UNIVERSITY OF LEEDS

DOCTORAL THESIS

---

# A Biomagnetic Field Mapping System for Detection of Heart Disease in a Clinical Environment

---

*Author:*
John William MOONEY

*Supervisor:*
Prof. Benjamin T.H VARCOE

*A thesis submitted in fulfilment of the requirements*
*for the degree of Doctor of Philosophy*

*in the*

Experimental Quantum Information Group
School of Physics and Astronomy

September 2018

# Declaration of Authorship

I, John William MOONEY , confirm that the work submitted is my own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. I confirm that appropriate credit has been given within the thesis where reference has been made to the work of others.

The chapter "Device Development" contains work from a jointly authored publication: Mooney JW, Ghasemi-Roudsari S, Banham ER, Symonds C, Pawlowski N, Varcoe BTH. A portable diagnostic device for cardiac magnetic field mapping. Biomed Phys Eng Express. 2017;3(1):015008. The paper was written by myself and I was involved with all aspects of the work. The calibration of the sensor response and all the data analysis were carried out by myself. Ben Varcoe optimised the coil geometry and was project supervisor. The printed circuit boards (PCB) were populated by Shima Ghasemi and myself. Initial testing of the ICM was performed by Chris Symmonds, including gradiometer configurations. The array mounting hardware was designed and built by Brian Gibbs. The COMSOL modelling of the sensor array was performed by Nick Pawlowski and myself. The Acquisiton user interface was coded by Edward Reade-Banham and myself.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.
The right of John William MOONEY to be identified as Author of this work has been asserted by John William MOONEY in accordance with the Copyright, Designs and Patents Act 1988.

*"Today I see, for the first time, the true coast of Ithaca."*

– Riccardo Fenici

A comment made by Professor Fenici at the end of the 2018 European Workshop on Magnetocardiography, where he gave a talk likening his nearly 4 decade journey in Magnetocardiography to Odysseus' arduous journey home to Ithaca.

*"Premature optimisation is the route of all evil."*

– Donald E. Knuth

A widely applicable meme. I put it here in relation to the field of magnetocardiography; to highlight the importance of evolving the whole system, and not getting entrenched on optimising single sensor performance.

# Acknowledgements

Primarily to my academic supervisor, Professor Ben Varcoe, for fascinating insights and countless stimulating discussions, and without whom this project would not have existed.

To all those who have shared the QIX lab experience! Beth Newton, Joseph Wilson, Anne Ghesquiere, Raoul Franky, Diego A. Quiñones Valles, Freya Louise, Matthew Everitt, Dastan Khalid, Rachel Thompson, Stephen Staples, Martin Jones. Especially my collaborator Shima Ghasemi-Roudsari, for her friendship over the years and her hard work running the clinical trials and deeply exploring the diagnostic capabilities of MCG. My masters students, Nick Pawlowski, Dan Wilson and Mike Smith; all of whom were a pleasure to work beside as we amplified our minds by sharing them.Thanks to James De-Lisle for sharing his knowledge of Simulated Annealing Monte Carlo optimisation. Thanks to those who helped, during the wild early development phase; Edward Reade Banham, Debbie Grimmond, Chris Symmonds, Mira Moran, Sophie Mann and Francis Ridgeon. A special thanks to the technical staff at the University of Leeds, particularly Phillip Thornton who was key to the development of the magnetometer electronics. Also, Lee Harris, Stuart Weston, Brian Gibbs, Richard Oliver, Andy Price, Brian Chesworth and Trevor Hayes. Thanks to Professor Gary Green at York Neuroimaging centre for support and access to the magnetically shielded room, which proved invaluable to our research. To the team at LGI whose assistance with the clinical trial was essential; L Falk, R Byrom, M Williamson, D Brettle, Professor M Kearney and S Smye.

To the Alderhillians! Peter, Noemi, Stephen, Tiia, Jack, Adam, Georgie, Rebecca and Sofia for your understanding and friendship. To Marta Digna Feal Franqueira for your care and valuable advice at critical times. To my parents who have always given me love and support, and let me play with soldering irons and scalpels as a kiddo. I am forever grateful. To my Grandfather, John Green for his caring and critical mind and our stimulating discussions.

I would like to thank the multitude of contributors to the open source projects that have been invaluable to me during this research, which include: GNU/Linux, Python (particularly the matplotlib, Scipy, Numpy and Scikit-Learn libraries), Jupyter, Tensorflow, Gimp, Inkscape...

I am grateful for the funding provided by the NHS National Innovation Center, the IKC for Medical Devices, the BHRC, HEIF, IPGroup, Creavo Medical Technologies and the University of Leeds who awarded a 110-scholarship.

UNIVERSITY OF LEEDS

# *Abstract*

Faculty of Maths and Physical Sciences
School of Physics and Astronomy

Doctor of Philosophy

## A Biomagnetic Field Mapping System for Detection of Heart Disease in a Clinical Environment

by John William MOONEY

This PhD was inspired by the clinical demand for a system to triage chest pain and the untapped diagnostic potential of magnetocardiography (MCG) to reliably detect silent ischaemic heart disease, which is responsible for the highest mortality rate of any single disease category. The aim was to develop a low cost and portable biomagnetic field mapping system capable of differentiating between healthy and diseased hearts within an unshielded hospital environment. This entailed the development of a system based on an array of magnetometers with sufficient sensitivity ($104\text{fT}/\sqrt{\text{Hz}}$ at 10Hz) and noise rejection performance ($68.4 \pm 3.9$ dB) to measure the small magnetic field associated with the heart beat, the magneocardiogram, within a much larger background noise.

The array of induction coil magnetometers (ICM) we developed had sufficient sensitivity and were robust to high amplitude noise. These sensors were also cheap to manufacture and capable of operating on battery power, allowing a low cost, portable device to be developed.

The key element that allowed us to achieve unshielded operation was the development of an algorithmic spatial filter, used as a substitute to operation within a magnetically shielded room. This coherent noise rejection (CNR) algorithm exploits the difference in spatial coherence between the local cardiac signals and the distant background noise sources. The observed coherence width during a clinical trial of the system within a hospital ward was $2.8 \pm 0.9 \times 10^6 mm^2$. This allowed us to capture MCG signals with a signal to noise ratio of $SNR_{QRS} = 0.93 \pm 4.43$dB. The performance of CNR was found to improve by 9dB per order of magnitude increase in environmental spatial coherence width. The coherence width can be increased by changes to hospital architecture, electromagnetic field regulation and device design optimisation.

The thesis also explores a variety of approaches to obtain binary diagnostic information from MCG, from traditional statistical learning on manually engineered features to machine learning. I found that machine learning techniques, in particular convolutional neural networks (CNN), were able to capture more diagnostic information than traditional techniques and achieved world class prediction accuracy of 88% on the clinical trial dataset.

# Contents

**Declaration of Authorship** iii

**Acknowledgements** vii

**Abstract** ix

**1 Introduction** 1
- 1.1 Motivation . . . 1
- 1.2 Background and Review . . . 3
  - 1.2.1 The Heart . . . 3
    - The Cardiac Cycle . . . 4
    - Cardiac Action Potential . . . 7
    - The Cardiac Magnetic Field . . . 8
    - Cardiac Diagnostic Information . . . 10
  - 1.2.2 History of Magnetocardiography Devices . . . 10
  - 1.2.3 Magnetometer Development at Leeds . . . 12
  - 1.2.4 Clinical Trials . . . 13
- 1.3 Thesis Outline . . . 15
  - 1.3.1 Criteria for success . . . 16

**2 Device Development** 17
- 2.1 The Induction Coil Sensor . . . 17
  - 2.1.1 Calibration of Sensor Response . . . 21
  - 2.1.2 Sensitivity . . . 24
- 2.2 Alternative Sensing Technologies . . . 24
- 2.3 System Architecture . . . 26
  - 2.3.1 Detector Array . . . 26
    - Array Design Improvements . . . 27
    - Pre-amplifier Noise . . . 29
    - Modelling Flux Linkage Between the Sensor Cores . . . 30
    - Procedures for Testing the Array . . . 33
  - 2.3.2 Digital Signal Acquisition . . . 34
    - Synchronous ECG Acquisition . . . 34
  - 2.3.3 Operator Interface . . . 35
    - Preparation of the Scan Participant . . . 36
    - Power-up Routine . . . 36
    - Positioning the Detector Array . . . 37
  - 2.3.4 Acquisition User Interface . . . 37
  - 2.3.5 Analytical Workstation . . . 38
- 2.4 Conclusion . . . 38

3 Environmental Noise & Digital Signal Processing 41
3.1 Introduction . . . 41
3.1.1 Electromagnetic Noise Sources . . . 42
3.1.2 Acoustic Noise . . . 43
3.1.3 Ballistocardiographic noise . . . 43
3.2 Noise Analysis . . . 47
3.2.1 Allan Variance . . . 48
3.2.2 Frequency Analysis . . . 50
Wavelet Decomposition . . . 50
3.2.3 Spatial Coherence . . . 51
3.3 Methods for Environmental Noise Rejection . . . 55
3.3.1 Magnetically Shielded Rooms . . . 58
3.3.2 Procedural Techniques . . . 59
3.4 Digital Signal Processing . . . 60
3.4.1 Cardiac-Cycle Averaging . . . 62
3.4.2 Lowpass Filtering . . . 63
3.4.3 Coherent Noise Rejection . . . 63
3.4.4 Interpolation . . . 64
Boundary Constraints . . . 65
3.4.5 Integration . . . 66
3.4.6 Exploration of Supplementary Denoising Algorithms . . . 68
Adaptive Bucket Subtraction . . . 68
Signal Space Separation . . . 68
3.5 Conclusion . . . 68

4 Diagnostic Information 71
4.1 Introduction . . . 71
4.2 Parameter Engineering . . . 72
4.2.1 Dipole Angle Measurement . . . 73
4.2.2 Labelling MCG Features . . . 74
4.3 Diagnostic Methods . . . 75
4.3.1 Statistical Machine Learning . . . 78
Logistic Regression . . . 79
Genetic Optimisation of Parameter Set . . . 80
4.3.2 Machine Learned Classifiers . . . 81
Support Vector Classifier . . . 81
SVC Results . . . 83
Convolutional Neural Network . . . 85
CNN Results . . . 86
4.3.3 Mapping Diagnostic Information . . . 87
dMap(MCG) . . . 87
Time-Frequency dMap(ECG) . . . 88
4.4 Conclusion . . . 88

5 Discussion 93
5.1 Future Research & Development . . . 95
5.1.1 Sensors . . . 95
5.1.2 Explore the Limitations of Spatial Coherence and CNR . . . 96
5.1.3 Digital Signal Processing Improvements . . . 96

High Frequency Components . . . . . 96
Fidelity Optimisation . . . . . 97
5.1.4 Clinical Application of CNN & dMap . . . . . 97
5.1.5 Clinical Ballistocardiography? . . . . . 98
5.1.6 Stress Test MCG . . . . . 98
5.1.7 Computational Cardiac Modelling, Standards for MCG & Personalised medicine . . . . . 99
5.2 Conclusion . . . . . 99

**A Appendix 101**
A.1 Electronics . . . . . 101
A.1.1 Circuit Diagrams . . . . . 101
A.1.2 Thermal Behaviour . . . . . 101
A.2 Array Hardware . . . . . 101
A.3 Analysis User Interface . . . . . 101

**B Equations 109**

**Bibliography 111**

# List of Figures

1.1 The helical arrangement of cardiac muscle tissue. [26] . . . . . . . . . . . . 4
1.2 Illustration of a section of myocyte. The myofibrils are the contractile elements, they are surrounded by extensive ion exchange network called the Sarcoplasmic Reticulum, which is connected to the extracellular space by T-tubules. A large number of Mitochondria are present, which provide the chemical energy necessary for contraction in the form of ATP. [28] . . . . . 5
1.3 The phases of the cardiac cycle, relating mechanical dynamics with electrophysiological presentation. During the depolarisation phase, atrial contraction occurs and a P-wave can be observed in ECG, followed by ventricular contraction with observation of QRS wave. Then the repolarisation phase occurs, where the muscle relaxes and a T-wave is observed. [26] . . 6
1.4 The conduction network of the heart and corresponding action potential timeseries shown relative to the ECG in the bottom right. The action potential propagates from the sinoatrial node, through the atria to the atrioventricular node, then after a delay, into the purkinje fibers for precise distribution into the ventricles. . . . . . . . . . . . . . . . . . . . . . . . . . . 7
1.5 The action potential of a cardiac myocyte, showing the stages ionic current flow and the membrane channels involved. The duration from initial spike to return to resting potential is approximately 200ms. . . . . . . . . . . . . . 8
1.6 A MFM of a healthy volunteer during the depolarisation phase of the cardiac cycle. The small dark blue circles show the position of the sensors used to acquire the MFM; the false colour MFM is interpolated from these discrete sensors measurements. . . . . . . . . . . . . . . . . . . . . . . . . . . 10
1.7 The mutual information of two variables across a noisy channel. [Image attribution: Jvstone] . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 11
1.8 The number of papers published per year that are concerning magnetocardiography. From 1967 to 2015, data from WebofScience. Before the introduction of commercial multi-channel SQUID MCG in 1992 less than 5 papers per year were published. It is currently $approx$50 per year. . . . . . 12
1.9 A flowchart showing the capture of MCG signals. The false colour map over the sensor array demonstrates the approximate scale of the magnetic field relative to the torso. MCG and ECG signals are synchronously recorded by an ADC and sent to a computer (blue box) for digital signal processing. The end result is a series of magnetic field maps, one for each sample acquired in the cardiac cycle. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 14
1.10 The MK1 device in operation during the first clinical trial at Leeds General Infirmary. Research Nurse Lorraine Falk can be seen operating the device. While the scan subject lies supine underneath the detector array. . . . . . 15

2.1 Annotated rendering of the induction coil sensor with pre-amplifier circuit board, showing an exploded view along the coil's length. . . . . . . . . . . 19
2.2 Photograph of the (7cm diameter) induction coil sensor. The 2cm diameter core is visible in the centre of the plastic coil bobbin. The pre-amplifier PCB is mounted to the coil bobbin with three silicone feet and its perimeter taped securely with kapton tape. DC battery power is provided via the white 4-way header. The amplified signals are output via the gold plated SMB connector into a coaxial cable. . . . . . . . . . . . . . . . . . . . . . . . . . 19
2.3 Diagram of detector system circuit. A complete diagram including the power supply components is located in the appendix A.1 . . . . . . . . . . 20
2.4 Photography of the "small coil" Induction Coil Magnetometer; 4.6cm diameter. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 21
2.5 Schematic of experimental setup used to measure the sensor response. When measuring the response of an air core coil the fluxgate could remain in the centre for simultaneous measurement. The two coils are wired in series. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 22
2.6 Bode diagram showing the frequency response of the RLC filter circuit. The response is constant from 0 Hz to 200 Hz. The resonant peak occurs at 1485Hz, which is within the frequency range of the ADC. . . . . . . . . . . 23
2.7 Bode diagram showing the frequency response of the small coil sensor RLC filter circuit. The resonant peak occurs at 6130Hz. The response is nearly constant throughout the sampling range of the ADC (2.4kHz). . . . 23
2.8 Noise power spectral density in the MSR at YNIC compared to the clinical environment at LGI. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 25
2.9 Exploded view CAD rendering of the detector array used in the MK1 MCG device. The magnetometers are shown in red, with bounding boxes for the power and signal shown. The main array holder is formed from solid Acetyl plastic. The acrylic hood sits on top of the array, with a small cut out for the cable glands. The whole assembly is mounted on two aluminium right-angle extrusions that are bolted to a floor standing aluminium frame. 27
2.10 Photograph of the Mk1 detector array. The pre-amplifier circuit boards are visible (green PCB's). Analogue power distribution is provided by the loose parallel wiring. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 28
2.11 Comparison of the Mk1 detector array (small circles) with the 37 sensor array (larger circles). The sensor spacing is 56mm accept in the centre row that has 52mm spacing. A dimensioned diagram is in the appendix A.4. . 28
2.12 Exploded view CAD rendering of the battery power supply for the pre-amplifiers. The key switch toggles between charging and scanning modes. The battery voltage can be checked by pressing the button and reading the LCD screen. The internal components are galvanically isolated from the rack-mount enclosure for safe hospital use. . . . . . . . . . . . . . . . . . . 29
2.13 Experimental setup for testing the efficacy of shielding the acquisition electronics from the sensors. A 6mm thick copper plate is placed between the sensor array and distribution board that contains the acquisition electronics. 30
2.14 Comsol simulation showing magnetic flux density streamlines captured by a touching pair of small iron cylinders (image center) magnetised by a magnet (bottom left). . . . . . . . . . . . . . . . . . . . . . . . . . . . . 30

2.15 Dependence of the flux through the centre sensor of a hexagonal array on the radial spacing between sensors. Each colour represents a different core radius starting at $r = 5$mm up to $r = 20$mm in steps of $\Delta r = 2\,5$mm. . . . 31
2.16 COMSOL physics simulation showing (a) a simulated dipole field and (b) the dipole field as measured by the sensor array (black circles) then interpolated. The reconstructed dipole has a visible distortion from elliptical and a rotation of the pole peaks; The difference between actual and observed dipole angle is $9.2^o$. (c) Shows the angle error in peak position over a $360^o$ rotation, can reach $\pm 15^o$. . . . . . . . . . . . . . . . . . . . . . . . . . . 32
2.17 The error in the position (in meters) of a magnet as seen by the detector array. The error within the array is $<$ 6cm, but grows quickly outside the array perimeter. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 32
2.18 Butterfly coil produces an annular magnetic field. A dynamic voltage is applied using an arbitrary function generator, with a resistor placed in series with the coil to limit the current. The field is imaged by the detector array for testing purposes, to check sensor configuration is correct and to test the achievable signal to noise ratio. . . . . . . . . . . . . . . . . . . . . 33
2.19 Synchronous recording of ECG and R Wave Trigger output from the Accusync 42 device. Showing the location of the fiducial on the rising edge of the R-wave. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 35

3.1 Allan deviations and overlapping Allan deviations (pale blue and orange) from magnetometer time series recorded by the array with the medical bed in ON and OFF states. The lower plot shows the difference between the ON and OFF Allan Deviations. The ON state introduced significant coloured noise with time periods up to approximately 1 sec. . . . . . . . . 44
3.2 The cycle averaged 3-axis accelerometer signals (X, Y, Z) of myself, shown with the synchronously recorded ECG signal. . . . . . . . . . . . . . . . . . . 45
3.3 The principal components of acceleration of myself and lab mate Beth, show comparable structure. . . . . . . . . . . . . . . . . . . . . . . . . . . . 46
3.4 The principal component of acceleration is shown along with the corresponding synchronously recorded ECG and MCG signals. The mean covariance between the acceleration PC and the MCG signals is shown. A high covariance is seen over the repolarisation phase of the cardiac cycle. 46
3.5 Time series of two 1 hour scans taken at LGI. Semi-periodic transient noise spikes are visible, these are due to the passing of lifts in the neighbouring lift shafts. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 49
3.6 Comparison of environmental noise at the University lab and LGI, and comparing measurements within the room at LGI; close to the neighbouring lift shafts and in the center of the room (approx 3m away). . . . . . . . 50
3.7 A Hanbury Brown and Twiss Intensity Interferometer records correlations in radiowave emissions from stars (a far-field, thermal source) between two spatially separated antennas. The mutual spatial coherence of emissions at the antennta locations is proportional to the time averaged correlator output. To determine the angular size of the source, the correlation is measured as a function of antenna separation, $d$. The width of the resulting Gaussian distribution is proportional to the stars angular size, $\theta$. . . . . . . 52

3.8 Comparison of HBT correlations. The raw background signals were all highly coherent. The MCG scans were slightly above the coherence detection limit of the array, and therefore become anti-correlated within a single sensor spacing. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 53
3.9 Schematic of experimental setup for long baseline measurement of spatial coherence. A hydraulically compressed concrete block rested upon a closed cell foam mat for vibration isolation, forming the test stand. The sensors were spaced in one of two configurations as shown (Lin, a=7). . . 53
3.10 Comparison of HBT correlations. The teststand experiments were performed with two spacing distributions, one linear with 336mm separation (lin) and the other (a=7) spaced over a geometrical sequence ($ar^{(n-1)}$; r=2, a=70mm). These multiple hour measurements were subdivided into 10 minute slices, the slice mean and SEM are shown. The MCG background shows the environmental coherence measurred by the device array. The MCG signal is the incoherent remainder after digital signal processing. . . 54
3.11 Comparison of a) Coherence and b) Attenuation (dB) spectrograms: A clear correlation is visible between the spectrograms at a fine-grained level. The maximum coherence and attenuation is seen in the 50Hz band and its harmonics. In general the coherence decreases with increasing frequency. Some broadband sources are present; such as the pulse at 140s. . . . . . . . 56
3.12 The CNR attenuation and coherence width of each scan in the clinical trial. A strong positive correlation between coherence width and attenuation. Cycle averaging removes both coherent and incoherent noise, spreading the coherence width; over the dataset we see a reduction in coherence after averaging, though in a small number of scans (15%) the coherence increased. The MCG plot shows the coherence of the signals before the final lowpass filter stage with the attenuation of the lowpass filter. As expected, there is no correlation between lowpass filter performance and coherence. 56
3.13 Allan variance comparison of YNIC Shielded environment and University lab . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 59
3.14 MCG signals from shielded and unshielded environments in sequential stages of the digital signal processing routine, with synchronously recorded ECG signals. The final unshielded signal i, still has a noise remainder, however the QRS region is well defined. . . . . . . . . . . . . . . . . . . . . . . . 61
3.15 Frequency response of the moving average filter using a boxcar window of width 20ms (blue) and 40ms (green). The 20ms window attenuates the 50Hz utility grid noise by 60dB. . . . . . . . . . . . . . . . . . . . . . . . . . 63
3.16 Illustration Comparing the spatial coherence of distant and local sources as observed by a detector array. The summing (mean) of the signals to produce the bucket is performed after digitisation. The anti-correlations are computed by subtracting the bucket signal as in equation 3.11 . . . . . 64
3.17 Elementwise comparison of the original boundary (top two rows) used to prevent extrapolation with the optimised version (bottom two rows. . . . 66
3.18 Signals recorded by the 7cm ICM array within the YNIC MSR. The MCG traces are shown in blue, the integrated signals in red and the ECG is shown in green. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 67

4.1 Example of decision boundary in the case of over-fitting (green line) and a good general representation (black line) [81]. . . . . . . . . . . . . . . . . . 72

4.2 Match templates imallx and imally on the left and right. The purple regions are -1, the yellow +1. 73

4.3 Comparison of dipole angle measured using pole peak to peak (P2P) and the correlation method (corrts). The time axis shows time relative to R peak. The corrts method is smoother, does not exhibit discrete steps, and captures a rotation through 180° that is missed by the P2P method at -50ms. 74

4.4 A Paul Wavelet decomposition of a single magnetometer time series. The wavlet transform peaks from the other magnetometers are shown in circles, and the mean peak position denoted with a black line. This technique measures the centroid of the depolarisation and repolarisation regions. The centroids are shown on the MCG butterfly plot above. 76

4.5 A peak finding algorithm applied to the DOG Wavelet decomposition of a single MCG time series. The peak finding area is constrained to the two boxes shown. The greyscale color of the peak markers indicates their magnitude. The centroid of all peaks is shown by the black cross hairs. The red vertical line shows the peak found from the Paul wavelet decomposition. 77

4.6 The proposed labelling system for MCG, shown annotating a set of MCG time series. The CorrTS trace is useful for reliably highlighting the MCG zero crossings. 78

4.7 Linear decision boundary, with 100% Sensitivity. Young healthy control (YHC) group are more polarised than the age matched HV of 6060. 79

4.8 Mean population fitness, optimised for 500 generations with a genetic algorithm. With a fitness criteria of $\frac{exp(AUC_{in}+AUC_{out})}{6}$ & $Spec_{out}$. 80

4.9 SA-MCMC search algorithm. Search path shown in black, with start and end indicated by the blue and red circles. The background heatmap is produced by an exhaustive grid search ($64 \times 64 = 4096$), of approximately equal number of search points as the SA-MCMC (4000). 83

4.10 SVC prediction accuracy as a function of training steps (x) and the degree of sub-sampling of the input data. 84

4.11 ROC curve for the 1st external 4-fold loop classifier. 84

4.12 ROC curve for the external 20-fold loop classifier. 85

4.13 Labelled example data used for the 1D and 2D CNN. The data shape is $2000 \times 15 = 30000$pixels. The QRS complex can be seen in the centre, and the T-wave (ballistocardiographic) signals to the right. 86

4.14 Mean ROC curve of a 10-fold cross validated CNN trained on the 6060 dataset. The light grey/blue curves show the individual validation runs. The $AUC = 88\%$ 87

4.15 Diagnostic map of 6060 MCG dataset. Each cross shows the AUC of a 3D CNN trained only on a window of data $\pm 40ms$ surrounding the cross. A 3 fold cross validation was performed (R,G,B crosses). 88

4.16 The lower plot shows the ECG (black) and its derivative (blue). The above spectral map above shows AUC diagnostic performance of a CNN trained on subsets of the ECG data, the width of each subset is varied over the y-axis. The T region is highly diagnostic, even when taking a small 10ms subset. 89

4.17 The lower plot shows the ECG (black) and its derivative (blue). The above spectral map above shows AUC diagnostic performance of a CNN trained on subsets of the ECG-derivative data, the width of each subset is varied over the y-axis. The early S-T region has no diagnostic information, since the gradient is approximately zero. . . . . . . . . . . . . . . . . . . . . . . . . 89

4.18 Spectral map showing the difference in diagnostic information content between the ECG gradient and the ECG. Positive valued regions correspond to superior ECG gradient diagnostic performance. . . . . . . . . . . . . . . 90

A.1 Full circuit diagram of the pre-amplifier. Power is supplied into J1. The coil is connected across J4 and J5 and the signal is output through J7 and J8. The SMB PHASE header was included to allow the possibility of balanced signal transmission, though this was never used. . . . . . . . . . . . . . . . 102

A.2 Full circuit diagram of the analogue power supply. When in scan mode the batteries are connected to the XLR output sockets via 15Ω resistors, and the red LED is illuminated. A 220nF filter capacitor is connected in parallel between the resistors and the output. When in scanning mode depression of SW2 causes the digital volt meter to display the battery voltage. In charge mode the batteries are connected in series with the 24V charger, and green LED turns on. A diode, D3 is present to prevent reverse current/discharge into the charger. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 103

A.3 Infrared and visible images of preamplifier PCB at thermal equilibrium in a 20 °C environment. Thermal equilibrium is reached within 20 minutes. The INA217 amplifier chip reached 37.7 °C, the linear voltage regulators reached 44.8 °C. The gain temperature drift coefficient of the amplifier is typically 40 ppm/°C, the resultant drift is likely to be insignificant for cycle averaging. However care should be taken to ensure all amplifiers in the array are identically thermally coupled to the environment, to ensure good matching for coherent noise rejection. . . . . . . . . . . . . . . . . . . . . . 104

A.4 Dimensioned drawing of the 37 coil array layout. . . . . . . . . . . . . . . 105

A.5 Mechanical simulation of the lightened 37 coil array showing the strain induced by its own weight. The stress was within the elastic deformation capabilities of the material. However a peak strain of 0.1mm displacement was observed and the strain was non uniform across the array. The design is functional, but perhaps a higher stiffness would be better to acoustically couple the sensors. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 106

A.6 User interface for analysis for clinical trial analysis. The user cycles through individual scans modifying the cursor positioins to match the standard MCG features and flagging up errors or noise in the scan. The MFM corresponding to each cursor position is shown in the lower left of the screen, these maps are updated whenever a cursor is moved, allowing realtime exploration. The scan shown is from the 37 sensor 'alpha' commercial prototype. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 107

A.7 Flowchart showing the electrical layout of the MCG device. The patient is isolated from the mains supply by layers of switches and fuses. . . . . . . 108

# List of Tables

4.1 Table of SVM performance. . . . . . . . . . . . . . . . . . . . . . . . . . . 83
4.2 Table of CNN performance. . . . . . . . . . . . . . . . . . . . . . . . . . . 86
4.3 Table comparing all classifier performance. . . . . . . . . . . . . . . . . . . . 90

# List of Abbreviations

| | |
|---|---|
| **MCG** | **M**agneto**c**ardio**g**raphy |
| **dMCG** | **D**erivative **MCG** |
| **MFM** | **M**agnetic **F**ield **M**ap |
| **ECG** | **E**lectro**c**ardio**g**raphy |
| **dECG** | **D**erivative **ECG** |
| **dMap** | **D**iagnostic **Map** |
| **BCG** | **B**allisto**c**ardio**g**raphy |
| **EMG** | **E**lectro**m**yo**g**raphy |
| **MMG** | **M**agneto**m**yo**g**raphy |
| **MEG** | **M**agneto**e**ncephalo**g**raphy |
| **BSPM** | **B**ody **S**urface **P**otential **M**apping |
| **MFM** | **M**agnetic **F**ield **M**ap |
| **ACS** | **A**cute **C**oronary **S**yndrome |
| **IHD** | **I**schemic **H**eart **D**isease |
| **CAD** | **C**oronary **A**rtery **D**isease |
| **NSTEMI** | **N**on-**ST**-**E**levated **M**yocardial **I**nfarction |
| **AF** | **A**trial **F**ibrillation |
| **SAN** | **S**ino**a**trial **N**ode |
| **AVN** | **A**trio**v**entricular **N**ode |
| **SR** | **S**arcoplasmic **R**eticulum |
| **EMF** | **E**lectro**m**otive **F**orce |
| **ELF** | **E**xtremely **L**ow **F**requency |
| **PCB** | **P**rinted **C**ircuit **B**oard |
| **IC** | **I**ntegrated **C**ircuit |
| **ADC** | **A**nalog to **D**igital **C**onverter |
| **DSO** | **D**igital **S**torage **O**scilloscope |
| **ICM** | **I**nduction **C**oil **M**agnetometer |
| **OPM** | **O**ptically **P**umped **M**agnetometer |
| **NV** | **N**itrogen **V**acancy |
| **CNN** | **C**onvolutional **N**eural **N**etwork |
| **SVM** | **S**upport **V**ector **M**achine |
| **PCA** | **P**rincipal **C**omponent **A**nalysis |
| **FFT** | **F**ast **F**ourier **T**ransform |
| **SSS** | **S**ignal **S**pace **S**eparation |
| **ZC** | **Z**ero **C**rossing |
| **CorrTS** | **Corr**elation **T**ime **S**eries |
| **MA** | **M**ap **A**ngle |
| **DS** | **D**ipole **S**ize |
| **MAD** | **M**ap **A**ngle **D**ynamics |
| **DSD** | **D**ipole **S**ize **D**ynamics |
| **MMR** | **M**aximum-**M**inimum **R**atio |

| | |
|---|---|
| **MMRD** | **M**aximum-**M**inimum **R**atio **D**ynamics |
| **HRV** | **H**eart **R**ate **V**ariability |
| **LGI** | **L**eeds **G**eneral **I**nfirmary |
| **YNIC** | **Y**ork **N**euro-**I**maging **C**entre |
| **QIX** | **Q**uantum **I**nformation **E**xperimentation **L**ab |

# Physical Constants

| | |
|---|---|
| Speed of Light | $c_0 = 2.997\,924\,58 \times 10^{8}\,\mathrm{m\,s^{-1}}$ |
| Permeability of free space | $\mu_0 = 12.57 \times 10^{-7}\,\mathrm{H\,m^{-1}}$ |
| Permittivity of free space | $\mu_0 = 8.854 \times 10^{-12}\,\mathrm{F\,m^{-1}}$ |
| Boltzmann's Constant | $K_B = 1.380\,648\,52 \times 10^{-23}\,\mathrm{m^{2}\,kg\,s^{-2}\,K^{-1}}$ |

*I dedicate this thesis to hearts everywhere, especially those close to me past, present and future. Particularly the memory of my late grandmother Mary Green who was full of love and charm and lived for decades longer than nature would have allowed thanks to cardiology.*

# Chapter 1

# Introduction

The aim of this thesis is to develop a low-cost, portable magnetocardiography (MCG) system for the acquisition of diagnostic information in an unshielded clinical environment, using sensitive induction coil magnetometers (ICM) and digital signal processing (DSP) to record the magnetic signals of the heart and separate them from the dominant environmental noise. To achieve this, a multidisciplinary approach was required, spanning physics, physiology, electronics, digital signal processing and machine learning. The key contribution this thesis provides is a sensor agnostic algorithm for spatial filtering of environmental noise called coherent noise rejection, which replaces the use of expensive and immobile magnetically shielded rooms.

The resulting MCG system has the potential to bring magnetocardiography into mainstream adoption. Our device in particular is designed for triage of chest pain in accident and emergency departments, since it detects the presence of heart disease quickly and without contact, and critically it is able to detect silent ischemic heart disease, which is currently difficult, expensive and time consuming to triage.

This chapter provides an introduction to magnetocardiography and its significance as a clinical diagnostic tool, followed by a history of magnetocardiographic devices, including our developmental work that began before the advent of this PhD. An outline of the thesis is provided at the end of this chapter.

## 1.1 Motivation

Heart disease is the leading cause of death globally; responsible for 18.3% of deaths in 2010. Of these deaths the majority (72.5%) are due to ischemic heart disease (IHD) [1]. IHD, also known as coronary heart disease (CHD) or coronary artery disease (CAD) is caused by a build up of fat deposits that causes a narrowing or blocking of the coronary arteries reducing blood flow through the heart muscle. When IHD is detected early it can be managed by a range of treatments and lifestyle changes to improve blood flow and thereby prevent or minimise infarction (cell death) [2, 3]. Currently early detection is problematic because low-level ischemic disease states, known as non-ST-elevated myocardial infarctions (NSTEMI), are not detectable with electrocardiographic techniques. Though NSTEMI can be diagnosed precisely with Angiography, this process is invasive, time consuming and expensive[4]. Blood marker testing for Troponin is sensitive to acute ischemic damage but requires serial monitoring over several hours to return a high specificity result and is prone to false positives since it measures the presence of myocardial injury, not the cause, which can be pericarditis, renal failure and even strenuous exercise[5]. No high throughput, economical technology currently exists for the detection of NSTEMI. However, several studies have shown that NSTEMI can be detected using Magnetocardiography (MCG) [6–12].

Magnetocardiography is the study of the magnetic field associated with the electrochemical action of the heart. The heart's magnetic field is the strongest in the body due to its helical geometry and the sustained ion exchange required to precisely coordinate contraction and so achieve efficient pumping. To capture the magnetic field, an array of sensitive magnetometers are placed in proximity to the heart and their signals are recorded at a rate sufficient to capture the heart's dynamics. The resulting MCG signals form a spatio-temporal magnetic field map (MFM) that encodes information on the motion of charges within the heart tissues, which occur due to the propagation of action potentials (AP). AP propagation is an active process that is powered by aerobic respiration. In IHD regions of cardiac tissue do not receive adequate blood supply, this prevents aerobic respiration and removal of waste products. Thus AP propagation is disrupted through ischemic regions. This disruption can be measured in the MFM.

Compared to electrocardiographic techniques such as body surface potential mapping (BSPM), the magnetic field is not distorted by the inhomogeneous conductivity of the bodily tissues. Though BSPM can meet or exceed the spatial resolution of MCG, it only encodes information of the surface potential of the epicardium (outer surface of the heart), not the internal cardiac currents that are observed with MCG, hence the inability of electrocardiographic techniques to detect silent ischemic heart disease (NSTEMI). Furthermore, to perform BSPM requires a disposable array of electrodes to be fitted to the patients chest, a skillful, time consuming and costly ($\sim 1000(\$)$) procedure. Whereas MCG is a non-contact technique offering the possibility of high clinical throughput.

Though the cardiac magnetic field is the largest in the body, it is a challengingly small field to measure. To reach the required sensitivity, the sensor system used for MCG must have very low inherent noise and a high response to magnetic fields. Typically SQUID sensors are employed; these sensors are composed of a superconducting loop containing one or two Josephson junctions, the voltage across these junctions is proportional to the magnetic flux through the loop. To maintain the superconducting state the SQUID must be cryogenically cooled. Multi channel SQUID arrays have enabled the field of magnetocardiography to exist, however their high system cost has prevented mainstream adoption.

The cardiac field is tiny compared to the Earth's magnetic field and the fields generated by the human built environment. This noise is often more than a million times more intense than the heart's magnetic field. Typically magnetically shielded rooms (MSR) are used to create a measurement environment. These rooms have several skins composed of high permeability metals to isolate external magnetic fields and thick conductive sheets to attenuate electromagnetic radiation down to low frequencies. The high installation cost of an MSR makes them uneconomical to install in many hospitals.

The ideal application for MCG is the triage of chest pain, the determination of whether or not the chest pain is attributable to an acute coronary syndrome. Currently a large portion of patients arrive at hospitals with chest pain, in the majority of cases the pain is of non-cardiac origin. Because no tool exists that can confidently rule out NSTEMI, all ECG negative chest pain patients are treated as having an NSTEMI until other diagnostic results can be obtained [13]. This results in many patients entering an expensive, invasive and time consuming diagnostic pathway only to find out they are healthy[14].

The challenge to widespread clinical adoption of MCG is achieving a sufficient cost/ benefit ratio. The commercial MCG devices currently available are expensive fixed installations that usually require magnetically shielded rooms (MSR) to remove environmental noise. This static nature reduces their throughput and makes them unsuitable for patients in delicate conditions that cannot be moved safely. Furthermore the cost of an MSR can be more than an MCG device. To reach a new cost/benefit paradigm the ideal device must be cheap (initial capital and running costs), ideally portable (capable of moving to a patients bedside) and capable of operating outside of a MSR. This is what I have strived to create during this PhD.

In order for MCG to be diagnostically useful, a translation is required from magnetic field map to heart health state. ECG diagnostics have developed over many decades, by a process of human pattern recognition and statistical analysis. Many people are now skilled in analysing ECG patterns to determine the health state of the heart. Though lots of research has been performed to find patterns in MCG, especially since the commercial introduction of multichannel SQUID in the early 1990's (see figure 1.8), the search is not as comprehensive as ECG and it could take several more decades to discover a reliable set of diagnostic patterns.

Recent developments in machine learning offer an alternative avenue for an MCG diagnostic tool. In particular, convolutional neural networks have proven to be excellent at image and movie classification and I have found them highly suitable for MCG diagnosis. I have focused on just a binary classification of disease state, as it is sufficient for triage. A reliable triage device will allow true cardiac patients to be forwarded on to the plethora of devices capable of precise classification and disease localisation (tomography, perfusion CT, MRI), whilst relieving those clinical pathways from healthy people with non-cardiac chest pain.

## 1.2 Background and Review

### 1.2.1 The Heart

The mammalian heart evolved about 210 million years ago [15, 16], since then its basic arrangement has not changed. It is a four chambered system with separate circuits for pulmonary and body circulation allowing independent control; A high system blood pressure can be achieved for oxygen and nutrient transfer whilst having a lower pressure in the pulmonary circuit for optimal oxygen exchange in the lungs [17]. The structure of the heart muscle has a helical shape, see figure 1.1. Its pumping action is closer to the wringing of a mop, than the contraction of a sack [18–20]. This enables an efficient vortex blood flow, which has a higher velocity for a given pressure than a non vortex flow [21–23]. This important aspect of cardiac function is just beginning to be exploited by a new vortex flow imaging modality that can provide an index of cardiac dysfunction based on vorticity of intraventricular and aortic blood flow. [24, 25].

The heart is surrounded by the pericardium, a double layered membrane that plays a role in ion exchange modulation, mechanical mounting, lubrication and prevention of infections. The heart is mostly composed of cardiomyocytes, specialised muscles cells, able to reliably synchronise their contraction and perform at high power output for long periods. To facilitate this, cardiomyocytes contain a large number of mitochondria. The bulk of tissue formed by the myocytes is an anisotropic fibrous network, with the cells connected only at their ends, forming a branched structure. The permeable nature of this structure allows intracellular fluid flow and diffusion to exchange ions and respiration components at the high rate required to sustain a coherent action potential. To enhance ion exchange with the interior, the cells exhibit microscopic tubes called T-tubules that penetrate the cell membrane, interfacing the extra-cellular space with the sarcoplasmic

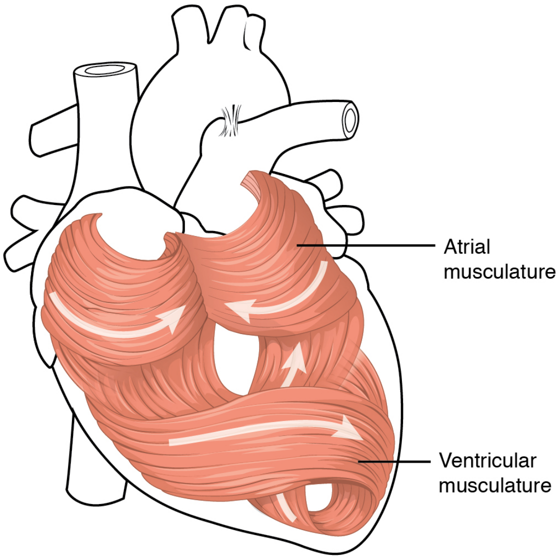

FIGURE 1.1: The helical arrangement of cardiac muscle tissue. [26]

reticulum, see figure 1.2. The sarcoplasmic reticulum is an organelle that consists of an internal network of tubules surrounding the myofibrils (the contractile units of the cell), it stores calcium ions ($Ca^{2+}$) that are used to trigger contraction.

The importance of $Ca^{2+}$ ion exchange is highlighted by the lethal mechanism of hydrofluoric acid toxicity[27]. Following exposure to hydrofluoric acid, it quickly dissolves into the blood where it reacts with with $Ca^{2+}$ ions causing hypocalcaemia. If the exposure is sufficient (>2.5% of skin surface) and not treated promptly, the result is fatal following cardiac arrest.

Interspersed with the myocyte tissue is an insulated network of specialised conductive tissues called the purkinje fibers, that distribute and help coordinate the pacemaker action potential; they also function as a backup pacemaker system, able to spontaneously generate action potentials. The initial pacemaker signal originates in the sinoatrial node (SAN), though capable of spontaneous pacemaking, the native rate is modified by the autonomic nervous system in response to demand. The autonomic response takes into account many factors including dissolved oxygen and carbon dioxide levels in the blood, arterial blood pressure, body temperature and emotional state. The pacemaker signal is divided into two distinct phases, the first phase initiated by the SAN is responsible for depolarising the atria. Then the atrioventricular node (AVN) momentarily delays the propagation, allowing blood to fill the ventricles, before relaying it to the ventricles via the purkinje fibers.

