

Bioelectric Signal Analysis to Expose Nervous
Control of the Human Heart

by
David George Western

UCL

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

Signed Declaration

I, David Western, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

This thesis describes the development of new methods to infer the nature of nervous control of the human heart using recordings of its electrical behaviour. Malfunctions of this control system are a leading cause of death, and can be triggered by a diverse range of influences including basic physiological factors and one's emotional state. However, the mechanisms of failure remain poorly understood, partly due to a lack of relevant human data. The principal purpose of the work described in this thesis is to improve the availability of such data.

A literature review was conducted, covering the current understanding of electrical activity in the heart and its control by the nervous system, as well as the techniques available to observe that behaviour. A variety of novel techniques were developed and implemented experimentally to demonstrate their utility. Specialised methods for the filtering and subsequent spectral analysis of electrocardiograph (ECG) signals were used to expose differences between psychologically distinct groups in terms of their response to emotional stimuli. Algorithms were developed to automatically process unipolar electrogram recordings with minimal human intervention, enabling the analysis of heterogeneous electrophysiological dynamics, which requires datasets of a size that would otherwise make in-depth analyses intractable. New indices were developed for measuring the timing of localised electrical activation and recovery from unipolar electrograms, in order to overcome the fact that conventional indices are not well suited to dynamic analyses. Experiments using these tools demonstrated that respiration induces heart-rate independent modulation of the ventricles' electrophysiological behaviour via the autonomic nervous system.

By improving the accessibility of human *in situ* data, the developed tools enable new research methodologies to study interactions between the heart and the nervous system, which may ultimately contribute to the development of new treatments to prevent thousands of deaths in the UK alone each year.

Acknowledgements

Out of the many people who have helped me in bringing this thesis together, I would like to thank my supervisor, Ben Hanson, first. He helped me to secure the funding for this project in the first place, and was a joy to work with throughout. He was consistently generous with his time, and repeatedly went out of his way to act in my interests. He was able to grant me confidence in my own ideas, while always providing the insightful criticism that they needed. Thank you, Ben.

Thank you also to Peter Taggart, whose guidance in physiological matters was invaluable, and without whom the key collaborations of this project may never have been initiated. The work of our collaborators is acknowledged in the thesis where appropriate, but I would like to say that it was a pleasure working with Martina Di Simplicio, Andrea Reinecke, Jas Gill, and Julian Bostock. I am grateful to Vanessa Diaz, my second supervisor, for contributing her perspective when needed and for taking the time to give feedback on a draft of this thesis. Thank you to Kevin Drake, Dara McCutcheon, Pete Johnson, James Burnett, Christian Klettner, Doug Lazenby, Adam Harman-Clarke, Kevin Lau, Ben Bhattacharya-Ghosh, Xiao Jie, Lei Xu, Malcolm Finlay, and Justine Bhar-Amato for all the enjoyable and useful discussions along the way.

I would never have been able to pursue this project without the funding provided by the EPSRC, and am grateful to the UCL Department of Mechanical Engineering for choosing to award that funding to me. The bursary awarded by the Worshipful Company of Scientific Instrument Makers was also much appreciated, as were the contributions from the UCL Graduate School towards conference expenses.

Thank you to Simon Neild, Rick Hyde, and Rosie Jones, and to Ben Hanson again, for their patience and faith when I chose to juggle more projects than I ought to.

Thank you to everyone who ever shared a lunch-time crossword with me. That routine disruption of my work so often turned out to be exactly what my work needed.

Thank you to Mum, Dad, Gareth, Hanne, Oliver, Mamgu, Tom, Jenni and Kester, and to my friends in Brixton and beyond for all the love, support, distractions and perspective. I am particularly grateful to wonderful Kate, whose love, patience, and understanding, especially through the later stages of this project, ensured that I never completely lost myself to it. Maybe I could have done it without you, but I certainly wouldn't have wanted to.

Table of Contents

Declaration.....	2
Abstract.....	3
Acknowledgements	4
Table of Contents.....	5
List of Figures	7
List of Tables.....	14
Nomenclature.....	15
Related Publications	18
I. Introduction.....	20
A. Background.....	21
B. Structure of the report.....	22
II. Literature Review.....	22
A. The mechanical structure and function of the heart.....	23
B. Electrical activity in the heart.....	24
C. The cellular basis of cardiac behaviour.....	31
D. Recording the heart's electrical activity	34
E. Unipolar electrograms	47
F. Control of the heart by the autonomic nervous system	64
G. Heart rate variability as an indicator of autonomic activity.....	71
III. Correlating Psychological Processes and ECG Measurements.....	80
A. Simultaneous use of MRI and ECG recording	81
B. Potential cardiac correlates of psychological activity	88
IV. Unipolar electrograms to measure dynamics of electrophysiology	105
A. Obstacles for the use of unipolar electrograms to measure dynamic behaviour.....	106
B. Summary of design requirements	109
C. Development environment	112
D. Description of developed algorithms	113
E. Evaluation of developed algorithms	125
F. Summary of achievements for chapter IV	136
G. Scope for improvement	137
V. Experimental applications	140
A. Can the ECG reveal evidence of emotionally generated asymmetry in autonomic outflow to the heart?.....	140
B. Can HRV analysis with spectral averaging be used to expose cardiac effects of transient emotional responses? – ‘The Oxford Experiments’	149

C.	Exploring respiratory effects on ventricular repolarisation.....	166
D.	Summary of findings	188
VI.	Conclusion	191
A.	Summary of achievements.....	191
B.	Plans and suggestions for future work	193
	References.....	194
Appendix I.	Deriving an Expression for Extracellular Potential, in Terms of Transmembrane Potential, Using the Core-Conductor Model	212
Appendix II.	Deriving an Expression for the Unipolar Electrogram, Incorporating the Lead Field Concept	214
Appendix III.	Deriving an Expression for the Extracellular Potential Induced by a Single Myocyte in an Infinite, Homogeneously Conductive Medium	217
Appendix IV.	Custom Software for Automated QRS Detection.....	220
Appendix V.	The Berger Method for Calculating QT Intervals	226
Appendix VI.	Matlab Code to Find and Remove Pacing Artefacts.....	228
Appendix VII.	Matlab Code to Find Stabilised Activation Times (A_S)	230
Appendix VIII.	Matlab Code to Find Stabilised Recovery Times (R_S).....	235

List of Figures

Figure 1: a) The chambers and blood-flow pathways of the heart. b) A schematic diagram of the circulatory system.....	23
Figure 2: Timing of the action potential and the resultant mechanical contraction in cardiac myocytes.....	24
Figure 3: Progression of the electrical signal through a healthy heart	25
Figure 4: Ventricular myocyte fibre orientation, as it varies through the myocardial wall. The main picture shows the ventricles from a left-posterior view, with the upper half of the heart removed for clarity.....	26
Figure 5: A mechanism leading to re-entry	29
Figure 6: Computer modelling of re-entry due to regional differences in refractory period.	30
Figure 7: Action potential morphologies for a ventricular myocyte (a), a myocyte in the sinoatrial node (b), and a purkinje fibre (c). The non-contractile myocytes (b and c) exhibit an upward drift in the resting action potential.....	32
Figure 8: The ion currents primarily responsible for each phase of the action potential.	33
Figure 9: The electrode configuration used for the three main leads of a standard ECG. The dashed lines show Einthoven’s triangle, which indicates the effective direction of each lead’s polarity.....	35
Figure 10: <i>Left</i> - The variation in the cardiac dipole vector during ventricular depolarisation. Diagrams of the ventricles show the spread of the depolarised region (grey, negative) at three different stages of the process, the timings of which are indicated relative to the start of ventricular depolarisation. <i>Right</i> – The morphology of the QRS complex in the three main ECG leads can be inferred based on the magnitude of the cardiac dipole vector and its alignment with the lead vector at each point in time.....	36
Figure 11: One beat of a healthy ECG trace.....	37
Figure 12: Left) A schematic of the fundamental arrangement for contact MAP recording. Right) The Franz configuration for MAP recordings.....	43
Figure 13: A comparison of MAP and transmembrane action potential (TAP) recorded simultaneously from adjacent sites.....	43
Figure 14: The relationship between the Monophasic Action Potential (MAP) and typical unipolar electrograms (UEGs) taken from the same region of myocardium. Commonly used repolarisation indices, explained in the text, are marked on the MAP and on the UEG for positive (solid) and negative (dashed) T-waves.	45
Figure 15: Two unipolar electrograms (UEGs), recorded simultaneously from adjacent sites, and a bipolar electrogram (BEG), which was calculated as the difference between the two unipolar electrograms. The deflection produced in the BEG by local activation is distinct and coincides with the downslope of the UEGs, as expected. The deflection produced by repolarisation is not distinct in the BEG, making these recordings unsuitable for repolarisation studies.....	46
Figure 16: The morphology of the UEG can be explained as the difference between a local action potential (solid lines in top row) and a smoother remote contribution (dashed lines in top row). In the second row, the local AP is subtracted from the remote contribution to give the simulated UEG, S. The third and fourth rows give the temporal derivatives of the traces in the first and second rows, respectively. Vertical lines mark the steepest downslope of the local action potential (solid red line) and the steepest upstroke of the UEG (dashed black line). t_R is found to coincide with T_{up} . Early-repolarising sites (column A) present positive T-waves, while late-repolarising sites (column C) give negative T-waves.	49
Figure 17: The discrete, linear core-conductor model for a single fibre.....	50
Figure 18: <i>Left</i> : The approximate lead field for ECG lead I across the human torso, ignoring the effects of variations in conductivity. <i>Right</i> : The Brody effect, used here as a demonstration of how the lead field concept facilitates intuitive predictions of the effects of variations in conductivity on bioelectric measurements	52
Figure 19: The negative of the remote component of the simple UEG model, calculated for five different electrode configurations using (5). The lead field was calculated using a	

realistic model of the distribution of conductivity throughout the heart. The dashed line shows the average of the five solid lines.....	53
Figure 20: Unipolar electrograms (UEGs) measured from various sites near the centre of a 30 mm x 20 mm slab of canine ventricular myocardium. Also shown are corresponding computed UEGs, determined using a bidomain model. Note that only the UEG segments corresponding to activation are shown; repolarisation is not shown. The 2-dimensional sequence of activation is shown by isochrones (dashed lines) with a spacing of 2 ms	56
Figure 21: Top – The time-course of a typical mammalian ventricular action potential, as predicted by the Luo-Rudy model ^[Luo & Rudy, 1991] . Lower left – A close-up of the depolarisation phase of the action potential, as demarcated by the dashed box in the top panel. Lower right – The second temporal derivative of the membrane potential during this phase. As described in the text, during depolarisation this morphology can be assumed to be the same as that of the second spatial derivative. Hence it indicates the distribution of membrane current along the myocyte, as indicated by the added ‘position’ axis.....	58
Figure 22: A schematic of the transmembrane and intracellular currents surrounding a depolarisation wavefront in a cardiac myocyte. Also depicted are the associated extracellular potentials.....	59
Figure 23: The transfer function H for a monopole current source (see equation (88) in Appendix III) calculated along one dimension of the myocardial surface. Intuitively, an electrode closer to the area of interest will provide a narrower focus.....	61
Figure 24: The action potential duration (APD) determined from intracellular recordings in canine ventricular myocardium. Substantial transmural heterogeneity of APD exists. M-cells exhibit markedly longer APDs than endocardial and epicardial myocytes, and vary more steeply with changes in heart rate.....	62
Figure 25: Recordings from guinea pig ventricular myocytes show that ischemia causes the action potential to decrease in amplitude and become more triangular in shape.....	63
Figure 26: Hierarchy of the nervous system.....	64
Figure 27: Simplified block-diagram representation of feedback control loop in Figure 28.....	65
Figure 28: Cardiovascular branches of the autonomic nervous system as a feedback control loop	65
Figure 29: Routes of sympathetic and parasympathetic innervation of the heart. It should be noted that this representation is simplified; the ventricles are also sparsely innervated by parasympathetic fibres, and all connections are subject to anatomic variation.....	66
Figure 30: The effects of parasympathetic stimulation of the guinea pig SA node, and effects of sympathetic stimulation of the frog SA node.....	67
Figure 31: Emotional stimuli are processed in various brain centres (the individual roles of which are beyond the scope of this report), leading to a response of the autonomic nervous system. Effects on the heart include altered tone in the blood vessels that supply the myocardium and alterations to myocyte action potentials. Autonomic effects on the various membrane ion channels may lengthen or shorten action potentials heterogeneously, or they may increase the occurrence of early and delayed after-depolarisations (EADs and DADs). All of these effects can lead to the development of ventricular fibrillation (VF).....	70
Figure 32: Two RR interval power spectra, with clear LF and HF peaks, along with pie charts illustrating the LF/HF ratio. The change in the peak magnitudes reflects the ANS response induced by moving from a horizontal resting position to a 90 degree head-up tilt.....	72
Figure 33: A plot of the sinc function, which determines the amplitude of leakage in the m^{th} frequency bin for a sinusoidal signal component with a non-integer number of cycles, k , in the signal.....	74
Figure 34: AR and FFT spectra for two different RR interval series (left and right).....	76
Figure 35: A discrete wavelet transform of a hypothetical signal composed of combined low and high frequency components followed by the separated components. The mother wavelet used was the Daubechies 4 wavelet, which is shown on the right, compressed according to the different values of s	77

Figure 36: A short example of a discrete event series representation of RR intervals, showing that it is irregularly sampled.....	78
Figure 37: The application of the spectral subtraction technique to an ECG recording corrupted by MRI artefact.....	82
Figure 38: Periodic components of a signal can be removed by dividing the signal into windows of an appropriate length, averaging those windows to identify the periodic component, then subtracting that component from the original windows.....	84
Figure 39: Results of the segment-mean method.....	85
Figure 40: Examples of the segment-mean technique and spectral subtraction, used to remove MRI artefact from test signals with different signal-to-noise ratios (SNRs).....	87
Figure 41: A conventional model of habituation in the biological response to a stimulus. <i>Top Panel</i> - The response decreases gradually as the stimulus is sustained. <i>Bottom Panel</i> - When the stimulus is applied intermittently, some recovery occurs. A longer rest period enables greater recovery, so that the magnitude of the response approaches that seen for the initial stimulus.....	90
Figure 42: The solid line is the mean Power Spectral Density (PSD) calculated from four 30-second blocks for a single subject. The dashed lines indicate the maximum and minimum values of the PSD found at each frequency. To reduce the effects of habituation, the 30-second blocks were non-continuous (separated by intervals of 38 seconds) but coincided with similar experimental stimuli. The vertical dashed lines indicate the boundaries of the VLF, LF, and HF frequency ranges. In this example, it can be seen that large peaks centred in the VLF range contribute significantly to the LF range, and may also make small contributions to the HF range.....	90
Figure 43: Partial derivatives (sensitivities) of HF_n with respect to LF (upper panel) and HF (lower panel), when spectral content above the HF band is neglected.....	93
Figure 44: Partial derivatives (sensitivities) of LF/HF with respect to LF (upper panel) and HF (lower panel).....	94
Figure 45: QT -RR relationships for six different healthy human subjects, demonstrating substantial differences between subjects.....	97
Figure 46: Typical weighting profiles for the recent history of RR intervals, used to account for QT -RR hysteresis.....	98
Figure 47: A comparison of the performances of the 16 models listed in Table 4.....	100
Figure 48: A repeat of the analysis in Figure 47, with the spline model excluded.....	102
Figure 49: ARI measurements from UEGs plotted against APDs measured from transmembrane potential recordings. The two are shown to be strongly correlated for measurements from 7 sites in 5 pigs during two different conditions (ischemia and normal myocardial perfusion).....	107
Figure 50: Assessing the stability of the point T_{up} as an indicator of local repolarisation time in the T-wave of the UEG.....	108
Figure 51: A schematic diagram of the proposed use of real-time UEG-based measurements for novel experimental or therapeutic procedures.....	111
Figure 52: A screenshot of Wave_Inspector, the interface designed to facilitate testing of prototype algorithms for UEG processing.....	112
Figure 53: T_{op} - A 3.5 s segment of a UEG recording. The dashed box indicates the region shown in detail in traces A-D. <i>Trace A</i> - A pacing artefact and activation wave, with no processing applied. <i>Trace B</i> - Application of a low-pass Butterworth filter (3 rd order, zero-phase, 100 Hz cut-off) does not effectively eliminate the pacing artefact and causes it to broaden and overlap with the activation time, potentially disrupting the detection of this event. The area in which the influence of the filtered artefact is perceptible is marked with a dashed line, which leads up to the original trace for reference. <i>Trace C</i> - An unfiltered version of the signal, with the pacing artefact removed by a custom algorithm. The area influence by the artefact (marked out by dotted lines) is noticeably smaller than in the filtered version, leaving a greater portion of the activation down-stroke unaffected. <i>Trace D</i> - The filter used to form trace B was applied to trace C. Because the pacing artefact was replaced with low-frequency content, the filtering does not noticeably influence the surrounding signal morphology.....	114

Figure 54: A UEG with vertical lines superimposed to mark activation times. The seven activation waves in view were all detected successfully, despite the fact that arrhythmia and intermittent pacing have caused substantial variations in signal morphology.....	118
Figure 55: Qualitative descriptions of the various criteria applied in the scanning window to identify activation waves.....	119
Figure 56: Heuristic formulae used to establish $T_{end}(i)$, the end of the window in which the T-wave should be located. Each line is labelled with the number of the equation used to construct it. The shaded area represents the time after $A(i)$ in which the T-wave can occur.	121
Figure 57: <i>Top panel</i> – In this example of rate adaptation, ARI switches sporadically between two extremes, obscuring the typical pattern of adaptation. <i>Middle panel</i> – Two consecutive beats in the UEG signal. \dot{V}_{min} and T_{up} are marked as, respectively, dashed and solid vertical lines. The timing of T_{up} switches between two ends of the T-wave upstroke, suggesting the kind of instability described in Figure 50. <i>Bottom panel</i> – Inspection of the signal’s derivative confirms that T_{up} is switching between two ends of a broad peak in dV/dt	123
Figure 58: Examples of cases in which the activation detection algorithm fails, taken from the electrograms used in the accuracy assessment that is summarised in Table 5. <i>Top trace</i> – A small activation wave is not detected due to the influence of the large subsequent wave in the broad window. <i>2nd trace</i> – Forces between the electrode and the myocardium at the catheter tip cause MAP like morphologies. The downward spikes associated with local activation are smaller than usual and are superimposed onto upward deflections, but can still be detected in some cases (for example, in the second beat here, but not the first). <i>3rd trace</i> – At other times, these forces and the resultant artefact appear only during ventricular contraction. The artefacts’ effect on the broad window cause the algorithm to overlook small, fractionated activation waves, and occasionally they are mistaken for activation waves themselves. Comparison with a simultaneous recording from a nearby electrode (bottom trace) reveals which deflections are true activation waves and which are artefact.....	126
Figure 59: A demonstration of the recovery time detection algorithm’s ability to exclude unreliable measurements.....	127
Figure 60: A sample of the electrogram used to produce the ARI measurements in Figure 61. In the first of the two beats shown, the boundary between the activation wave and T-wave upstroke is unclear. This ambiguous morphology ultimately leads to an erroneous identification of the recovery time.....	128
Figure 61: <i>Top panel</i> – Three versions of the same ARI series, taken from electrogram signal of which an example is shown in Figure 60. The first was achieved by switching off the algorithm’s auto-vetting features. The second was achieved by reinstating auto-vetting, so that ambiguous morphologies such as that seen in Figure 60 would not substantially distort the overall dynamics in ARI. Beats that were unnecessarily rejected are marked “X”. The third series was produced by manually correcting the first series. <i>Bottom Panel</i> – Power spectra calculated from these 80-second ARI series using the autoregressive method. The errors associated with ambiguous morphologies lead to substantial distortion of the series’ measured frequency content. In contrast, rejecting ambiguous beats and interpolating across the resultant gaps in the ARI series preserves the key features of the frequency spectrum, even though more beats than necessary were rejected.....	129
Figure 62: APD changes in response to step changes in the cycle length during artificial pacing of an <i>in vivo</i> human heart. APD (measured from MAP recordings) responds with a step change followed by a more gradual convergence. Black arrows indicate the times at which the step change in pacing rate occurs.....	131
Figure 63: The behaviour described by in Figure 62 is evoked at the onset of pacing and exposed using automated detection of activation-recovery intervals from unipolar electrograms. The top trace shows that the cycle length changes suddenly from its natural value of roughly 1230 ms to the paced cycle length (500 ms). The bottom trace shows that the activation-recovery interval responds with a sudden initial change, followed by a more gradual convergence towards a settled value.....	132

Figure 64: Compare with Figure 57. <i>Top panel</i> – When ARI is calculated from the stabilised indices A_S and R_S , the switching behaviour is eradicated, leaving a clearer impression of the rate adaptation profile. <i>Middle panel</i> – Two consecutive beats in the UEG signal. The shaded box indicates the portion of the signal used to calculate R_S , which sits stably near the centre of the T-wave upstroke. <i>Bottom panel</i> – Inspection of the signal’s derivative confirms that the timing of R_S is stable relative to the broad peak in dV/dt . R_S is the centroid of the shaded region between dV/dt and a threshold of 0.5 times the maximum in dV/dt	133
Figure 65: When the conventional (unstabilised) indices of activation and recovery time are used, the sensitivity of the ARI measurement to alternans is extremely inconsistent. Use of the stabilised indices on the same signal reveals that alternans is in fact present throughout this section.	134
Figure 66: An example of switching behaviour in the conventional activation time index, marked by solid arrows for each beat. The extent to which this switching corrupts the dynamics of the index can be seen by plotting the intervals between the pacing artefact and the index (dashed line, bottom panel). When the stabilised index A_S is used (solid line, bottom panel), it is revealed that the changes in activation wave morphology are driven by consistent alternans.	135
Figure 67: The ‘wiring’ of the optic nerve causes visual information from different sides of the visual field to be transmitted to different sides of the brain	141
Figure 68: Timeline illustrating the protocol used for half the subjects in the Brain-Heart Laterality experiments. For the other half, the left-right order was reversed	142
Figure 69: <i>Upper panel</i> – Average response of QT intervals to the IAPS images. The plot was formed by averaging all the epochs in which the negative image was displayed on the left in one series (blue), and all the epochs in which the negative image was displayed on the right in the other (red). <i>Lower panel</i> – P-values calculated using Welch’s t-test, indicating the statistical significance of any difference between the ‘left’ and ‘right’ groups of measurements at each point in time. The dashed line marks the conventional significance threshold, $p = 0.05$. At no time does the measurement cross fall below this threshold to achieve statistical significance.	144
Figure 70: Same as Figure 69, with the exception that these plots compare the 1 st and 2 nd IAPS epochs for each subject, regardless which side images were presented on. P-values notably decrease for the first two minutes of images, achieving statistical significance ($p < 0.05$) for much of this period.	145
Figure 71: Average response of RR intervals to the IAPS images. The plot in the top panel was formed by averaging all the epochs in which the negative image was displayed on the left in one series (blue), and all the epochs in which the negative image was displayed on the right in the other (red). The plot in the lower panel was formed by averaging the first (blue) and second (red) IAPS epochs, regardless of which sides the images were presented.	146
Figure 72: Average heart-rate corrected repolarisation response to IAPS images, measured as QT_C	147
Figure 73: A graphical presentation of the IAPS protocol used in the emotion regulation experiments, with time running from top-left to bottom-right.....	150
Figure 74: Average subjective affect ratings after each of the image blocks, with error bars representing the standard errors. A significant difference exists between the ‘maintain’ and ‘suppress’ conditions, but not between the two groups.....	152
Figure 75: Mean values of LF/HF, with error bars representing the standard error of the samples. The two groups show opposite tendencies in the difference between maintaining and suppressing. Taken individually, this tendency is only significant for the panic group ($p = 0.044$), but the difference between the two groups’ tendencies is more significant ($p = 0.036$).....	154
Figure 76: In low-neurotics, the power in the HF band of the HRV spectrum was found to be higher when suppressing one’s emotional response to an aversive stimulus. In high-neurotics, a non-significant tendency of the opposite sign was found.	156
Figure 77: A significant group \times task interaction was found in HF. Although neither within-group difference is statistically significant (p-values shown from t-tests), the difference between these tendencies is significant ($p = 0.043$).....	159

- Figure 78: Gender differences in HF during the image blocks. The drug SSRI appears to reduce HF when suppressing, relative to when maintaining ($p=0.043$). This effect is most significant in women.159
- Figure 79: Correlations between *VLF* and the two basic indices of autonomic activity, *LF* and *HF*. Although positive correlations exist, they cannot be wholly explained by low-frequency leakage due to the use of short signals; the correlations still exist when longer recordings are used.....161
- Figure 80: *Left* – Diagram showing the positioning of the electrode catheters within the heart. RA and LA are the right and left atria. RV and LV are the right and left ventricles. *Right* – X-ray image of the catheters in place.168
- Figure 81: Histograms showing the difference in median ARI between 6 breaths-per-minute and three other breathing rates. These calculations were performed for 20 electrodes in 16 different subjects. Electrodes were excluded if fewer than 90 percent of beats produced successful recovery time identification. Also shown are the mean difference and associated p-value (from Welch’s t-test) for each graph. It must be noted that, as explained in the text, these are not true t-tests because the data are not truly independent readings.....172
- Figure 82: Examples of oscillations in ARI at three different respiratory frequencies, along with spectra calculated from the respiratory trace (abdomen circumference) and ARI series using the Fourier Transform. The spectra confirm that the dominant frequency of the ARI oscillations matches the respiratory frequency precisely.....173
- Figure 83: Examining the phase relationships between respiration, ARI, and blood pressure at five different respiratory frequencies for a single subject. All ARI series are calculated from the same electrode (RV 5). Dashed lines mark the onset of inspiration, and solid vertical lines mark the onset of exhalation. In the bottom-right pane, the electrograms corresponding to the two circled beats in the 12 BrPM pane are superimposed for comparison, temporally aligned so that the pacing artefacts (the first deflection seen) are coincident. The measured activation times (A_S) are also coincident, whereas the measured recovery time (R_S , marked by vertical black arrows) occurs earlier on inspiration than on exhalation.....175
- Figure 84: Oscillations in ARI and ABP during a breath-hold.....176
- Figure 85: Averaged power spectral density (PSD) in ARI (left) and ABP (right) in the range 0.08 Hz – 0.6 Hz, at the site with the clearest oscillations, shown for five subjects and all three drug conditions. The solid black squares indicate the mean value for each condition177
- Figure 86: Time-frequency representation of ARI oscillations, using the wavelet transform. *Panel A* – Wavelet transform applied to the respiratory signal. Clear ridges coincide with the different respiratory frequencies (15 BrPM = 0.25 Hz, 9 BrPM = 0.15 Hz, 30 BrPM = 0.5 Hz, 6 BrPM = 0.1 Hz) The respiratory to which the transform was applied is presented on the back wall. *Panel B* – Wavelet transform applied to an ARI series, with clear ridges occurring in the same places as in panel A. On the back wall, F_D is the dominant frequency in the ARI (black) and RESP (red) signals, calculated as the frequency with the largest wavelet coefficient at each point in time (F_D for the RESP signal is calculated from the wavelet transform in panel A.). This trace confirms that the oscillations in ARI track the respiratory frequency closely. The RESP signal is presented again on the back wall for reference, along with the ARI series to which the wavelet transform was applied. *Panel C* – Same as panel B, except that the ARI and RESP signals used are from after the administration of metoprolol to block sympathetic activity. From the colour map, it can be seen that the ridges corresponding to ARI oscillations are, on average, reduced in magnitude. *Panel D* – Same as panels B and C, but using RESP and ARI signals from after parasympathetic blockade. In this subject, parasympathetic blockade has caused further reduction in the magnitude of ARI oscillations, to the extent that the dominant frequency F_D for ARI no longer tracks the respiratory frequency consistently.....178
- Figure 87: Oscillations in ARI and ABP observed during Cheyne-Stokes respiration, a form of CSA. ARI and ABP oscillate at the respiratory frequency (~ 0.3 Hz here) and the slower CSA frequency (~ 0.015 Hz). The degree to which each of the two frequencies

is expressed in ARI differs substantially between electrodes. Differences can also be observed in the phase of the slower CSA oscillation, e.g. between LV 10 and LV 1.....	179
Figure 88: During each hyperpnea, as respiratory volume increased an increase was also seen in the time delay between each onset of inspiration and the subsequent peak in systolic blood pressure. The two events are marked for each beat with white dots in the top two panels. In the bottom trace, the solid line shows the varying delay between the two events. The relationship cannot be explained by changes in the duration of each breath, as indicated by the dashed line in the bottom panel.....	180
Figure 89: The shifting temporal relationship between ABP and respiration, as introduced in Figure 88, was observed in all instances of CSA. From plots similar to Figure 88, representative values of the delay between onset of inspiration and peak systolic pressure were chosen for the beginning, middle, and end of each hyperpnea cycle. These values were then averaged across all cycles observed in a particular subject (min. 2 clear cycles, max. 10) to produce the values presented in this figure. T-tests indicated that the values for the start and end of each hyperpnea differed from the values in the middle with statistical significance (p-values shown on graph).....	181
Figure 90: Two unipolar electrograms from the same electrode in the left ventricle (LV 6) at different stages of the breathing cycle, aligned according to the pacing artefact. The measured times of local activation (A_s) and recovery (R_s) are marked on each trace with vertical lines. The delay between pacing stimulus and activation time was often found to vary with respiration.	183
Figure 91: An illustration of the R wave detection method described in steps II-IV above, applied to an ECG signal containing blood-flow artefacts caused by the concurrent use of MRI.	221
Figure 92: RRvarGUI, developed for ECG processing.....	223
Figure 93: A GUI developed to allow quick and easy correction of detected QRS complexes.....	224
Figure 94: Where the end of a T-wave is indistinct the intersection of the steepest downslope with an isoelectric baseline may be used. In this particular case the region of uncertainty is approximately 30 ms long.....	226
Figure 95: Illustration of the Berger method	227

List of Tables

Table 1: A quantified comparison of the performances of the segment-mean and spectral subtraction techniques, using the “residual variance” variable.....	86
Table 2: A quantified comparison of the performances of the segment-mean and spectral subtraction techniques, using the “harmonic peak reduction” variable.....	86
Table 3: p-values calculated using Welch’s t-test to show the statistical significance of differences in the various HRV measurements between the ‘stimulation’ and ‘rest’ conditions. Overall, the statistical significances yielded by the spectral averaging method were stronger than those yielded by the use of only the first instances of the stimulation and rest periods.	91
Table 4: A list of models that were tested for their ability to describe the <i>QT-RR</i> relationship.....	99
Table 5: Accuracy percentages for the activation wave detection algorithm applied to 10 electrodes in each ventricle.....	125
Table 6: Summary of outcomes when the analysis of Figure 61 was extended across all 20 simultaneous electrogram recordings from the same subject.....	128
Table 7: Quantitative comparison of the distortion of ARI dynamics (spectral content) with and without auto-vetting switched on.....	130
Table 8: A summary of p-values found using Welch’s t-test to describe the distinction between various experiment conditions in terms of six HRV based measurements.....	153
Table 9: p-values from a repeated measures ANOVA to expose group effects and interactions between group and task (‘Maintain’ vs. ‘Suppress’).....	156
Table 10: p-values from a repeated measures ANOVA examining the differences in ‘resting’ HRV between the initial screening and the pre-task ECG recording (after the seven day course of SSRI or placebo)	158
Table 11: p-values from a repeated measures ANOVA examining the differences in HRV between the ‘maintain’ and ‘suppress’ image blocks. Pre-task ‘resting’ values of the HRV parameters were entered as covariates to control for individual differences.....	158
Table 12: p-values from ANOVAs used to expose statistically significant effects in the orienting response.....	160
Table 13: Summary of oscillations in activation-recovery intervals (ARI) and arterial blood pressure (ABP) across 16 subjects.....	174

Nomenclature

Symbols

A_S	The stabilised activation-time index. An alternative to \dot{V}_{\min} in UEG analysis.
$d\vec{S}$	Outward normal vector to an infinitesimal surface element, with magnitude equal to the element's area.
dV/dt	The slope (first time-derivative) of the UEG.
I_e	Extracellular current.
I_i	Intracellular current.
\bar{I}_m	Transmembrane current per-unit-length of myocyte.
L	The local component of a UEG.
$r(x, p)$	Distance from x to p .
\bar{R}_e	In the core-conductor model, resistance per-unit-length for the extracellular domain.
\bar{R}_i	In the core-conductor model, resistance per-unit-length for the intracellular domain.
R_S	The stabilised recovery-time index. An alternative to T_{up} in UEG analysis.
t_R	The steepest downward slope during phase 3 of an action potential, used as an index of the timing of repolarisation.
t_{ST}	In recovery-time detection, the beginning of the window in which T_{up} can be found.
T_{end}	In recovery-time detection, the end limit of the window in which T_{up} can be found.
T_{up}	The steepest up-stroke of the UEG T-wave, used as an index of the timing of local repolarisation.
T_{down}	The steepest downward slope of a positive T-wave in the UEG, sometimes used as an alternative index of the timing of local repolarisation.
T_{d2}	The timing of the first local-minimum of the UEG's second derivative after T_{up} .
v_m	Transmembrane electrical potential.
\dot{V}_{\min}	The steepest downward slope of an activation wave in the UEG, used as an index of the timing of local depolarisation.
x	Distance along the length of a cell.
Γ	The remote component of a UEG.
σ_e	Extracellular conductivity.
σ_i	Intracellular conductivity.
Θ	The outer boundary surface of the heart.
ϕ_e	Extracellular electrical potential.
ϕ_i	Intracellular electrical potential.
∇Z	Lead field.

