavior of the ECG into account, the template may be updated by replacing the contribution from the oldest cycles with the incoming ones. New cycles that mismatches the template can be rejected in order to avoid misdetections. To that end, Pearson's correlation coefficient between an incoming cycle and the template cycle has been used[3]:

$$r = \frac{\sum_{i=1}^n (\mathbf{t}_i - \bar{\mathbf{t}}) (\mathbf{c}_i - \bar{\mathbf{c}})}{\sqrt{\sum_{i=1}^n (\mathbf{t}_i - \bar{\mathbf{t}})^2} \sqrt{\sum_{i=1}^n (\mathbf{c}_i - \bar{\mathbf{c}})^2}},$$

where \mathbf{t} and \mathbf{c} are the template and incoming mECG cycles respectively, and $\bar{\mathbf{t}}$ and $\bar{\mathbf{c}}$ are their respective means.

The central issue with TS based techniques is their dependence on accurate mQRS detection (state-of-the-art adult QRS detectors usually have a 99 % accuracy over a variety of diverse databases[3]). In most cases (unless in the refractory period of the fQRS), a missed mQRS detection will end up yielding a false positive when the fQRS detection is applied. TS methods are also intrinsically sensitive to ectopic beats, since the template approach by its nature assumes that there are great morphological and temporal similarities between nearby heart beat cycles. One way to handle this issue would be to construct templates for different beat types (e.g. normal sinus rhythms and premature ventricular contraction) and then switch template based on the currently most probable beat type. Similar attempts have recently been made for templates based on Kalman filters[85].

3.2.2 Adaptive filtering

Two kinds of linear adaptive filters, see Part III Section 1.1, will be considered here, the least mean square (LMS) and the recursive least square (RLS) adaptive filters.

They are widely used, including multiple attempts related to fECG extraction, but one of their main disadvantages in the context of fECG extraction is their susceptibility to reference thoracic signals of poor quality.

3.2.2.1 Least mean square adaptive filter

The LMS filter was implemented in NI-fECG extraction methods already in 1975[65], although no quantitative results were reported. LMS searches for filter coefficients that minimize the mean square error $e^2(n)$ between the output of the filter $\hat{y}(n)$ and the desired response $y(n)$. Consider the N last samples of the input signal, $\mathbf{u}(n) = [u_1(n - N + 1), \dots, u_1(n)]^T$, $\forall n > N$, let $\mathbf{w}(n) = [w_1(n), \dots, w_N(n)]$ be the filter weights and $e(n) = y(n) - \mathbf{w}^T \mathbf{u}(n)$ the error rate at time step n . The optimal weight vector is then obtained through the so-called Wiener solution $\mathbf{w}_o = \mathbf{R}^{-1} \mathbf{P}$, where \mathbf{R} is the correlation matrix of the input $u(n)$ and \mathbf{P} is the cross correlation between $u(n)$ and $y(n)$.

However, calculation of the optimal solution is rarely practical in practice[86]. It requires computation of the auto correlation and cross correlation matrices and matrix inversion, which is computationally expensive (and thus time consuming). Furthermore, if the signals are non-stationary both \mathbf{R} and \mathbf{P} will be time dependent and therefore have to be recomputed more or less continuously. In real time applications, an update of the weights on a sample by sample basis update that still converges towards the optimal Wiener solution is required. This can be done through a steepest gradient descent based technique that adaptively regulates the weight values every time step[65]. The new estimate of the weight vector $\mathbf{w}(n + 1)$ is then equal to the current estimate $\mathbf{w}(n)$ plus a term proportional to the negative gradient of the error at the n th iteration[87]. The weight update rule then becomes $\mathbf{w}(n + 1) = \mathbf{w}(n) - \mu \nabla E [e^2(n)]$. By making the assumption that the true gradient of the error can be estimated by the momentary gradient of the error, i.e. $\nabla E [e^2(n)] \approx \nabla e^2(n) = -2e(n)\mathbf{u}(n)$, the update rule can then be rewritten $\mathbf{w}(n + 1) = \mathbf{w}(n) + 2\mu e(n)\mathbf{u}(n)$. If the factor 2 is included in μ , the LMS adaptive algorithm may be summarized as:

$$\begin{aligned}\hat{y}(n) &= \mathbf{w}^T(n - 1)\mathbf{u}(n), \\ e(n) &= y(n) - \hat{y}(n), \\ \mathbf{w}(n) &= \mathbf{w}(n - 1) + \mu e(n)\mathbf{u}(n).\end{aligned}$$

Here the first equation yields the filter prediction, the second equation the error evaluation, and the third equation updates the filter weights at time step n . Note how the filter length N and the gradient descent step size μ needs to be set, preferably through optimization on a training set.

3.2.2.2 Recursive least square adaptive filter

The RLS method aims to minimize the total squared error between the filter output and the desired response, i.e. as opposed to the LMS algorithm, it takes the total error from the beginning of the signal to the incoming sample into account. A forgetting factor, $\lambda \in [0, 1]$ regulates to what extent old data contributes to new updates of the filter coefficients. For $\lambda = 1$ all past samples would have the same impact, and for $\lambda \rightarrow 0$ merely the most recent samples contributes. The objective may be defined as finding the parameters that minimizes the “loss-function” $\epsilon(n)$:

$$\epsilon(n) = \sum_{i=1}^n \beta(n, i) e(i)^2 = \sum_{i=1}^n \beta(n, i) [y(i) - \hat{\eta}(n, i)]^2.$$

Here $\hat{\eta}(n, i) = \mathbf{w}^T(n) \mathbf{u}(i)$ and $\beta(n, i) = \lambda^{n-i}$ (for the exponentially weighted least squares solution). The RLS method employs the following update procedure at each time step [88, 89]:

$$\begin{aligned} e(n) &= y(n) - \mathbf{w}^T(n-1) \mathbf{u}(n), \\ \mathbf{k}(n) &= \frac{\mathbf{P}(n-1) \mathbf{u}(n)}{\lambda + \mathbf{u}(n)^T \mathbf{P}(n-1) \mathbf{u}(n)}, \\ \mathbf{P}(n) &= (\mathbf{I} - \mathbf{k} \mathbf{u}(n)^T) \mathbf{P}(n-1) \frac{1}{\lambda}, \\ \mathbf{w}(n) &= \mathbf{w}(n-1) + \mathbf{k}(n) e(n), \end{aligned}$$

where \mathbf{I} is the identity matrix and $\mathbf{P} = \varphi(n)^{-1}$ with $\varphi(n) = \sum_{i=1}^n \lambda^{n-i} \mathbf{u}(i) \mathbf{u}(i)^T$.

The filter length N and the forgetting factor λ has to be chosen carefully, preferably through optimization on a training set. Normally, RLS converges faster than LMS and yields more accurate results, but it is also more computationally demanding.

3.2.3 Kalman filtering

In recursive Bayesian filtering, the posterior distribution $P(\mathbf{x} | \mathbf{y}_{1:k})$ of a hidden state random variable \mathbf{x} at time k is recursively estimated using measurements $\{\mathbf{y}_k\}$ and a prior estimate of \mathbf{x} based on knowledge of the system. The parameters are assumed to be random variables, evolving in accordance with an evolution equation (the prior estimate model), and are observed via measurements related to the current state \mathbf{x} through a measurement equation. The system dynamics are formulated as

$$\begin{cases} \mathbf{x}_k = \mathbf{G}_{k-1} \mathbf{x}_{k-1} + \mathbf{w}_{k-1} & \text{(evolution equation),} \\ \mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{v}_k & \text{(measurement equation).} \end{cases}$$

Here \mathbf{G}_{k-1} denotes the state transition matrix and \mathbf{H}_k the observation matrix, which are the noiseless connections between the current and previous state vector and the measurement vector and the state vector respectively. $\{\mathbf{w}_k\}$ and $\{\mathbf{v}_k\}$ are the respective associated noises, which are assumed to be white, zero-mean, uncorrelated ($E[\mathbf{v}_k, \mathbf{w}_j^T] = 0$), and have the stationary covariance matrices $\mathbf{Q}_k = E[\mathbf{w}_k, \mathbf{w}_k^T]$ and $\mathbf{R}_k = E[\mathbf{v}_k, \mathbf{v}_k^T]$ respectively. Furthermore, it is assumed that \mathbf{Q}_k and \mathbf{R}_k are diagonal, i.e. that the components of the noise models are uncorrelated.

A Kalman filter (KF) makes an estimation of the state \mathbf{x}_k based on the noisy measurements $\{\mathbf{y}_k\}$ and knowledge about the dynamics of the system. In 2003, McSharry et al. introduced a dynamical model of the ECG[90] that uses a set of Gaussian functions to approximate ECG cycles. In a discretized form with a small sampling period $\delta \rightarrow 0$, it was used for a KF[3]:

$$\begin{cases} \theta_{k+1} & \equiv (\theta_k + \omega\delta) \bmod 2\pi \\ z_{k+1} & = z_k - \sum_{i=1}^N \delta \frac{\alpha_i \omega}{b_i^2} \Delta\theta_{i,k} \exp\left(-\frac{\Delta\theta_{i,k}^2}{2b_i^2}\right) + \eta_k. \end{cases}$$

Here θ_k and z_k denote the discrete phase and amplitude, α_i , b_i , and χ_i the peak amplitude, width, and center parameters of the N Gaussian functions, $\Delta\theta_{i,k} = \theta_k - \chi_i$, and η_k is a perturbation term representing random additive noise that models the error of the ECG model as compared to a real ECG. A version of KF for non-linear systems, the extended Kalman Filter (EKF), was implemented in order to filter the mECG from e.g. the fECG, which was regarded as noise. The filtered mECG could then be subtracted from the AECG. To find the Gaussian parameters, non-linear curve fitting in combination with a random search was conducted until the RMS error between a template mECG cycle and the cycle obtained through the Gaussian functions was below 5%.

Both for the TS and the EKF frameworks an estimated mECG cycle is subtracted from the AECG. Unlike TS, however, there are no constraints on P, QRS, and T wave lengths for the EKF. While TS methods subtract the mECG cycle on a window around the mQRS with a specified length, the EKF continuously estimates the mECG. Moreover, the EKF is more adaptive and would better handle cases with a very high level of non-stationarity. Similar to TS, however, EKF requires prior and (particularly) precise information about the location of the mQRS. This makes it especially vulnerable to a noisy mECG.

3.3 Performance

Unfortunately, direct comparison of the respective performances of different fECG extraction methods is in many cases very difficult since various evaluation methods and data sets have been used. However, the ‘‘PhysioNet Computing in Cardiology Challenge 2013 - Noninvasive Fetal ECG’’[91] was held in an attempt to compare and evaluate current methodologies and remediate the deficiency of public high standard databases with expert annotations and/or reference signals. Although a lot of the

data used in the challenge is still not open to public, variants of some of the methods listed in this chapter were implemented and benchmarked during the challenge and ensuing studies. In both cases below, a Pan-Tompkins (concept described in Part III Section 5.1) based fQRS extraction method was used.

When comparing various TS methods, a technique where it was combined with PCA[80], outperformed the others[3]. It was definitely the most adaptive, suggesting that the mECG cycle was subtracted to a greater extent with more adaptive methods.

In a comparison of the best performing TS version, LMS, RLS, and an echo state neural network (the concept is described thoroughly in Part III Section 4.3.3) on a commercial database, the results presented in Table 3.1 were obtained[3]. Note that the ± 5 BPM criterion was used for F_1 during this comparison.

Method	F_1 (%)
TS_{pca}	89.32
LMS	87.88
RLS	88.23
ESN	90.17

Table 3.1: List of methods and their respective performance on a commercial database. The results are reproduced from Behar[3].

As seen in Table 3.1, the relative performances were fairly even, making it difficult to draw confident conclusions on whether it was the specific content of the data set or the quality of the respective methods that led to the moderate variations in outcome. When the methods were applied to another database (basically the same as the one used in this thesis), similar levels of performance were obtained once more, yet again with a slight victory for the ESN ($F_1 \sim 97\%$ compared to $F_1 \sim 95\%$ for the others)[3]. Since one of the core differences is that ESNs, as opposed to LMS and RLS, can handle non-linear relationships between the chest and AECG, the quite even results implies that the mapping was mainly linear. However, looking into the performance on each record of the commercial database, the ESN achieved a much higher F_1 score than the RLS and LMS techniques on two of them. Further inspection indicated that this pronounced difference was due the ESN being better at removing the mECG and not due to the signal quality. It appears reasonable that the mapping is mostly (but not always) linear, and that the more complex nature of the ESNs mainly pays off when that is not the case.