### The Cardiac Cycle

The cardiac cycle describes the sequence of actions that occurs during a heart beat, see figure 1.3. At the start of the cycle the heart muscle is in a relaxed state. Pressure in the venous network causes blood to flow into the major chambers of heart; the atria and ventricles. Then the SAN produces an action potential it spreads first to the atria causing atrial systole; the contraction of the atria, which forces blood into the ventricles. See figure 1.4. After a delay in the AVN, the action potential continues to propagate through the purkinje fibers into the ventricles, causing them to depolarise and contract; ventricular

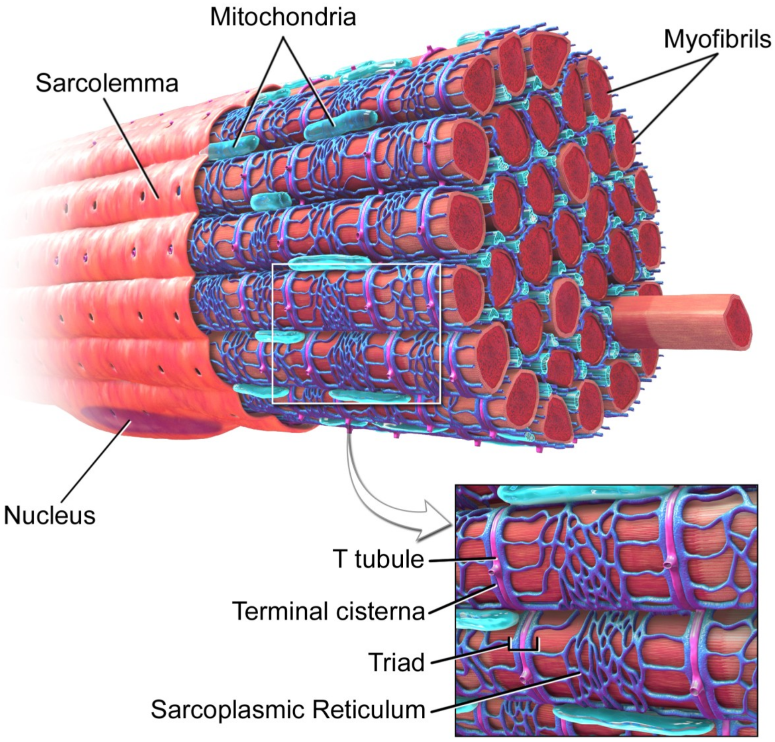
Mitochondria
Sarcolemma
Myofibrils
Nucleus
T tubule
Terminal cisterna
Triad
Sarcoplasmic Reticulum

FIGURE 1.2: Illustration of a section of myocyte. The myofibrils are the contractile elements, they are surrounded by extensive ion exchange network called the Sarcoplasmic Reticulum, which is connected to the extracellular space by T-tubules. A large number of Mitochondria are present, which provide the chemical energy necessary for contraction in the form of ATP. [28]

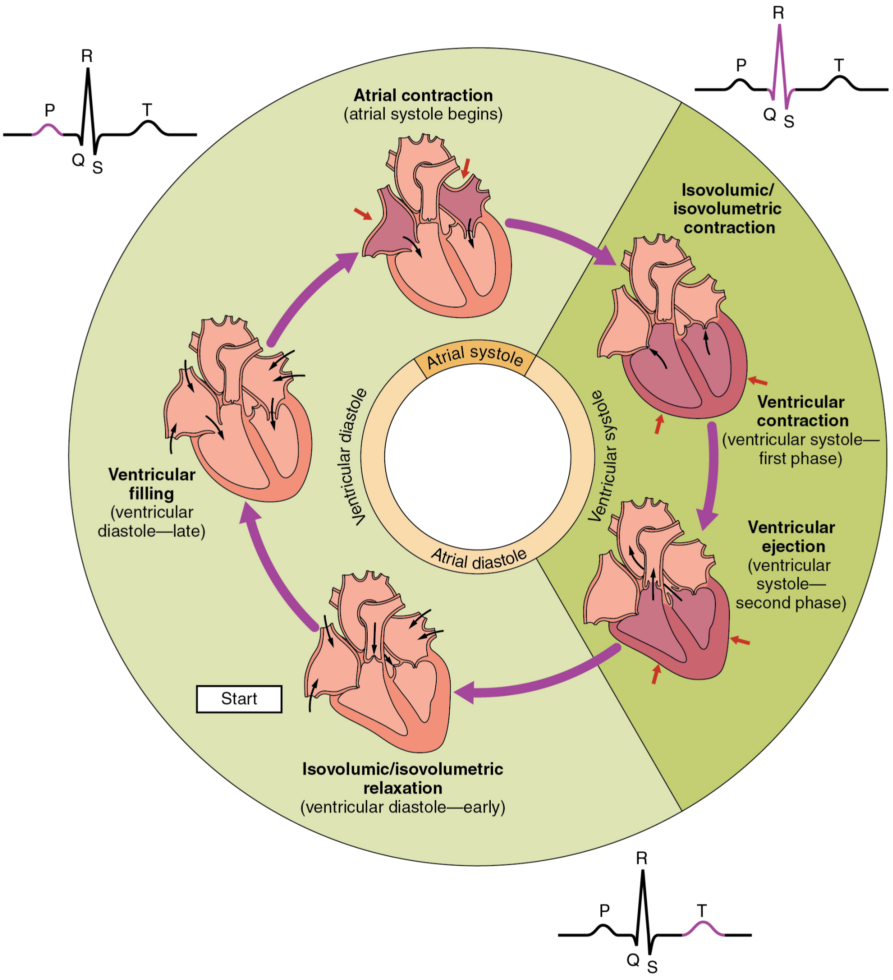
R
P
T
Q
S
Atrial contraction
(atrial systole begins)
R
P
T
Q
S
Isovolumic/
isovolumetric
contraction
Ventricular
contraction
(ventricular systole—
first phase)
Ventricular
ejection
(ventricular
systole—
second phase)
Ventricular
filling
(ventricular
diastole—late)
Atrial systole
Ventricular diastole
Ventricular systole
Atrial diastole
Start
Isovolumic/isovolumetric
relaxation
(ventricular diastole—early)
R
P
T
Q
S

FIGURE 1.3: The phases of the cardiac cycle, relating mechanical dynamics with electrophysiological presentation. During the depolarisation phase, atrial contraction occurs and a P-wave can be observed in ECG, followed by ventricular contraction with observation of QRS wave. Then the repolarisation phase occurs, where the muscle relaxes and a T-wave is observed. [26]

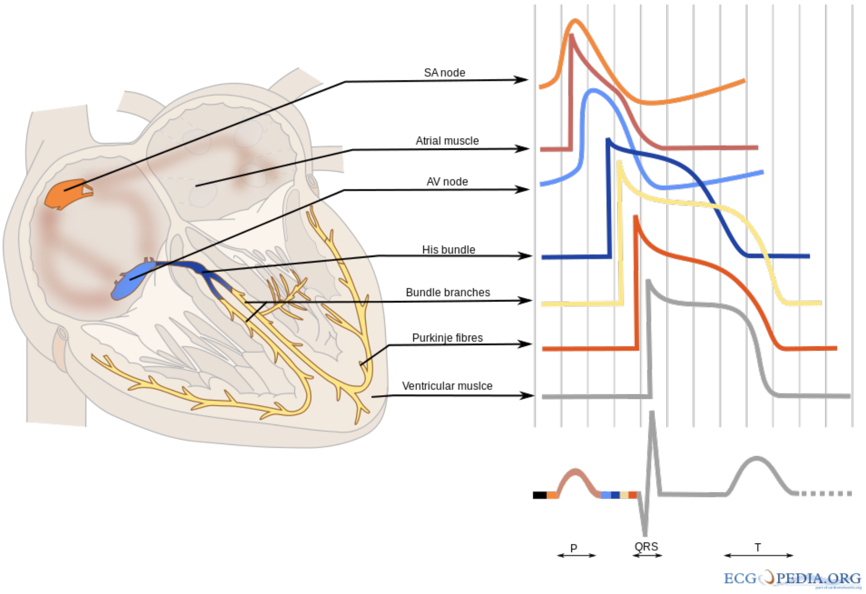

FIGURE 1.4: The conduction network of the heart and corresponding action potential timeseries shown relative to the ECG in the bottom right. The action potential propagates from the sinoatrial node, through the atria to the atrioventricular node, then after a delay, into the purkinje fibers for precise distribution into the ventricles.

systole. A resting phase follows (atrial and ventricular diastole) where the muscle relaxes and repolarises in preparation for the next cardiac cycle.

### Cardiac Action Potential

The action potential is an evolved biological mechanism for communicating information between cells, it is responsible for the transmission of signals along nerve cells and for coordinating the contraction of cardiomyocytes across the heart. The precision of synchronisation is critical for efficient cardiac output. Many heart diseases manifest in the disruption of this action potential, including IHD and arrhythmias. A measurement of the heart's magnetic field can reveal action potential abnormalities and therefore detect disease.

The action potential is a change in voltage across the cell membrane, due to the flow of ions between the intra and extra-cellular spaces. These flows are controlled by a set of membrane channel proteins. The coordinated action of these membrane channels drives the voltage impulse to travel along the cell. In the resting state an imbalance of ions stores energy in electrochemical potential, which is used to power the next impulse. The power for this is provided by membrane pumps, which use the chemical energy stored in the ATP molecules created by the mitochondria by (under non-ischemic conditions) aerobic respiration.

At rest a cardiomyocyte has a negative resting membrane potential of $-90$mV due to the existence of large negatively charged proteins within the cell, and the action of the $Na^+/K^+$-ATPase Pump, which actively pumps sodium ions ($Na^+$) outside the cell whilst

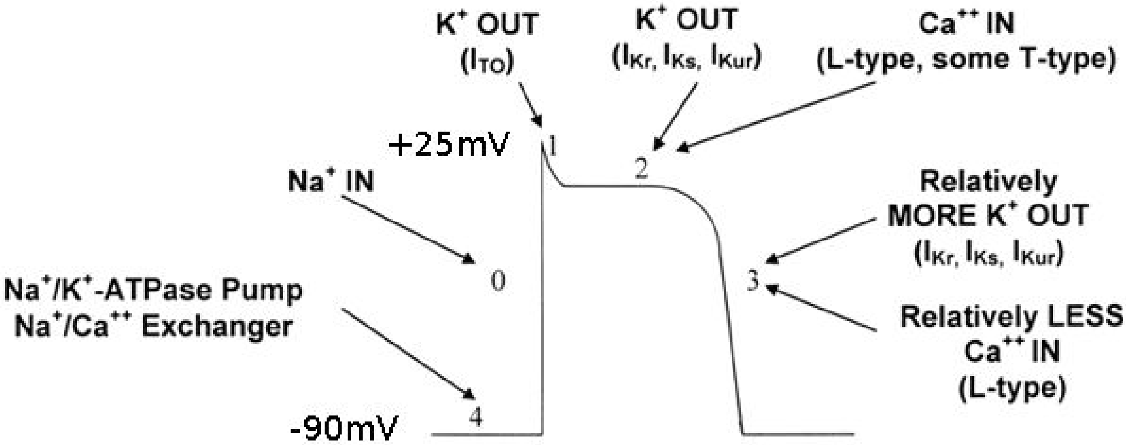

FIGURE 1.5: The action potential of a cardiac myocyte, showing the stages ionic current flow and the membrane channels involved. The duration from initial spike to return to resting potential is approximately 200ms.

bringing in potassium ions ($K^+$). If the membrane potential is increased above the threshold potential ($-70$mV); such as by the action of a pacemaker cell, the voltage sensitive membrane channel $Na_s$ opens allowing $Na^+$ ions to flow in down the concentration gradient causing a further rise in membrane potential to around $+25$mV.

Once inside the cell these ions are free to diffuse transversely, and do so, spreading the region of increased membrane potential and triggering further $Na_s$ channels to open. The depolarised region propagates along the cell. This positive membrane voltage triggers the $Na_s$ to close and a $K^+$ voltage sensitive channel to open, causing an outflux of $K^+$, reducing the membrane potential.

In a nerve cell this would cause repolarisation, a return to the resting state, within a few milliseconds. However, in a cardiomyocyte a plateau phase is maintained by the influx of $Ca^{2+}$, which balances the $K^+$ outflow. This plateau can last for 200ms allowing synchronisation of the entire ventricular tissue. The $Ca^{2+}$ bind to ryanodine receptors on the SR, causing the SR to release more $Ca^{2+}$, creating a 'calcium spark'; a simultaneous increase in $Ca^{2+}$ across the cell that binds to the protein Troponin causing myofibril contraction.

Eventually the L-type $Ca^{2+}$ channels close, and a number of $K^+$ channels open, causing the repolarisation of the cell. The $Na^+/Ca^{2+}$ Exchanger and the $Na^+/K^+$ pump then work to restore ion concentrations to the resting state configuration. As the $Ca^{2+}$ concentration within the cardiomyocytes decreases the myofibrils relax.

### The Cardiac Magnetic Field

The heart's magnetic field is generated by the movement of ions during AP propagation through the cardiac tissues. There is also a magnetohydrodynamic component due to the flow of blood, since blood is a conductive fluid, however this component is much smaller than the AP induced magnetic field.

The electromagnetic field of our universe has two components; the electric field and the magnetic field. When we treat the electric and magnetic components separately we invoke Helmholtz's theorem, which states that a vector field can be decomposed into two components a divergent component (electric) and a curl component (magnetic). In general the electric field is effected by the distribution of charged particles and the magnetic is effected by the movement of charged particles.

However, the two field components also interact directly, forming electromagnetic waves. The energy in an electromagnetic wave oscillates between the electric and magnetic fields. Though the energy in each component is equal, in the far-field, the force that acts on a charged particle due to an electromagnetic wave is predominantly due to the electric field, the magnetic force is weaker by a factor of c. The distinction between far-field and near-field is approximately characterised by being more or less than one wavelength from the wave source. In the near-field (especially $< \lambda/2\pi$) the relationship between E and B is more complex, and in general the magnetic force is increased.

In bio-electro-magnetism, the organ-scale electromagnetic waves generated have frequencies of <1kHz. Therefore wavelengths ($\lambda = c/f$), on the order of $10^7$m (@ 30Hz). Thus the measurement of bio-electro-magnetic signals is always near-field and the magnetic components are significant.

The inhomogeneous conductivity of the body tissues causes a distortion in the surface electric potential measured by ECG. There is no significant distortion of the magnetic field, though there will be an extremely small distortion due to the inhomogeneous distribution of paramagnetic (oxygen, iron, magnesium...) matter within the body.

The electromagnetic field is described by Maxwell's equations. The perturbation of the magnetic field due to the motion of charges is described by Maxwell's 4th equation (Ampere's Law):

$$\nabla \times \mathbf{B} = \mu_0 \left( \mathbf{J} + \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right) \tag{1.1}$$

Which states that the curl of a magnetic field around an axis is proportional to the component of the current density plus the change in electric field along the axis. Gauss's Law states that the magnetic field forms closed loops (since $\nabla \cdot \mathbf{B} = 0$), a measurement of B outside the heart contains information on the current density within the heart, that is revealed by the curl operator ($\nabla \times \mathbf{B}$).

The complex 3D helical cardiac structure results in a complex, subtle magnetic field, which has an unknown micro/nano structure, as no mapping has been performed below a spatial resolution of 100µm [29, 30]. Such nano-scale measurements are needed to create and validate the next generation computational cardiac models, which will take into account cardiomyocyte fiber anisotropy[31].

The measured field outside the torso is usually dipolar unless there are serious conduction abnormalities, see figure 1.6. The information content available is determined by the number and distribution of sensors. Some work has been done to calculate optimal sensor positions, specific for localisation of current distributions [32]. A planar array in front of the chest with magnetometers aligned to the superior-posterior axis (normal to the chest), is reasonably well aligned with the principal current vector. For example, Fenici's 37 channel planar array device is able to localise arrhythmic regions to within 6-7mm [33]. My aim is not to maximise localisation accuracy; as this is suited to a specific use case of MCG as an aid to ablation of fibrillation regions, but to capture sufficient information to reliably determine the presence of cardiac disease.

The amplitude of a dipole field decreases with radial distance from the source as $1/r^3$. Therefore it is essential to position the sensors as close as possible to the heart. In the case of a living animal, this is usually limited to placement outside the torso/thorax. Though

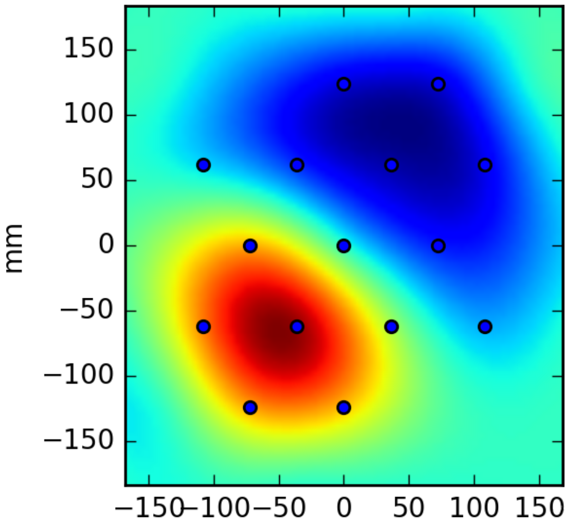

FIGURE 1.6: A MFM of a healthy volunteer during the depolarisation phase of the cardiac cycle. The small dark blue circles show the position of the sensors used to acquire the MFM; the false colour MFM is interpolated from these discrete sensors measurements.

the development of magnetrodes, sharp electrodes with micron scale magnetometer arrays on the tip, or mounted to a catheter probe, could be placed extremely close, even within the heart.

#### Cardiac Diagnostic Information

In general, the aim of MCG is to maximise the mutual information $I(X;Y)$ between the set of digital voltage time series that constitute the MCG $Y$ and the cardiac magnetic field, $X$.

$$I(X;Y) = H(X) - H(X|Y) \tag{1.2}$$

Where $H(X)$ is the entropy of the magnetic field and $H(X|Y)$ is the entropy (or uncertainty) of $X$ given knowledge of the magnetic field via the MCG, $Y$.

To increase the mutual information, $H(X|Y)$ can be decreased by the addition of extra magnetometers to capture spatial information, increasing the sampling rate to capture more temporal information and reducing environmental noise. $H(X|Y)$ is decreased by reducing the noise in the sensor system.

The broader aim of this PhD is to consider $Y$ to be the disease state of the heart (my concern at this point is binary; healthy or not). The diagnostic algorithm would then be considered part of the channel.

### 1.2.2 History of Magnetocardiography Devices

The first sensor used to detect the heart's magnetic field was a planar gradiometer composed of two 30cm long induction coils each containing two million turns of copper wire, Baule and McFee 1963. They conducted measurements in a field, "several dozen yards from the nearest source of interference"; mainly the 60Hz hum of the electrical energy distribution grid. This allowed them to observe an MCG albeit with a lot of noise, even

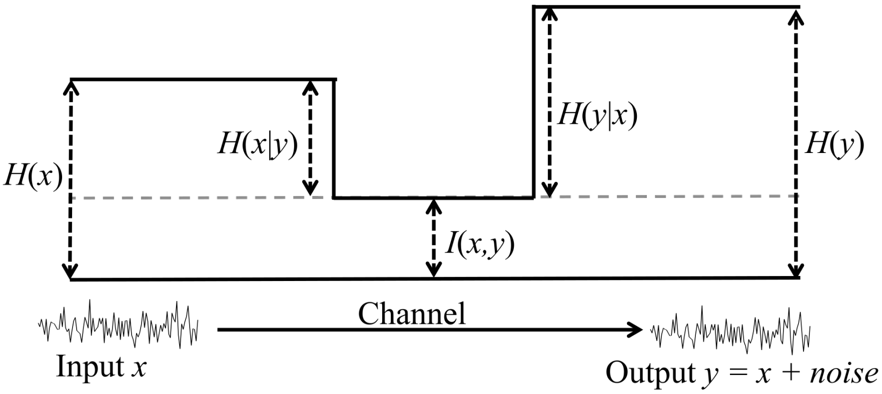

FIGURE 1.7: The mutual information of two variables across a noisy channel. [Image attribution: Jvstone]

after signal averaging. Cohen set out to improve the signal quality. He began by employing shielded rooms for his measurements, initially with copper coils and slightly lower noise electronic amplifiers[34, 35]. Then in 1969 he performed the first measurement with a SQUID sensor, and achieved an unprecedented signal to noise ratio (SNR) [36].

In 1982 induction coils were briefly revisited by Estola & Malmivuo who optimised an induction coil to capture the cardiac field [37]. They realised that advances in instrumentation amplifiers could allow comparable noise performance to SQUIDs, and achieved this with a sensitivity of $130\text{fT}/\sqrt{\text{Hz}}$ at 20Hz. However, their coil had a diameter of 112mm, making it impractical for forming a multichannel device.

A few years later the first magnetic field maps were recorded, by moving a single sensor to an array of positions over the chest [38]. However, these single sensor systems required precise repeated movement of the sensor, which was error prone and took a long time to acquire a magnetic field map, since the sampling time was a function of the spatial resolution. This made them impractical for clinical or deeper research use, so the field of magnetocardiography stagnated, awaiting the development of a multichannel system. The first commercially available multichannel device was developed by Siemens in 1990[39]. Following this, and a quick succession of other multichannel devices, the field of magnetocardiography grew. This can be seen by the sharp increase in the number of MCG papers published since 1990, see figure 1.8.

There are now approximately 50 papers published per year. A wealth of medical knowledge has been created regarding how to harness MCG for clinical diagnostic use. Multichannel SQUID devices for MCG are commercially available, however their economic paradigm has not changed significantly since the 1990's. They still have a high initial capital outlay in the millions of dollars and high running costs due to their liquid helium cryogenics. This cost is comparable with magnetic resonance imaging and x-ray computed tomography machines, which are useful for both diagnosis and treatment of heart disease; currently MCG cannot compete with the cost to benefit ratio of MRI and CT. However, if the cost of MCG can be lowered it has the potential to become a widely used clinical technology.

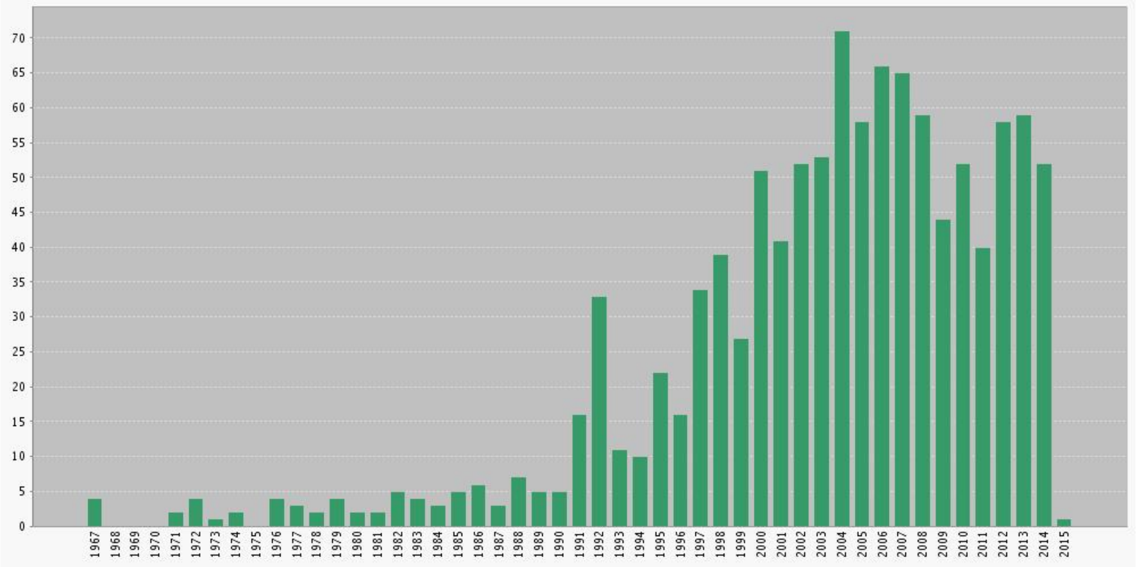

FIGURE 1.8: The number of papers published per year that are concerning magnetocardiography. From 1967 to 2015, data from WebofScience. Before the introduction of commercial multi-channel SQUID MCG in 1992 less than 5 papers per year were published. It is currently $approx$50 per year.

### 1.2.3 Magnetometer Development at Leeds

Our research began back in 2011, in the experimental quantum information lab where I was an undergraduate student.

Our first attempts at measurement involved a digital storage oscilloscope (DSO) and some commercially available coils. We also created our own coils, I 3D printed a bobbin with a square section centre bore in which we mounted a soft iron laminate for flux concentration. We wound the bobbin by hand with the aid of a winding machine in the departments mechanical workshop. It was extremely difficult to achieve a reasonable winding precision, which meant the resulting coils had a high variance in inductance, severely reducing performance when used for gradiometry. Our attempts to observe cardiac fields with these coils failed. Although we were deceived for a while by ballistocardiographic signals from the swivel chair we sat in during testing.

We needed better coils with a higher inductance and greater manufacturing precision. Ben devised an analytical procedure for optimal coil design [40] and arrived at a proposal geometry, which we had manufactured by Mundorf GmbH. These were the 7cm induction coils that are discussed most extensively in this thesis (see figure 2.1). To improve the data acquisition precision we employed a National Instruments DAQ (USB-6216), which could achieve 16-bit measurement, which was an improvement over the 12-bits we were getting from the DSO. To control this device Edward Reade-Banham and I wrote an acquisition program using Labview, including a user interface.

The amplifiers were designed by Phillip Thornton, based on his expertise in audio electronics he designed a circuit based around a low impedance amplifier, the INA217. In order to minimise the signal path to the amplifier input the amplifiers were mounted as close as possible to the coil sensors. Shima soldered the boards to the coils and we had our first useful induction coil magnetometers (ICM).

Since gradiometer arrangements are the typical approach for background noise suppression in MCG, we attempted to create a gradiometer with a pair of ICM. The coils were held coaxially with a 160mm separation between the coil centroids; within an acetal plastic tube, which had a thread machined into its outside surface as a guide for winding a global coil. Signals could be fed through this global coil, which, because of its symmetrical arrangement, would be picked up by both coils equally. Chris Symmonds created an active coil matching setup using a signal generator to create a 1khz signal, which was fed through the global coil, and a DSO to capture the detection coil signals and perform FFT. The idea was to modify the individual coil gain until measured amplitudes of the 1kHz global field were matched. However, this approach was not fruitful; the relative gain changed too quickly to allow a stable match and the resulting background noise suppression performance was poor.

The next evolutionary step was the creation of a planar array. Four ICM were mounted in a close packed square arrangement within a block of engineering plastic. Our lab was an adversely noisy environment, which made capture all the more difficult. However, this was as much a blessing as a curse, since we were forced to confront noise mitigation early on. To determine the performance of the ICM in the absence of environmental noise, we journeyed to York Neuro-imaging centre (YNIC) to test our 4-coil device in their Magnetically shielded room (MSR). The results were fantastic and unambiguous; with these coils, amplifiers and ADC system we could observe the heart's magnetic field in real time. From this point on, development became mostly a matter of handling the external noise.

Based on the work of others it was clear that we needed a higher spatial resolution to achieve useful magnetic field maps. We acquired a 16 channel ADC from national instruments and created a 15 channel detector array; leaving one channel for the synchronous recording of an ECG. I suggested a hexagonal packing arrangement since this achieves the highest packing efficiency for the cylindrical sensors, a schematic representation of the result is given in figure 1.9. Brian from the mechanical workshop then designed an array, which he created by milling 19 bores into a chunk of acetal plastic, including accommodations for wiring. This block was capped off underneath with a 5mm sheet of acetyl that supported the coil's mass. The assembly was covered with a custom acrylic hood manufactured by Barton Plastics, Leeds. This assembly rested on aluminium angle brackets connected to a floor mounted frame constructed from square section (t-slot) aluminium extrusion that rolled on lockable castor wheels to allow positioning. As seen in figure 1.10.

We began to extract MCG's reliably in our lab. Though it was a difficult environment due to the presence of a large mains distribution board and un-interruptible power supply in the same room, and a series of high power (10-100KW) switch mode power supplies in the room opposite. We focused more on environmental noise removal routines, and created a user interface. A second trip to YNIC yielded some excellent 12 channel signals (faulty wiring causing the loss of some sensors)

After preliminary pilot studies in our lab with young healthy volunteers, the next step was to perform a clinical trial, which included cardiac patients.

### 1.2.4 Clinical Trials

The first clinical trial was carried out at Leeds General Infirmary. Permission to use human subjects in the collection of data was granted by the University of Leeds ethics committee (ref. no. MEEC 12–034). This trial is known as the 6060 trial since 60 patients and 60 healthy volunteers participated, each having four scans, creating a total of 480 scans. The aim of the trial was to assess the predicitive capability of the device and to

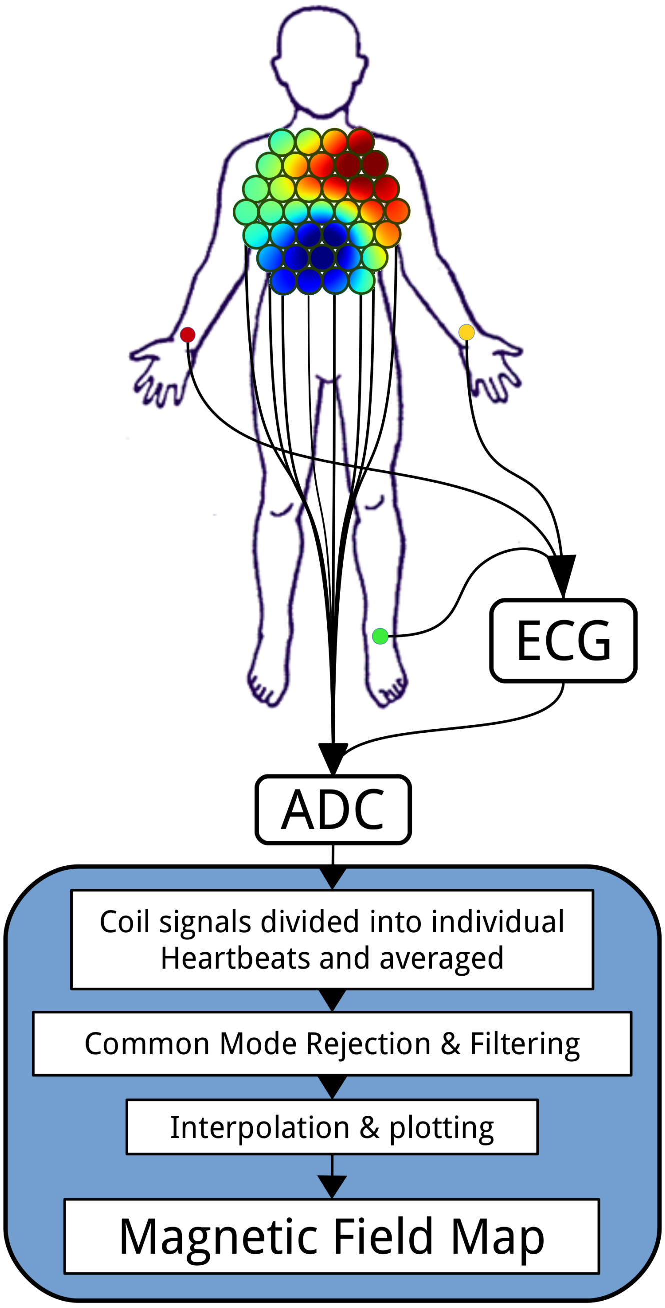
ECG
ADC
Coil signals divided into individual Heartbeats and averaged
Common Mode Rejection & Filtering
Interpolation & plotting
Magnetic Field Map

FIGURE 1.9: A flowchart showing the capture of MCG signals. The false colour map over the sensor array demonstrates the approximate scale of the magnetic field relative to the torso. MCG and ECG signals are synchronously recorded by an ADC and sent to a computer (blue box) for digital signal processing. The end result is a series of magnetic field maps, one for each sample acquired in the cardiac cycle.

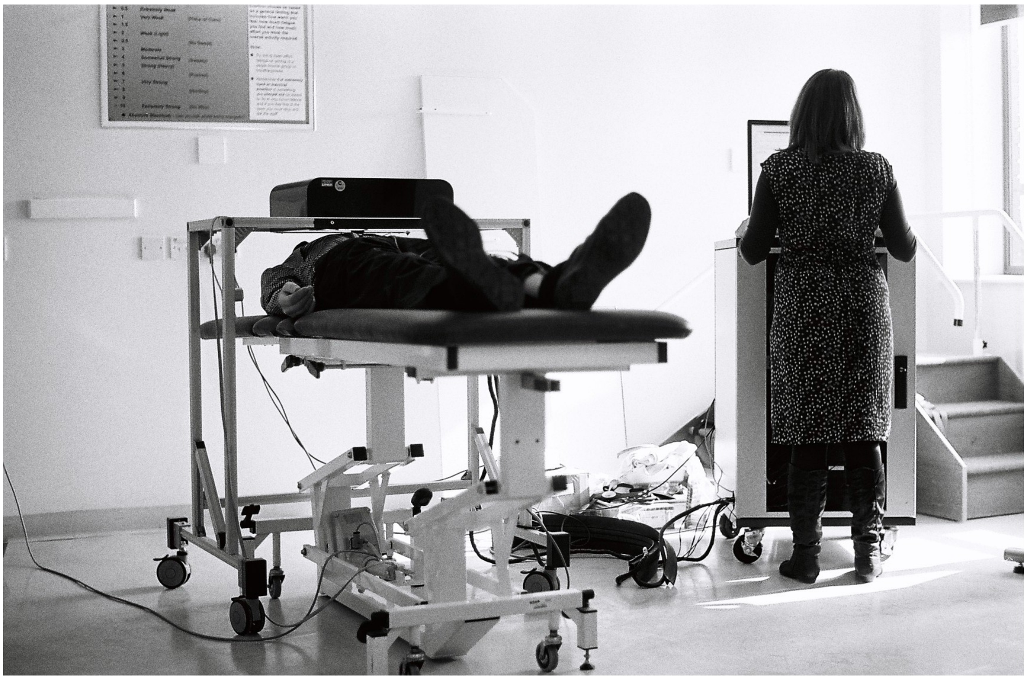

FIGURE 1.10: The MK1 device in operation during the first clinical trial at Leeds General Infirmary. Research Nurse Lorraine Falk can be seen operating the device. While the scan subject lies supine underneath the detector array.

understand the variability between multiple scans of the same person. Each person was scanned twice by Loraine and twice by Roo, allowing the inter-operator variability to be assessed. The scans were performed on two separate occasions, separated by a week or more. The trial took over a year to complete, since the patients were recruited from the cardiology clinical pathway at the hospital and had to commit to a second scan one week later. A second trial was carried out in the QIX lab, taking scans of 50 healthy young students. This trial was known as the young healthy controls (YHC), and provided a highly polarised view of healthy cardiac condition; in comparison with the age-matched healthy volunteers of the 6060, many of whom were not in optimal health.

Shima Ghasemi-Roudsari was responsible for managing the data collection process, which she writes about in her thesis and a recently published paper [41, 42]. The device operation was performed by two research nurses, Loraine and Roo. I was there for backup and provided the technical support. A couple of years later, with this dataset in hand I began this PhD.

## 1.3 Thesis Outline

The thesis is divided into three main chapters, plus a final discussion chapter. The first chapter documents the development of induction coil magnetometers, their optimised integration into a detector array and the formation of a complete magnetocardiography device including user interface. The second chapter details the environmental noise encountered by an MCG device in a clinical environment and develops techniques for analysis and filtering; A strong focus is given to the spatial coherence of environmental noise and coherent noise rejection. The third chapter covers the extraction of diagnostic information from the obtained MCG signals, using a spectrum of techniques from classical parameter extraction to convolutional neural networks.

### 1.3.1 Criteria for success

A successful thesis will document the creation of a low cost, portable MCG system suitable for unshielded operation in a clinical environment. The performance of this system will be determined primarily by the overall diagnostic performance. In particular, the thesis focuses on environmental noise rejection performance and will provide an understanding of the specific noise handling requirements for MCG, and effective techniques for noise rejection.

# Chapter 2

# Device Development

This chapter covers the architecture of the MCG system, from the individual magnetometers through their assembly into a magnetic field mapping array to a complete medical device, including software for capture and analysis.

The development of the device would not have been possible without the technicians in the electronics and mechanical workshops of the physics department. In particular, Phil Thornton created the pre-amplifier circuits and their power supplies; choosing to use a low-impedance amplifier that was ideally suited. The printed circuit boards (PCB) were populated by Shima Ghasemi and myself. Initial testing of the ICM was performed by Chris Symmonds, including gradiometer configurations. The array mounting hardware was designed and built by Brian Gibbs. The COMSOL modelling of the sensor array was performed by Nick Pawlowski and myself. The Acquisiton user interface was coded by Edward Reade-Banham and myself. The rest of the work was carried out by myself in collaboration and under the supervision of Prof. Ben Varcoe.

## 2.1 The Induction Coil Sensor

Induction coils were the first sensors employed for magnetocardiography, by Baule and McFee in 1963. They used a pair of 30cm long coils, each with 2 million turns, wired together in a gradiometer configuration. These coils needed to be large in order to achieve a high response, so that the signals could be observed with the electronics available at the time; the magnetocardiographic coil voltages were around 30μV before amplification. However, the large individual detector size made the devices impractical; due to the high mass and the distortion caused by the permeable cores. Also the hardwired planar gradiometer configuration was not suited to magnetic field mapping. Induction coils were soon superseded by superconducting quantum interference detectors (SQUID), which can achieve excellent spatial resolution and sensitivity. However, their requirement for cryogenics and shielding makes SQUID sensors unattractive.

Due to advances in electronic technology since the 1960's, amplifiers and analogue to digital converters (ADC) are now capable of operating in the nanoVolt regime. This capability allows the operation of induction coil magnetometers with a lower response and therefore smaller size. It also allows signal processing to be performed digitally, allowing virtual gradiometer arrangements, zero-distortion filtering and advanced noise removal algorithms.

To design a new induction coil magnetometer we first determined the response required to match the cardiac signal to the electronic detector system. Then with the response as a constraint, we optimised the geometry to achieve sufficient response with minimal inherent noise, high directionality and small diameter (for high spatial resolution). The inherent noise determines the sensitivity limit of a detector. For real-time signal capture the sensitivity limit must be smaller than the signal of interest. When signal averaging the inherent noise reduces in proportion to the number of averages, n, by $\sqrt{n}$. In the case of a typical 10 minute scan containing 600 heart beats, the inherent noise is attenuated by a factor of 25x. This could potentially relax the sensitivity requirement. But to operate in an unshielded clinical environment, the averaging must be used to attenuate the superimposed environmental noise, therefore the sensitivity must still be a fraction of the signal amplitude.

An induction coil magnetometer will have an output voltage determined by equation 2.1.

$$V = AN\frac{dB(t)}{dt} = ANB2\pi f \tag{2.1}$$

Where $N$ is the number of windings, $A$ is the effective cross sectional area of the coil, $B(t)$ is the time varying magnetic field with magnitude $B$ and $f$ is the frequency of oscillation of the field. We know approximate values for $B$ and $f$, this provides a constraint on $A$ and $N$.

The sensitivity limit $S$ of the coil is determined by the Johnson-Nyquist noise resulting from the thermal distribution of the electrons within the windings[43, 44]):

$$S = \frac{\sqrt{4K_BTR_w}}{2\pi fNA} \tag{2.2}$$

Where $K_B$ is Boltzmanns constant, $T$ is the temperature, $R_w$ is the winding resistance given by 2.3:

$$R_w = \frac{2N\rho r_{coil}}{a^2} \tag{2.3}$$

Where $a$ is the wire radius and $\rho$ the resistivity of wire used for the windings of a circular coil of average winding radius $r_{coil}$. A combination of equations 2.1 and 2.2 can be used to determine the coil structure with the lowest noise level given the design constraints.

If the coil parameters are length $L$, the coil outer diameter $D$, and the coil inner diameter $D_i$, the dimensions that give the lowest noise level have the ratio $D_i : D = 0.425 : 1$. In addition to minimising the noise we primarily want to measure the component of the magnetic field aligned to the axis of the coil. This is because the radial magnetic field components are larger than the tangential, and correspond to tangential EMF that are mostly silent in ECG [45]. The optimal structure to measure the axial component is achieved when $L/D = 0.69$ for the above ratio $D_\mathrm{i}/D$. The absolute value of the coil diameter is determined according to the desired device resolution, leaving the radius of the wire as the only remaining free parameter in the coil design.

We chose a 7cm coil diameter, which allowed a hexagonal arrangement of 19 coils over the active area of the torso. This resulted in a coil of 5.4cm length, 2cm core diameter, with approximately 4000 turns and weighing about 1kg, as seen in figures 2.1 and 2.2. The soft iron core used gave a gain in response of about 3x. The resulting coil had an inductance of $1148 \pm 15$mH, a resistance of $60 \pm 2\Omega$ and a theoretical sensitivity of $197\mathrm{fT}/\sqrt{\mathrm{Hz}}(@10\mathrm{Hz})$.

The coil was connected directly to a pre-amplifier circuit with the INA217 integrated

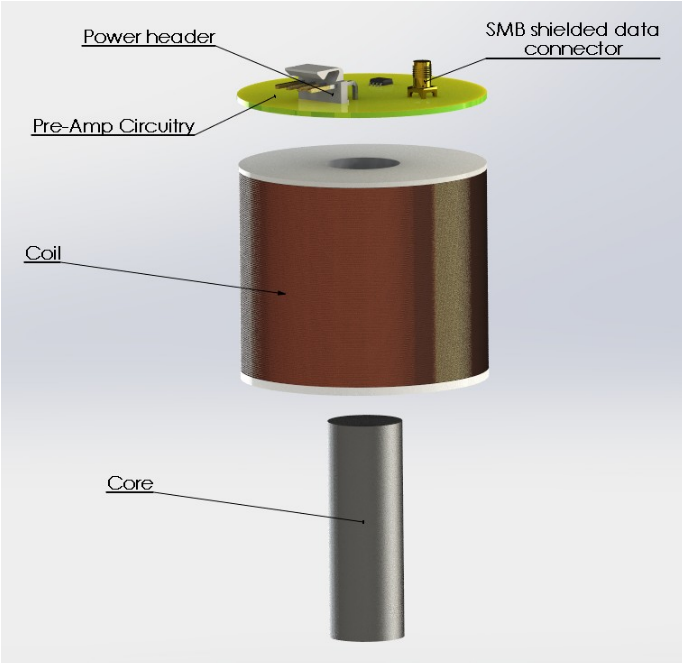
Power header
SMB shielded data connector
Pre-Amp Circuitry
Coil
Core

FIGURE 2.1: Annotated rendering of the induction coil sensor with pre-amplifier circuit board, showing an exploded view along the coil's length.

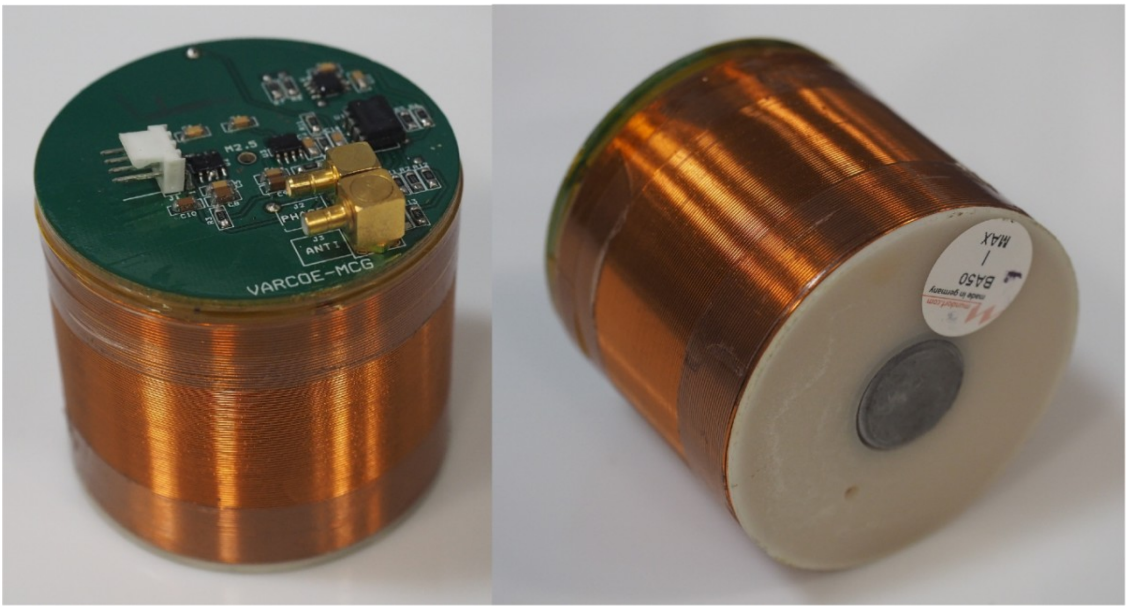
VARCOE-MCG
ANT1
BA50
MAX

FIGURE 2.2: Photograph of the (7cm diameter) induction coil sensor. The 2cm diameter core is visible in the centre of the plastic coil bobbin. The pre-amplifier PCB is mounted to the coil bobbin with three silicone feet and its perimeter taped securely with kapton tape. DC battery power is provided via the white 4-way header. The amplified signals are output via the gold plated SMB connector into a coaxial cable.

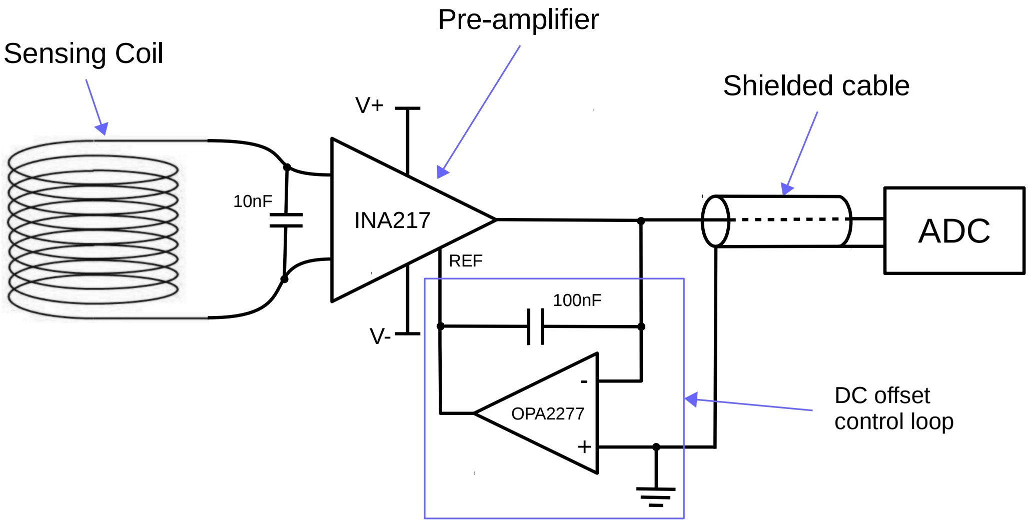

FIGURE 2.3: Diagram of detector system circuit. A complete diagram including the power supply components is located in the appendix A.1

circuit (IC) at its core, see figure 2.3. Current induced in the coil creates a potential difference across the high impedance amplifier input. This input voltage is magnified by the amplifier, which creates an output signal with a 1000x gain. In order to minimise the un-amplified signal path the pre-amplifier was located on top of the magnetometer. The output signals were carried by shielded coaxial cables to the ADC inputs.

The amplifier was powered using batteries and a linear regulator IC (LM7805), which provided a stable, low noise voltage with an estimated 40µV of output noise, equivalent to 40nV of input noise. This noise level is about equal to the ADC input noise and a fraction of the sensing coil noise (see section 2.1.2), though this is acceptable, there are now equivalent IC available with an order of magnitude lower noise, such as the TPS7A4700. A 10µF capacitor was used in parallel with the output, increasing this capacitance reduces the regulator noise but increases the inrush current. This capacitance should be increased as much as possible within the stable operation of the circuit. Also capacitors should be chosen that are immune to acoustic noise, that can occur via piezo-electric transduction.

We developed a 2nd generation smaller sensor that had a 4.6cm diameter, see figure 2.4. Each sensor weighed 330g, which considerably reduced the total mass of the array. This smaller diameter allowed us to create a device with 36 sensors in the same area as 19 large coils, which matches the typical resolution used by conventional SQUID MCG. The same pre-amplifier circuit was used, with a compressed PCB layout to fit the smaller area. I reflow-soldered the surface mount components, then tested each circuit before soldering the through hole connections to the coils. The PCBs were affixed to the coil in the same manner as the large coils; with three silicone pads between the board and the coil and kapton tape around the circumference. These sensors had a reduced response, but also a lower inherent noise. We have found them sufficient for mapping the QRS MCG, though the achievable SNR makes T-wave observation infeasible.

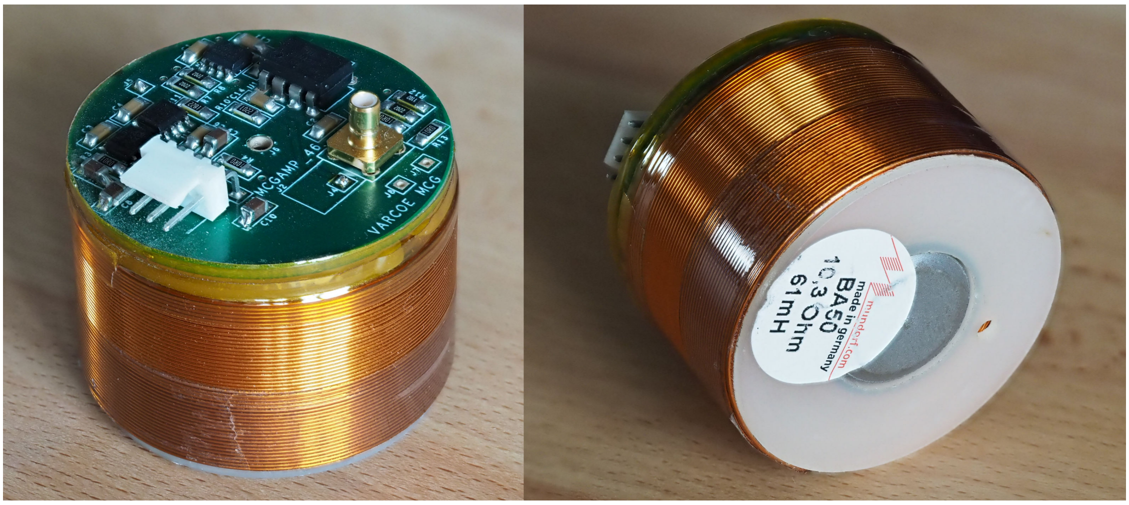

FIGURE 2.4: Photography of the "small coil" Induction Coil Magnetometer; 4.6cm diameter.