Abbreviations

ABP	Arterial Blood Pressure.
ANOVA	Analysis of Variance.
ANS	Autonomic Nervous System.
AOI	Area of Influence of a pacing artefact after processing.
APD	Action Potential Duration.
AR	Auto-regressive (model).
ARI	Activation-Recovery Interval; taken from UEGs as a surrogate measure of APD.
AV node	Atrioventricular node
BEG	Bipolar Electrogram; a recording of electrical potential difference between two electrodes positioned within a few millimetres of each other, inside the myocardium or on its surface.
BHL	Brain-Heart Laterality; the hypothesis that left-right asymmetric processing in the brain may lead to laterally unbalance autonomic drive to the heart, leading to unstable cardiac electrical activity.
BP	Blood Pressure.
BrPM	Breaths Per Minute.
CL	Cycle length. The inverse of heart rate.
CNS	Central Nervous System.
CSN	Cardiac Sympathetic Nerve.
CSA	Central Sleep Apnea.
DI	Diastolic Interval.
ECG	Electrocardiograph.
ERP	Effective Refractory Period; determined experimentally as the shortest period after depolarisation at which it is possible to induce another depolarisation using an artificial stimulus at the site.
FFT	Fast Fourier Transform.
FIR	FIR filter; Finite Impulse Response.
fMRI	Functional Magnetic Resonance Imaging.
GUI	Graphical User Interface.
HDIG	History Dependent Inverse Gaussian (model).
HF	The high-frequency (0.15 Hz – 0.4 Hz) band of the HRV spectrum.
HF	The power in the high-frequency (0.15 Hz – 0.4 Hz) band of the HRV spectrum.
HF_n	The normalised version of HF. See equation (15), page 73.
HRV	Heart Rate Variability.
IAP	Intracellular Action Potential; a recording of a single cell's electrical behaviour, using an electrode positioned inside the cell body and another just outside the membrane.

IAPS	International Affective Picture System.
IFT	Inverse Fourier Transform.
LF	The low-frequency (0.04 Hz – 0.15 Hz) band of the HRV spectrum.
LF	The power in the low-frequency (0.04 Hz – 0.15 Hz) band of the HRV spectrum.
LF_n	The normalised version of LF. See equation (15), page 73.
LV	Left ventricle.
MAP	Monophasic Action Potential; a recording of electrical behaviour in a local area of myocardium, serving as an approximation of the action potential of a single cell in that region.
MAP_{90}	The time, during repolarisation, at which the MAP recovers 90 percent of the difference between the resting value and the peak (depolarised) value.
MRI	Magnetic Resonance Imaging.
NS	Nervous System.
HR_{\min}	A measure of the ‘orienting’ response; the minimum heart rate in the first 2 seconds after stimulus onset.
HR_{\max}	A measure of the ‘orienting’ response; the maximum heart rate occurring 2-4 seconds after stimulus onset.
PSD	Power Spectral Density.
QT	QT interval; the time between the Q-wave and T-wave in a single beat of the ECG.
QT_c	QT corrected; the QT interval, normalised to account for heart-rate.
RR	RR interval; the time between two consecutive R-waves in the ECG.
RSA	Respiratory Sinus Arrhythmia; oscillations in heart-rate at the respiratory frequency.
RV	Right ventricle.
SA node	Sino-atrial node.
SCD	Sudden Cardiac Death.
SSRI	Selective Serotonin Reuptake Inhibitor; a class of antidepressant.
UEG	Unipolar Electrogram; a recording of electrical potential difference between two electrodes, one of which is positioned inside the myocardium or on its surface, while the other is positioned more remotely but within the conductive medium.
VF	Ventricular Fibrillation.
VLF	The very-low-frequency (0 Hz – 0.04 Hz) band of the HRV spectrum.
VLF	The power in the very-low-frequency (0 Hz – 0.04 Hz) band of the HRV spectrum.

Related Publications

The following publications were generated from this project.

Journal Papers:

Hanson, B., Western, D., Bostock, J., Gill, J., Boyett, M., Zhang, H., Coronel, R., and Taggart, P., 2012. Cyclical modulation of human ventricular repolarization by respiration. **Recently submitted** (May 2012) to *Frontiers in Cardiac Electrophysiology*.

Di Simplicio, M., Costoloni, G., Western, D., Hanson, B., Taggart, P., and Harmer, C., 2011. "Decreased Heart rate variability during emotion regulation in subjects at risk for psychopathology." *Psychological Medicine*, Available through Cambridge Journals Online [Accessed 7 June 2012].

Conference Papers:

Western, D., Hanson, B., and Taggart, P., 2011. Oscillatory patterns of respiration: consequences for the stability and control of cardiac electrophysiology. Proceedings of *The 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1949-1952, Boston, USA, 30 Aug-3 Sep 2011.

Western, D., Taggart, P., and Hanson, B., 2010. Real-time feedback of dynamic cardiac repolarization properties. Proceedings of *The 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 114-117, Buenos Aires, Argentina, 1-4 Sep 2010.

Conference Posters:

Gill, J., Hanson, B., Bostock, J., Western, D., Aldo Rinaldi, C., and Taggart, P., 2012. The effect of emotional stress induced by movie clips on local endocardial activation and repolarization. Abstract to be published in the proceedings of *Heart Rhythm 2012*, Boston, USA, 9-12 May 2012.

Di Simplicio, M., Western, D., Hanson, B., Taggart, P., and Harmer, C., 2012. Short-term antidepressant administration modulates cardiovascular response during emotion regulation in subjects at risk for psychopathology. Abstract in the proceedings of *The 67th Annual Meeting of the Society of Biological Psychiatry*, Philadelphia, USA, 3-5 May 2012.

Western, D., Gill, J., Hanson, B., Bostock, J., and Taggart, P., 2011. Endocardial ventricular repolarization oscillates at respiratory frequencies in humans: insights on entrainment from wavelet analysis. *Journal of the American College of Cardiology*, v. 57, no. 14, p. E125, 2011.

Gill, J., Hanson, B., Western, D., Bostock, J., and Taggart, P., 2010. Respiration exerts phasic modulation of ventricular repolarisation. Abstract in the proceedings of *Heart Rhythm 2010*, Denver, USA, 12-15 May 2010.

Di Simplicio, M., Costoloni, G., Western, D., Hanson, B., Taggart, P., and Harmer, C., 2010. Cardiovascular autonomic nervous system alterations during emotion regulation in subjects at risk for psychopathology. Abstract in the proceedings of the *65th Annual Meeting of the Society of Biological Psychiatry*, New Orleans, USA, 20-22 May 2010.

I. Introduction

I.A. Background

Each contraction of the heart occurs in response to an electrical signal that is passed between adjacent cardiac muscle cells. The orchestrated progression of this signal through the heart is essential to efficient propulsion of blood around the body, and disruptions to the signal can lead to fatal heart malfunctions. In fact, the proximal cause of most cardiac-related deaths is some form of disruption in electrical activity. In the UK alone, around 100,000 such cases – known as Sudden Cardiac Deaths – occur annually^[NHS Choices, 2011]. Currently, no causal mechanism can be identified in approximately 1 in 20 of these cases^[Bowker et al, 2003]. Consequently, research and clinical attention is directed towards the cardiac control mechanisms that exist to regulate the spread of electrical activity. The human heart incorporates several such mechanisms to ensure a sufficient and reliable supply of blood to the body's organs. High-level mechanisms, driven by nervous input and hormonal influences, adapt the heart's behaviour to fluctuating demands. Meanwhile, low-level mechanical, electrical, and chemical interactions maintain the beat-to-beat stability of cardiac activity.

The engineer's role in researching these control mechanisms is rapidly expanding. Computer technology has become an integral part of modern medicine, to the extent that such applications now make up a substantial branch of the computing industry. These advances allow physiological data to be gathered and processed in far greater quantities and detail than were previously possible, and to be interpreted through manipulations that were previously prohibited by their computational expense.

This thesis focuses on the development of novel signal processing methods to elucidate the control mechanisms of the heart. Cardiac electrical activity can be recorded using electrodes positioned on the skin or, for greater spatial specificity, inside the heart. Such recordings have been used for over a century to gain insight into the distribution of electrical activity during each cardiac cycle. Analysis of this behaviour and the time-course of any changes or fluctuations in response to experimental stimuli can expose the nature of the underlying control mechanisms, as is the aim of the work described in this report. New techniques and modifications of existing techniques were developed to investigate the nature of the autonomic nervous system's control of the heart. These techniques were implemented in several different experiments, yielding novel physiological observations that expand the current understanding of nervous interactions with the heart.

The overarching purpose of this work was to improve the accessibility of human *in vivo* data relating to the interactions between the nervous system and the heart. It is hoped that the developed tools will enable the expansion of a 'systems identification' approach to exploring these physiological mechanisms, employing established engineering approaches to the characterisation of unknown systems.

I.B. Structure of the report

This report begins with an overview of the workings of the heart, the means by which its behaviour is controlled by the nervous system, and the techniques used to monitor its electrical activity, focussing on those techniques that are suitable for assessing nervous input. This overview is provided in the form of a literature review in section II.

Sections III and IV describe the signal processing tools and techniques developed during this project. Section III focuses on techniques for the analysis of non-invasive electrocardiograph (ECG) signals, with special relevance to experiments investigating modulation of cardiac behaviour by the central nervous system. Section IV describes techniques for analysing the invasive recordings known as unipolar electrograms (UEGs). New measures of localised electrical behaviour are defined. These new measures enable the use of UEGs to characterise dynamic variations in cardiac electrophysiology, thus providing insights into the nature of the underlying control mechanisms. In contrast, pre-existing measures are only suitable for quasi-static analyses.

In section V, experimental implementations of the developed methods are discussed. Their performance is assessed and several new physiological insights are described.

Section VI offers a summary of the work carried out. Plans and recommendations for future work are also described. Further details on various aspects of this project can be found in the appendices and in the publications listed on page 18.

II. Literature Review

This chapter explains the present understanding of nervous control of the heart and existing techniques for measuring and characterising cardiac electrical activity. The various topics covered in each section and their relevance to this project are outlined below. Sections II.A-II.C cover the basic workings of the heart and cardiac electrophysiology to orient readers from a conventional engineering background to the subject. They include descriptions of common mechanisms by which normal cardiac function breaks down. An awareness of these pathologies is necessary to appreciate the relevance of the novel analytical tools described later in this thesis; the main incentive for developing these tools was to enable research to improve our present understanding of the heart's mechanisms of failure. Furthermore, the reasoning on which the developed methods are based relies on an understanding of the behaviours that they are designed to elucidate. Nevertheless, readers who are already familiar with cardiac electrophysiology may choose to begin at section II.D, page 34.

Sections included in the Literature Review:

- **II.A. Mechanical structure and function of the heart**
- **II.B. Electrical activity in the heart**
- **II.C. The cellular basis of normal cardiac behaviour**
- **II.D. Recording the heart's electrical activity:** This section outlines the various approaches that exist to record cardiac electrical activity, including invasive and non-invasive methods. By clarifying the limitations of the available methods, it justifies the development of novel methods to overcome those limitations, described in later chapters.
- **II.E. Unipolar electrograms:** This particular recording technique is explored in greater detail because some of the most important novel contributions described later in the thesis are signal processing techniques designed specifically for these recordings.
- **II.F. Control of the Heart by the Autonomic Nervous System:** The novel signal processing techniques developed later in this thesis are designed specifically to enable research into autonomic cardiac control. This section outlines those control mechanisms to the extent that is necessary in order to appreciate how they informed the design process.
- **II.F. Heart rate variability as a window on autonomic activity:** This section acknowledges the substantial work that has been carried out previously in this particular field of cardiac signal processing. In doing so, it contextualises the novel contributions to that field, described in section III.

II.A. The mechanical structure and function of the heart

To properly understand the control mechanisms of the heart, they must be considered in the context of the healthy cardiac cycle, which they exist to maintain. To that end, this section briefly outlines the structure and function of the heart.

The role of the heart is to provide the pressure gradient that drives blood around the body's circulatory system. Therefore its continual operation is crucial to ensuring that all regions of the body are supplied with the necessary oxygen, nutrients, and warmth. The mammalian heart can be seen as two separate pumps beating simultaneously; the right side propels deoxygenated blood through the pulmonary system to be oxygenated, and the left side pumps the oxygenated blood around the body. As shown in Figure 1, each of these pumps is further divided into two chambers: a ventricle, which provides the majority of the pressure increase, and an atrium, which pre-loads the ventricle to improve its efficiency.

The relative complexity and capacity of the systemic circulation compared to the pulmonary circulation results in a greater drop in pressure as the blood passes through. As blood exits through the aorta, it varies in pressure between 80 and 120 mmHg (10.7-16.0 kPa), and returns to the right atrium at less than 5 mmHg (0.7 kPa). In the pulmonary circulation, blood typically leaves the right ventricle at roughly 10-23 mmHg (1.3-3.1 kPa) and enters the left atrium at around 7 mmHg (0.9 kPa)^[Noble et al, 2005]. To overcome this obstacle, the left ventricle has considerably greater muscle mass than the right ventricle; its walls are around three times as thick^[Levick, 2003].

The wall dividing the two ventricles is known as the septum. Contraction of the right ventricle manifests itself as a movement of the right ventricle's outer wall towards the septum. In contrast, the left ventricle reduces in both diameter and length upon contraction. The asymmetry in the structure and contraction of the heart cause it to twist forwards with each beat. A fine degree of coordination between all cardiac cells is required to ensure that the contractions are mechanically efficient. That coordination is dependent on the specific manner in which the wave of activity

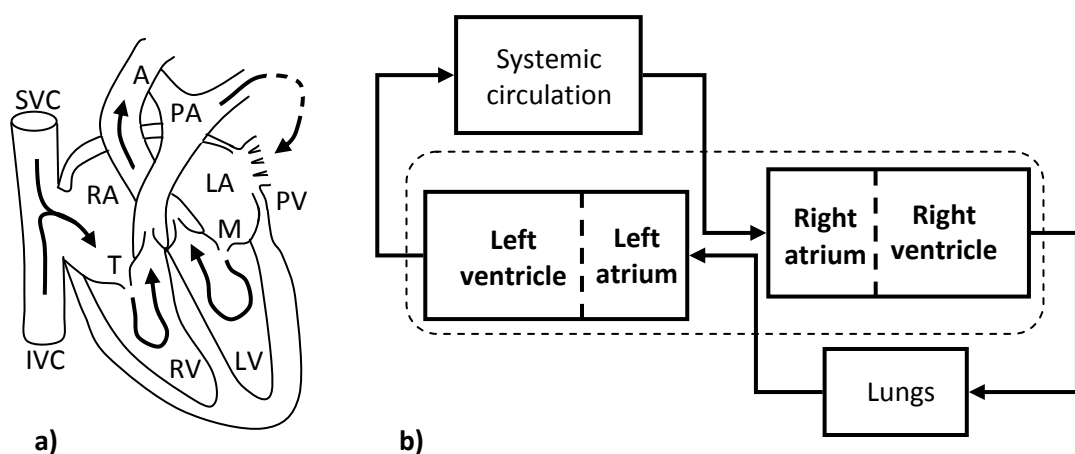


Figure 1: a) The chambers and blood-flow pathways of the heart: RA, LA – left and right atria; RV, LV – left and right ventricles; S – septum; SVC, IVC – superior and inferior vena cavae; T, M – tricuspid and mitral valves; A, PA – aorta and pulmonary artery; PV – pulmonary veins. The aortic valve and pulmonary valve are not labelled, but can be seen at the entrance to the aorta and pulmonary artery, respectively. The arrows indicate direction of blood flow. b) A schematic diagram of the circulatory system. The dashed line indicates the boundaries of the heart.

progresses through the heart, as well as the speed of each cell's response to that stimulus. Both of these facets of coordination are subject to various modulating influences, presenting a complex network of control mechanisms that contribute to cardiac stability and efficiency.

When considering the timing of control signals, it is important to recognise that the atria and ventricles do not contract in synchrony. The atria contract first to fill the ventricles, then the ventricles contract to drive blood around the body. The state of contraction in each chamber is referred to as systole, and relaxation is referred to as diastole. The elastic properties of the collagen in the cardiac muscle cells encourage the ventricles to expand to their resting volume after systole. Once ventricular pressure drops below atrial pressure, the atrioventricular valves open and blood flows freely from the veins through the atria to the ventricles even before atrial systole occurs. In fact, in a patient at rest, atrial systole provides only the last 15-20% of the blood volume that fills the ventricle. It is only at times of increased cardiac output that the atrial boost becomes important^[Levick, 2003].

II.B. Electrical activity in the heart

As described in the introduction, cardiac contractions are initiated by the spread of electrical activity through specialised fibre-like muscle cells, which are called myocytes. Maintaining normal, efficient heart function requires careful control of the progression of this wave of activity through the heart. This section describes the healthy spread of electrical excitation through the myocardium, the collective name for the muscular tissue of the heart. Some common, related pathologies are also introduced to illustrate the vulnerabilities of the system.

Like nerve cells, myocytes maintain an electrical potential difference across their membranes such that they are more negatively charged inside than outside while 'at rest'. A slight deflection of this potential difference in the positive sense, beyond a certain threshold, causes the cell to depolarise; the potential difference becomes slightly positive for a short time before the cell




Figure 2: Timing of the action potential and the resultant mechanical contraction in cardiac myocytes. Adapted from [Levy & Pappano, 2006].

repolarises. The depolarisation-repolarisation pattern is referred to as the action potential. Contraction of the cell occurs in response to depolarisation.

The currents drawn by a depolarising myocyte are sufficient to trip the threshold potential of the adjacent myocytes, and in this manner the signal is passed from one cell to the next. In healthy hearts, the signal does not get passed back to a myocyte that has just fired because it will be briefly refractory (unresponsive). This trait causes the signal to progress only in the forward direction^[Noble et al, 2005]. Action potential duration (APD) is therefore a key feature in determining cardiac stability. Much of this thesis focuses on the measurement and analysis of APDs, their spatiotemporal variations, and the controlling mechanisms.

The signal is initiated in the sino-atrial (SA) node, which is located on the posterior wall of the right atrium. From here it spreads through the right and left atria, provoking atrial systole. The atria are electrically insulated from the ventricles by a ring of tissue known as the annulus fibrosus. The only electrical pathway through this insulating barrier is the atrioventricular (AV) node. The AV node imposes a slight delay on the transmission of the signal to ensure that the atria contract fully before the ventricles contract, as described in the previous section. This delay is caused by the slow conduction velocity and complex arrangement of the fibres in the AV node. From there, the signal is passed on to the bundle of His, which divides into two branches (one for each ventricle) running down the septum. The cells that make up these branches are known as Purkinje fibres and have high conduction velocities. They carry the signal quickly to the heart's apex, the lower tip of the ventricles. From there, the fibres spread up through the ventricles. Contraction begins in the septum and spreads up the ventricles from the apex and from the inner ventricular walls (endocardium) to the outer walls (epicardium), squeezing blood up and out of the arterial valves^{[Levick, 2003], [Noble et al, 2005]}. In the thick wall of the left ventricle, the fibre orientation varies gradually between the endocardium and epicardium, as shown in Figure 4^{[Greenbaum et al, 1981], [Nielsen et al,}

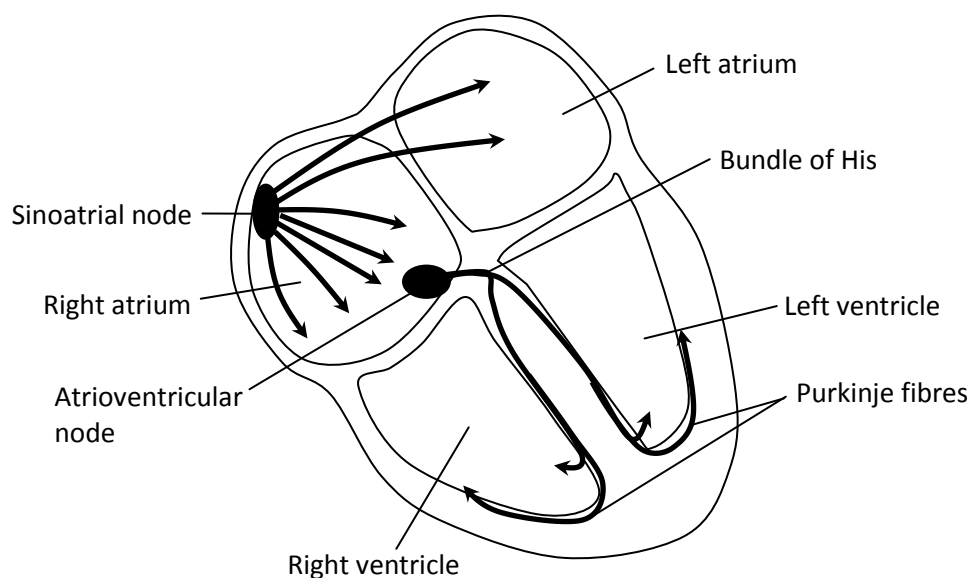


Figure 3: Progression of the electrical signal through a healthy heart. Adapted from [Levy & Pappano, 2006].




Figure 4: Ventricular myocyte fibre orientation, as it varies through the myocardial wall. The main picture shows the ventricles from a left-posterior view, with the upper half of the heart removed for clarity. The inset (bottom-left) shows the plane through which this cross-section was taken. A piece has been removed from the outer wall of the left ventricle, and arrows are drawn to show the approximate fibre orientation at the endocardium, mid-myocardium, and epicardium. Based on data from [Scollan et al, 2000].

1991], [LeGrice et al, 1995], [Scollan et al, 2000]. Endocardial fibres run roughly in the longitudinal (apex-base) direction while spiralling around the longitudinal axis. In the mid-myocardium, the fibres are roughly oriented in the plane transverse to that axis. The epicardial orientation is similar to that of the endocardium, except that the spiralling around the heart occurs in the opposite sense. This overall arrangement causes the left-ventricle to contract with a twisting, wringing motion, which is believed to be desirable for efficiently pumping blood^[Wei, 1997]. An understanding of the subtleties of cardiac fibre orientation is important for the interpretation of invasive electrical recordings; the orientation of the fibres relative to the direction in which excitation spreads significantly influences the morphology of the signal recorded from a nearby electrode, as will be explained in section II.E.

The conduction velocity of the activation wave varies between different types of myocyte. In humans, the fastest conduction occurs in the Purkinje fibres at around 130-170 cm/s^[Kupersmith et al, 1973]. In ventricular myocytes, the conduction velocity is typically around 60-70 cm/s in the longitudinal direction of the fibres^{[Taggart et al, 2000], [Conrath & Opthof, 2006]} and roughly 20 cm/s in the transverse direction^{[Taccardi et al, 1998], [Conrath & Opthof, 2006]}. Much slower conduction occurs in regions such as the AV node.

Arrhythmias

The correct spatiotemporal propagation of electrical activity is crucial to the effective mechanical operation of the heart and, as mentioned previously, disruptions to the normal progression of the electrical signal can be fatal. Such disruptions are known as arrhythmias or, more

accurately but less commonly, as dysrhythmias^[Bennett, 2002]. Due to the chain-reaction nature of the signal transmission, spurious activity in a localised region of the heart can lead to recurring unwanted activity of a more global nature. The following summary of the mechanisms of failure outlines the boundaries within which cardiac behaviour must be maintained, by nervous control or other influences, to ensure survival. There are several types of arrhythmia, each distinguished by the region of the heart that is affected and by the mechanisms involved:

Respiratory sinus arrhythmia:

A common and benign condition, sinus arrhythmia is the regular fluctuation of heart rate with breathing. It is caused by the influence of the autonomic nervous system on the heart. A more detailed description is provided in section II.F.2, and previously unobserved effects on ventricular behaviour were exposed in the experiments described in section V.C

Ectopic beats and escape beats:

As described previously, the heartbeat is normally initiated in the SA node. Other regions of the heart are also capable of initiating beats, but they fire at a slower rate so they are normally pre-empted by the SA node. If the SA node fails to fire in time, another part of the heart will initiate an escape beat. This spontaneous activity, known as automaticity, ensures that such problems with the SA node are survivable. The intrinsic frequency of automaticity is normally lower for cells further along the natural conduction path. Hence, if the SA node fails, the escape beat is more likely to be triggered in the AV node than, say, the ventricles. This hierarchal chain of command is referred to as dominance^[Mangoni & Nargeot, 2008]. It maximises the likelihood that the escape beat will follow the natural conduction path through the ventricles, so that the crucial contraction is well coordinated.

Although it normally contributes to cardiac stability, automaticity may pose a threat in some situations, because any maladjustment of automaticity may give rise to ectopic (out of place) beats. Ectopic beats are essentially escape beats that occur in the absence of any failure of the SA node^[Bennett, 2002]. An ectopic is likely to replace the next normal beat because the affected myocytes will be in their refractory (unresponsive) period when the signal initiated by the SA node arrives. If this refractoriness is not absolute, the subsequent activation wave may be disrupted but not eliminated, leading to more chaotic activity with potentially fatal consequences.

Heart block:

Heart block is the partial or total blockage of the electrical conduction path between the atria and the ventricles, and is further classified by three degrees of severity. First-degree heart block is a slowing of this conduction path, and causes an increased interval between atrial and ventricular systole. Second-degree heart block is the occasional

failure of the signal to reach the ventricles. Third-degree heart block is the consistent failure of the signal to reach the ventricles. In this case, automaticity ensures that the ventricles continue to beat of their own accord at a low frequency, but the atria beat at the normal frequency regulated by the SA node. Cardiac output is severely limited, but can be restored by an artificial pacemaker^[Noble et al, 2005]. The ‘bottle-necking’ of the activation wave at the AV node makes any conduction problems in this region more survivable and easier to diagnose. Similar abnormalities occurring in a small area of the ventricles, for example, would disrupt an activation wave without eliminating it. The undesirable consequences of an abnormally progressing activation wave have been explained in previous scenarios.

Circus:

If a wave of activity somehow returns to part of the myocardium that has already fired but has emerged from its refractory period, the cells will fire again. This phenomenon is known as re-entry, and can lead to a circular pattern of excitation that pre-empts the signal from the SA node. Circus arrhythmias can be treated using drugs that globally prolong the refractory period, such as quinidine and procainamide^[Bennett, 2002]. This simple treatment illustrates the importance of APD and the refractory period in maintaining cardiac stability. While global adjustment of APD can be effective in cases such as this, cases involving abnormal spatial heterogeneities are more complex. Hence it is important to understand the mechanisms that influence APD on a local basis. Novel methods to enable investigations into the nature of such mechanisms are described in section IV, and experimental implementations of these methods have yielded new physiological insights, as described in section V.C

Fibrillation:

Instances of re-entry can cause the heart to enter a state known as fibrillation, in which it repeatedly beats in an uncoordinated manner. Atrial fibrillation is typically survived, albeit with reduced cardiac output. Ventricular fibrillation, however, is fatal unless interrupted within a few minutes. Proneness to sustained ventricular fibrillation is largely dependent on the relative durations of the action potential and diastolic interval (DI, measured from the time of repolarisation to subsequent activation) in each myocyte. This fact illustrates the need for researchers to identify any mechanisms that might destabilise or inappropriately modulate the APD-DI relationship. A mechanism particularly deserving of scrutiny is the autonomic nervous system, which will be introduced in section II.F.

All instances of fibrillation begin with an occurrence of re-entry and can be seen as repeating occurrences of re-entry. Hence, it is also important to understand the specific mechanisms of re-entry. A conceptual outline of the necessary conditions for re-entry to occur is presented in Figure 5.




Figure 5: A mechanism leading to re-entry. Adapted from [Levy & Pappano, 2006].

A bundle of fibres, S, splits into two branches, L and R, which are later connected by C. Panel A shows a ‘healthy’ progression of electrical activity; the signal spreads down branches L and R simultaneously and enters C from each side. Where the two signal-fronts meet, they are both blocked because they reach myocytes that have just entered their refractory period. The signal progresses only in the intended direction, further down L and R. Panel B shows a case in which regions of L and R have, for some reason, not emerged from their refractory period. The consequence is heart block, as discussed previously. In panel C, only one of the two branches is refractory. The signal from the other branch spreads across C and is blocked by the refractory region from the other side. However, because the signal arriving via branch L takes longer to reach this region, there is a chance that it will no longer be refractory, as illustrated in panel D. In this case, the signal from branch L crosses C and travels up R in the reverse direction. If S has also emerged from its refractory period, the signal will continue in the reverse direction. This qualifies as re-entry and may lead to continuous fibrillation. By this reasoning, Levy & Pappano deduce the following two necessary conditions for re-entry, which were first laid out in [Schmitt & Erlanger, 1928]:

1. “At some point in the loop the impulse can pass in one direction but not in the other.”
2. “The effective refractory period of the re-entered region must also be less than the propagation time around the loop.”