In the 2013 challenge, the results listed in Table 3.2 were obtained (on another data set, but also with the ± 5 BPM criterion) for implementations of a regular TS, the TS-PCA, Extended Kalman filter (denoted TS_{EKF}), PCA, and ICA[3].

Method	<i>Se</i> (%)	<i>PPV</i> (%)	<i>F1</i> (%)
TS	81.8	81.7	81.6
TS _{pca}	88.1	84.5	86.1
TS _{EKF}	83.0	81.1	81.9
PCA	57.4	47.9	51.6
ICA	69.1	60.0	63.7

Table 3.2: List of methods and their respective performance in the PhysioNet Computing in Cardiology Challenge 2013. The results are reproduced from Behar[3].

Clearly, based on the results presented in Table 3.2, PCA and ICA do not seem to be viable alternatives on their own. Moreover, TS_{EKF} did not outperform simpler TS methods (some of which are not displayed in the table). This could potentially be a result of the automated Gaussian initialization procedure that was employed (examples were found where the different waves were not well matched by the Gaussian fitting), or because the covariance matrices were set as stationary. It was found that gain factors governing the relative trust in the ECG model and the observations was of paramount importance to the EKF performance. The best results were obtained for a much higher trust in the ECG model than the observations, which resembles a TS technique with the additional perks of some further adaptability and no constraints on the P, QRS, and T wave lengths[3].

It is also noteworthy that it has been estimated that a representation of a typical ECG signal requires 4-6 statistically uncorrelated dimensions[92]. Consequently, it is probable that the four abdominal channels available during the benchmarking are not sufficient for optimal BSS performance.

Another important result from the challenge was that certain combinations of methods (e.g. TS and ICA) outperformed single methods. It is clear that one should not necessarily focus entirely on finding the perfect single fECG extraction method, but rather also look into combinations of methods in order to take advantage of their respective strengths and weaknesses.

Chapter 4

Experimentally evaluated methods - FECG extraction

The feasibility of three different fECG extraction methods has been scrutinized more thoroughly in this thesis. In the case where no maternal chest reference electrode is available, the viability of different kinds of ICA have been investigated. For the case where one has access to a thoracic signal, two kinds of recurrent neural networks have been evaluated, LSTMs (long short-term memory) and ESNs (echo state neural networks). Useful preprocessing steps are also discussed briefly in this chapter.

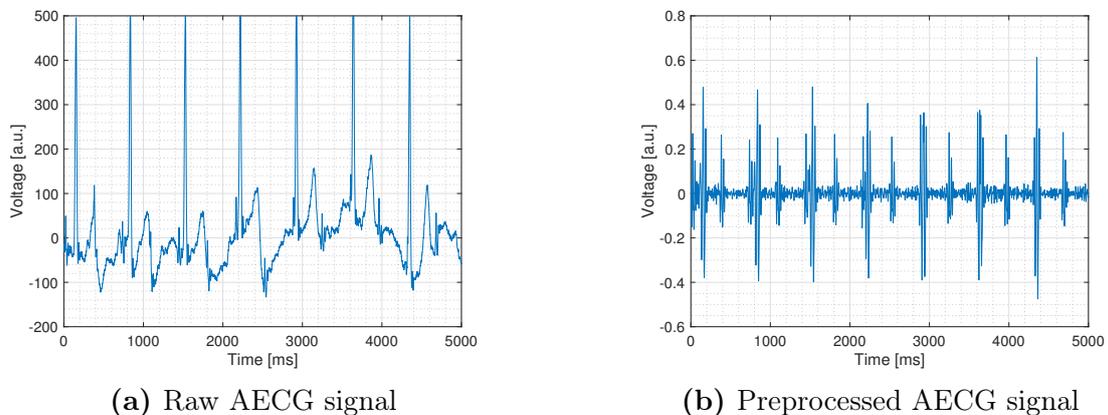


Figure 4.1: A typical result of applying the preprocessing steps on an AECG signal.

4.1 Preprocessing

Preprocessing within the field of NI-fECG extraction has been investigated to a fair extent in earlier studies. For instance, as compared to adult ECG filtering, it is not unusual to apply a higher low frequency cutoff when solely considering fQRS extraction and not other morphological features[82, 93]. The digital preprocessing steps applied to the signals in this thesis followed the procedure in Behar[3], although

the values of the upper and lower cutoff frequencies were optimized (in the fQRS extraction sense) during the GA procedure, see Part III Section 4.3.4.

To remove baseline wander and irrelevant high frequency content on each of the abdominal channels, two zero phase Butterworth filters were applied: a third order high pass filter and a fifth order low pass filter. Normalization of the thoracic and abdominal channels was conducted by: 1) using the first five seconds of each recording to determine the amplitude range of the ECG and then divide the signal with this amplitude, 2) subtract the mean of the first five seconds from the signal, and 3) transform the signal through the hyperbolic tangent to avoid that reservoir states in the ESN assumed unexpected values due to outliers. The last, non-linear step may not be suitable for morphological studies, but improves the robustness of the fQRS detection procedures.

An example of the result of the preprocessing procedure on an AECG signal is showed in Figure 4.1.

4.2 No maternal chest reference

This section considers the case where the fECG extraction method has to rely completely on information from abdominal channels. Since no mECG reference signal is available, one might instead try to take advantage of spatial differences in the appearances of the two ECGs as governed by their respective ECG net dipole map (see Part I Section 2.4. The signals of all the abdominal channels can be mapped onto a multidimensional domain where each axis corresponds to one of the electrodes, before BSS (in this case different versions of ICA) is applied to find independent/periodic components of this signal.

4.2.1 Variants of ICA

Linear ICA is formulated as given an observation vector $\mathbf{x} = x_1, x_2, \dots$, find the mixing matrix \mathbf{A} and sources $\mathbf{s} = s_1, s_2, \dots$ that transforms

$$\mathbf{x} = \mathbf{A}\mathbf{s} \tag{4.1}$$

given some measure on the sources \mathbf{s} . In other words, the observations \mathbf{x} are constructed from a linear combination of sub-components from the sources, using a constant mixing matrix \mathbf{A} . This an underlying issue with linear ICA based methods. The observations and sources are related through a time dependent linear transform which can be seen in Figure 2.4. Large loops arise from the mECG and small loops from the fECG. A linear transform can not transform this data in such a way that the sources are independent.

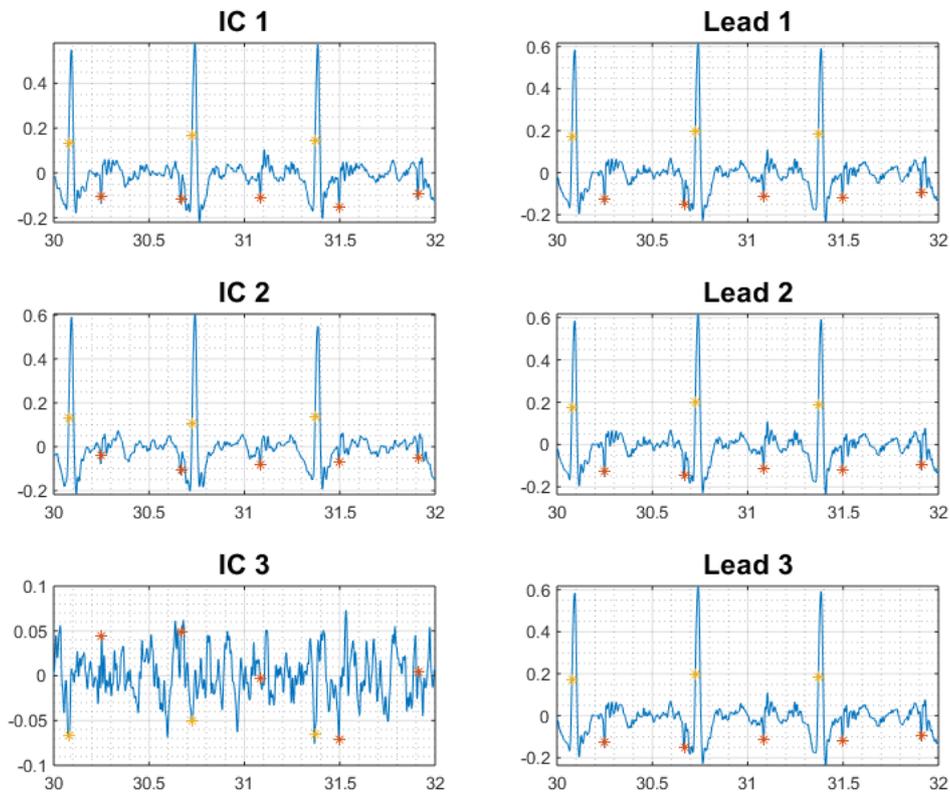


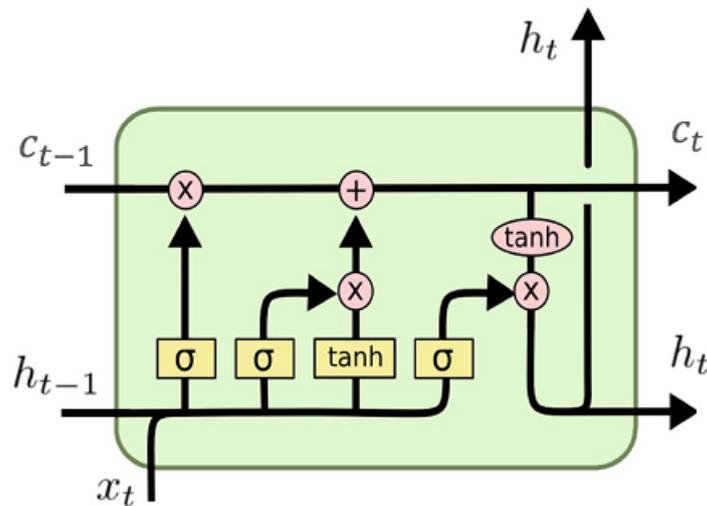
Figure 4.2: Typical independent components from abdominal electrodes. The figures to the left shows three independent components and the figures to the right shows three abdominal leads. The yellow stars are maternal QRS annotations and the red stars are fetal QRS annotations. Looking at the independent components, the ICA algorithm does not separate the fECG and mECG well.

Attempts were made using linear ICA on abdominal signals. A typical result is shown in Figure 4.2. On the left side are three independent components, IC1, IC2 an IC3. On the right side are the leads that were used in the ICA. As can be seen, the algorithm does not manage to separate the fECG and mECG well. However, it should be noted that the ICA was applied directly to preprocessed data, and better results can be obtained by first reducing the mECG through e.g. TS, see Part III Section 3.3. Nevertheless, due to this fundamental discrepancy between the problem at hand and the algorithm, linear ICA algorithms were deemed unappealing and thus disregarded in the choice of signal separation methods to evaluate further experimentally.

4.3 Maternal chest reference

Similarly to many temporal methods, see Part III Section 3.2, the following techniques will also rely on removal of the maternal ECG component within the AECG mixture. However, these methods will utilize the signal from a maternal thoracic (chest) reference electrode in doing so. The core principles are that the thoracic signal should be fECG free and that there is some relationship between the chest and abdominal mECG signal. Thus should the mapping of the thoracic signal that best matches the AECG correspond to the maternal ECG contribution in the latter, which then could be subtracted from the composite abdominal signal. This mapping could be performed in multiple ways. In this thesis two different kinds of recurrent neural networks are investigated; Echo state neural networks and LSTM neural networks.

4.3.1 The long short-term memory network



LSTM
(Long-Short Term Memory)

Figure 4.3: Figure showing the standard LSTM cell. c_t , x_t and h_t are vectors with cell state, input and output. The yellow boxes are neural network layers with sigmoid or tanh activation functions. Red circles are point-wise operations. The two sigmoid units to the left control the information flow to the cell state and are called the forget gate and external input gate. Data from the current input x_t and previous output x_t are presented to the input gate through a tanh activation layer. The sigmoid unit to the right is called the output gate and controls to what extent the current cell state is presented to the output.