### 2.1.1 Calibration of Sensor Response

The sensor outputs a voltage proportional to the change in magnetic field. In order to determine the magnetic flux density from this voltage we must know the sensor response. All of my diagnostic research has used sensor voltage units, because I have not had to transfer a diagnostic algorithm between different device datasets. Knowing the absolute magnetic field magnitude is useful for creating device agnostic diagnostics and comparing the sensor performance with different sensor technologies.

The induction coil outputs a signal proportional to the number of coil turns, its effective cross sectional area and the gradient of the magnetic field, see equation 2.4. The effective cross-sectional area is difficult to calculate accurately since it is dependent on permeability of the core and the relative angle of the incident magnetic field. Also there is a non-linear frequency dependence due to the resonance between the inductance and parasitic capacitance, however the response is constant in the typical frequency range used for MCG. Therefore, to determine the magnetic flux density from the measured voltage time series the coil response, R, must be mapped out.

$$V = \frac{fB}{R(f)} \tag{2.4}$$

To map the response I applied a series of sinusoidally varying uniform magnetic fields of known magnitude and frequency to the sensor and recorded the resulting amplitudes of the induced voltages. To ensure a uniform field the sensor was placed in the centre of a pair of 30cm diameter Helmholtz coils. A sinusoidal voltage output from a signal generator was applied to the helmholtz coils and current limited using a series resistor. The resulting Helmholtz field was measured using a calibrated fluxgate magnetometer (Bartington Mag-01H with B-probe). The fluxgate probe was replaced with the induction coil sensor and its output voltage was measured at several applied frequencies using a DSO. The measured response was 290 fT $\mu$V$^{-1}$ at 30 Hz and 813 fT$\mu$V$^{-1}$ at 10 Hz.

The response was constant up to ~350Hz, where the signal begins to resonate with the tank circuit, as can be seen in the theoretical bode plot for this circuit shown in figure 2.6. The tank circuit is a resonant circuit created by the $C$ = 10nF capacitor in parallel

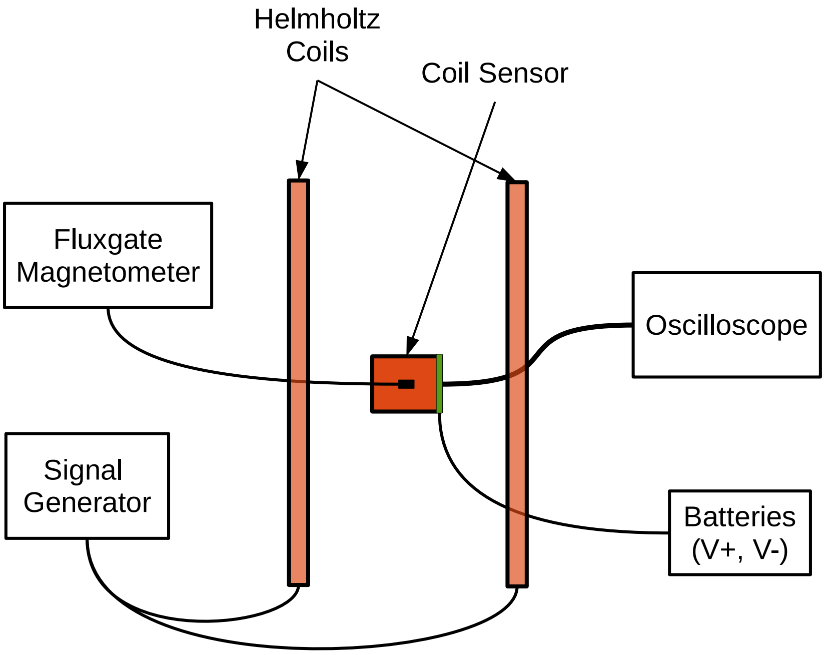

FIGURE 2.5: Schematic of experimental setup used to measure the sensor response. When measuring the response of an air core coil the fluxgate could remain in the centre for simultaneous measurement. The two coils are wired in series.

with the coil and pre-amplifier, see figure 2.3. It acts as a low pass filter, with a cut-off frequency equal to the resonant frequency of the circuit, $f_0$:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \tag{2.5}$$

Signals above this cut-off frequency are attenuated. The inductance of the 7cm coils is $L = 1148 \pm 15$mH. Therefore they have a resonance at $f_0 = 1485 Hz$, using equation 2.5. There is also a parasitic capacitance that exists between the coil windings, which has an associated resonance. However, I estimate it to be on the order of 10pF, with a corresponding resonance at 225kHz, making it insignificant compared to the effects of the filter capacitor.

The frequency response of the tank circuit can be mapped with a bode plot, see figure 2.6. The bandwidth of the resonant peak is dependent on the series resistance, which for the large coils is $60 \pm 2\Omega$. Higher resistances increase the filter bandwidth, reducing the sharpness of the cut-off. The Q-factor for the 7cm coils is $Q = 178$. The response is approximately flat up to 200Hz, which is sufficient to capture cardiac signals without distortion, though might cause a minor amplification of the highest frequency components in QRS, which can reach 300Hz.

In the small coil sensor (4.6cm) the same 10nF capacitor was employed. However, the coil inductance was much lower at approximately 67 4mH. This results in a higher resonant frequency of $f_0 \approx 6130 Hz$. The bode plot for the small sensor can be seen in figure 2.7 This puts the cut-off frequency above the sampling rate of the ADC. As opposed to the large coil, where the cut-off is within the ADC sampling frequency. Though I have

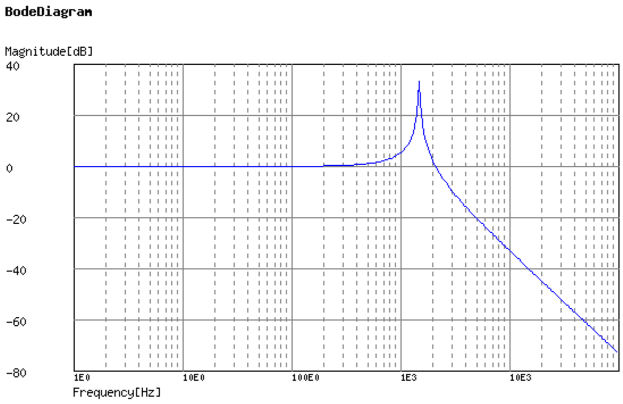

(c)okawa-denshi.jp

FIGURE 2.6: Bode diagram showing the frequency response of the RLC filter circuit. The response is constant from 0 Hz to 200 Hz. The resonant peak occurs at 1485Hz, which is within the frequency range of the ADC.

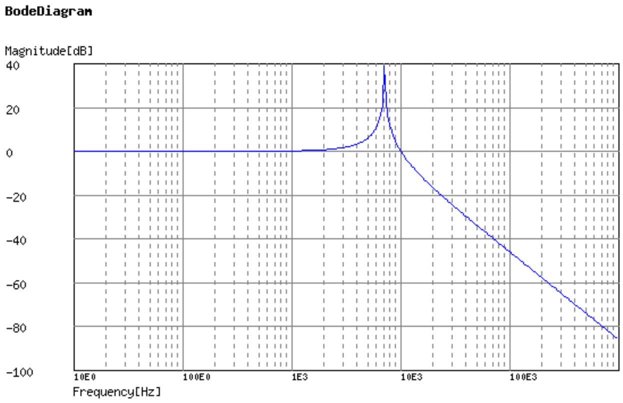

(c)okawa-denshi.jp

FIGURE 2.7: Bode diagram showing the frequency response of the small coil sensor RLC filter circuit. The resonant peak occurs at 6130Hz. The response is nearly constant throughout the sampling range of the ADC (2.4kHz).

not directly compared the two approaches, it seems better to maintain a linear response within the acquisition range and avoid the resonant peak, since it could reduce the useful dynamic range by amplifying resonant environmental noise.

At higher frequencies amplifier slew can cause distortion. The slew rate defines the maximum rate at which the amplifier output can track the input. The amplifier chip we used has a slew rate of 15V/µs. The maximum frequency that can be measured without slew rate error given by equation 2.6:

$$f = \frac{SlewRate}{2\pi VoltageRange} \tag{2.6}$$

Thus the maximum frequency that can be measured over the full voltage range of the amplifier ($-5$ V to 5 V) is 240kHz. This is well above the typical MCG frequencies and the sampling rate of the ADC 2400Hz. Though high frequency noise is not directly measured by the ADC it can still effect the sampled voltage because the ADC input contains capacitors that can integrate the high frequency components. Since slew distortion can be non-linear, the distortion of high frequency noise will reduce the coherence of the high frequency noise. However, the presence of the low pass filter formed by capacitor in parallel with the amplifier input prevents the formation of slew distortion.

### 2.1.2 Sensitivity

The sensitivity of a detector is limited by its inherent noise. In the case of an induction coil detector system the inherent noise is due to the thermal movement of electrons within the conductors of the wire, pre-amplifier and the analogue components of the ADC. This thermal electronic noise is known as Johnson-Nyquist noise [43, 44].

The noise contribution from the ADC is $28\mathrm{nV}/\sqrt{\mathrm{Hz}}\,(@10\mathrm{Hz})$. The pre-amplifier noise is a negligible $3.5\mathrm{nV}/\sqrt{\mathrm{Hz}}\,(@10\mathrm{Hz})$. The theoretical sensitivity of the coil is $241\mathrm{nV}/\sqrt{\mathrm{Hz}}\,(@10\mathrm{Hz})$, as given by equation 2.2. The terms in his equation are well defined except for the effective cross sectional area. It is approximately equal to the cross sectional area of the coil, but is increased by the addition of the soft iron core that has a high permeability. The actual effective cross sectional area is difficult to measure, it is dependent on the presence of neighbouring permeable material and the relative angle of the magnetic field to the core. However, if the sensor is isolated from environmental noise then the inherent noise can be measured.

The sensitivity was measured experimentally in the magnetically shielded room (MSR) at York Neuro-Imaging Center (YNIC). This MSR is typically used for Magnetoencephalography (MEG) measurements, its noise level is therefore very low, probably around $5\mathrm{fT}/\sqrt{\mathrm{Hz}}$. The sensor was recorded for 30 minutes. From this time series a FFT was computed to determine the noise power spectral density, see figure 2.8. The sensitivity of the detector system was $127\mathrm{nV}/\sqrt{\mathrm{Hz}}\,(@10\mathrm{Hz})$ or converted to units of Tesla: $104\mathrm{fT}/\sqrt{\mathrm{Hz}}$. The ADC contributed to about 25% of the noise. This is sufficient for MCG, since the QRS peak is approximately 50pT.

## 2.2 Alternative Sensing Technologies

There are many types of magnetometers in addition to the ICM. SQUID have of course proven themselves as excellent sensors for biomagnetism. However, their high cost, and the cryogenic requirement is unlikely to change any time soon. If we want to see a step

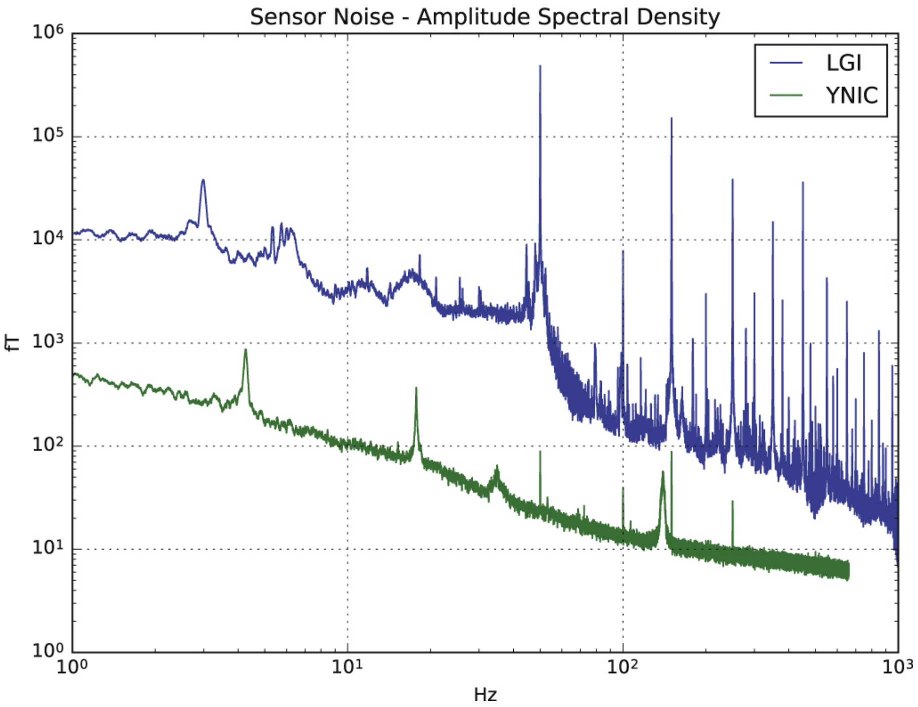

FIGURE 2.8: Noise power spectral density in the MSR at YNIC compared to the clinical environment at LGI.

change in the field of magnetocardiography, sensors that can produce a lower cost system are needed. Even if the individual sensor cost is higher, the ability to operate in a portable device, outside of a MSR would change the economics of MCG, potentially facilitating to a new wave of adoption.

Recent developments in miniaturisation and packaging of optically pumped magnetometers (OPM) are beginning to excite biomagnetism researchers. OPM works by probing a vapor of alkali atoms with a laser beam. The applied magnetic field causes energy level splitting in the atoms (zeeman effect), which can be detected by a shifting of the resonance peak as probed by the frequency dithered laser [46, 47]. In this configuration OPMs can just about reach sufficient sensitivity for magnetocardiography. A second operating regime called spin exchange relaxation free (SERF) magnetometry is much more sensitive but must be operated in zero magnetic field, which requires a shielded room with active cancellation of the Earth's field.

Another type of optical magnetometer is the diamond nitrogen vacancy (NV) magnetometer. This consists of a diamond molecule, with pairs of carbon atoms removed, one replaced by a nitrogen atom and the other left vacant, but with an extra electron implanted. This system is similar to the alkali vapor OPM, the NV system has energy levels that are susceptible to the zeeman effect and can be probed by a laser. However, the robust diamond lattice surrounding the NV allows orders of magnitude higher sensitivity. This technology is new however, and it has not yet been implemented in a package suitable for magnetocardiography, though this is being researched [48].

Tunnelling magneto-resistance (TMR) magnetometers are a class of solid state device that is commonly used in the read heads of hard drives to measure the information encoded in magnetic domains on the hard-disk. They have recently been used for MCG in a MSR. The high inherent noise gave a SNR of about 1, with respect to the R wave peak.

Signal averaging was required to observe the MCG clearly [49, 50].

Flux gate magnetometers are based on induction coils, but use an applied dithering field in order to measure absolute magnetic field. The extra electronics involved with this means that their noise is generally higher than regular ICM. On the other hand, over a narrow frequency band, the lock in amplification effect provides excellent noise rejection. Recently a commercial three axis flux gate magnetometer from Bartington was used to measure an MCG within a shielded room [51].

I experimented with the creation of a planar induction coil. The planar geometry makes it sensitive to a broad range of field vector angles, which could be seen as a benefit or not. The coils I developed were a spiral with 4µm traces. My plan was to have these laser micro-machined onto PCB material, and stack several layers (with interlinks between the layers), to produce a coil with several thousand turns but weighing grams. Sadly I was unable to have them manufactured, as the available equipment was only capable of raster movement, it did not have a rotating stage that would allow precision manufacture of the circular micron scale geometry. The benefit of planar coils would have been lightness, potentially lower manufacturing cost, especially when considering system integration. Most of all, the ability to produce highly matched coils for excellent background subtraction.

## 2.3 System Architecture

This section documents the construction of the device from the magnetometer array to the user interface, and covers the important design elements that we have found by iterative prototyping and theoretical analysis. The primary focus is on the Mk1 system, which uses fifteen 7cm coil sensors (2.2). It also covers the evolution into the Mk2, which used the smaller 46mm sensors (2.4) and further on the commercial prototypes, Alpha and Beta.

All the devices share a similar overall architecture, with the detector array separate from the other hardware. The computer, acquisition system and power electronics are housed in a standalone enclosure unit, with analogue power and shielded signal cables connecting it to the detector array (with integrated pre-amplifiers). This design exists to isolate the sensing region from the electromagnetic noise of acquisition components and the vibration from the operator interacting with the computer system. In addition to optimising the measurement performance, the device must operate in a clinical environment, therefore its design must consider hygiene, electrical safety, physical safety.

First I will cover the development of the detector array, then the analogue power supply and finally the digital acquisition system.

### 2.3.1 Detector Array

The induction coil sensors have already been covered in detail in 2.1. But the formation of the individual sensors into a detector array, their power supply and signal transmission warrant further detail.

For the Mk1 device we created the 19 coil sensor array in a hexagonal arrangement, see figure 2.9. Hexagonal packing provides the highest packing density, and therefore the highest spatial resolution for magnetic field mapping. Despite having room for 19 magnetometers, the ADC we were able to acquire had only 16 channels, one of which was used for synchronous ECG recording. The 15 coil limitation enforced an asymmetrical

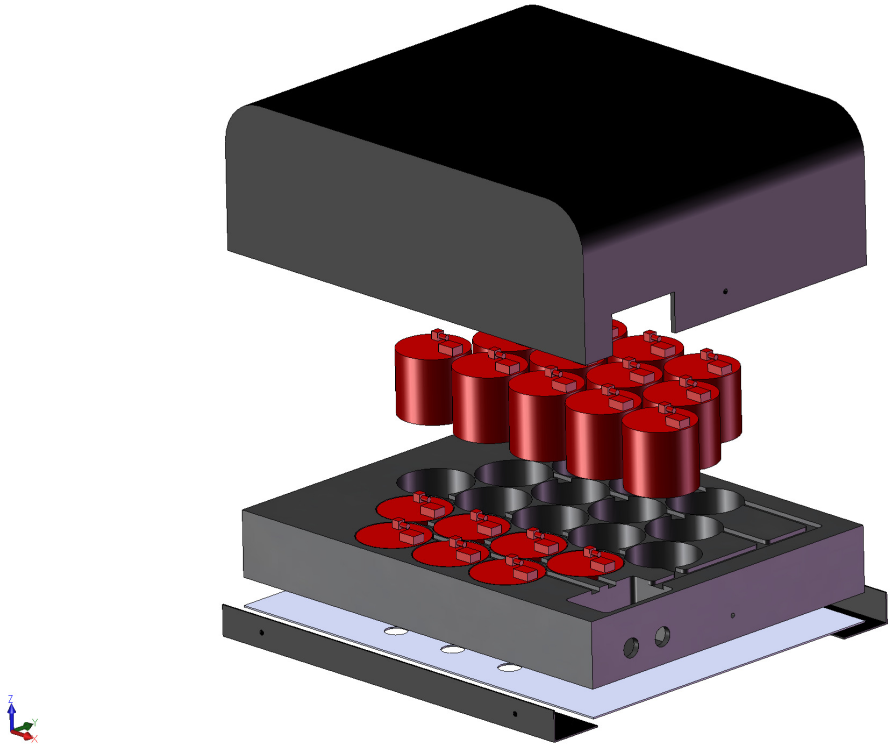

FIGURE 2.9: Exploded view CAD rendering of the detector array used in the MK1 MCG device. The magnetometers are shown in red, with bounding boxes for the power and signal shown. The main array holder is formed from solid Acetyl plastic. The acrylic hood sits on top of the array, with a small cut out for the cable glands. The whole assembly is mounted on two aluminium right-angle extrusions that are bolted to a floor standing aluminium frame.

sensor arrangement; we chose to give the highest sensor density along the axis perpendicular to the principal vectorcardiographic (ECG) component, which is approximately the principal axis of magnetic activity.

The array was constructed from Acetyl; a strong and stable engineering plastic (that is dielectric). The sensors were closely packed with gaps of only about 1mm, and shimmed with kapton tape to ensure they fit tightly in their bores. This setup results in all the sensors becoming acoustically coupled, which is essential for acoustic noise rejection. The array was covered with an acrylic hood, for physical protection and hygiene (cleanable surface). It also ensured isolation from air currents and solar induced heating, creating a stable temperature environment, ideal for keeping a consistent gain matching between the sensors.

### Array Design Improvements

The move to smaller (4.6cm) sensors allowed the creation of a new detector array with a higher spatial resolution; both in density and total number of sensors. I designed a new array containing 37 coil sensors, again in a hexagonally packed arrangement, see figure 2.11 for a comparison with the Mk1 array. The array extent was slightly larger in x and y, with a slight squeezing of the width to fit within typical torso widths.

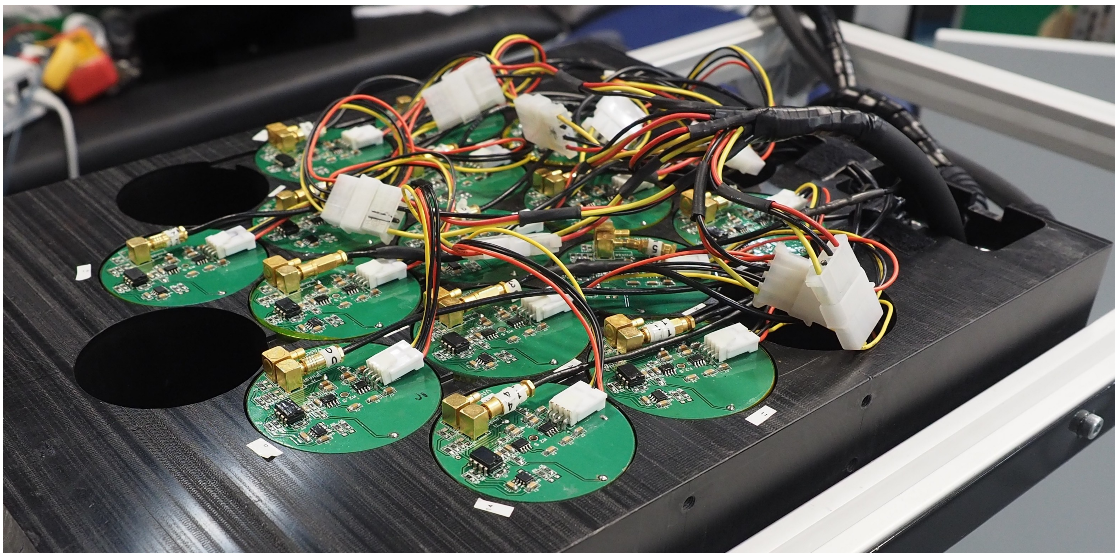

FIGURE 2.10: Photograph of the Mk1 detector array. The pre-amplifier circuit boards are visible (green PCB's). Analogue power distribution is provided by the loose parallel wiring.

The array was designed to be mounted on the end of an arm, because of this it was vital to minimise the array mass in order to reduce the tip-over angle. My intention was to 3D print it from Nylon, to maintain stiffness whilst including internal voids to reduce mass. I created a 3D CAD model and used a finite element analysis simulation in order to ensure a reasonable level of stiffness, see figure A.5. Further mass optimisation could have been performed, using an evolutionary algorithm, however the overall design showed minimal projected weight savings; since the mass of the copper within the coils constitutes the majority of the array mass. Since the increased cost of this approach did not provide sufficient benefit it was decided to machine the array out of Acetyl

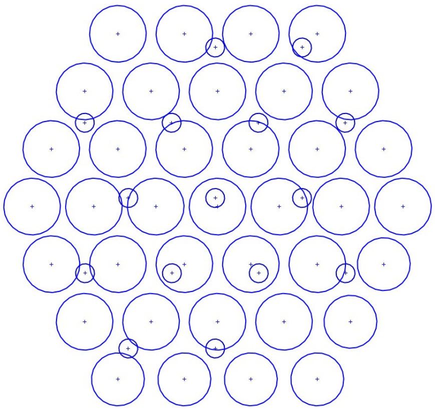

FIGURE 2.11: Comparison of the Mk1 detector array (small circles) with the 37 sensor array (larger circles). The sensor spacing is 56mm accept in the centre row that has 52mm spacing. A dimensioned diagram is in the appendix A.4.

### Pre-amplifier Noise

The ICM pre-amplifiers were powered by a common dual rail ±12V power supply, connected to the power source via a screened multi-core umbilical cable with final distribution through loosely parallel conductors as seen in figure 2.10. The power source consisted of a pair of 12V lead acid batteries located in the standalone enclosure, see figure 2.12. A key switch on the front of the enclosure was used to switch between scanning and charging modes. In scanning mode the batteries could power the amplifiers for about 4 hours.

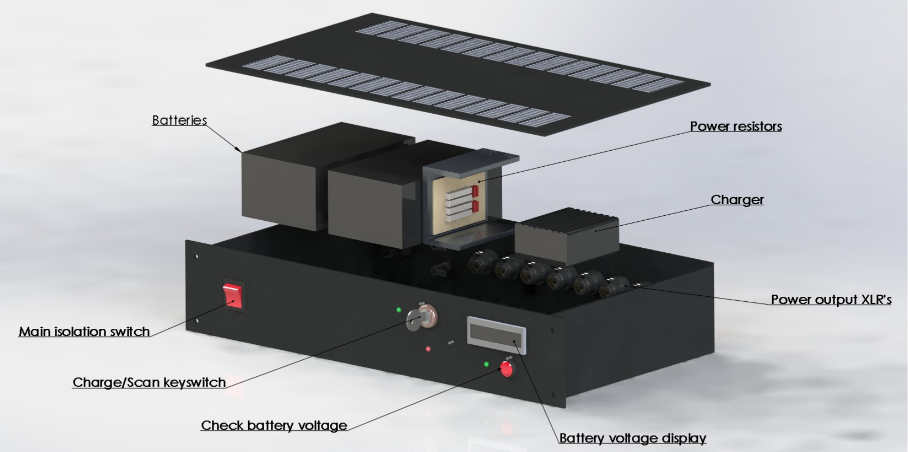

FIGURE 2.12: Exploded view CAD rendering of the battery power supply for the pre-amplifiers. The key switch toggles between charging and scanning modes. The battery voltage can be checked by pressing the button and reading the LCD screen. The internal components are galvanically isolated from the rack-mount enclosure for safe hospital use.

The current required for each pre-amplifier varies up to $\sim$ 100mA, depending on the output signal amplitude. This current induces a magnetic field that can be a significant noise source for the ICM. The magnitude of this induced field can be estimated using equation B.2. For tightly coupled parallel conductors the induced magnetic field is $\sim$ 800pT at 10cm distance. However, the wires were not tightly bundled, and were spread apart to enter molex connectors (2.54mm spacing on the small end, 5mm spacing on the large connectors); Where the wires become separated the fields produced by the opposing currents no longer cancel as effectively and the net magnetic field increases, with a 5mm gap the field at 5cm increases by a factor of 5x. This field has a large enough magnitude to effect ICM measurement.

The chaotic arrangement of conductors used for analogue power distribution in the Mk1 allowed the induced magnetic fields to destructively interfere, with a pseudorandom remainder, see figure 2.10. Ideally twisted pair (or coaxial) conductors should be employed to reduce the induced field by an estimated 3-100x, dependent on the ratio of the helix pitch to wire separation [52, 53].

In the Alpha device power and signal distribution were handled by a single large PCB located directly above the array. This PCB split the incoming power rail into 37 individual sensor supplies. The long straight traces radiated the power rail noise, producing a structure on the MFM. In an attempt to shield this local source, I manufactured a 6mm copper plate and placed it between the sensors and distribution PCB, see figure 2.13. Though this removed the distribution board pattern, the SNR was greatly worsened as

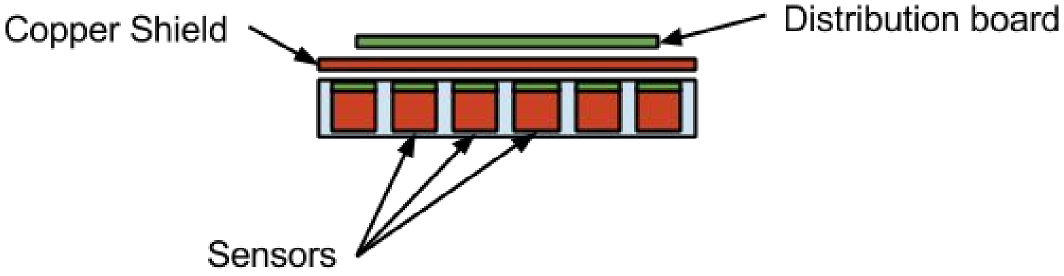

FIGURE 2.13: Experimental setup for testing the efficacy of shielding the acquisition electronics from the sensors. A 6mm thick copper plate is placed between the sensor array and distribution board that contains the acquisition electronics.

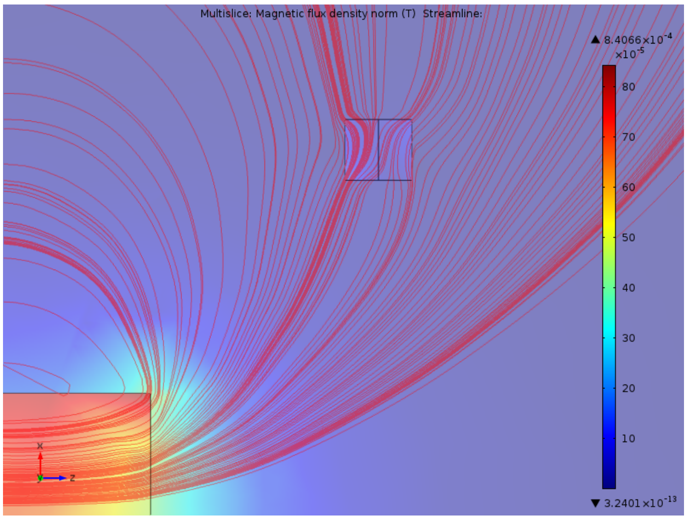

FIGURE 2.14: Comsol simulation showing magnetic flux density streamlines captured by a touching pair of small iron cylinders (image center) magnetised by a magnet (bottom left).

distant field sources were now being locally re-radiated, drastically reducing their spatial coherence.

To reduce power supply noise pick-up and decoherence, the pre-amplifiers could be moved away from the array. The resulting noise reduction is likely to outweigh the noise increase due to the increase in unamplified signal pathway length, which can itself be reduced by using twisted pair or coaxial conductors.

### Modelling Flux Linkage Between the Sensor Cores

The work in this subsection was the result of a collaboration with my student Nick Pawlowski, who also wrote it up in his Masters Thesis, which was submitted to the University of Leeds [54].

The effective area of the induction coil is increased by the presence of the high permeability soft iron core, which becomes magnetised by the applied field and concentrates magnetic flux. When multiple cores are brought together their effective areas begin to

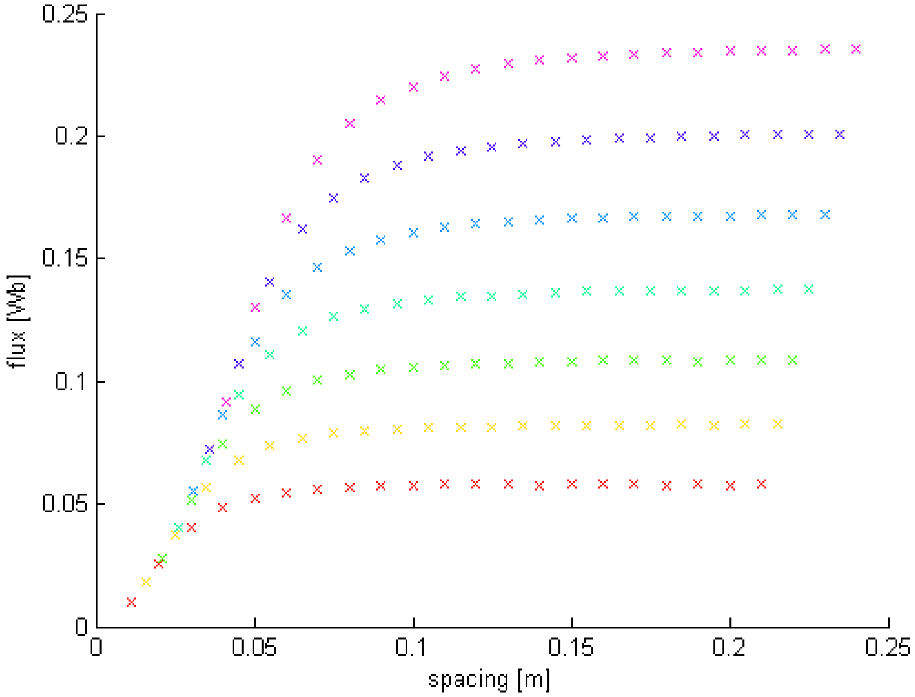

FIGURE 2.15: Dependence of the flux through the centre sensor of a hexagonal array on the radial spacing between sensors. Each colour represents a different core radius starting at $r = 5$mm up to $r = 20$mm in steps of $\Delta r = 2\,5$mm.

overlap. To determine the effect this has on coil magnetisation Comsol was used to simulate in 2D, a pair of iron cores in the presence of a dipole magnetic field of a similar relative scale to the heart field, see figure 2.14.

We created a series of experiments within the simulation. In order to determine the optimal sensor spacing we simulated a 7 coil hexagonal array; one center coil surrounded by six others. The spacing between the cores was varied in $\Delta r = 5$mm steps, whilst measuring the flux through the centre coil, figure 2.15. At close spacing the center coil flux increased approximately linearly with increasing spacing, until reaching a saturation point where it plateaued. For the $r = 20$mm core that we are using in the magnetometers, the optimal spacing occurs in the region of 8 cm to 12 cm.

We then looked at the spatial measurement error in the Mk1 array. Firstly we located a magnet underneath the centre coil and rotated it from $0^o$ to $355^o$ in $1^o$ steps. The resulting core flux measurements were used to construct magnetic field maps that were then interpolated. Due to the limited spatial resolution the reconstructed field is distorted, see figure 2.16. We measured angular errors of up to $15^o$, which is a quarter of the angular distance between sensors. The minimum errors occur when the dipole is aligned with the sensors or halfway between a pair. These particular error magnitudes are dependent on the size of the dipole field relative to the array spacing, however the lensing characteristics are universal. The dipole angle is measured from the pole peak positions. However, when an integrating method is used; such as a centroid or convolutional method 4.2.1, the error is reduced substantially.

The positional accuracy of the sensor array was tested by moving a smaller magnet

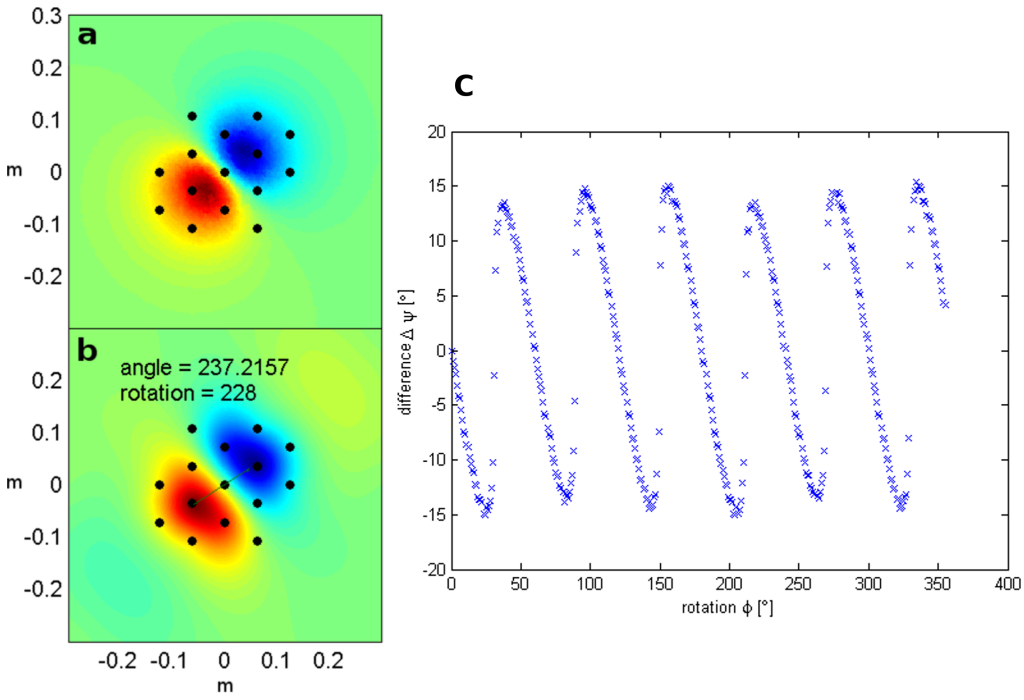

FIGURE 2.16: COMSOL physics simulation showing (a) a simulated dipole field and (b) the dipole field as measured by the sensor array (black circles) then interpolated. The reconstructed dipole has a visible distortion from elliptical and a rotation of the pole peaks; The difference between actual and observed dipole angle is $9.2^o$. (c) Shows the angle error in peak position over a $360^o$ rotation, can reach $\pm 15^o$.

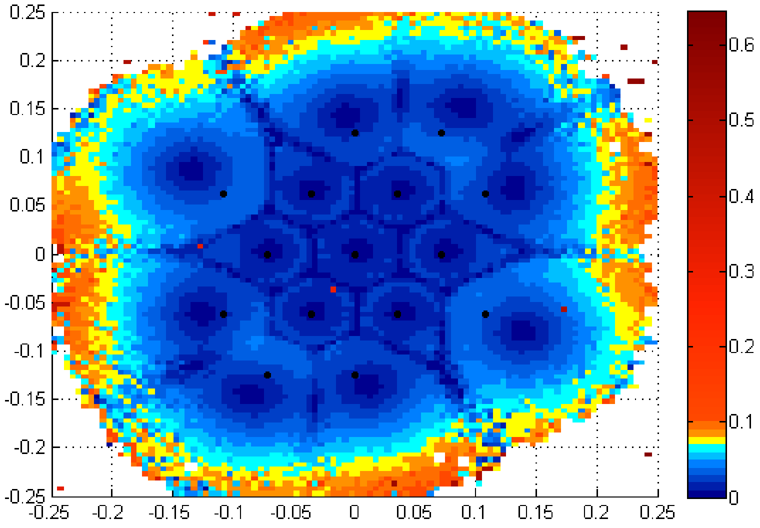

FIGURE 2.17: The error in the position (in meters) of a magnet as seen by the detector array. The error within the array is $< 6$cm, but grows quickly outside the array perimeter.

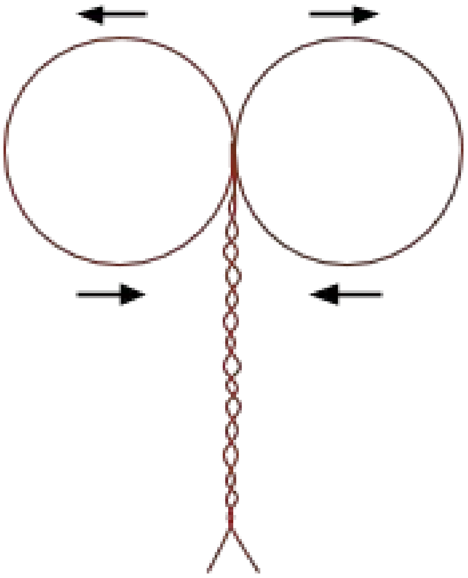

FIGURE 2.18: Butterfly coil produces an annular magnetic field. A dynamic voltage is applied using an arbitrary function generator, with a resistor placed in series with the coil to limit the current. The field is imaged by the detector array for testing purposes, to check sensor configuration is correct and to test the achievable signal to noise ratio.

around then measuring the reconstructed dipole centre. From this we obtained a map of the positional error, see figure 2.17. Within the array the error is $\leq$ 5cm. Outside the array the error grows drastically. This is due to the interpolation algorithm attempting to extrapolate. To prevent this wild extrapolation error we can employ a boundary constraint by creating a ring outside the array perimeter that is tied to the magnetic ground. In this experiment ground would be set to zero flux. But for MCG ground is set to the mean sensor value to match the normalisation.

### Procedures for Testing the Array

To test the mapping capability of MCG devices I built a 'phantom' heart. It produced a dipole magnetic field with an adjustable power level. It consisted of a figure-8 loop of wire with the ends forming a twisted pair, see figure 2.18. An arbitrary signal generator was connected in series, with a resistor to limit the current, the voltage across the coil determines the power output. A recorded ECG waveform was output from the signal generator to create a time varying magnetic field with similar frequency components to an MCG. By placing the figure-8 under the the detector array the dipole field could be imaged. This was used to perform a setup sanity check; With the phantom power set high its signal is present in all sensors. The structure of the imaged dipole was compared with the placement of the phantom, to ensure the sensors are wired correctly.

The phantom was also used to measure the fidelity of noise removal algorithms. The emitted power of the phantom was lowered into the noise. Because the field structure is known, the signal captured after DSP could be compared with the expected signal to check for distortion and extract a measure of fidelity. This is important for advanced noise removal algorithms that function by decomposing the signal, separating the noise components then recomposing the signal, since the recomposition process can create arbitrary structure. This is especially problematic with learning algorithms that look for an expected signal since they can end up synthesising a signal from noise.

### 2.3.2 Digital Signal Acquisition

The pre-amplified signals are transmitted to the ADC via coaxial/shielded cables with BNC connectors. The Mk1 device used a 16 channel ADC provided by National Instruments, the PXI4498. This device has an internal gain of 30dB, which we used to provide a dynamic range of 0.316V. Operated at 24bit resolution the smallest perceivable voltage increment was 18.8nV. When scanning the ADC synchronously sampled the 15 magnetometer voltages and the ECG voltage at regular intervals of 0.5ms (2kS/s). The sampled voltages were streamed from an FIFO buffer directly to the computers random access memory (via DMA protocol) for processing and storage.

For the Alpha device we decided to build an ADC based on the ADS1298 from Texas Instruments. This chip is similar to the 4498 with 24bit digitisation and more than sufficient sample rate. The benefit of using it is a reduction in system cost and the ability to scale the number of inputs by daisy chaining multiple chips (80+ channels is achievable).

The N sensors are synchronously recorded to produce a time series $V_t^i$:

$$V_t^i = \{V_t^i : i \in N : t \in T\} \tag{2.7}$$

Where the time index set T are the natural numbers, from zero to T (following the convention used to index NumPy arrays [55]). A 10 minute scan will contain $24_{bit} * i * T = 57.6MB$ of information.

In the Mk1 device, the digital signal acquisition components were separated from the detector array. For the alpha prototype I experimented with mounting the ADC in proximity to the array. The idea was to create a head with completely integrated electronics, which would simplify the connections to the device body; from 37 analogue connections + analogue power, to a single digital connection (perhaps cat5 cable) + analogue and digital power. This is desirable from a manufacturing and maintenance point of view. However, the noise field created by the ADC was too intense and incoherent, and completely prevented MCG acquisition due to poor SNR. The latest prototypes prioritise the separation of the electronics from the induction coil detector array.

Digital signal processing is covered in section 3.4.

#### Synchronous ECG Acquisition

One of the ADC channels is used to synchronously record an ECG from the scan subject. Unlike the noisy MCG signals, the ECG has a positive SNR; because of this the ECG is used as a fiducial to locate the heart beats in time. This allows a cardiac cycle average to be composed from the MCG signals, which is essential for extracting the signal from the noise; see section 3.4.1. The precision of the fiducial is important for cycle averaging performance. Especially for high frequency features in the depolarisation region such as HIS bundle activation in the P-R segment [56].

The ECG device used was an Accusync 42 [57]. This mains powered 3-lead ECG outputs a filtered and amplified ECG signal through a BNC cable. It also outputs a gate signal; this TTL binary signal switches on at the rise of the ECG R-wave and off 400ms later.

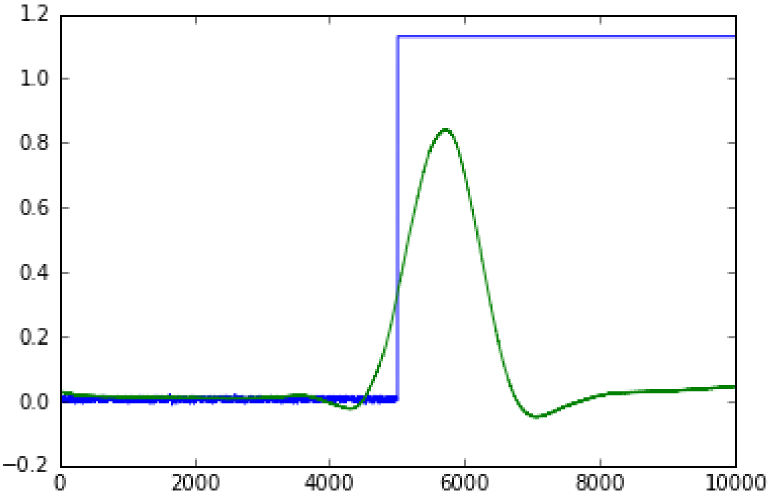

FIGURE 2.19: Synchronous recording of ECG and R Wave Trigger output from the Accusync 42 device. Showing the location of the fiducial on the rising edge of the R-wave.

The Mk1 device had only one ADC channel available for recording the ECG signal, therefore we had to choose between recording the trigger signal or the full ECG waveform. In the early experiments and for the first clinical trial, we chose to record the gate signal, since we trusted it to have a higher accuracy than our own R wave trigger extraction algorithm, which consisted of a high pass filter to remove the baseline wander followed by an amplitude threshold to detect the R-wave. Whilst simple to implement and fast to compute, this technique was prone to give a fiducial phase error because it relies on perfect baseline wander removal and a consistent R wave shape.

However, the gate signal of Accusync device is not perfect, it must also tackle noise, though because it is proprietary we do not know its fiducial phase error. We do know that it employs a real-time filter (IIR) and that such filters have an inherent phase error. In contrast, post processing filters can be made with zero phase error (FIR). Also, higher complexity fiducial extraction algorithms using wavelets can be used in post processing to gain further accuracy. In retrospect, it would have been better to record the full ECG waveform, rather than the gate signal.