The situations described in Figure 5 are useful for understanding the mechanisms for re-entry, but it is important to note that re-entry can also occur when the signal is spreading in three dimensions through the myocardium, not only in a network of one-dimensional fibres.

A common cause of fibrillation is a myocardial infarction, also known as a heart attack. This term refers to the death of an area of myocardium due to ischaemia, an interruption in blood supply. The affected region, the infarct, will no longer transmit the electrical signal, so it disrupts the normal conduction path^{[[alife, 2000](#)]}. This obstruction can create a scenario similar to that shown in Figure 5 above, in which the signal can arrive at a certain region by more than one path, increasing the chances of re-entry.

Such an obstruction is not necessary for re-entry to occur, however. Clayton & Taggart demonstrate, using a computational model, how re-entry can occur in a two-dimensional sheet of myocardium in which one region has a longer refractory period than an adjacent region. Regional differences in APD and the APD-DI relationship exist in healthy cardiac tissue^{[[Clayton & Taggart, 2005](#)]}. Unhealthy differences might arise from any phenomenon that heterogeneously modulates APD, such as localised ischaemia due to a blocked coronary artery, but even healthy differences can lead to re-entry if a sudden rate change occurs. In Figure 6, region R1 has a longer refractory period than R2. Wave-front S2 passes through the region, but S3 is initially blocked because R1 is still refractory. When S3 passes R1, the electrical activity spreads into the region from the other side, so re-entry has occurred. Note that the conditions described by Schmitt and Erlanger are met, although the “loop” is dissimilar to that described in Figure 5.

As demonstrated, spatial variations in electrical properties are of crucial importance to cardiac stability. Hence it is important to be able to measure the properties of myocardium locally. Available techniques for doing so are described in section II.D. The implementation and




Figure 6: Reproduced from [[Clayton & Taggart, 2005](#)]. Permission not required. Computer modelling of re-entry due to regional differences in refractory period, which is longer for R1.

advancement of such techniques, especially to infer nervous control of the heart, form the central topic of this thesis. The work to be described has provided new insights into the control regimes from which undesirable heterogeneities in cardiac properties might arise, and has suggested new methods of investigation to further expose the underlying mechanisms.

Various hypotheses and partially-understood mechanisms exist to be explored. Any local ischemia, or restriction of blood supply, is known to influence myocytes' electrical behaviour in a number of ways^[Shaw & Rudy, 1997]. Furthermore, if the ischemia (or any other insult) causes cell death, the unexcitable dead cells will alter the progression of the activation significantly. The resultant abnormal distribution of repolarisation may interfere with subsequent beats in a manner that has been described previously. It has also been suggested that an individual's emotional state might induce heterogeneous changes in cardiac behaviour via a lateral imbalance of input from the autonomic nervous system^[Lane & Jennings, 1996]. Furthermore, it is known that nervous input to the heart is modulated by respiratory behaviour, as in the benign sinus arrhythmia introduced on page 27. It may be assumed that this influence has manifestations throughout the heart, rather than being confined to the easily observable effect on the SA node, but a detailed understanding of the widespread cardiac effects of respiration has yet to be obtained. The experiments carried out in this project focus on the effects of emotion and respiration on cardiac behaviour. Both of these sources of nervous activity input may have important implications for cardiac stability, particularly when considering that any of the described mechanisms may coexist and form complex interactions. Achieving a more thorough understanding of these interactions could expose previously unknown risk factors related to the 100,000 Sudden Cardiac Deaths in the UK each year^[NHS choices, 2011].

II.C. The Cellular Basis of Cardiac Behaviour

In the previous section, the problem of re-entry was explained in terms of the spatiotemporal relationship between excitation and refractoriness, hence this section describes the cellular behaviour associated with these states.

As mentioned previously, myocytes are the muscle cells of the heart. This section will focus on the ion flows that initiate mechanical contraction, and how these are related to the changes in electrical potential. The variations in electrical properties of myocytes are evident in Figure 7, which denotes the four phases of the action potential for the most important contractile cells, ventricular myocytes. The profiles of the action potential for two non-contractile myocytes are also shown.

The electrical currents in the myocyte are caused by the flow of ions across the cell membrane. The ions in question are, primarily, Sodium (Na^+), Potassium (K^+), and Calcium (Ca^{++}). The movement of these ions is the result of two influences, chemical potential and electrical potential. The electrical potential is the voltage difference measured across the cell membrane, as featured in Figure 7. The positively charged ions are attracted towards areas of more negative electrical potential. Chemical potential refers to the concentration gradient for a particular ion. The random movement of particles tends to result in them spreading into areas that have a lower




Figure 7: Action potential morphologies for a ventricular myocyte (a), a myocyte in the sinoatrial node (b), and a Purkinje fibre (c). The non-contractile myocytes (b and c) exhibit an upward drift in the resting action potential. The values of E marked on the vertical axis are the Nernst potentials for the principal ions involved in the action potential. Representative data from [Levick, 2003].

concentration of those particles. When ions of a single type are allowed to move freely, they will tend towards an equilibrium state in which the electrical potential opposes the chemical potential so that there is no net flow. This equilibrium potential is known as the Nernst potential. The Nernst potentials for the three ions that dominate the action potential are shown in Figure 7. These do not vary significantly because changes in the relative concentrations of ions during the action potential are negligible.

Because ion flows are not solely dictated by electrical potential gradients, theoretical analyses of the electrical fields generated by the myocardium typically model the membrane's behaviour as a current source, rather than a voltage source. The behaviour of the current source can be inferred from the familiar action potential morphology in order to predict the resultant electric field. Such analyses are important in relating recordings of cardiac electrical activity to the underlying electrophysiological activity, as will be demonstrated in section II.E.

Ions are not always free to move across the myocyte membrane because they must pass through ion-specific gates that open and close in response to the membrane potential. When a channel opens, the resultant ion flow causes the membrane potential to shift towards the Nernst potential for that ion. There are also exchangers that work to keep K^+ in the cell and Na^+ and Ca^{++} out. The exchangers ensure that myocytes always have a relatively high concentration of K^+ inside with respect to outside and a relatively low concentration of Na^+ and Ca^{++} [Philipson, 1997].

Figure 8 indicates the main ion flows at each phase of the action potential. To infer the behaviour of cardiac control mechanisms from electrophysiological behaviour, it is important to understand the separate contributions of the various ion flows to action potential morphology; different controlling influences affect different ion channels selectively. Hence the probable mechanism responsible for any change in cardiac behaviour can often be identified by observing the change in the action potential.

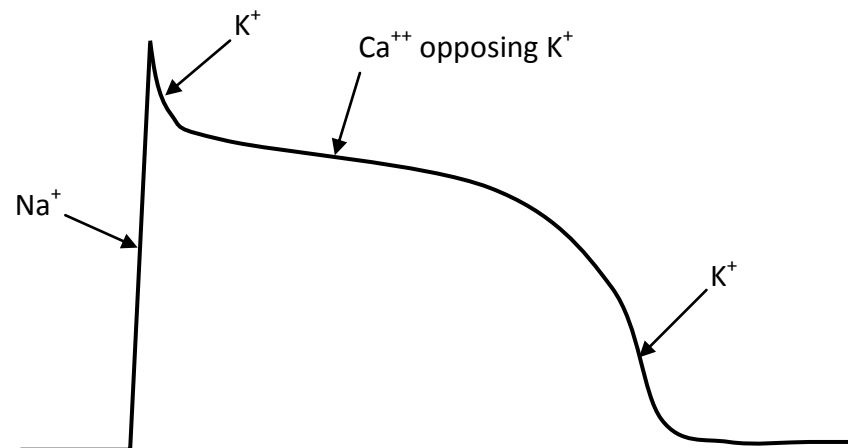


Figure 8: The ion currents primarily responsible for each phase of the action potential.

If some outside influence, such as electrical activity in an adjacent myocyte, depolarises the resting myocyte beyond its threshold potential, roughly -60 mV to -65 mV, Na^+ channels open to allow a sudden influx of Na^+ . As a result, the cell quickly depolarises further to a membrane potential of $+20$ mV to $+30$ mV. This is phase 0 of the action potential, as labelled in Figure 7. The fast sodium channels are said to be time-dependent as well as voltage dependent; they close soon after opening, regardless of membrane potential. They cannot be re-opened until the myocyte has repolarised beyond about -70 mV, and the myocyte is refractory as long as the Na^+ channels are unable to be opened^[Weiss, 1997]. An understanding of this relationship between the action potential and the refractory period is essential to the development of appropriate experimental measures of refractory period, since the change in extracellular potential caused by the action potential offers a convenient window on cardiac electrical activity. Section II.D will elaborate on this point.

The change in membrane potential during phase 0 leads to a transient outward flow of K^+ , i_{to} , which brings the potential quickly back towards negative values (phase 1) when the fast Na^+ channels close. An additional class of K^+ channels briefly opens to accelerate this flow. In fact, there are at least ten different categories of K^+ channel that open and close at different times in response to different stimuli^[Weiss, 1997]. These channels will not be described in detail, but their variety is relevant to this thesis in that it illustrates the complexity of mechanisms available to be modulated to adjust a myocyte's behaviour.

The action potential of a cardiac myocyte differs from that of nerve cells in that it has a plateau phase (phase 2). This is caused by a sustained influx of Ca^{++} that counterbalances the efflux of K^+ . During the plateau phase, the number of open Ca^{++} channels gradually decreases, and a subclass of K^+ channel, the 'delayed rectifier' channel, correspondingly reopens. These effects produce the downward curvature of the plateau phase. The action potential ultimately terminates when the inward rectifier K^+ channels suddenly open (phase 3), allowing full repolarisation. Sympathetic nerve stimulation is known to increase K^+ flow, thereby shortening the action potential in order to allow an increased heart rate^[Levy & Martin, 1979]. The opposite effect can be used to provide added security against spurious activity during slower heart rates. Nervous input can also modulate APD independent of heart rate, with serious implications for cardiac stability.

It was noted earlier that for the Na^+ channels to open and initiate an action potential, something must cause the myocyte to depolarise beyond its threshold potential. For contractile myocytes, this stimulation is normally provided by an adjacent myocyte. Non-contractile myocytes, which are found in the SA node, AV node, bundle of His, and Purkinje fibres, do not necessarily need to be triggered by an outside influence, as explained in the description of automaticity on page 27. They have a special class of Na^+ channel that opens in response to very negative membrane potentials, allowing a small inward current. They also allow a slow influx of Ca^{++} and a reduction in K^+ conductance during the resting phase^[Weiss, 1997]. The resting potential of the myocyte is therefore unstable. It steadily becomes more positive until it reaches its threshold potential, at which point Na^+ channels open, firing an action potential. It was noted previously that this automaticity prevents upstream blockages of electrical activity from causing total heart failure. The upward drift of the resting potential can be observed in Figure 7.b for a Purkinje fibre and an SA node myocyte. It is the rate of this upward drift that determines the hierarchy of dominance described on page 27. The opening and closing of the ion channels that control this drift rate are subject to modulation from the nervous system. In the SA node for example, such modulation is responsible for any changes in heart rate. Maladjustment of the drift rates throughout the heart, relative to each other, may threaten cardiac stability by increasing the likelihood of ectopics and escape beats (see page 27).

The wide variety of ion transfers across the myocyte membrane form an array of low-level control mechanisms that combine to help regulate the timing and force of cardiac contractions. This control is manifested within single cells as well as in the coordination between cells to maintain stable cardiac activity. Furthermore, these mechanisms can be influenced by nervous and hormonal input to manifest high-level control of the heart. Hence cardiac electrophysiological behaviour can be taken as an indicator of activity in the nervous system. The timings of the ion transfers can be measured in terms of the effects on the electrical field in and around the heart, as explained in the following sections. Numerous existing methods exploit this principle to monitor the behaviour of cardiac cells and their controlling influences, with various levels of detail ranging from that of the non-invasive electrocardiogram to that of delicate *in vitro* studies using glass microneedles.

II.D. Recording the Heart's Electrical Activity

II.D.1. The Electrocardiogram

II.D.1.a. The genesis of the electrocardiogram

When a large collection of myocytes is activated at roughly the same time, as normally occurs in the atria and ventricles, the potential difference across the heart is significant enough to be measured simply by placing electrodes on the skin. This recording is known as the electrocardiogram (ECG), and is extremely useful in diagnosing various cardiac abnormalities. As a non-invasive measure of cardiac electrophysiology, it also offers a convenient window on nervous input to the heart. This section explains how the signal arises from the collective behaviour of

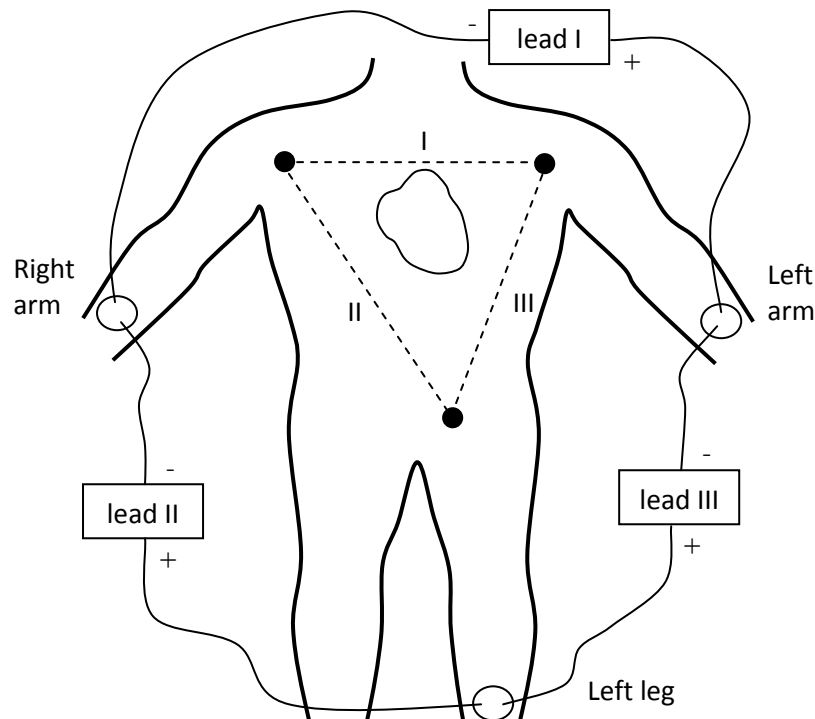


Figure 9: The electrode configuration used for the three main leads of a standard ECG. The dashed lines show Einthoven's triangle, which indicates the effective direction of each lead's polarity.

cardiac cells. Section III describes the design of signal processing techniques to analyse the ECG. An understanding of the signal's genesis is crucial to ensure that any newly developed measurements offer meaningful descriptions of the signal's source.

ECGs are normally produced using the standard electrode placements shown in Figure 9. Although the electrodes are placed at the wrists and ankle of the patient, the limbs can be seen as a direct, passive electrical connection to the torso. These electrodes can therefore be taken as an indicator of the potential at the nearest corner of the torso, and the lines between those points make up Einthoven's triangle^[Levick, 2003]. The polarity of the leads is chosen so that the main ECG events will produce positive deflections in each lead of Einthoven's triangle for a healthy patient.

Geselowitz^{[Geselowitz, 1989],[Geselowitz, 1967]} shows theoretically that the ion flows around a depolarising/repolarising cardiac myocyte amount to an electrical current source. This current is conducted through the body and thus gives rise to a small difference in the electrical potential at the skin surface on opposite sides of the heart. When large numbers of myocytes act approximately in unison, the change in this potential difference is large enough to be measured in a lead connected across the two points. Various researchers have invested significant effort in computational models to show that these mechanisms give rise to the commonly observed ECG^{[Wei, 1997],[Geselowitz, 1989]}. The nature of the ECG interpretation involved in this thesis does not necessitate such detailed theoretical analysis. However, the theory will be introduced in section II.E in the context of unipolar electrograms, another form of cardiac electrical recording employed in this project.

For a less rigorous but more intuitive understanding of how the ECG signal arises, one can consider a single cardiac dipole, which acts across the boundary between polarised and depolarised regions. This dipole can be represented by a vector that changes in magnitude and direction as the

wave of excitation progresses through the heart. Figure 10 shows how the dipole vector changes during depolarisation of the ventricles, which account for the main bulk of the myocardium. The overall change in the cardiac dipole during ventricular depolarisation is a rotation in the anti-clockwise direction (when facing the patient) and a gradual increase then decrease in magnitude.

Figure 10 also shows the influence of the perspective of different leads on the ECG recordings of those leads. The magnitude of the potential difference recorded in a particular lead depends not only on the magnitude of the cardiac dipole, but also on the alignment of that dipole with respect to that lead. A lead is most responsive when the dipole runs parallel to its line of perspective. When the two are perpendicular, the lead will not detect any potential differences. As a result the shape, magnitude, and timing of the observed QRS complex, which represents ventricular depolarisation, depends on the lead being monitored. Comparing the manifestations of an event in each of the leads allows a spatial awareness of the progression of electrical activity through the heart at that time. For example, Frank proposes a mathematical transform that may be applied to a configuration of seven electrodes to yield an approximation of the dipole vector in three orthogonal dimensions^[Frank, 1956]. This approach is known as Vector Cardiography.

Repolarisation of the ventricles is also detectable in the ECG recording, and is represented by the T-wave. It may seem strange, however, that the T-wave and the QRS complex both produce upward deflections, as seen in Figure 11, considering they represent opposite effects. The generally accepted explanation for this is that repolarisation occurs in the reverse direction to depolarisation, so the resultant dipole is of a similar sense; the areas of ventricular myocardium that tend to be excited latest also tend to have the shortest action potentials, so the later they depolarise, the sooner they repolarise^[Levick, 2003]. Although the latter statement is true, recent studies challenge the accuracy




Figure 10: *Left* - The variation in the cardiac dipole vector during ventricular depolarisation.

Diagrams of the ventricles show the spread of the depolarised region (grey, negative) at three different stages of the process, the timings of which are indicated relative to the start of ventricular depolarisation. *Right* - The morphology of the QRS complex in the three main ECG leads can be inferred based on the magnitude of the cardiac dipole vector and its alignment with the lead vector at each point in time. Data on progression of activation from [Levick, 2003].




Figure 11: One beat of a healthy ECG trace.

of the interpretation that the differences in action potential duration cause repolarisation to spread in the opposite direction to depolarisation^{[Franz et al, 1987], [Conrath & Opthof, 2006], [Hanson et al, 2009]}. When electrodes are positioned in direct contact with the myocardium to achieve greater spatial specificity, it is typically found that repolarisation spreads roughly in the same sense as depolarisation for the transmural direction but in the opposite sense for the apico-basal (longitudinal) direction. The importance of invasive recordings to overcome the limited spatial specificity of the ECG will be revisited at the end of this section.

As a non-invasive measure of the spatiotemporal distribution of electrical activity in the heart, the ECG has proven to be useful in various applications. In the clinical setting, it is already a well-established diagnostic tool, and can be used to identify a broad range of maladies. These include physical injury to the heart, changes in gross shape or size, defects in the conduction of electrical activity, abnormal rates or rhythms, and changes in the ionic activity that generates the action potential or in the concentrations of the chemicals that influence those ion flows^[Geselowitz, 1989]. Short-term variations in ECG features can be taken as evidence of changes in nervous activity or other controlling influences, making the ECG useful in broader physiological studies, beyond cardiology. From this project, the experiments discussed in sections V.A and V.B fall into this category.

II.D.1.b. Typical ECG features

An important aspect of modern ECG analysis is the automatic identification of the various features that typically appear in the signal. Algorithms to perform this task are a crucial component of the tools developed in section III. The development of such algorithms requires an awareness of not just the standard ECG features, but also the variability in those features, so that any algorithms implemented are applicable to a wide range of signal morphologies. The following paragraphs will identify the most important features of the healthy ECG and give some examples of common deviations from the norm and their causes, although this is certainly not to be taken as an exhaustive description of the range of abnormalities to be found.

Figure 11 shows a typical ECG recording acquired from an experiment that will be discussed later in this report. The principle features of a single beat are listed below:

P-wave:

The P-wave represents atrial depolarisation. The ECG does not normally show evidence of the atria repolarising because this activity is more asynchronous than the repolarisation of the ventricles^[Debbas et al, 1999].

PR interval:

A healthy ECG trace should show an isoelectric segment between the P wave and the QRS complex. Known as the PR interval, this feature corresponds to the conduction delay imposed by the AV node. In healthy subjects, it should last no longer than 0.22 seconds, although longer intervals will occur in cases of first degree heart block^[Fleming, 1980]. In ECG processing tasks, it is important to be aware of the full range of morphologies that the P wave and PR segment might take, because the voltage measured during the PR interval is often taken as an isoelectric reference^[Martinez & Olmos, 2005]. Although this method has often proved useful, it is important to note that the implicit assumption that no transmembrane or extracellular currents flow at this time is false; the notion that transmembrane voltage is constant and spatially homogeneous during the plateau phase of an action potential is only an approximation.

QRS complex:

Some aspects of the QRS complex have already been discussed, and Figure 10 demonstrates how the shape of the complex varies between leads. The normal magnitude of the complex varies from patient to patient, but if the relative magnitudes and durations of the Q, R and S waves are not as expected for a particular lead, the cause may be a myocardial infarction. The damaged region alters the normal direction in which the wave of excitation progresses through the ventricles, so the direction and perhaps also the magnitude of the cardiac dipole changes^[Fleming, 1980].

ST segment:

The ST segment occurs during the plateau phase of the ventricular action potentials. It is normally isoelectric, but in some cases it may be shifted upwards or downwards from the baseline. This effect is often attributed to ‘injury currents’. Myocytes in a damaged or ischemic region do not generate the same membrane potentials as healthy myocytes, and the local differences in electrical potential cause anomalous currents to flow, which influence the ECG readings^[Samson & Scher, 1960]. The abnormal electrical behaviour disrupts the normal progression of activation in and around the ischemic region, posing a threat to cardiac stability.

T-wave:

The manner in which the ventricles repolarise is evidenced in the T-wave and is dependent on a variety of influences. In fact, it is not uncommon to see day-to-day variations in the T-wave shape for a healthy patient. Changes in the overall timing of the T-wave with respect to the preceding QRS complex indicate global modulation of ventricular repolarisation characteristics.

Typically, a healthy T-wave is a deflection in the same direction as the QRS complex, and is asymmetrical in that the slope is steeper on the right side than the left. Subtle changes in T-wave shape indicate heterogeneous modulation of repolarisation properties. Changes in T-wave timing or shape occurring over a time-scale of seconds or minutes can serve as measurable evidence of nervous input and hormonal influences on cardiac repolarisation. Inverted or abnormally shaped T-waves may be caused by infarcts or localised ischaemia altering the progression of the cardiac dipole. Also, T-wave inversion and depression of the ST segment are both commonly observed in patients in an anxious state^[Fleming, 1980]. This observation offers a basic example of how nervous activity generated by psychological behaviour can induce measurable changes in cardiac electrophysiology.

II.D.1.c. Limitations of the ECG

As outlined in the preceding paragraphs, the ECG can be used to identify a wide range of pathologies and to characterise normal cardiac electrical activity. It is particularly useful as an investigative tool because the technique is non-invasive. However, the approach is clearly limited in that it cannot give a detailed impression of the extent to which specific regions of myocardium contribute to an observed deflection. To gain this level of insight, electrodes must be placed in close proximity or direct contact with the myocardium. It was mentioned earlier that studies of this nature reveal subtle aspects of cardiac conduction and repolarisation patterns that are not discernible in the ECG^{[Franz et al, 1987], [Taccardi et al, 2005], [Hanson et al, 2009]}. A further limitation of ECGs, more pertinent to this thesis, is rooted in the fact that nervous control of the heart is anatomically and functionally heterogeneous^[Pauza et al, 2000]. This heterogeneous control is necessary to maintain the overall stability and efficiency of cardiac behaviour in the face of varying physiological scenarios. It follows that a more complete impression of the nature of nervous control of the heart would be achieved from recordings that enable the analyst to distinguish between activity from different regions of myocardium with greater spatial specificity than is available from the ECG.

Magnetocardiography^[Koch, 2004] and Body Surface Potential Mapping^[Messinger-Rapport & Rudy, 1990] are two non-invasive techniques that provide greater spatial detail than the ECG by using a large number of electrodes or, in the case of magnetocardiography, magnetic field sensors positioned around the body. However, both techniques require accurate models of the impedance properties of the tissue throughout the torso. Uncertainties in these models limit the reliability of those more advanced non-invasive methods. Hence invasive measurements are invaluable tools for accurate observations of localised cardiac electrical behaviour. The following section summarises the invasive techniques available.

II.D.2. Invasive recording techniques

Several techniques are available to measure localised cardiac electrical activity. The following sections characterise the most common methods. Each method is considered with respect to the following three criteria, which characterise an ideal method of recording cardiac

electrical activity to infer nervous control of the human heart: (1) suitability for human *in vivo* studies, (2) ability to characterise local repolarisation as well as depolarisation, and (3) spatial specificity. Through this reasoning, it will be shown that the recordings known as unipolar electrograms have the potential for much wider use in exploring nervous control of the heart in humans. The development of novel signal processing techniques to exploit this potential is described in section IV.

- **Suitability for human *in vivo* studies:** In animal studies, artificial stimulation may be supplied to nerves leading to the heart in order to broadly characterise their influence. However, studies of this kind have limited potential to improve our understanding of nervous control of the human heart in a way that acknowledges the link between psychological factors and the development of fatal arrhythmia^[Taggart et al, 2011]. To probe these mechanisms, one must take into account the nature of normal, physiologically generated nervous inputs. Such experimentation requires the entire central and autonomic nervous system to be intact, which implies that *in vivo* experimentation is essential. Human *in vivo* experiments are particularly desirable not only because they are more clinically significant than animal studies, but also because the high level of communication and cooperation between the experimenter and the subject (if conscious) makes it possible to induce nervous inputs to the heart through natural pathways using psychological stimuli.
- **Ability to characterise local repolarisation as well as depolarisation:** The majority of recent studies that invasively explore cardiac electrical activity and the genesis of arrhythmias *in vivo* give greater attention to the spatiotemporal distribution of depolarisation than that of repolarisation. It generates a more distinct artefact in the recording, and any blockage of the spread of depolarisation by a refractory region can be observed without directly measuring repolarisation behaviour. However, the dynamic interaction between repolarisation and depolarisation is the most fundamental aspect of arrhythmia development; to properly understand how nervous modulation of these behaviours might influence the risk of arrhythmias developing, one must be able to characterise the modulation of repolarisation properties even when the spread of depolarisation has not yet been disrupted. Therefore any investigation of this nature requires a system of recording and measurement that directly characterises local repolarisation properties.

An additional advantage of studying repolarisation properties is that they are more directly attributable to the local region of myocardium, whereas the timing of depolarisation is dependent on varying conduction velocities along the entire path that the activation wave has travelled. The three-dimensional complexity of this path hinders accurate determination of local conduction properties from a sparse set of depolarisation times,

hence repolarisation is a more appropriate indicator of any heterogeneous modulation of cardiac electrophysiology.

- **Spatial specificity:** It has been explained previously that the nervous system innervates the heart heterogeneously. Not only is it essential for experiments to involve physiologically natural nervous activity, the investigator must consider the coordination of nervous input and its effects across the vast network of cardiac nerve endings. When considering the arrhythmic potential of a nervous stimulus, one must be able to consider the spatiotemporal distribution of depolarisation and repolarisation with spatial resolution that is considerably finer than the size of a single chamber of the heart, since various experiments and modelling studies have shown that arrhythmias that eventually spread across the whole heart can originate from relatively small areas of myocardium.

Naturally, the insight that might be gained from a set of recordings improves as the spatial resolution becomes finer and the number of electrodes simultaneously deployed increases, but the practicality of the experiment correspondingly decreases. Selecting a recording method to use in an invasive study requires the experimenter to find the appropriate compromise between spatial resolution and the practicality of implementation. The following sections describe various techniques for recording localised cardiac electrical activity, arranged approximately in ascending order of practicality and descending order of spatial specificity.

II.D.2.a. Transmembrane Action Potentials

It is possible to directly record the electrical activity of a single cardiac myocyte by penetrating the cell with a glass microelectrode^{[Draper & Weidmann, 1951], [Ogden, 1994]}. However, this approach is not suitable for *in vivo* experiments because of the delicate placement required and the susceptibility of the electrodes to breaking. Even when this risk is minimised by careful design of the electrode, it is difficult to maintain stable recordings^[Steinhaus, 1989]. These recordings are therefore not useful for studying nervous input to the heart because the tissue to be experimented on must be excised. However, the availability of this accurate but awkward technique is crucial for the validation of more practical methods.

II.D.2.b. Optical Action Potentials

Various chemical compounds have been discovered that bind to the membranes of cardiac myocytes and exhibit changing fluorescent properties that are voltage-sensitive^{[Cohen et al, 1974], [Salama et al, 1987]}. These dyes effectively reproduce the time-course of the membrane voltage for the cells they bind to, and can be used in combination with image processing techniques to achieve a useful measure of cardiac electrical activity. Because it allows the behaviour of a large area of tissue to be captured with high spatial resolution, this approach is particularly useful when studying the pattern

by which depolarisation and repolarisation spread across the heart. However, despite efforts to extend the method's usefulness, it remains limited to observing activity near the surface of the myocardium^[Efimov et al, 2004]. Furthermore, irreversible aspects of the procedure prohibit its application to human *in situ* studies.

II.D.2.c. Monophasic Action Potentials

Because bioelectric signals in general are relatively weak, *in vivo* recordings require that a reference electrode be positioned inside or in contact with the subject's body; using an external 'ground' reference yields signals that are unusable due to low signal-to-noise ratios^[Plonsey & Barr, 2007]. Unfortunately, locating the reference electrode on the subject complicates the nature of the measurement. As explained in section II.D.1, the ECG signal is measurable throughout the body. Hence, no stable reference electrode is available. It is not possible to acquire an *in vivo* electrical recording that purely reflects activity at a single site, and this makes it difficult to achieve fine spatial specificity.

The recordings known as Monophasic Action Potentials (MAPs) offer the finest level of spatial specificity that is currently practically achievable for *in vivo* measurements^{[Franz, 1991], [Franz, 1999]}. In its modern form, the MAP technique involves pressing an electrode against the myocardium while a reference electrode rests on or near the surface (see Figure 12). The pressure exerted via the first electrode causes the underlying myocytes to partially depolarise and become inactive. Myocytes under the reference electrode continue to function as normal, and each action potential drives extracellular ionic currents between the active and inactive regions. The potential difference measured between the two electrodes closely follows the shape of the action potentials in the region, as validated *in vitro* by comparison with transmembrane microelectrode recordings (see Figure 13)^{[Franz et al, 1986], [Ino et al, 1988]}. The term 'monophasic' refers to the fact that these signals resemble the single up-then-down deflection of the action potentials themselves. It distinguishes MAPs from the more complex morphologies discussed in the next section.