A challenge with recurrent neural networks is solving the issue that gradients tend to either vanish or explode as they are back-propagated in time, making it difficult

to learn long time dependencies. This issue has been addressed by many authors and include units that skip layers and leakage units where the state of the unit is integrated over time[94]. One such networks is the LSTM network. This network has proven to be very successful in a range of applications, e.g. speech recognition[95] and machine translation[96], and are fit for sequence to sequence mapping. The standard LSTM unit is shown in Figure 4.3. In contrast to the regular recurrent network unit, Figure 1.3, the unit has two recurrent connections. These are the cell state c_t and output h_t at time t . The information in the cell state is regulated by two gates, the forget gate and the external input gate. The forget gate (left-most sigmoid unit in Figure 1.3) combines information about the previous output and current input to decide how much of the cell state that should be forgotten. The input gate (middle sigmoid unit) decides to what extent information presented through the tanh activation layer should be added to the cell state. Finally, the output gate (right-most sigmoid unit) regulates what information that should be presented to the output from the cell state. Using this architecture, the network can learn when to forget information and when to keep information, and to learn long term dependencies in data[94].

The state $c_i^{(t)}$ of the LSTM unit is updated according to

$$c_i^{(t)} = f_i^{(t)} c_i^{(t-1)} + g_i^{(t)} \sigma \left(b_i + \sum_j U_{ij} x_j^{(t)} + W_{ij} h_j^{(t-1)} \right), \quad (4.2)$$

where f_i is the state of the forget gate, g_i the state of the external input gate, $x_j^{(t)}$ the current input, b_i is the bias and U_{ij} is the connection between other LSTM units in the net. The forget gate is updated according to

$$f_i^{(t)} = \sigma \left(b_i^f + \sum_j U_{ij}^f x_j^{(t)} + W_{ij}^f h_j^{(t-1)} \right) \quad (4.3)$$

where h_j is the state of all other LSTM units in the network. The activation of the external input gate

$$g_i^{(t)} = \tanh \left(b_i^g + \sum_j U_{ij}^g x_j^{(t)} + W_{ij}^g h_j^{(t-1)} \right). \quad (4.4)$$

the output of the unit

$$h_i^{(t)} = \tanh(s_i^{(t)}) q_i^{(t)}. \quad (4.5)$$

where $q_i^{(t)}$ is the activation of the external input gate, which is updated as

$$q_i^t = \sigma \left(b_i^o + \sum_j U_{ij}^o x_j^{(t)} + W_{ij}^o h_j^{(t-1)} \right). \quad (4.6)$$

As the LSTM network can learn long time dependencies in the data it might be appropriate for tracing the transformation between thoracic and abdominal electrode in time.

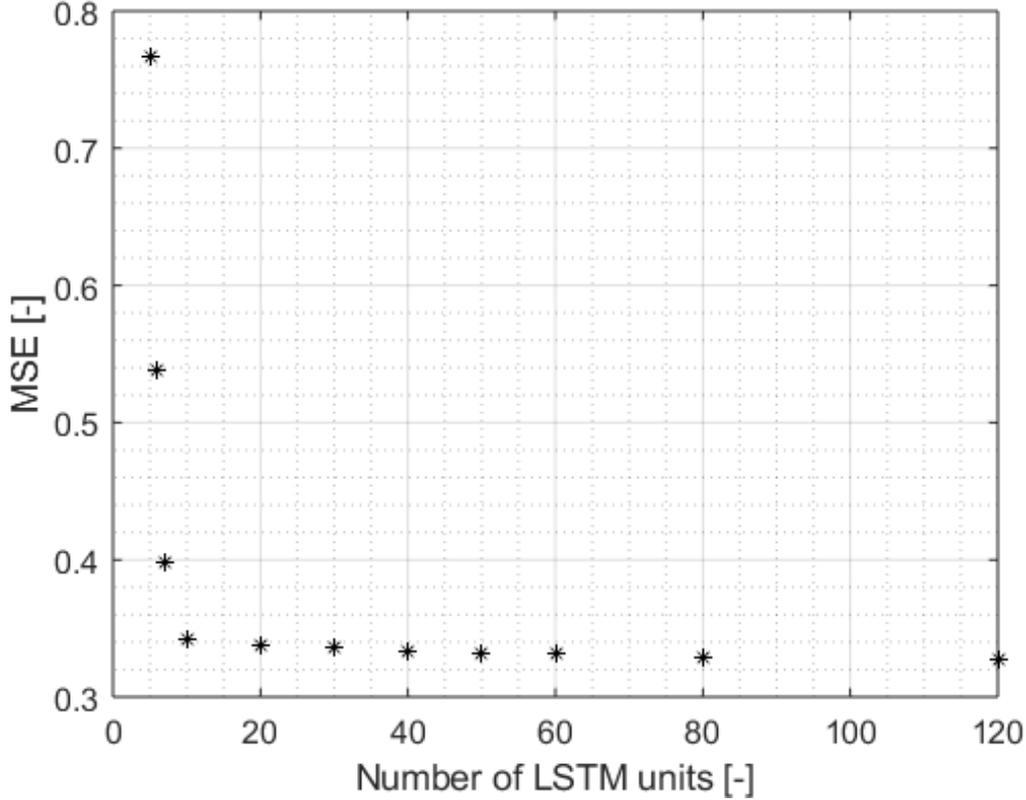


Figure 4.4: Figure showing the MSE between the LSTM prediction and an abdominal lead as a function of the number of units in the LSTM network. The number of units govern how much information that is remembered between each time step. As seen, the networks reaches a small MSE at only 10 units.

4.3.2 Application and optimization of LSTM network

A simple LSTM network was implemented with one input, a thoracic lead, connected to a dropout layer and fully connected layer with one output, an abdominal lead. In a dropout layer, input elements are randomly set to zero with a given probability. This effectively modifies the underlying network architecture between iterations and is used to prevent the network from over-fitting. The mean squared error (MSE) between the abdominal lead and network output was used as loss function. In this way, the networks removes as much as possible from an abdominal lead using a thoracic lead, which contains no information from the fECG. If the network does not over-fit, the prediction at the abdominal lead should contain the contribution from the maternal heart and no contribution from the fECG. The networks were trained on the first 10 seconds and validated on the remaining 80 seconds for each file. The networks were trained until the validation error reached a minimum and the adaptive moment estimation (Adam) optimization algorithm was used.

The first file in the data-set was used to find an appropriate number of units, an intrinsic property of the LSTM cell governing how much information that is remem-

bered between each time step. The number of units and the final validation MSE is shown in Figure 4.4. As can be seen, the network learns to approximate the transform well with as little as 10 units. However, the validation error keeps reducing and have not reached a minimum at 120 units. Unfortunately, the time it takes to train the network increases substantially with more units. 60 units were used as a compromise between time and performance. Each training session was completed in approximately 11 min on an Intel(R) Core(TM) i5-4690 3.5 GHz CPU.

4.3.3 Echo state neural networks

Echo state neural networks, ESNs, are a fairly recent subclass of recurrent neural networks. Originally introduced in the beginning of this millennium by Jaeger[97], ESNs quickly proved to offer a very practical approach to understanding and training of RNN networks. One of their key merits was that they showed that RNNs can perform tasks in an efficient and accurate manner without training all network weights. For a classic ESN, the RNN part (denoted reservoir) is generated randomly and sparsely, and merely the output from the reservoir is trained. The reservoir constitutes the dynamical system that maps the system input to a higher dimension[98]. As a side note, this concept was first introduced in a neuroscientific model of the corticostriatal processing loop[99].

While being both computationally inexpensive and conceptually simple, ESNs still achieved superb performance levels in various benchmark tasks within many different fields such as finance, energy conservation, and speech recognition, see for instance [100, 101, 102, 103].

Although being ostensibly simple, it can be difficult to apply ESNs in a successful manner. In particular, the initial construction of the reservoir is governed by multiple hyperparameters that have to be chosen sagaciously.

The versatile and computationally cheap non-linear modelling of dynamical systems offered by ESNs has attracted attention also within the field of fECG extraction. In fact, some of the organizers of the PhysioNet Computing in Cardiology Challenge 2013 implemented a version of it that achieved promising results[104]. Obtaining an F1 score of around 97 %, the ESN appeared to be quite successful in removing the mECG contribution from the abdominal mixture. However, as mentioned above, determining the optimal values for the numerous global parameters in an ESN is not a trivial task, especially since the size of the total search space rules out attempts of exhaustive, multidimensional searches. In the competition entry random and grid searches were tested, where the latter had to be done in an iterative manner with most parameters kept fixed at presumptively reasonable values while one or two parameters were optimized at a time. In this thesis, an echo state neural network based upon the competition entry and a seminal ESN tutorial written by the originator Jaeger[97] is developed, where the hyperparameters are optimized using an evolutionary algorithm with F1 as a fitness function. The intention of the choice of optimization method is to successively drive the parameter values towards various local maxima of F1 in order to investigate most interesting regions in the total

search space in a fraction of the iterations required to cover them with a random or a multidimensional grid search.

The purpose of the ESN is to find the (potentially non-linear) mapping of the input function $u(n)$ (the thoracic signal) that best matches the target signal $y(n)$ (the AECG), where n denotes the time step. The reservoir, passively excited by the input signal, acts as a memory and offers temporal context by preserving a nonlinear transformation of the input history in its state vector. Propagating through the nonlinear reservoir, the chest recording gives rise to a high-dimensional dynamical “echo response” from the former. The response is regarded as a projection of the thoracic signal onto a set of non-orthogonal basis functions which are used to reconstruct the desired output. [105]. The output signal is obtained by passing the reservoir output through a linear readout layer, which maps the reservoir states to the desired output (the AECG) through weights trained by linear regression.

The network architecture consisted of K inputs that shaped the input signal $\mathbf{u}(n) = [u_1(n), u_2(n), \dots, u_K(n)]$, M internal units composing the reservoir state vector $\mathbf{x}(n) = [x_1(n), x_2(n), \dots, x_M(n)]$, and L output units forming the output signal $\hat{\boldsymbol{\eta}}(n) = [\hat{\eta}_1(n), \hat{\eta}_2(n), \dots, \hat{\eta}_L(n)]$. Commonly, the internal units in ESNs are updated according to

$$\mathbf{x}(n+1) = \mathbf{x}(n) + f(\mathbf{W}\mathbf{x}(n) + \mathbf{W}_i\mathbf{u}(n+1) + \mathbf{W}_b\hat{\boldsymbol{\eta}}(n)),$$

where $\mathbf{W} \in \mathfrak{R}^{M \times M}$ makes up the reservoir weight matrix and $\mathbf{W}_i \in \mathfrak{R}^{M \times K}$ constitutes the input weight matrix. Both are randomly generated initially and kept fixed. $\mathbf{W}_b \in \mathfrak{R}^{M \times L}$ is the back projection weight matrix that links previous network outputs to future internal unit states and f is the activation function of the reservoir neurons, taken to be the hyperbolic tangent. This is a typical and practical choice since it is a strictly increasing bijection from \mathfrak{R} to $(-1, 1)$.

However, considering the dynamical pattern recognition task at hand as purely input-driven (the processed version of the mECG signal in the network output should not have any impact on future reservoir states), the rule may in this case be simplified by setting \mathbf{W}_b to zero. Furthermore, due to the impulse-based nature of the input signal, it makes sense to limit the temporal memory span of the internal states. In other words, very old states should not have a significant impact on future ones. This can be achieved with the leaky integrator neuron model[97], which includes a leakage rate (forgetting factor) $\alpha \in [0, 1]$. Implementing these changes, the new update rule becomes

$$\mathbf{x}(n+1) = (1 - \alpha)\mathbf{x}(n) + f(\mathbf{W}\mathbf{x}(n) + \mathbf{W}_i\mathbf{u}(n+1)). \quad (4.7)$$

As seen in (4.7), $\alpha = 1$ would mean that the neurons would not retain any information about their earlier states, whereas $\alpha = 0$ would return the ordinary update rule (except for the \mathbf{W}_b term).

The network output, in turn, is calculated through

$$\hat{\boldsymbol{\eta}}(n+1) = g(\mathbf{w}_o(n)\mathbf{z}(n)), \quad (4.8)$$

where $\mathbf{z}(n) = [\mathbf{x}(n), \mathbf{u}(n)]$ is the extended system state and g is the output activation function (also set to the hyperbolic tangent). Note how the extended system state enables the input states to have a direct impact on the network output without going through the RNN (the reservoir). The extent of the impact of specific elements within $\mathbf{z}(n)$ is governed by the weights of the readout layer, $\mathbf{w}_o(n)$. These weights may be adaptive, but are kept fixed in this ESN adaptation.