Later I found an excellent library for this called Biosppy [58], which has excellent tools for ECG fiducial extraction, utilising wavelet decomposition. Another advantage of recording the ECG is the potential to classify individual heart beats, so that intermittent cardiac phenomenon can be separated from regular beats, and beats containing a large amount of myoelectric noise (from skeletal muscle) can be rejected. An ideal MCG device is necessarily a hybrid MCG-ECG.

### 2.3.3 Operator Interface

The first level of interaction with the device is through its physical interface. During initial setup the position of the medical bed and device body within the room was chosen to minimise the influence of local noise sources. Then for every scan the operator routine was as follows:

- Participant preparation: Ensure participant understands the procedure, has removed metal objects from their person and is prepared to lie still for the scan duration. Also must have completed consent forms.
- Power-up: Turn on the computer system and switch batteries from charging to operating mode.
- Software Initialisation: Load the scanning software, operator log-in, entry of scan subject details.
- Scan positioning: With the participant lying on the medical bed the detector is positioned over their heart. Then the medical bed powered off.
- ECG Acquisition: ECG electrodes are connected to the skin and acquisition of stable ECG trigger signal confirmed.
- Commence MCG acquisition.

### Preparation of the Scan Participant

Interaction with the subject was important to prepare them for the scan. They were asked to remove metal objects from their person; especially those worn on the torso such as necklaces and bras that contain underwires, since they induce fields from the motion of the chest and cause decohrence. The subjects were instructed to lie still on the bed. However, they should have an understanding that their muscle tension, movement and talking is a source of noise. Perhaps a somatic meditative primer would assist the participant into entering a relaxed state.

During the first clinical trial this was not consistently communicated to the participants, as often I observed conversation between the the research nurses and the participant during the scan. I observed that laughter in particular was a huge source of noise, because of the vibration conducted into the metal framed medical bed. I have not tested the effect on a non-ferrous bed, the magnetomyographic signals from the torso muscle contraction may also be significant.

### Power-up Routine

In the Mk1/Mk2 devices the operator had manual control of the battery system powering the amplifiers. To conduct a scan they were required to switch from charging to operating mode, where the batteries are isolated from the charging system, this was done via a keyswitch that operated a relay inside the battery box. Because of the temperature dependency of the pre-amplifier gain, the switch-over was made several minutes before running a scan, this provided time for the amplifiers to warm up to equilibrium temperature. This resistive heating occurs due to the current passing through the linear regulator ICs and the amplifier IC, see figure A.3 for an IR spectrum image of the PCB at equilibrium temperature, which is approximately 15 °C above ambient. At the end of a scanning session it was essential for the device to be switched back to charging mode to prevent damage to the batteries from deep discharge. The charging process took approximately 8 hours from empty. The management of this charging process was a vector for human error, during the clinical trial there were multiple occasions when the device was left in operating mode overnight, causing the batteries to discharge and delaying acquisition because of the need to recharge.

In the Alpha device, the isolation relay was controlled via firmware. The operation mode was only engaged when the device was switched to scanning mode. A higher capacity battery bank was also employed that allowed longer scanning sessions. The Alpha device has a greater risk of topple over because the detector array was mounted on an arm connected to the main body of the device. To counter this risk, the extra mass of the batteries was located in the body base to minimise the centre of gravity thereby reducing the topple over angle to within regulatory limits.

#### Positioning the Detector Array

In the MK1 and MK2 devices the position of the detector array was controlled in the horizontal plane using the castor wheels rolling on the floor. The center of the array was aligned to 4th intercostal space, to the subjects left of the sternum, as this was found to be the average center of the MCG field. The distance between the chest and the array was controlled using the electric actuator built into the medical bed. Observation of several breathing cycles was required to ensure no contact during maximal inhalation. In the Alpha device the array was supported by an arm from the main device enclosure.

### 2.3.4 Acquisition User Interface

The device contained a monitor screen with a keyboard and mouse for interaction. The interface was password protected to ensure access only to registered operators. In addition to controlling the MCG acquisition process, the interface also allowed subject meta data to be entered, which was then stored in association with the MCG data. This metadata included:

- Operator identification.
- Pseudonymous subject registration.
- Age
- Gender
- Timestamp
- Miscellaneous notes

This meta information was stored in a text file called scan_info_dict, in the same folder as the MCG data. To anonymize each scan, the pseudonym was not directly stored, instead the stored subject I.D. was a alpha numeric string generated from the hash of the patient pseudonym plus a pseudo-random salt with the trial key, this gave the dataset complete anonymity for anyone without the trial key. The pseudonym cannot be directly reconstructed from the hash. In order to recreate the pseudonym from the hash during analysis the software used the key and a list of possible pseudonyms to create hashes, then matches the hashes with the stored subject I.D. to identify the scans. One issue with this system was that the pseudonyms were entered by the operator without a constraint mechanism, leading to a lack of consistency. For example, healthy volunteer number one could be specified as: HV01, HV1, Hv01, hv01, hv1... Which to a human provide equivalent information, but each hash to completely different values.

During the scan the operator was provided with a visualisation of the recorded information:

- Clock recording scan duration.
- Heart rate information.
- Real time stream of the ECG and MCG signals.
- Running total MCG cycle average.
- Magnetic field maps at particular points in the cardiac cycle.

The real time ECG signal was used to check for a consistent ECG acquisition. The streaming MCG signals were used to check noise level and to observe transient noise events. The running total MCG cycle averaging could be used to determine when sufficient SNR had been achieved.

### 2.3.5 Analytical Workstation

To visualise and experiment with diagnostic algorithms on the clinical trial dataset I wrote an analysis workstation software. The purpose of the user interface was:

- To facilitate exploration of the dataset.
- To error check the automatic feature labelling allocation.
- To provide human input to the diagnostic process: Marking positions in the MCG, flag position error, flag low SNR, Reject based on noise or other error.

Because the device was used primarily by us for research, the main requirement of the analytics software was the ability to apply experimental diagnostic algorithms. Because the dataset is quite large ($150MB \times 480 = 72GB$), and the algorithms are somewhat computationally intensive, the time to compute results can be several hours. The results should be computed as fast as possible because a short feedback time increases the number of design iterations possible. In order to achieve this I precomputed the interpolated MFM time series, which took an extra 60GB of storage, but sped up certain algorithms.

Most of the algorithmic experimentation took place within Jupyter notebooks, which are cellular python environments that share a common kernel and memory space but are divided into cells that can be run separately. Graphs can be plotted within the notebook and even be made interactive with widgets. However, a standalone program written in QT and Python was more appropriate for cycling through the clinical trial. This program was later translated into C++. As C++ is a compiled language it performs faster than Python, which is mostly interpreted (though elements of compilation exist in e.g. the numpy libraries).

## 2.4 Conclusion

We successfully created a MCG using induction coil magnetometers. The ICM performance was validated in a magnetically shielded room and found to be sufficient for cycle averaged MCG, with $104\mathrm{fT}/\sqrt{\mathrm{Hz}}$ at 10Hz of inherent noise. This unprecedented sensitivity is largely due to the low noise amplifier and ADC electronics, which were unavailable during previous attempts to create ICM for MCG.

We built a complete MCG system based around a detector array of 15 ICM, that was suitable for use in a clinical research environment. The device had a separate enclosure for the electronics excluding the pre-amplifiers that were located on top of the coils in the detector array. The array was supported on an aluminium frame, which could be positioned above a scan subject. This separation reduced the noise imposed on the array by the switching and power supply components. However, the placement of the pre-amps in close proximity was a significant source of noise due to the peak power requirements of 100mA per coil. We did not notice this noise source until the next prototype device, which had a regular arrangement of analogue power wiring; since the wiring structure imposed a detectable structure on the measured noise field.

The permeable cores of the ICM increased the coils effective area, increasing their reponse by $\approx 3\times$. However, when the ICM are arranged in close proximity, as in the detector array, their effective areas overlapped causing a distortion in the perceived magnetic field map. This leads to an error in the measured dipole angle.

This device was used to measure 480 scans in an unshielded clinical trial at LGI. The results were quite noisy, but definitely contained MCG signals. The encountered environmental noise and diagnostic performance are the subject of the following chapters.

## Chapter 3

# Environmental Noise & Digital Signal Processing

The contemporary built environment contains a multitude of sources emitting electromagnetic energy. This environmental noise dominates the biomagnetic fields generated by our organs. In order to map biomagnetic fields we must be able to isolate them from the onslaught of environmental noise. This is the critical challenge for clinical adoption of biomagnetic field mapping.

This chapter begins with an overview of environmental noise sources, followed by an analysis of noise using frequency decomposition, Allan variance and spatial coherence techniques. Then an exploration of noise rejection methods and their performance, both physical sheilding & digital signal processing.

The most profound finding is that the separability of the background environmental noise from the local incoherent cardiac signals is proportional to their relative degree of spatial coherence. This was exploited to create a coherent noise rejection algorithm that provides lower noise than a gradiometer array and is responsible for the devices ability to operate without a magnetically shielded room.

## 3.1 Introduction

Man-made electromagnetic noise is now present almost everywhere on the Earth's surface. In urban areas it dominates natural noise (of atmospheric and cosmic origin) [59]. It is common to see noise levels reach nanotesla (nT) levels, which is on the order of a $10^5$ times larger than the cardiac biomagnetic field. Whilst we could retreat to rural areas to perform our measurements within a quieter environment, this would prevent access to the technology for the majority of the human population who now live within cities. Therefore we must be able to operate in an urban environment.

The effect of noise on measurement performance is dependent on the noise structure and the measurement device. Primarily, if the noise amplitude is too high (outside the dynamic range of the detector system), then information will be lost. If the noise is much greater than the signal amplitude then we will not be able to perceive the signal since they are superimposed in the amplitude domain. However, we can employ techniques in the frequency and spatial domains to separate them, and use signal averaging to reduce the random noise. Coloured or periodic noise that is in rhythm with the heart beat will be constructively interfere in the signal average, so it is advantageous to have a non-periodic heart rate!

One guaranteed method for isolating environmental noise is to use a magnetically shielded room (MSR). These rooms act as spatial filters with shielding factors of 40dB-90dB. Offering $\mathrm{fT}/\sqrt{\mathrm{Hz}}$ noise levels within, suitable even for real time observation of MEG. See section 3.3.1 for a review of shielding types. However, MSR are expensive and throughput limited; the ability to perform MCG in an unshielded environment would drastically decrease the cost/benefit ratio, giving more people access to MCG. An alternative spatial filter is required.

The spatial coherence of environmental noise is usually higher than that of the cardiac signals, since the majority of noise sources are further from the detector array than the heart. We have developed an alternative coherent noise rejection algorithm based on inverse ghost imaging; its noise rejection performance is lower than an MSR, but we found it was sufficient for observation of depolarisation phase signals in a clinical environment, and there is potential for improved performance by increasing environmental spatial coherence and reducing device decoherence.

Typically axial gradiometers are used to subtract the spatially coherent components, this requires precise matching of the sensor gain in each pair and reduces the sensitivity by a factor of 2 due to superposition of each sensors inherent noise into the gradiometer signal. In contrast our method uses a bucket signal subtraction that reduces the sensitivity by only $(1 + \frac{1}{\sqrt{n}})$, where n is the number of sensors in the bucket signal.

### 3.1.1 Electromagnetic Noise Sources

There are many sources of ambient electromagnetic noise, arising from communications, computation, switch-mode power regulation, air conditioning, transport etc. Together these cover the most of the non-ionising radiation band [59]. The most pervasive source is the utility grid, where power is transmitted using an alternating current at a frequency of 50Hz, which induces a corresponding oscillating magnetic field at 50Hz (60Hz in the Americas and parts of Asia). The conductors of the utility grid run through almost every room in every building and power the majority of devices, because of this it is difficult to achieve more than a few meters of separation from the measurement region. The phase of the grid is necessarily maintained to a high precision, by synchronisation with master atomic clocks. This results in a highly temporally coherent 50/60Hz signal, with a white intensity distribution due to the thermal nature of the electronic conduction. The frequency distribution of the cardiac field ranges from a few hertz to about 150Hz. This overlap means that the utility grid cannot be filtered out using traditional filtering techniques without some loss of physiological information; whilst acceptable for many diagnostics, this can make it difficult to observe high frequency fibrillation components.

Electrically powered transport systems such as subways induce large ELF signals due to the large currents they employ for power distribution [60]. These sources typically occur with a Poissonian distribution, because of this they can be observed with Allan variance analysis as having an effect on a discrete range of averaging times. This effect can be seen in the observation of lift shaft noise at Leeds General Infirmary (LGI) in figure 3.6.

The magnetocardiography device is itself a source of noise, since it contains switching components for power supply voltage regulation, computation and the display interface. Any switch-mode device will create broadband interference due to the square nature of the pulses emitted. This can be minimised by using battery power and linear regulators. However, the digital nature of the signal processing components is unavoidable, therefore shielding must be employed and the digital components separated from the measurement region as far as practicable.

Hospital beds typically have an electronic height adjustment powered by AC mains. Even when the motor is not activated parts of the control electronics are still live, with mains current flowing through. This causes an emission of mains 50Hz hum, with associated colored noise, that is a large noise source due to its proximity to the detector array, see figure 3.1 for an Alan variance analysis of this noise. If the bed is isolated by switching it off at the wall then these signals disappear.

### 3.1.2 Acoustic Noise

The Earth's magnetic field is pervasive throughout the Earth's surface where it varies in strength between $25 - 65\,\mu\text{T}$. It extends to a minimum of 60,000km from the Earth's surface, where it is squeezed by the solar wind. Certain types of magnetometers such as SERFs cannot operate with such a large applied field and require active cancellation. Induction coil magnetometers (ICM) measure the time derivative of the magnetic field, not its absolute magnitude, and therefore are not directly affected. However, ferromagnetic objects within the Earth's field become magnetised and due to their hysteresis will retain that magnetisation for a short time after the applied magnetic field changes. Their movement therefore creates magnetic perturbations of comparable flux density to the Earth's field. This hysteresis perturbation noise was also visible through Allan variance analysis (3.2.1) to be occurring due to the passing of lifts through the lift shaft at LGI. Other researchers have observed hysteresis perturbations from passing vehicles several 10's of meters away. I found that even small movements could be detected by our devices, for example the movement of a swivel chair across the room, picking up a screwdriver and even the movement of the heart perturbing the medical bed frame.

The perturbation induced signal is dependent on the vibration mode; Rotational modes will induce varying signals dependent on the distance of the sensor from the axis of rotation and the sign of the rotation vector. Also important are the relative angle of the incident magnetic field with the detector array and the resulting flux linkage between sensors in the array. If there were no linkage then translational vibration modes would induce only common mode signals, but this is unlikely to occur in practice, since there is a high degree of flux linkage see section 2.3 for a magnetic flux linkage simulation of the array. To counter this noise the detector array must be mounted stiffly, and isolated from any sources of vibration such as the patient and any vibrations that may travel through the floor.

### 3.1.3 Ballistocardiographic noise

The mechanical action of the heart causes the ejection of blood from the ventricles. The recoil from this action is measured in a ballistocardiogram (BCG); The first ballistocardiography instruments consisted of a suspended table that the subject would lie upon in supine position and the table deflection recorded [61, 62]. The medical bed used in the clinical trial had a frame that was supported by a central column, the stiffness of this joint was sufficiently low that movement of the subject would cause oscillations about this joint. It was possible to visually perceive the recoil of ones heartbeat whilst lying on the bed, by parallax; using a fixed distant visual reference to compare the bed position. Since the bed frame was composed of steel it was magnetised by the Earth's field

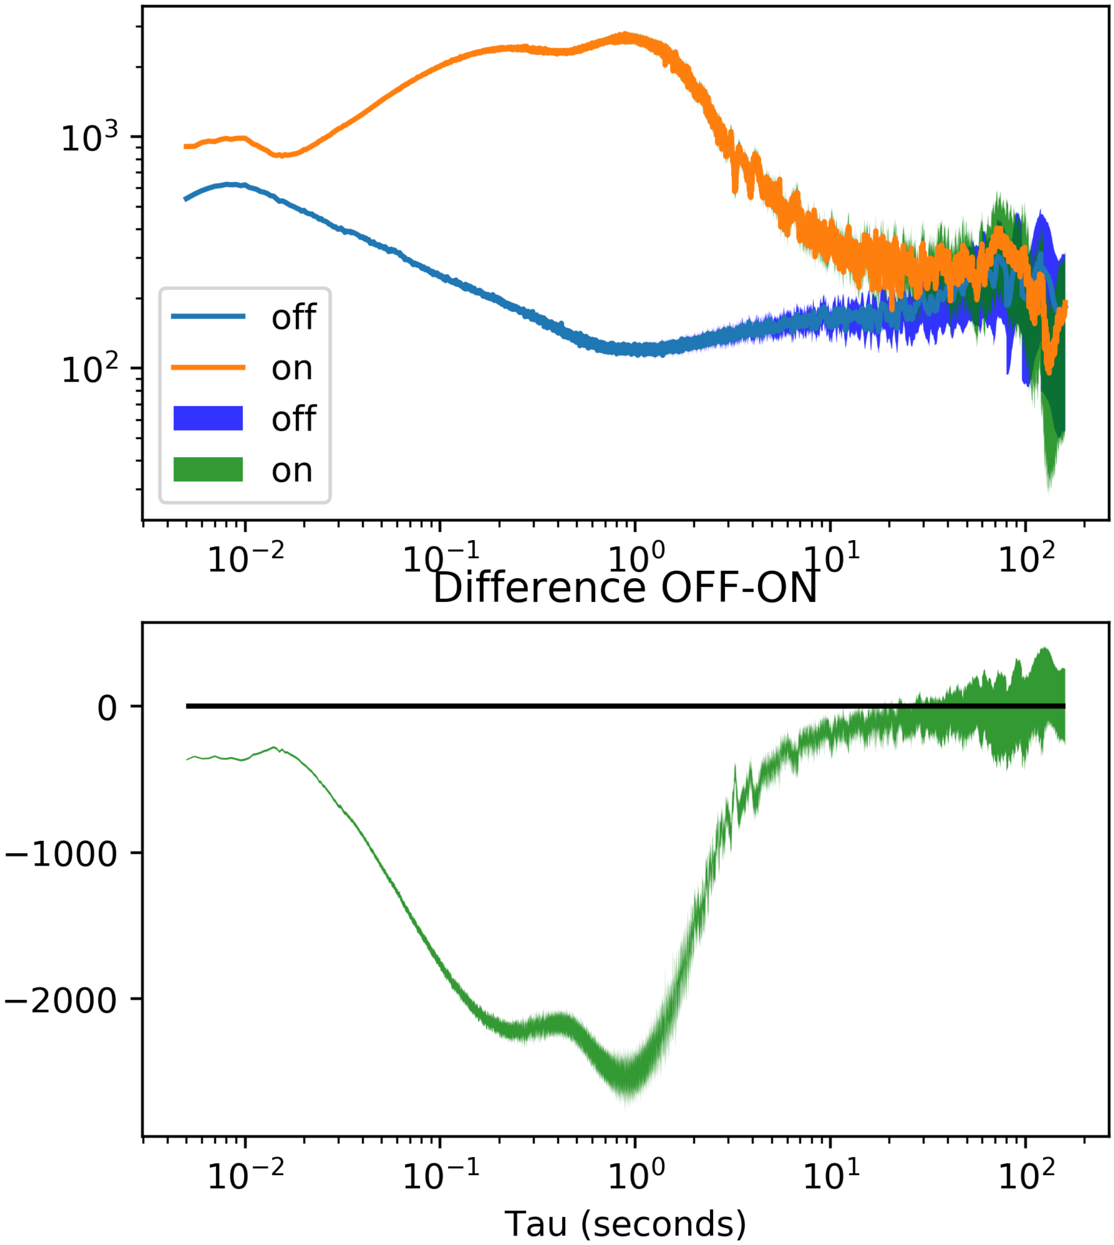

FIGURE 3.1: Allan deviations and overlapping Allan deviations (pale blue and orange) from magnetometer time series recorded by the array with the medical bed in ON and OFF states. The lower plot shows the difference between the ON and OFF Allan Deviations. The ON state introduced significant coloured noise with time periods up to approximately 1 sec.

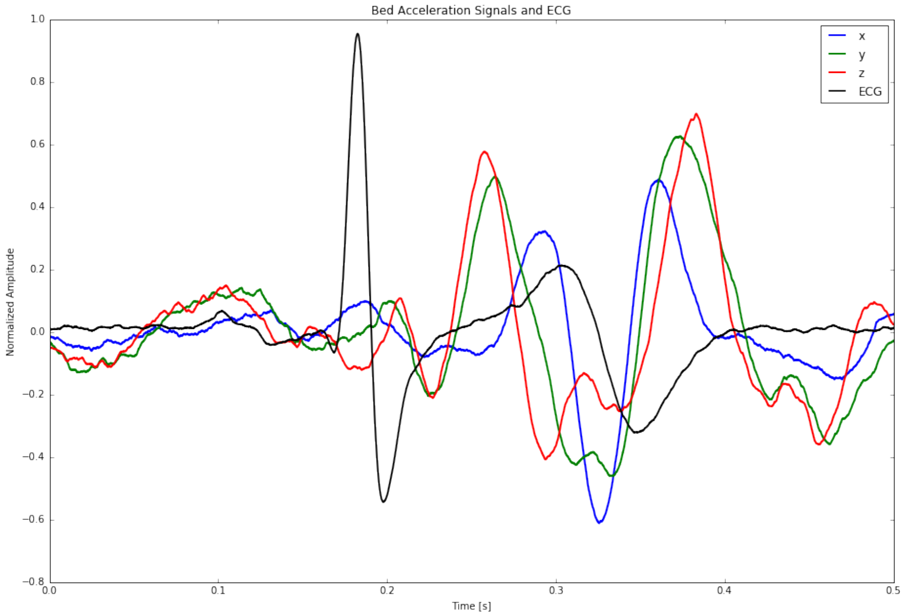

FIGURE 3.2: The cycle averaged 3-axis accelerometer signals (X, Y, Z) of myself, shown with the synchronously recorded ECG signal.

so the heart driven oscillations induced a hysteresis perturbation, a ballistocardiographic signal.

This ballistocardiographic signal went unidentified by us for many months; since it appeared in the repolarisation region of the cardiac cycle with a magnitude that was similar to the magnetocardiographic depolarisation signals, we therefore assumed it was a magnetocardiographic signal. Our suspicions were raised when we performed MCG scans on a wooden bed and the repolarisation signals disappeared into the noise. Since the repolarisation MCG signals are lower frequency than the depolarisation, hence lower amplitude due to the linear frequency response of ICM.

To prove the ballistocardiographic origin of the repolarisation phase signals I conducted an experiment to measure the acceleration of the bed using a 3-axis MEMS accelerometer (ADXL337Z). The accelerometer was affixed to the bed frame, directly beneath the measurement region. To minimise noise the accelerometer was powered by a battery via a twisted pair of wires and a linear voltage regulator (LM317) and the output signals were transmitted to a DSO using a coaxial cable. The signal to noise ratio was too low for real time observation of ballistocardiographic signals. Therefore the ECG signal was used as a fiducial for cycle averaging. I performed a scans of myself and a colleague and obtained synchronous ECG, MCG and BCG recordings, see figure 3.2. These signals are almost identical to published acceleration BCG signals [62]. The structure shows similar absolute amplitude dynamics in all axes, therefore I computed the principal acceleration component (PAC), using principal component analysis. The PAC can be compared with the MCG signals, visually and by computing the covariance.

The PCA of myself and my colleague display a very similar structure, particularly in the large amplitude oscillations observed in the repolarisation phase, see figure 3.3. A similar but inverse wave exists in the region preceding QRS, this could either be atrial contraction or a ringing of the bed that has continued from the previous heart beat. Since the heart rate is approximately regular this periodic noise will to some extent be captured

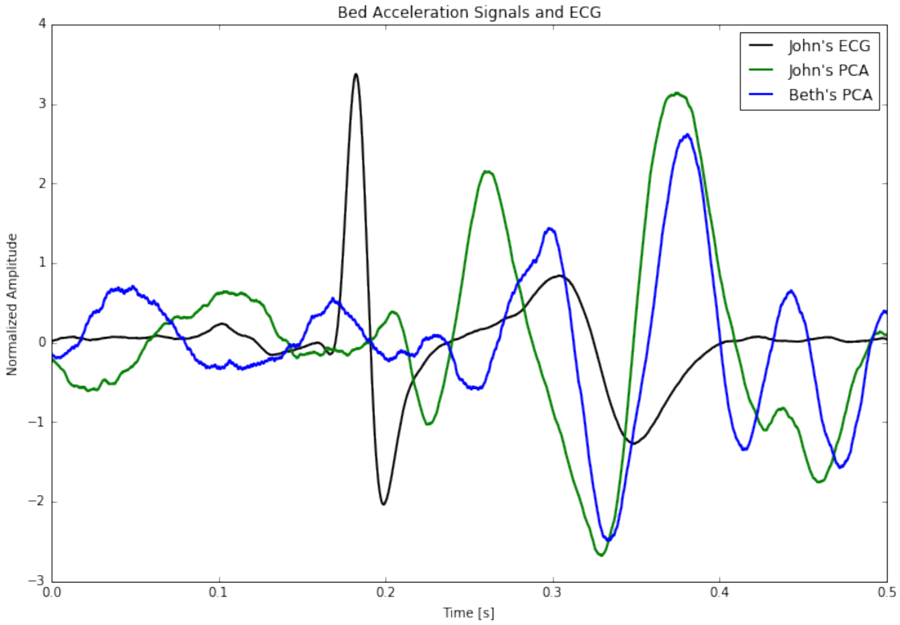

FIGURE 3.3: The principal components of acceleration of myself and lab mate Beth, show comparable structure.

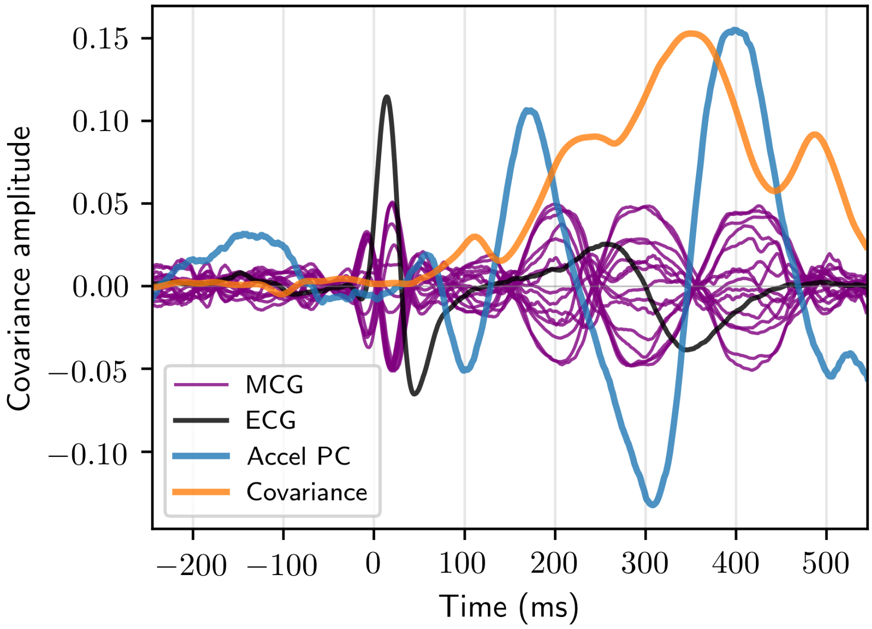

FIGURE 3.4: The principal component of acceleration is shown along with the corresponding synchronously recorded ECG and MCG signals. The mean covariance between the acceleration PC and the MCG signals is shown. A high covariance is seen over the repolarisation phase of the cardiac cycle.

in the cycle average. This could be confirmed by acquiring a real-time signal, perhaps using a laser balance type system, which would measure displacement directly, to a higher sensitivity, this is certainly feasible given the efficacy of the direct parallax observation obtained by the subject.

However, to prove the ballistocardiographic nature of the MCG signals the cycle averaged method is sufficient. The mean covariance of the PAC with each of the MCG detector signals was computed, see figure 3.4. A strong covariance is observed in the repolarisation phase but not the depolarisation phase, therefore our PQRS region measurements are MCG and the T region measurements are BCG. To detect pure MCG signals, the subject must be supported on a non-ferromagnetic bed. However, this is not necessarily ideal, I later found evidence that the BCG signals contain diagnostic information, see section 4.3.3. Therefore it could be clinically useful to measure these bed associated BCG signals with the MCG; particularly with ICM sensors, which have reduced sensitivity to the repolarisation phase signals, forming a hybrid MCG-BCG device.

## 3.2 Noise Analysis

Through noise analysis we can learn more about the composition of the noise environment, this information can be used to inform noise rejection strategies. Noise sources can be classified and tracked down giving the possibility to eliminate them or their effect, improving the signal to noise ratio.

The useful information that can be extracted from a measurement is proportional to the ratio of signal to noise power, which is equivalent to the ratio of variances or the ratio of amplitudes:

$$SNR = \frac{P_{signal}}{P_{noise}} = \frac{\sigma^2_{signal}}{\sigma^2_{noise}} = \frac{A_{signal}}{A_{noise}} \tag{3.1}$$

The Shannon-Hartley theorem explicitly defines the maximum information rate (mutual information, see eq:1.2) obtained through a measurement channel, subject to gaussian noise, as the channel capacity C (in bits per second):

$$C = Blog_2(1 + SNR) \tag{3.2}$$

Where $B$ is the bandwidth of the signal (Hz). For more complex coloured noise the theory still applies following frequency decomposition, treating each frequency component as a separate gaussian noise.

If the $SNR \neq 0$ then $C > 0$, and the total mutual information, $I$ between the cardiac magnetic field and the digitised magnetometer time series, $I = C * t$ can be increased by increasing the averaging time. In the rest of this chapter the mutual information is between the cardiac magnetic field and the digitised MCG signals at various stages of signal processing. We will discuss methods for maximising $C$, to capture maximum $I$ with minimal scan time.

In the case of cycle averaged MCG in an unshielded environment the important measure of scan time is the number of cardiac cycles acquired. The time available is dependent on how long the subject is able to be scanned. We found that clinicians and patients prefer a maximum duration of about 10 minutes. I have personally had scans of around 1 hour long that were quite uncomfortable because I was lying still with relaxed muscles, reducing blood flow and eventually leading to numb limbs. It is possible to concatenate multiple scan periods together, therefore a longer total scan time can be achieved by having breaks. This introduces a re-positioning error which in the current system would reduce the SNR, but if the position was known precisely then multiple measurements could be used to increase the spatial resolution. Though the total information captured increases in proportion to the number of acquired heart beats, $n$. The final diagnostic information captured after cycle averaging is proportional to $\sqrt{n}$. Therefore a 20 minute scan ($n \approx 1200$) captures only 30% more information than a 10 minute scan ($n \approx 600$).

In general, the SNR is increased by amplifying only the signal and reducing the noise. There are limits to the level of signal amplification since the inherent input noise of the amplifer is amplified, along with the signal, in proportion with gain. Inherent device noise is explored in chapter 2.

In general, the more information obtained about the noise, the more it can be separated from the signal, by reduction of $H(X|Y)$, see equation 1.2. This can be achieved by an algorithmic process, a change to the device hardware or a change of the environment. For example, if there was a person swinging a magnet around their head next to the device, we could observe them visually, then based on that knowledge we could ask them to stop. In order to increase our knowledge of the noise we need tools to analyse the noise structure. It can be analysed by its temporal, frequency and spatial properties.

A direct analysis of the time domain is useful for observing changes in the noise power/intensity. Transient and periodic (pulsatile) signals can be observed directly. The MCG device user interface included a real-time scrolling plot of the raw magnetometer time series. This was useful for visually correlating events with the observed noise. We learned a great deal about the vulnerability of the device to human origin noise; particularly the movement of ferromagnetic objects, quartz watches and subject movement (especially when they touch the array). However, the frequency range of perceivable events is limited by human perception and the scroll rate. Post analysis of the whole time series was also useful. For example, the scanning room we used for the clinical trial at LGI was next to a lift shaft. The passing of the lifts was too slow to perceive in real time, but was clearly visible when viewing the MCG timeseries over a longer interval, as seen figure 3.5.

### 3.2.1 Allan Variance

Allan variance (AV) was originally devised to measure the frequency stability of clocks [63, 64]. But it has proven to be useful as a general tool for the study of unstable noise sources, indicating the performance of signal averaging over long time periods. To compute its Allan variance a time series is split into a set of chunks with width $\tau$, the difference between the mean values of adjacent chunks is calculated, squared and averaged. To map the Allan variance the averaging time $\tau$ is varied from a single sample width up to a maximum of half the length of the timeseries.

$$\sigma_y^2(\tau) = \frac{1}{2}\langle(\bar{y}_{n+1} - \bar{y}_n)^2\rangle \tag{3.3}$$

Allan Deviation is the square root of the Allan variance. $\sigma_y(\tau) = \sqrt{\sigma_y^2(\tau)}$

The slope of the Allan variance curve can be used to identify the noise colour (or power spectrum). In figure 3.6 as the averaging time period is increased, the Allan variance decreases linearly with a slope that is typical of white noise. This linear trend reaches a knee that indicates the transition from stable white noise to unstable noise. The Lab

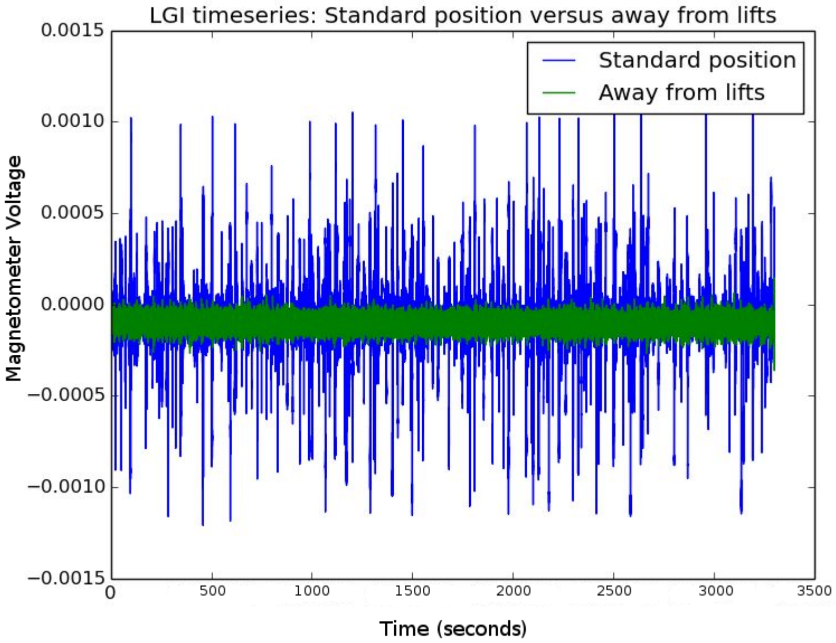

FIGURE 3.5: Time series of two 1 hour scans taken at LGI. Semi-periodic transient noise spikes are visible, these are due to the passing of lifts in the neighbouring lift shafts.

background Allan variance has the highest noise level but because the noise is stable for longer the knee occurs later than the LGI scans. This indicates that for averaging times greater than 100 seconds, the lab will have lower noise after averaging. The presence of a 2nd knee on the LGI - Standard position plot is due to the action of the lift shafts, which occurred with a Poissonian distribution approximately every few seconds. I discovered this knee during the clinical trial at LGI, which prompted us to change the position of the device to the opposite side of the room (where the "Away from lifts" scan was taken).

The presence of coloured power spectral noise is clear in the plot comparing the Allan variance with the medical bed isolated and powered; figure 3.1. The plot shows the overlapping Allan deviation, which does not contain the artefacts from the interaction of the 50Hz signal with the Allan variance windows. With the bed connected to the mains the averaging time increases until an averaging time of nearly 1 second. This is due to the 50Hz harmonics and other power-law noise present in the utility grid and amplified by its proximity to the detector array.

In figure 3.13: The unshielded signal has a higher Allan variance, that is consistently about two orders of magnitude higher than the shielded signals throughout the frequency range. From 20ms (timeperiod of 50Hz waves) to 1 second, fringes appear where the AV window width is a multiple of the 50Hz signal timeperiod. These fringes disappear when 50Hz has been notched out, because the mean of a single cycle of a sinusoid is zero. These fringes are not present in the shielded plot since the shielding is effective at attenuating utility grid noise. The shielded signal has a much lower Allan variance, approximately two orders of magnitude lower than the unshielded signals. The notched unshielded AV is actually lower than the un-notched shielded at the highest frequencies (above 1kHz approximately). This is because overall the notch is a low pass filter, whose response can be seen in section 3.4.2.

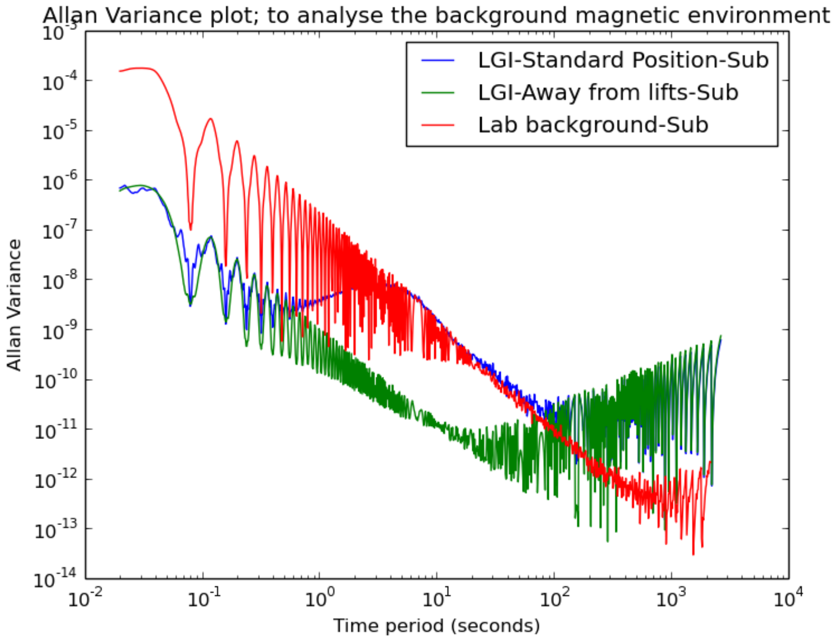

FIGURE 3.6: Comparison of environmental noise at the University lab and LGI, and comparing measurements within the room at LGI; close to the neighbouring lift shafts and in the center of the room (approx 3m away).

### 3.2.2 Frequency Analysis

A time series can be decomposed into its constituent frequencies using the Fast Fourier Transform (FFT) algorithm. The resulting frequency domain representation is composed of sinusoids with particular frequencies, amplitudes and phases. Noise sources can be identified in the frequency domain, particularly continuous wave narrow band oscillations such as the mains electricity grid. I used a FFT to determine the sensitivity of the large coil magnetometers by measuring the power spectral density of the inherent sensor noise, see figure 2.8.

Time-frequency analysis can be used to observe dynamic signals; the time series is sliced into a set of overlapping windows and a FFT computed from each window. The result is a waterfall plot or spectrogram. The frequency range and sensitivity to transients are dependent on the window width. I have taken this principle of moving window analysis and adapted it to observe the time-frequency dynamics of the spatial coherence in figure 3.11, and to perform diagnostic mapping in section 4.3.3.

#### Wavelet Decomposition

A signal can be decomposed into arbitrary wavelet components, instead of periodic sinusoidal components. To compute this wavelet transform an array of wavelets are convolved with the signal; the array is composed of frequency shifted and relative phase shifted variations of a fundamental wavelet. This fundamental wavelet is chosen to match particular features in the signal.

I used a pair of wavelet transforms to find peak features in an MCG, see section 4.2.2 for the details.

### 3.2.3 Spatial Coherence

Spatial coherence describes the statistical similarity of a field at two points in space. The concept of coherence was originally conceived from observations of direct wave interference (Young's double-slit experiment). The degree of spatial coherence of two waves is quantified by the cross-correlation function.

A 2D example can give an intuitive understanding of spatial coherence; visualise the surface waves produced by a stone dropped into a lake, the wave fronts close to the buoy have a high curvature, as the waves spread out their curvature decreases. The movement of an array of floats placed close to the source will be uncorrelated (low degree of coherence/incoherent). As the array is moved further away the movement of the floats gradually becomes more correlated (high degree of coherence), see figure 3.16.

The inverse relationship between coherence and distance can be exploited to separate the distance sources of background noise from local cardiac signals, using a technique akin to inverse ghost imaging. By subtracting the signals that are common-mode to all detectors we can remove environmental noise and preserve the cardiac signals, see section 3.4.3 for more detail on the coherent noise rejection (CNR). First, it is useful to measure the spatial coherence.

In 1956 Hanbury Brown and Twiss discovered that the degree of coherence could be quantified by correlating measurements of intensity fluctuations see equation 3.5; without requiring full knowledge of the wave frequency, amplitude and phase [65]. They used this to measure the spatial coherence of starlight from which they calculated the angular size of the originating star, see figure 3.7. It would have been impractical to use a direct photonic interferometer, due to the large variable baseline (10 to 188m) between their detectors at the Narrabri Observatory [66].

For a pair of detectors with outputs:

$$J_s = J_0 + \Delta J_s \tag{3.4}$$

Where $J_0$ is the (potentially uncorrelated) signal offset and $\Delta J_s$ is background intensity noise. Spatial correlations in the noise can be characterized via the normalized correlation coefficient:

$$g^{(2)}(r_1, r_2, t) = \frac{\langle \Delta J_1(r_1, t)\Delta J_2(r_2, t)\rangle}{(\langle[\Delta J_1(r_1, t)]^2\rangle)(\langle[\Delta J_2(r_2, t)]^2\rangle)} \tag{3.5}$$

Where the subscripts refer to detectors 1 and 2 separated by a distance $d = |r_2 - r_1|$.

In observation of both radio wave emissions from stars and the ambient noise detected during an MCG scan, fluctuations in the detected signals arise because the incident field on the detector is produced by a (relatively) statistically stationary chaotic/thermal source field. Fluctuations are very rapid compared with the detector response times but as the effective photon number n of the field will be very large for extremely low frequency (ELF) electromagnetic waves, an integration over the detector response time produces a $\frac{1}{\sqrt{n}}$ variance, which contributes to the background noise. However, for a field with a high degree of spatial coherence, the random effects are correlated by thermal field photon bunching.

Wang et al [67] considered an arrangement similar to Figure 3.7 with a spherical near-field thermal source and calculated the observed correlation as a function of detector

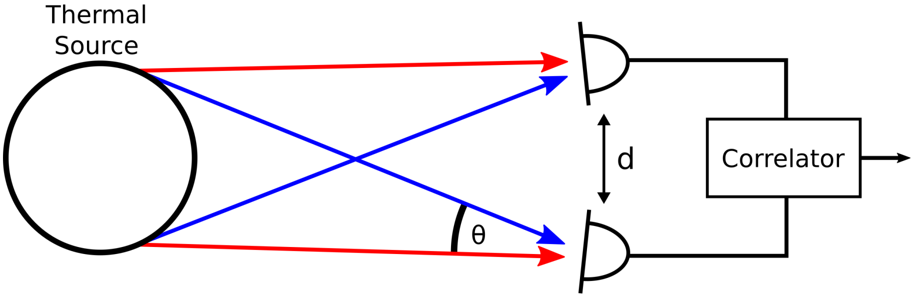

FIGURE 3.7: A Hanbury Brown and Twiss Intensity Interferometer records correlations in radiowave emissions from stars (a far-field, thermal source) between two spatially separated antennas. The mutual spatial coherence of emissions at the antennta locations is proportional to the time averaged correlator output. To determine the angular size of the source, the correlation is measured as a function of antenna separation, $d$. The width of the resulting Gaussian distribution is proportional to the stars angular size, $\theta$.

separation $d$;

$$C(d) = exp\left[-\frac{d^2}{\sigma_s^2 + (\frac{X_s^2}{4\pi w_s^2})(\frac{\sigma_s^2}{4w_s^2}+1)}\right] \tag{3.6}$$

Where $X_s$ is the distance to the source, $w_s$ is the spatial cross-sectional width of the source, and $\sigma_s$ is the spatial coherence length of the source. For chaotic light, the condition $X_s, w_s >> \sigma_s$ is always satisfied, in which case we have:

$$C(d) = exp\left[-\frac{d^2}{\sigma_s^2 + \frac{X_s^2}{4\pi w_s^2}}\right] \tag{3.7}$$

This is a gaussian distribution with a width proportional to the sources width/distance ratio, the degree of spatial coherence. In MCG we detect a superposition of cardiac source and background sources. The measured correlation $g^{(2)}(d)$, is an incoherent combination of these sources given by the equation:

$$C_T(d) = \alpha\left[\sum_i C_B^i(d)\right] + \beta C_L(d)\ ; \qquad \alpha + \beta = 1 \tag{3.8}$$

Where the subscripts T, B and L refer to the total, background and local-source (cardiac) correlations, respectively.

To measure the spatial coherence we can take multiple measurements of $g^{(2)}(d)$ at an array of distances to map out the curve $C_T(d)$; The degree of coherence is quantified by the width, $W$ of gaussian fitted to this curve:

$$F(d) = e^{\frac{-d^2}{W}} \tag{3.9}$$

The distance was specified in millimeter units, therefore the coherence width is in units of mm$^2$. I measured the environmental spatial coherence at LGI with the MCG detector array and found it to be highly coherent, with the signals more than 97% correlated across the array; a mean coherence width of $2.8 \pm 0.9 \times 10^6$mm$^2$. The resulting $C_T(d)$ is shown in comparison with longer baseline measurements in figure 3.10.