As labelled in Figure 12, the electrode at the deactivated region is normally referred to as the "different" or "exploring" electrode, while the other is referred to as the "indifferent" or "reference" electrode. This nomenclature may be inappropriate in light of some recent debate;^[Kadish, 2004] a point of contention exists as to whether an MAP recording should be taken as a representation of the behaviour below the "reference" electrode or the "exploring" electrode. After the introduction and widespread adoption of the Franz configuration for MAP recording,^[Franz, 1983] the predominant view was that the MAP was generated by the currents flowing between the inactivated region and the active cells at the circumference of that region, implying that the observed change in potential occurs at the "exploring" electrode. It should be noted that although this view has been predominant, it does not necessarily amount to a consensus because the literature on the subject is dominated by Franz's own research group. Kondo and colleagues recently challenged Franz's interpretation; they argued that the MAP is dominated by activity at the "reference"

electrode^[Kondo et al, 2004]. This claim was supported by *in vitro* experiments using a configuration like that in the left panel of Figure 12. Interventions to change the action potential duration (APD) near the “reference” electrode caused matching changes in the MAP, whereas similar interventions in the region below the “exploring” electrode caused no significant changes in the MAP. Franz argued that a number of methodological flaws corrupted the experiments by Kondo and colleagues^[Franz, 2005]. Many of these claims were refuted in ^[Nesterenko et al, 2005]. The remainder were unquantifiable in terms of their effect on the observed MAP. Subsequent contributions in the form of analytic^[Vigmond, 2005] and computational^[Colli-Franzone et al, 2007a] modelling suggest that, when the Franz configuration is not used, the activity at the “reference” electrode dominates the MAP. As Franz asserts^[Franz, 1999], some extracellular ‘injury current’ undoubtedly flows between the active and inactive regions of myocardium. However, the contribution of this current to the MAP via the “exploring” electrode is relatively small because an opposing current can be assumed to arise in the intracellular domain at the same site.

Conveniently, this re-interpretation of the origins of the MAP does not invalidate previous studies that have used the Franz electrode configuration because the region of active myocardium closest to the “reference” electrode is the same as that which surrounds the inactive region. Hence MAPs recorded by this method can be taken as an accurate representation of the surrounding active




Figure 12: Left) A schematic of the fundamental arrangement for contact MAP recording. Right) The Franz configuration for MAP recordings.

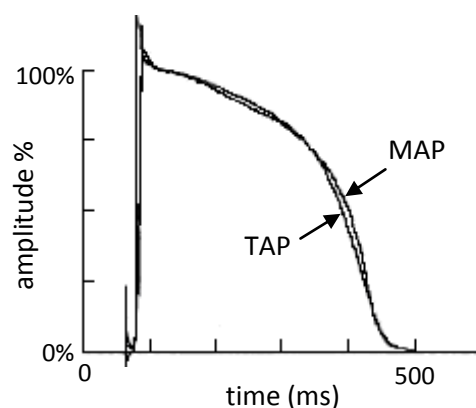


Figure 13: Adapted from ^[Franz et al, 1986] by permission of Oxford University Press. A comparison of MAP and transmembrane action potential (TAP) recorded simultaneously from adjacent sites. Their magnitudes were scaled to match at the point marked by the open arrow. The scaled MAP gives a close approximation of the time-course of the IAP.

tissue. However, the need to apply delicate pressure limits the number of simultaneous recordings that can be practically acquired.

II.D.2.d. Electrograms

It is not actually necessary to inactivate a region of tissue to obtain useful *in vivo* recordings from the heart with spatial specificity. Any configuration of electrodes in or around the beating heart will present some change in potential difference corresponding to the depolarisation-repolarisation sequence. These signals do not resemble action potentials in the way that the MAP does (see Figure 14 for an example). However, by incorporating *a priori* knowledge of the electrode configuration and some basic assumptions regarding the spatio-temporal distribution of electrical activity, one can infer the activity of a region of myocardium with reasonable spatial specificity. Recordings of this general nature are referred to as electrograms in order to distinguish them from the ECG. They have been used for almost as long as ECGs, originally through *in vitro* studies of isolated animal hearts and more recently as a tool for guiding surgical procedures [Kimber et al, 1996], [Kusumoto, 1999]. Electrogram recordings typically take one of two forms: unipolar or bipolar.

Unipolar electrograms:

Unipolar electrograms (UEGs) are measured as the difference in electrical potential between an exploring electrode, positioned adjacent to or even embedded in the area of myocardium in question, and a reference electrode positioned away from the myocardium, usually outside the heart. The name is used to distinguish UEGs from bipolar electrograms (BEGs), which are obtained by placing the second electrode in close proximity to the first. The main deflections commonly observed in a UEG strongly resemble those in ECGs, so they are often referred to by the same lettering system; a QRS complex reflects depolarisation and a T-wave reflects repolarisation. There are important differences, however, in the way that these features are interpreted and the nature of the conclusions that can be drawn from them.

During phase 0 of the myocyte action potential (Figure 14), the flow of Na^+ ions into the cell results in a decrease in the electrical potential outside the cell. When a group of myocytes in close proximity to the exploring electrode depolarise at approximately the same time, a substantial negative deflection appears in the UEG. Hence, the nominal time of local activation at a particular electrode is taken as the point in the activation wave with the greatest downward slope [Durrer et al, 1961]. This point will be referred to in this thesis as \dot{V}_{\min} . Similarly, the rapid change in transmembrane voltage during phase 3 of the action potential leads to an increase in extracellular potential at a nearby exploring electrode. Hence, the nominal time of local repolarisation is commonly taken as the point in the T-wave with the maximum upward slope (T_{up}), a convention referred to as the Wyatt method [Wyatt, 1980]. By this method, T_{up} shows good correlation with t_R , the point of maximum downward slope in the action potential [Haws & Lux, 1990]. Due to key physiological differences between




Figure 14: The relationship between the Monophasic Action Potential (MAP) and typical unipolar electrograms (UEGs) taken from the same region of myocardium. Commonly used repolarisation indices, explained in the text, are marked on the MAP and on the UEG for positive (solid) and negative (dashed) T-waves.

the mechanisms of activation and repolarisation, the interpretation of local recovery time measured by the Wyatt method is much less straightforward than that of the local activation time. The transmembrane voltage of any myocyte changes far more slowly during repolarisation, and the junctions between neighbouring cells introduce have a synchronising effect during this process, so the surrounding gradients in electrical potential are smaller than during depolarisation. Consequently, the rates-of-change in voltage measured by a nearby electrode are less predictable and harder to detect, making repolarisation measurements vulnerable to distortion of the UEG trace by artefacts of activity in remote regions.

Bipolar electrograms:

An alternative to the UEG is the bipolar electrogram (BEG), measured as the potential difference between two electrodes positioned a few millimetres apart on the myocardium. The primary advantage offered by the BEG over the UEG is that it is far less sensitive to remote activity; the close proximity of the two electrodes implies that any electric field generated by a distant source will be similar at both sites and will therefore be approximately cancelled out when the difference is measured. For this reason, BEGs are widely used for mapping the progression of activation wavefronts. However, deflections in the BEG show unpredictable morphology. They are also small in magnitude when compared with the UEG, especially during repolarisation, hence BEGs are only useful in studies confined to the activation phase. The unpredictability of BEG morphology stems from its dependence on the orientation of the local dipole with respect to the line between the two electrodes; a wavefront progressing perpendicular to this line may produce no discernible deflection^[Kimber et al, 1996]. The largest BEG deflections tend to occur when the local action potentials

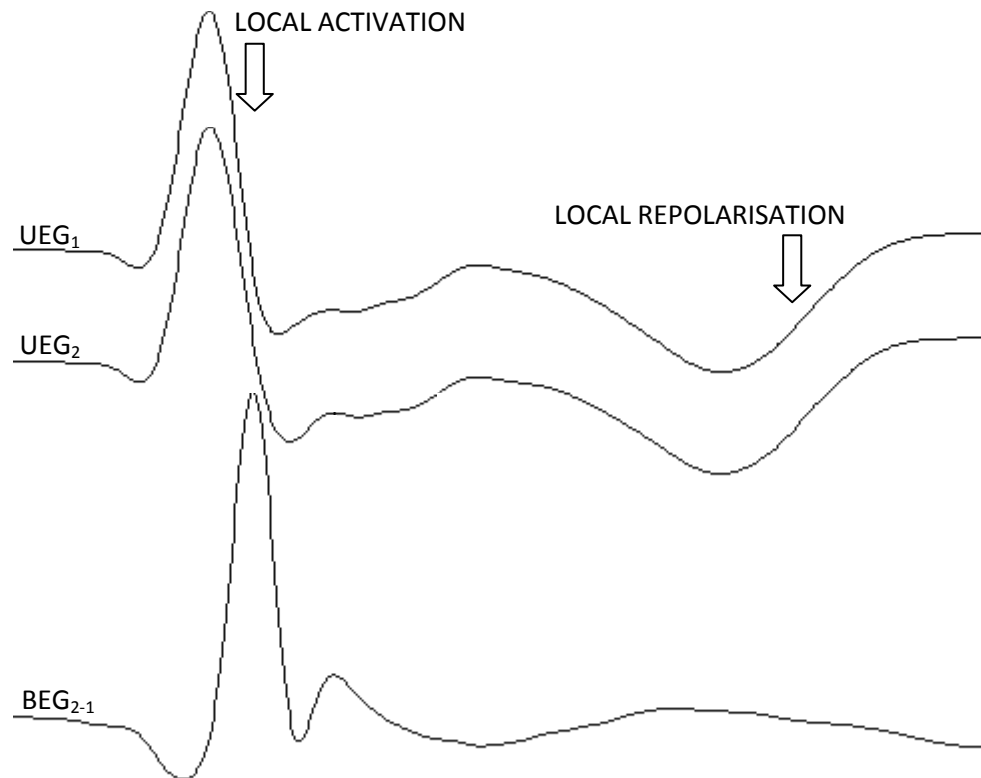


Figure 15: Two unipolar electrograms (UEGs), recorded simultaneously from adjacent sites, and a bipolar electrogram (BEG), which was calculated as the difference between the two unipolar electrograms. The deflection produced in the BEG by local activation is distinct and coincides with the downslope of the UEGs, as expected. The deflection produced by repolarisation is not distinct in the BEG, making these recordings unsuitable for repolarisation studies.

are changing rapidly, as during phase 0 and phase 3; a large potential difference between two close points on the myocardium can only occur if the slopes of the local action potentials are large relative to the speed at which the wavefront progresses. One might therefore expect two large deflections in the BEG, although either may be polyphasic. The local activation and repolarisation (t_R) times are taken as the points with the largest absolute magnitude on each of these deflections, regardless of polarity. Where such *a priori* knowledge is available, electrode pairs should be oriented in the direction along which activation is expected to progress, to maximise the likelihood of recording substantial deflections.

II.E. Unipolar electrograms

The UEG has proven to be a useful tool in characterising cardiac repolarisation properties, and modulations in those properties provide insight into the mechanisms controlling the heart's behaviour^{[Millar et al, 1985], [Haws & Lux, 1990], [Coronel et al, 2006], [Colli Franzone et al, 2007b], [Scacchi et al, 2009]}. Section V.C describes human *in vivo* experiments conducted to explore the nature of nervous control of the heart. Due to the advantages offered in repolarisation measurements, UEGs were chosen in favour of BEGs for these experiments. Novel signal processing techniques were developed to enable the analysis of dynamic changes in repolarisation behaviour, as described in section IV. To assess the validity of these new methods, a thorough understanding of the genesis of the UEG is required. To that end, this section offers a more detailed interpretation of UEG signals and how they relate to local cardiac electrical behaviour.

II.E.1. Debate concerning the reliability of UEG indices of local repolarisation

A confounding factor in the identification of local recovery times from UEGs is that the point t_R on the action potential (see Figure 13, page 43) does not coincide with any discrete event at the cellular level. Its appeal as a measurement of cellular behaviour is that it produces a distinct artefact that broadly coincides with phase 3 of the action potential. In assessing the suitability of UEG measurements as a representation of cardiac behaviour, one must specify which aspect of cardiac behaviour is of interest. In some studies, the UEG is used to monitor the electrophysiological properties of the underlying region. In other cases, it is used to measure the overall behaviour of that region and its impact on the global behaviour of the heart.

When considering a heart's arrhythmic potential, the predominant motivation for identifying a local recovery time is to mark out the transition of the local myocardium from the refractory state to the excitable state. In actual fact, this transition is not instantaneous. By the time the action potential has repolarised to around -50 mV, some of the fast Na^+ channels responsible for depolarisation have already returned to an excitable state^[Weiss, 1997]. The cell as a whole is thus partially excitable. This state is referred to as the relative refractory period to distinguish it from the preceding absolute refractory period. The rate at which Na^+ channels return to excitability is voltage- and time-dependent,^{[Goldman, 1995], [Roden et al, 2002]} so the degree of refractoriness in a cell at a given instant is dependent on the preceding time-course of repolarisation rather than on the timing of an instantaneous event. The ambiguity of the relative refractory period seems to defy the use of a single point to represent the end of the refractory period. However, phase 3 of the action potential is sufficiently sudden for such a simplification to be viable, especially in ventricular myocytes. In dogs, activation-recovery intervals (ARIs) calculated as the time between \dot{V}_{\min} and T_{sp} (Figure 14, page 45) have shown good correlation with the effective refractory period (ERP), determined experimentally as the shortest period after which it is possible to induce another depolarisation using an artificial stimulus at the site^[Millar et al, 1985]. It should be emphasised, however, that the study did not

show ARI to be a direct approximation of ERP; instead, there was a distinct tendency for ARI to underestimate ERP, particularly in positive T-waves, with differences typically in the range of 0-15ms.

This observation gave rise to the argument that a later point in the UEG might give a better estimate of the refractory period. For example, some studies have compared ARIs with the interval between activation and *MAP90* (Figure 14, page 45), the time at which the action potential has achieved 90% repolarisation from plateau potential (phase 2) to resting potential^{[Chen et al, 1991], [Gepstein et al, 1997], [Yue et al, 2004]}. These studies suggested that, to achieve a good agreement between ARI and *MAP90*, the ARI should be measured differently for positive T-waves compared with negative and biphasic T-waves; the authors advocated the use of T_{downs} , the point of most negative downslope (Figure 14), in place of T_{up} for calculations involving positive T-waves.

Although this alternative method for measuring ARIs may statistically give a more reliable approximation of when *MAP90* occurs, it lacks the theoretical foundation of the Wyatt method, which in any case is designed to indicate t_R , not *MAP90*. It has been shown that the downstroke of a positive T-wave does not consistently correspond with local repolarisation^[Coronel et al, 2006]. Furthermore, T-waves are known to present themselves on a continuous spectrum from positive to negative, via biphasic morphologies. Coronel and colleagues showed that, when a small area of myocardium is cooled to lengthen local ARIs, the T-waves of UEGs from that area can be found to move through the spectrum from positive to negative, while the shape of T-waves observed in a more remote region changes very little^[Coronel et al, 2006]. The alternative method is not compatible with this impression of continuously varying T-waves because, at whatever point the distinction between biphasic and positive T-waves is made, the measured ARI will undergo a sudden and dramatic change in value without physiological explanation. The Wyatt method is therefore preferable in any study into modulations of action potential duration (APD). If a measure of *MAP90* is needed, recent simulation studies^{[Colli Franzone et al, 2007b], [Potse et al, 2009]} have indicated that a good indicator is T_{d2} , the time after T_{up} at which the downward curvature of the T-wave is at a maximum (Figure 14). However, the use of the second time-derivative in this method makes it somewhat impractical outside of computational modelling, because each successive differentiation reduces the signal-to-noise ratio of a real signal.

II.E.2. Predicting unipolar electrogram morphology from the transmembrane potential

II.E.2.a. An intuitive model

To better understand the reliability and limitations of UEG-based repolarisation measures, one must consider in greater detail the relationship between the UEG and the familiar transmembrane potential. Potse and colleagues recently proposed a simple model that offers an intuitive explanation for how typical UEG morphologies arise from cardiac action potentials^[Potse et al, 2009]. The model calculates the UEG as the difference between a “local component”, arising from

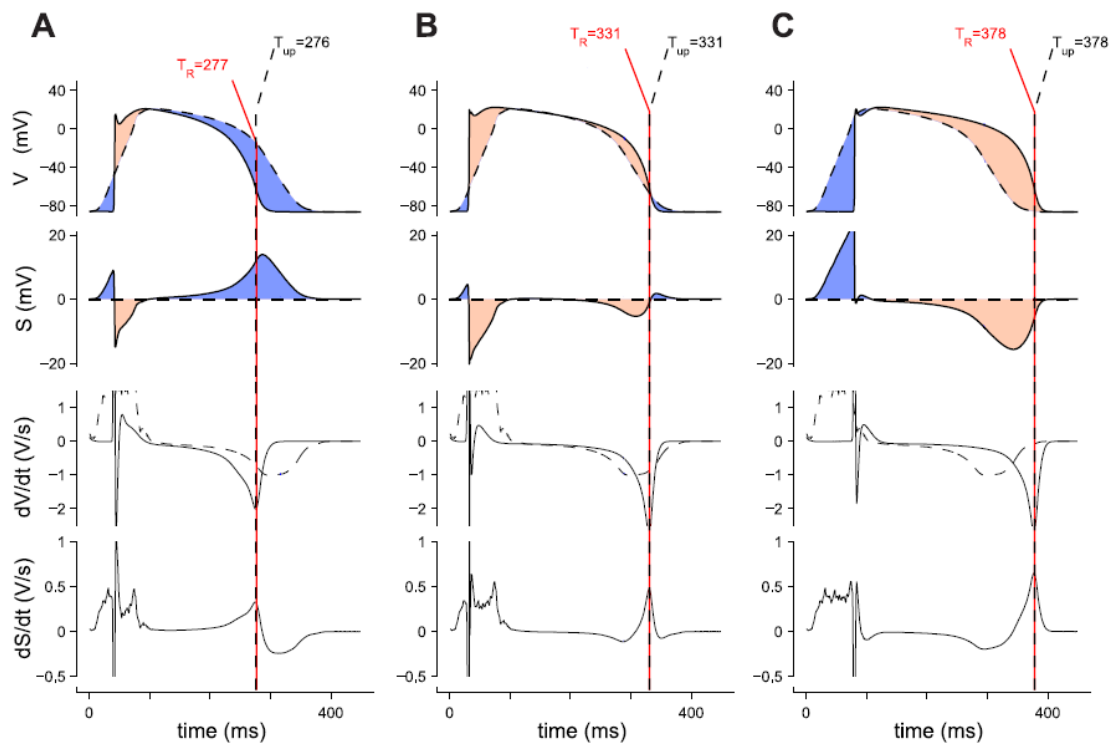


Figure 16: Adapted from [Potse et al, 2009]. Permission not required. The morphology of the UEG can be explained as the difference between a local action potential (solid lines in top row) and a smoother remote contribution (dashed lines in top row). In the second row, the local AP is subtracted from the remote contribution to give the simulated UEG, S . The third and fourth rows give the temporal derivatives of the traces in the first and second rows, respectively. Vertical lines mark the steepest downslope of the local action potential (solid red line) and the steepest upstroke of the UEG (dashed black line). T_R is found to coincide with T_{up} . Early-repolarising sites (column A) present positive T-waves, while late-repolarising sites (column C) give negative T-waves.

the point on the myocardium closest to the exploring electrode, and a “remote component” that accounts for the electrode configuration and the conductivity of the surrounding volume.

Figure 16 shows that the model proposed by [Potse et al, 2009] can explain the full spectrum of typical basic UEG waveforms (second row) – from positive T-waves (column A) to negative T-waves (column C), via biphasic T-waves (column B). Similar variations appear in the polarity of the activation waves. The figure also shows that the steep slopes in the local action potential upon depolarisation and repolarisation dominate the slope of the UEG at these times, supporting the use of \dot{V}_{min} and T_{up} as temporal markers of local behaviour.

Before this model can be accepted as a useful representation of the genesis of the UEG, the origins of the local and remote components must be explained theoretically.

The local component can be estimated as the extracellular potential ϕ_e immediately outside a myocyte using the well-known core-conductor model,^[Plonsey & Barr, 2007] the discrete form of which is shown in Figure 17. Note that the model assumes a restricted extracellular space, which is not a valid assumption for many UEG measurements. The implications of this discrepancy are discussed in a later section.

From the core-conductor model, the extracellular potential can be related to the transmembrane potential v_m as shown in (1). The full derivation for this expression is given in Appendix I.

$$\phi_e = -\frac{\bar{R}_e}{\bar{R}_i + \bar{R}_e} v_m. \quad (1)$$

where \bar{R}_i and \bar{R}_e are the resistance per-unit-length for the intracellular and extracellular domains. Since \bar{R}_i and \bar{R}_e are taken as constants, (1) shows that the time-course of the potential immediately outside a myocyte is morphologically similar to the time-course of the transmembrane potential, inverted and scaled by a constant factor. The morphology of the transmembrane potential is well known, hence we have an expression that yields ϕ_e as the “local component” in the model proposed by [Potse et al, 2009].

To characterise the “remote component” of their model of the UEG, Potse and colleagues refer to an expression achieved by Geselowitz^[Geselowitz, 1992]. Seen in (2), this expression describes the UEG using the bidomain model, which is the most realistic model of cardiac electrophysiology for which whole-heart simulations are currently computationally tractable^{[Miller & Geselowitz, 1978], [Geselowitz & Miller, 1983], [Geselowitz, 1989], [Vigmond et al, 2008], [Plank et al, 2008]}. This model treats the intracellular and extracellular domains as continuous media that both occupy the whole of the myocardium, with their effective properties scaled by a factor to account for the actual geometric relationship between the two




Figure 17: The discrete, linear core-conductor model for a single fibre. \bar{R}_i and \bar{R}_e are the resistance per unit length for the intracellular and extracellular domains. Δx is the incremental distance along the length of the cell. \bar{I}_m is the membrane current per-unit-length. I_i and I_e are the intracellular and extracellular currents. ϕ_i and ϕ_e are the respective potentials. The black squares represent membrane elements. These elements can be described in greater detail using a Hodgkin-Huxley type model,^{[Hodgkin & Huxley, 1952], [Plonsey & Barr, 2007]} which is not necessary for the present explanation.

domains. The two domains interact by the exchange of membrane current. From the bidomain model, it can be shown that the voltage between an electrode in the myocardium at point p and another elsewhere in the body is

$$UEG = -\frac{\sigma_i}{\sigma_i + \sigma_e} v_m(p) - \int_{\Theta} \sigma_i v_m \nabla Z \cdot d\vec{S}. \quad (2)$$

Here, σ_i and σ_e are the effective intracellular and extracellular conductivities. Θ is the outer boundary surface of the heart. The full derivation for this equation is given in Appendix II. Note that the first term is equivalent to the right-hand side of (1). This is the local component L of the UEG.

$$L = -\frac{\sigma_i}{\sigma_i + \sigma_e} v_m(p) \quad (3)$$

Before the remaining undefined variables in (2) are introduced, note that the second term gives the desired remote component Γ because [Potse et al, 2009] define their model as

$$UEG = L - \Gamma \quad (4)$$

Hence we have an expression for Γ that can be examined to achieve a more thorough understanding of UEG morphology:

$$\Gamma = \int_{\Theta} \sigma_i v_m \nabla Z \cdot d\vec{S}. \quad (5)$$

$d\vec{S}$ is the outward normal vector to each infinitesimal surface element of the heart's outer boundary, Θ . Its magnitude is equal to the area of that element. ∇Z is known as the lead field for the electrode configuration. It describes the overall distribution of conductivity within the torso, relative to the positioning of the electrodes, including any anisotropy and heterogeneity. In doing so, it yields an understanding of the relationship between a source and a voltage measurement in this complicated environment. The lead field can be defined more precisely as follows. Consider passing a unit current between two electrodes, A_1 and A_2 , in an arbitrary volume conductor while measuring the resultant potential difference between another pair of electrodes, B_1 and B_2 . The principle of reciprocity, a fundamental of electromagnetic theory, states that if a unit current is instead applied between B_1 and B_2 , the same potential difference will be observed between A_1 and A_2 [Malmivuo & Plonsey, 1995], [Jackson, 1998]. This principle can also be applied to describe the medium surrounding a distribution of current sources, such as the myocardium [McFee & Johnston, 1953], [Geselowitz, 1989]. If ∇Z is the negative of the electric field induced by passing a unit current between the UEG electrodes, then the voltage measured between those electrodes by a particular distribution of current source density \vec{J} can be written as a volume integral throughout the heart's volume, H ,

$$UEG = \int_H \vec{J} \cdot \nabla Z dV. \quad (6)$$

It is shown in Appendix II that these two different expressions for the UEG – (2) and (6) – are in agreement with each other.

In (6), ∇Z is seen to act as a transfer impedance; its value at each point in space expresses the sensitivity of the lead voltage to a source element at that point. The direction and magnitude of the source element at each point is described by \vec{J} . The concepts of the lead field and reciprocity facilitate an intuitive understanding of the gross effect of the surrounding tissues and their various properties. The relative sensitivity of a lead to different points within a source distribution can be roughly assessed by considering the electric field lines that a current applied to the lead would follow. The separation between these field lines indicates the strength of the lead field. Figure 18 illustrates the usefulness of this approach, using the Brody effect as an example. The Brody effect refers to a bias in the sensitivity of ECG recordings towards transmural currents compared with currents travelling tangential to the cardiac wall [Malmivuo & Plonsey, 1995]. This bias results from the approximately spherical shape of the heart and the relative conductivities of the myocardium, intracardiac blood, and the surrounding tissue. The intracardiac blood is more conductive than myocardium, and the heart is mostly surrounded by the lungs, which have a very low effective conductivity. As a result of this arrangement, the lead field is stronger through the centre of the heart, where the field lines cross the heart in the transmural direction.

The ECG has been used here to illustrate the lead field concept because the subtleties of the ECG, such as the Brody effect, have been well researched. However, the concept is equally applicable in analysing the subtleties of UEG recordings, including manifestations of the Brody effect, as will be discussed in later sections of this thesis. For now, the most important implication to note is that the dot product $\nabla Z \cdot d\vec{S}$ in (5) can be generally be expected to be negative. A unit current applied between the exploring electrode inside the heart and the reference electrode outside would flow across the heart boundary in the outward direction. ∇Z was said to be defined as the

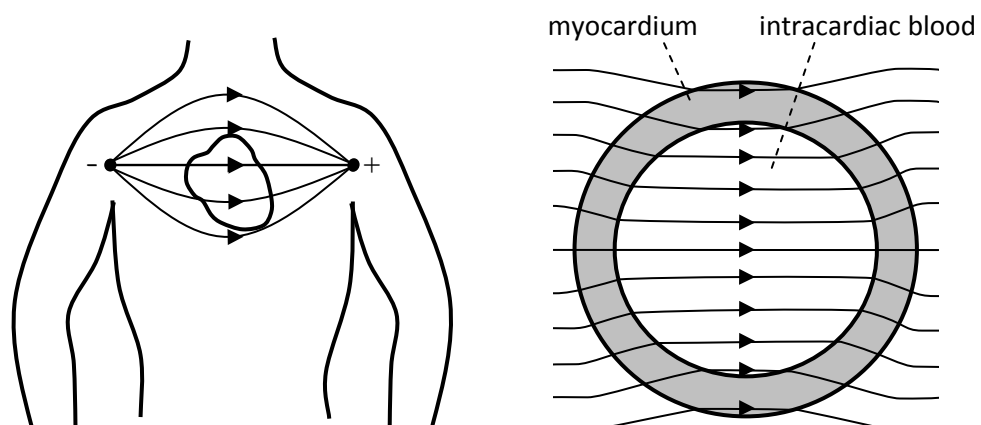


Figure 18: *Left:* The approximate lead field for ECG lead I across the human torso, ignoring the effects of variations in conductivity. *Right:* The Brody effect, used here as a demonstration of how the lead field concept facilitates intuitive predictions of the effects of variations in conductivity on bioelectric measurements. As proposed in [Malmivuo & Plonsey, 1995], the heart is modelled as a spherical blood volume enclosed by myocardium, with a uniform lead field imposed. Given that the intracardiac blood and myocardium are more conductive than the surrounding tissue, it can be predicted that an ECG lead is more sensitive to transmural (radial) currents than tangential currents, because the lead field lines are diverted towards the more conductive pathway.

negative of this electric field, so it is directed inwards and its dot product with the outward vector $d\vec{S}$ will be negative.

Now that the factors contributing to the remote component, Γ , of the UEG model have been identified and explained, the typical morphology of Γ can be assessed. The remarkable fact that v_m , in the integrand of (5), is evaluated only at the surface of the heart allows an intuitive interpretation of how Γ will vary as each wave of activity spreads across the surface of the myocardium. During the activation phase, as an increasing proportion of the heart's surface area depolarises to a more positive transmembrane potential, Γ will generally decrease and cause an upward trend in the UEG. Taccardi and colleagues noticed a component of this nature in UEG activation waves and referred to it as “drift of the reference potential” [Taccardi et al, 1998]. During the recovery phase, an opposite trend can be expected. Potse *et al* confirmed this in their validation of the model by computing Γ for five different electrode configurations using a realistic bidomain model of the heart, as seen inverted in Figure 19 [Potse et al, 2009]. On average (dashed line) Γ resembles a smoothed action potential, as predicted in the preceding analysis. In a few cases, however, the upward slope is interrupted and/or preceded by a downward notch. These deflections may be caused by phase 1 of the action potential (as seen in Figure 14, page 45) or by the possibility of the lead field crossing the heart surface in the outward direction at some points and changing the sign of $\nabla\mathbf{Z} \cdot d\vec{S}$ in (5). The likelihood of the latter explanation is increased by the fact that Potse *et al* always positioned their reference electrode very close to the myocardium, making it conceivable that a single field line would cross the surface multiple times.

To summarise, (5) expresses the UEG as the difference between a local component, which tracks the negative of the transmembrane potential at the site of the exploring electrode, and a remote component that tracks the spread of activity around the surface of the heart. The sensitivity of the UEG lead to this surface activity differs for each point on the surface according to the




Figure 19: Adapted from [Potse et al, 2009]. Permission not required. The negative of the remote component of the simple UEG model, calculated for five different electrode configurations using (5). The lead field was calculated using a realistic model of the distribution of conductivity throughout the heart. The dashed line shows the average of the five solid lines.

electrode configuration and the surrounding tissue properties. Figure 16 (page 49) shows, by this interpretation, that the polarity of the activation and recovery waves in UEG traces is dependent on the timing of local activity relative to global activity. For example, positive T-waves are recorded at sites that repolarise early and negative T-waves are recorded at the sites that repolarise later.