\mathbf{W} and \mathbf{W}_i are both initialized sparsely from a uniform distribution on the interval $[-1, 1]$, where the sparsity parameter $\psi \in [0, 1]$ determines the fraction of matrix elements to be non-zero. However, \mathbf{W} is then scaled to a spectral radius ρ in the range $[0, 1]$, which has been shown empirically to more or less guarantee the existence of echo states. The echo state property basically states that for every internal signal $x_i(n)$ there exists an echo function e_i that maps input/output histories to the current state. This property has been shown to be crucial for the function and performance of ESNs[97].

For each 90 second recording in the data set (and each abdominal channel separately), the first 30 were used for training the readout layer weights. This was done through linear regression, which aimed to find the vector \mathbf{w}_o that minimized the difference between the RHS of (4.8) (the predicted maternal contribution to the AECG) and the AECG itself. In the rare cases where a solution could not be found, the beginning of the training span was moved by one second until convergence.

Multiple parameters have significant effects on the performance of the network. For instance, an excessive number of internal units M increases the risk of overfitting, i.e. it impairs the networks ability to generalize to slightly varying inputs. In addition to ensuring the echo state property by a very high probability, spectral radius ρ influences the time scale over which an earlier excitation keeps “echoing” in the reservoir; a small ρ makes the output more dominated by recent inputs[97].

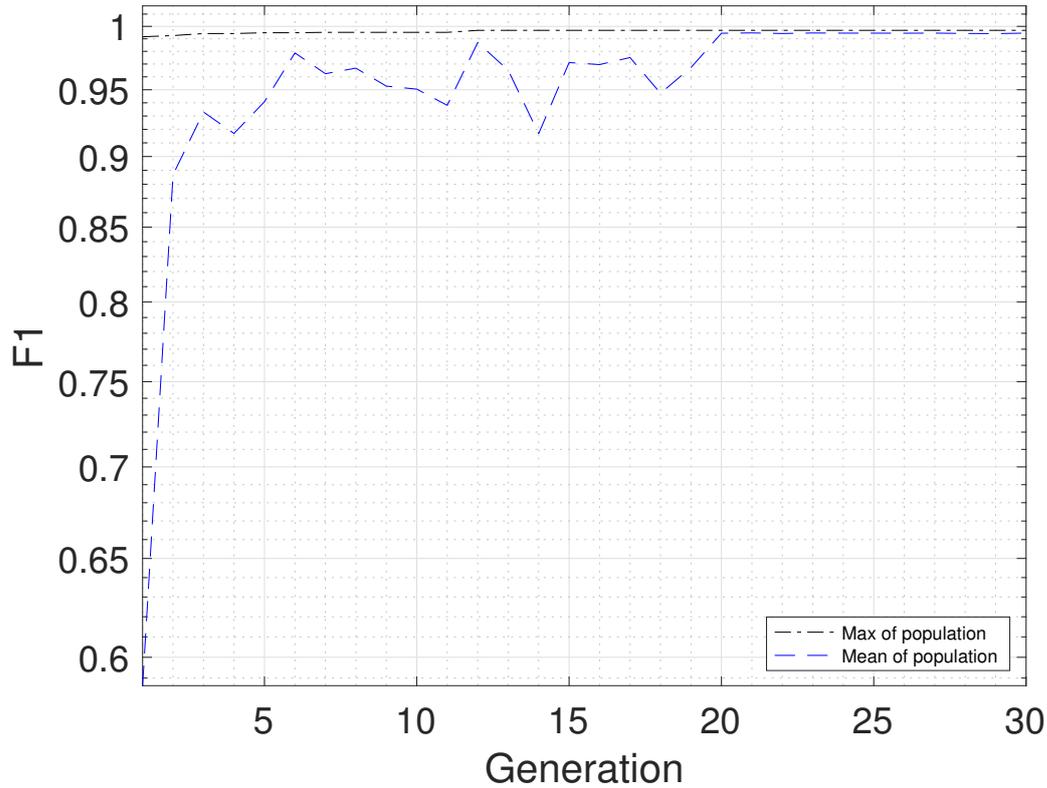


Figure 4.5: The mean and maximum F1 values of the population for each generation. Both measures had seemingly converged properly after the 30 generations that the GA was run for. The fact that the mean F1 value did not converge immediately indicates that the chosen values of the crossover and tournament operators did not favor stronger individuals excessively, which could have got the whole population stuck in a local maximum almost instantly.

4.3.4 Optimization of ESN performance using a genetic algorithm

To optimize the ESN, a genetic algorithm was implemented that used the obtained F1 score, see Part III Section 2.2, of 10 of the 14 recordings in the database (randomly selected initially and then kept fixed for all evaluations) as fitness measure. Since the main goal of the signal processing endeavor in this thesis is to develop an algorithm with a very high beat detection reliability, the F1 score with a 50 ms acceptance tolerance was used. The continuous wavelet transform method described in Part III Section 5.2 was used for QRS detection. As previously mentioned, the first 30 seconds of every recording was used for training of the ESN readout layer, and the last 60 seconds for fECG extraction. Consequently, only the latter part of each recording was used for the fitness measure. The used GA parameters are listed in Table 4.1. The total search space and the obtained optimized configuration of the variables are specified in Table 4.2. Note that the lower and upper cutoff frequency are part of the preprocessing applied to the AECG before it is fed to the ESN.

For the best performing set of variables, an F1 score of 0.9968 was acquired for the training data. Figure 4.5 displays the convergence procedure of the mean and maximum F1 values of the individuals of each generation. Considering the successive convergence of the average F1 value, it seems like the values of the crossover and tournament operators did not impede initial exploration of the search space to an excessive extent.

To check whether substantial overfitting had occurred, the remaining four recordings of the database were used as a test set containing previously unseen data. The test set comprised $\sim 29\%$ of the total data. The F1 score on the test set for the optimized variables was 1, i.e. there were no misdetections at all on the test set according to the used figure of merit.

Parameter	Value
Population size	40
Number of genes	60
Crossover probability	0.4
Mutation probability	1/60
Tournament selection parameter	0.75
Tournament size	3
Number of variables	6
Number of generations	30
Number of copies of best individual	1

Table 4.1: List of used parameter values in the GA.

Variable	Search range	Optimized value
Lower cutoff frequency, f_b	0-40 Hz	26
Upper cutoff frequency, f_h	47-142 Hz	104
Leakage parameter, α	0-0.8	0.7554
Neurons in reservoir, M	45-225	149
Sparsity parameter, ψ	0.05-0.25	0.1851
Spectral radius, ρ	0.2-1	0.4604

Table 4.2: List of preprocessing and ESN parameters that were optimized by the GA and their respective search ranges and obtained optimized values.

Chapter 5

Experimentally evaluated methods - fQRS extraction

Once the mECG (presumably) had been removed from the AECG through an fECG extraction procedure, two different fQRS extraction methods were tested to acquire the fHR. The Pan-Tompkins (PT) method[106] was developed in the 1980s to detect adult QRS complexes and performs digital analyses of slope, amplitude, and width. A simple adaptation of it, based upon the one done by Behar[3], was considered here. The second method, a continuous wavelet transform based QRS extraction method, was developed as part of this thesis and looks for impulses of frequency packets with certain characteristics.

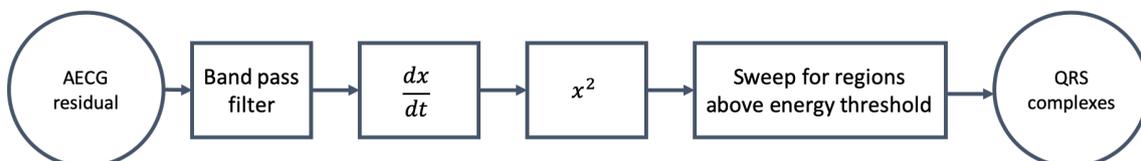


Figure 5.1: A block diagram describing the main features of the Pan-Tompkins QRS detection algorithm, as applied to an AECG where the maternal contribution has been removed.

5.1 Pan-Tompkins

In short, the original PT algorithm begins with a band pass filter to reduce noise and then takes the derivative of the signal to enhance the R wave slopes. Thereafter the square of the signal is computed to further emphasize the high frequency characteristics of the QRS complex and convert all data points to positive values. Finally, a window (with a size slightly greater than or equal to the widest QRS complex) is swept along the signal looking for regions above a specified energy threshold. A block diagram of these main features of the algorithm is displayed in Figure 5.1.

In the fQRS adaptation used in this thesis, a refractory period of 150 ms was used. Given the detected energy peaks, the position of each R-peak was adjusted so that the sign of all R-peaks was all either positive or negative (based upon the average of the local optima of each high energy segment). This was performed in order to avoid undesirable fHR changes as a result of some R-peaks being negative and some being positive. It is also crucial if one would like to use TS based methods for the fECG extraction, since construction of a template mECG cycle to subtract from subsequent cycles entails same sign peaks.

In this thesis, Pan-Tompkins was applied and considered separately for each of the electrodes and no efforts were made to combine the detections of the different electrodes into a combined detection based on e.g. detection similarities, differences, and some electrode reliability measure.

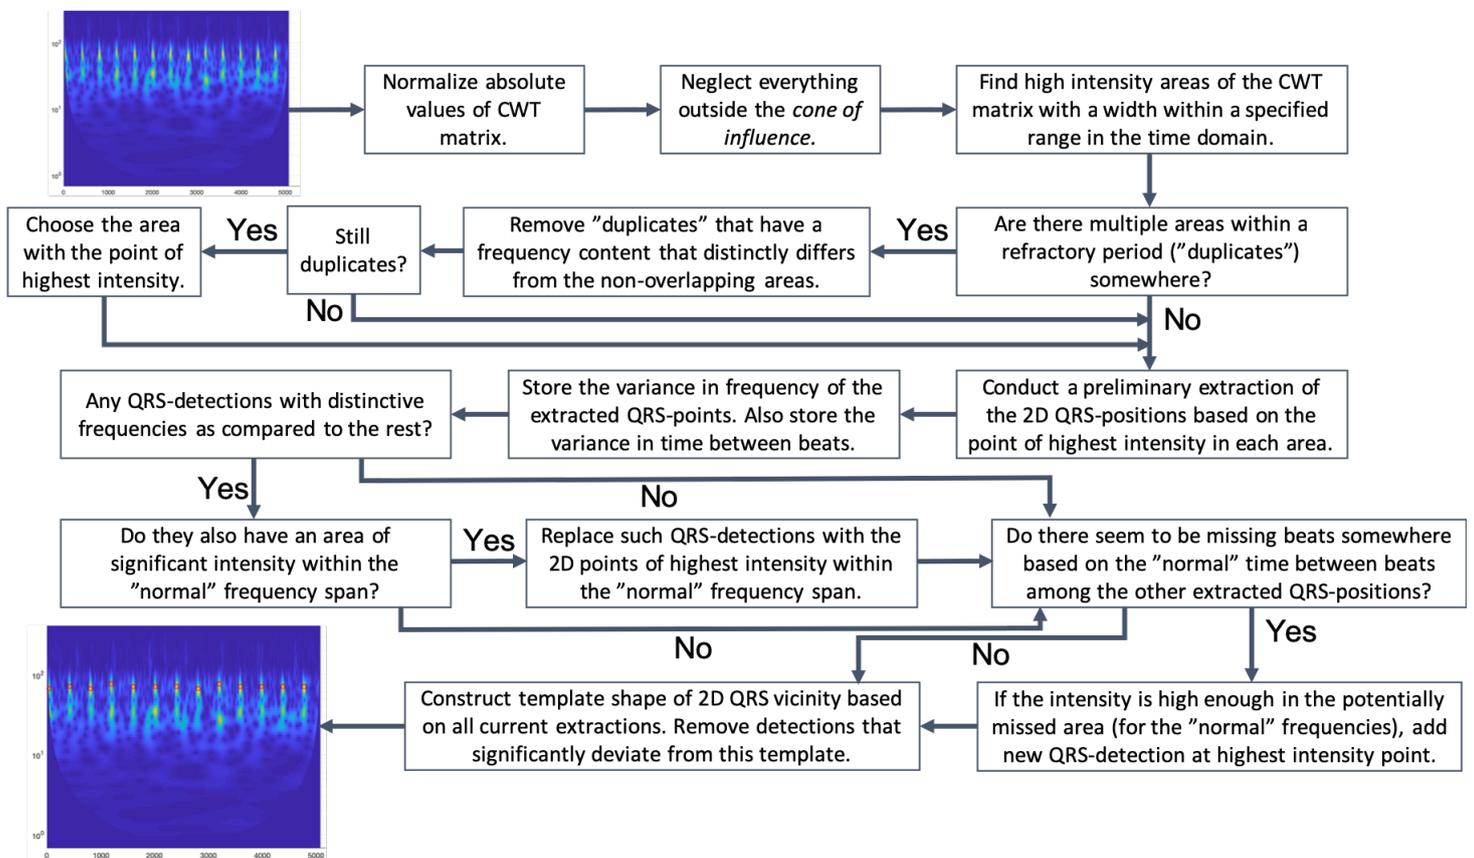


Figure 5.2: Flowchart of the algorithm that makes predictions of the QRS peak positions (the red circles) from a CWT of a signal.