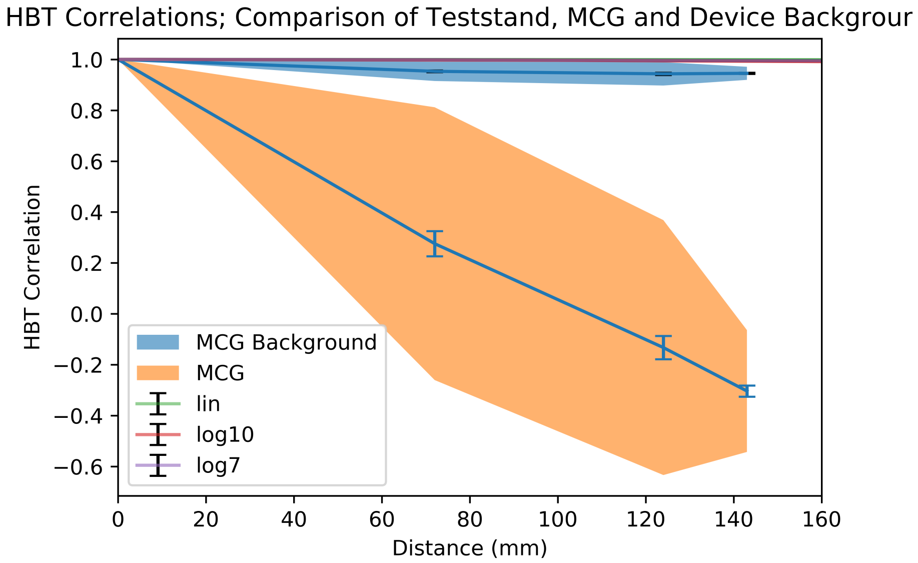

FIGURE 3.8: Comparison of HBT correlations. The raw background signals were all highly coherent. The MCG scans were slightly above the coherence detection limit of the array, and therefore become anti-correlated within a single sensor spacing.

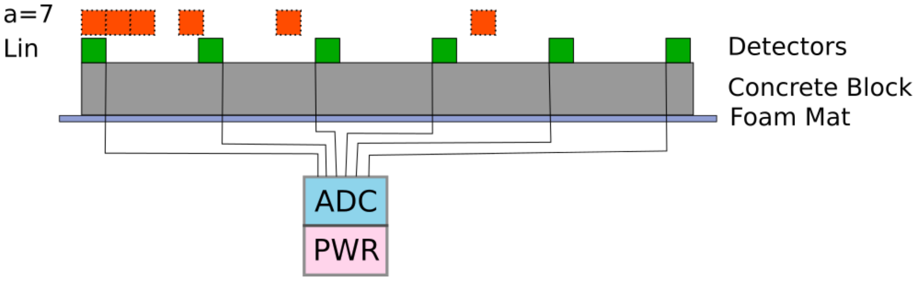

FIGURE 3.9: Schematic of experimental setup for long baseline measurement of spatial coherence. A hydraulically compressed concrete block rested upon a closed cell foam mat for vibration isolation, forming the test stand. The sensors were spaced in one of two configurations as shown (Lin, a=7).

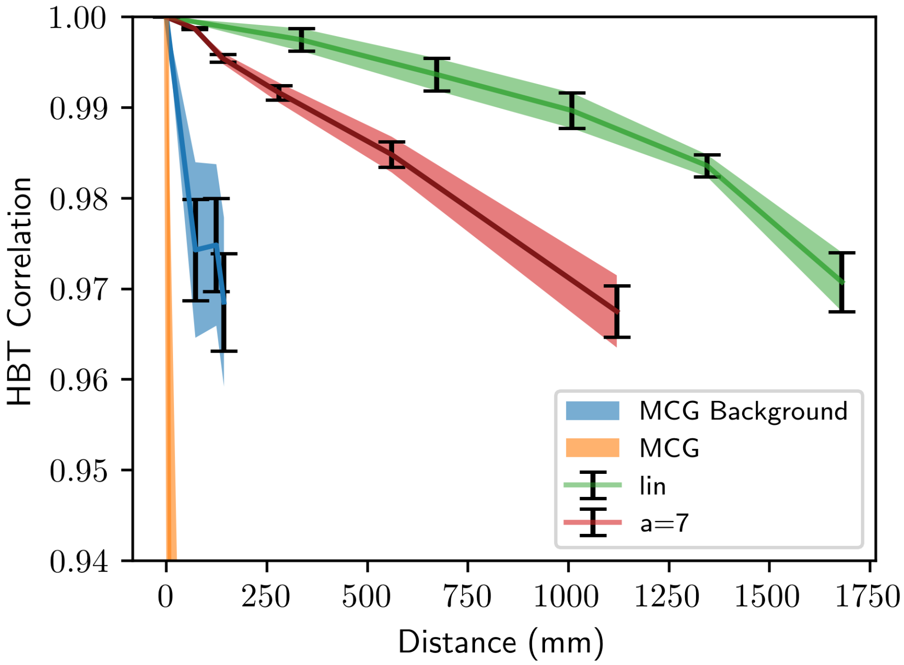

FIGURE 3.10: Comparison of HBT correlations. The teststand experiments were performed with two spacing distributions, one linear with 336mm separation (lin) and the other (a=7) spaced over a geometrical sequence ($ar^{(n-1)}$; r=2, a=70mm). These multiple hour measurements were subdivided into 10 minute slices, the slice mean and SEM are shown. The MCG background shows the environmental coherence measurred by the device array. The MCG signal is the incoherent remainder after digital signal processing.

The minimum detectable coherence is defined by zero correlation between neighbouring sensors, which is determined by the minimum sensor spacing of 72mm. The coherence of the cardiac signals (signal processed MCG) was computed, and found to be on the limit of detectable coherence. Their mean correlation at one sensor spacing was 0.3 and became negative within two sensor spacings, see figure 3.8. The cardiac signal coherence is slightly above the detection limit because the curvature of the cardiac dipole is shorter than the width of the bucket baseline (the array width).

To measure the spatial coherence over a longer baseline I built a dedicated detector array apparatus, see figure 3.9. Six magnetometers (7cm ICM) were arranged on a compressed concrete block that sat on a foam mat for vibration isolation. The sensor pre-amplifiers were powered using batteries. The signals were recorded with a Labjack T7-Pro ADC at 24bit, 1kHz. I recorded the background environment over several one hour measurements, with the ICM arranged in either, a linear spacing (uniform magnetometer distribution) or a geometric sequence $(ar^{(n-1)}; r = 2; a = 7cm)$ provided a test of proximity effects.

Both setups exhibited a higher spatial coherence than observed with the MCG device array, see figure 3.10. The linear distribution had the highest coherence width of $4.9 \pm 0.9 \times 10^7 \mathrm{mm}^2$, 15× more coherent than the MCG background. A kink is visible in the initial part of the geometric sequence curve, which is likely due to proximity effects; decoherence caused by either flux linkage between the cores, scattering of EM waves by the conductive mass of the sensors, it could also be due to cross-talk from the amplifier output or power supply noise.

A coherence spectrogram was computed by segmentation and frequency decomposition of the sensor time-series followed by application of coherence analysis to the time/freq components, see figure(3.11). The coherence spectrogram shows the dominant source of coherent noise at 50Hz (and its harmonics), originating from mains power infrastructure. Some broadband pulses are visible. In general the coherence decreases with frequency; this frequency gradient remains even if the raw data used in each time/frequency bin is individually normalised, hence the gradient is not due to the frequency response of the induction coil magnetometers. The attenuation spectrogram shows the attenuation achieved by the coherent noise rejection algorithm, which is explained in section 3.4. A visual comparison with the attenuation spectrogram shows that these spectrograms are highly correlated down to the fine-grained level, suggesting that increasing background spatial coherence allows a greater SNR to be achieved.

## 3.3 Methods for Environmental Noise Rejection

The attenuation of the coherent noise rejection (CNR) algorithm, which is detailed in the next section 3.4, is dependent on the spatial coherence of the environmental noise. To determine the precise relationship I calculated the background spatial coherence width (by fitting a gaussian) of each scan in the clinical trial, and also measured the CNR attenuation. Plotting these together in figure 3.8 reveals a nearly direct linear relationship between environmental coherence and attenuation, of 9dB per order of magnitude increase in spatial coherence.

The observed spatial coherence of a chaotic noise source is proportional to its angular size, which is due to the structure and positioning of the source. The device itself can be the least coherent noise source, either through electromagnetic interference or by decoherence of the ambient noise field. Decoherence occurs from scattering of photons

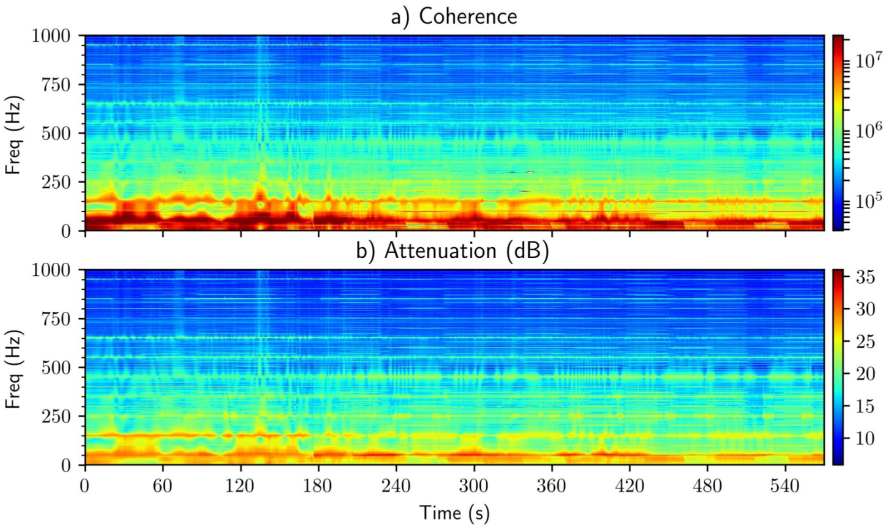

FIGURE 3.11: Comparison of a) Coherence and b) Attenuation (dB) spectrograms: A clear correlation is visible between the spectrograms at a fine-grained level. The maximum coherence and attenuation is seen in the 50Hz band and its harmonics. In general the coherence decreases with increasing frequency. Some broadband sources are present; such as the pulse at 140s.

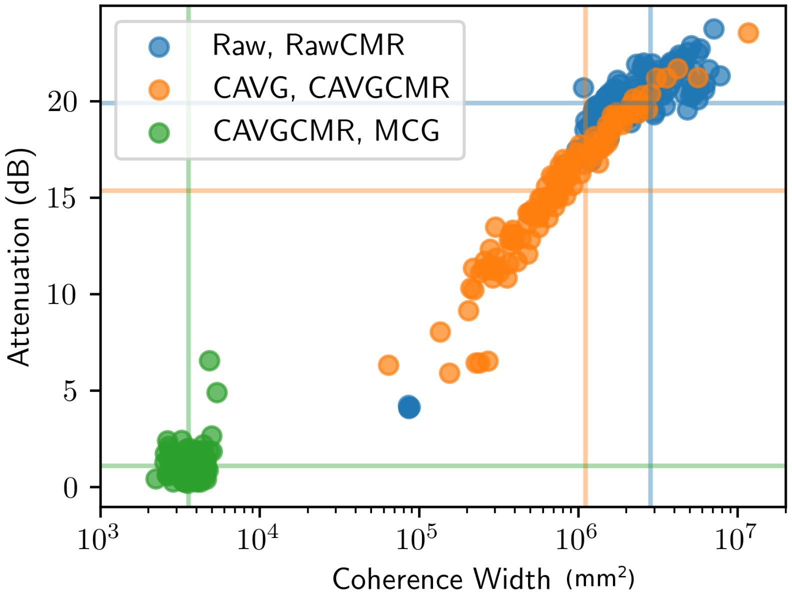

FIGURE 3.12: The CNR attenuation and coherence width of each scan in the clinical trial. A strong positive correlation between coherence width and attenuation. Cycle averaging removes both coherent and incoherent noise, spreading the coherence width; over the dataset we see a reduction in coherence after averaging, though in a small number of scans (15%) the coherence increased. The MCG plot shows the coherence of the signals before the final lowpass filter stage with the attenuation of the lowpass filter. As expected, there is no correlation between lowpass filter performance and coherence.

from the conductive material of the detector array and inductive coupling between the permeable cores; This inductive coupling was simulated in COMSOL, as seen in figure 2.17. The structure of the detector array material imposes structure on the reflected field. In certain situations the coherence of a scattered signal can be locally increased due to focusing effects, but for a thermal/chaotic field the reflected field is generally less coherent than the ambient. The clinical electromagnetic environment is largely chaotic, certainly from the perspective of a limited planar sensor array. Devices with a higher spatial measurement complexity such as MEG helmet arrays are able to capture more spatial information about the noise field. This allows them to employ signal space separation (SSS) algorithms that estimate the spatial origin of signal components. The complexity of the environment is still much higher than the array capability. Typically only a binary separation is performed, to differentiate between signals originating within or outside of the helmet array. SSS algorithms have been shown to work with planar arrays [68], though they performs best with curved arrays.

To minimise decoherence there should be no conductive material in the proximity of the measurement area. The ICM themselves are an unavoidable source of decoherence since they are composed of copper and iron. In any case the detector array and supporting structure should be composed of plastic or non-conductive composite material such as glass fiber (carbon fiber should be avoided since it is conductive). To prevent acoustic noise the magnetometers must be stiffly coupled, and the whole detector array must be isolated from vibration travelling through the floor. A high inertia is ideal, however this reduces portability. The structure of the bed that supports the scan subject has been a large source of acoustic noise, transducing the movement of the patient inducing ballistocardiographic signals. To remove this acoustic noise it should be manufactured from a non-ferrous material. It should also ideally be non-conductive to decrease decoherence.

The enclosure containing the device electronics is an unavoidable source of noise due to the switching components required for digital signal handling. The enclosure should be separated from the detector array. The optimal separation is difficult to determine. The greater the separation the longer the analogue signal pathway, which leads to noise pickup dependent on the intensity of the nose field. The amplified analogue signals are less prone to noise pickup. It is probably optimal to position them at some middle point between the enclosure and the array.

The optimal positioning of the measurement region within a given room is dependent on nearby noise sources. In LGI where we performed the first clinical trial, the room was next to a set of lift shafts. The difference between the sides of the room adjacent and opposite to the lifts had a detrimental effect on averaging time as measured by Allan variance, of approximately 10x. Before using a room for clinical MCG a room survey should be performed. Multiple magnetometer arrays could be placed in the room, perhaps one per corner and one in the centre. The magnetometer signals would be recorded over the period of a week, in order to capture the full temporal variability. Noise analysis on these signals would reveal the suitability of the room and optimal position within the room.

Given the costly alternative of installing MSR there is a potential incentive to drastically modify hospital architecture to create rooms that have improved spatial coherence, which can be dedicated to biomagnetic field mapping. Potential architectural changes include:

- Building the hospital from wooden structural components such as CLT, which is a material rising in popularity for large (18+ stories) buildings [69]. Removing the steel rebar mesh that permeates typical steel-concrete buildings.

- Location on upper levels away from traffic. Also away from high power machinery such as MRI or power distribution centres.
- Cable routing should be optimised such that the cable network has a minimal width/distance ratio. Twisted pair cabling and shielded cabling should be used wherever possible.

Twisted pair conductors should be used for all power delivery in the vicinity. Since they offer a reduction in emitted power that is proportional to the ratio of helix pitch to conductor separation. Compared to untwisted parallel conductors the emitted power can be 3-100x less [52, 53, 70]. The mains electricity wiring within the room and its neighbours should use twisted pair cables. One study by VIRGO (the gravitational wave detection experiment) assessed the impact of twisted mains cables on magnetic field noise, in particular transients carried on the mains network, and found a $\approx 10x$ reduction in the induced magnetic field amplitude over a range of distances from the wire (5cm to 2m) [71].

The maximum spatial coherence in a clinical environment could be sufficient to rival the performance of magnetically shielded rooms (40-90dB shielding factor). An increase in coherence of $10^2$ would increase CNR by $18dB$ to $38dB$, which would rival the low end shielded rooms. Experimentation or modelling is required to determine the practical limit of spatial coherence within a clinical environment.

### 3.3.1 Magnetically Shielded Rooms

Magnetically shielded rooms (MSR) have multiple layered walls consisting of conductive material for shielding electric fields and high permeability material such as mu-metal (a Nickel-Iron soft feromagnetic alloy) for magnetic field screening. They are typically used to reach noise levels suitable for the real-time measurement of biomagnetic signals with minimal signal processing, on the order of $\approx$ 10fT. The performance of a MSR is usually specified by the shielding factor, which is the ratio of external ambient field to the resulting attenuated internal field. However, their performance is not constant over frequency. Typically passive shielding performs best with higher frequencies. Lower frequencies are more difficult to shield against due to skin depth effects. Often active shielding is employed to tackle low frequency noise. Active shields consist of large room scale coils and capacitors that are driven to cancel out the incident field. Whilst effective for certain applications, they also introduce their own thermal noise.

Small scale shields have been developed to enclose the minimum volume required for a person and the measurement apparatus [46]. These shields should be drastically cheaper than room scale solutions, due to volumetric scaling. They can also be semi-portable, allowing retrofitting into an existing hospital. The coherence of the residual noise within the enclosure will be low due to the proximity of the walls, drastically reducing the performance of CNR. Therefore the shielding factor must be sufficient to achieve a good final SNR without CNR.

We tested a prototype 4 sensor array of 7cm ICM in the MSR at York Neuro-Imaging Centre. The detector array was placed inside the room, with the signal and analogue power cables passing through a small screened hole in the wall to the instrumentation outside. The MSR had a very low noise floor as it was typically used for MEG measurements. This allowed us to observe MCG signals in real-time, providing validation that our ICM had sufficient sensitivity. I later used the background recordings from YNIC to

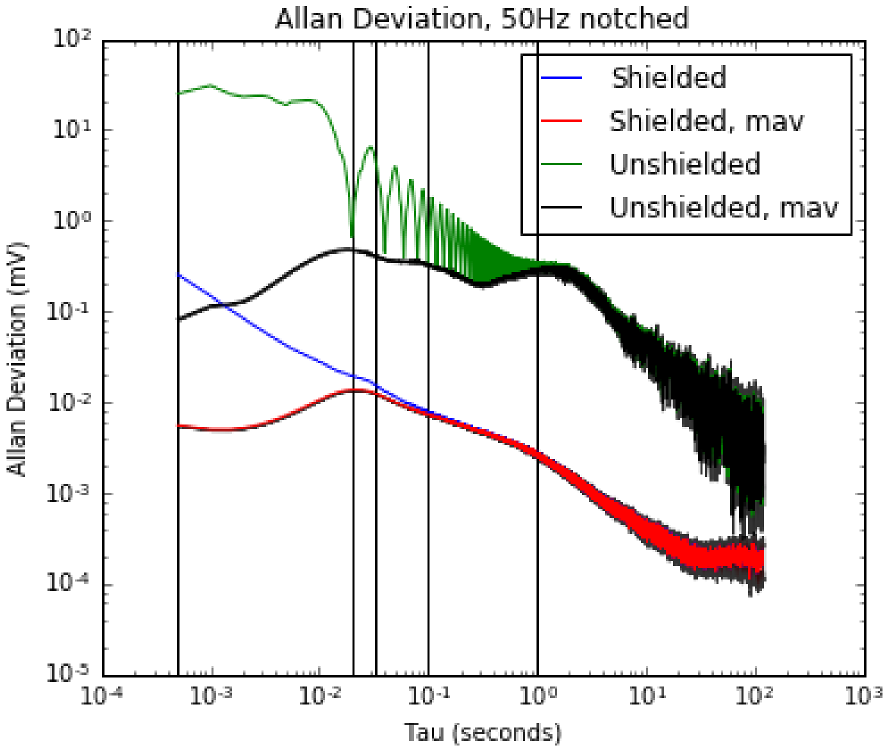

FIGURE 3.13: Allan variance comparison of YNIC Shielded environment and University lab

measure the inherent noise of the 7cm ICM and combined with a measure of its frequency response, measure its spectral sensitivity, see figure 2.8.

I found the signals recorded within the YNIC MSR to be spatially incoherent. However, this was because the inherent sensor noise was greater than the ambient YNIC field. Allan variance analysis found the noise field to be almost gaussian, with no 50Hz mains noise present. A knee occurs at approximately 20 seconds, which could be representative of an instability in the pre-amplifiers.

### 3.3.2 Procedural Techniques

There are several actions that can be performed whilst operating the device that can improve the quality of the scan data.

- The subject should remove all metal items from their person, such as jewellery, watches and bras that contain an underwire.
- The subject should be briefed to relax all muscles and avoid moving and talking. It can be difficult for some people to effect this due to pain or distraction, so a guided somatic meditation could be used.
- The device operators and other people in the vicinity should avoid movement, especially of ferromagnetic objects and electronic devices as far away from the measurement region as possible.

It is common for medical beds to have an electrically operated height adjustment. The beds we used in our lab and for the clinical trial were both electric. We observed high amplitude noise in the real time MCG signals when the bed was turned on. During a scan we needed to electrically isolate it by switching it off at the wall after height adjustment was performed, before the scan began, else it conducted mains noise.

## 3.4 Digital Signal Processing

The discretised and quantised signals are synchronously recorded by the computer, encoding the magnetometer voltages in time series of digital values, as covered in section 2.3.2. When operating in an unshielded environment these time series contain many orders of magnitude of environmental noise from a diverse range of sources. To remove this noise we have developed a digital signal processing (DSP) routine, a basic flowchart of this routine is given in figure 1.9. This DSP routine provided an average of $68.4 \pm 3.9$ dB of noise rejection when applied to the clinical trial dataset.

The noise present during the trial was estimated by comparing the ratio of variances between the captured raw signals (noise) and final stage output of the digital signal processing (DSP); the MCG signal with some noise remainder. A slice containing just the QRS region was used to represent the signal. The DSP stages were applied in the order stated and shown in figure 1.9. The mean SNR was approximately $-69.3 \pm 3.9$ dB. Averaging (approximately 600 cardiac cycles) & low pass filtering provided $50.5 \pm 2.2$ dB of attenuation. Then coherence noise rejection provided $16.7 \pm 1.6$dB of attenuation. A final low pass filter provided a further $1.21 \pm 0.7$ dB of attenuation. The final SNR was estimated by summing the attenuations; $SNR_{QRS} = 0.93 \pm 4.43$dB.

Digital signal processing has several advantages over analogue electronic filtering. Digital filters are available with zero phase distortion (FIR filter) such as the moving boxcar window average filter we used for low pass filtering. In contrast, electronic filters cause phase distortion and the electronic components used have their own noise that is superimposed on the signal. Another benefit of DSP is the ability to experiment with many routines. With electronic filtering, the routine is hard coded and cannot be changed without repeating the recording.

The effect of each DSP step is shown in figure 3.14, I have found that the end result is independent on the order of applying each step, however the compuational requirements of the remaining stages are reduced by cycle averaging first. There are four algorithms used in the standard DSP routine that are combined to create a MCG;

- Cardiac Cycle Averging (CAVG) –[Temporal Filter]
- Coherent Noise Rejection (CNR) –[Spatial Filter]
- Low Pass Filtering –[Frequency filter]
- Interpolation –[Spatial-Frequency filter]

The notch filter primarily acts to attenuate the utility grid signal but is also applied a second time as a final smoothing step to attenuate high frequency jitter. The coherent noise rejection algorithm removes common mode components, which originate from sources distant from the detector array. Cardiac cycle averaging attenuates any noise that is not synchronous with the heart beat. The cardiac cycle average is then normalised by subtracting the mean signal, which removes any DC offset.

The smallest digitisation level is 18nV, which is equivalent to 5.2fT at 10Hz. The magnetometers inherent noise is 104fT at 10Hz, which is 20 digitisation steps. In theory we

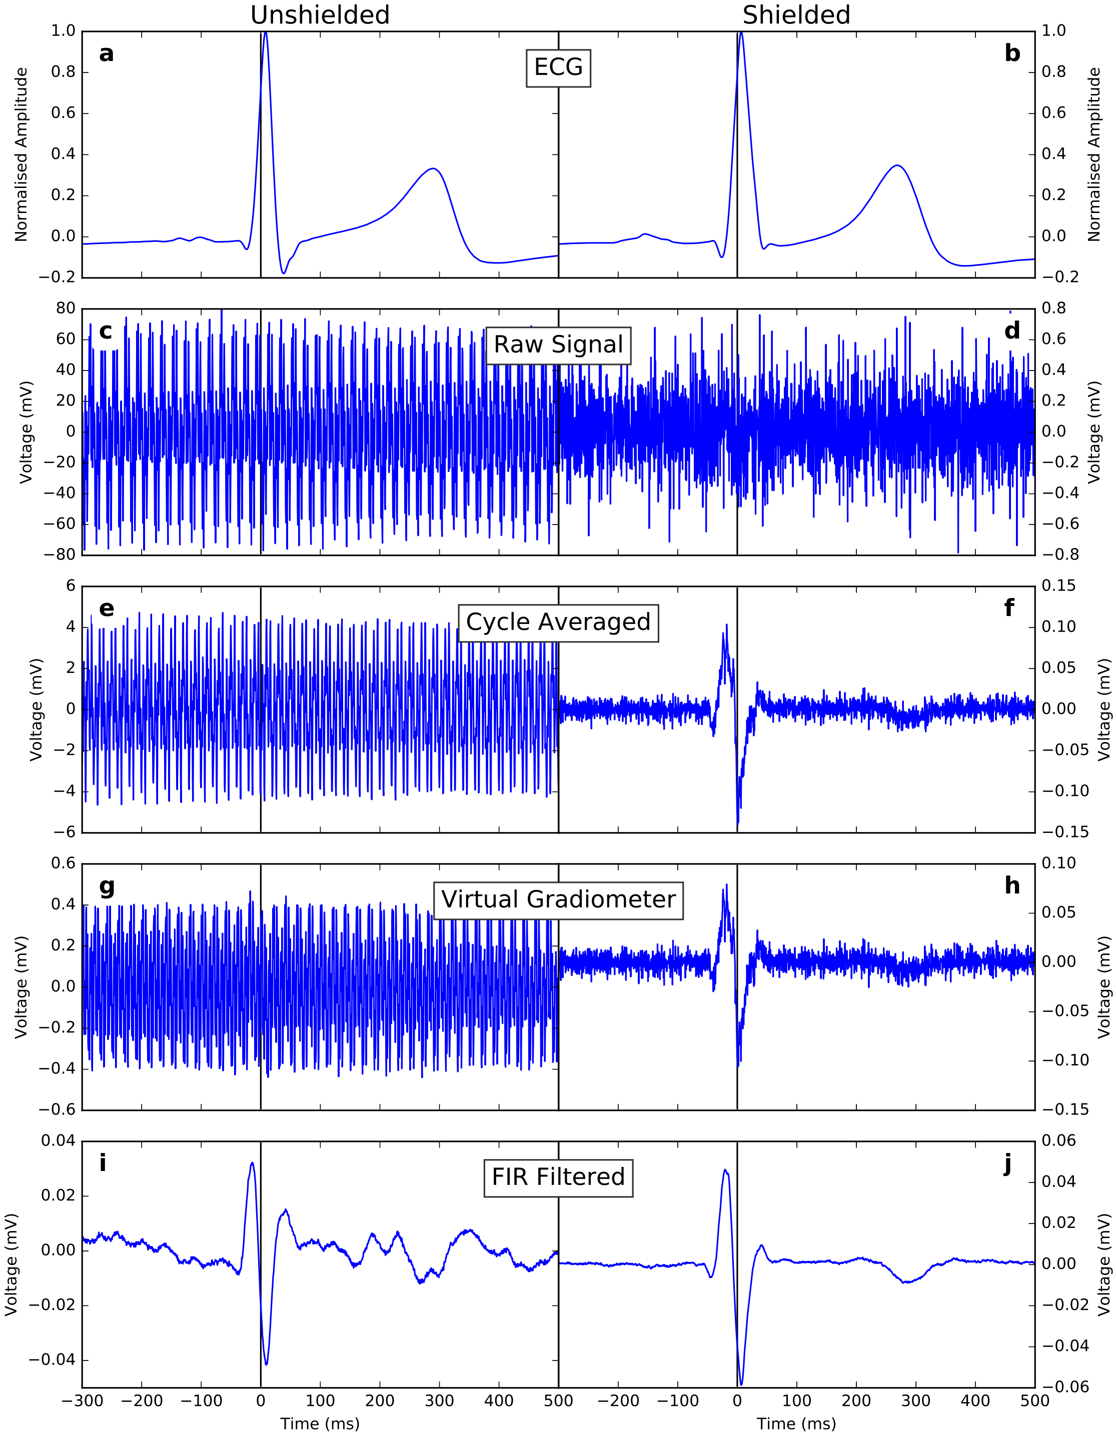

FIGURE 3.14: MCG signals from shielded and unshielded environments in sequential stages of the digital signal processing routine, with synchronously recorded ECG signals. The final unshielded signal i, still has a noise remainder, however the QRS region is well defined.

could use an ADC with 4bits less resolution and still be without effecting the measurement capability. Or to approach the problem from the opposite end; If we assume that the peak of QRS is $\approx 50pT$, then with the current 24bit system we are encoding the MCG signals with 13bits. This is more than sufficient to capture the detailed curvature. We could probably manage with around 8 bits of encoding, which would require 5bits less for a total of 19bits (for the an ADC of the same dynamic range). So in theory a 19bit or 20bit ADC should be sufficient, perhaps less with stochastic resonance effects[72–74].

### 3.4.1 Cardiac-Cycle Averaging

The cardiac cycle average is a single cardiac cycle obtained by constructively interfering multiple heart beats. This process acts as a temporal filter, attenuating components that are not synchronous with the heart beat. The ECG r-wave is used as a fiducial to create $n$ slices of MCG time series, $S^i_t$. A slice of $\{-300 < t < 500ms\}$ about the fiducial is sufficient to capture the whole cardiac cycle. The single cycle average $X$, is obtained by superimposing these slices and computing the element-wise mean:

$$X = \frac{1}{n}\left(\sum_{i=1}^{n} S^i_t\right) \tag{3.10}$$

A 10 minute measurement will contain approximately $n = 600$ heart beats, reducing the random noise amplitude by $\frac{1}{\sqrt{n}}$, by about $24\times$. However, the noise is not purely random, some noise components are periodic and will constructively interfere in the average. The mean attenuation performance over the LGI trial was $20\times$. The best performance occurs when there is a high heart rate variability, leading to an irregular beat cycle, which is not co-periodic with regular, clock-like sources. There are sources that are always synchronous with the heart beat, such as the ballistocardiographic signals transduced by the metal medical bed. It is common for our heart beats to synchronise with music; therefore it is advisable to avoid listening regular tempo beats during a measurement.

Our use of CAVG makes the assumption that all heart beats are identical, which is not necessarily true. Most people have a positive heart rate variability (HRV), which can cause the form of the cardiac cycle to change. Since the duration of the cardiac cycle shortens with tachycardia (high heart rate > 100bpm); by compression of the P-R and S-T intervals. A stress MCG would entail increasing the subjects heart rate before they lie supine for a scan, where their heart rate will gradually decrease. It is not valid to average high HR beats with resting state beats. We performed a preliminary test of Stress-MCG at LGI and noticed a severe change in dipole angle >60°, which slowly decayed to normal angle with the heart beat. This is corroborated by Hanninen's study [75, 76]. There are also transient cardiac events such as ectopic beats and other arrhythmias.

A synchronous ECG recording is required for CA. In section 2.3.2 I discussed the importance of capturing a full ECG time series, and not only a gate signal, in order to obtain maximum fiducial accuracy and so preserve the high frequency cardiac signals in the CA. It would also be useful to perform an ECG analysis to categorise the heart beats. Categories could be separately CA or rejected.

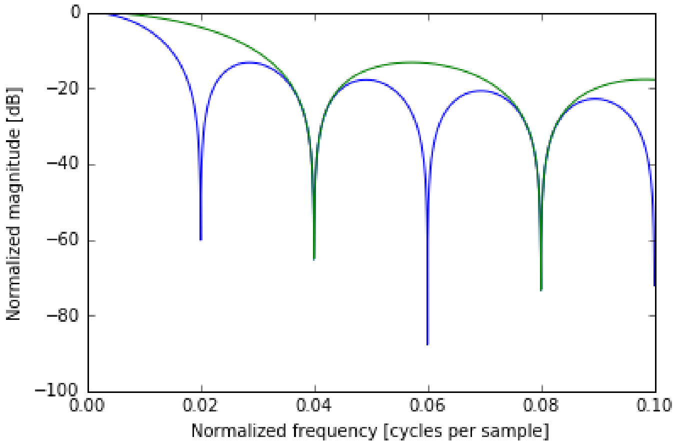

FIGURE 3.15: Frequency response of the moving average filter using a boxcar window of width 20ms (blue) and 40ms (green). The 20ms window attenuates the 50Hz utility grid noise by 60dB.

### 3.4.2 Lowpass Filtering

The lowpass filter primarily acts to attenuate the utility grid signal but it is also applied as a final smoothing step to attenuate high frequency noise. The filter is formed by a rolling window convolution of a boxcar window with the time series. The width of the window is equal to the time period of the utility grid; since the mean of a single cycle sinusoid is zero, the utility grid noise is nearly perfectly attenuated. The efficacy of the filter relies on the window width being an integer multiple of the sampling period; The optimal sample rate is 2.4kS/s since this is divisible by both 50/60Hz, which gives a window width of 48 samples. The frequency response of the filter can be seen in figure 3.15. The lowpass filter function written in Python:

```
def movingaverage(signal, window_size=48):
        window = np.ones(int(window_size))/float(window_size)
        return np.convolve(signal, window, 'same')
```

This filter is applied twice to sufficiently attenuate the high frequency noise. For smoothing purposes the window shape could be modified; Gaussian windows are common. Though the boxcar window works sufficiently well and is faster to compute.

The output of the filter often contains distorted edges where the convolution overlaps the time series boundary. These outermost regions of the time series (within one window width of the edges) must be trimmed or set equal to the mean; to account for this the width of the cycle average can be extended.

### 3.4.3 Coherent Noise Rejection

To remove spatially coherent noise we use an inverse ghost imaging approach. A reference sensor records the noise field in the measurement region but does not record cardiac signals. Subtraction of the reference signal will remove the correlated noise, leaving the uncorrelated components.

Axial gradiometers achieve this by having a reference sensor placed at a greater radial distance from the heart, since the amplitude of the cardiac dipole reduces with $1/r^3$ the relative portion of cardiac signal in the reference is small. The performance of this

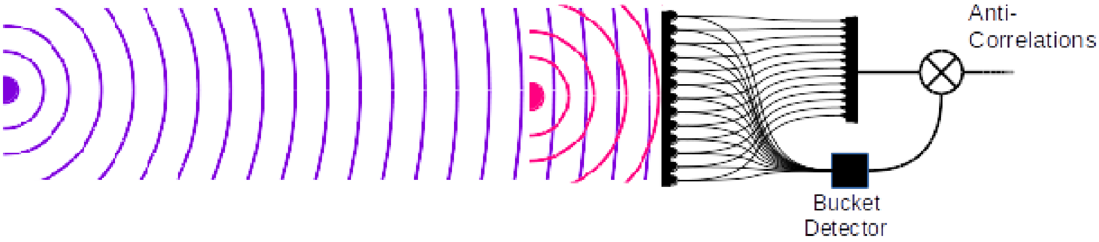

FIGURE 3.16: Illustration Comparing the spatial coherence of distant and local sources as observed by a detector array. The summing (mean) of the signals to produce the bucket is performed after digitisation. The anti-correlations are computed by subtracting the bucket signal as in equation 3.11

method is dependent on having well matched magnetometers. We avoided use of axial gradiometers because it is difficult to manufacture ICM with identical responses. The inherent noise of the gradiometer is double the magnetometer noise (assuming identical magnetometers), since the inherent noise adds during subtraction.

Our approach was to create a virtual reference sensor composed of the mean value of all sensors in the array, called the synthetic bucket signal. The inherent noise of the bucket is reduced by spatial averaging of the n sensors by a factor of $\frac{1}{\sqrt{n}}$. The bucket signal contains minimal cardiac signal; since the cardiac field is a dipole, it spatially integrates to zero over the measurement region (assuming symmetric intersection of the array with both poles). The algorithm is illustrated in figure 3.16. Recalling equation 3.4, the uncorrelated signal offset $J_0$ that contains the cardiac signals is approximately obtained by subtraction of the bucket signal:

$$J_0^n \simeq J_s^n - \frac{\sum_{n=0}^{N} J_s^n}{N} \tag{3.11}$$

The uncorrelated signal offset contains a factor of $\frac{1}{\sqrt{n}}$ of the inherent sensor noise.

### 3.4.4 Interpolation

Radial Basis Function Interpolation is used to create a high resolution magnetic field map from a smaller set of spatially separated samples, by fitting a smoothly varying surface between the original samples [77, 78]. To achieve this the algorithm assumes that the field does not contain information of a higher spatial frequency than the array spacing frequency. This is effectively a low-pass spatial frequency filter, excluding high frequency components above the original spatial Nyquist frequency. There are several algorithms available. A basic linear algorithm interpolated with linear gradients between the samples, which is suitable for certain 1D applications. But it creates a faceted surface in 2D which is not a realistic physical magnetic field; since the transitions between the facets have infinite curvature. A cubic interpolation produces a smoothly varying field that is more realistic.

My python implementation took the sensor values with their x, y coordinates and used the Rbf function from scipy.interpolate to interpolate them onto a numpy meshgrid. The extent and resolution of the meshgrid were set to give a 92x64 grid of 4mm spaced pixels.

```
from scipy.interpolate import Rbf
def interpolateSensors(x, y, z):
    """Interpolates a scatter of sensor values, z with coords x, y.
    Returns interpolated Map as a tuple; xi, yi, zi
    """
    xi = np.arange(mcg.xmin, mcg.xmax, mcg.gridspacing)
    yi = np.arange(mcg.ymin, mcg.ymax, mcg.gridspacing)
    xi, yi = np.meshgrid(xi, yi)
    rbf = Rbf(x, y, z, function='cubic')
    zi = rbf(xi, yi)
    return xi, yi, zi
```

Extrapolation occurs when the interpolation algorithm is used to generate points outside of the minimum bounding box formed by the outermost magnetometers. Gradients at the boundary are continued without bound. Especially when using a cubic interpolation, which can form convex features outside the bounding box. It is not appropriate to extrapolate ad infinitum, but it is also not representative to draw a bounding box between the centres of the perimeter sensors, since the effective area of those sensors extends further out. As a compromise a virtual set of sensors is created surrounding the array. Their magnitudes are set to the mean value of the sensor array. In this way the extrapolation is curtailed with the least assumption.

### Boundary Constraints

The first boundary I implemented was a rectilinear arrangement of 8 sensors, see figure 3.17. Whilst this was better than having no boundary, it was sub optimal; because its density was not high enough it caused distortion and the radial distance of the boundary was too high, allowing too much extrapolation. I later developed an improved boundary, see figure 3.17. The virtual sensors were arranged on the same lattice as the actual magnetometers which reduced distortion. Also a double layer of sensors was employed to provide a ring of flat gradient, which curtails extrapolation of inner gradients. A third layer could be employed to prevent curvature (2nd order gradient) from being transmitted, but this does not seem necessary.

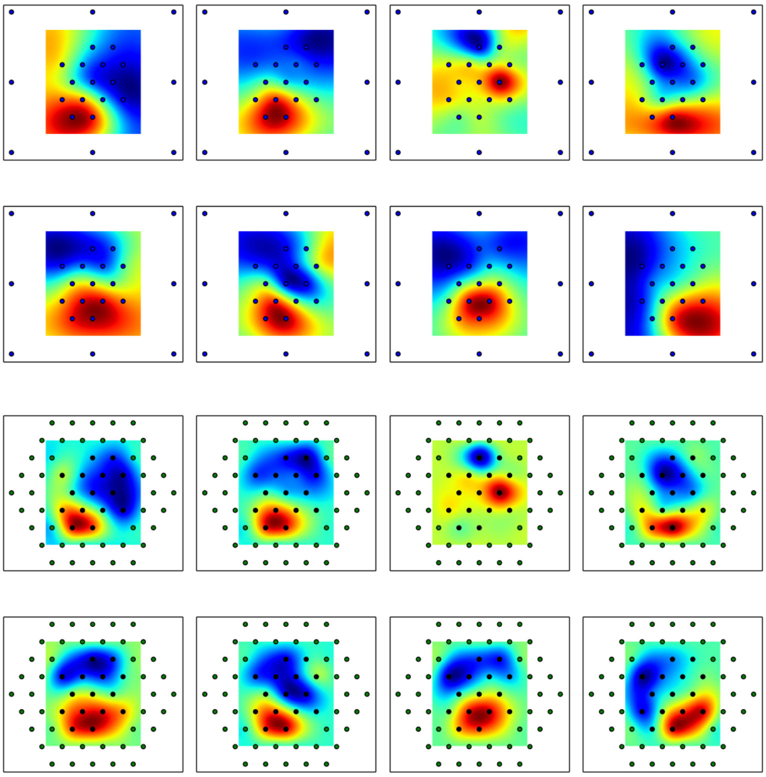

FIGURE 3.17: Elementwise comparison of the original boundary (top two rows) used to prevent extrapolation with the optimised version (bottom two rows.

### 3.4.5 Integration

ICM magnetometers record the time derivative of the magnetic field $dB(t)/dt$. SQUID magnetometers measure the magnetic field time series $B(t)$. The magnetometer signals can be integrated to yield $B(t)$. This would be useful for comparing MCG temporal dynamics recorded by our device with other MCG devices, it would also allow use of the same diagnostic algorithms.

Integration acts as a low-pass filter. However, there is little low frequency information recorded, because ICM's have a linear frequency response starting at zero for static magnetic fields and increasing until their resonant frequency (LC - Due to parasitic capacitance). Also a DC offset control loop exists in the pre-amplifier circuit to prevent baseline wander. This acts as a high-pass filter with a cut-off of 1.6Hz. Furthermore, the signals contain electronic flicker noise, which dominates at low frequencies since its intensity varies inversely with frequency $1/f$. Hence integration of ICM signals results in a signal dominated by low frequency noise; baseline wander, figure 3.18 shows the highest quality integrated signals I was able to obtain from within the MSR at YNIC.

I think that it is better not to integrate. The derivative contains the same information, but shifted in phase and frequency. I directly observed this shift in diagnostic information, see section 4.3.3. Performing diagnostics in the derivative has the advantage of greater sensitivity to peak activity. Since a maximum in the integral becomes a zero crossing in the derivative, and zerocrossings are higher definition features; low sensitivity to noise.

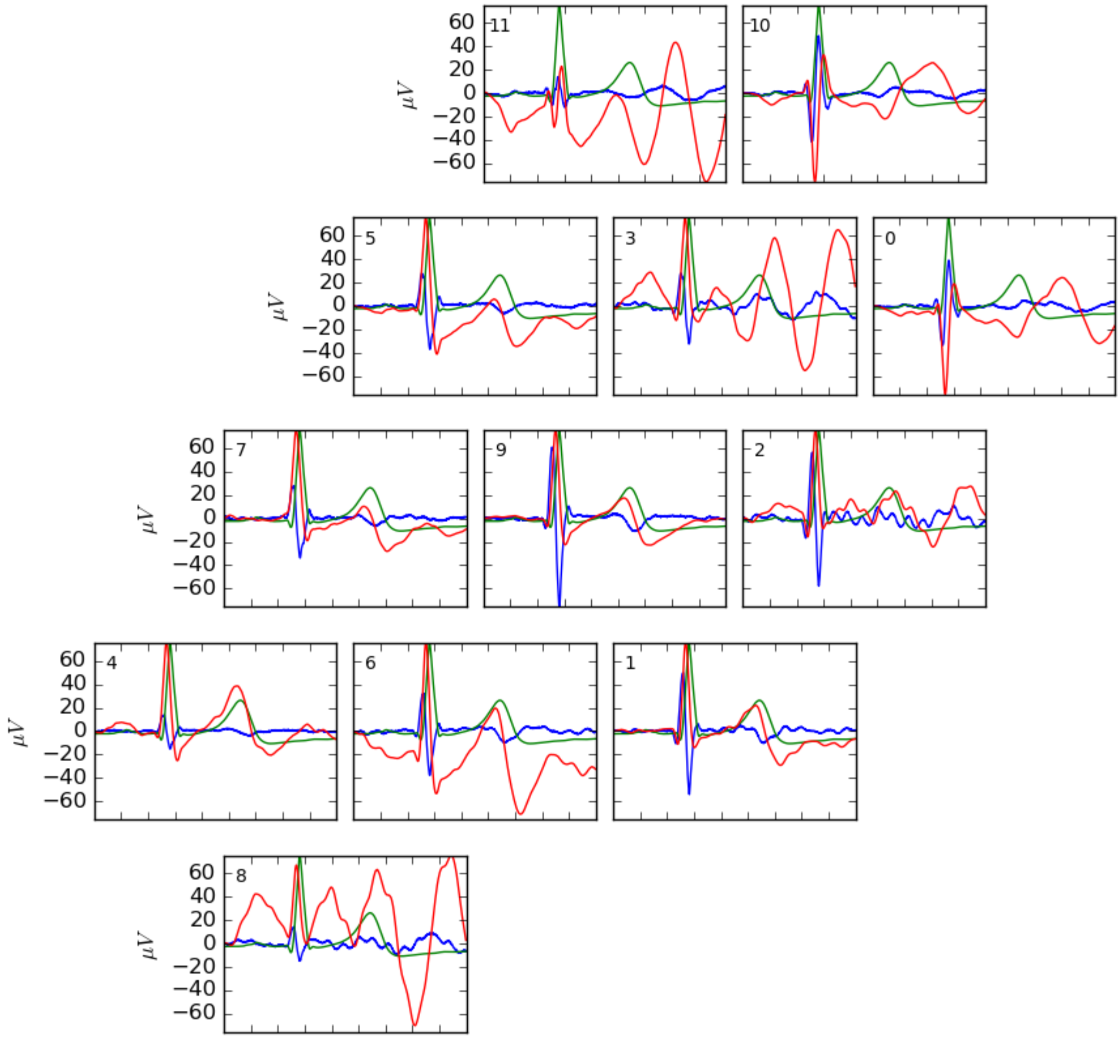

FIGURE 3.18: Signals recorded by the 7cm ICM array within the YNIC MSR. The MCG traces are shown in blue, the integrated signals in red and the ECG is shown in green.

### 3.4.6 Exploration of Supplementary Denoising Algorithms

I have investigated a several other signal filtering algorithms including an adaptive bucket subtraction that I conceived of, Empirical Mode Decomposition[79] and Reference denoising [80]. None of the methods I explored significantly improved SNR accept for reference denoising. However, it caused significant distortion to the MCG. I concluded that it was not productive to explore these methods further at this development stage due to the risk of synthesising signal from noise. In order to develop supplementary methods further it is essential to assess them using fidelity; This could be achieved using the phantom to produce a known signal that can be detected at high amplitude (high SNR) and at a similar magnitude to the cardiac field.