Potse and colleagues use this model to support the use of T_{up} in the UEG as a marker of t_R in the local action potential^[Potse et al, 2009]. However, the model also clarifies a limitation of this measure. Previously, it was noted that ARIs measured using T_{up} from positive T-waves tend to underestimate the effective refractory period by up to 15 ms^[Millar et al, 1985]. Note in Figure 16 that the upstroke of a positive T-wave occurs at a time when the remote component imposes a slight downward curvature on the UEG. This imposed curvature can be expected to cause the inflection point T_{up} to occur at a slightly earlier time than the inflection point in the local action potential. An opposite effect occurs for negative T-waves because the remote component imposes an upward curvature on the upstroke. The resultant delay in T_{up} would tend to cancel out the physiological difference between action potential duration (APD) and effective refractory period (ERP), whereas in positive T-waves this difference would be exaggerated by the artificial shortening of the ARI. The exaggerated difference explains the observation by Millar *et al* that the discrepancy between ARI and ERP is greater in positive T-waves than in negative T-waves. Further evidence for this effect appears in [Scacchi et al, 2009]. In the results of this simulation study, it can be seen that at sites with relatively short APDs (and, presumably, relatively early repolarisation times), the ARI tended to underestimate APD slightly. At sites with long APDs, the ARI tended to overestimate APD. The curvature of the remote component does not cause a significant difference between T_{up} and t_R in Figure 16, but here Potse *et al* have used a particularly rectangular action potential morphology, which makes t_R temporally distinct. In other models or in experimental measurements, the transitions between phase 2, phase 3, and the resting phase of the action potential may be less distinct so that the timing of T_{up} is more susceptible to the effect of curvature in the UEG's remote component.

To conclude this section, it should be noted that (2) describes the effective conductivities of the intracellular and extracellular domains using scalar quantities. It would be more accurate to account for the anisotropy of the myocardium by treating conductivity as a vector with separate components for the longitudinal and transverse directions, with respect to the fibre orientation. However, this treatment would not allow such a succinct expression for the UEG as that seen in (2). [Geselowitz & Miller, 1983] show that (2) holds (where the scalar conductivities are replaced by the magnitudes of the conductivity tensors) as long as the ratio between the longitudinal and transverse conductivities is the same for the two domains. The implications of deviations from this condition are discussed in the next section.

II.E.2.b. Local effects of anisotropy

Equation (2) grants useful insight into the broad morphology of UEGs, in which activation and repolarisation waves typically appear within the spectrum of positive, negative, and biphasic forms presented in Figure 16; the part of this spectrum into which a particular UEG falls is dependent on the timing of local activity relative to the surface activity. However, this interpretation does not explain the full range of features that are observed in practice.

For example, Figure 20 shows several activation waves measured from early-depolarising sites in an excised slab of canine myocardium. Many of these activation waves show substantial positive deflections and/or gradual negative deflections preceding the sharp down-stroke of local activation. The model of Potse *et al*^[Potse et al, 2009] predicts only a sharp down-stroke followed by a gradual upward slope for these sites. Figure 20 includes “computed” UEGs, determined by implementing the bidomain model as a set of differential equations, which successfully predict all features of the recorded UEGs. The bidomain model is capable of predicting these additional features, but they are not apparent from the form of (2).

For the purposes of understanding the genesis of typical UEG features, perhaps the most useful formulation of the bidomain model is that achieved by Colli Franzone and colleagues^[Colli Franzone et al, 1998]. They decompose the right hand side of (6) (page 52) into four components, each of which has a visually identifiable contribution to the typical UEG waveform. This is accomplished by first decomposing the source term \vec{J} into an axial component, \vec{J}_a , and a ‘conormal’ component, \vec{J}_c , such that $\vec{J} = \vec{J}_a + \vec{J}_c$. The axial component acts parallel to the fibre direction. The conormal component acts in the direction of $M\nabla v_m$, where M is the conductivity tensor, a 3x3 array describing the summed intracellular and extracellular conductivities in all three dimensions. Following similar manipulations to those used in reaching (2) (Green’s Theorem and other vector identities; see Appendix II), the resulting expression is

$$UEG = \int_{\ominus} \alpha v_m M \nabla Z \cdot d\vec{S} - [\alpha v_m(p) - \alpha v_m(q)] + \int_H v_m M \nabla \alpha \nabla Z dV + \int_H \vec{J}_a \cdot \nabla Z dV \quad (7)$$

The four components of this expression of the UEG are discussed individually below.

Surface component, UEG_S :

$$UEG_S = \int_{\ominus} \alpha v_m M \nabla Z \cdot d\vec{S} \quad (8)$$

Colli Franzone and colleagues^[Colli Franzone et al, 1998] derive the conormal source term as $\vec{J}_c = -\alpha M \nabla v_m$. The constant α relates the conductivities of the intracellular (i) and extracellular (e) domains in the direction transverse (t) to the fibre orientation: $\alpha = \sigma_{t,i} / (\sigma_{t,i} + \sigma_{t,e})$. UEG_S is equivalent to the remote component Γ (see (5)) in the simplified model described by Potse and colleagues^[Potse et al, 2009].




Figure 20: Reproduced from [Spach et al, 1979] with permission. Unipolar electrograms (UEGs) measured from various sites near the centre of a 30 mm x 20 mm slab of canine ventricular myocardium. Also shown are corresponding computed UEGs, determined using a bidomain model.

Note that only the UEG segments corresponding to activation are shown; repolarisation is not shown. The 2-dimensional sequence of activation is shown by isochrones (dashed lines) with a spacing of 2 ms. Activation was initiated by a pacemaker stimulus at the centre of the region, near electrode-site 1. The slab has been divided into four regions defined by the typical UEG morphology they presented. Region A, near the site of activation, is characterised by an immediate downward deflection (UEG 1). Region B (UEGs 2 and 3) is an area of fast conduction along the axial direction of the myocytes. Activation in this region presents a sharp downstroke in the UEG, preceded by a more gradual upward deflection. Region C is an area of slower conduction in the transverse direction. In UEGs from this region (4-6), a gradual negative deflection precedes the sharp downstroke associated with local activation. Between regions B and C lies the transitional region D, in which the direction of activation propagation cannot be comfortably categorised as either axial or transverse. Correspondingly, the UEGs from this region (7 and 8) bear ambiguous resemblance to those from both region B and region C.

Local component, UEG_L :

$$UEG_L = -[\alpha v_m(p) - \alpha v_m(q)] \quad (9)$$

When the position q of the reference electrode is outside the heart, $v_m(q) = 0$. Hence UEG_L is ordinarily a scaled, inverted version of the action potential at the location of the

exploring electrode, p . It is equivalent to the local term in the Potse model. It should be noted that Colli Franzone and colleagues refer to UEG_L as the ‘jump component’ because its morphology resembles a Heaviside step function when studying activation only. However, their model is also valid during repolarisation, when UEG_L can be expected to vary more smoothly over time.

Tissue component, UEG_T :

$$UEG_T = \int_H v_m M \nabla \alpha \nabla Z dV \quad (10)$$

Note that the term $\nabla \alpha$ refers to a spatial variation in tissue conductivity. Colli Franzone and colleagues remark that α can be expected to vary smoothly in healthy myocardium^[Colli Franzone et al, 1998]. Hence the contribution of UEG_T to UEG will be small in most cases. However, if the exploring electrode is positioned close to an area of pathological myocardium, UEG_T is likely to become non-negligible.

Axial component, UEG_a :

$$UEG_a = \int_H \vec{J}_a \cdot \nabla Z dV \quad (11)$$

Colli Franzone and colleagues derive the axial source term as $\vec{J}_a = (-\vec{\beta} \cdot \nabla v_m) \hat{a}$. \hat{a} is a unit vector in the direction of the myocardial fibres. $\vec{\beta}$ acts in the same direction with a magnitude that, similar to α , relates the conductivities of the intracellular (i) and extracellular (e) domains in the axial (a) and transverse (t) directions: $|\vec{\beta}| = \sigma_{a,i} - \sigma_{t,i} [(\sigma_{a,i} + \sigma_{a,e}) / (\sigma_{t,i} + \sigma_{t,e})]$.

The morphology of UEG_a can be more easily inferred by converting (11), which is standard form for a dipole source distribution, to its equivalent monopole formulation^[Plonsey & Barr, 2007]:

$$UEG_a = \int_H \frac{d(\vec{\beta} \cdot \nabla v_m)}{dx} Z dV \quad (12)$$

Assume $\vec{\beta}$ to be constant throughout H and consider the case in which ∇v_m is parallel to the fiber direction i.e. $\nabla v_m = \hat{a}(dv_m/dx)$. We then have

$$UEG_a = |\vec{\beta}| \int_H \frac{d^2 v_m}{dx^2} Z dV \quad (13)$$

Here we see that that the axial component is proportional to the second derivative, in the axial direction, of the membrane potential. Consider the spatial variation in that term. The top panel of Figure 21 shows a ventricular action potential simulated using the Luo-Rudy model of membrane behaviour,^[Luo & Rudy, 1991] as implemented by Benjamin Battacharya-Ghosh from the Department of Mechanical Engineering, University College London. Because the depolarisation of a myocyte occurs quickly and the wave of activation spreads quickly, the spatial distribution of the membrane potential, $v_m(x)$, during this phase can be approximated by stretching the time scale of $v_m(t)$ according to an assumed conduction velocity (around 65 cm/s)^{[Plonsey, 1969], [Plonsey & Barr, 2007]}. Clearly the resultant morphology of the second spatial derivative, d^2v_m/dx^2 , will follow that of d^2v_m/dt^2 , which is seen in the lower-right panel of Figure 21. Note that the activation sequence can be simplistically modelled as two closely positioned monopole current sources of opposite sign. As the action potential propagates down a fibre towards an extracellular electrode (from right to left

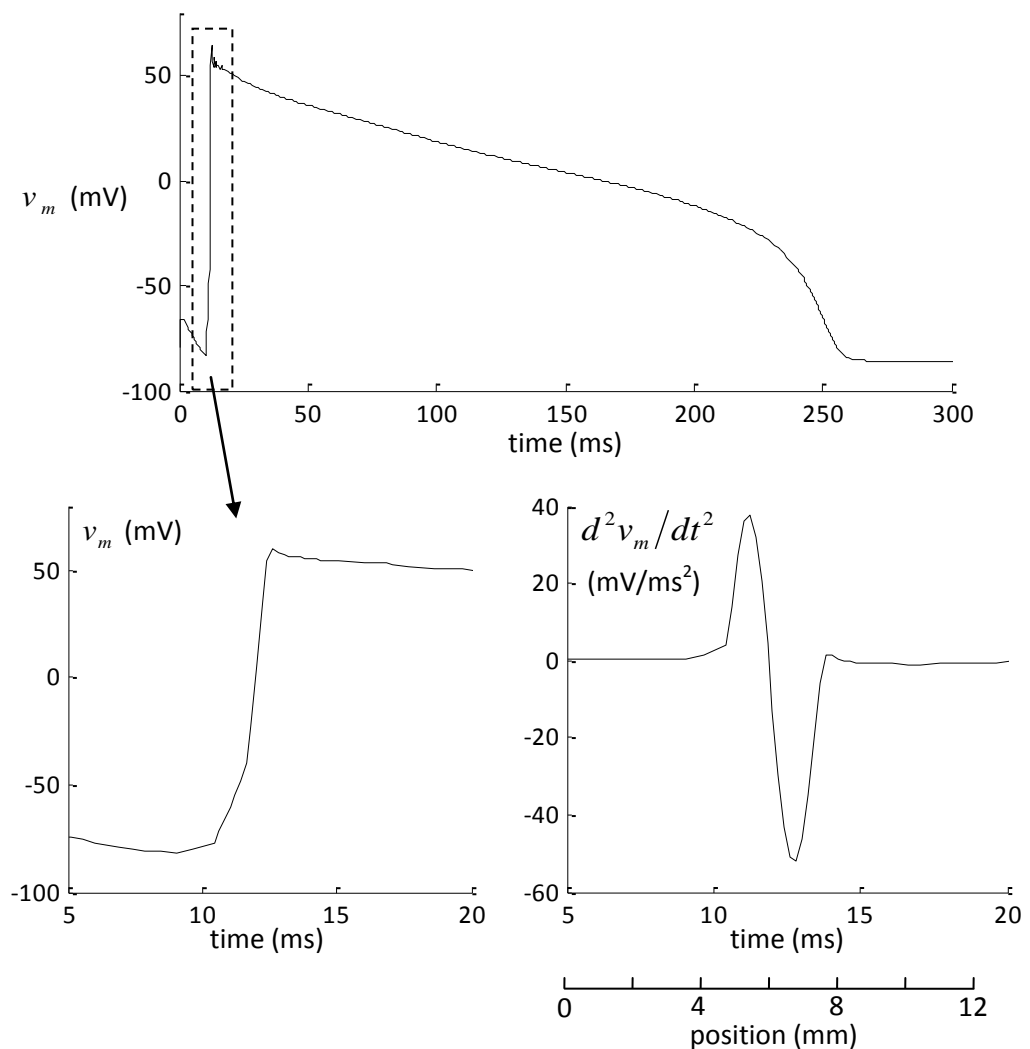


Figure 21: Top – The time-course of a typical mammalian ventricular action potential, as predicted by the Luo-Rudy model^[Luo & Rudy, 1991]. Lower left – A close-up of the depolarisation phase of the action potential, as demarcated by the dashed box in the top panel. Lower right – The second temporal derivative of the membrane potential during this phase. As described in the text, during depolarisation this morphology can be assumed to be the same as that of the second spatial derivative. Hence it indicates the distribution of membrane current along the myocyte, as indicated by the added ‘position’ axis.

in Figure 21), the electrode will first record an upward deflection from the positive monopole, followed shortly by a downward deflection produced by the negative monopole. When its appearance is well defined, this biphasic feature is typically referred to in the literature as the “intrinsic deflection” of local activation.

The emergence of the positive (*outward*) current leading the wavefront may seem counterintuitive, given that depolarisation was said to be characterised by an *influx* of Na^+ ions. The physical explanation is that the current associated with this influx dissipates down the inside of the fibre towards the inactive region ahead of the wavefront, where the intracellular potential is lower. As this intracellular potential is thus raised, current is driven outwards and through the extracellular medium to feed the depolarising current. These currents are depicted by the yellow arrows in Figure 22, which also gives the associated extracellular potentials determined experimentally by [Spach et al, 1973].

The derivation of (13) assumed ∇v_m to be parallel to the fibre orientation. As the angle between the two increases to 90° , the magnitude of the local contribution to UEG_a decreases to 0. This component is responsible for the overall UEG morphology’s dependence on the direction of propagation relative to the fibre orientation, as seen in Figure 20.

The reason given for introducing (7) as an expression for the UEG was that the simplified model proposed by Potse and colleagues, though useful as an intuitive explanation of gross UEG morphologies, could not explain the full range of features presented in Figure 20. (7) is easily related to the Potse model, because the local and surface components are common to both.

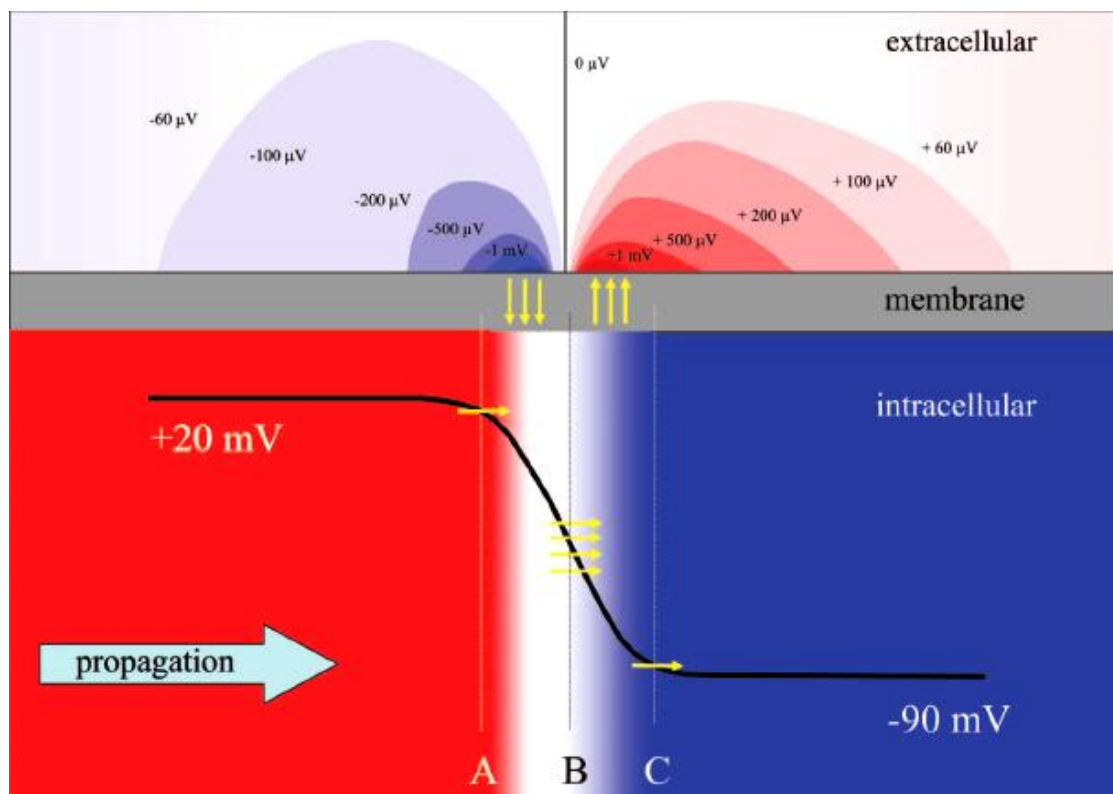


Figure 22: Reproduced from [de Bakker & Wittkamp, 2010] with permission. A schematic of the transmembrane and intracellular currents (yellow arrows) surrounding a depolarisation wavefront in a cardiac myocyte. Also depicted are the associated extracellular potentials.

However, (7) includes two additional terms: the tissue component, which can usually be neglected, and the axial component, which accounts for the more complex features seen in Figure 20. For example, note that the axial component induces a large positive deflection before the down-stroke of local activation at site 3, towards which activation spread in the direction of fibres. At other sites the effect is small either because the myocytes have depolarised before the outward currents have developed (e.g. sites 1 and 2) or because activation has propagated transverse to the fibres, such that intracellular currents could not be driven ahead of the depolarising influx (e.g. site 4). At these sites, the sharp down-stroke of local depolarisation is pre-empted by a more gradual down-slope due to the potential field induced by inward currents nearby. Sites 5-8 exhibit a blending of these two effects. The preceding analysis demonstrates the importance of tissue anisotropy in determining the morphology of the activation wave in the UEG.

II.E.2.c. Spatial specificity

For a monopole current source in an infinite, homogeneous, isotropic conductive medium, the transfer function relating the source to the potential field is inversely proportional to the distance between the source and the observation point, as shown in Appendix III. For a dipole source, the transfer function is inversely proportional to the square of that distance. Figure 23 plots the transfer function for the monopole-source case, in order to give an impression of the field of view of the UEG and its dependence on electrode positioning. When the exploring electrode is closer to the myocardial surface, the relative influence of neighbouring surface regions is reduced compared to that of the region immediately under the electrode. As mentioned above, the transfer function H is the inverse of the distance between the source (x', y', z') and the observation point (x, y, z) . In Figure 23, the coordinate system has been defined such that $z = 0$ across the myocardial surface and $y = 0$ in the plane of view. The transfer function relating the current source to the observed potential is then (ignoring a constant factor)

$$H(x - x') = \frac{1}{\sqrt{(x - x')^2 + (y')^2 + (z')^2}} \quad (14)$$

As mentioned previously, the true electrocardiographic transfer function ∇Z (or Z when a monopole-source formulation is used) is dependent on the heterogeneity and anisotropy of the extracellular conductivity of the myocardium and surrounding volume. The lead field's dependence on the overall anatomy of the body makes it difficult to achieve a generalised description of the spatial specificity of the UEG. Various studies have employed computational modelling to explore the susceptibility of the UEG to artefacts of remote activity^{[Steinhaus, 1989], [Taccardi et al, 1998], [Colli Franzone et al, 2007b], [Scacchi et al, 2009]}. Such studies have made invaluable contributions to our understanding of the field of vision of the UEG, but they are necessarily prescriptive in terms of electrode configurations and the spatiotemporal distribution of activity that they consider. The infinite, homogeneous conductor scenario described in the previous paragraph can be used for a more general impression of the UEG's spatial specificity, as in Figure 23. However, this interpretation is subject to the




Figure 23: The transfer function H for a monopole current source (see equation (88) in Appendix III) calculated along one dimension of the myocardial surface, using an electrode positioned above the point $x = 0$ at a distance of $z = 2\text{mm}$ (thick dashed line) and $z = 0.2\text{mm}$ (thick solid line).

Vertical lines mark the point at which H becomes half that at the point directly beneath the electrode ($x = 0$). Intuitively, an electrode closer to the area of interest will provide a narrower focus.

caveat that the body's confined nature and variations in conductivity will cause the true transfer function to deviate from this form. The nature of this deviation can be explored analytically by treating the heart as existing in an infinite medium of homogeneous conductivity and replacing the major discontinuities of conductivity (e.g. the skin and the surface of the lungs) with equivalent dipole-source layers^[Wilson & Bayley, 1950]. Such analysis cannot produce useful, generalised results regarding the spatial specificity of the UEG. However, the interested reader should note that this concept is a useful complement to the lead field concept as a means of qualitatively assessing a UEG measurement's dependence on movement in the surrounding tissues.

II.E.2.d. Repolarisation in the UEG

The preceding interpretation of the genesis of UEG morphology focuses on depolarisation because the chain-reaction nature of activation propagation makes for a relatively predictable spatiotemporal distribution of activity, hence a predictable UEG morphology. Prediction of the repolarisation portion of the UEG requires additional considerations.

The spatiotemporal distribution of repolarisation is more dependent on the behaviour of individual cells and, in particular, on the gating behaviour of the different ion channels. Some degree of synchronisation is provided by electrotonic effects^[Conrath & Opthof, 2006]; the gap junctions that join myocytes at their ends allows some continuity of the intracellular fluids, such that the drop in intracellular potential of a repolarising myocyte draws current from its neighbour. Thus, the intracellular potential of the neighbour also drops slightly, and the voltage-sensitive ion channels respond by initiating full repolarisation (phase 3 of the action potential). Experimental evidence of this synchronisation comes from the observation that myocytes in the intact myocardium are found to repolarise earlier than when the individual cell is isolated^[Clayton & Holden, 2004].

Although some direct synchronisation exists, there is substantial variation between the intrinsic APDs of myocytes from different regions of the heart. In normal circumstances this variation actually yields closer synchronisation of repolarisation times, because sites that depolarise later under the normal activation sequence tend to exhibit shorter action potentials than early-activating sites: APD is generally shorter at the base of the ventricles than at the

apex,^[de Bakker & Opthof, 2002] and is shorter at the epicardium than at the endocardium^[Yan & Antzelevitch, 1998]. The fact that the spatial gradient of APD tends to be opposite to the spread of activation yields closer synchronisation of the timing of repolarisation across the ventricles as long as activation propagates by the normal sequence. An exception to this tendency occurs in the form of a recently discovered class of myocytes known as M-cells^[Sicouri & Antzelevitch, 1991]. These cells are found in the middle layer of the ventricular wall and exhibit much longer APDs, as well as increased dependence on heart rate, as shown in Figure 24.

In practice, ventricular repolarisation is found to behave as a progressing wave, with similarities to the depolarisation wave. However, the direction and speed of the repolarisation wave are much more variable, owing to the heterogeneity of APDs. Also, phase 3 of the action potential typically lasts much longer than phase 0 (tens of milliseconds, compared with 2-4 ms) so the repolarisation wavefront could not be treated as a thin layer of activity unless the wave progressed very slowly. As a result, the second spatial derivative of membrane potential is very small, so repolarisation does not present additional UEG features akin to the early positive deflections seen in the activation waves of Figure 20.

II.E.3. Sources of error in measurements from unipolar electrograms

Although modelling studies and analytic work have shown that the UEG provides useful indicators of the timing of local depolarisation and repolarisation, practical implementation of these




Figure 24: Reproduced from [Sicouri & Antzelevitch, 1991] with permission.

The action potential duration (APD) determined from intracellular recordings in canine ventricular myocardium. Substantial transmural heterogeneity of APD exists. M-cells exhibit markedly longer APDs than endocardial and epicardial myocytes, and vary more steeply with changes in heart rate.




Figure 25: Recordings from guinea pig ventricular myocytes show that ischemia causes the action potential to decrease in amplitude and become more triangular in shape. Data from [Louch et al, 2002].

techniques is subject to several confounding factors, particularly when conditions do not match the ‘ideal’ conditions assumed in the preceding analysis. These factors are briefly detailed in two separate categories below.

Deviation from the assumed action potential morphology: It was noted in section II.C that the precise morphology of cardiac action potentials varies throughout the heart and is subject to varying physiological conditions. Those drawn in the top row of Figure 16 are particularly rectangular in shape; the steep slope during phase 3 (repolarisation) ensures a distinct artefact in the UEG. Certain conditions such as ischemia and rapid pacing (artificial control of heart rate) can cause the action potentials to become more triangular (i.e. with a more gradual downslope from phase 1 to full repolarisation) [Steinhaus, 1989]. In such cases, the upward UEG deflection associated with local repolarisation would be less sharp. In this situation the measurement of repolarisation time as T_{up} becomes less meaningful or more easily corrupted by the behaviour of the remote component of the UEG, as discussed on page 54.

Deviation from the idealised pattern of propagation: In justifying the use of \dot{V}_{min} and T_{up} to indicate local membrane behaviour, it was assumed that activation and repolarisation behave as single, steadily propagating waves passing close to the exploring electrode. It has been demonstrated through the modelling and experimental studies that, when the spread of activity does not match that assumption, \dot{V}_{min} and T_{up} can differ substantially from the membrane events they are said to represent [Steinhaus, 1989], [Scacchi et al, 2009], [de Bakker & Wittkamp, 2010]. Such exceptions can occur in the following scenarios:

- During fibrillation.
- At locations where propagation is initiated or terminated.
- Where two regions of myocardium, separated by an insulating layer, are both within the exploring electrode’s field of view. For example, the Purkinje fibres are separated from the

surrounding myocardium by a sheath of collagen^[de Bakker & Wittkampf, 2010]. The two regions will depolarise and repolarise at different times, so that the UEG appears as the superposition of two “normal” UEGs whose features will not necessarily be easily distinguished from each other. Such electrograms are often referred to as being “fractionated”. As a second example, the base of the ventricles is very close to parts of the atria, but is separated by the annulus fibrosus and the AV node.

- At sites of discontinuous tissue properties. These might include the boundaries of diseased/damaged tissue regions or regions where contractile myocytes meet specialised conductive cells. At these sites the effective intracellular conductivity or membrane behaviour may change sharply, selectively altering the action potential morphology of some cells and their manifestation in the UEG. Natural variations in action potential duration, such as those occurring in through the ventricular wall, may also cause a sudden change in the speed at which the repolarisation wave propagates, corrupting the measurement of T_{up} ^[Scacchi et al, 2009].

II.F. Control of the Heart by the Autonomic Nervous System

II.F.1. Overall structure and function

The heart is innervated at various points by the autonomic nervous system (ANS) (Figure 29). The heart rate and contractility (strength of contraction) at any given time are largely dependent on the balance between sympathetic and parasympathetic nervous activity. Sympathetic activity can loosely be interpreted as encouragement; in the heart, it generally contributes to an increase in rate and contractility. Conversely, parasympathetic nervous activity has an inhibitive effect.

Activity in the ANS varies in response to physiological factors such as blood pressure and muscle work rate^[Spyer, 1994]. Sensory measurements of these factors are processed by the brain, which

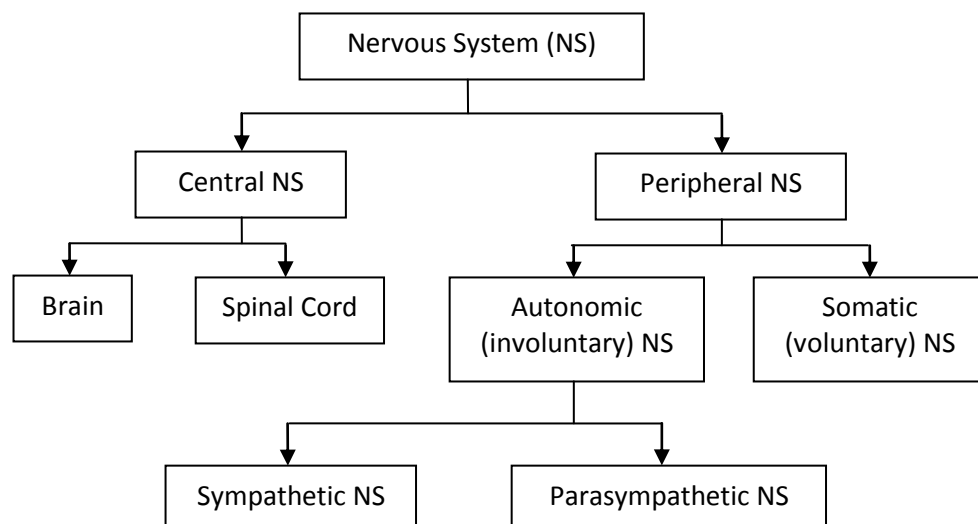


Figure 26: Hierarchy of the nervous system.

adjusts sympathetic and parasympathetic (a.k.a. vagal) activity appropriately. ANS activity is also known to vary according to the emotional state of the subject, and this has a measurable effect on cardiac parameters [Napadow et al, 2008], [Critchley et al, 2005]. The resultant change in cardiovascular behaviour alters the blood flow and hence influences the physiological state of the body. In that sense, the ANS is a control system whose purpose is to ensure that cardiac output constantly adapts to changing demands. Its response time is short enough to address sudden changes, such as the change in blood pressure induced when a person stands up. The feedback control loop for the ANS is illustrated in Figure 27 and Figure 28.

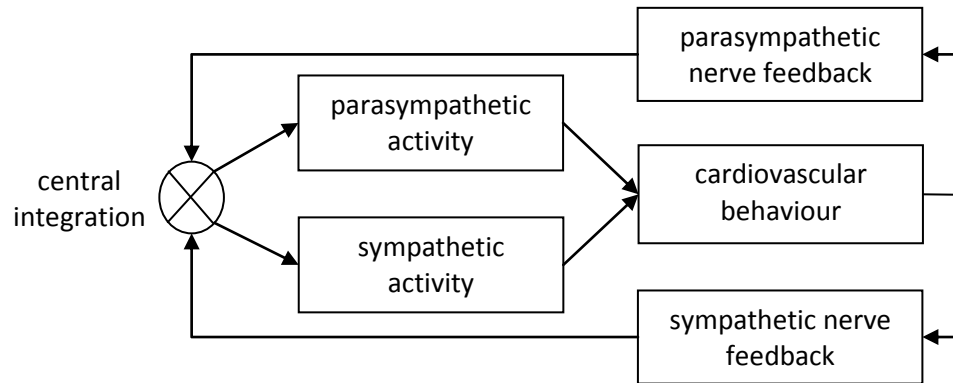


Figure 27: A simplified block-diagram representation of the feedback control loop in Figure 28.