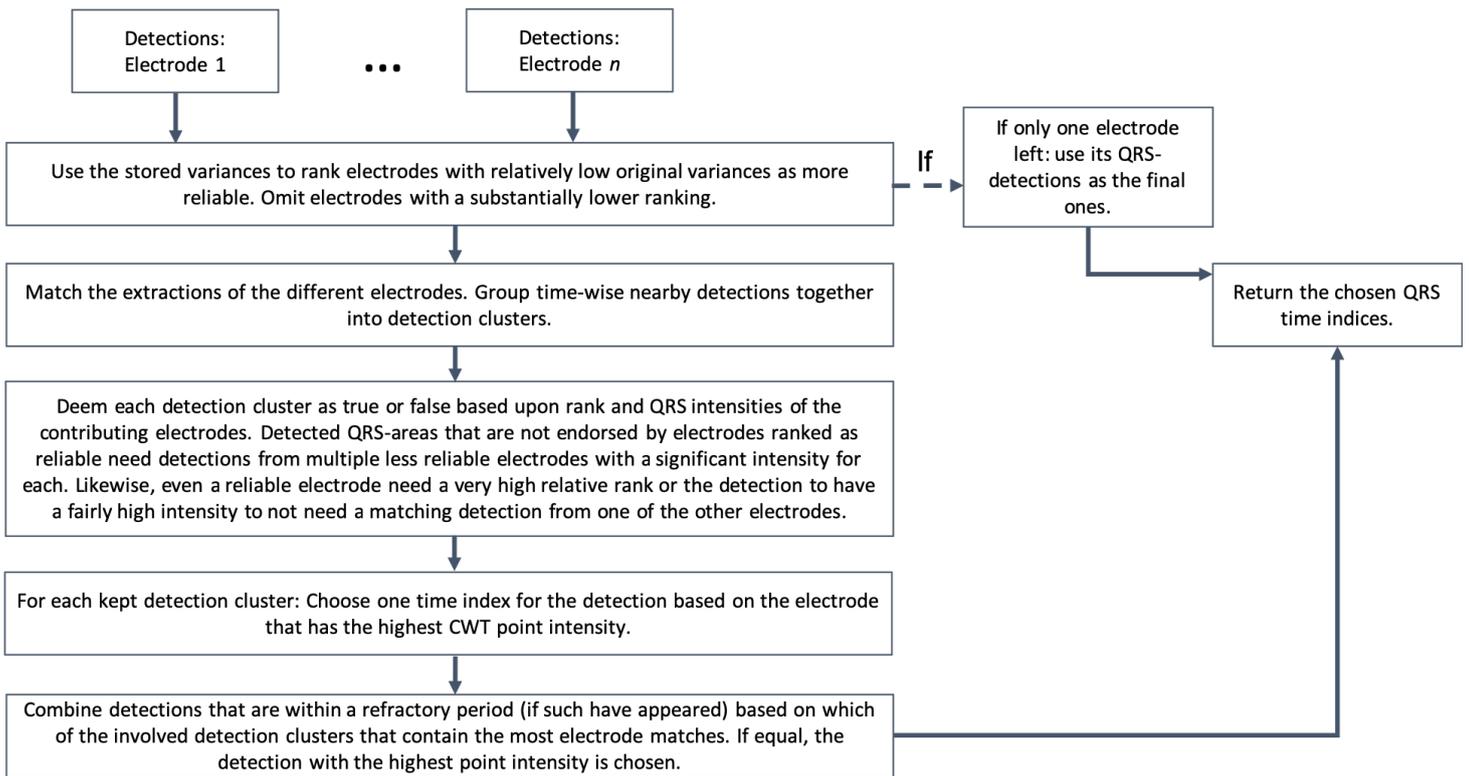


Figure 5.3: Flowchart of the algorithm that combines predictions of the QRS peak positions of multiple channels into a final set of predictions for the given time period.

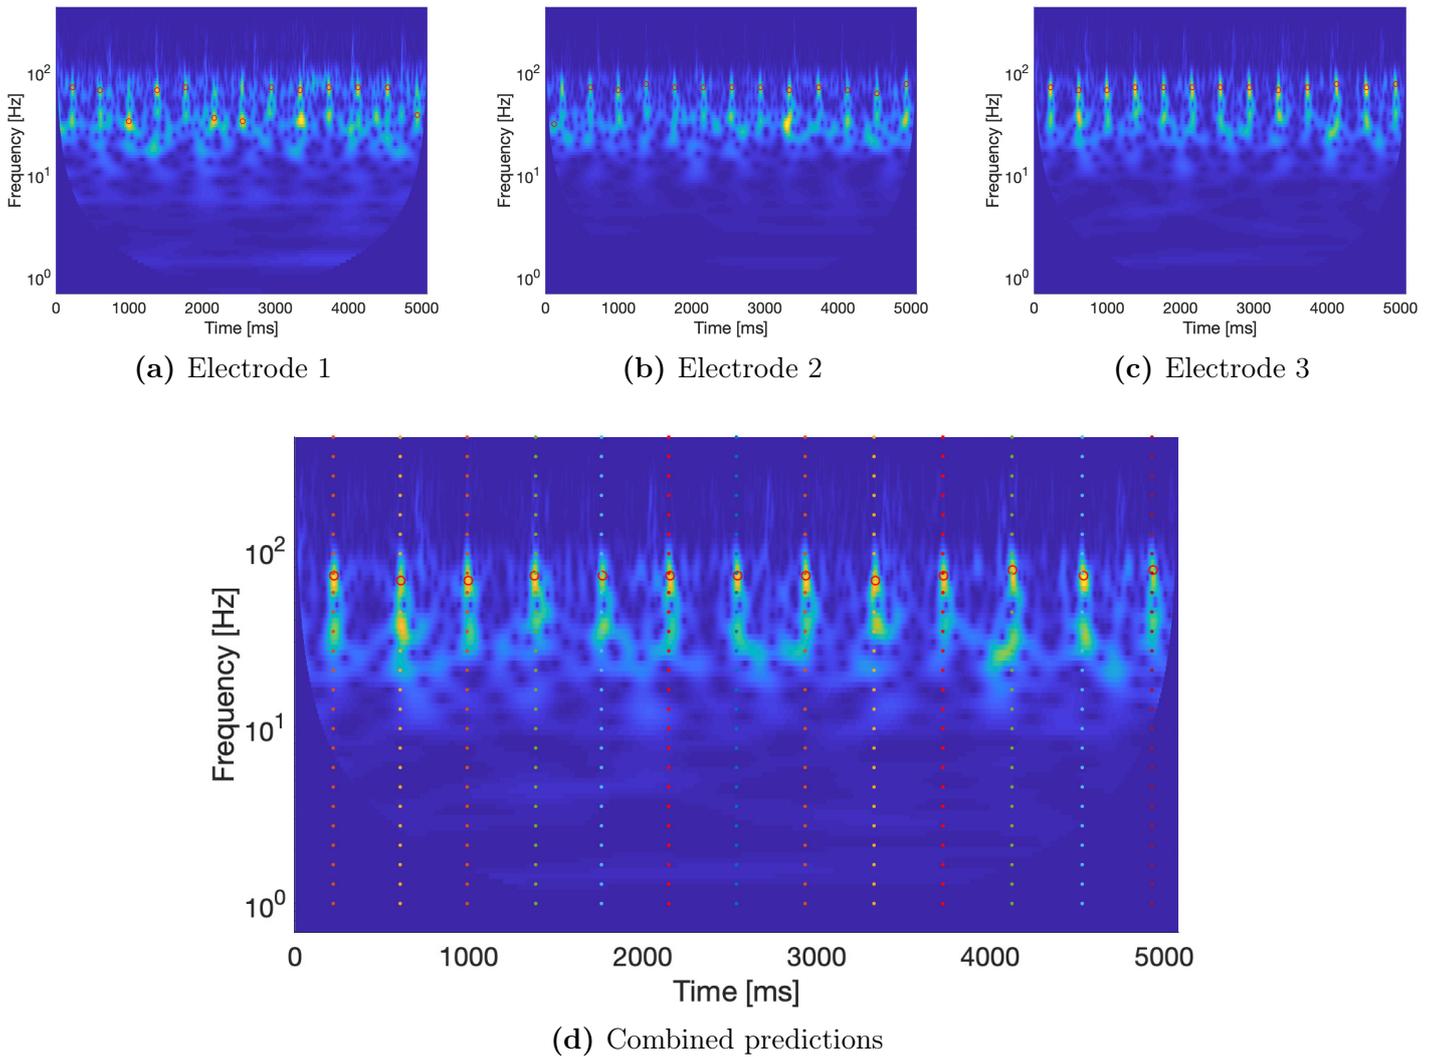


Figure 5.4: a) to c) show an example of QRS peak predictions (red circles), based upon the procedure described in Figure 5.2, made on CWTs of the signals from the three abdominal channels for a chosen five second period with an erroneous prediction (the leftmost in b)). The combined, final predictions based upon the predictions for each channel and the algorithm described in Figure 5.3 are shown in d). The dotted vertical lines denote the corresponding annotations that are used as reference. Note how the erroneous prediction from the second electrode has been omitted in the combined version because the algorithm deemed it as unreliable.

5.2 QRS detection from a scalogram of a continuous wavelet transform

After the removal of the mECG component in each of the three abdominal electrode channels of a recording, roughly five seconds of the AECG residual of each electrode are transformed by a CWT at a time and plotted in a scalogram. This simulates an online setup, where only the signal during the last couple of seconds should have an

impact on the momentary beat-to-beat detection. Once transformed, QRS detection is performed on each channel through a series of steps summarized in Figure 5.2. In short, the algorithm searches for high-intensity regions of the CWT scalogram, whilst taking e.g. size, shape, and frequency content of the region and the refractory period of the fetal heart into account. The QRS “predictions” of each channel are then combined to form the final set of predicted QRS peak times. The combination is based on ranking of the electrodes, intensity maxima in the respective delineated regions of each channel, and whether multiple electrodes make the same prediction and is described briefly in Figure 5.3. The reliability ranking system of the electrodes is based on the relative variances in frequency and in time between the preliminary predicted QRS positions in the CWT. The variance values are calculated before most of the steps intended to manage erroneous/missing detections have been applied, thus supposedly conveying a flavor of how prominent the important features of the CWT of the signal from the respective electrode are.

An example of the tentative and final results of the CWT-based QRS extraction procedure can be seen in Figure 5.4.

Chapter 6

Results

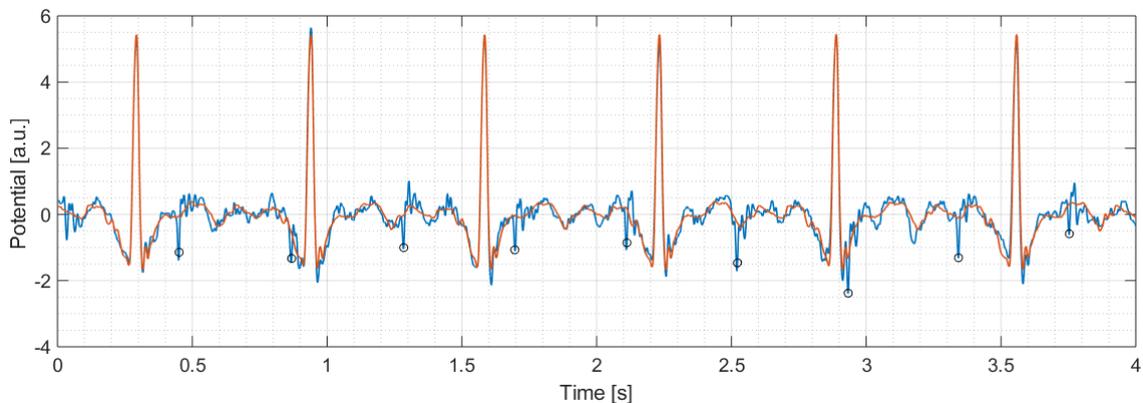


Figure 6.1: Figure showing the potential from an abdominal lead (blue) and the LSTM mapping from thoracic to abdominal electrode (red). The red circles show annotations for the fQRS complexes. The maternal R peaks (i.e. the high peaks) are more or less completely mirrored by the mapping (the tip of the blue underlying line can be seen slightly for the second peak). Meanwhile, as desired, the fetal QRS complexes are not included in the mapping.