#### Adaptive Bucket Subtraction

The purpose of the bucket detector is to capture the coherent background. I was inspired to improve its performance by the fine grained, transient variations in spatial coherence that were uncovered by time-frequency analysis (see figure 3.11). Another inspiration was gradient subtraction; by having separate weights for each sensor in the array during summation of the bucket signal, gradients in the background can be mapped. I hypothesised that by optimising the weights such that the spatial coherence is maximised, the bucket detector should be improved. I implemented this adaptive bucket subtraction algorithm, on a cardiac cycle basis. Each cardiac cycle in the scan had its spatial coherence optimised using an Adam Optimiser.

The result was a small increase in ELF and high frequency noise rejection, but a small but broadband decrease in noise rejection in the mid-spectrum. Perhaps the result would be better if the weight optimisation was performed on smaller windows, or on overlapping windows. However, that would increase the computation time, which was already on the order of a day to compute for the trial dataset.

#### Signal Space Separation

Signal space separation (SSS) is usually applied to MEG helmet array data in order to separate the signals from within the subjects head from those outside. The algorithm fits a set of spherical harmonic expansions to the data. One set from within the helmet, the other from outside. The curved geometry of the helmet array is suited to this. However, Sekihara recently showed that it is possible to apply SSS to planar arrays, provided careful choice of the separation boundary position [68]. This method could be applied to our devices MCG signals, and might result in an improvement over the standard coherent noise rejection algorithm.

## 3.5 Conclusion

A suite of noise analysis techniques were employed to understand the noise environment. This informed the development of digital signal processing routines for removal of different noise types, and allowed particular noise sources to be detected and combated. The ability to perform filtering and analysis in the digital realm, as opposed to analogue hardware filtering, increased the iterative development rate since testing of routines could be performed repeatedly and separate to measurement. The digital signal processing routine that we developed was able to separate magnetocardiographic signals from the environmental noise present in a clinical environment. The clinical trial achieved a final SNR of $SNR_{QRS} = 0.93 \pm 4.43$dB. The key to achieving this without using a magnetically shielded room or gradiometry (which halves the sensitivity of each element in the MFM array) was the creation of an alternative spatial filter that provides coherent noise rejection.

The coherent noise rejection algorithm has an excellent attenuation of background noise, providing an average of $1.21 \pm 0.7$ dB of attenuation during the 6060 trial, where the average environmental spatial coherence was $2.8 \pm 0.9 \times 10^6 mm^2$. The CNR performance was found to increase by 9dB per order of magnitude increase in spatial coherence. The required level of coherence could be specified by the coherence width ($mm^2$), which would allow it to be regulated. The attenuation performance could be increased to rival MSR provided a sufficiently spatial coherent environment. To achieve this, the mains wiring distribution could adopt use of twisted pair conductors and be used sparingly or isolated from the biomagnetic measurement locale. Also other devices could be isolated from the environment or required to have minimal decoherence properties. The practical limits to spatial coherence within a working hospital are unknown. It would be interesting to conduct a complete spatial coherence survey of a hospital in order to identify potential areas for improvement. A useful approximation could be achieved using a simulation based on architectural plans.

Biomagnetic devices themselves can be optimised to reduce their decoherence effect. I observed a decoherence effect in the ICM sensors, detectable via an extra reduction in HBT correlation as the ICM were brought into proximity. This could be improved by removing the permeable cores to reduce inductive coupling, removing the pre-amplifiers. However, the decoherence from the copper windings is unavoidable. All-optical sensors could achieve a lower decoherence, however optical sensors typically rely on a control field induced by an active coil, that will cause decoherence. In general the material composing the detector array and surrounding it should ideally be non conductive (and non-ferromagnetic).

An adapted Allan variance analysis was used to observe the long term stability of environmental noise, it proved useful for predicting the expected cycle averaging performance over a scan and detecting particular noise sources such as the ELF intermittent noise from the passing lift shafts in LGI. The frequency domain was evaluated using FFT and wavelet decomposition techniques. The spatial coherence of the noise was evaluated using HBT correlations.

The diverse array of electromagnetic noise sources present in the clinical environment includes other medical devices. Though medical devices must conform to electromagnetic compatibility standards to ensure they do not interfere adversely with other devices, these standards do not regulate the low frequency range occupied by biomagnetic signals ($<$ 1kHz). This results in a hostile and unpredictable environment, which requires robust noise rejection techniques to operate within. In the future, new regulations controlling spatial coherence and low frequency emission could foster the widespread use of biomagnetic field mapping and magnetocardiography.

## Chapter 4

# Diagnostic Information

This chapter covers the extraction of diagnostic information from MCG using binary classifiers. A range of classification techniques are employed, from traditional statistical learning approaches to recently developed machine learning techniques. Excellent diagnostic performance is found in the clinical trial data, especially when using convolutional neural networks. I developed a diagnostic information mapping technique that reveals the average concentration of diagnostic information within an MCG dataset, this was useful for validating the learning process and could inform future device optimisations. The support vector classifier was developed in collaboration with my student Dan Wilson. The convolutional neural network was developed in collaboration with my student Mike Smith.

## 4.1 Introduction

The primary goal of this research has been to extract diagnostic information from the magnetic field of the heart, in order to determine the presence of heart disease. I have limited the detection to a binary classification of heart disease (the determination of whether a heart is healthy or unhealthy), rather than a more nuanced multi-class classification. Because binary classification provides an overall measure of diagnostic performance and is suited to the devices primary application in triage of chest pain.

In general a binary classifier is a function that produces a single bit output from a number of inputs. These inputs form a multi-dimensional space, which the function divides into two regions, each belonging to a possible output state. The boundary between these two regions is called the decision boundary. The performance of a classifier function is dependent on its ability to describe a representative decision boundary. There are two broad approaches to creating this decision function, statistical modelling and machine learning. In general the machine learning techniques are able to model higher complexity decision boundaries. This is excellent for capturing complex relationships between the input variables. However, to achieve a good general performance (the ability to predict unseen data), this boundary must not over-fit the training data. Over-fitting occurs when the boundary is over complex and represents the noise in the training data, see figure 4.1. It is easy to produce a classifier that can predict with 100% accuracy the class of its training set, yet perform randomly on unseen data. Some techniques exist that can be applied during training of machine learned classifiers to reduce over-fitting. The principal technique is to reduce the dimensionality of the input data.

The first step in creating a diagnostic is to simplify the input data, by reducing its dimensionality (the number of dimensions). This is due to Hughes' phenomenon, which

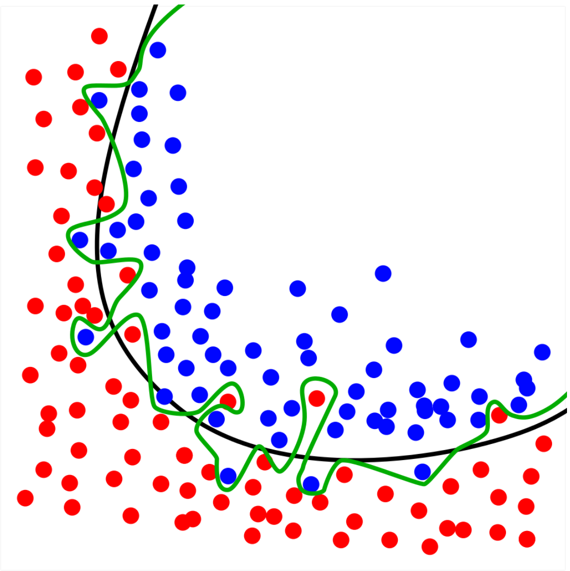

FIGURE 4.1: Example of decision boundary in the case of over-fitting (green line) and a good general representation (black line) [81].

states that as the dimensionality of a pattern classifier increases above the sufficient minimum, its predictive power decreases [82]. Each MCG in the 6060 clinical trial consists of $2000_{samples} \times 15_{sensors} = 30000$ individual measurements each of which is a separate dimension. These are interpolated to produce a magnetic field map time series: $2000_{samples} \times 92 \times 87 = 16 \times 10^6$. The dimensionality can be reduced spatially and temporally. The spatial resolution of the interpolated MFM is much higher than required to describe the observed dipole features and therefore can immediately be reduced to $17 \times 19$, giving $6.5 \times 10^5$ dimensions. This dimensionality is suitable for machine learning techniques, but still to high for statistical modelling techniques. When training the support vector classifier we found better performance by sub-sampling the time series data every x samples, where x is the with of the moving average filter window; this resulted in a time series of $50 \times 15 = 750$;

For statistical modelling we must perform manual feature engineering; extracting a set of single valued parameters that capture the important diagnostic information. These parameters are also useful outside of the diagnostic process, as they can be related to biological knowledge for a wider understanding of disease physiology. This understanding can then feedback into the device design to improve measurement performance. There have been previous attempts to parameterise MCG, which we can borrow and expand upon.

## 4.2 Parameter Engineering

The aim of parameter engineering is to obtain a compressed representation of MCG that captures the diagnostic information with minimal dimensionality. This representation consists of a set of parameters that can be computed from an MCG timeseries, which can then be used in a statistical learning to produce a diagnostic classifier. Furthermore, by using previously created parameters we could compare the output of our device with other MCG devices. Shima Ghasemi-Roudsari performed a thorough literature review

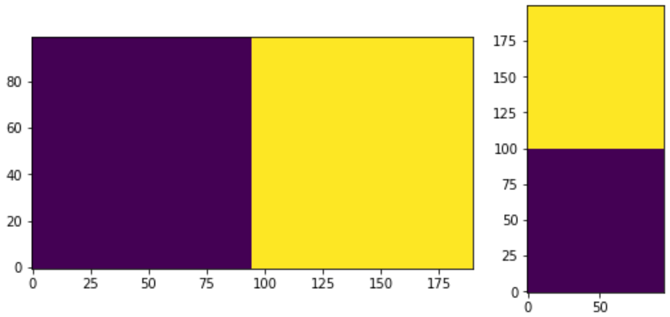

FIGURE 4.2: Match templates imallx and imally on the left and right. The purple regions are -1, the yellow +1.

to find parameters used by other groups[41, 83, 84]. I created algorithms to extract those parameters, and some novel ones, from our data and together we applied them via an analysis workstation program that I wrote. Shima provides a detailed analysis of this process and our results in her thesis [41]. To assist in parameter extraction the workstation contained a number of automated extraction routines to extract standard temporal and spatial features.

Extracting parameters requires algorithms, some are simple to implement, other parameters involving dynamic features are more difficult to acquire; since they require the identification of features within the cardiac cycle. For example the interval between R-T peaks varies from person to person, therefore the peaks must be found and identified. Using human pattern recognition is one option for dealing with this, however, compared to computers humans have limited throughput rate, produce variable output and are error prone.

In the end we used a composite method, using a human interface to validate and adjust the automated methods. The automated temporal feature detection algorithm is detailed in section 4.2.2. A typical healthy MFM is a dipolar field, though in certain abnormalities a quadrupole field can be observed. Spatial feature extraction was achieved by treating the MFM as a basic dipole field, entirely described by the 2D locations of its poles; or equivalently the center position, pole separation distance and relative pole angle. The dipole center position relative to the heart geometry can only be approximately known due to the MCG array positioning method. Likewise, the absolute pole separation is dependent on the normal distance between the array and heart. The relative dipole angle is known to a comparatively low uncertainty and is an important, established diagnostic parameter.

### 4.2.1 Dipole Angle Measurement

The standard method for measuring the dipole angle is to find the peaks of the positive and negative poles using a peak finding algorithm, the dipole angle is the angle of the vector between those peaks. This technique suffers from a macroscale Barkhausen noise [85]. Since the flux is captured unequally by the discrete number of cores in the array, as seen in the flux modelling results, figure 2.17. Also the effect is not removed and potentially amplified by the non-linear interpolation algorithm.

The dipole angle measurement can be improved by using the total area of each pole in the calculation of the dipole vector. To achieve this a pair of feature masks, imallx and imally are convolved with the MFM and the maximum convolution of each is recorded,

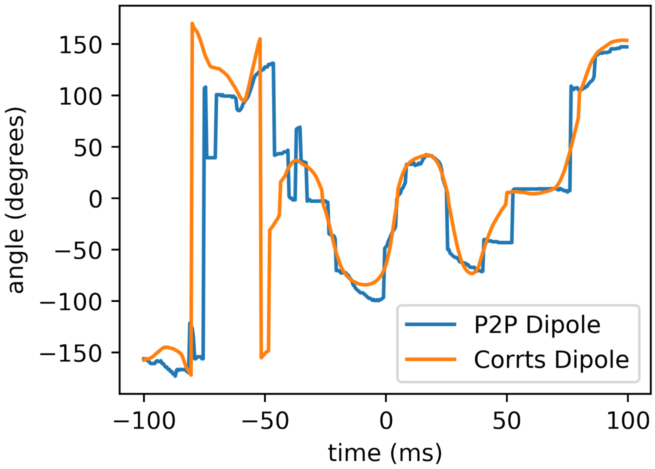

FIGURE 4.3: Comparison of dipole angle measured using pole peak to peak (P2P) and the correlation method (corrts). The time axis shows time relative to R peak. The corrts method is smoother, does not exhibit discrete steps, and captures a rotation through 180° that is missed by the P2P method at -50ms.

see figure 4.2. The imallx mask is translated from left to right and the imally mask from top to bottom. The position at which the convolution is maximumised is the centre of the dipole. The relative value of these maximum convolutions (x,y) determines the dipole angle, which is calculated using $180arctan2(y,x)/\pi$. Applying this to every MFM in the cardiac cycle produces a dipole angle time series (CorrTS), see figure 4.3. Comparing the P2P and CorrTS approaches, the CorrTS is much smoother and does not exhibit discrete steps like the P2P (barkhausen style noise). The CorrTS also captures a rotation through 180° that is missed by the P2P method.

This corrts could be used to extract dipole angle parameters (including dynamic parameters) as an alternative to pole peak to peak measurement. It is also useful for detecting zero crossings (ZC); positions where the MCG traces all cross through zero amplitude as the dipole is angled out of the plane of the array, or the dipole amplitude reaches zero; because we measure derivative MCG (dMCG) the zerocrossing is an extrema in MCG. ZC are not always easy to define on the MCG time series, since there is often dispersion of the point. Looking for ZC in the corrts provides a spatially averaged method of detection, which seems to give a higher precision. ZC are a common manifest in the dMCG we observe, and were used in our standard parameter analysis to analyse MCG, see figure 4.6.

### 4.2.2 Labelling MCG Features

The MCG we observe typically have a 'butterfly' structure of lobes between the ZC, as seen in figure 4.6. The peaks of these lobes correspond to the maximum gradient of magnetic field, which is proportional to action potential velocity. These peak positions were used in the parameter analysis we developed [41]. I wrote a program to allow a standard set of markers to be moved around by a human operator to fit each MCG. The process of aligning the markers to features in the MCG was time consuming and had a subjective element. To aid the time efficiency I created an automatic marker alignment system using wavelet decomposition.

A wavelet decomposition is similar to a fourier transform, but instead of decomposing into sinusoids of varying phase, frequency and amplitude, the signal is decomposed into a phase, frequency and amplitude components of an arbitrary wavelet. The wavelet transform is performed by a convolution type operation, for more information see section 3.2.2. Different wavelets will resonate with different features in the MCG time series; In particular I found two wavelets that were effective at MCG feature detection when used in combination, the Paul and DOG wavelets.

A Paul wavelet has multiple oscillations, causing the Paul wavelet transform to produce a decomposition shown in figure 4.4. The depolarisation and repolarisation phases manifest as separate regions, each region appears as a single entity. The peak of each entity is representative of its centroid, as such I use the distance between depolarisation and repolarisation peaks as a measure of R-T interval. Another definition we used was based on zerocrossings, see figure 4.6

The DOG wavelet is effectively a native peak finding transform. It was used to find peaks within the MCG, as seen in figure 4.5. A DOG transform was performed on each sensor time series in the MCG, and the mean peak position within a predefined area was taken as the peak. This was more difficult for the multi-lobed repolarisation region, where the peak positions were first clustered around the centroid found in the Paul wavelet transform.

These peak positions were used to align the standard lobe markers (QR, RS, T1, T2, T3, T4), followed by a human validation process, for error checking. This standard we created is shown in figure 4.6. It is inherently flawed due to the ballistocardiographic nature of the repolarisation signals that we later discovered. It also fails to account for the S-J peak that usually appears.

## 4.3 Diagnostic Methods

The aim was to produce a diagnostic binary classifier, able to differentiate between healthy (disease-negative) and unhealthy (positive) hearts from the MCG data. In this section I discuss the methods I have explored to build binary classifiers, starting with the traditional statistical methods and then covering machine learning methods. The use of statistical learning methods are common in medicine since their working is transparent and a large, well understood body of knowledge exists to determine the reliability of their predictions. The machine learning methods by contrast use iterative, self learning algorithms to discover diagnostic information in the input data. Their prediction performance is capable of surpassing the statistical methods. However, it is non-trivial to observe the workings of machine learned classifiers, especially with high dimensionality. This makes it difficult to obtain a reliable measure of uncertainty.

A receiver operating characteristic (ROC) curve is a plot of sensitivity (true positive rate, $\frac{TP}{P}$) versus 1-specificity (false positive rate, $\frac{TN}{N}$). It is mapped out by varying the classifier decision threshold. When operating a classifier the threshold must be fixed, this decision is a trade off between sensitivity and specificity. The area under this curve (AUROC) is a measure of the overall prediction accuracy. A random chance prediction has an AUROC of 0.5, the line $y = x$. The threshold that gives maximum classification accuracy is determined by the knee of the curve. However, it is sometimes preferable to bias for a high sensitivity or high specificity depending on the application. For example, using the MCG for triage of chest pain as a 'rule-out' device for heart disease would favour high sensitivity at the expense of an increased false positive rate.

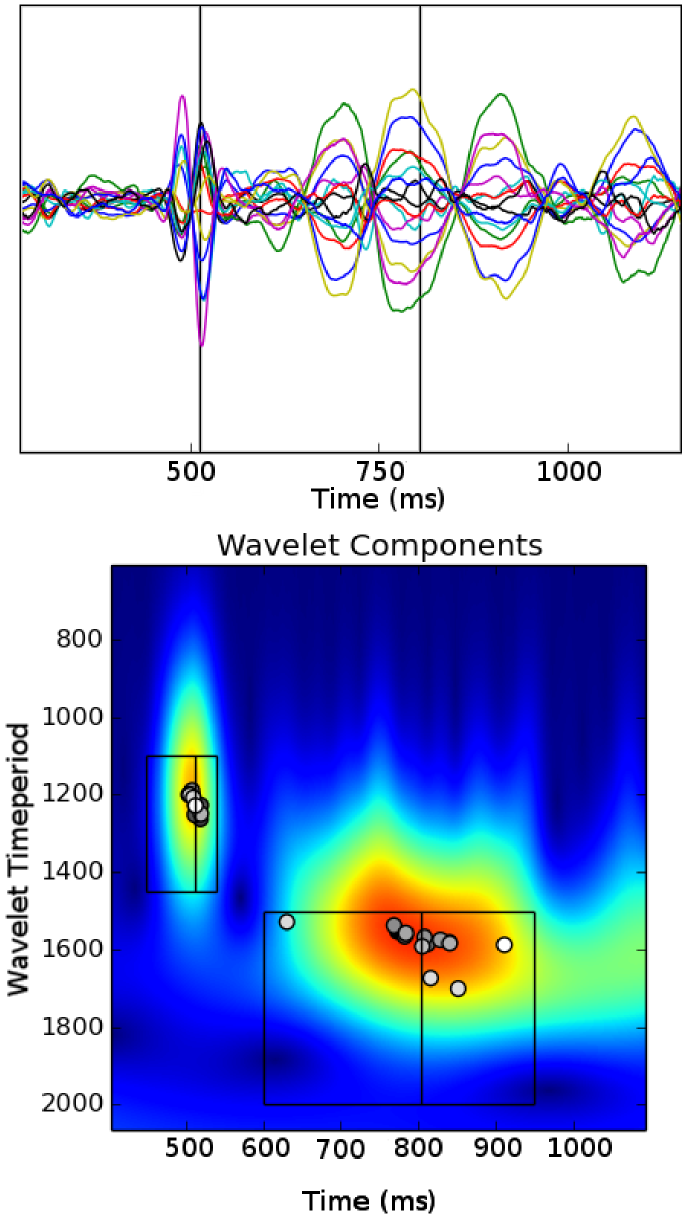

FIGURE 4.4: A Paul Wavelet decomposition of a single magnetometer time series. The wavlet transform peaks from the other magnetometers are shown in circles, and the mean peak position denoted with a black line. This technique measures the centroid of the depolarisation and repolarisation regions. The centroids are shown on the MCG butterfly plot above.

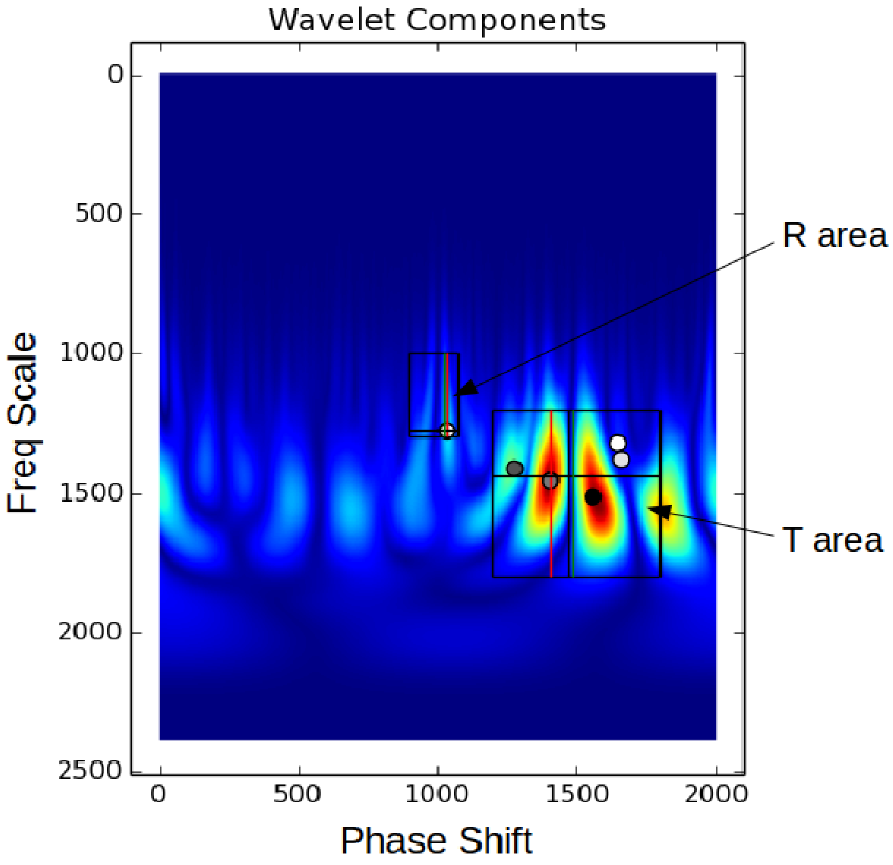

FIGURE 4.5: A peak finding algorithm applied to the DOG Wavelet decomposition of a single MCG time series. The peak finding area is constrained to the two boxes shown. The greyscale color of the peak markers indicates their magnitude. The centroid of all peaks is shown by the black cross hairs. The red vertical line shows the peak found from the Paul wavelet decomposition.

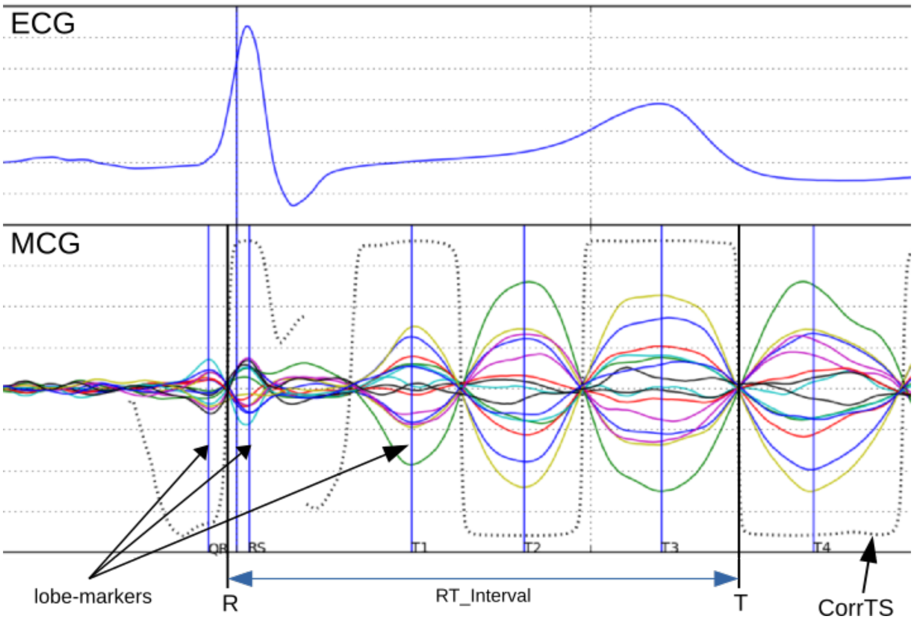

FIGURE 4.6: The proposed labelling system for MCG, shown annotating a set of MCG time series. The CorrTS trace is useful for reliably highlighting the MCG zero crossings.

To measure the general learning performance, a classifier must be tested on unseen data; data that was not used for training (cross-validation). Without this there is a risk of overfitting; were the model learns the specifics of the training data, including the noise, but fails to build a general model. There is less risk of this with statistical methods, but for the machine learning approaches used later on, the complexity of the decision boundary is great enough to map all the noise in the data, achieving perfect overfitting. Regularisation techniques and dimensionality are both used to help learn a general model.

Typically a k-fold Cross-validation is used. This method allows all of the data to be used for training, by rotating batches, whilst always validating on unseen data. Each batch uses $1/k$ of the available data for validation and the remainder for training. The training process can be repeated k times with different, non-overlapping validation subsets chosen each time.

The class imbalance problem is encountered often in medicine [86, 87]; An in balance in the distribution of training samples between the classes leads to high sensitivity with little specificity. To counter this, during optimisation of the training function the loss function (which computes the penalty for incorrect classification) can be biased to counter the imbalance by weighting the loss of each class based on its prevalency. The 6060 trial is a well balanced dataset, but the random subsets used for training in cross-validation could have been unbalanced if not sampled equally from each class.

### 4.3.1 Statistical Machine Learning

The classical approach to creating a diagnostic is statistical modelling. Patterns are discovered in the data using statistical methods such as determining p-values, testing for the assumption of Gaussianity, and computing correlation coefficients between parameters. Using this information a number of parameters can be rejected based on the minimum redundancy maximum relevance (mRMR) approach. This calculates the parameters that

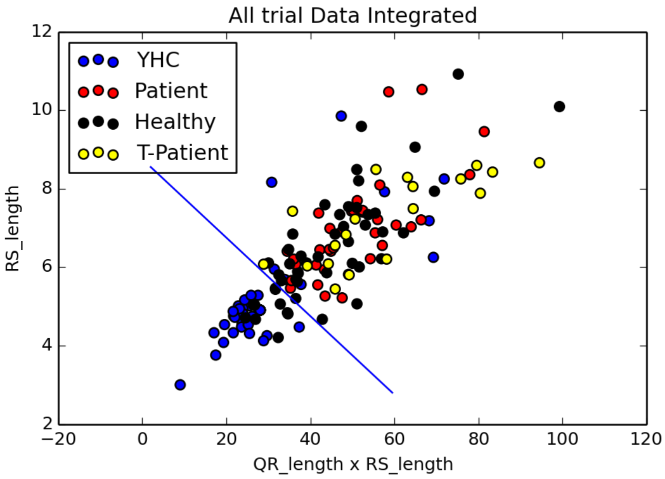

FIGURE 4.7: Linear decision boundary, with 100% Sensitivity. Young healthy control (YHC) group are more polarised than the age matched HV of 6060.

provide the highest diagnostic information (response variable), without redundant encoding.

These parameters are then used in construction of a classifier. The most basic classifier conceivable is a set of independent parameter thresholds. These can be combined using logic in a variety of ways, such as; if any parameter is above its threshold then class as unhealthy (positive). The shape of the decision boundary produced by this classifier can only be rectilinear. As such it is not able to capture relationships between parameters.

An upgrade to the independent constant threshold approach is with a linear decision boundary. The input parameters are treated as a vector, the dot product of this vector with a weight vector is calculated. This vector product is thresholded to produce a binary linear decision boundary. Figure 4.7, shows a linear decision boundary of two parameters, chosen to achieve 100% sensitivity. The young healthy control group (YHC) is composed of 50 scans of students from the University, all under the age of 30 years. The YHC group is much more polarised with the patient group compared with the age-matched healthy volunteers. This indicates an average deterioration of cardiac health with age and potentially reflects insufficient quality control in the healthy volunteer selection process.

### Logistic Regression

A common approach in medicine is to use a logistic regression classifier, which has a decision boundary that consists of a polynomial constructed from the parameters, this polynomial is the input to an exponential function. When non-linear polynomials are used this function can represent curved decision boundaries.

Abbas Al-Shimary and Paul Wathall performed a logistic regression analysis on the clinical trial data, choosing 4 parameters from the set of 43. The resulting model had a predictive accuracy, as determined by the area under the ROC curve of $AUROC = 71.2\%$. When setting the sensitivity to 100% (all disease detected), the resulting specificity was 30%. Meaning that 70% of healthy participants are deemed patients.

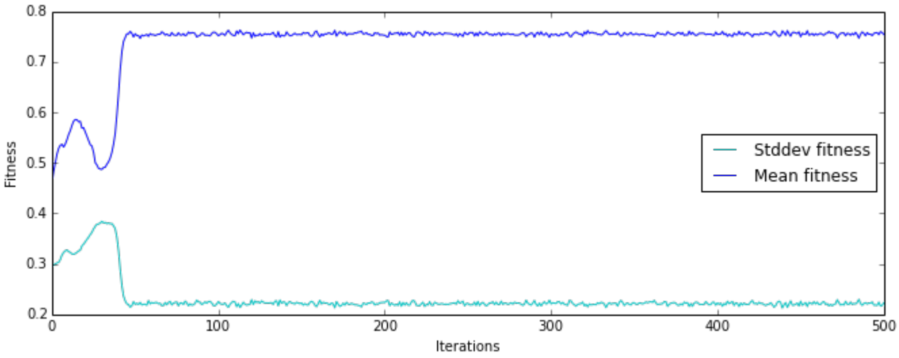

FIGURE 4.8: Mean population fitness, optimised for 500 generations with a genetic algorithm. With a fitness criteria of $\frac{exp(AUC_{in}+AUC_{out})}{6}$ & $Spec_{out}$.

Whilst this level of discrimination could be clinically useful; especially when in combination with ECG, such that MCG is used for detection of NSTEMI, a higher level of discrimination would be ideal. Increasing the number of parameters in the classifier could improve the prediction accuracy. However, if the number of parameters in the model is increased, the number of parameter permutations increases. From $n = 43$ total parameters, the number of possible combinations of $k$ parameters is $\frac{n}{k} = \frac{n!}{k!(n-k)!}$. For 10 features there are 1917334783 possibilities. It is infeasible to test all of these combinations, due to the inefficiency of the logistic regression optimisation process.

### Genetic Optimisation of Parameter Set

I used a genetic algorithm to find the optimal set of 10 parameters (Using the DEAP evolutionary toolkit for python). The algorithm uses an evolutionary approach in an iterative fashion, to improve the set of parameters according to a fitness criterion. Initially a population of 4200 individuals is randomly generated; each individual represents a 10 parameter set. Each iteration the individuals are tested:

- Train a classifier on the training dataset using the parameters of the individual.
- Apply individual model to validation dataset.
- Evaluate performance according to the fitness criteria.

At the end of each iteration the best individuals are selected from the population, Cross breed using a two point crossover and mutated to form the next generation.

The first fitness criterion I used was to simultaneously minimise the Sensitivity and Specificity of the validation set ($Sens_{out}$ & $Spec_{out}$). The optimal set of parameters it found was: $RS_{length}, R - T_{mean}, T4_{corr}, T4cursor, T_{meancorrvar}, T2_{angle}, T3cursor, T_{stdev}, T1_{length}$.

$$\begin{array}{cc} Sens_{in} = 0.7894 & Spec_{in} = 0.95 \\ Sens_{out} = 1.0 & Spec_{out} = 0.4490 \\ AUC_{in} : 0.8184 & AUC_{out} : 0.5978 \end{array}$$

However, the algorithm struggled to optimise both simultaneously.

The best approach I developed included AUROC in the fitness function. A two parameter fitness function was used. First parameter was $\frac{exp(AUC_{in}+AUC_{out})}{6}$, this acts to maximise both in and out of sample sensitivity and specificity. The second parameter was out of sample specificity, $Spec_{out}$.

The resulting optimal individual was $RS_{angle}$, $R_{p2p}$, $QR_{length}$, $T2cursor$, $RS_{length}$, $QRcursor$, $T4_{angle}$, $R - T_{mean}$, $T4_{length}$, which yielded:

$$\begin{array}{cc} Sens_{in} = 0.947 & Spec_{in} = 0.95 \\ Sens_{out} = 0.29 & Spec_{out} = 0.79 \\ AUC_{in} = 0.992 & AUC_{out} = 0.687 \end{array}$$

The selection pressure for out of sample sensitivity was not high enough. The $Spec_{out}$ maximisation was too high. Instead this should be swapped for $Sens_{out}$. I preliminarily tried this parameter previously, but the GA found a local minima in with $Sens_{out} = 1.0$, $Spec + out = 0.0$. Seeing the above plot of average population fitness over time, it seems that at least 40 generations (with a population of 4200 individuals) are needed to find the global maxima.

The addition of extra parameters, did not improve the prediction accuracy compared with the 4 parameter results obtained by Paul and Abbas.

### 4.3.2 Machine Learned Classifiers

The machine learning approach uses algorithms that can learn from data without explicit programming. The field began in the 1950's [88, 89]. But grew slowly due to the lack of computing power available. Recent decades have seen an acceleration in the development and performance of machine learning techniques. In 2016 many eyes watched a milestone in the development of artificial intelligence as DeepMind's AlphaGo, a deep neural network based system, achieved super human performance by defeating world champion Go player Lee Sedol, [90].

In general their are two types of machine learning, supervised & unsupervised learning. In supervised learning the algorithm is provided with labelled input data indicating its class, and is trained (generally some form of gradient descent) to correctly assign a class to input data. A review of supervised learning techniques is found in [91]. An unsupervised learning system is provided with unlabelled data and creates its own labels to discover a descriptive function. It would be interesting to apply an unsupervised learning algorithm such as T-SNE, as it may reveal multiple classes in the dataset, that correlate with our knowledge of the patients.

Machine learned classifiers, including BNN, BPNN, PNN & SVM, have been applied to MCG before [92, 93]. And have been shown to out perform Cardiologists in MCG diagnosis [33]. However, machine learned classifiers are rarely used in clinical practice. Medical devices that contain ML algorithms usually come with provisos that the information is not to be trusted and must be independently verified by an expert. This is because of the lack of clarity of the inner workings of these algorithms, more precisely the lack of uncertainty or reliability of a single classification [94]. Comparatively, human devised diagnostic parameters (statistical models) have robust associated uncertainty values that allow the diagnostician to calculate risk. Despite this limitation, I use them as a tool for determining the diagnostic information within the MCG dataset we have gathered.

#### Support Vector Classifier

The first approach I tried was a classifier based on a support vector machine (SVM). Support vector machines were originally conceived in 1963 by V. Vapnik, and upgraded in 1992 using the kernel trick to allow non-linear classification [95]. Since this technique is able to create arbitrarily complex decision boundaries, it was logical to see if it could improve the diagnostic performance compared to LR with its limited decision boundary structure. The work in this subsection was the result of a collaboration with my student Dan Wilson, he wrote up this work in his Masters Thesis, which was submitted to the University of Leeds [96]. We used the scikit-learn library for python which contains a support vector classifier (sklearn.svm.SVC) [97].

A support vector classifier (SVC) finds a hyperplane decision boundary that is able to divide the input vector space by class. The distance of a sample from the decision boundary is a measure of its classification uncertainty. We chose to use the gaussian radial basis function (RBF) kernel, in order to create a non-linear decision boundary.

The hyperplane is defined by support vectors. Training is performed to discover the optimal support vectors, based on the maximum separability between classes and a pair of hyper-parameters $C$ and $\gamma$, which determine the shape of the decision boundary. $C$ is a cost function that determines the penalty for misclassification of a sample. Reducing C allows some misclassification to a smoother fit that is likely to produce a more general model. $\gamma$ determines the broadness of the RBF kernel gaussian; effectively determining the curvature of the decision boundary. With a high curvature and a high misclassification penalty it is easy to produce a SVC that has 100% accuracy on the training data. Hence it is essential to optimise the hyperparameters (HP) for validation accuracy, not training accuracy.

The simplest method for hyperparameter (HP) optimisation is an exhaustive grid search. However, this method is time consuming and is limited to a pre-defined grid step resolution. We wrote a simulated-annealing monte-carlo markov-chain (SA-MCMC) search algorithm to find the optimal HP. The functioning of a SA-MCMC algorithm is analogous to a particle bouncing around a landscape, getting stuck in valleys. Initially the particle has a high temperature and is able to move out of small valleys without getting stuck, over time the temperature drops (simulated annealing) and the particle ends up trapped in the (ideally) global maximum. Each iteration in the algorithm is a new position in HP space and is probabilistically chosen and evaluated by training a SVC with those HP and noting the resulting prediction accuracy. If the result is an improvement over the last position the position will be selected, else that move will be rejected. This is a gradient descent function, that will move to a local maximum. In order to find the global maxima the algorithm must be able to leave local maxima. To achieve this, a global variable, the temperature, determines the ability of the particle to move against the gradient. When the temperature is high, moves that decrease the accuracy are allowed. As the temperature is slowly lowered, and the final position should be the global maximum.

The SA-MCMC temperature and step size parameters must be tuned. The temperature must have a start, end and an evolution curve (linear in this case). The step parameter is dynamic based on the search path gradient, but an initial, minimum and maximum value must be defined. This hyper-HP tuning is not required for grid search. It was quite a labour intensive process, though it would not be impossible to automate.

To evaluate the SA-MCMC approach compared to grid search, we performed an exhaustive grid search of $64 = 4096$ positions, and a 4000 iteration SA-MCMC in the same space. The SA-MCMC spent about 600 more iterations in the high performance area than the grid search. This ability to focus and fine tune, without changing the grid resolution, is what makes SA-MCMC superior.

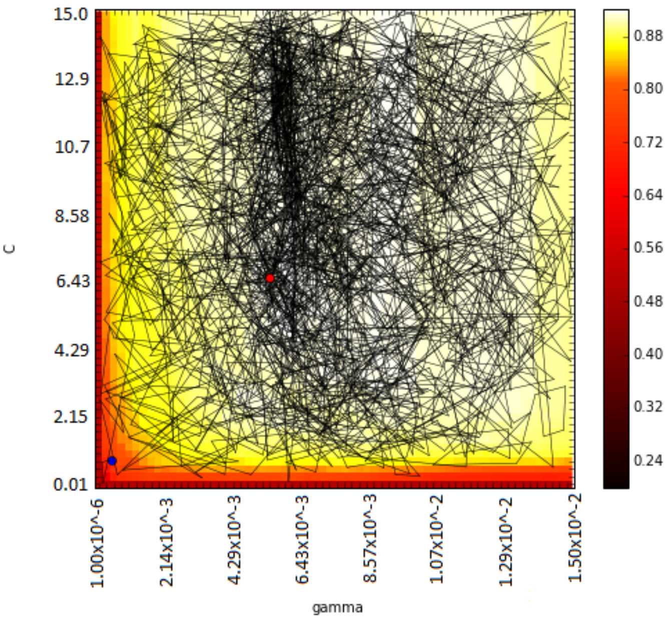

FIGURE 4.9: SA-MCMC search algorithm. Search path shown in black, with start and end indicated by the blue and red circles. The background heatmap is produced by an exhaustive grid search ($64 \times 64 = 4096$), of approximately equal number of search points as the SA-MCMC (4000).

| Folds | Accuracy(%) | Sensitivity(%) | Specificity(%) | Mean AUROC |
|---|---|---|---|---|
| 20 | 73.8 | 94.1 | 53.1 | 0.92 |
| 4(1) | 79.3 | 86.7 | 71.2 | 0.93 |
| 4(2) | 76.5 | 87.7 | 63.1 | 0.94 |
| 4(3) | 73.8 | 94.1 | 53.1 | 0.93 |

TABLE 4.1: Table of SVM performance.

### SVC Results

When training the support vector classifier we found optimal performance by sub-sampling the time series data every 40 samples; creating a time series of $50 \times 15 = 750$. This was expected because the low frequency information is attenuated by the low-pass filter (40 sample moving average). Figure 4.10 shows the mean classification accuracy with number of training iterations and varying levels of sub-sampling.

Because of the HP optimisation stage it is essential to perform a nested cross validation. The inner loop for training the HP and the outer loop for validation. We performed several combinations, all had 4 fold cross validation in the internal loop. The outer loops varied, we tried one 20 fold loop to evaluate smaller training sets, and three runs of 4 fold outer loops, see table 4.1. Each 4 fold outer loop employed different HP search parameters (an increase in the step size of SA-MCMC).

The 20-fold performed as well as the 4-fold on average, however the variance was much higher due to the smaller training set. This variance is visible in the comparison of the ROC results, figures 4.11 and 4.12. The 1st 4-fold run produced the 4 SVC with the highest mean classification accuracy.

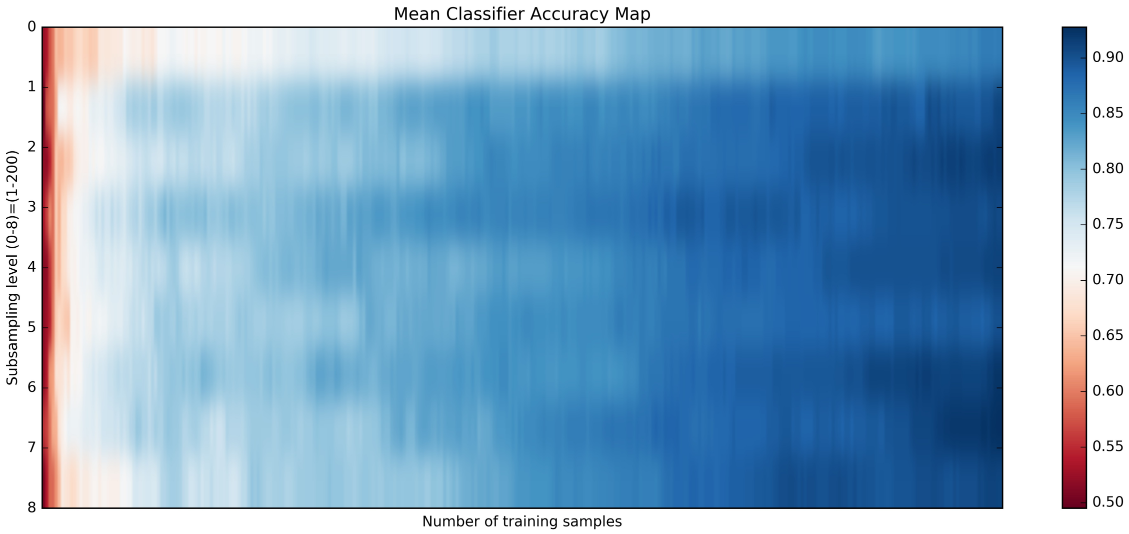

FIGURE 4.10: SVC prediction accuracy as a function of training steps (x) and the degree of sub-sampling of the input data.

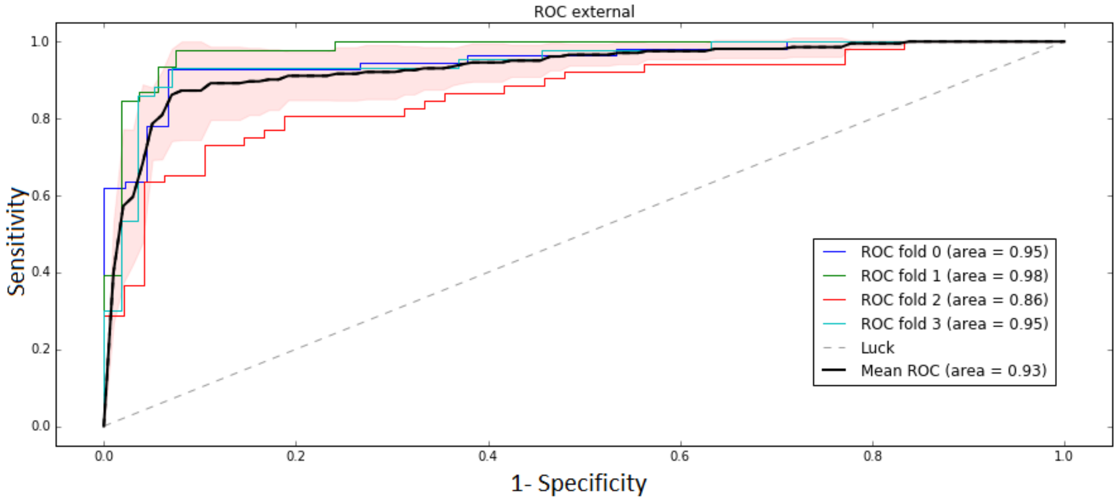

FIGURE 4.11: ROC curve for the 1st external 4-fold loop classifier.