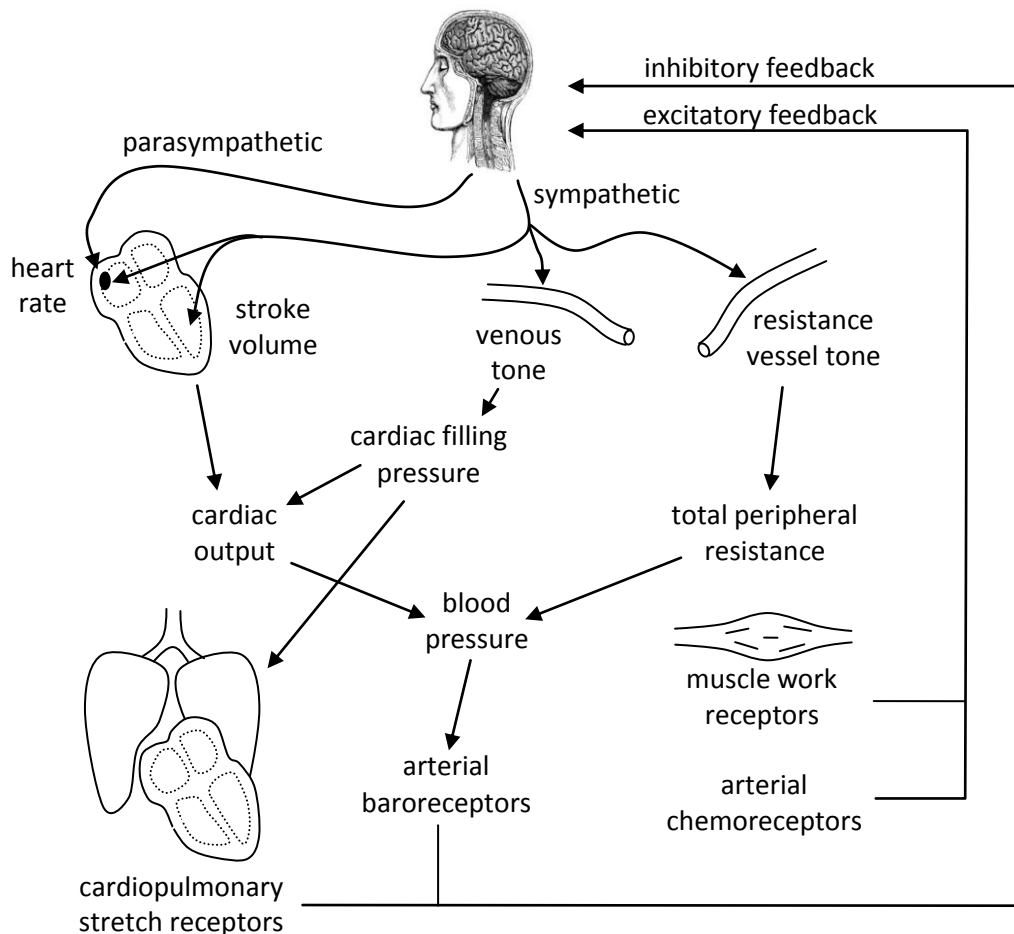


Figure 28: Cardiovascular branches of the autonomic nervous system as a feedback control loop. Physiological interactions as described by [Levick, 2003].




Figure 29: Adapted from [Levick, 2003]. Routes of sympathetic and parasympathetic innervation of the heart. It should be noted that this representation is simplified; the ventricles are also sparsely innervated by parasympathetic fibres, and all connections are subject to anatomic variation.

As mentioned previously, the ANS controls both heart rate and contractility. As illustrated in Figure 29, the parasympathetic NS innervates the sino-atrial node via the right vagus nerve. Increased activity in this nerve causes a reduced heart rate. This effect is achieved by the release of the neurotransmitter acetylcholine, which opens certain K^+ channels in sino-atrial myocytes and closes some Na^+ and Ca^{++} channels. The balance of electrochemical potentials therefore changes, hyperpolarising the cell and decreasing the slope of the resting potential's upward drift, as illustrated in Figure 30. The membrane potential takes longer to reach its threshold value, so the time between heart beats increases [Bolter et al, 2001].

The atrio-ventricular node is innervated by the parasympathetic nervous system via the left vagus nerve, which induces a response similar to that induced in the SA node by right vagus nerve. The likelihood of the AV node's automaticity (see p. 27) pre-empting that of the SA node is therefore reduced.

Activity in sympathetic fibres leads to the release of noradrenaline (also known as norepinephrine), rather than acetylcholine [Zipes & Jalife, 1999]. Increased levels of noradrenaline influence cardiac myocytes in various ways. Sympathetic fibres on the right side of the body innervate both the SA node and the AV node. Increased activity in these fibres causes an increased heart rate by mechanisms that are essentially opposite to those by which parasympathetic activity lowers heart rate [Hutter & Trautwein, 1956], [Levy & Martin, 1979]. On the left side, sympathetic fibres innervate the atria and ventricles. Increased noradrenaline levels indirectly lead to an increased probability and duration of Ca^{++} channels opening. The increased presence of Ca^{++} in the myocyte increases the force of contraction, as explained in appendix 2. Although the increased influx of Ca^{++} has a correspondingly increased depolarising effect, this is overcompensated by an increased outward K^+ current, so the action potentials of contractile myocytes are in fact shortened, allowing more frequent excitations. To the same end, sympathetic nerve activity on the right side leads to an

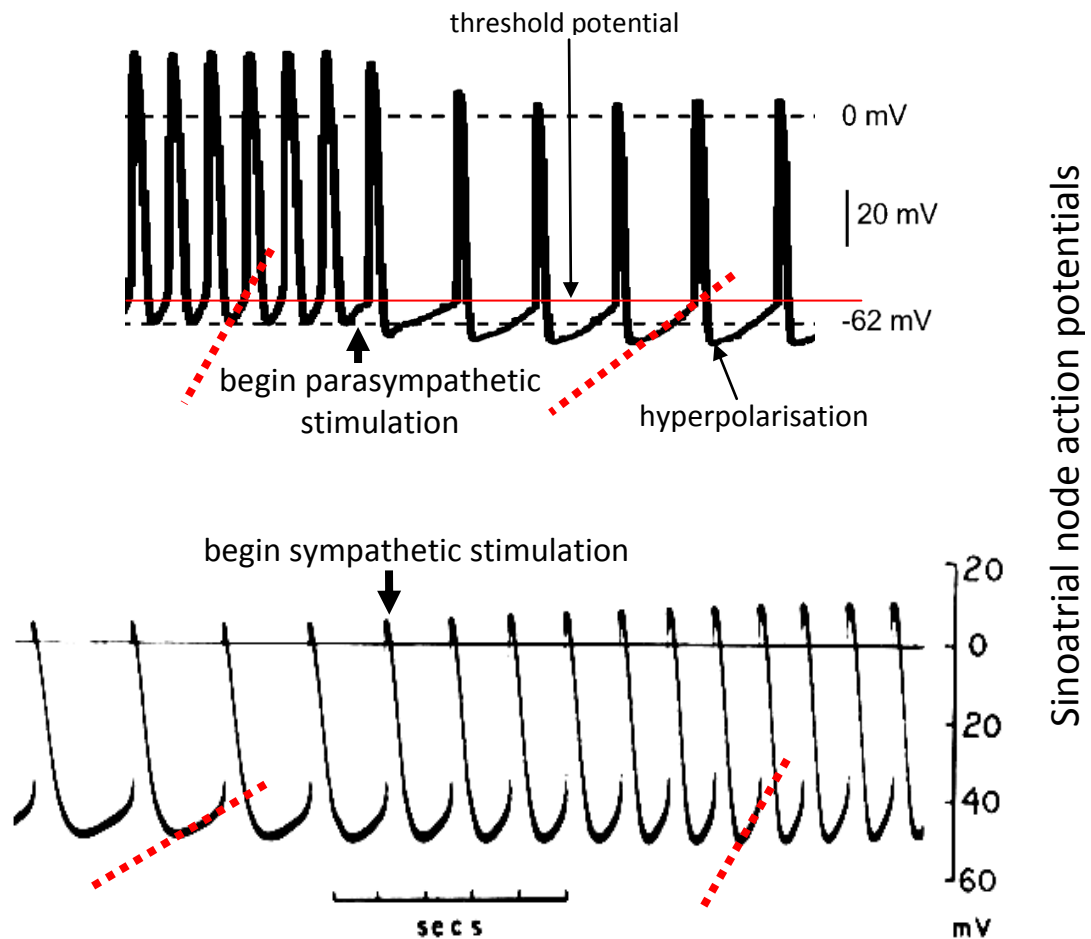


Figure 30: The effects of parasympathetic stimulation of the guinea pig SA node (adapted from [Bolter et al, 2001] with permission), and effects of sympathetic stimulation of the frog SA node (Adapted from [Hutter & Trautwein, 1956]; permission not required). Both types of autonomic input affect the rate of upward drift in the resting potential. Parasympathetic stimulation also affects the magnitude of the resting potential (hyperpolarisation).

increase in conduction velocity through the AV node, reducing the delay between atrial and ventricular contraction [Levy & Martin, 1979].

Furthermore, various genetic channelopathies (diseases of the ion channels) exist which increase a patient's likelihood of developing a fatal arrhythmia. Each channelopathy affects a particular ion channel's sensitivity to either sympathetic or parasympathetic input. For example, in the channelopathies known as Long QT Syndrome (types 1 and 2), increased sympathetic activity modulates the behaviour of certain Ca^{++} channels to prolong APD, which is opposite to the normal effect of increased sympathetic tone and incompatible with the concomitant increase in heart rate [Taggart et al, 2011]. In Brugada Syndrome, which affects a class of Na^+ channel, Nakazawa and colleagues found that patients with increased parasympathetic tone were more susceptible to arrhythmia [Nakazawa et al, 2003]. Patients with such channelopathies as these have selectively increased vulnerability to different environmental/behavioural stimuli. This observation leads to two important points: 1) in isolated cases, the task of inferring autonomic activity from cardiac electrophysiology can be complicated by the emergence of effects opposite to the expectations on which the inference depends, and 2) ANS behaviour that would be considered normal in most

subjects can have disastrous effects when acting on an abnormal heart. This second point emphasises the importance of gaining an understanding of even subtle aspects of autonomic interactions with the heart, as ordinary physiological mechanisms may in some cases give rise to catastrophic behaviour. With this motivation, the experiments discussed in sections V.A and V.C were designed to further the general understanding of physiological mechanisms that have not previously been linked by evidence to the emergence of arrhythmia.

II.F.2. Respiratory sinus arrhythmia

It is commonly observed that heart-rate fluctuates with breathing; this phenomenon has been introduced previously as ‘respiratory sinus arrhythmia’ (RSA). Blood pressure is also known to oscillate at the respiratory frequency^[Zhang et al, 2002]. Numerous candidate mechanisms have been proposed to explain these oscillations, including mechanical and autonomic effects. One might then hypothesise that cardiac electrophysiological properties beyond the sinus node also oscillate with respiration; myocyte behaviour around the heart is known to be influenced by ANS activity, as discussed in the previous section, and by mechanical inputs^{[Kohl et al, 2006], [Taggart & Sutton, 2011]}. Such manifestations of RSA and their significance with regard to the stability of cardiac behaviour have not been thoroughly explored in the literature, but novel observations of this nature were produced in the experiments described in section V.C. To contextualise those results, this section briefly introduces the mechanisms by which RSA is thought to occur.

Respiration is known to act as a mechanical input to blood pressure^[Cohen & Taylor, 2002]. The expansion of the chest during inspiration reduces pressure within the chest’s thoracic cavity. This pressure-drop acts on all four chambers of the heart, as well as the veins and arteries within the thoracic cavity, to varying degrees. One important result is that the left-ventricular stroke volume is reduced^[Innes et al, 1993]. The reduction in stroke-volume can be assumed to contribute to a reduction in arterial blood pressure. Also, the influence of intrathoracic pressure on the stretching of the SA node is thought to cause small accelerations in heart rate during inspiration^[Larsen et al, 2010]. However, it is found that the magnitude of RSA in heart-transplant recipients is just 2-8% of normal levels, suggesting that autonomic nervous input is responsible for the major part of normal RSA^[Bernardi et al, 1989].

Debate persists regarding the role of the ANS in generating RSA^{[Eckberg, 2009], [Karemaker, 2009], [Larsen et al, 2010]}. The key point of contention is whether the dominant mechanism is the baroreflex effect or a separate oscillatory input from the central nervous system. The baroreflex argument states that the decrease in arterial blood pressure due to the mechanical effects of inspiration triggers an increase in heart rate so that cardiac output increases, counteracting the pressure-drop. The competing argument is that sympathetic and parasympathetic outflow to the heart, including baroreflex signals, are modulated in the central nervous system by breathing. This phenomenon is known as respiratory gating, and yields increased sympathetic outflow and decreased

parasympathetic outflow during inspiration^[Eckberg, 2003]. Changes of this nature in the ANS outflow to the heart are expected to influence the behaviour of myocytes throughout the heart. The specific effects of ANS activity on the various membrane channels has been studied extensively. One can therefore expect numerous RSA-related electrophysiological effects in addition to the change in heart rate, such as modulation of action potential durations.

II.F.3. Psychological Effects

II.F.3.a. The pathway from emotional stimuli to cardiac effects

As well as maintaining stable physiological behaviour in response to the varying demands of the body's organs, the ANS also responds to commands from higher processing centres. Such responses are required to prepare the body for forecast requirements, as in the fight-or-flight response. The potential for this reaction to destabilise cardiac behaviour is colloquially familiar ("Calm down before you give yourself a heart-attack!"), but only recently has it been possible to investigate such mechanisms in detail. The possibility that the autonomic response to a stressor could have fatal consequences are not exaggerated. Taggart and colleagues summarise a range of evidence linking psychological influences to the occurrence of sudden cardiac death^[Taggart et al, 2011]. Thayer & Brosschot also list recent studies that have found evidence for an ANS-mediated link between negative emotional states and various forms of ill-health^[Thayer & Brosschot, 2005]. Figure 31 outlines the pathway by which emotional stimuli might evoke a fatal cardiac response.

II.F.3.b. The brain-heart laterality hypothesis

As described previously (see Figure 29, page 66), the ANS breaks out from the spinal cord in separate branches for the left and right sides of the body, and these branches do not innervate the heart symmetrically in terms of either functionality or spatial distribution of nerve endings. Furthermore, it is known that cerebral hemispheric asymmetries exist in the processing of emotions^[Critchley et al, 2005]. These observations led to the Brain-Heart Laterality (BHL) Hypothesis, that "the minority of people who are more lateralised to the left hemisphere during emotion channel this activation ipsilaterally to induce a lateralised imbalance in sympathetic input to the heart, increasing the likelihood of ventricular fibrillation and sudden death."^[Lane & Jennings, 1996] In simpler terms, asymmetric processing of emotions in the brain may generate asymmetric sympathetic outflow to the heart. The left sympathetic branch predominantly controls the behaviour of contractile myocytes while the right branch controls heart rate and the AV node's behaviour. Asymmetric input might therefore result in a scenario where ventricular action potential duration is not appropriately matched to heart-rate, increasing the risk of arrhythmia developing.

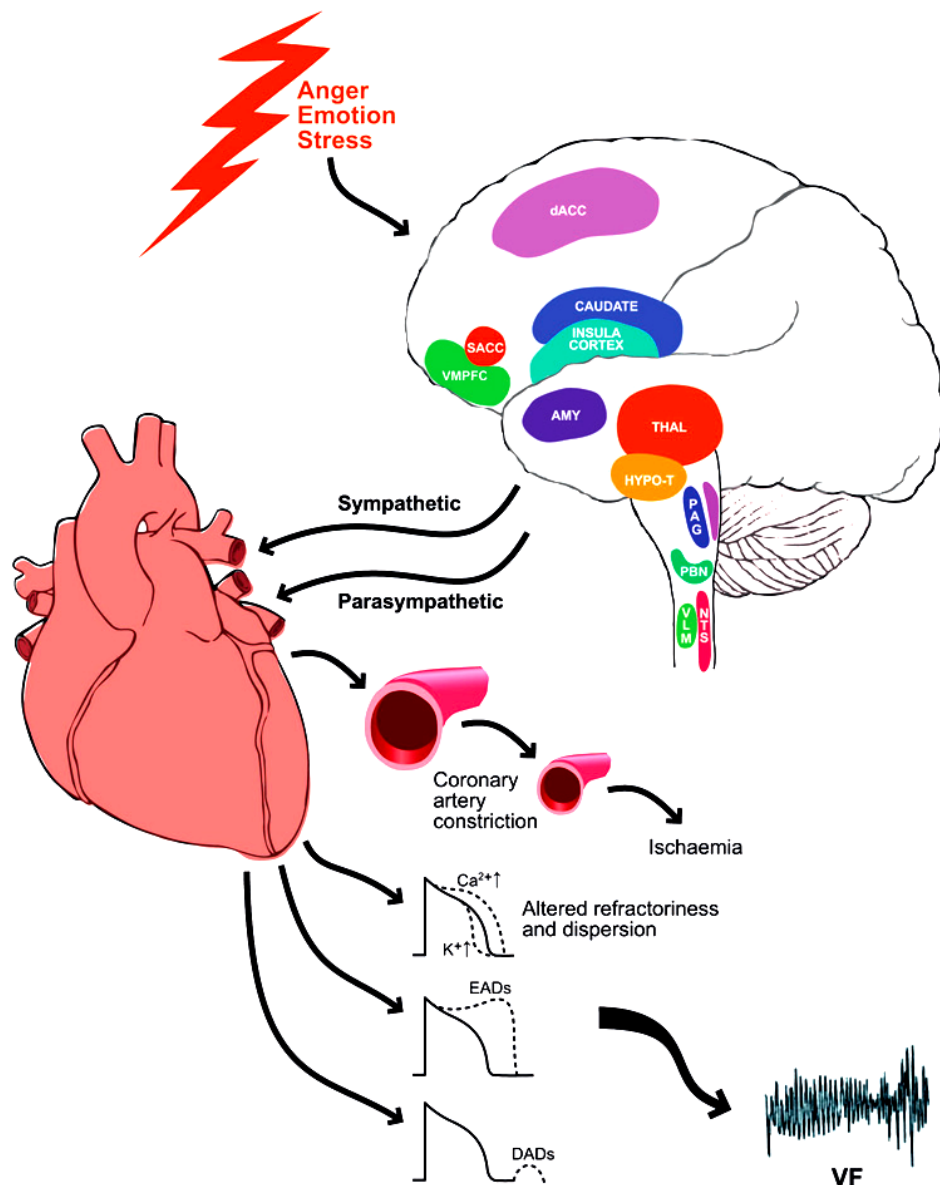


Figure 31: Reproduced from [Taggart et al, 2011] with permission. Emotional stimuli are processed in various brain centres (the individual roles of which are beyond the scope of this report), leading to a response of the autonomic nervous system. Effects on the heart include altered tone in the blood vessels that supply the myocardium and alterations to myocyte action potentials. Autonomic effects on the various membrane ion channels may lengthen or shorten action potentials heterogeneously, or they may increase the occurrence of early and delayed after-depolarisations (EADs and DADs). All of these effects can lead to the development of ventricular fibrillation (VF).

To summarise, the ANS innervates the heart at various sites to modulate heart rate and contractility. The localised effects of this innervation vary between sites, adjusting the behaviour of different ion channels to maintain beat-to-beat stability. An awareness of these mechanisms is essential in any investigation concerned with how interactions with the nervous system influence cardiac electrical stability. Previous sections described how experimental observations of electrical fields produced by the heart can be related to the morphology and spatiotemporal distribution of action potentials. It was also noted that an understanding of the key ion transfers that contribute to the action potential allows insight into the physiological mechanisms behind any observed changes in action potential. This section has briefly outlined how changes in ANS input to the heart can

induce changes in these cellular mechanisms. By combining these concepts, recordings of cardiac electrical behaviour can be used to infer the ANS input to the heart. This approach underlies all of the experimental work carried out in this project.

II.G. Heart Rate Variability as an Indicator of Autonomic Activity

As has been explained previously, heart rate continuously varies as the different control mechanisms that influence it are fine-tuned to the body's physiological state. The patterns of variation in heart rate, rather than just the rate itself, can therefore give insight into variations in the controlling influences both within and outside of the heart. For short term responses, the most significant of these influences is the ANS, although the endocrine (hormonal) system also plays a significant role that is beyond the scope of this thesis. Heart rate variability (HRV) is often taken as a useful, non-invasive measure of ANS activity. Various methods of quantifying HRV have been implemented in the past^[TFESCNASPE, 1996] and recent developments continue to improve upon these measurements^{[Barbieri et al, 2005], [Barbieri & Brown, 2006], [Chen et al, 2009]}. In general these methods are applied to time-series formed from the intervals between heart beats. These intervals are most commonly measured from ECG signals as the time between consecutive R waves, and in this case they are referred to as 'RR intervals'.

II.G.1. Interpretation of HRV measurements

The most widely accepted measures of HRV are based on frequency-domain measurements. They are popular because two distinct peaks tend to occur in the power spectrum of RR intervals. One of these peaks usually occurs in the range 0.04 Hz – 0.15 Hz (low frequency, or LF, band), and the other occurs in the range 0.15 Hz – 0.4 Hz (high frequency, or HF, band), as seen in Figure 32. Note that the HF band roughly matches the normal range of human respiratory frequencies; these oscillations are caused by respiratory sinus arrhythmia, as explained on page 68. Much like for RSA, some uncertainty persists regarding the origins of the oscillations manifested in the LF peak^[Julien, 2010]. In particular, it is unclear whether they arise from a resonance effect in the baroreflex mechanism or from a centrally mediated oscillation in sympathetic activity.

The magnitudes of the LF and HF peaks are often cited as indicators of sympathetic and parasympathetic activity, although caution must be exercised in such interpretations. The HF peak has been taken as a measure of baseline parasympathetic activity because the respiratory gating effect is more substantial in this branch of the ANS^[Katona, 1970]. However, more recent studies^{[Eckberg et al, 1985], [Taylor et al, 2001]} have shown that sympathetic activity also operates at the respiratory frequency and has a lesser, modulating influence on RSA. The nature of the LF peak is more ambiguous. Many take its power (area) to be an indicator of fluctuations in sympathetic activity,^{[Malliani et al, 1991], [Kamath & Fallen, 1993], [Rimoldi et al, 1990], [Montano et al, 1994]} while others believe that sympathetic and parasympathetic activity both contribute^{[Akselrod et al, 1981], [Appel et al, 1989]}. Hence, in some studies, the

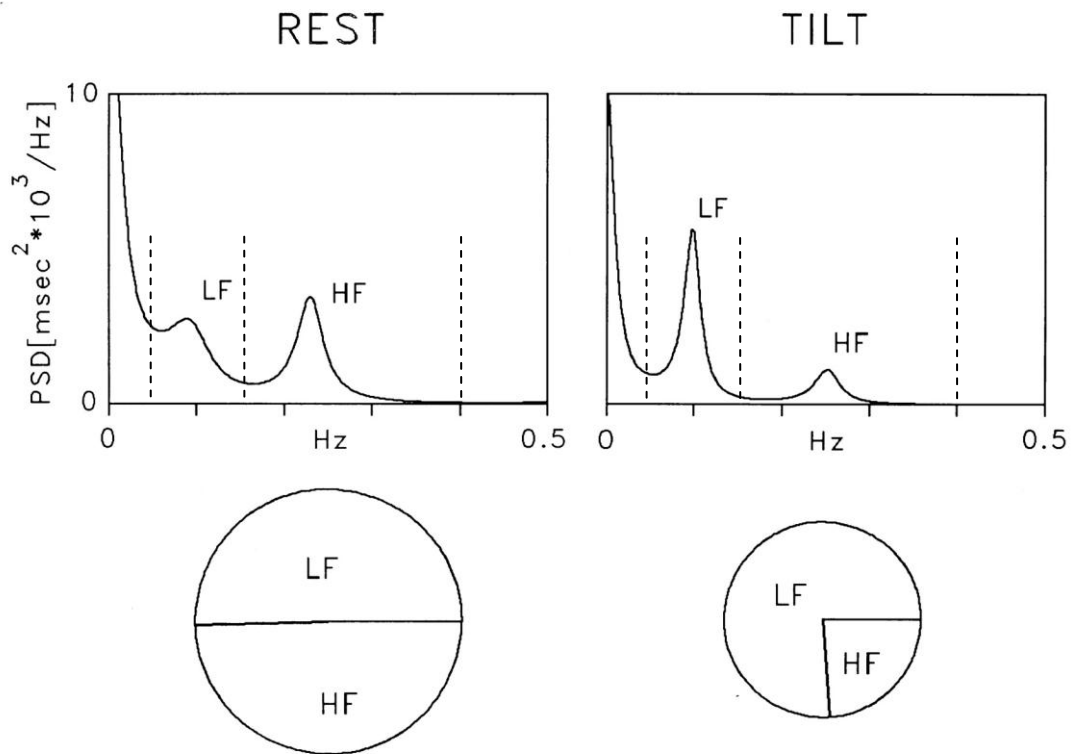


Figure 32: Adapted from [TFESCNASPE, 1996] with permission. Two RR interval power spectra, with clear LF and HF peaks, along with pie charts illustrating the LF/HF ratio. The change in the peak magnitudes reflects the ANS response induced by moving from a horizontal resting position to a 90 degree head-up tilt.

powers of the HF and LF peaks are taken as separate indicators of parasympathetic and sympathetic oscillations, while others use the ratio of the two peak powers, LF/HF, as an indicator of sympathovagal (sympathetic vs. parasympathetic) balance.

Levels of sympathetic and parasympathetic activity are known to vary in response to each other [Katona & Jih, 1975]. Upon collective consideration of this and the preceding observations, it seems reasonable to conclude that, rather than attributing a particular branch of the ANS to either the LF or the HF peak, one can assume that both branches exhibit a continuous frequency response to their respective inputs, with the sympathetic response predominant at lower frequencies (< 0.15 Hz) and the parasympathetic response predominant at higher frequencies (> 0.15 Hz). This paradigm is consistent with the observation that parasympathetic activity tends to respond more quickly than sympathetic activity to a step-change in stimulus. For example, the parasympathetic response to a sudden change in blood pressure occurs within 1 second, whereas the accompanying sympathetic response takes 2-3 seconds [Borst & Karemaker, 1983]. The LF peaks in RR interval spectra emerge as a resonant behaviour when this baroregulation system is modelled with parasympathetic components acting on shorter time-scales than sympathetic components [deBoer et al, 1987], [Deutsch & Deutsch, 1993].

Another important consideration when interpreting these spectra, is that the powers of the peaks do not necessarily correlate directly with the level of autonomic activity. They represent only the extent to which variation in heart rate can be attributed to ANS *variations* about a nominal baseline level [Malik & Camm, 1993]. That said, the LF/HF ratio is commonly found to agree with the presumed changes in sympathovagal balance according to changing physiological states in humans

and animals^[Malliani et al, 1991]. Also, in dogs the mean level of parasympathetic control has been found to correlate linearly with the amplitude of heart-rate variations at the respiratory frequency^[Katona & Jih, 1975]. However, some cases have been found to defy this correlation, with large HF oscillations observed while mean parasympathetic activity was assumed to be low^[Larsen et al, 2010].

A final consideration to note with regard to the interpretation of heart rate variability measures is that the variations in the LF and HF peaks are sometimes exposed only when normalised units (nu) are used to describe the power of each spectral component^[TFESCNASPE, 1996]. Without normalisation, the power is expressed in units of milliseconds-squared. When the same variable is expressed in nu, it represents the power of the spectral component relative to the total power of the spectrum, excluding VLF power (very low frequency, <0.04 Hz), as described in (15).

$$nu = \frac{ms^2}{total\ power - VLF\ power} \quad (15)$$

II.G.2. Existing methods for the measurement of HRV

There are various ways of calculating the spectra from which these measurements are taken. The choice of method is dependent on the specific nature of the RR interval series to be examined. Hence several popular techniques are outlined in this section. The most commonly used are the Fast Fourier Transform (FFT), Auto-regressive models, and wavelet transforms.

Fast Fourier Transform:

The familiar Fast Fourier Transform (FFT) is popular because it is computationally efficient and easy to implement. For a signal $x(n)$, the Fourier transform is defined as^[Lyons, 2004]

$$X(m) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi mn/N} \quad (16)$$

where N is the total number of samples in x , and m has integer values 0 to $N-1$.

The time-frequency uncertainty principle states that as the length of a signal decreases accuracy with which its frequency content can be defined also decreases. This principle can be illustrated by taking the Fourier Transform as an example. For a signal with sampling frequency f_s , the centre frequency of the m th frequency bin can be written as mf_s/N . When a signal component has a frequency kf_s/N , where k is not an integer, the component cannot be contained by a single frequency bin. Instead, its energy distributed across all bins with a weighting $|w|$ that is approximated by the sinc function, given in (17) ^[Lyons, 2004]. This phenomenon is known as ‘leakage’.

$$w(m) = \frac{N \sin[\pi(k-m)]}{2 \pi(k-m)} \quad (17)$$




Figure 33: A plot of the sinc function (equation (17)), which determines the amplitude of leakage in the m^{th} frequency bin for a sinusoidal signal component with a non-integer number of cycles, k , in the signal. Black dots indicate the amplitude of the function at the frequencies corresponding to integer values of m . When k is an integer these points align with the zero-crossings of the sinc function except at $k=m$.

As can be seen in Figure 33, the magnitude of this leakage tends to decrease for frequency bins further from k , although substantial variations can occur between neighbouring bins. By increasing the signal length N increases, these frequency bins draw closer together, such that the leakage is concentrated into a narrower range of frequencies. This is a manifestation of the time-frequency uncertainty principle. Capturing more cycles at each frequency yields a better-defined spectrum.

Auto-regressive (AR) spectral analysis:

A significant drawback of the FFT is that the amount of leakage into each frequency bin can vary substantially for small changes in the source frequency. As a result, the spectrum can contain peaks whose appearance is not easily interpreted.

An attractive alternative is auto-regressive (AR) spectral analysis. An AR model approximates a signal's behaviour, predicting each successive point based on the values of a few preceding points. The general formulation for an autoregressive model representing a series $x(n)$ is as follows:

$$x(n) = \sum_{i=1}^M a_i x(n-i) + \varepsilon(n) \quad (18)$$

... where M is the model order (the number of points used to predict the next) and the coefficients a_i are the weightings by which each previous point is multiplied to predict the next. The series $\varepsilon(n)$ is the model's error in predicting each n^{th} point. Strictly speaking, $\varepsilon(n)$ has no values

for $1 \leq n \leq M$. The coefficients a_i are chosen to minimise $\sum \varepsilon(n)^2$, then the model's power spectrum $P(\omega)$ can be calculated as

$$P(\omega) = \frac{\sigma^2}{\left(1 - \sum_{i=1}^M a_i e^{-j i \omega}\right)^2} \quad (19)$$

where σ is the standard deviation of the prediction error ε , and ω is any frequency expressed as a fraction of $f_s/2\pi$, f_s being the sample frequency in Hertz.