6.1 Removal of the mECG from the AECG mixture

The task of removing the mECG contribution from the abdominal mixture appeared to be solved quite successfully by both LSTM and ESN. An example can be found in Figure 6.1, which shows the measured potential at an abdominal lead and the LSTM mapping from the thoracic to the abdominal electrode. It also includes annotations of the fQRS positions. Note e.g. how the maternal R peaks (i.e. the high peaks) are more or less completely mirrored by the mapping, whereas the fetal QRS complexes are not included in the mapping.

6.2 LSTM and ESN in combination with Pan-Tompkins

This section presents the results of the LSTM and ESN networks when the Pan-Tompkins implementation was used for fQRS detection.

6.2.1 LSTM

The positive predictive value and sensitivity for all test files and electrodes are shown in Table 6.1 and 6.2. The average predictive value was 0.9847 and average sensitivity 0.9723. The corresponding F1 measure is 0.9785.

File \ Electrode	1	2	3
154	0.9437	0.9864	0.9932
192	0.9933	1.0000	0.9932
244	0.9934	0.9935	0.9934
274	0.9379	0.7877	1.0000
290	1.0000	0.9793	0.9862
323	0.9936	0.9936	1.0000
368	0.9797	0.9797	0.9795
444	1.0000	1.0000	1.0000
597	0.9871	1.0000	0.9870
733	0.9873	0.9935	0.9936
746	0.9932	0.9932	1.0000
811	0.9935	0.9871	0.9868
826	0.9871	0.9871	0.9872
906	0.9926	1.0000	1.0000

Table 6.1: Positive predictive value for the case of LSTM in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9847.

File \ Electrode	1	2	3
154	0.9054	0.9797	0.9797
192	0.9933	0.9933	0.9800
244	0.9869	0.9935	0.9869
274	0.8531	0.6497	0.9548
290	1.0000	0.9660	0.9728
323	1.0000	1.0000	1.0000
368	0.9864	0.9864	0.9728
444	1.0000	1.0000	1.0000
597	0.9871	1.0000	0.9806
733	0.9750	0.9563	0.9750
746	0.9932	0.9932	1.0000
811	0.9935	0.9871	0.9613
826	0.9745	0.9745	0.9809
906	0.9926	1.0000	1.0000

Table 6.2: Sensitivity for the case of LSTM in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9723.

6.2.2 ESN

The positive predictive value and sensitivity for all test files and electrodes are shown in Table 6.1 and 6.2. The average predictive value was 0.9635 and average sensitivity 0.9667. The corresponding F1 measure is 0.9651.

File \ Electrode	1	2	F1
154	0.9122	0.9669	0.9726
192	0.9671	0.9548	0.8675
244	0.9934	0.9935	0.9934
274	0.7102	0.5682	0.9886
290	1.0000	0.9726	0.9865
323	1.0000	0.9936	0.9935
368	0.9799	0.9864	0.9730
444	1.0000	1.0000	1.0000
597	1.0000	0.9936	1.0000
733	0.9810	0.9554	0.9936
746	0.9866	0.9866	0.9932
811	0.9808	0.9872	0.9623
826	0.9565	0.9500	0.9809
906	0.9853	1.0000	1.0000

Table 6.3: Positive predictive value for the case of ESN in combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9847.

File \ Electrode	1	2	F1
154	0.9122	0.9865	0.9595
192	0.9800	0.9867	0.9600
244	0.9869	0.9935	0.9869
274	0.7062	0.5650	0.9774
290	1.0000	0.9660	0.9932
323	1.0000	1.0000	0.9935
368	0.9932	0.9864	0.9796
444	0.9938	0.9812	0.9875
597	0.9935	1.0000	0.9935
733	0.9688	0.9375	0.9688
746	0.9932	0.9932	0.9932
811	0.9871	0.9935	0.9871
826	0.9809	0.9682	0.9809
906	0.9853	1.0000	1.0000

Table 6.4: Sensitivity for the case of ESNin combination with Pan-Tompkins. Mean over all electrodes and used files is 0.9723.

6.3 LSTM and ESN in combination with a CWT based fQRS detection

This section presents the results of the LSTM and ESN networks when the CWT based method was used for fQRS detection.

6.3.1 LSTM

The PPV, sensitivity, and F1 score for all test files are shown in Table 6.5. The average predictive value was 0.9939 and the average sensitivity 0.9835. The corresponding F1 measure is 0.9885.

File	PPV	Se	F1
154	0.9801	1.0000	0.9900
192	1.0000	1.0000	1.0000
244	0.9869	0.9869	0.9869
274	0.9888	0.9944	0.9916
290	0.9932	0.9932	0.9932
323	1.0000	1.0000	1.0000
368	0.9797	0.9864	0.9830
444	1.0000	0.9938	0.9969
597	1.0000	1.0000	1.0000
733	1.0000	0.9750	0.9873
746	1.0000	1.0000	1.0000
811	1.0000	0.9290	0.9632
826	0.9862	0.9108	0.9470
906	1.0000	1.0000	1.0000
Mean	0.9939	0.9835	0.9885

Table 6.5: Table showing the positive predictive value, sensitivity and F1 measure for LSTM as the fECG extraction technique and the CWT method for the fQRS detection.

6.3.2 ESN

The PPV, sensitivity, and F1 score for all test files are shown in Table 6.6. The average predictive value was 0.9976 and the average sensitivity 0.9968. The corresponding F1 measure is 0.9972.

File	PPV	Se	F1
154*	1.0000	1.0000	1.0000
192*	1.0000	1.0000	1.0000
244	0.9869	0.9869	0.9869
274	1.0000	1.0000	1.0000
290	1.0000	1.0000	1.0000
323	1.0000	1.0000	1.0000
368	0.9865	0.9932	0.9898
444	1.0000	1.0000	1.0000
597	1.0000	0.9935	0.9968
733	1.0000	0.9938	0.9969
746	1.0000	1.0000	1.0000
811*	1.0000	1.0000	1.0000
826	0.9936	0.9873	0.9904
906*	1.0000	1.0000	1.0000
Mean	0.9976	0.9968	0.9972

Table 6.6: Table showing the positive predictive value, sensitivity and F1 measure for ESN as the fECG extraction technique and the CWT method for the fQRS detection. Files marked with * were part of the test set, and thus not included in the GA optimization procedure. Despite being previously unseen by the network, flawless performance was obtained for these files.

6.4 Performance with ± 5 BPM condition

As mentioned in Part III Section 2.2, to comply with industry standards a heart rate tolerance of ± 5 BPM would be desirable. Due to issues discussed in that section, this measure was not prioritized in this thesis. However, the current performance of the different techniques, abiding by this tolerance level, can be seen in Table 6.7.

Method	<i>Se</i> (%)	<i>PPV</i> (%)	<i>F1</i> (%)
LSTM - PT	0.7493	0.7570	0.7693
LSTM - CWT	0.8781	0.8691	0.8736
ESN - PT	0.7690	0.7664	0.7677
ESN - CWT	0.8808	0.8815	0.8812

Table 6.7: The positive predictive value, sensitivity and F1 measure for the different fECG and fQRS extraction method combinations when the rule of a maximum heart rate difference of ± 5 BPM is applied as tolerance during the grading procedure. The evaluation was performed on the full database.

Chapter 7

Discussion

Both the LSTM and ESN in combination with the CWT achieved excellent detection percentages on the database considered in this thesis, even though only three abdominal electrodes were used (the electrode combination part of the CWT based method is expected to improve with the number of electrodes). For the former, around 1 in every 100th heartbeat was misdetected and even fewer misdetections were made for the ESN. That level of reliability should be more than enough in an online implementation, since a single, rare error is quickly corrected by the subsequent beats. Furthermore, efforts to define and implement metrics that give a confidence measure based upon the detection performances have been initiated, although further work is required within this subject.

However, the data used in this thesis, although of varying signal quality, is not sufficient for a general claim about the reliability of the methods to be made. This since it merely contains recordings of fairly healthy fetuses and does not include samples of the various abnormal and potentially very dangerous conditions that might occur during gestation or labor. Consequently, one might conclude that the developed methods would perform well during normal conditions, but more and different data would be required in order to prove that reliable detection also would occur during abnormal circumstances.

The Pan-Tompkins method did not fully achieve the same level of accuracy as the CWT method, but it is important to emphasize that the electrodes were treated independently for this technique and their detections were not combined in any way to allow the final detections to be based upon a weighing of the information from multiple sources. Future methods that enable e.g. electrode ranking based upon information in the time domain could potentially increase this performance, although this was not prioritized in this thesis.

In general, even though Pan-Tompkins has achieved high performance levels (in particular for QRS extraction of adult ECGs), it is based upon a couple of empirical rules and operations. Supposedly there should be non-linear transforms and decision trees that outperform this manually developed method. Consequently, e.g. a regular FFNN or an RNN would presumptively be able to conduct an even better signal

pattern recognition, even in noisy areas, in order to detect the QRS complexes, provided that it is trained on enough data of various character. Personally the authors of this thesis believe that such a detection method, if applied after the CWT algorithm has detected the regions of each heart beat, would be able to set accurate time stamps of each beat in a very reliable way. This provided that a sufficient amount of data with correct and accurate annotations is available, which unfortunately was not the case during this thesis. Multiple examples were found where the exact positions of the annotations seemed inconsistent, even though they were placed in the right region. It has also previously been reported that some annotations of the manually annotated parts of the NI-FECG database should be discarded since inadequate accuracy has been found during visual inspections[3]. If data with correct and accurate annotations could be made accessible by using both abdominal electrodes and SECG (as reference) simultaneously, we believe that sufficient performance also for the ± 5 BPM condition should be within reach with a CWT-FFNN or an CWT-RNN combination.

Regarding robustness of the methods, the different algorithms were tested independently in this thesis. One could however imagine that usage of a combination of multiple QRS detectors relying on different detection methods (such as CWT and PT) potentially could improve the total extraction method's ability to handle bigger variations in the signal shape and also reliably detect declines in signal quality. This has not been evaluated in this thesis, but could turn out to be rewarding if investigated in the future.

Looking into improvability of the assessed methods, the weights of the readout layer of the ESN were kept fixed after the initial training. It is however also possible to allow them to evolve continuously in an online adaptation. A non-adaptive approach was chosen in this thesis for simplicity and because it actually performed slightly better when a comparison was made in the PhysioNet Computing in Cardiology Challenge 2013[104]. However, there is no guarantee that the same would be true for the ESN implementation developed in this thesis. Furthermore, there are multiple ways to implement such adaptive methods, so further investigation of adaptive output weight updates in the future could prove fruitful.

In general, the ESN and its hyperparameters appear to have been sufficiently optimized for the problem at hand, although it needs to be evaluated on more data of varying character to safely conclude that the hyperparameters do not need further optimization to improve the networks ability to generalize to new input patterns. In terms of computational expensiveness, no optimization has been performed with regard to this. However, full evaluation (training of the ESN and subsequent fQRS detection made through the CWT method with one thoracic and three abdominal electrodes) of the whole database, i.e. analysis of $14 \times 90 = 1260$ seconds of recordings, took on average around 150 seconds on a 3.1 GHz Intel Core i5 processor. It thus appears as if the current algorithmic efficiency is sufficient for online applications, even though it should still be optimized in order to reduce power consumption and possibly enable the computation to be integrated into an embedded system instead of an external computer.

The LSTM was not optimized in any way with respect to the F1 performance on the database. Instead, only the effects of the number of units and the training time was studied. By varying the low pass and high pass filters, which were fixed at 95 and 20 Hz respectively, the network might have achieved a better performance. Note that information regarding movement of the electrodes due to e.g. breathing was filtered out in the high pass filter.

Also, the networks were quite computationally heavy to train. A network with 60 units took approximately 11 min to train on the 10s signal. This would be significantly lower if a modern GPU was used instead of a CPU. However these can be quite expensive, which would revoke the benefit that the electrode system would be cheap. To resolve this issue, one could use a pre-trained network on specific electrode positions and slightly adapt the model for each specific case.

One could also imagine applying the LSTM architecture for a completely different purpose. During experimentation with the LSTM nets, it was noted that the net performed well at predicting the next value at an abdominal signal given the current value in a thoracic signal. This could serve as an improved and time adaptive model in e.g. a Kalman filtering network.