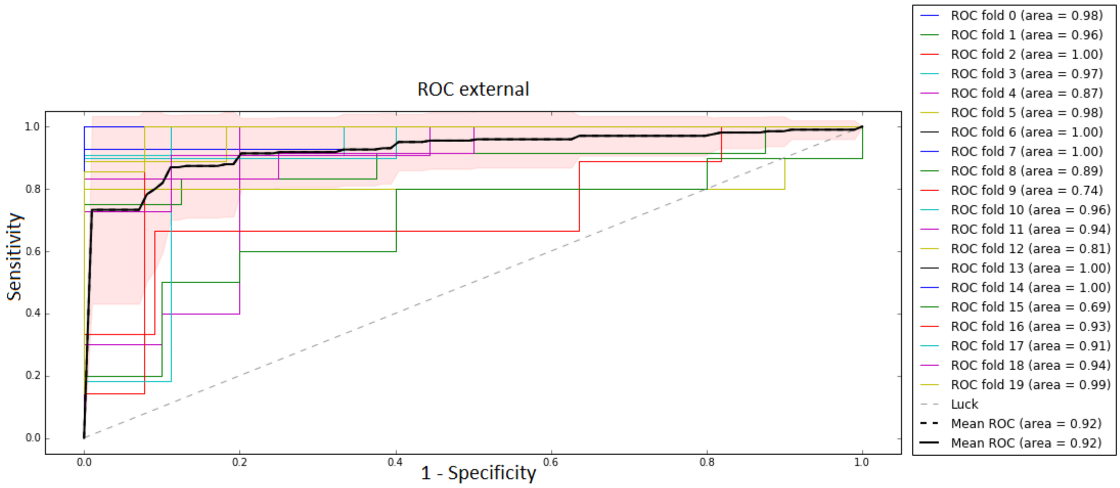

FIGURE 4.12: ROC curve for the external 20-fold loop classifier.

### Convolutional Neural Network

Convolutional neural networks (CNN) are a fairly recent development, originating with LeNet5 in 1998, which was developed for image classification [98]. The key advances since that time are due to the development of backpropagation training using GPU's and the invention of rectified linear units; ReLU neurons that were more computationally efficient to train compared to the tanh and logit functions previously used, but provided the same non-linearity. CNN are not limited to image classification tasks, they can be applied to arbitrary dimensional problems. I have applied CNN to two MCG representations, a single 2D image of time versus sensor and a 3D MFM movie. These approaches resulted in nearly identical performance, however the 3D (interpolated) MFM movie format is device agnostic, allowing the same CNN to diagnose various detector array arrangements. The work in this subsection was the result of a collaboration with my student Mike Smith, who also wrote it up in his Masters Thesis, which was submitted to the University of Leeds [99]. We used the TFLearn library, which is a high level wrapping of tensorflow, suited for easier creation of CNN. I chose to use tflearn and tensorflow because they are well integrated with python, especially interactive jupyter notebooks and were therefore easy to integrate with the rest of the MCG code.

A classic neural network consists of multiple layers of neuron units. Each neuron is connected to every neuron in the layer before it, and produces a single output. This output is produced by an activation function that takes a weighted sum of the input connections as its input. The activation function is a leaky-ReLU, which is small and constant for negative values of the weighted sum, and linearly increases with positive values.

A multiple layered neural network is known to be a universal approximator, able to represent an arbitrary continuous function; able to learn any diagnostic pattern. However, it is difficult to train a classic neural network due to Hughes' phenomenon, which states that as the dimensionality of a neural network increases, its predictive power decreases. This curse of dimensionality can be diminished by having a convolutional front end to the neural network.

The convolutional layers function similarly to the first stages of the human vision system; raw pixel distributions are converted to maps of higher level features such as edges,

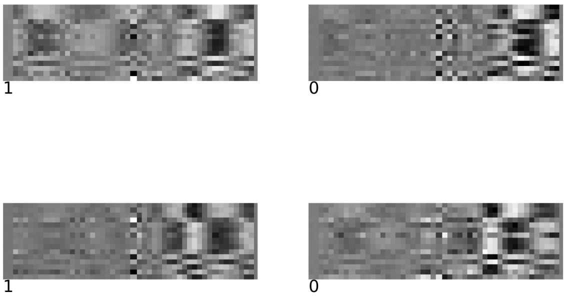

FIGURE 4.13: Labelled example data used for the 1D and 2D CNN. The data shape is $2000 \times 15 = 30000$pixels. The QRS complex can be seen in the centre, and the T-wave (ballistocardiographic) signals to the right.

| CNN | Accuracy(%) | Sensitivity(%) | Specificity(%) | Mean AUROC |
|---|---|---|---|---|
| 1D | $81 \pm 3$ | $77 \pm 2$ | $84 \pm 3$ | $0.88 \pm 0.01$ |
| 2D | $88 \pm 3$ | $87 \pm 3$ | $89 \pm 3$ | $0.93 \pm 0.02$ |
| 3D | $88 \pm 3$ | $86 \pm 3$ | $90 \pm 3$ | $0.94 \pm 0.02$ |

TABLE 4.2: Table of CNN performance.

which are abstracted further by the following layers into higher complexity features such as faces. The CNN convolution layer consists of several convolutional filters, each filter is a feature mask whose weights are learned along with the neuron weights during SGD. Each filter is applied at multiple positions in the input, without overlapping the input edges. The output of the convolutional layer is a set of feature maps. A pooling layer follows each convolution layer. The pooling layer compresses the feature map by sub-sampling, choosing the highest activation within a subset of the map.

Because the feature masks are learnt, there is no requirement to perform manual feature engineering. The downside is that the learnt masks encode patterns that are not necessarily human perceivable, in our experience the masks looked like noise. However, it is common in image classification tasks to see structured masks develop, producing recognisable features. These computationally evolved features can inform our understanding of biologically evolved systems, such as the retinal layers in the human eye.

### CNN Results

We trained three CNN with differing convolution layer architectures; 1D, 2D and 3D convolution masks. The 1D and 2D CNN were trained on the same data that was used for the SVC, but without sub-sampling $2000 \times 15 = 30000$, as can be seen in figure 4.13. The 3D CNN was trained on interpolated MFM time series, $19 \times 17 \times 2000 = 642000$. The convolutional masks used for the 1D were $1 \times 5$, the 2D used $15 \times 5$ and the 3D used $5 \times 5 \times 5$. A 10 fold cross-validation was used to ensure validation on unseen data. The 1D net lacked the relative spatial information that was available to the 2D and 3D nets, this accounts for the disparity in classification accuracy, seen in table 4.2. The extra spatial information provided by interpolation did not improve the classifier accuracy much, only a 1% improvement in AUROC.

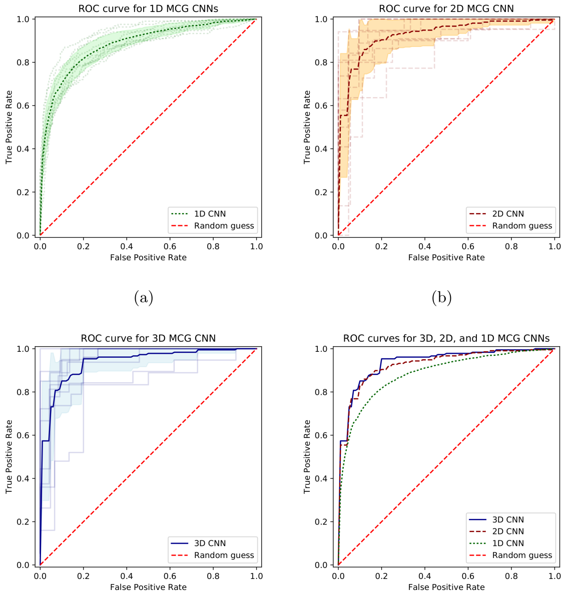

(a) (b)

(c) (d)

FIGURE 4.14: Mean ROC curve of a 10-fold cross validated CNN trained on the 6060 dataset. The light grey/blue curves show the individual validation runs. The $AUC = 88\%$

## 4.3.3 Mapping Diagnostic Information

I devised a method of mapping diagnostic information (dMap) by training CNN on subsets of the data and measuring the classification performance of each subset. The compute time required for this was massive, though each individual CNN was faster to train, 100 CNN were required for the MCG diagnostic map (dMap(MCG)), see figure 4.15. To improve compute time I built a system with four 1080ti GPU's that were watercooled to provide stable high power operation. With this system I was able to perform a time-frequency mapping, which required approximately 30× more compute time. With the new system I was able to complete a map in one week, see figures 4.16,4.17.

I did not have time to apply full spectroscopic mapping to the MCG dataset due to the computation time required and reaching the end of the PhD. However, I performed a non-spectroscopic diagnostic map to the MCG clinical trial data and found the region containing the ballistocardiograhic information to be highly diagnostic.

### dMap(MCG)

Moving through the cardiac cycle, the dMap(MCG) shows some diagnostic information is present in the P-wave region, with a peak of 0.8 AUC. This drops in the P-Q region to $< 0.6$, before rising in the QRS complex to $\approx 0.8$ AUC. It continues increasing into the repolarisation region peaking at 0.9 AUC. The high diagnostic performance in the repolarisation region indicates that the ballisocardiographic signals contain diagnostic information.

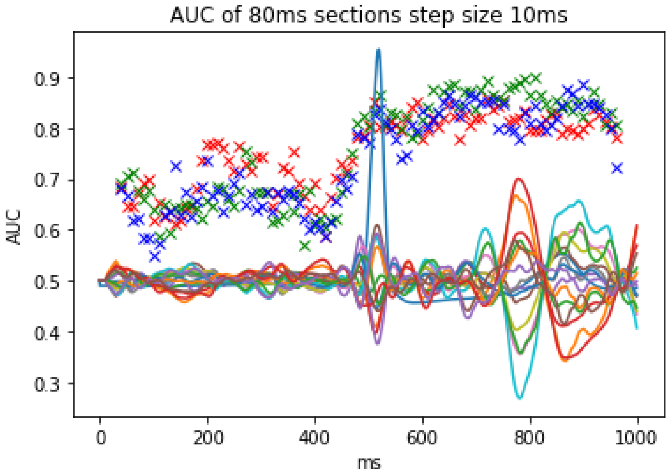

FIGURE 4.15: Diagnostic map of 6060 MCG dataset. Each cross shows the AUC of a 3D CNN trained only on a window of data ±40$ms$ surrounding the cross. A 3 fold cross validation was performed (R,G,B crosses).

### Time-Frequency dMap(ECG)

Using the PTB ECG database (20000 scans of Frank lead 'Vx') I performed a spectroscopic map of the ECG diagnostic information. The spectrogram was created by using multiple window widths, from 8 to 120ms. The spectrum of window sizes reveal the true width of diagnostic features. The results corroborate with contemporary understanding of ECG diagnostics, showing the highest diagnostic value in the S-T region, with some information in the R wave and P wave, see figure 4.16.

In order to investigate the effect of the derivative on diagnostic information. I repeated the diagnostic mapping on the derivative of each ECG (dECG), 4.17. The result showed no change in the overall diagnostic performance, but at a smaller scale the information content changed. In order to compare the two I subtracted the dMap(dECG) from the dMap(ECG), see figure 4.18. The result showed a loss of AUC from the S-T interval, with a corresponding gain in AUC in the S-J region. This shift in diagnostic information from S-T to S-J indicates that the diagnostic performance of the S-T interval (used to detect IHD) is not lost in the derivative, but instead shifts to the preceding gradient.

## 4.4 Conclusion

The diagnostic performance that was achieved using a range of classifier techniques confirms that the MCG device captured diagnostic information during the 6060 clinical trial, despite the appearance of high noise, and the presence of ballistocardiographic signals. In fact these ballistocardiographic signals contained diagnostic information, as was observed by diagnostic mapping using CNN.

The LR classifier created by Paul and Abbas used 4 parameters, it performance was better than the 10 parameter LR that I genetically optimised. This could be a manifestation of Hughes' phenomenon, or simply a poor optimisation by the genetic algorithm.

The machine learning techniques (SVC and CNN) outperformed the statistical machine learning approach (LR), see table 4.3. The accuracy of the SVC, 79.3% outperformed

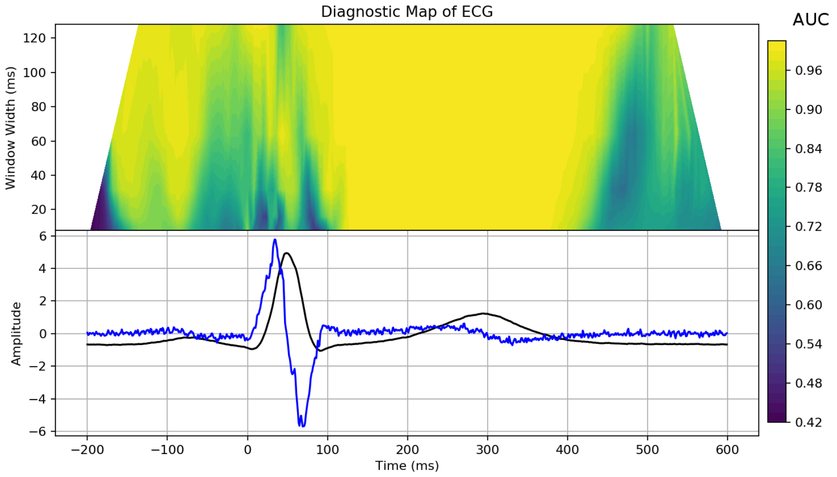

FIGURE 4.16: The lower plot shows the ECG (black) and its derivative (blue). The above spectral map above shows AUC diagnostic performance of a CNN trained on subsets of the ECG data, the width of each subset is varied over the y-axis. The T region is highly diagnostic, even when taking a small 10ms subset.

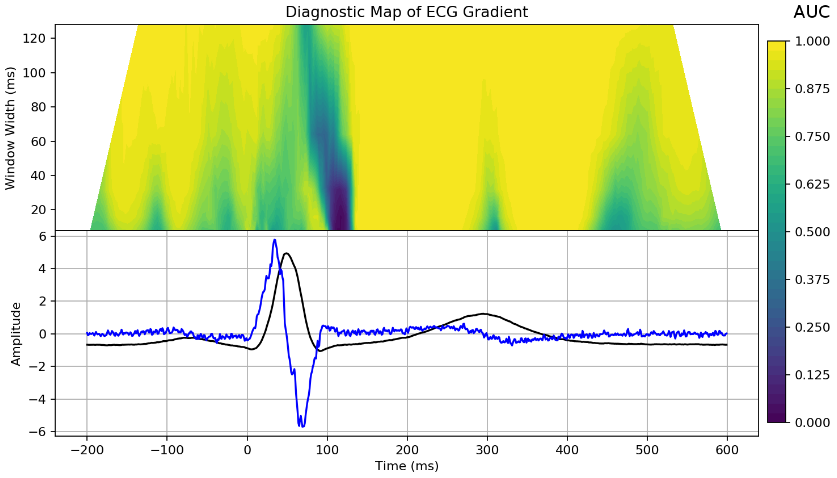

FIGURE 4.17: The lower plot shows the ECG (black) and its derivative (blue). The above spectral map above shows AUC diagnostic performance of a CNN trained on subsets of the ECG-derivative data, the width of each subset is varied over the y-axis. The early S-T region has no diagnostic information, since the gradient is approximately zero.

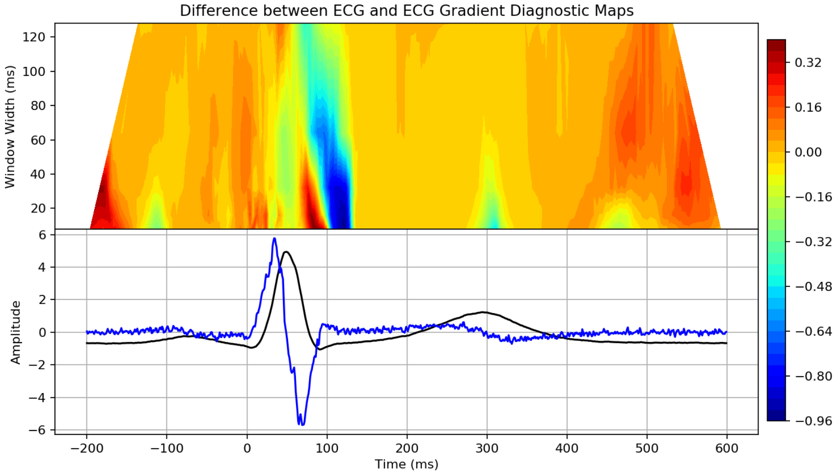

FIGURE 4.18: Spectral map showing the difference in diagnostic information content between the ECG gradient and the ECG. Positive valued regions correspond to superior ECG gradient diagnostic performance.

| Classifier | AUROC |
|---|---|
| LR (Genetic) | 0.69 |
| LR (Paul Abbas) | 0.71 |
| SVC | 0.93 |
| CNN | 0.94 |

TABLE 4.3: Table comparing all classifier performance.

the linear SVC used by Kangwanariyakul et al in 2010 [93], but was bested by their back-propagation neural network (BPNN), Bayesian neural network (BNN), and probabalistic neural network (PNN) that had accuracies of 78.4%, 78.4%, and 70.6% respectively. The CNN had the highest classification accuracy of all the approaches at 88%, this is higher than any published result of machine learned MCG classification. The highest published result was by Tantimongcolwat et al who used a direct kernel self-organising map (DK-SOM) to achieve an accuracy of 80.4% [92]. Fenici et al obtained a similar accuracy of 80% using a direct kernel partial least squares (DK-PLS) method. This achievement highlights the performance of CNN.

The history of diagnostics can be seen as a huge, multi-decade exercise in pattern recognition, with the work carried out by human scientists exploring and validating patterns with clinical trials. With the advent of machine learning, we can now construct algorithms that can recognise arbitrary patterns in datasets, provided sufficient labelled training data and computation is available. The efficacy of these techniques is startling. However, they suffer from some drawbacks, it is difficult to know how and why an algorithm arrived at an answer, and they are vulnerable to incorrect classification resulting from adversarial examples with small input perturbations. From the point of view of a diagnostician, knowing the uncertainty of a predicition is vital to their decision making process. Though most ML methods have a probabilistic output, the probability does not reliably represent the uncertainty. Some approaches have been proposed to reliably extract uncertainty, such as the drop-out based Bayesian method [100]. In addition to this, the networks must be made robust to adversarial examples, using for example differential privacy[101]. Adversarial examples exist due to the limited training experience of neural networks. Their differentiability allows adversarial examples to be found via a gradient descent procedure.

It is difficult to understand what features were used by the CNN to achieve classification. The learnt convolutional features were analysed but they appeared to be random. Similarly the weights of the fully connected layers appear random to human observation. It is common in image classification with CNN to see learnt feature masks that look similar to those developed by humans for convolutional filtering of images (such as edge detection filters).

The diagnostic mapping technique I developed is able to assess which regions of the input data contain diagnostic information. This is useful for validating the algorithm to ensure it is not learning on some common noise features. Furthermore this information can be used to inform future device improvements. For example, capturing the derivative S-J component is likely to be useful for detection of IHD based on the dECG results. This is exciting news since ICM MCG measures the derivative and cannot observe the static magnetic field in the S-T segment. With integration it might be possible, but in my experience its too noisy, also our device filters out frequencies below 2Hz so there may not be sufficient information to compute an accurate integral.

When designing the DSP routine we were expecting to observe signals similar to those acquired by DC SQUID. However, the information in a signal is represented by higher frequency components in the derivative. The pass band of our DSP is about $2 - 40$Hz, some higher frequency components are allowed through but generally attenuated by the low pass filter, as seen in figure 3.15. Arguably we should try to extract more of the higher frequency components, since there are a number of cardiac abnormalities that are detectable by 'fragmentation' high frequency components in the depolarisation phase. For example, atrial fibrillation can be detected in the P-wave with fragmented components of $200 - 300$Hz. It would be interesting to use an alternative low pass filter, or only apply it once, then use a second low pass filter with a higher cut-off. This would allow a greater proportion of high frequency components into the dMCG. It would have a noisy appearance but is potentially a rich diagnostic resource.

## Chapter 5

# Discussion

The induction coil magnetometers have proven to be suitable sensors for magnetocardiography, converting magnetic flux into voltage with sufficient sensitivity and a robustness to high amplitude noise perturbations. Measurement of the ICM sensitivity (using the recordings taken in the MSR at YNIC) was an important step in the project since it enabled scientific communication of the their performance, and thus comparison with alternative sensing technologies. Their sensitivity was sufficient for MCG in a clinical environment with an inherent noise of $104\text{fT}/\sqrt{\text{Hz}}$@10Hz. This was quite unexpected by the field of magnetocardiography as the last ICM developed for MCG; back in 1982, required a large (12cm diameter) sensing coil to achieve a similar sensitivity [37]. Essential to our ICM performance is the improvement in low noise amplification and digitisation electronics that has occurred in the intervening decades. The pace of electronics development is high, since the beginning of this research project new amplifier chips have become available that have even lower noise performance, which could enable smaller ICM and therefore higher resolution ICM MCG.

Whilst the ICM had sufficiently sensitivity, removing the dominant environmental noise encountered during unshielded operation was a significant challenge, which has been somewhat successful. The digital signal processing routine we developed was able to remove enough environmental noise to routinely observe cardiac depolarisation signals in an unshielded clinical environment. The key element of our DSP routine that allowed unshielded MCG was the coherent noise rejection algorithm. This algorithm is sensor agnostic but performs best when the sensors have well matched gain; much like gradiometry but CNR offers a reduced stringency of matching and a higher sensitivity. The performance of the CNR algorithm is dependent on the environmental noise possessing a sufficiently high spatial coherence width.

Using HBT correlations I measured the spatial coherence width encountered during the 6060 trial and compared it with the CNR attenuation and found a linear relationship of 9dB attenuation per order of magnitude increase in spatial coherence. By extrapolating this curve there is a potential for CNR to rival the spatial filtering performance of magnetically shielded rooms. However, when comparing the spatial coherence measured by the MCG device with a longer baseline experiment, I observed a decoherence effect due to the ICM proximity with their neighbours. This is likely due to a combination of inductive coupling, photon scattering and amplifier noise. If this decoherence effect could be reduced to the level of the teststand experiment it would obtain a 15× increase in coherence, resulting in a 10dB attenuation of environmental noise. Theoretically an array of optical magnetometers could achieve the lowest decoherence effect. Though typically optical magnetometers employ actively driven field coils, which will create a decoherence effect; A fully optical magnetometer would be ideal.

ICM measure the time derivative of the magnetic field, dMCG and as such are most sensitive to higher frequencies (response increases linearly with frequency). However, in much of our analysis the higher frequency components were filtered out (above 40 Hz) to create smooth butterfly plots that had similar appearance to SQUID MCGs and to concentrate on the fundamental frequencies of the cardiac cycle. This was potentially a mistake since it is possible to detect cardiac abnormalities by the presence of high frequency 'fragmentation' components. A desirable improvement to the DSP routine would be to capture the higher frequencies whilst still rejecting mains grid noise.

Ballistocardiography is one of the earliest cardiac diagnostic techniques, the technology to acquire a BCG has existed for over a century, however interpretation of the resulting signals has been difficult. The measurement of ballistocardiographic signals transduced by the bed frame was unexpected and initially perceived as an unwanted noise source. However, the diagnostic mapping technique revealed the presence of diagnostic information within the BCG signals, which could be used by the CNN to improve prediction accuracy. A hybrid MCG/BCG device may be ideal, with the MCG detecting in the depolarisation phase and the BCG detecting in the repolarisation phase. This is likely to require a consistent bed frame and orientation to the Earth's magnetic field for consistency of BCG signals, though a permanent magnet affixed to the bed could substitute the Earth's field.

The repolarisation signals were masked by the ballistocardiographic noise from the metal frame of the medical bed. Scans performed on a wooden bed could sometimes observe T-wave but with a reduced SNR, due to the reduced sensitivity to lower frequencies that is somewhat inherent to induction coils (unless operated as a flux gate magnetometer), but exacerbated by our use of a DC offset loop circuit, which acted as a 1.6Hz high pass filter.

It is important for the MCG to perform well with the use of statistical learning classifiers, since these are well adopted in medicine and have established methods for measuring the uncertainty of a given prediction. The performance of such methods is dependent on the set of parameters extracted from the data. One critical parameter is the MFM dipole angle. Measuring the dipole angle using the correlation method resulted in a higher accuracy than the original pole peak to peak method we employed. The CorrTS time series produced were also useful for extraction of standard temporal features since it enabled a more precise location of the zero crossings.

The localisation of amplitude peaks within the MCG was improved by applying a peak finding algorithm in the wavelet domain. The resulting peaks were used to automatically assign MCG temporal features to a standard template that Shima and I developed. I also wrote a analysis program with a user interface that allowed the feature marker positions to be adjusted, when the automatic system made an error. Using this methodology we extracted 43 parameters from each of the 480 scans in the 6060 trial. Individually these parameters held little to no diagnostic value, but combining several together revealed some linearly separable structures that were further polarised with the YHC data. The use of statistical machine learning provided a further improvement in diagnostic performance. Interestingly a Logistic Regression classifier using 4 parameters outperformed a LRC that used 10 parameters. This is either a manifestation of Hughes' phenomenon or poor performance of the genetic optimisation algorithm, which was used to choose the best combination of 10 parameters from 43.

A great leap in diagnostic performance was achieved using machine learning methods. The 79.3% achieved by the SVC is comparable to other published results, and provides a good justification for the diagnostic information capturing capability of the device. Remarkably the convolutional neural network achieved 88% classification accuracy on unseen data, better than any published accuracy achieved in MCG. It remains to be seen whether other groups will observe similar benefits from CNN. Clinical application of machine learning diagnostic will require creating a robust measure of prediction uncertainty and a level of transparency for the clinical practitioners using the classifiers. Some transparency can be gained by using the diagnostic mapping technique, it shows for a given dataset, which areas of the signal are used to reach a prediction.

The technique I developed using CNN to map diagnostic information content was applied to map the ECG and dECG. It revealed that the diagnostic information contained in the plateau of the S-T interval is shifted backwards into the preceding gradient (S-J). This exciting evidence suggests that IHD can be detected in the depolarisation phase and escalates the inducement to develop the analysis of higher frequency components.

## 5.1 Future Research & Development

A simple but important improvement to the device is to provide feedback to the operator on the environmental noise and the progression in scan quality. This would allow the operator to proactively address extreme noise environments, by removing sources of noise or moving the device to a quieter environment. Certain types of noise such as the scan subject impacting the array during measurement, have a characteristic structure that could be automatically detected and removed. The spatial coherence could be continuously computed and communicated to the operator via an indicator showing the coherence width relative to the acceptable limit.

Surprisingly persons with highly regular heart beats are difficult to scan unshielded. Since their heart beat is co-periodic with periodic environmental noise, the cycle averaging performance decreases. It is simple to calculate heart rate variability, therefore the operator could be made aware of the reason for reduced noise rejection performance. Strategies to increase heart rate variability could be performed. The least ethically problematic of these could be to use erratic music to disrupt the hearts rhythm, though this might not work as well as heart rate synchronisation to regular tempo music.

An extremely low HRV can be symptomatic of diseases in conduction pathways or pacemaker node disease. Therefore it is useful for the clinical decision maker to have access to HRV information in addition to the MCG. This suggests that the MCG analysis system should be a hybrid ECG and MCG. Since the ECG is already required by the system for cardiac cycle averaging and it provides a host of diagnostic information that may be complimentary to MCG, especially if using machine learning for classification.

### 5.1.1 Sensors

MCG devices that have up to 64 channels within a similar area to our device are commercially available, they are able to observe smaller MFM features and therefore achieve higher localisation accuracy. A smaller ICM sensor would enable an increased spatial resolution, at the cost of reduced SNR to magnetic field. However the reduced SNR could be offset by spatial averaging, by modifying the extent of the spatial averaging a dynamic effective spatial resolution can be achieved, which can be adjusted in response to changing environmental noise conditions or resolution requirements. Furthermore increasing the array sensor count would increase the performance of CNR by reducing the inherent noise in the bucket signal (which is proportional to $(1 + \frac{1}{\sqrt{n}})$.

ICM have an inherent decoherence due to the conductive and permeable material. This could be reduced somewhat by removing the soft iron core, at the expense of response, but this effect is limited. An optical magnetometer array may achieve lower decoherence, especially if it uses completely optical components; locating the electronic elements distant from the measurement region. With OPM this is difficult to achieve since the Alkali vapor must be heated optically (typically heated to 90°). An active coil is usually used to apply a known magnetic field to the vapor, this could be replaced by a spin aligning optical pulse. It would be easier to achieve an all optical magnetometer based on diamond NV centres. However, these magnetometers are not yet capable of the required sensitivity for MCG, though it appears the engineering challenges are addressable.

### 5.1.2 Explore the Limitations of Spatial Coherence and CNR

To exploit the increased CNR noise rejection of 9dB attenuation per order of magnitude increase in coherence width device design changes could be made to reduce decoherence. For example, any conductive material near the detector array should be replaced with an alternative dielectric material, and electronics placed as distant as possible. As an alternative to investing in a MSR, the spatial coherence of measurement environment could be increased by modifying hospital architecture. In order to determine the optimal placement of the device, it would be useful to conduct an extensive study of a hospital before using it for unshielded MCG. The practical limitations to spatial coherence in an urban area are unknown. To better understand the limits a mixture of surveying and simulation will be required.

The coherent noise rejection algorithm that we developed requires a planar array. It would be interesting to compare its performance with a SSS algorithm using a curved array. Furthermore, I observed transient coherent events in a time-frequency coherence analysis of a background recording at LGI. This suggests that an adaptable CNR algorithm could be created to reject extremely incoherent events. Though the adaptive algorithm I created had inferior noise rejection performance than the simple CNR.

### 5.1.3 Digital Signal Processing Improvements

There is not much room for improvement in the basic elements of the digital signal processing; the averaging time can be increased, however this gives diminishing returns $(\sqrt{n})$. The low pass filter is optimally tuned for 50Hz removal, and I am not aware of any filter with comparable performance. It would be useful to obtain some components above 50Hz, but this would require a different low pass filter. The advanced denoising algorithms that I experimented with either did not provide a significant improvement in SNR or caused distortion. However, it would be interesting to see the performance of a SSS type algorithm, which are commonly used for MEG.

#### High Frequency Components

The primary use case for this device is in triage of subjects presenting at the emergency department with chest pain. We know from previous work that MCG is sensitive to a broad range of cardiac abnormalities. Of a particular interest, due to their prevalence, are IHD and Atrial Fibrillation (AF). IHD can be detected in the QRS complex. It is known in ECG that fragmentation in the high frequency components of QRS is indicative of IHD [102, 103]. AF can be detected by observation of fragmentation in the high frequency components of the P-wave in ECG and MCG [104–106].

I have observed some diagnostic performance in the P-wave region using diagnostic mapping. Though the frequency range of the map was not sufficient to confirm whether the diagnostic information was contained in fragmentation components or the broader P-wave morphology. The diagnostic information shift observed between the ECG and dECG is further evidence that IHD can be detected in the depolarisation region; specifically the S-J region. This corroborates the evidence gathered from the clinical trial that dMCG is indeed sensitive to IHD.

In theory dMCG is ideal for observation of these high frequency components. The difficulty arises because of the need to attenuate 50Hz without also attenuating the higher frequency cardiac signals, for which we employ a FIR filter. This filter also attenuates multiples of 50Hz and higher frequencies in general, see figure 3.15. The amplitude of the utility grid could be lowered by using twisted pair cabling, and by enforcing low frequency emission standards below the current limit of 150kHz, down into the range of organscale biomagnetic signals. However, even with these measures it would likely be prevalent. Perhaps an alternative filter can be found.

#### Fidelity Optimisation

The difficulty with using advanced, adaptive denoising algorithms is analogous to the phrase "with great power comes great responsibility". These algorithms deconstruct the noisy signal, select and modify components, then reconstruct a "noiseless" signal. However, the knowledge these systems have of the noise field is inherently limited, therefore they must make 'educated guesses', which can lead to erroneous output that looks convincingly like the expected signal.

Therefore in the training or tuning of denoising algorithms it is important to quantify their performance with a fidelity metric. This can be achieved using a phantom; a known signal that is detected with both high and very low SNR. The output of the denoising algorithm is then compared with the noiseless signal using for example correlation or total least squares difference. This method can be used to train and validate a zero distortion denoising algorithm.

My correlation based algorithm for measuring the dipole angle provided a lower noise method for dipole angle extraction; effectively by harnessing spatial averaging. Further development of correlation algorithms tailored for MCG could be useful for noise suppressing extraction of other parameters and for device agnostic diagnostic algorithms. As an alternative to manual feature engineering, convolutional filters could be learnt by sparse coding.

### 5.1.4 Clinical Application of CNN & dMap

In this thesis I have been concerned with maximising the classifier accuracy. However, for clinical application of machine learned classifiers, it is not sufficient to achieve high accuracy. A measure of prediction uncertainty is also required in order to assess the risk when decision making. This uncertainty can be quantified by the conditional entropy between the true disease state and the classifier output [107, 108]. The challenge lies in measuring the conditional entropy for a high dimensional classifier system like a CNN.

Another issue with CNN, that is known from their application to image classification, is that imperceptible adversarial input changes can be found that can wildly change the output class. Such adversarial examples can be searched for using a gradient descent algorithm. A mammalian mind would not be so easily fooled, but such minds are much deeper than computational neural networks, and have millions of years of evolutionary polishing. There are techniques to harden CNN to these vulnerabilities, such as differential privacy[101].

The next steps for CNN development involve mapping the conditional entropy and ’smoothing’ it to remove the influence of adversarial inputs, and provide a reliable risk along with the classification.

### 5.1.5 Clinical Ballistocardiography?

The BCG signals occur only in the repolarisation phase, therefore it is possible to simply separate MCG and BCG temporally. However, it may be desirable to obtain the underlying MCG repolarisation. A technique such as frequency filtering could delete the BCG signal but would not preserve the underlying repolarisation phase MCG information. It might be possible to acquire an inertial measurement of the bed, that would be transformed using knowledge of the local Earth’s magnetic field to produce an expected magnetic BCG signal, which could then be subtracted from the MCG. A 6 degree of freedom inertial measurement (3-axis gyroscopic and 3-axis acceleration) would be required to fully constrain the movement of the bed. The transform could be a magnetodynamic model of the bed and detector array or perhaps a machine learned algorithm, which could be trained to work with specific bed and array combinations.

This is a non trivial task that may not be clinically appropriate given the variety of beds, differences in the orientation of the local vector of Earth’s magnetic field, the potential for non uniform external fields, and the difficulty of achieving sufficient SNR with the inertial measurements. However, an added benefit of doing this is the ability to remove the hysteresis noise created from arbitrary movements of the patient, which would improve the overal MCG signal to noise ratio, especially for certain patients (such as those with fluid build-up in the lungs) for whom it is difficult to remain still. Alternatively a permanent magnet could be affixed to a non ferrous (and non-conductive to minimise decoherence) bed, the field from this magnet could be chosen to dominate the Earth’s field and thus provide a BCG that is independent of bed orientation.

### 5.1.6 Stress Test MCG

The purpose of stress testing is to diagnose IHD. By stressing the cardiovascular system, the demand on ischemic regions of the heart is increased, amplifying the effect of narrowed coronary arteries, allowing otherwise silent ischemia to be observed. Typically stress testing is performed with ECG. Several groups have applied stress testing to MCG and found extreme changes (60°) in the dipole angle of the S-T segment [75, 76]. Kanzaki et al found evidence for ischemia in the QRS segment after exercise [109]. Stress test MCG could be useful for measuring cardiac decline. By comparing a persons MCG over time and over stress the mutual information between the heart and MCG is increased simply by increasing the entropy of the heart in a known manner; perhaps quantifying stress by heart rate or ejection fraction, pulse wave characteristics.

### 5.1.7 Computational Cardiac Modelling, Standards for MCG & Personalised medicine

Multi-model multi-scale physics simulations of the heart model the propagation of an action potential through a heart, and compute the resulting magnetic field. Ischemic regions can be simulated, and the resulting changes in magnetocardiogram and cardiac output calculated. Patient specific bio-electro-mechanical models could be used to localise heart disease (ischemic or fibrillatory regions), and predict treatment outcomes [110].

Currently these models are based on 100-200 micron voxels. They simulate the action potential using either a monodomain or bidomain model. However, the minute anisotropic fibrous structure is not modelled, though it is known to have an effect in the bulk [31]. Including microscopic anisotropy, reaching the individual cell scale in future models will be important to improve accuracy. To validate these models experimental data is required to understand the electromagnetic behaviour of the action potential in minute detail. This entails constructing micron scale magnetometer arrays. Scaling down the experiment by Wikswo, which achieved 100 micron pixel mapping of the apex of a rabbit heart [29].

## 5.2 Conclusion

The ability to perform magnetocardiography with a low cost, portable device is important for its mainstream adoption. The previous paradigm shift was caused by the release of multi-channel SQUID devices, it enabled a wave of clinical science to be performed, which discovered the excellent potential for MCG as a diagnostic tool. However, the high cost and fixed nature of SQUID MCG limited their widespread clinical adoption. The development of our induction coil magnetometer based MCG system demonstrates that a low cost, portable MCG system is achievable.

Whilst ICM are robust to high amplitude noise, achieving sufficient signal to noise ratio in such circumstances is difficult, for any sensor. However, by exploiting spatial coherence for noise removal, we have shown that unshielded operation can succeed in hostile hospital environments, replacing gradiometers and magnetically shielded rooms. The extent to which the environment must be modified and regulated to reliably achieve the required spatial coherence width remains to be seen; such engineering of the spatial coherence can be achieved by quantifying the coherence width $(m^2)$ using HBT correlations. The maximum achievable spatial coherence within a clinical environment is unknown, but I have estimated that an order of magnitude increase is possible, given a reduction in the devices own decoherence effect.

The native measurement of the derivative MCG (dMCG) should be embraced as a new category of MCG. It has its own advantages (high frequency sensitivity and therefore high sensitivity to the depolarisation phase of the cardiac cycle) and disadvantages (no sensitivity to absolute magnetic field, reduced sensitivity to low frequency and therefore to the plateau phase in repolarisation). Though the two categories contain equivalent information, transforming between them is inherently lossy due to the frequency filtering effects of integration. Therefore processing the native dMCG signals offers an information advantage.

Manual feature engineering can produce reliable diagnostic parameters that are recognisable by humans. Convolutional algorithms are excellent for extraction of these parameters, as demonstrated by my algorithm for measuring the dipole angle (CorrTS). They are also used by convolutional neural networks, which learn their own convolutional filters in order to abstract information content from the high dimensional input. This technique is extremely effective, when applied to our clinical trial data from an unshielded hospital environment, it achieved an 88% classification accuracy on unseen data. This fantastic result validates the diagnostic information capturing ability of our device, and according to the literature is the highest ever achieved in MCG. It will be exciting to see if CNN can make the transition from research tool to clinical use, the key will be a reliable measure of prediction uncertainty.

There is a great potential for our MCG system to change the clinical practice of chest pain triage, providing a quick, non contact detection of heart disease. The sensitivity of MCG to a broad array of heart disease, especially silent ischaemic disease means it could be relied upon to 'rule-out' the presence of heart disease. Further validation with clinical trials is required for clinical adoption.

# Appendix A

# Appendix

## A.1 Electronics

### A.1.1 Circuit Diagrams

### A.1.2 Thermal Behaviour

## A.2 Array Hardware

## A.3 Analysis User Interface

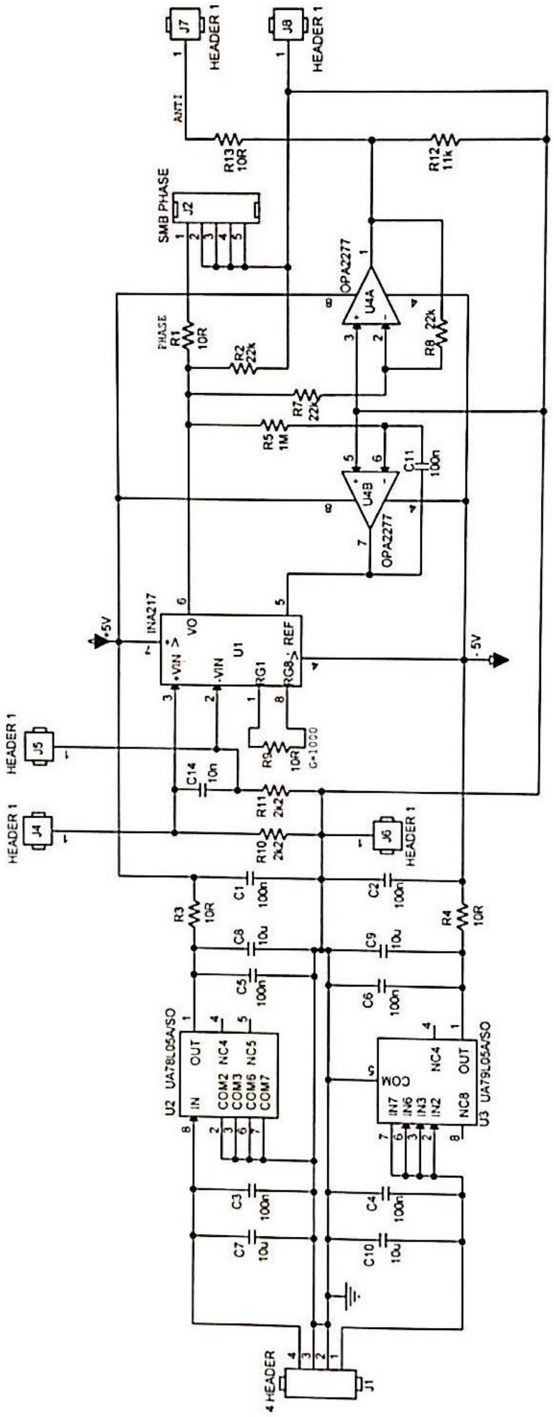
J7
HEADER 1
J8
HEADER 1
ANTI
R13
10R
R12
11k
SMB PHASE
J2
OPA2277
U4A
R8 22k
PHASE
R1
10R
R2
22k
R7
22k
R5
1M
C11
100n
U4B
OPA2277
INA217
U1
+5V
-5V
+VIN
-VIN
VO
RG1
RG8
REF
RG
10R
G=1000
C14
10n
R11
2k2
R10
2k2
J4
J5
J6
HEADER 1
C1
100n
C2
100n
R3
10R
R4
10R
C8
10u
C9
10u
C5
100n
C6
100n
U2 UA78L05A/SO
IN
OUT
COM2
COM3
COM6
COM7
NC4
NC5
U3 UA79L05A/SO
COM
IN7
IN6
IN3
IN2
NC8
NC4
OUT
C3
100n
C4
100n
C7
10u
C10
10u
J1
4 HEADER

FIGURE A.1: Full circuit diagram of the pre-amplifier. Power is supplied into J1. The coil is connected across J4 and J5 and the signal is output through J7 and J8. The SMB PHASE header was included to allow the possibility of balanced signal transmission, though this was never used.

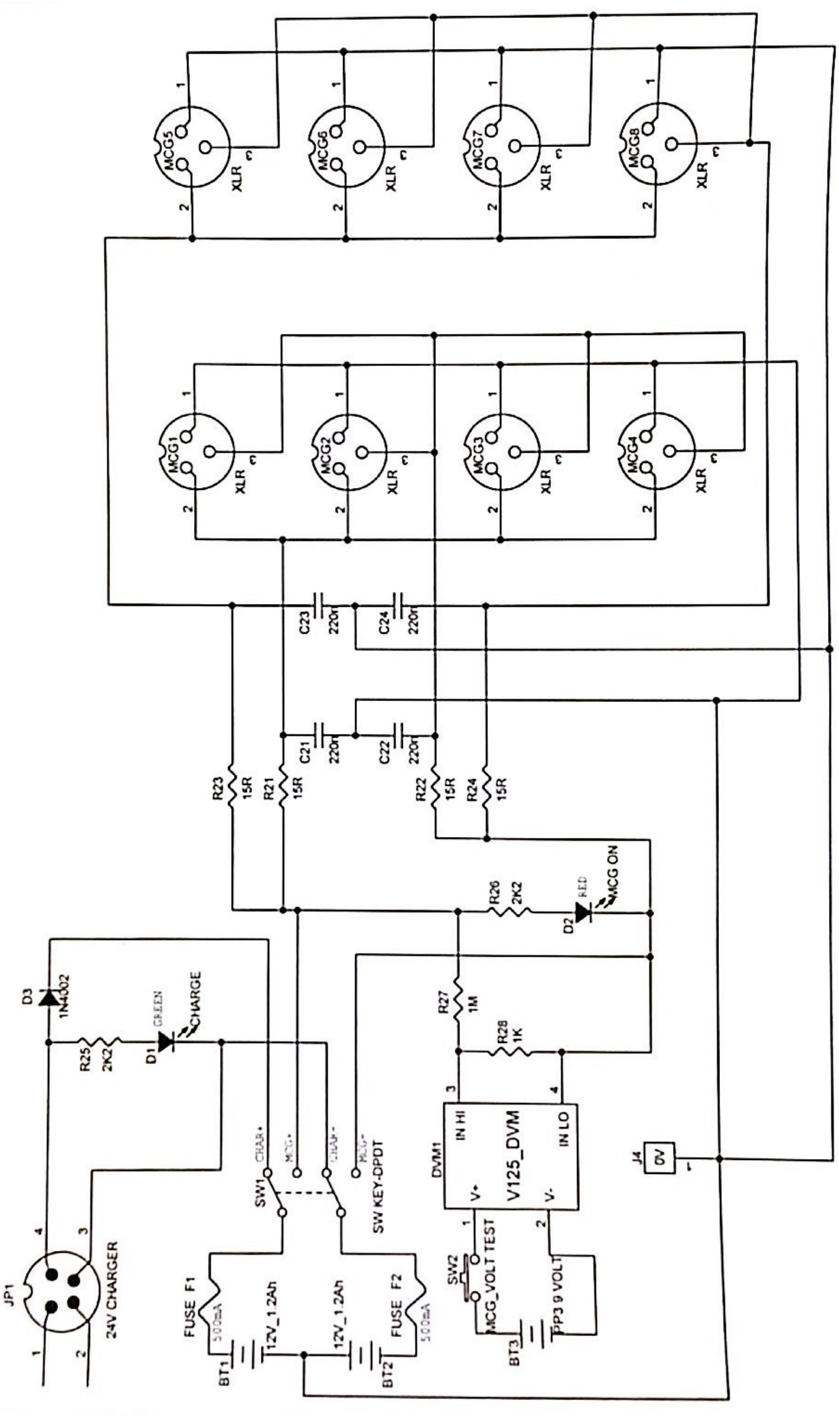
MCG1
MCG2
MCG3
MCG4
MCG5
MCG6
MCG7
MCG8
XLR
1
2
3
C21 220n
C22 220n
C23 220n
C24 220n
R21 15R
R22 15R
R23 15R
R24 15R
R25 2K2
R26 2K2
R27 1M
R28 1K
D1 GREEN CHARGE
D2 RED MCG ON
D3 1N4002
JP1
24V CHARGER
SW1
CHAR+
MCG+
CHAR-
MCG-
SW KEY-DPDT
SW2
MCG_VOLT TEST
FUSE F1 500mA
FUSE F2 500mA
BT1
BT2
BT3
12V_1.2Ah
12V_1.2Ah
PP3 9 VOLT
DVM1
V125_DVM
IN HI
IN LO
V+
V-
J4
0V

FIGURE A.2: Full circuit diagram of the analogue power supply. When in scan mode the batteries are connected to the XLR output sockets via 15Ω resistors, and the red LED is illuminated. A 220nF filter capacitor is connected in parallel between the resistors and the output. When in scanning mode depression of SW2 causes the digital volt meter to display the battery voltage. In charge mode the batteries are connected in series with the 24V charger, and green LED turns on. A diode, D3 is present to prevent reverse current/discharge into the charger.