Unlike the FFT, this method does not constrain the spectrum to pre-selected frequency bins. Instead, a continuous spectrum is automatically calculated from the transfer function of the AR model, leading to “smoother” and more easily interpretable spectra. AR spectra and FFTs are similarly sensitive to signal length in the sense that using shorter signals broadens the peaks in the AR spectrum, much like it stretches the sinc function associated with FFT leakage (shown in Figure 33). Unlike the FFT however, the AR spectrum's magnitude does not vary erratically between neighbouring frequencies due to leakage (assuming an appropriate selection of model order, discussed in the next paragraph).

The validity of the AR spectrum is of course dependent on the validity of the AR model chosen. Of particular importance is the choice made for the ‘order’ of the model, the number of points used to predict the next point. Various methods are available to simplify and standardise the process of selecting the model's order, the most commonly used being the Akaike information criterion^[Takalo et al, 2005]. AR spectra are generally said to be preferable to FFT spectra for identifying central (dominant) frequencies when only a small number of data points is available, because the frequency resolution of the AR spectrum is not dependent on signal length^[Parati et al, 1995]. Hence they are slightly preferable in situations for which the quasi-stationarity of the signal is a tenuous assumption; using a shorter signal strengthens that assumption. Still, as signal length decreases the sensitivity of the spectrum to the choice of model order becomes increasingly important; if the order is too high spurious peaks may be produced, and if it is too low the spectrum will be over-smoothed, with separate peaks losing their distinction. [Kay, 1999] recommends that, for a signal of length N , all model orders in the range $N/3$ to $N/2$ should be tested. To ensure the model captures the signal's frequency content without bias, the prediction error $\varepsilon(n)$ should be required to pass a whiteness test. This test is implemented by examining the autocorrelation of the prediction errors. According to [Box et al, 2008], we expect this series to have zero mean and variance $1/n$, where $n = N-m$. We would then expect 95% of the autocorrelation series to lie within the bounds $\pm 1.96/\sqrt{n}$.

AR spectra are often used in HRV studies because they tend to yield clear peaks, allowing easier interpretation than spectra produced by the FFT. Figure 34 compares the results produced by both of these techniques when applied to two different RR interval series.

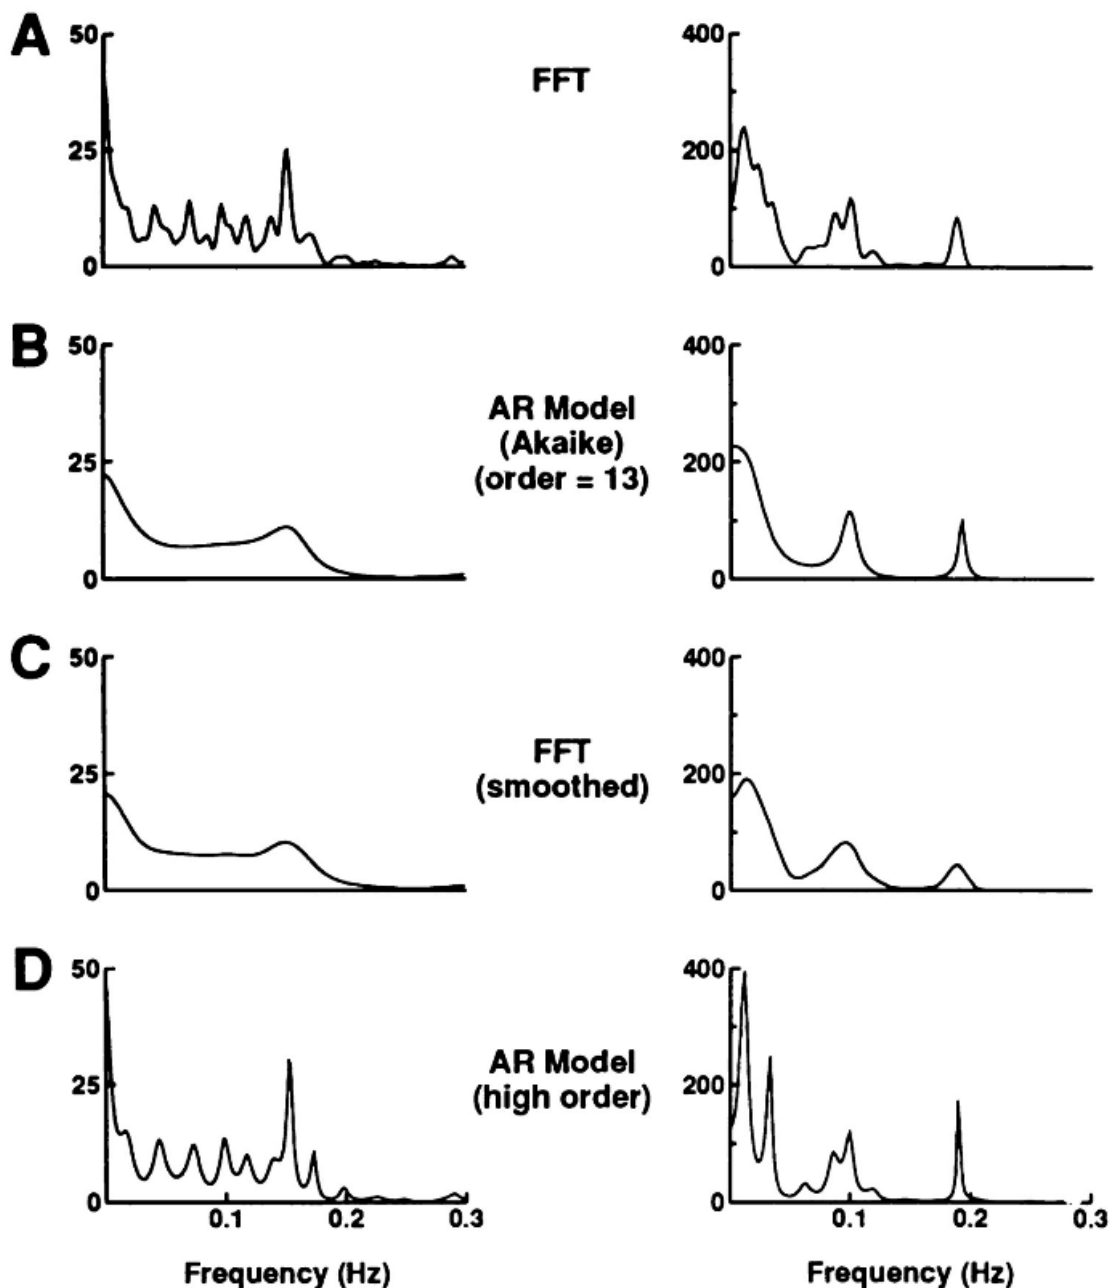


Figure 34: Reproduced from [Parati et al, 1995] with permission. AR and FFT spectra for two different RR interval series (left and right). An AR model with order selected by Akaike's Information criterion (order = 13 for both cases in row B) produces similar spectra to an FFT that has been smoothed with a Gaussian window (row C). When a higher model order is deliberately chosen (order = 30 for both cases in row D), the AR spectrum resembles the unsmoothed FFT (row A).

Wavelet Transform:

The effectiveness and applicability of both the FFT and AR models are heavily dependent on the assumption that the signal is 'stationary'. According to Malik and colleagues, this assumption is secure only "if mechanisms responsible for heart period modulations of a certain frequency remain unchanged during the whole period of recording"^[TFESCNASPE, 1996]. This condition implies that AR models and the FFT are not appropriate for measuring varying autonomic activity, which is one of the main justifications for using the wavelet transform instead of these other methods. The wavelet transform has no theoretical limit to its resolution in either the time- or frequency-domain. The basic method is to compare each section of the signal with a 'mother

wavelet' such as the Morlet wavelet or the Mexican-hat wavelet^[Mallat, 1999]. As shown in Figure 35, the wavelet can be stretched or compressed in time, allowing different frequencies to be examined.

The continuous wavelet transform of a function $x(t)$ is mathematically defined as

$$W(s, \tau) = \int x(t) \psi_{s,\tau}^*(t) dt \tag{20}$$

where * denotes complex conjugation. The basis function $\psi_{s,\tau}(t)$ is calculated from the mother wavelet as

$$\psi_{s,\tau}(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t-\tau}{s}\right) \tag{21}$$

It can be seen from (21) that τ imposes a translation of the mother wavelet along the time axis, while s stretches or compresses it along that axis. The factor of s^{-2} is used to normalise the




Figure 35: Reproduced from [Pichot et al, 1999]. Permission not required. A discrete wavelet transform of a hypothetical signal composed of combined low and high frequency components followed by the separated components. The horizontal τ axis is equivalent to time in the signal, while the vertical s axis is equivalent to period, with high-frequencies represented at low values of s . It can be seen that this representation of the signal's frequency content deals well with the non-stationarity of the signal. The mother wavelet used was the Daubechies 4 wavelet, which is shown on the right, compressed according to the different values of s .




Figure 36: A short example of a discrete event series representation of RR intervals, showing that it is irregularly sampled.

energy of the basis function^[Valens, 2004]. Figure 35 presents an example of the discrete wavelet transform applied to a hypothetical signal. The discrete wavelet transform differs from the continuous transform in that only a finite number of values for s and τ are tested.

The implementation of the FFT, AR spectra, and wavelet transforms all require some prior manipulation of the RR interval time-series to overcome the fact that it is irregularly sampled (see Figure 36). Typically, this step simply involves interpolating between the available data, with cubic-splines perhaps, then resampling this continuous function. For AR modelling, resampling is done at around 1 Hz or higher,^[Takalo et al, 2005] approximating the slow point-process nature of the original RR interval signal. For FFT and the wavelet transform, higher resampling frequencies are used to approximate a continuous signal. An additional benefit of the interpolation process is that it helps minimise the disruption caused by ectopic beats (introduced on page 27). These phenomena corrupt HRV measurements by appearing as large, short-lived changes in heart rate while not directly representing autonomic control of heart rate. Ideally, a signal used for HRV measurements should be free of ectopic beats, but rejecting all signals with ectopics from a study can often reduce the data set substantially and may introduce a population bias. Removing each ectopic beat and its immediate successor then interpolating across the resultant gap in the data series is a simple and effective way of minimising their distortion of the HRV measurement^[Parati et al, 1995].

Point-Process Modelling:

Despite its practical advantages, the use of interpolation to infer a continuous-time measure from the RR interval series is not physiologically meaningful. [Barbieri et al, 2005] suggest that a more meaningful continuous function can be achieved by recognising the point-process nature of the original series, the fact that it is a one-dimensional sequence of binary events rather than a two-dimensional series. Their approach treats the occurrence of heart-beats as a stochastic process. This treatment is justified because, according to their references, the drift of sino-atrial membrane potential towards its threshold value has previously been modelled as a Gaussian

random walk with drift. The probability density function of the time at which the next beat is expected to occur is modelled as a Gaussian distribution. At each point in time, the mean and standard deviation of this distribution are calculated from a set of continuously varying parameters, θ , which are estimated from the recent RR intervals, hence it is “history dependent”. θ is continuously varying because it takes into account the influence of the time elapsed since the last R-wave, even though the actual duration of the current interval is not yet known. Spectral components can be calculated from θ in the same way as they are calculated from an AR model’s parameters, but because θ varies continuously, this method does not assume stationarity. Furthermore, the “history dependent inverse Gaussian (HDIG) model” has a theoretical advantage over the wavelet transform in that it does not violate causality^[Napadow et al, 2008]; its value at any given time is dependent only on previous events. The method has been cross-validated against conventional HRV measurements^{[Barbieri et al, 2005], [Barbieri & Brown, 2006]}.

Spectral analysis using an HDIG model therefore seems to be a promising technique for observing short-term changes in autonomic activity in a way that is well-suited to the nature of the signal. However, in spite of its theoretical appropriateness this method has not yet been widely adopted, presumably due to its recent introduction and its mathematical complexity. To overcome these obstacles, it must be shown that the method offers distinct advantages over more widely used alternatives in terms of its accuracy as a reflection of direct measurements of ANS activity.

III. Methods for Correlating Psychological Processes and ECG Measurements

Sections III and IV of this thesis cover the methods developed over the course of the project to infer the nature of autonomic cardiac control from bioelectric signals.

As described in section II.F.3, a growing body of evidence implicates psychological and emotional factors in cases of Sudden Cardiac Death and various forms of ill health^[Taggart et al, 2011], but the mechanisms responsible for these psychosomatic effects are not well understood. The main objective of the work described in the present section was to develop analytic tools that would permit simultaneous measurement of electrocardiographic and cerebral responses to stimuli, thereby enabling novel research into the interactions between these two physiological systems. A more detailed set of objectives, targeting specific obstacles to progress in this field, is listed below.

- To develop novel signal processing methods to allow the use of ECG recordings in combination with continuous fMRI (functional magnetic resonance imaging):

Such studies are normally inhibited by the electrical artefact that fMRI induces in the ECG. The availability of such methods to overcome this problem would enable the use of ECG to record cardiac behaviour while fMRI is used to simultaneously assess brain activity, based on the level of blood flow to different areas of the brain. These efforts are discussed in section III.A.

- To determine whether it is possible to use spectral measurements of heart rate variability (HRV) to assess a subject's autonomic response to a psychological stimulus:

Psychological responses to any stimulus fade quickly, whereas HRV measurements require relatively long ECG signals. Efforts to overcome this apparent incompatibility are described in section III.B.1.

- To determine whether T-wave measurements can be used to assess autonomic input to the ventricles:

In essence, conventional HRV measurement techniques use the firing rate of the sino-atrial node as an indication of overall autonomic input to the heart. However, autonomic input to the heart may be heterogeneous, which has important implications for cardiac stability. This concept is central to the Brain-Heart Laterality hypothesis described in section II.F.3.b. Hence it would be useful to have a non-invasive measure of nervous input to other parts of the heart. The possibility of using measurements based on the T-wave to perform this role is explored in section III.B.2.

III.A. Simultaneous use of MRI and ECG recording

It is possible to noninvasively measure the spatial distribution of activity throughout the brain using positron electron tomography (PET) or functional magnetic resonance imaging (fMRI). These techniques can be used in combination with ECG recording as a means of investigating interactions between the heart and the brain based on their correlated response to various stimuli. fMRI is typically preferred over PET for several reasons, including superior spatial resolution. However, the concurrent use of ECG and MRI is hampered by the fact that the varying magnetic field of the MRI device produces an artefact in the ECG. This artefact may distort the ECG to the extent that automatic detection of features such as QRS complexes becomes much less reliable. Although the artefact can be attenuated by the use of conventional linear filters^[Critchley et al, 2003], its frequency content typically overlaps that of the normal ECG features. Hence it is not possible to remove the artefact by linear filtering without distorting the ECG substantially. Alternative methods were developed for the removal of the MRI artefact, exploiting the available *a priori* knowledge of the nature of the artefact.

In psychological experiments discussed in section V.B.3, an MRI machine was programmed to perform a full scan of the subject's brain every 3 seconds. Each full scan divided the brain into 45 adjacent slices, which were imaged consecutively. The frequency spectrum of the ECG artefact therefore contained significant components at frequencies of 15 Hz and higher. The important characteristics of an ECG signal are generally considered to exist within the 0.05-100 Hz band^[Imich, 1985], ^[Nearing et al, 1996]. These data confirm that there is significant overlap between the desired signal and the artefact in terms of frequency composition, ruling out the sole use of conventional linear filters. Two alternative approaches were developed. These are referred to as 'spectral subtraction' and 'the segment-mean technique', and are discussed in the following sections.

It should be noted that an additional artefact is produced by the flow of blood in the magnetic field; it can be shown from Faraday's Law that the flow of a conducting fluid in a magnetic field will produce an electrical potential^[Dimick et al, 1987]. This potential becomes superimposed on the ECG recording, distorting features such as T-waves. The form of this blood-flow artefact is less predictable, and therefore less easily removed, than the artefact caused by variations in the magnetic field alone. Nevertheless, useful measurements of QRS complexes can still be attained from these signals.

III.A.1. Spectral subtraction

This method requires a separate recording of the unwanted artefact only, without ECG features. The frequency spectrum (FFT) of this artefact signal is multiplied by a constant so that the dominant frequency component matches the magnitude of its counterpart in the spectrum of the original signal. The scaled artefact spectrum is then subtracted from the spectrum of the original signal, ideally leaving a spectrum that represents only the ECG components and random noise. The inverse Fourier transform (IFT) of this spectrum gives the ECG signal with the fMRI artefact removed. Figure 37 illustrates the implementation and results of this method. In this case, the "artefact-only" recording was obtained by positioning the ECG electrodes on the subject's leg and

carrying out an MRI scan of the leg. A quantitative analysis of the performance of the spectral subtraction technique is provided in Table 1 and Table 2 (page 86) for comparison with the performance of the segment-mean technique. It should be noted that there are several ways in which different implementations of the spectral subtraction technique could be improved or altered for other applications. These are discussed below, and the specific choices made for the present application are explained:

The artefact recording:

It is impossible to obtain an artefact-only recording that perfectly matches the artefact in the ECG recording, because the nature of the artefact is dependent on the positioning of the electrodes and the leads themselves with respect to the magnetic fields. The artefact is also dependent on the conductive properties of the subject's tissues. In this case, the artefact-only signals were collected by placing the electrodes on the subject's leg in a similar arrangement to that used on the chest, and running a separate MRI scan. The conductive properties of the leg are thus taken as an approximation of the chest.

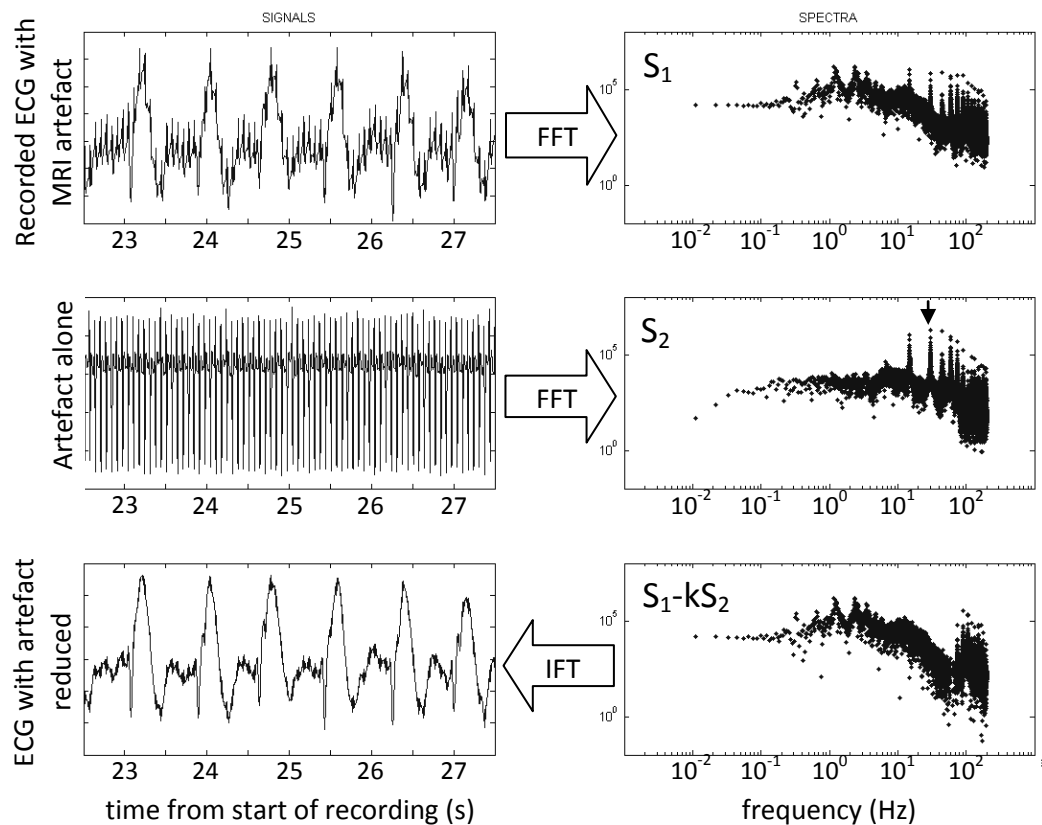


Figure 37: The application of the spectral subtraction technique to an ECG recording corrupted by MRI artefact. In the top-left panel, the presence of the artefact makes it difficult to distinguish crucial ECG features. As expected, the artefact is characterised by large components at 15 Hz and multiples thereof (harmonics). The artefact spectrum was scaled by k , the ratio between its dominant component (black arrow) and the same component in the original recording, then the scaled spectrum was subtracted from the original. The inverse Fourier transform was then used to generate a filtered signal. In the bottom-left panel, the artefact has been attenuated to the extent that sharp, downward QRS complexes can be identified. These are followed by a large blood-flow artefact, which cannot be removed.

Scaling:

The artefact signal was scaled so that the magnitude of the 15 Hz peak of the spectrum matched that of the 15 Hz peak in the spectrum of the raw signal, thus ensuring that these peaks cancel out completely. Unfortunately, the relative heights of the harmonic peaks are not the same in the artefact-only recording as in the raw signal, so other peaks are not completely removed, particularly the higher frequency harmonics.

Phase preservation:

The phases of the various frequency components of the artefact-only spectrum will differ the corresponding phases of the original spectrum. Hence one must subtract the magnitudes of the complex spectra, rather than the complex spectra themselves, to ensure that the artefact is attenuated. However, this step removes the phase information from the spectra. To obtain a meaningful signal from the inverse transform however, we must restore the lost phase information by assigning a suitable phase to each frequency component. The ‘cleaned’ signal in Figure 37 was produced by assigning the phase of each component of the raw signal’s spectrum to the corresponding component of the ‘cleaned’ spectrum. It should be noted, however, that these phase values are influenced by the MRI artefact in the original signal. The method still works reasonably well because the scaling mentioned above ensures that at most frequencies the artefact-only spectrum is much lower in magnitude than the original spectrum.

An alternative approach, which has not been explored, would be to use a cross-correlation to align the artefact-only signal in time with the artefact of the raw signal. Their spectral phases should then approximately match, so that the influence of the artefact on the phase of the original signal could be accounted for.

Averaging the artefact-only spectrum:

A more reliable impression of the artefact-only spectrum might be obtained by averaging it over several smaller time windows^[Vaseghi, 2000]. The potential benefits of this approach as a part of the spectral subtraction technique were not tested. However, the segment-mean technique exploits the same basic principle as Vaseghi’s suggestion but is confined to the time-domain.

III.A.2. The segment-mean technique

This method does not require a separate artefact-only recording. Instead, it takes advantage of the fact that the artefact has a known periodicity to identify it and then subtract it from the signal.

Figure 38 helps to illustrate this concept. The top panel shows a signal that is composed of two sine waves: one has a magnitude of 5 and a frequency of 1.3 Hz, and the other has a magnitude of 1 and a frequency of 15 Hz. Suppose that one wishes to remove any components with a frequency of $F_A = 5\text{Hz}$ or some multiple thereof, including the 15 Hz sine wave. To begin, the signal is divided into windows with a length of $1/F_A = 0.2s$. When these windows are averaged together, only the components that repeat identically in each window are preserved. All other

components cancel each other out if enough windows are used in the average. In this example, only the 15 Hz wave repeats identically in every window, so it is the only component preserved in the average, as seen in the second panel.

This component can then be removed from the original signal by simply subtracting the averaged window from each of the original windows. This is equivalent to subtracting a signal that is constructed by concatenating the window average onto itself multiple times, as illustrated in the third panel. The final panel shows the result of this subtraction. The 15 Hz component is no longer perceptible, and the signal appears as a pure 1.3 Hz sine wave.

As described, the only prerequisite information for this method is the period of the artefact to be removed. For the contaminated ECG signals seen earlier, choosing this period is not as straightforward as it might seem. If the artefact from each slice taken in a full scan can be assumed to be the same, then the period of the MRI artefact can be taken as 1/15 s. Otherwise, assuming that each full scan produces an identical artefact, the period of the artefact is 3 s. The artefact was taken as having a period of 1/15 s. This decision is justified by the observation that the signal spectra do not show any significant components at 1/3 Hz, implying that there is not a significant variation between the 45 spikes in a full scan.

Figure 39 presents an example of the segment-mean technique applied to the same MRI-corrupted ECG signal as used to illustrate spectral subtraction in Figure 37. For both methods, it is impossible to ascertain the extent to which the processing has distorted the desired signal, particularly when considering the unpredictable nature of the remaining blood-flow artefact, since




Figure 38: Periodic components of a signal can be removed by dividing the signal into windows of an appropriate length, averaging those windows to identify the periodic component, then subtracting that component from the original windows.




Figure 39: Results of the segment-mean method.

no comparable uncontaminated signal exists. Hence, in order to compare the performance of the two approaches, artificial signals were constructed by adding uncontaminated ECGs together with the “artefact-only” recordings from a patient’s leg, as used for spectral subtraction. Although these test-signals do not include any blood-flow artefact, they allow the methods’ performance to be assessed objectively because the processed signals can be compared with the original uncontaminated signal to ascertain the extent to which the MRI artefact has been removed and the extent to which the ECG features have been distorted.

III.A.3. Validation of the developed filtering techniques

Two variables were defined to quantify the methods’ performance: the residual variance (RV) and the harmonic peak reduction (HPR). The RV is the squared standard-deviation of the residual signal formed by subtracting the original ECG from the processed test-signal. RV is expressed as a percentage of the variance of the original ECG in order to normalise for differences in signal magnitude. HPR is the change in the summed power of the first five harmonics of the MRI artefact (i.e. 15 Hz, 30 Hz, 45 Hz, 60 Hz, 75 Hz), calculated from the FFT of the test-signal and the processed signal. It is expressed as a percentage of the summed power of those peaks in the test-signal. HPR is intended specifically to measure the attenuation of the MRI artefact, while ignoring any concomitant distortion of the signal’s other frequency components. RV provides a more general measure of the difference between the uncontaminated ECG and the processed signal.

The variables were calculated for a catalogue of artificial test signals. These were formed from uncontaminated 12-channel ECG recordings from 12 different patients, combined with 2 “artefact-only” leg recordings. For each of the 288 possible combinations of these signals (12 x 12 x 2 = 288), 3 different versions were constructed using different signal-to-noise ratios (SNRs). SNR was defined as the ratio of the standard deviations of the uncontaminated ECG and the “artefact-only” recording. The amplitude of the “artefact-only” recording was scaled linearly to achieve the desired SNR. Care was taken to ensure that the “artefact-only” recording used during the spectral subtraction method was not the same recording as used to construct the test-signal, as this could be expected to unfairly exaggerate the effectiveness of the technique.

Table 1 and Table 2 provide statistical comparisons of the performance of the two techniques across the full catalogue of test signals. These data confirms that both techniques are effective; the RV is consistently low, even for extreme noise-levels (SNR = 1), and the HPR is consistently high except when the noise level is low to begin with (SNR = 10). Figure 40 provides a visual impression of these values by presenting an example of the test-signal for each noise level and the corresponding results of each processing technique. Table 2 indicates that the performance of the two methods is almost identical in terms of HPR. However, Table 1 shows that in terms of RV the segment-mean technique provides a more faithful reproduction of the uncontaminated ECG in high-noise scenarios (SNR = 1), while spectral subtraction performs more favourably in low-noise scenarios (SNR = 10).

Residual variance (RV)

SNR	Technique	Mean	Median	Standard deviation	Min	Max
1	Segment-mean	1.67 %	1.66 %	0.21 %	1.15 %	2.30 %
	Spectral subtraction	4.30 %	3.88 %	1.67 %	1.68 %	9.10 %
3	Segment-mean	0.65 %	0.67 %	0.19 %	0.23 %	1.19 %
	Spectral subtraction	0.68 %	0.61 %	0.27 %	0.28 %	2.06 %
10	Segment-mean	0.56 %	0.57 %	0.19 %	0.14 %	1.09 %
	Spectral subtraction	0.19 %	0.17 %	0.09 %	0.07 %	0.77 %

Table 1: A quantified comparison of the performances of the segment-mean and spectral subtraction techniques, using the “residual variance” variable. Each of the statistical values is calculated from a set of 288 test-signals. SNR = signal-to-noise ratio, the ratio of the standard deviations of the uncontaminated ECG and “artefact-only” recording used to construct each test signal.

Harmonic Peak Reduction (HPR)

SNR	Technique	Mean	Median	Standard deviation	Min	Max
1	Segment-mean	99.56 %	99.57 %	0.20 %	99.07 %	99.95 %
	Spectral subtraction	99.62 %	99.63 %	0.18 %	99.10 %	99.96 %
3	Segment-mean	95.28 %	95.38 %	2.22 %	89.05 %	99.45 %
	Spectral subtraction	95.54 %	95.76 %	2.18 %	89.41 %	99.54 %
10	Segment-mean	34.76 %	2.53 %	37.60 %	-5.60 %	95.23 %
	Spectral subtraction	36.12 %	4.49 %	37.23 %	0.14 %	95.52 %

Table 2: A quantified comparison of the performances of the segment-mean and spectral subtraction techniques, using the “harmonic peak reduction” variable.




Figure 40: Examples of the segment-mean technique and spectral subtraction, used to remove MRI artefact from test signals with different signal-to-noise ratios (SNRs).

For the experimental application described in section V.B.3, the aim was to allow QRS complexes to be reliably identified even in the presence of large MRI artefacts. Custom software was developed for automated detection of QRS complexes in ECG recordings, as described in Appendix IV. For low-artefact signals, it was found that the complexes can typically be identified without any processing. Hence the segment-mean technique was deemed more useful for this project. If the aim was to characterise less distinct ECG features such as T-waves in the presence of small amounts of noise, the spectral subtraction technique would be preferable. Note, though, that the feasibility of such studies is doubtful, because neither processing technique would prevent the blood-flow artefact from obscuring the ECG's smaller features.

III.B. Cardiac correlates of psychological activity

III.B.1. Heart Rate Variability (HRV)

This section introduces the obstacles and caveats related to the use of HRV measures in psychological experiments. Taking these considerations into account, adaptations of conventional HRV methods are developed to allow wider use of HRV measures in experiments of this nature.