Revisiting ICA methods, it was concluded in Part III Section 4.2.1 that stand-alone linear ICA methods had intrinsic drawbacks that made it unappealing to the problem at hand. Furthermore, it had insufficient detection levels in the performance review in Part III Section 3.3. That being said, there have recently been several reports on ICA methods that utilize temporal information to separate subsignals using a non-linear transform, e.g. Hyvärinen[107]. Such extensions of regular ICA methods might be promising for separation of the fECG and mECG.

Another future improvement is to take pregnancies with more than one fetus into account. This thesis solely considers single pregnancies, but the prevalence of multiple pregnancies has been estimated to be around 3%[108] which cannot be deemed as negligible. The mECG removal methods of this thesis would probably still be successful, but the fQRS detection techniques would have to be adapted for such cases since the AECG residual still would comprise heart beats from multiple sources. Moreover, the refractory period condition could not be used in the same way as for the current methods. Intuitively, a BSS algorithm would probably be suitable to distinguish the sources based upon spatial information. As mentioned previously, the performance of a BSS based algorithm is expected to improve with the number of abdominal electrodes. A case study of multiple pregnancies has been conducted by Sameni[20], where different versions of linear ICA were applied for the fECG source separation after the mECG had been removed.

Chapter 8

Conclusion

Multiple fECG extraction and fQRS detection methods from previous works have been reviewed and compared in terms of performance, advantages, and disadvantages. Inspired by key ideas and findings from earlier works, novel methods have been developed that for the data available detect fetal heart beats with the same level of reliability as state-of-the-art QRS detectors do for adult ECGs. Merely one thoracic electrode and three abdominal electrodes were used during the evaluation, but the robustness and reliability of the methods is expected to improve even more for a higher number of electrodes. Additional measurement data with reliable annotations is required to work towards the desired resolution of ± 5 BPM for a commercial application. The performance of the developed methods also needs to be assessed on data that includes abnormal cardiac behavior.

Part IV

Conclusion and future work

This thesis has mainly consisted of two different parts. The first part looked into the development and assessment of a high impedance non-contact electrode, the second part into the development and assessment of fECG extraction and fQRS detection methods. Some of the highlights, key findings, and areas of further improvement are listed below.

- A non-contact electrode was designed and optimized with respect to noise, stability and resistance to artifacts.
- The electrode can reliably measure an ECG signal and should be able to measure the fECG with similar signal to noise ratios as state of the art contact electrodes.
- The electrodes are suitable for stationary implementation in stretchy fabric, which facilitates an easy application procedure, re-usability and long term fECG monitoring.
- A full measurement setup has been implemented and is ready to make recordings on the abdomen of pregnant women, although further improvements are still possible.
- The total production cost of an array of electrodes should be a fraction of the cost of current CTG apparatuses.
- In the future, input neutralization should be implemented to reduce noise and to make the electrode more stable with respect to variation in input capacitance. Also, changing the front end amplifier would most certainly reduce both noise and the price of the electrode.
- The issue of large DC offsets on the output could be mitigated by referencing the instrumentation amplifier input to a high passed version of itself rather than to ground.
- Either contact or non-contact right leg drive should be added to the system to reduce common mode signals such as PLI.
- Multiple fECG extraction and fQRS detection methods from previous works have been reviewed and compared in terms of performance, advantages, and disadvantages.
- Inspired by key ideas and findings from earlier works, novel methods have been developed that for the data available detect fetal heart beats with the same level of reliability as state-of-the-art QRS detectors do for adult ECGs.
- Different kinds of recurrent neural networks, LSTMs and ESNs, have displayed good abilities to mirror time-dependent transformations of the thoracic mECG to the mECG component of the AECG mixture.
- The hyperparameters of the ESN network were successfully optimized using a genetic algorithm.
- Special care has not been taken to optimize the algorithmic efficiency, but the ESN-CWT combination already has a low enough computation time for online

monitoring through an external computer. Further optimization is however desirable to lower energy consumption and enable online computation in an embedded system.

- Merely one thoracic electrode and three abdominal electrodes were used during the evaluation, but the robustness and reliability of the methods is expected to improve even more for a higher number of electrodes.
- Multiple fQRS extraction of different characters, e.g. Pan-Tompkins and the CWT based method, could be combined in the future for additional robustness and easier detection of leads with poor signal quality.
- The NI-fECG in combinations with the developed algorithms would be suitable for quantifying beat-to-beat variations.
- Additional measurement data with reliable annotations is required to work towards the desired resolution of ± 5 BPM for a commercial application. The performance of the developed methods also needs to be assessed on data that includes abnormal cardiac behavior.
- An FFNN or an RNN that considers QRS regions in the time domain that have been pre-detected by the CWT based method could probably be used for accurate determination of the exact time of the R peak. Hopefully this additional step should be sufficient to make the total algorithm reach the desired resolution.
- Cases of multiple pregnancies occur too often (around 3% of all births) to ignore completely and should be considered in the future. A blind source separation based method is probably required.

Bibliography

- [1] AL Goldberger, LAN Amaral, L Glass, JM Hausdorff, PCh Ivanov, RG Mark, JE Mietus, GB Moody, C-K Peng, and HE Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 101(23):e2, 2000.
- [2] Ian M. Symonds, Sabaratnam. Arulkumaran, E. M. (Edwin Malcolm) Symonds, and E. M. (Edwin Malcolm). Symonds. *Essential obstetrics and gynaecology*. Churchill Livingstone/Elsevier, 2013.
- [3] Joachim Behar. Extraction of clinical information from the non-invasive fetal electrocardiogram. 2016.
- [4] Papa November. Capacitor schematic with dielectric, 2008.
- [5] Stanisław Skowron. Understanding the faraday cage [Public domain].
- [6] M A Hasan, M B I Reaz, M I Ibrahimy, M S Hussain, and J Uddin. Detection and Processing Techniques of FECG Signal for Fetal Monitoring. *Biological procedures online*, 11(1):263–95, 2009.
- [7] Cburnett. Artificial neural network, 2006.
- [8] F. Deloche. Recurrent neural network unfold, 2017.
- [9] The World Bank. Mortality rate, neonatal (per 1,000 live births).
- [10] WHO | Stillbirths. *WHO*, 2016.
- [11] Gari D Clifford, Ikaro Silva, Joachim Behar, and George B Moody. Non-invasive fetal ECG analysis. *Physiological Measurement*, 35(8):1521–1536, 8 2014.
- [12] Duncan R. Neilson, Roger K. Freeman, and Shelora Mangan. Signal ambiguity resulting in unexpected outcome with external fetal heart rate monitoring. *American Journal of Obstetrics and Gynecology*, 198(6):717–724, 6 2008.
- [13] First Marx, Online Training, Practical Training, C Darin, Rank Online Training, M Kimberly, G Deepa, Ethics Board, Enter Principal, Investigator Primary, Food Systems, Emu Behaviour Study, and New Co-investigator. *Molecular Biology of the Cell 6 edition by Bruce Alberts*. 2014.

- [14] Martin Morad and Leslie Tung. Ionic events responsible for the cardiac resting and action potential. *The American Journal of Cardiology*, 1982.
- [15] John E. Hall. *Guyton and Hall Textbook of Medical Physiology (12e)*. 2011.
- [16] Russell K. Hobbie. The Electrocardiogram as an Example of Electrostatics. *American Journal of Physics*, 2005.
- [17] David B. Geselowitz. On the Theory of the Electrocardiogram. *Proceedings of the IEEE*, 1989.
- [18] Russell K. Hobbie. Improved explanation of the electrocardiogram. *American Journal of Physics*, 2005.
- [19] J. G. Nijhuis. Fetal motility and fetal behavior. *van Geijn HP, and Copray FJA (eds)*, page 183–87, 1994.
- [20] Reza Sameni. *Extraction of Fetal Cardiac Signals from an Array of Maternal Abdominal Recordings*. PhD thesis, 2008.
- [21] Stephanie Pildner von Steinburg, Anne-Laure Boulesteix, Christian Lederer, Stefani Grunow, Sven Schiermeier, Wolfgang Hatzmann, Karl-Theodor M. Schneider, and Martin Daumer. What is the “normal” fetal heart rate? *PeerJ*, 2013.
- [22] Peter van Leeuwen, Silke Lange, Anita Klein, Daniel Geue, and Dietrich H W Grönemeyer. Dependency of magnetocardiographically determined fetal cardiac time intervals on gestational age, gender and postnatal biometrics in healthy pregnancies. *BMC Pregnancy and Childbirth*, 2004.
- [23] Maria Peters, John Crowe, Jean Francois Piéri, Hendrik W P Quartero, Barrie Hayes-Gill, David James, Jeroen Stinstra, and Simon Shakespeare. Monitoring the fetal heart non-invasively: A review of methods. *Journal of Perinatal Medicine*, 2001.
- [24] J.G. Stinstra. The reliability of the fetal magnetocardiogram, 11 2001.
- [25] J. F. Piéri, J. A. Crowe, B. R. Hayes-Gill, C. J. Spencer, K. Bhogal, and D. K. James. Compact long-term recorder for the transabdominal foetal and maternal electrocardiogram. *Medical and Biological Engineering and Computing*, 2001.
- [26] E. M. Graatsma, B. C. Jacod, L. A J Van Egmond, E. J H Mulder, and G. H A Visser. Fetal electrocardiography: Feasibility of long-term fetal heart rate recordings. *BJOG: An International Journal of Obstetrics and Gynaecology*, 2009.
- [27] E. M. (Edwin Malcolm) Symonds, Daljit. Sahota, and Allan. Chang. *Fetal electrocardiography*. Imperial College Press, 2001.

- [28] G. HENSON, G. S. DAWES, and C. W.G. REDMAN. Characterization of the reduced heart rate variation in growth-retarded fetuses. *BJOG: An International Journal of Obstetrics & Gynaecology*, 1984.
- [29] H P van Geijn, F J Copray, D K Donkers, and M H Bos. Diagnosis and management of intrapartum fetal distress. *European Journal of Obstetrics, Gynecology, and Reproductive Biology*, 1991.
- [30] Martijn A. Oudijk, Anneke Kwee, Gerard H.A. Visser, Sofia Blad, Erik J. Meijboom, and Karl G. Rosén. The effects of intrapartum hypoxia on the fetal OT interval. *BJOG: An International Journal of Obstetrics and Gynaecology*, 2004.
- [31] Barrie Hayes-Gill, Sarmina Hassan, Fadi G. Mirza, Sophia Ommani, John Himsworth, Molham Solomon, Raymond Brown, Barry S. Schifrin, and Wayne R. Cohen. Accuracy and Reliability of Uterine Contraction Identification Using Abdominal Surface Electrodes. *Clinical Medicine Insights: Women's Health*, 2012.
- [32] F Ahlfeld. Über intrauterine Atmungsbewegungen des Kindes. *Verh. dsch. Ges. Gynak*, 1888.
- [33] John J. Greer. Control of breathing activity in the fetus and newborn. *Comprehensive Physiology*, 2012.
- [34] J C Dornan, J W Ritchie, and S Ruff. The rate and regularity of breathing movements in the normal and growth-retarded fetus. *British journal of obstetrics and gynaecology*, 91(1):31–6, 1 1984.
- [35] P. M. Boyce and A. L. Todd. Increased risk of postnatal depression after emergency caesarean section. *Medical Journal of Australia*, 1992.
- [36] Annika Karlström, Regina Engström-Olofsson, Karl Gustaf Norbergh, Mats Sjöling, and Ingegerd Hildingsson. Postoperative pain after cesarean birth affects breastfeeding and infant care. *JOGNN - Journal of Obstetric, Gynecologic, and Neonatal Nursing*, 2007.
- [37] S. B. Barnett and D. Maulik. Guidelines and recommendations for safe use of Doppler ultrasound in perinatal applications. *Journal of Maternal-Fetal and Neonatal Medicine*, 2009.
- [38] J. A. Crowe, A. Harrison, and B. R. Hayes-Gill. The feasibility of long-term fetal heart rate monitoring in the home environment using maternal abdominal electrodes. *Physiological Measurement*, 1995.
- [39] S. L. Bloom, C. Y. Spong, E. Thom, and et al. Fetal Pulse Oximetry and Cesarean Delivery. *Obstetric Anesthesia Digest*, 2008.
- [40] Christine E. East, Shaun P. Brennecke, James F. King, Fung Yee Chan, and Paul B. Colditz. The effect of intrapartum fetal pulse oximetry, in the presence