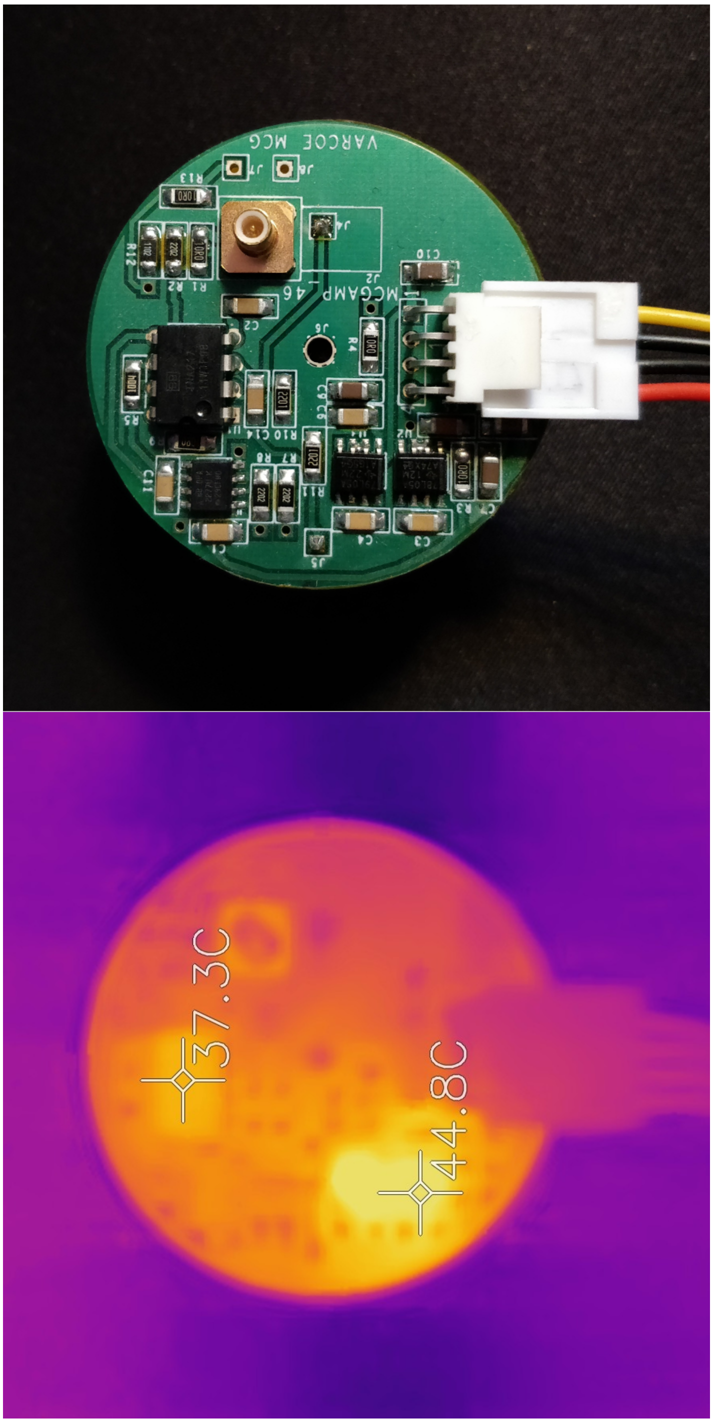

FIGURE A.3: Infrared and visible images of preamplifier PCB at thermal equilibrium in a 20 °C environment. Thermal equilibrium is reached within 20 minutes. The INA217 amplifier chip reached 37.7 °C, the linear voltage regulators reached 44.8 °C. The gain temperature drift coefficient of the amplifier is typically 40 ppm/°C, the resultant drift is likely to be insignificant for cycle averaging. However care should be taken to ensure all amplifiers in the array are identically thermally coupled to the environment, to ensure good matching for coherent noise rejection.

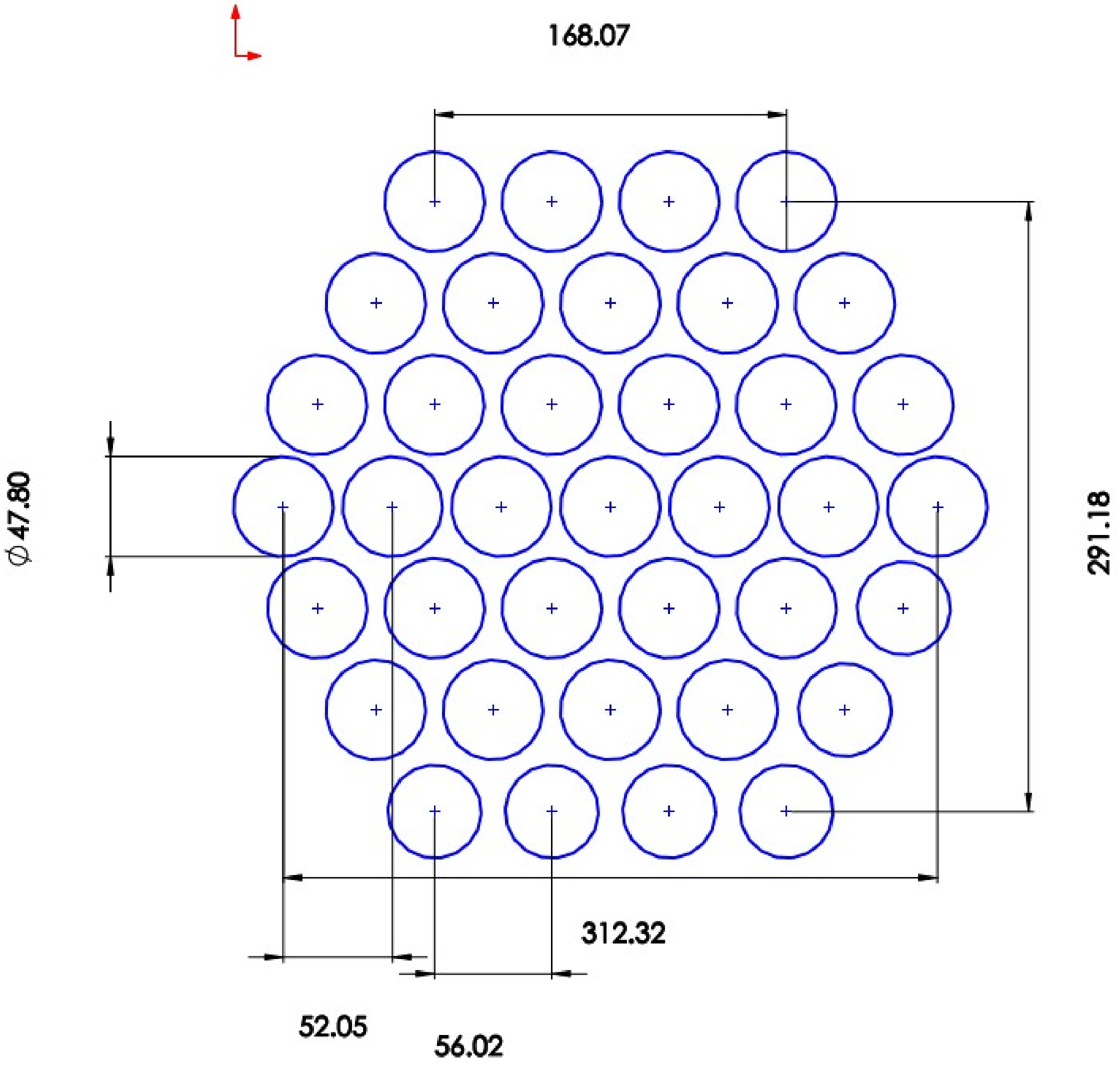
168.07
Ø 47.80
291.18
312.32
52.05
56.02

FIGURE A.4: Dimensioned drawing of the 37 coil array layout.

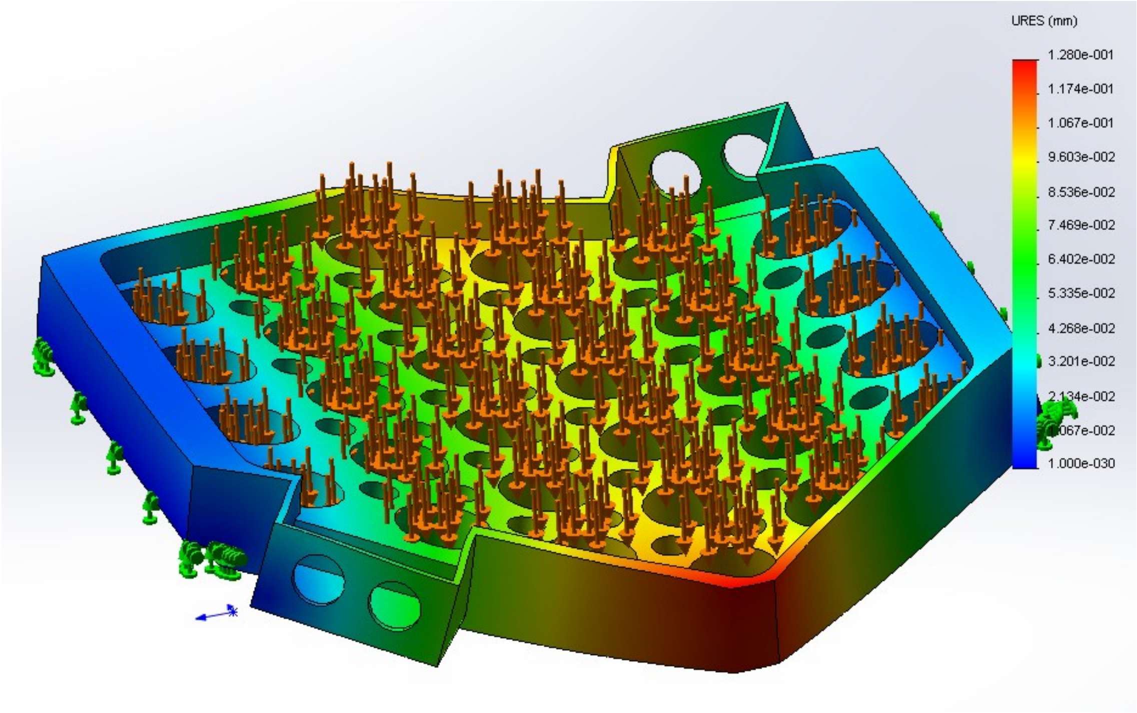
URES (mm)
1.280e-001
1.174e-001
1.067e-001
9.603e-002
8.536e-002
7.469e-002
6.402e-002
5.335e-002
4.268e-002
3.201e-002
2.134e-002
067e-002
1.000e-030

FIGURE A.5: Mechanical simulation of the lightened 37 coil array showing the strain induced by its own weight. The stress was within the elastic deformation capabilities of the material. However a peak strain of 0.1mm displacement was observed and the strain was non uniform across the array. The design is functional, but perhaps a higher stiffness would be better to acoustically couple the sensors.

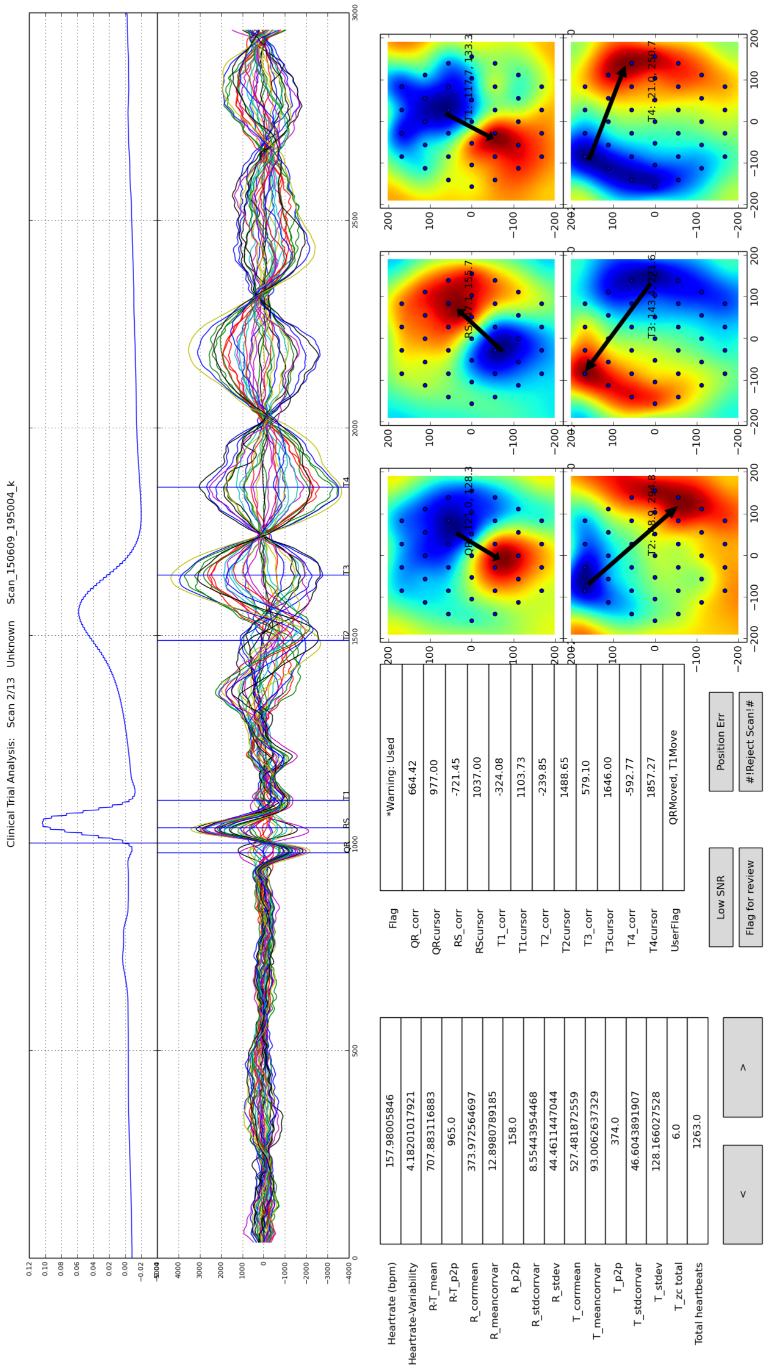

Clinical Trial Analysis: Scan 2/13 Unknown Scan_150609_195004_k

| | |
|---|---|
| Heartrate (bpm) | 157.98005846 |
| Heartrate-Variability | 4.18201017921 |
| R-T_mean | 707.883116883 |
| R-T_p2p | 965.0 |
| R_corrmean | 373.972564697 |
| R_meancorrvar | 12.8980789185 |
| R_p2p | 158.0 |
| R_stdcorrvar | 8.55443954468 |
| R_stdev | 44.4611447044 |
| T_corrmean | 527.481872559 |
| T_meancorrvar | 93.0062637329 |
| T_p2p | 374.0 |
| T_stdcorrvar | 46.6043891907 |
| T_stdev | 128.166027528 |
| T_zc total | 6.0 |
| Total heartbeats | 1263.0 |

| | |
|---|---|
| Flag | *Warning: Used |
| QR_corr | 664.42 |
| QRcursor | 977.00 |
| RS_corr | -721.45 |
| RScursor | 1037.00 |
| T1_corr | -324.08 |
| T1cursor | 1103.73 |
| T2_corr | -239.85 |
| T2cursor | 1488.65 |
| T3_corr | 579.10 |
| T3cursor | 1646.00 |
| T4_corr | -592.77 |
| T4cursor | 1857.27 |
| UserFlag | QRMoved, T1Move |

FIGURE A.6: User interface for analysis for clinical trial analysis. The user cycles through individual scans modifying the cursor positioins to match the standard MCG features and flagging up errors or noise in the scan. The MFM corresponding to each cursor position is shown in the lower left of the screen, these maps are updated whenever a cursor is moved, allowing realtime exploration. The scan shown is from the 37 sensor 'alpha' commercial prototype.

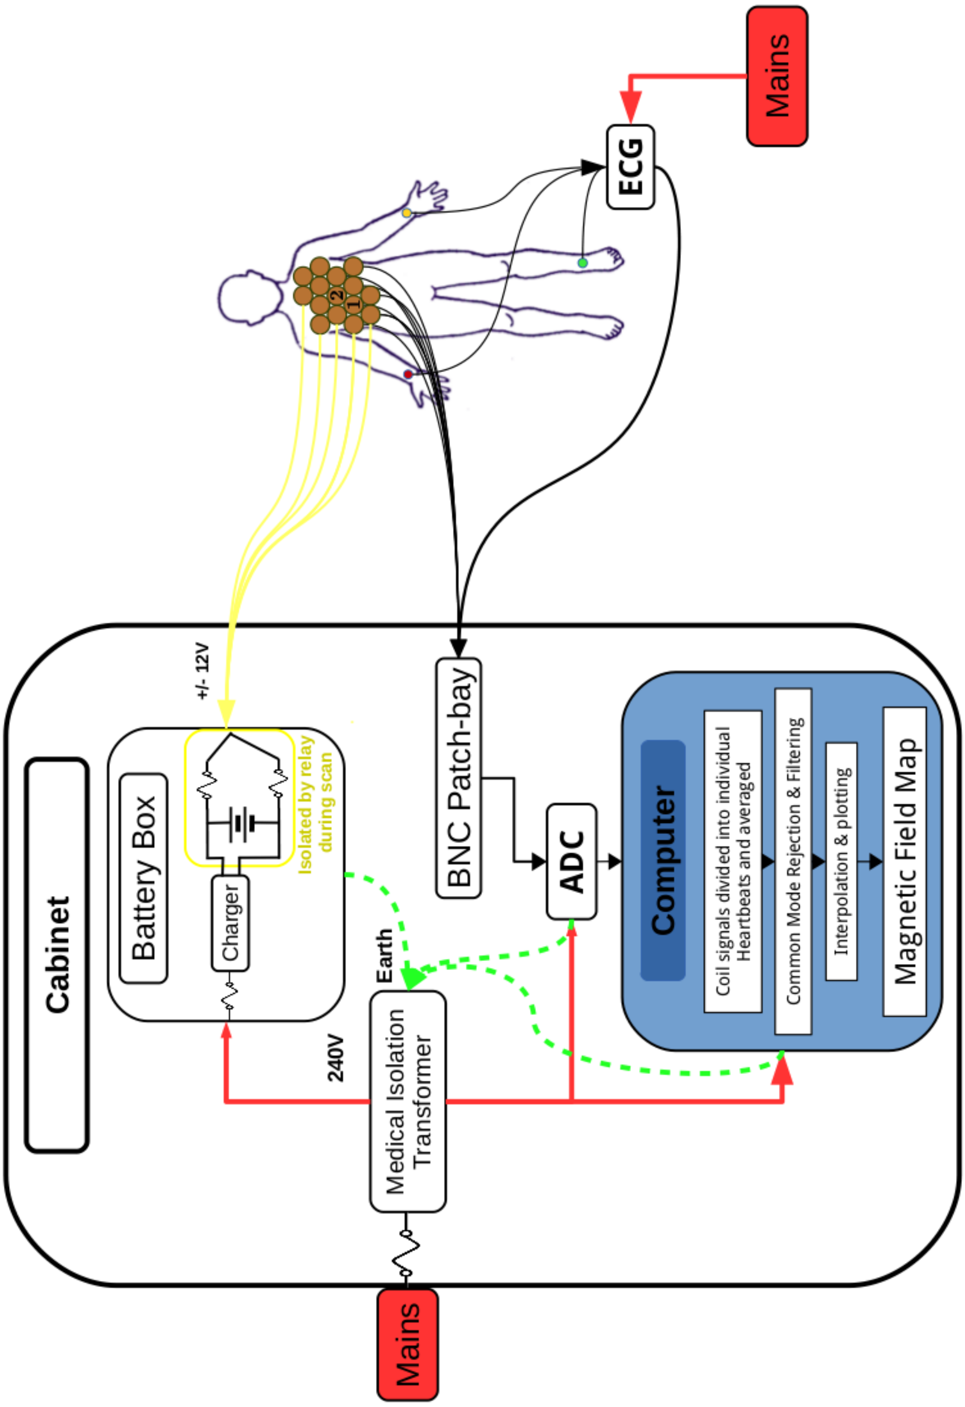
Mains
ECG
1
2
+/- 12V
Cabinet
Battery Box
Charger
Isolated by relay during scan
BNC Patch-bay
ADC
Computer
Coil signals divided into individual Heartbeats and averaged
Common Mode Rejection & Filtering
Interpolation & plotting
Magnetic Field Map
Earth
240V
Medical Isolation Transformer
Mains

FIGURE A.7: Flowchart showing the electrical layout of the MCG device. The patient is isolated from the mains supply by layers of switches and fuses.

# Appendix B

# Equations

Document equations:

Amperes law for the magnetic field generated by an infinitely long wire, at a radial distance, $r$:

$$B = \frac{\mu_0 I}{2\pi r} \tag{B.1}$$

For a pair of wires separated by $d$ with currents travelling in opposite directions the resultant magnetic field at a radius $R$ and angle $\alpha$ from the centre of the pair is [52]:

$$B = \frac{\mu_0 I d}{\pi r^2} cos(\alpha) = 2B\frac{d}{R} cos(\alpha) \tag{B.2}$$

For a twisted pair of wires the magnetic field reduces proportionally with the ratio between pitch of the helix and the diameter of the wire; $\alpha = p/d$. If the wire is tightly twisted, reaching a real world limit of around $\alpha = 25$ the magnetic field can be attenuated by 100x compared to an untwisted pair of wires [53].

# Bibliography

1. Et al, L. Global and regional mortality from 235 causes of death for 20 age groups in 1990 and 2010: a systematic analysis for the Global Burden of Disease Study 2010. en. *The Lancet* **380,** 2095–2128. ISSN: 01406736 (Dec. 2012).
2. Serruys, P. W. Comparison of Coronary-Artery Bypass Surgery and Stenting for the Treatment of Multivessel Disease. en. *The New England Journal of Medicine,* 8 (2001).
3. Ornish, D. *et al.* Can lifestyle changes reverse coronary heart disease?: The Lifestyle Heart Trial. English. *The Lancet* **336,** 129–133. ISSN: 0140-6736, 1474-547X (July 1990).
4. Khan, M. H. *et al.* Comparison between Angiographic Findings of Coronary Artery Disease in STEMI and NSTEMI Patients of Bangladesh. eng. *Mymensingh medical journal: MMJ* **25,** 221–225. ISSN: 1022-4742 (Apr. 2016).
5. Tanindi, A. & Cemri. Troponin elevation in conditions other than acute coronary syndromes. en. *Vascular Health and Risk Management,* 597. ISSN: 1178-2048 (Sept. 2011).
6. Kyoon Lim, H., Kim, K., Lee, Y.-H. & Chung, N. Detection of non-ST-elevation myocardial infarction using magnetocardiogram: new information from spatiotemporal electrical activation map. eng. *Annals of Medicine* **41,** 533–546. ISSN: 1365-2060 (2009).
7. Lim, H. K. *et al.* Usefulness of magnetocardiogram to detect unstable angina pectoris and non-ST elevation myocardial infarction. eng. *The American Journal of Cardiology* **103,** 448–454. ISSN: 1879-1913 (Feb. 2009).
8. Park, J. W. & Jung, F. Qualitative and quantitative description of myocardial ischemia by means of magnetocardiography. eng. *Biomedizinische Technik. Biomedical Engineering* **49,** 267–273. ISSN: 0013-5585 (Oct. 2004).
9. Shin, E.-S. *et al.* Incremental diagnostic value of combined quantitative and qualitative parameters of magnetocardiography to detect coronary artery disease. en. *International Journal of Cardiology* **228,** 948–952. ISSN: 01675273 (Feb. 2017).
10. Lim, H. K. *et al.* Can magnetocardiography detect patients with non-ST-segment elevation myocardial infarction? *Annals of Medicine* **39,** 617–627. ISSN: 0785-3890 (Jan. 2007).
11. Park, J.-W., Hill, P. M., Chung, N., Hugenholtz, P. G. & Jung, F. Magnetocardiography Predicts Coronary Artery Disease in Patients with Acute Chest Pain. en. *Annals of Noninvasive Electrocardiology* **10,** 312–323. ISSN: 1542-474X, 1082-720X (July 2005).
12. Hailer, B. & Van, P. L. Detection of coronary artery disease with MCG. eng. *Neurology & clinical neurophysiology : NCN* **2004,** 82–82. ISSN: 1538-4098 (2004).
13. Backus, B. E. *et al.* A prospective validation of the HEART score for chest pain patients at the emergency department. *International Journal of Cardiology* **168,** 2153–2158 (2013).

14. Fothergill, N. J., Hunt, M. T. & Touquet, R. Audit of patients with chest pain presenting to an accident and emergency department over a 6-month period. en. *Emergency Medicine Journal* **10,** 155–160. ISSN: 1472-0205 (Sept. 1993).
15. Bishopric, N. H. Evolution of the Heart from Bacteria to Man. en. *Annals of the New York Academy of Sciences* **1047,** 13–29. ISSN: 1749-6632 (June 2005).
16. Jensen, B., Wang, T., Christoffels, V. M. & Moorman, A. F. Evolution and development of the building plan of the vertebrate heart. en. *Biochimica et Biophysica Acta (BBA) - Molecular Cell Research* **1833,** 783–794. ISSN: 01674889 (Apr. 2013).
17. Stephenson, A., Adams, J. W. & Vaccarezza, M. The vertebrate heart: an evolutionary perspective. en. *Journal of Anatomy* **231,** 787–797. ISSN: 1469-7580 (Dec. 2017).
18. Torrent-Guasp, F. *et al.* Towards new understanding of the heart structure and function. eng. *European Journal of Cardio-Thoracic Surgery: Official Journal of the European Association for Cardio-Thoracic Surgery* **27,** 191–201. ISSN: 1010-7940 (Feb. 2005).
19. Buckberg, G. D. The Structure and Function of the Helical Heart and Its Buttress Wrapping. II. Interface Between Unfolded Myocardial Band and Evolution of Primitive Heart. en. *Seminars in Thoracic and Cardiovascular Surgery* **13,** 320–332. ISSN: 10430679 (Oct. 2001).
20. Buckberg, G. D. Basic science review: the helix and the heart. eng. *The Journal of Thoracic and Cardiovascular Surgery* **124,** 863–883. ISSN: 0022-5223 (Nov. 2002).
21. Ku, D. N. BLOOD FLOW IN ARTERIES. en, 37 (1996).
22. Numata, S. *et al.* Blood flow analysis of the aortic arch using computational fluid dynamics. en. *European Journal of Cardio-Thoracic Surgery* **49,** 1578–1585. ISSN: 1010-7940, 1873-734X (June 2016).
23. Spiczak, J. v. *et al.* Quantitative Analysis of Vortical Blood Flow in the Thoracic Aorta Using 4D Phase Contrast MRI. en. *PLOS ONE* **10,** e0139025. ISSN: 1932-6203 (Sept. 2015).
24. Jang, J. *et al.* A Reconstruction Method of Blood Flow Velocity in Left Ventricle Using Color Flow Ultrasound. en. *Computational and Mathematical Methods in Medicine* **2015,** 1–15. ISSN: 1748-670X, 1748-6718 (2015).
25. Gharib, M., Rambod, E., Kheradvar, A., Sahn, D. J. & Dabiri, J. O. Optimal vortex formation as an index of cardiac health. en. *Proceedings of the National Academy of Sciences* **103,** 6305–6308. ISSN: 0027-8424, 1091-6490 (Apr. 2006).
26. College, O. *OpenStax: Anatomy and Physiology* 2013. <`https://openstax.org/details/anatomy-and-physiology`>.
27. Sheridan, R. L. *et al.* Emergency management of major hydrofluoric acid exposures. *Burns* **21,** 62–64. ISSN: 0305-4179 (Feb. 1995).
28. staff, B. *"Medical gallery of Blausen Medical 2014"* WikiJournal of Medicine 1 (2). 2014. <`https://en.wikipedia.org/wiki/Skeletal_muscle#/media/File:Blausen_0801_SkeletalMuscle.png`>.
29. McBride, K. K., Roth, B. J., Sidorov, V., Wikswo, J. P. & Baudenbacher, F. J. Measurements of Transmembrane Potential and Magnetic Field at the Apex of the Heart. *Biophysical Journal* **99,** 3113–3118. ISSN: 0006-3495 (Nov. 2010).

30. Lamata, P. *et al.* Images as drivers of progress in cardiac computational modelling. *Progress in Biophysics and Molecular Biology* **115,** 198–212. ISSN: 0079-6107 (Aug. 2014).
31. Dibb, R., Qi, Y. & Liu, C. Magnetic susceptibility anisotropy of myocardium imaged by cardiovascular magnetic resonance reflects the anisotropy of myocardial filament alpha-helix polypeptide bonds. *Journal of Cardiovascular Magnetic Resonance* **17,** 60. ISSN: 1532-429X (July 2015).
32. Lau, S., Petković, B. & Haueisen, J. Optimal Magnetic Sensor Vests for Cardiac Source Imaging. en. *Sensors* **16,** 754 (May 2016).
33. Fenici, R., Brisinda, D., Meloni, A. M., Sternickel, K. & Fenici, P. *Clinical Validation of Machine Learning for Automatic Analysis of Multichannel Magnetocardiography* en. in *Functional Imaging and Modeling of the Heart* (eds Frangi, A. F., Radeva, P. I., Santos, A. & Hernandez, M.) (Springer Berlin Heidelberg, 2005), 143–152. ISBN: 978-3-540-32081-4.
34. COHEN, D. & CHANDLER, L. Measurements and a simplified interpretation of magnetocardiograms from humans. *Circulation* **39,** 395 (1969).
35. Cohen, D. Detection and analysis of magnetic fields produced by bioelectric currents in humans. *Journal of Applied Physics* **40,** 1046–1048 (1969).
36. Cohen, D., Edelsack, E. A & Zimmerman, J. E. Magnetocardiograms taken inside a shielded room with a superconducting point-contact magnetometer. *Applied Physics Letters* **16,** 278–280 (1970).
37. Estola, K. P. & Malmivuo, J. Air-core induction-coil magnetometer design. en. *Journal of Physics E: Scientific Instruments* **15,** 1110–1113. ISSN: 0022-3735 (Oct. 1982).
38. Cohen, D. & McCaughan, D. Magnetocardiograms and their variation over the chest in normal subjects. *The American journal of cardiology* **29,** 678–685 (1972).
39. Hoenig', H. E. *et al.* KRENIKON 1990. en, 9.
40. Mooney, J. W. *et al.* A portable diagnostic device for cardiac magnetic field mapping. en. *Biomedical Physics & Engineering Express* **3,** 015008. ISSN: 2057-1976 (2017).
41. Roudsari, S. G. *Magnetocardiography A Novel Cardiac Diagnostic Technology* en. PhD thesis (University of Leeds, 2016).
42. Ghasemi-Roudsari, S. *et al.* A portable prototype magnetometer to differentiate ischemic and non-ischemic heart disease in patients with chest pain. en. *PLOS ONE* **13,** e0191241. ISSN: 1932-6203 (Jan. 2018).
43. Johnson, J. B. Thermal Agitation of Electricity in Conductors. *Physical Review* **32,** 97–109 (July 1928).
44. Nyquist, H. Thermal Agitation of Electric Charge in Conductors. *Physical Review* **32,** 110–113 (July 1928).
45. Baule, G. M. & McFee, R. The magnetic heart vector. en. *American Heart Journal* **79,** 223–236. ISSN: 00028703 (Feb. 1970).
46. Borna, A. *et al.* A 20-channel magnetoencephalography system based on optically pumped magnetometers. en. *Physics in Medicine and Biology* **62,** 8909–8924. ISSN: 1361-6560 (2017).

47. Lembke, G. *et al.* Optical multichannel room temperature magnetic field imaging system for clinical application. EN. *Biomedical Optics Express* **5,** 876–881. ISSN: 2156-7085 (Mar. 2014).

48. Dale, M. W. & Morley, G. W. Medical applications of diamond magnetometry: commercial viability. en. *arXiv:1705.01994 [cond-mat, physics:physics, physics:quant-ph].* arXiv: 1705.01994. <`http://arxiv.org/abs/1705.01994`> (visited on 09/10/2018) (May 2017).

49. Fujiwara, K. *et al.* Magnetocardiography and magnetoencephalography measurements at room temperature using tunnel magneto-resistance sensors. en. *Applied Physics Express* **11,** 023001. ISSN: 1882-0786 (Jan. 2018).

50. Chaves, R. C., Freitas, P. P., Ocker, B. & Maass, W. MgO based picotesla field sensors. en. *Journal of Applied Physics* **103,** 07E931. ISSN: 00218979 (2008).

51. Sengottuvel, S. *et al. Feasibility study on measurement of magnetocardiography (MCG) using fluxgate magnetometer* en. in (Mumbai, India, 2018), 060018. doi:`10.1063/1.5028788`. <`http://aip.scitation.org/doi/abs/10.1063/1.5028788`> (visited on 09/24/2018).

52. Baltag, O., Rosu, G. & Rau, M. C. *Magnetic field of parallel and twisted wire pairs* en. in *2017 10th International Symposium on Advanced Topics in Electrical Engineering (ATEE)* (IEEE, Bucharest, Romania, 2017), 324–329. ISBN: 978-1-5090-5160-1. doi:`10.1109/ATEE.2017.7905020`. <`http://ieeexplore.ieee.org/document/7905020/`> (visited on 08/26/2018).

53. Yang, C.-F., G. Lai, G., Su, C.-T. & M. Huang, H. Mitigation of magnetic field using three-phase four-wire twisted cables: MITIGATION OF MAGNETIC FIELD BY THREE-PHASE TWISTED CABLES. en. *International Transactions on Electrical Energy Systems* **23,** 13–23. ISSN: 20507038 (Jan. 2013).

54. Pawlowski, N. *Test and optimisation of the current used MCG by COMSOL simulations* en. MPhys Research Project (University of Leeds, 2014).

55. *Indexing — NumPy v1.15 Manual* <`https://docs.scipy.org/doc/numpy/reference/arrays.indexing.html`> (visited on 09/04/2018).

56. Senthilnathan, S. *et al.* Enhancing the Reliability in the Noninvasive Measurement of the His Bundle Magnetic Field Using a Novel Signal Averaging Methodology. en. *Annals of Noninvasive Electrocardiology* **17,** 186–194. ISSN: 1542-474X (July 2012).

57. Accusync. *AccuSync 42 - Accusync Medical Research Corporation* <`http://www.accusync.com/accuSync42.html`> (visited on 09/04/2018).

58. Carreiras, C. *et al. BioSPPy: Biosignal Processing in Python* 2015. <`https://github.com/PIA-Group/BioSPPy/`> (visited on 09/08/2018).

59. Bianchi, C. & Meloni, A. Natural and man-made terrestrial electromagnetic noise: an outlook. en. *ANNALS OF GEOPHYSICS* **50,** 11 (2007).

60. Hong-Fang, G. *et al.* Measurements, characteristics, and origin of new electromagnetic interference on magnetocardiographic measurements. en. *Chinese Physics B* **21,** 040702. ISSN: 1674-1056 (2012).

61. Scarborough, W. R. & Baker, B. M. Editorial: Ballistocardiography Appraisal of Current Status. en. *Circulation* **16,** 971–975. ISSN: 0009-7322, 1524-4539 (Dec. 1957).

62. Pinheiro, E., Postolache, O. & Girão, P. Theory and Developments in an Unobtrusive Cardiovascular System Representation: Ballistocardiography. *The Open Biomedical Engineering Journal* **4,** 201–216. ISSN: 1874-1207 (Oct. 2010).
63. IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology - Random Instabilities. *IEEE Std 1139-1999,* 0_1– (1999).
64. Allan, D. Statistics of atomic frequency standards. en. *Proceedings of the IEEE* **54,** 221–230. ISSN: 0018-9219 (1966).
65. Hanbury Brown, R. & Twiss, R. Q. A Test of a New Type of Stellar Interferometer on Sirius. en. *Nature* **178,** 1046–1048 (Nov. 1956).
66. Hanbury Brown, R., Davis, J. & Allen, L. R. The Angular Diameters of 32 Stars. en. *Monthly Notices of the Royal Astronomical Society* **167,** 121–136. ISSN: 0035-8711, 1365-2966 (Apr. 1974).
67. Wang, L.-G., Qamar, S., Zhu, S.-Y. & Zubairy, M. Hanbury Brown–Twiss effect and thermal light ghost imaging: A unified approach. en. *Physical Review A* **79.** ISSN: 1050-2947, 1094-1622. doi:`10.1103/PhysRevA.79.033835`. <`http://link.aps.org/doi/10.1103/PhysRevA.79.033835`> (visited on 01/09/2015) (Mar. 2009).
68. Sekihara, K. Signal Space Separation Method for a Biomagnetic Sensor Array Arranged on a Flat Plane for Magnetocardiographic Applications: A Computer Simulation Study. en. *Journal of Healthcare Engineering* **2018,** 1–19. ISSN: 2040-2295, 2040-2309 (2018).
69. Kuilen, J. D., Ceccotti, A., Xia, Z. & He, M. Very Tall Wooden Buildings with Cross Laminated Timber. en. *Procedia Engineering* **14,** 1621–1628. ISSN: 18777058 (2011).
70. McDonald, K. T. Low-Frequency Electromagnetic Waves on a Twisted-Pair Transmission Line. en, 13.
71. Daniel Vander--Hyde, F. P. *Testing the magnetic fields produced from twisted wires* <`http://www.phys.ufl.edu/ireu/IREU2013/pdf_reports/Daniel_Vander-Hyde_2_experimental.pdf`> (visited on 09/10/2018).
72. Benzi, R, Sutera, A & Vulpiani, A. The mechanism of stochastic resonance. en. *Journal of Physics A: Mathematical and General* **14,** L453–L457. ISSN: 0305-4470, 1361-6447 (Nov. 1981).
73. Gammaitoni, L., Hänggi, P., Jung, P. & Marchesoni, F. Stochastic Resonance: A remarkable idea that changed our perception of noise. en. *The European Physical Journal B* **69,** 1–3. ISSN: 1434-6028, 1434-6036 (May 2009).
74. Rosso, O. A. & Masoller, C. Detecting and quantifying temporal correlations in stochastic resonance via information theory measures. en. *The European Physical Journal B* **69,** 37–43. ISSN: 1434-6028, 1434-6036 (May 2009).
75. Hänninen, H. *et al.* Detection of Exercise-Induced Myocardial Ischemia by Multichannel Magnetocardiography in Single Vessel Coronary Artery Disease. en. *Annals of Noninvasive Electrocardiology* **5,** 147–157. ISSN: 1542-474X (Apr. 2000).
76. Hänninen, H. *et al.* Features of ST segment and T-wave in exercise-induced myocardial ischemia evaluated with multichannel magnetocardiography. eng. *Annals of Medicine* **34,** 120–129. ISSN: 0785-3890 (2002).
77. Wright, G. *Radial Basis Function Interpolation: Numerical and Analytical Developments* PhD thesis (University of Colorado, Department of Applied Mathematics, 2003).

78. *scipy.interpolate.Rbf — SciPy v1.1.0 Reference Guide* <https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.Rbf.html> (visited on 09/11/2018).
79. Colominas, M. A., Schlotthauer, G., Torres, M. E. & Flandrin, P. Noise–Assisted EMD Methods in Action. en. *Advances in Adaptive Data Analysis* **04,** 1250025 (2012).
80. Sternickel, K., Effern, A., Lehnertz, K., Schreiber, T. & David, P. Nonlinear noise reduction using reference data. en. *Physical Review E* **63.** ISSN: 1063-651X, 1095-3787. doi:10.1103/PhysRevE.63.036209. <http://link.aps.org/doi/10.1103/PhysRevE.63.036209> (visited on 09/14/2016) (Feb. 2001).
81. (wikimedia user), C. *Overfitting.png* <https://commons.wikimedia.org/wiki/File:Overfitting.svg> (visited on 09/22/2018).
82. Hughes, G. On the mean accuracy of statistical pattern recognizers. en. *IEEE Transactions on Information Theory* **14,** 55–63. ISSN: 0018-9448 (Jan. 1968).
83. Kandori, A. *et al.* Space-Time Database for Standardization of Adult MCG Making Standard MCG Parameters. en. *Pacing and Clinical Electrophysiology* **31,** 422–431. ISSN: 1540-8159 (Apr. 2008).
84. Koch, H. *et al.* A reference dataset for verifying numerical electrophysiological heart models. *BioMedical Engineering OnLine* **10,** 1–15. ISSN: 1475-925X (2011).
85. Willcox, M. & Mysak, T. An Introduction to Barkhausen Noise and its Applications. en. *Insight NDT.* <www.insightndt.com/papers/technical/t013.pdf> (visited on 09/23/2018) (2004).
86. Foster, K. R., Koprowski, R. & Skufca, J. D. Machine learning, medical diagnosis, and biomedical engineering research - commentary. *BioMedical Engineering OnLine* **13,** 94. ISSN: 1475-925X (July 2014).
87. Zhang, G. P. Neural Networks for Classification: A Survey. *Trans. Sys. Man Cyber Part C* **30,** 451–462. ISSN: 1094-6977 (Nov. 2000).
88. Turing, A. M. Computing Machinery and Intelligence. *Mind, New Series* **59,** 433–460 (1950).
89. Kononenko, I. Machine learning for medical diagnosis: history, state of the art and perspective. eng. *Artificial Intelligence in Medicine* **23,** 89–109. ISSN: 0933-3657 (Aug. 2001).
90. Silver, D. *et al.* Mastering the game of Go with deep neural networks and tree search. en. *Nature* **529,** 484–489. ISSN: 1476-4687 (Jan. 2016).
91. Kotsiantis, S. B. Supervised Machine Learning: A Review of Classification Techniques. en. *Informatica* **31,** 249–268 (2007).
92. Tantimongcolwat, T., Naenna, T., Isarankura-Na-Ayudhya, C., Embrechts, M. J. & Prachayasittikul, V. Identification of ischemic heart disease via machine learning analysis on magnetocardiograms. eng. *Computers in Biology and Medicine* **38,** 817–825. ISSN: 0010-4825 (July 2008).
93. Kangwanariyakul, Y., Naenna, T., Nantasenamat, C. & Tantimongcolwat, T. Data mining of magnetocardiograms for prediction of ischemic heart disease. <https://eldorado.tu-dortmund.de/handle/2003/27470> (visited on 08/19/2015) (2010).

94. Kukar, M. Quality Assessment of Individual Classifications in Machine Learning and Data Mining. *Knowl. Inf. Syst.* **9,** 364–384. ISSN: 0219-1377 (Mar. 2006).
95. Ben-Hur, A., Horn, D., Siegelmann, H. T. & Vapnik, V. Support Vector Clustering. en. *Journal of Machine Learning Research,* 125–137 (2001).
96. Wilson, D. *Optimising a Support Vector Classifier for use with a Magnetocadiogram* MPhys Research Project (University of Leeds, 2016).
97. *1.4. Support Vector Machines — scikit-learn 0.19.2 documentation* <`http://scikit-learn.org/stable/modules/svm.html`> (visited on 09/22/2018).
98. Lecun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86,** 2278–2324. ISSN: 00189219 (Nov. 1998).
99. Smith, M. *Machine Learning with Convolutional Neural Networks in Medical Diagnosis* en. MPhys Research Project (University of Leeds, Apr. 2017).
100. Leibig, C., Allken, V., Ayhan, M. S., Berens, P. & Wahl, S. Leveraging uncertainty information from deep neural networks for disease detection. *Scientific Reports* **7.** ISSN: 2045-2322. doi:`10.1038/s41598-017-17876-z`. <`https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5736701/`> (visited on 09/22/2018) (Dec. 2017).
101. Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D. & Jana, S. Certified Robustness to Adversarial Examples with Differential Privacy. *arXiv:1802.03471 [cs, stat].* arXiv: 1802.03471. <`http://arxiv.org/abs/1802.03471`> (visited on 09/24/2018) (Feb. 2018).
102. Tsutsumi, T. *et al.* Time–frequency analysis of the QRS complex in patients with ischemic cardiomyopathy and myocardial infarction. *IJC Heart & Vessels* **4,** 177–187. ISSN: 2214-7632 (Sept. 2014).
103. Abboud, S. & Zlochiver, S. High-frequency QRS electrocardiogram for diagnosing and monitoring ischemic heart disease. English. *Journal of Electrocardiology* **39,** 82–86. ISSN: 0022-0736, 1532-8430 (Jan. 2006).
104. Censi, F. *et al.* P-wave Variability and Atrial Fibrillation. en. *Scientific Reports* **6,** 26799. ISSN: 2045-2322 (May 2016).
105. Yoshida, K. *et al.* Ability of magnetocardiography to detect regional dominant frequencies of atrial fibrillation. en. *Journal of Arrhythmia* **31,** 345–351. ISSN: 18804276 (Dec. 2015).
106. Hiraki, T. *et al.* Frequency- and time-domain analysis of P wave in patients with paroxysmal atrial fibrillation. eng. *Pacing and clinical electrophysiology: PACE* **21,** 56–64. ISSN: 0147-8389 (Jan. 1998).
107. MacKay, D. J. C. Information Theory, Inference, and Learning Algorithms. en, 640.
108. Meyen, S. Relation between classification accuracy and mutual information in equally weighted classification tasks. de, 68.
109. Kanzaki, H., Nakatani, S., Kandori, A., Tsukada, K. & Miyatake, K. A new screening method to diagnose coronary artery disease using multichannel magnetocardiogram and simple exercise. en. *Basic Research in Cardiology* **98,** 124–132. ISSN: 0300-8428, 1435-1803 (Mar. 2003).
110. Krishnamurthy, A. *et al.* Patient-specific models of cardiac biomechanics. *Journal of Computational Physics* **244,** 4–21. ISSN: 0021-9991 (July 2013).