- of a nonreassuring fetal heart rate pattern, on operative delivery rates: A multicenter, randomized, controlled trial (the FOREMOST trial). *American Journal of Obstetrics and Gynecology*, 2006.
- [41] Janusz Jezewski, Janusz Wrobel, and Krzysztof Horoba. Comparison of Doppler ultrasound and direct electrocardiography acquisition techniques for quantification of fetal heart rate variability. *IEEE Transactions on Biomedical Engineering*, 2006.
- [42] Frédéric Vrins, Christian Jutten, and Michel Verleysen. Sensor Array and Electrode Selection for Non-invasive Fetal Electrocardiogram Extraction by Independent Component Analysis. pages 1017–1024. 2004.
- [43] Fawwaz T. (Fawwaz Tayssir) Ulaby. *Fundamentals of applied electromagnetics*. Prentice Hall, 1999.
- [44] David K. (David Keun) Cheng. *Fundamentals of engineering electromagnetics*. Addison-Wesley Pub. Co, 1993.
- [45] John D. (John Denis) Enderle. *Bioinstrumentation*. Morgan & Claypool Publishers, 2006.
- [46] R J Prance, a Debray, T D Clark, H Prance, M Nock, C J Harland, and a J Clippingdale. An ultra-low-noise electrical-potential probe for human-body scanning. *Measurement Science and Technology*, 11(3):291–297, 2000.
- [47] Burr-Brown Corporation. INA116 Datasheet. 1995.
- [48] Yu Mike Chi, Tzyy Ping Jung, and Gert Cauwenberghs. Dry-contact and noncontact biopotential electrodes: Methodological review. *IEEE Reviews in Biomedical Engineering*, 2010.
- [49] Enrique Spinelli, Federico Guerrero, Pablo García, and Marcelo Haberman. A simple and reproducible capacitive electrode. *Medical Engineering and Physics*, 2016.
- [50] A. Searle and L. Kirkup. A direct comparison of wet, dry and insulating bioelectric recording electrodes. *Physiological Measurement*, 2000.
- [51] E. Huigen, Abraham Peper, and C. A. Grimbergen. Investigation into the origin of the noise of surface electrodes. *Medical and Biological Engineering and Computing*, 2002.
- [52] M. S. Spach, R. C. Barr, J. W. Havstad, and E. C. Long. Skin-electrode impedance and its effect on recording cardiac potentials. *Circulation*, 1966.
- [53] *Applied Biomedical Engineering*. 2012.
- [54] Yu M. Chi, Patrick Ng, and Gert Cauwenberghs. Wireless noncontact ECG and EEG biopotential sensors. *ACM Transactions on Embedded Computing Systems*, 2013.

- [55] Enrique Spinelli and Marcelo Haberman. Insulating electrodes: A review on biopotential front ends for dielectric skin-electrode interfaces, 2010.
- [56] Gustavo Castro and Scott Hunt. Designing High Performance Systems with Low Noise Instrumentation Amplifiers. Technical report, 2012.
- [57] Yu M. Chi, Christoph Maier, and Gert Cauwenberghs. Ultra-high input impedance, low noise integrated amplifier for noncontact biopotential sensing. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2011.
- [58] Benjamin C. Kuo and M. F. Golnaraghi. *Automatic control systems*. John Wiley & Sons, New York, 8th ed. edition, 2003.
- [59] Anthony J. Portelli and Slawomir J. Nasuto. Design and development of non-contact bio-potential electrodes for pervasive health monitoring applications. *Biosensors*, 2017.
- [60] University of Oslo Østby, Joar Martin. Lecture no 9 (Mot 6) Noise in field effect transistors.
- [61] A. J. Clippingdale, R. J. Prance, T. D. Clark, and C. Watkins. Ultrahigh impedance capacitively coupled heart imaging array. *Review of Scientific Instruments*, 1994.
- [62] C. J. Harland, T. D. Clark, and R. J. Prance. Electric potential probes-New directions in the remote sensing of the human body. *Measurement Science and Technology*, 2002.
- [63] Texas Instruments. INA12x Precision, Low-Power Instrumentation Amplifiers.
- [64] Ko Keun Kim and Kwang Suk Park. Effective coupling impedance for power line interference in capacitive-coupled ECG measurement system. In *5th Int. Conference on Information Technology and Applications in Biomedicine, ITAB 2008 in conjunction with 2nd Int. Symposium and Summer School on Biomedical and Health Engineering, IS3BHE 2008*, 2008.
- [65] B. Widrow, J.R. Glover, J.M. McCool, J. Kaunitz, C.S. Williams, R.H. Hearn, J.R. Zeidler, Jr. Eugene Dong, and R.C. Goodlin. Adaptive noise cancelling: Principles and applications. *Proceedings of the IEEE*, 1975.
- [66] Jean François Cardoso. Blind signal separation: Statistical principles. *Proceedings of the IEEE*, 1998.
- [67] V.E. Bondarenko. Artificial neural networks. *Salem Press Encyclopedia of Science*, 2018.
- [68] B. Mehlig. *Artificial Neural Networks*. Department of Physics, University of Gothenburg., Göteborg, 25 edition, 2018.

- [69] M. (Mattias) Wahde. *Biologically inspired optimization methods : an introduction*. WIT Press, 2008.
- [70] John H. (John Henry) Holland. *Adaptation in natural and artificial systems : an introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan Press, 1975.
- [71] F Gray. Pulse code communication. Technical report, 1953.
- [72] Continuous Wavelet Transform and Scale-Based Analysis, 2019.
- [73] Sofia C. Olhede and Andrew T. Walden. Generalized Morse wavelets. *IEEE Transactions on Signal Processing*, 2002.
- [74] Jonathan M. Lilly and Sofia C. Olhede. Generalized morse wavelets as a superfamily of analytic wavelets. *IEEE Transactions on Signal Processing*, 2012.
- [75] American National Standards Institute. Testing and reporting performance results of cardiac rhythm and ST segment measurement algorithms: ANSI/AAMI EC57. *Association for the Advancement of Medical . . .*, 2012.
- [76] Yutaka Sasaki. The truth of the F-measure. *American Review of Respiratory Disease*, 1982.
- [77] ANSI/AAMI. Cardiac monitors, heart rate meters, and alarms. *American National Standard (ANSI/AAMI EC13:2002)*, 2002.
- [78] J.F. Cardoso and A. Souloumiac. Blind beamforming for non-gaussian signals. *IEE Proceedings F Radar and Signal Processing*, 1993.
- [79] A. Hyvärinen and E. Oja. Independent component analysis: Algorithms and applications. *Neural Networks*, 2000.
- [80] Partha Pratim Kanjilal, Sarbani Palit, and Goutam Saha. Fetal ECG extraction from single-channel maternal ECG using singular value decomposition. *IEEE Transactions on Biomedical Engineering*, 1997.
- [81] Sergio Cerutti, Giuseppe B. Baselli, Silvia Civardi, Enrico Ferrazzi, Anna Maria Marconi, Massimo Pagani, and Giorgio Pardi. Variability analysis of fetal heart rate signals as obtained from abdominal electrocardiographic recordings. *Journal of Perinatal Medicine*, 1986.
- [82] S. M M Martens, Chiara Rabotti, Massimo Mischi, and Rob J. Sluijter. A robust fetal ECG detection method for abdominal recordings. *Physiological Measurement*, 2007.
- [83] Mihaela Ungureanu, Johannes W.M. Bergmans, Swan Guid Oei, and Rodica Strungaru. Fetal ECG extraction during labor using an adaptive maternal beat subtraction technique. *Biomedizinische Technik*, 2007.

- [84] R. Vullings, C. H L Peters, R. J. Sluijter, M. Mischi, S. G. Oei, and J. W M Bergmans. Dynamic segmentation and linear prediction for maternal ECG removal in antenatal abdominal recordings. *Physiological Measurement*, 2009.
- [85] Julien Oster, Joachim Behar, Omid Sayadi, Shamim Nemati, Alistair E.W. Johnson, and Gari D. Clifford. Semisupervised ECG Ventricular Beat Classification with Novelty Detection Based on Switching Kalman Filters. *IEEE Transactions on Biomedical Engineering*, 2015.
- [86] G. V.P.ChandraSekharYadav, B. Ananda Krishna, and M. Kamaraju. Performance of Wiener Filter and Adaptive Filter for Noise Cancellation in Real-Time Environment. *International Journal of Computer Applications*, 2014.
- [87] Simon S. Haykin and Bernard Widrow. *Least-mean-square adaptive filters*. John Wiley, 2003.
- [88] Karl J. (Karl Johan) Astrom and Bjorn. Wittenmark. *Adaptive control*. Dover Publications, 2008.
- [89] A. Vahidi, A. Stefanopoulou, and H. Peng. Recursive least squares with forgetting for online estimation of vehicle mass and road grade: Theory and experiments. *Vehicle System Dynamics*, 2005.
- [90] Patrick E. McSharry, Gari D. Clifford, Lionel Tarassenko, and Leonard A. Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*, 2003.
- [91] Ikaro Silva, Joachim Behar, Reza Sameni, Tingting Zhu, Julien Oster, Gari D Clifford, and George B Moody. Noninvasive Fetal ECG: the PhysioNet/Computing in Cardiology Challenge 2013. *Computing in cardiology*, 40:149–152, 3 2013.
- [92] Reza Sameni, Christian Jutten, and Mohammad B. Shamsollahi. What ICA provides for ECG processing: Application to noninvasive fetal ECG extraction. In *Sixth IEEE International Symposium on Signal Processing and Information Technology, ISSPIT*, 2007.
- [93] Sebastian Zaunseder, Fernando Andreotti, Marcos Cruz, Holger Stepan, Claudia Schmieder, Hagen Malberg, Alexander Jank, Niels Wessel, Alexander Jank, and Hagen Malberg. Fetal QRS Detection by means of Kalman Filtering and using the Event Synchronous Canceller. *Int. J. Bioelectromagn.*, 2013.
- [94] Yoshua Bengio. *Deep Learning Book*. MIT Press, 2015.
- [95] Alex Graves, Abdel Rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2013.
- [96] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to Sequence Learning with Neural Networks Ilya. *arXiv:1409.3215v3 [cs.CL]*, 2016.

- [97] Herbert Jaeger. A tutorial on training recurrent neural networks , covering BPPT , RTRL , EKF and the " echo state network " approach. *German National Research Center for Information Techonlogy*, 2005.
- [98] Mantas Lukoševičius. A practical guide to applying echo state networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2012.
- [99] Peter Ford Dominey and Franck Ramus. Neural network processing of natural language: I. Sensitivity to serial, temporal and abstract structure of language in the infant, 2000.
- [100] I Ilies, H Jaeger, and O Kosuchinas. Stepping forward through echoes of the past: forecasting with echo state networks. 2007.
- [101] Herbert Jaeger. The “ echo state ” approach to analysing and training recurrent neural networks. *Tecnical report GMD report*, 2001.
- [102] Herbert Jaeger and Harald Haas. Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication. *Science*, 2004.
- [103] D. Verstraeten, B. Schrauwen, and D. Stroobandt. Reservoir-based techniques for speech recognition. 2008.
- [104] Joachim Behar, Alistair Johnson, Gari D. Clifford, and Julien Oster. A comparison of single channel fetal ecg extraction methods. *Annals of Biomedical Engineering*, 2014.
- [105] Sotirios P. Chatzis and Yiannis Demiris. The copula echo state network. *Pattern Recognition*, 2012.
- [106] Jiapu Pan and Willis J. Tompkins. A Real-Time QRS Detection Algorithm. Technical Report 3, 1985.
- [107] A. Hyvärinen and H. Morioka. Unsupervised Feature Extraction by Time-Contrastive Learning and Nonlinear ICA. In *30th Conference on Neural Information Processing Systems (NIPS 2016)*, 2016.
- [108] B. Boyle, R. McConkey, E. Garne, M. Loane, M. C. Addor, M. K. Bakker, P. A. Boyd, M. Gatt, R. Greenlees, M. Haeusler, K. Klungsøyr, A. Latos-Bielenska, N. Lelong, R. McDonnell, J. Métneki, C. Mullaney, V. Nelen, M. O’Mahony, A. Pierini, J. Rankin, A. Rissmann, D. Tucker, D. Wellesley, and H. Dolk. Trends in the prevalence, risk and pregnancy outcome of multiple births with congenital anomaly: A registry-based study in 14 European countries 1984-2007. *BJOG: An International Journal of Obstetrics and Gynaecology*, 2013.