III.B.1.a. Habituation – A non-stationary process

It was mentioned previously that the most widely implemented HRV methods, namely the FFT and AR spectra, assume the RR interval series to be a stationary signal (i.e. frequency content does not change). This assumption is invariably false, but serves as an acceptable approximation as long as the physiological state does not alter substantially within the time-frame of the recording. When measuring the effects of physical stimuli (e.g. exercise, changes in body orientation) or states of consciousness (e.g. sleep vs. wakefulness) this condition can be satisfied. Appelhans & Luecken^[Appelhans & Luecken, 2006] summarise numerous studies investigating the HRV response to long-term psychological conditions. In experiments involving emotions, however, a sustained and consistent response is not easily generated due to a process known as ‘habituation’^{[Bradley et al, 1993], [Rankin et al, 2009]}. The psychological reaction to a stimulus is subject to continuous re-evaluation through interactions between different areas of the brain, making the response highly time-dependent. Repetitive stimuli may be used in an attempt to prolong the response, but it is generally found that any response to a repetitive stimulus attenuates over time.

The rate of attenuation cannot be formally predicted because it is highly dependent on the nature of the stimulus and of the response^[Rankin et al, 2009]. However, an approximate impression of the rate at which habituation occurs for a particular stimulus is reflected in the experimental procedures employed by experienced psychology researchers. For example, the experiments discussed in section V.B follow a procedure introduced by Phan and colleagues to study neural and psychological responses to graphic images^[Phan et al, 2005]. Phan *et al* present their subjects with images in blocks of 20 seconds duration, with the expectation that using longer blocks would cause the intensity of the measured response to be diluted by habituation. The experiments discussed in section V.B incorporate HRV measurements into the procedure used by Phan *et al* in order to assess autonomic activity. Ordinarily it is recommended that spectral HRV analysis be conducted on an RR interval series with a duration of two minutes or longer^[TFESCNASPE, 1996]. Clearly this guideline must be violated in the context of psychological experiments of this kind, and any such violation must be justified with consideration for the variables of interest and for the reasoning behind the guideline itself.

The reason for establishing a guideline minimum signal length is that capturing multiple cycles of a signal component attenuates the leakage between neighbouring frequencies (see page 74) and thereby achieves a more precise identification of the dominant frequencies. Close examination of the wording provided by TFESCNASPE (The Task Force of the European Society of Cardiology and the North American Society of Pacing Electrophysiology) in their justification for the 2-minute recommendation indicates that it is not a precisely defined limit:

“The recording should last for at least 10 times the wavelength of the lower frequency bound of the investigated component... Thus, recording of approximately 1 minute is needed to assess the HF components of HRV, while approximately 2 minutes are needed to address the LF component.”^[TFESCNASPE, 1996]

When HRV is used as an indication of ANS activity, TFESCNASPE themselves confirm that the frequencies of interest are the LF range (0.04 Hz – 0.15 Hz) and the HF range (0.15 Hz – 0.4 Hz). A signal component for which 2 minutes is “10 times the wavelength” would have a frequency of 0.083 Hz, more than double the “lower frequency bound” of the LF range. Similarly, 10 wavelengths in 1 minute would imply a frequency of 0.167 Hz, rather than 0.15 Hz. The imprecise nature of these guidelines suggests that it would be acceptable to breach them, provided that the likely consequences (leakage between neighbouring frequencies and unstable spectral peaks) are taken into account.

III.B.1.b. ‘Spectral Averaging’ to stabilise spectra from short-time recordings

Proposed technique:

In that same article, TFESCNASPE point out that “Averaging of spectral components obtained from sequential periods of time is able to minimise the error imposed by the analysis of very short segments”^[TFESCNASPE, 1996]. This technique is well known in the field of signal processing as the Bartlett Method^[Vaseghi, 2000]. It is subject to the caveat that the physiological scenario may change from one segment to the next, so that they do not each represent the same state.

The proposed technique extends the Bartlett Method to non-sequential periods of time, with the assumption that, in well-controlled experimental conditions, the spectral content within those periods does not vary significantly. For example, in psychological experiments where significant habituation is expected to occur in less than 1 minute, spectra from multiple 30 second blocks of stimulation can be combined by taking the average value of each frequency bin. Thus an ‘average spectrum’ is formed, providing a stabilised impression of the average response to the stimulus. This approach, referred to henceforth as ‘spectral averaging’, was implemented in several experiments described in section V.B and in a published journal paper^[Di Simplicio et al, 2011]. The inclusion of resting periods between stimulus blocks allows for some recovery of the response, as illustrated in Figure 41, so that the averaged response is higher when compared with that of a continuous stimulation period.

Given that spectra calculated from short-time signals provide a less reliable representation of a signal’s low-frequency content, one can expect that measurements of this low-frequency content will be less repeatable (more variable) across several similar measurements. This low-frequency variability, or instability, was observed in the experiments described in section V.B. Visual inspection of the spectra revealed that large peaks centred in the VLF range (<0.04 Hz) would occasionally appear. The power associated with these peaks would often overlap into the LF range, and sometimes into the HF range. An example is presented in Figure 42.




Figure 41: A conventional model of habituation in the biological response to a stimulus. *Top Panel* - The response decreases gradually as the stimulus is sustained. *Bottom Panel* - When the stimulus is applied intermittently, some recovery occurs. A longer rest period enables greater recovery, so that the magnitude of the response approaches that seen for the initial stimulus.




Figure 42: The solid line is the mean Power Spectral Density (PSD) calculated from four 30-second blocks for a single subject. The dashed lines indicate the maximum and minimum values of the PSD found at each frequency. To reduce the effects of habituation, the 30-second blocks were non-continuous (separated by intervals of 38 seconds) but coincided with similar experimental stimuli. The vertical dashed lines indicate the boundaries of the VLF, LF, and HF frequency ranges. In this example, it can be seen that large peaks centred in the VLF range contribute significantly to the LF range, and may also make small contributions to the HF range.

Validation:

Table 3 serves as validation of the spectral averaging method by showing that it yields clearer distinctions between HRV measurements associated with different experiment conditions. The low-frequency spectral instability caused by the use of short-time recordings can be seen as noise corrupting the HRV parameters. Spectral averaging provides a more stable measurement, counteracting the added noise and thereby allowing subtle stochastic phenomena to be discerned more clearly.

The experiment from which these data were produced is explained in detail in section V.B.3. In short, ECGs were recorded from 25 subjects while they were presented with slideshows of unpleasant images for periods of 30 seconds, interleaved with periods in which they rested for 30 seconds while staring at a stationary crosshair. The various standard HRV measurements explained in section II.G were calculated in two different ways: 1) using the HRV spectra calculated from only the first 30-second stimulation period and the first 30-second rest period, and 2) using the spectral averaging across four instances of the stimulation period and across four instances of the rest period.

Welch’s t-test^[Bain & Engelhardt, 2000] was used to evaluate the statistical significance of the differences between the ‘stimulation’ and ‘rest’ conditions, in terms of the HRV measurements. The output of a t-test is a p-value, which represents the probability that two distributions of measurements have the same mean. A p-value of less than 0.05 is conventionally considered to be statistically significant, since it suggests a 95 percent chance that the two distributions are different.

In Table 3, the p-values produced using the spectral averaging method are notably lower than those produced by using only the first instances of the ‘resting’ and ‘stimulation’ periods. These results indicate that spectral averaging improves the ability of these HRV measurements to discriminate between the ANS responses generated by the two experiment conditions. Hence spectral averaging is a more useful method for experiments of this nature.

III.B.1.c. Novel interpretations of HRV parameters

As described in section II.G.2, the normalised HRV parameters LF_n and HF_n are often used in place of LF and HF to account for the effects of broadband contributions to the spectrum. In the literature, the normalised variables are granted the same physiological interpretation as their non-

Statistical significance of differences in HRV measurements between experiment tasks

Method	VLF (0-0.04 Hz)	LF (0.04- 0.15 Hz)	HF (0.15- 0.4 Hz)	LF_n	HF_n	LF/HF
First instance only	0.84	0.098	0.88	0.020	0.041	0.11
Spectral averaging	0.45	0.001	0.075	<0.001	0.001	0.025

Table 3: p-values calculated using Welch’s t-test^[Bain & Engelhardt, 2000] to show the statistical significance of differences in the various HRV measurements between the ‘stimulation’ and ‘rest’ conditions. Cases deemed to be statistically significant ($p < 0.05$) are marked in bold. Overall, the statistical significances yielded by the spectral averaging method were stronger than those yielded by the use of only the first instances of the stimulation and rest periods.

normalised counterparts; HF_n , like HF , is taken as a reflection of parasympathetic activity, while LF_n , like LF , is seen as either a reflection of sympathetic activity or the combined effects of sympathetic and parasympathetic activity. The following analysis employs a mathematical approximation to expose the analytic relationship between the normalised variables and their non-normalised counterparts. Through this manipulation, it will be shown that the normalised HRV parameters can more meaningfully be interpreted as a reflection of sympathetic-parasympathetic balance with mathematical properties that are distinct from those of LF/HF (the conventional measure of sympathetic-parasympathetic balance).

The spectral power above the HF range is typically negligible; see, for example, Figure 32 (page 72) and Figure 42 (page 90). If the power above the HF band is taken to be zero then, from the original definition in (15) (page 73), LF_n and HF_n can each be approximated as

$$LF_n \approx \frac{LF}{LF + HF} \quad \text{and} \quad HF_n \approx \frac{HF}{LF + HF} \quad (22) \text{ and } (23)$$

Both measurements are clearly sensitive to changes in each of the frequency ranges related to ANS activity, but relating them to the sympathetic-parasympathetic balance is less intuitive than for LF/HF . The normalised variables' behaviour can be better understood by examining the partial derivatives of the approximations in (22) and (23). These are given below.

$$\frac{\partial LF_n}{\partial LF} \approx \frac{1}{LF + HF} \left(1 - \frac{LF}{LF + HF} \right) \quad \text{and} \quad \frac{\partial LF_n}{\partial HF} \approx -\frac{LF}{(LF + HF)^2} \quad (24) \text{ and } (25)$$

$$\frac{\partial HF_n}{\partial LF} \approx -\frac{HF}{(LF + HF)^2} \quad \text{and} \quad \frac{\partial HF_n}{\partial HF} \approx \frac{1}{LF + HF} \left(1 - \frac{HF}{LF + HF} \right) \quad (26) \text{ and } (27)$$

For comparison, the partial derivatives of LF/HF are:

$$\frac{\partial(LF/HF)}{\partial LF} = \frac{1}{HF} \quad \text{and} \quad \frac{\partial(LF/HF)}{\partial HF} = -\frac{LF}{HF^2} \quad (28) \text{ and } (29)$$

These last four relationships are plotted in Figure 43 and Figure 44. The nature of $\partial LF_n/\partial LF$ and $\partial LF_n/\partial HF$ can be inferred from the plots of $\partial HF_n/\partial HF$ and $\partial HF_n/\partial LF$, respectively, given the similarity of the equations.




Figure 43: Partial derivatives (sensitivities) of HF_n with respect to LF (upper panel) and HF (lower panel), when spectral content above the HF band is neglected. In both panels, the line $LF=HF$ is superimposed on the surface as a thick black line. In the upper panel, this line marks the peak in the partial derivative for each value of LF . In the lower panel, the line marks the peak partial derivative for each value of HF . LF is typically greater than HF , so most physiological cases can be expected to occupy the side of the line marked “ $LF>HF$ ”. See equations (26) and (27).

To achieve an impression of where physiological values of LF and HF typically fall on these graphs, consider that Bigger *et al.*^[Bigger et al, 1995] found, in a study of 274 healthy subjects, that the mean \pm standard deviation of LF/HF values was 4.61 ± 2.33 . In other words, LF is typically greater than HF , so most physiological cases can be expected to occupy the regions marked “ $LF>HF$ ” in Figure 43. This figure supports the proposal that LF_n and HF_n should be interpreted as measures of sympathetic-parasympathetic balance, rather than as separate indicators of the LF band and HF band, respectively; note that, for each pairing of LF and HF values within the typical physiological range, the sensitivities of HF_n to the two bands can be seen to be similar (roughly within an order of magnitude). The same holds for LF_n . In recognition of this relationship, neither of the normalised




Figure 44: Partial derivatives (sensitivities) of LF/HF with respect to LF (upper panel) and HF (lower panel). See equations (28) and (29).

variables can be said to be a focussed measure of their nominal frequency bands, especially when observing subtle changes evoked by stimuli.

Figure 43 and Figure 44 also demonstrate that the sensitivities of LF_n , HF_n , and LF/HF all vary substantially across the physiological range of LF and HF values, and the morphologies of the plots in Figure 43 are very different from the morphologies in Figure 44. Hence, while all three of these HRV parameters are indicators of sympathetic-parasympathetic balance, they are likely to provide different results in studies examining changes evoked by stimuli, because the parameters will have different sensitivities to those changes. The general profile of these the parameters' sensitivities should therefore be considered when analysing results from any such experiments. This principle is invoked in section V.B.6 in the context of an experimental application of the 'spectral averaging' approach.

III.B.1.d. The orienting response

Another important consideration when using HRV analysis in psychological experiments is the effect of the ‘orienting response’, a subject’s initial physiological reaction to a new stimulus^{[Sanchez-Navarro et al, 2006], [Adenauer et al, 2010]}. For example, the presentation of unpleasant pictures is known to cause a deceleration of mean heart rate, but this is preceded by a small deceleration then a sharp acceleration^{[Lang et al, 1997], [Sanchez-Navarro et al, 2006]}. Hence the signal within the first 4-10s of a new stimulus is highly non-stationary. The first few seconds after the onset of a new stimulus should therefore be excluded from the signal to which spectral HRV analysis is applied. Although methods such as the wavelet transform and point-process modelling are not as heavily dependent on the quasi-stationarity assumption as other methods, such as FFT and AR modelling, (see section II.G) no method can be reliably interpreted during step changes in ANS activity because such changes are not oscillatory in nature.

The orienting response can itself be used as an indicator of autonomic cardiovascular regulation. [Hodes et al, 1985] proposed that the response be quantified in terms of two simple indices, the minimum heart rate in the first 2 seconds after stimulus onset and the maximum heart rate in the next 2 seconds. These or similar indices have been shown to distinguish between control subjects and sufferers of panic disorders, including post-traumatic stress disorder^[Adenauer et al, 2010] and Generalised Anxiety Disorder^[Thayer & Lane, 2000].

Provided a suitably quasi-stationary segment of the signal can be extracted, AR modelling is a useful method when using HRV as a measure of psychological effects on ANS activity. As described in section II.G.2, AR spectral analysis yields spectra that are smoother, and thus more easily interpretable, than spectra produced by the FFT. Although smoothness itself is not a scientific goal, interpretability is an important consideration when results are to be presented to an audience outside cardiology and engineering, because many researchers will be unfamiliar with HRV measures or spectral methods in general. Hence AR modelling is the method of choice for the experiments presented later in this thesis, except where indicated otherwise.

III.B.2. T-wave measurements

Measurements of the ECG’s T-wave can be useful as an additional indication of ANS input to the heart, because they reflect the behaviour of the ventricles, rather than just the SA node. This level of information is crucial when assessing the arrhythmic potential of any input to the heart. Furthermore, it has been hypothesised that T-wave measurements might offer a means of testing the brain-heart laterality hypothesis (page 69), since sympathetic control of the ventricles is dominated by the left sympathetic branch, whereas sympathetic input to the SA and AV nodes comes from the right branch^[Lane & Schwarz, 1987]. This section considers the most widely used category of T-wave measurements, the QT interval, and assesses its appropriateness as a measure of ANS input to the ventricles.

III.B.2.a. QT intervals

It has been explained previously that, when analysing unipolar electrograms, ARI is measured to approximate local action-potential duration. When analysing the ECG, the QT interval is used to approximate global action-potential duration. It is well known that the heart exhibits spatial heterogeneity of action-potential duration^{[Weissenburger et al, 2000], [Okin et al, 2000]}. Nonetheless, the QT interval has been proven to be a useful diagnostic tool for identifying global abnormalities in repolarisation properties, whether their origins are genetic, hormonal, damage-related, or electrophysiological^[Kautzner, 2002]. The specific techniques used to define it vary but, in general, the QT interval refers to the timing between the QRS complex and the T-wave of the ECG. The technique that was ultimately chosen to measure QT intervals in the experiments described in section V was the Berger method^[Berger et al, 1997].

The Berger method characterises small changes in QT intervals by segmenting the ECG into individual beats and stretching/compressing the time axis for each beat to fit it to a template beat, as shown in Figure 95, page 227. The factor by which the time axis was multiplied for a fitted beat can then be multiplied by the template QT interval to give a representative QT interval for that beat. Other techniques are more widely used for diagnostic purposes, but are not well suited to measuring small, dynamic variations in action potential duration, for reasons explained in Appendix V. The Berger method was designed as a more accurate representation of repolarisation dynamics, taking into account the shape of the whole T-wave, and is therefore more appropriate for measuring the subtle changes associated with fluctuating input from the ANS. For completeness, Appendix V provides a more detailed description of the Berger method.

III.B.2.b. Correcting for heart rate

Action potential duration (APD) is known to be highly dependent on heart rate, with QT intervals shortening as heart rate increases (see section II.B). This relationship prevents healthy activation wavefronts from being broken by refractory zones when the cardiac cycle length is short, while allowing sufficiently long refractory periods to block arrhythmic wavefronts at longer cycle lengths. APD may also be modulated independently of heart rate, and many anti-arrhythmic drugs achieve their effect by adjusting the relationship between the two. The expected clinical convenience of being able to compare QT intervals measured at different heart rates inspired attempts to develop a variable QT_C , in which the QT interval is “corrected” to offset its heart rate dependence. For the present project, measurements of this nature were explored as a potential means of exposing the components of ventricular ANS control that are independent of ANS control of heart rate. Unfortunately, the implementation of QT_C measurements is confounded by a lack of consensus as to what model should be used to describe the QT -RR relationship. To resolve this issue a comparison of previously proposed models was carried out.

The most commonly used formulae for QT_C are the following:

$$\text{Bazett's formula}^{\text{[Bazett, 1920]}}: \quad QT_C = \frac{QT}{\sqrt{RR}} \quad (30)$$

$$\text{Fridericia's formula}^{\text{[Fridericia, 1921]}}: \quad QT_C = \frac{QT}{RR^{1/3}} + b \quad (31)$$

In all of the above formulae, QT and RR (cycle length) are to be given in seconds. RR is typically taken as an average of about ten RR intervals preceding the T-wave measured, because repolarisation properties do not adapt instantaneously to changes in heart rate. However, recent studies^{[Pueyo et al, 2004], [Malik et al, 2008]} have shown that the time-scale of this hysteresis varies substantially between individuals. Furthermore, Malik and colleagues found that where parametric models are fitted to individual QT - RR relationships, the parameter values differ substantially between subjects^[Malik et al, 2002]. Figure 45 illustrates the inter-individual variability of the QT - RR relationship.

Based on this recent evidence, it was decided that reliable assessment of any heart-rate-independent repolarisation changes would require subject-specific models of the QT - RR relationship. For each subject, on the advice of ^[Malik et al, 2008] the hysteresis in the relationship was accounted for by using \overline{RR} in place of RR , where \overline{RR} is a weighted average of recent RR intervals:

$$\overline{RR}(i) = \sum_{j=-N_i+1}^0 w(j) RR_{i-j} \quad (32)$$

$$\text{where } w(j) = K\gamma(1-\gamma)^{-j} \quad (33)$$

$$\gamma = \frac{\lambda}{1+N_i} \quad \text{and} \quad K = \frac{1}{1-(1-\gamma)^{N_i}} \quad (34) \text{ and } (35)$$

At each i^{th} beat, N_i is chosen as the number of beats in the trailing window of 200s. λ is a subject-specific parameter chosen to optimise the QT - RR relationship. For illustrative purposes, Figure 46 shows the weighting profile produced using $N_i = 251$ and three different values of λ .




Figure 45: Reproduced with permission from ^[Malik et al, 2002]. QT - RR relationships for six different healthy human subjects, demonstrating substantial differences between subjects.

A code to fit QT - RR models to experimental data was constructed using Matlab® (2007a, The MathWorks, Natick, MA, USA). This code contains the following steps, incorporating the above-mentioned hysteresis modelling and the optimisation scheme proposed by [Pueyo et al, 2004]:

- I. QT intervals were determined using the Berger method.
- II. \overline{RR} was determined using an initial assumption of $\lambda = 10$.
- III. The parameters of a chosen QT - RR model were determined by using Matlab's `fminsearch` function to minimise the residual mean-squared-error in the model's prediction of QT values from \overline{RR} values. The termination criterion for this optimisation was that the residual should converge to within a tolerance of 1 ms^2 .
- IV. The hysteresis variable λ was then optimised for the new QT - \overline{RR} model in the same way.
- V. The QT - RR parameters were optimised once more, using the new value of λ .
- VI. QT_C was then calculated by subtracting from each QT interval the value predicted by the optimised model. Note that this gives a series with a mean of approximately zero. For ease of interpretation, it is conventional to calculate QT_C in such a way as to give values in the range of expected values for QT at a resting heart rate. Hence a constant value, equal to the QT interval of the template used in the Berger method, was added to the series.

Note that steps III and IV could be repeated many more times to improve the accuracy of the relationship, but [Pueyo et al, 2004] found that the improvement from further iterations is not substantial.

Various parametric models were considered in order to describe the QT - \overline{RR} relationship. Two prior studies [Malik et al, 2002], [Batchvarov et al, 2002] compared a large number of models (6 and 10,




Figure 46: Typical weighting profiles for the recent history of RR intervals, used to account for QT - RR hysteresis. Refer to equation (32).

respectively), each having two parameters, and found no single model that provided the optimal fit for all subjects. Testing multiple models on every subject improves the likelihood of finding a close fit, but substantially increases the computational expense of the procedure. Furthermore, this approach prevents meaningful comparisons of parameter values across subjects and encumbers the communication of results.

To ascertain whether the use of multiple models was worthwhile for the purposes of the present project, the following test was carried out. A total of 16 models were fitted to the $QT-\overline{RR}$ data from 27 ECG recordings. These recordings were gathered during the psychological experiments discussed in section V.A. The quality-of-fit for each model was measured as the residual squared error after optimisation. The models tested are listed in Table 4. These include all of the models tested by Malik *et al* and Batchvarov *et al*, the more widely used Bazett and Fridericia models ((30) and (31)), the exponential parametric model proposed by Sarma *et al*, and three tri-parametric models that have not been previously proposed. The tri-parametric models were added to test the hypothesis that bi-parametric models are prone to over-constraining the $QT-\overline{RR}$ relationship. In particular, the cubic spline model was designed to avoid any unnecessary assumptions about the form of the $QT-\overline{RR}$ relationship. Such assumptions, implicit in the other

Model name	Description
Bazett ^[Bazett, 1920]	$QT = a\overline{RR}^{-1/2}$
Fridericia ^[Fridericia, 1921]	$QT = a\overline{RR}^{-1/3}$
Sarma ^[Sarma et al, 1984]	$QT = a + be^{c\overline{RR}}$
Linear ^[Malik et al, 2002]	$QT = a + b\overline{RR}$
Hyperbolic ^[Malik et al, 2002]	$QT = a + b\overline{RR}^{-1}$
Parabolic ^[Malik et al, 2002]	$QT = a\overline{RR}^b$
Logarithmic ^[Malik et al, 2002]	$QT = a + b\ln(\overline{RR})$
Shifted logarithmic ^[Malik et al, 2002]	$QT = \ln(a + b\overline{RR})$
Exponential ^[Malik et al, 2002]	$QT = a + be^{-\overline{RR}}$
Arctan ^[Batchvarov et al, 2002]	$QT = a + b\arctan(\overline{RR})$
Hyperbolic tan ^[Batchvarov et al, 2002]	$QT = a + b\tanh(\overline{RR})$
Inverse hyperbolic sine ^[Batchvarov et al, 2002]	$QT = a + b\operatorname{arcsinh}(\overline{RR})$
Inverse hyperbolic cosine ^[Batchvarov et al, 2002]	$QT = a + b\operatorname{arccosh}(\overline{RR} + 1)$
Parabolic with QT offset	$QT = a\overline{RR}^b + c$
Parabolic with RR offset	$QT = a(\overline{RR} + c)^b$
Three-point cubic spline	<i>Defined by three points with x-values (\overline{RR} axis) set to $\mu - \sigma$, μ, and $\mu + \sigma$, where μ is the mean \overline{RR} and σ is the standard deviation. The y-values (QT axis) are the parameters to be optimised. For each set of three points, the full curve can be described using Matlab's <code>interp1</code> function.</i>

Table 4: A list of models that were tested for their ability to describe the $QT-RR$ relationship.

models, reflect empirical observations but overstretch the present understanding of the physiological mechanisms that govern the true relationship.

For the purpose of comparison, each model's performance on each ECG recordings was quantified in terms of the percent difference in residual, $\% \Delta r$:

$$\% \Delta r = \frac{r - r_{opt}}{r_{opt}} \times 100\% \quad (36)$$

r is the residual (least mean-squared-error) of the model in question, and r_{opt} is the lowest residual out of all the models for the signal in question. The results of this comparison are summarised in Figure 47. This graph shows that the majority of the models perform similarly to one another. Understandably, the mono-parametric Bazett and Fridericia models showed the poorest ability to describe the various $QT-RR$ relationships.




Figure 47: A comparison of the performances of the 16 models listed in Table 4. The results are expressed in terms of the percent difference in residual, as explained in (36). Each column of the graph represents a distribution of 27 points, one for each signal. The red line is the median of the distribution, the blue box is bound by the 25th and 75th percentiles, and the whiskers give the maximum and minimum values.

The three-point cubic spline stands out as the most appropriate model; it presented the lowest residual in all but seven cases, where it differed from the optimal choice by less than 1 percent. This result suggests that the reasoning given for the spline model's proposal – that it would be more widely accurate because it avoids unnecessary assumptions about the overall form of the data – was sound. The three-point cubic spline might therefore offer an improvement on existing models to account for the heart-rate dependence of a QT interval series.

However, this test relied on a data-set that was limited in terms of the number of patients involved (27) and the range of the QT and RR data from each recording (patients were sedentary throughout). Hence a more extensive study will be required to confirm the advantage of the 3-point cubic spline model. Although the present test was useful in suggesting a direction for further work, it was decided that the spline model should not be used for the experiments described in section V.A due to the need for more extensive validation. In particular, the fact that no approximate physiological meaning can be attributed to the model's individual parameters makes it an unwieldy basis for discussions of individual patient characteristics.

To more closely compare the performances of the other models, the analysis of Figure 47 was repeated with the spline model excluded. These results are presented in Figure 48. The maximum values of $\% \Delta r$ are below 2 percent for 11 of the 15 models in this analysis. The fact none of these 11 models differed substantially from the others for any of the 27 recordings indicates that they are largely redundant with one another. Although no single model provides the optimal residual in every case, any of these 11 models is consistently adequate. Hence it was decided that the use of multiple models was not necessary to characterise the QT - \overline{RR} relationship; a single model could be applied in all cases, and any of these 11 would suffice. The Sarma model offered the lowest mean $\% \Delta r$. However, the Sarma model is not widely used and accepted compared with the Bazett and Fridericia models. Hence the model chosen for use in the experiments of section V.A was the second strongest performer, the parabolic model with QT offset, which is essentially a more flexible version of the Bazett and Fridericia models, which are already widely accepted.

III.C. Summary of achievements for chapter III

This chapter described several achievements relating to the stated objective, which was to develop analytic tools enabling non-invasive research into the interactions between cerebral processes and cardiac electrophysiology. These achievements are summarised below.

Removing MRI artefact from the ECG:

Two novel signal processing techniques (spectral subtraction and the segment-mean technique) were devised to extract the MRI artefact from a simultaneously recorded ECG signal. These methods were compared using a set of artificial test signals, and the segment-mean technique was deemed to be the more useful of the two, based on its performance in high-noise scenarios.




Figure 48: A repeat of the analysis in Figure 47, with the spline model excluded. Among the 11 most accurate of these 15 models, $\% \Delta r$ never differed by more than 2 percent.

Applied to experimental data, both techniques were found to attenuate the MRI artefact sufficiently to allow the use of automatic detection of QRS complexes. Hence they enable observations of heart rate and HRV, as surrogate measures of autonomic input to the heart, while fMRI is used simultaneously to record the brain activity associated with that autonomic state. An experiment of this kind is described in section V.B.3.

A major limitation of the developed techniques is that neither is capable of removing the distortion caused by the blood-flow artefact. Consequentially, they cannot be used with confidence to assess those ECG features that are less distinct than the QRS complex, such as T-waves, which would otherwise grant more detailed insight into electrophysiological behaviour than just the intervals between heart beats. Furthermore, it should be noted that measurement of these intervals concurrent with MRI might also be achievable using alternatives to the ECG, such as pulse oximetry or plethysmography, which essentially measure the blood flow associated with each heart beat rather than the bioelectric signal. Nevertheless, the wide availability of ECG recording devices constitutes a significant advantage of the proposed approach over such alternatives. Also, blood-flow-based

measurements are likely to become unreliable during arrhythmic episodes, when blood is not always being pumped effectively. The developed signal processing techniques therefore offer relatively convenient and reliable means of measuring heart rate and HRV during simultaneous MRI recording.

Overcoming the obstacle of habituation:

The technique referred to as ‘spectral averaging’ was developed as a means of enabling the use of HRV to assess cardiac autonomic input during psychological experiments. The habituation phenomenon typically attenuates psychological responses too quickly for reliable spectral measurements to be taken at the HRV frequencies. The spectral averaging approach overcomes this difficulty by combining spectra from several non-consecutive stimulation blocks in a way that counteracts the side-effects of using short RR interval series. Experimental applications of the method will be described in greater detail in section V.B. Results from these experiments were presented in the present chapter to demonstrate that, by reducing the random variations in the HRV measurements, the proposed technique increases the statistical significance of an observed difference in HRV response between two experimental cohorts. This result indicates that the method provides a more stable measure of HRV, as intended, and demonstrates the utility of the method for research purposes.

Novel interpretations of HRV variables:

Theoretical analysis was used to provide novel insight into the relationships between the primary HRV variables (LF , HF) and the secondary variables (LF_n , HF_n , LF/HF). It was shown that HF_n has significant sensitivity to LF (likewise, LF_n to HF). This observation warns against the conventional interpretation of the normalised variables as alternatives to LF and HF . They are more meaningfully interpreted as measures of sympathetic-parasympathetic balance since they reflect changes in both of the main HRV frequency bands. In that sense, they can be viewed as alternatives to LF/HF . When interpreting experimental values of LF_n , HF_n , and LF/HF , one must consider the varying sensitivities of the indices (as depicted in Figure 43 and Figure 44) and how they might effect the conclusions drawn. This caveat has special relevance to the experiments described in section V.B, as will be discussed.

Insights into model selection for QT-RR relationships:

Existing models of the relationships between RR intervals and QT intervals (i.e. between heart rate and ventricular action potential durations) were reviewed, with the aim of identifying a heart-rate independent measure of repolarisation behaviour that could be used to monitor autonomic input to the ventricles. The novel three-point spline model was found to consistently outperform the numerous previously proposed models in its ability to accurately describe the overall

relationship between QT and RR intervals. However, the spline model was noted as having a special limitation in terms of its utility in clinical and research implementations; the model is not conducive to discussions of individual patient characteristics because its individual parameters (the positioning of the points used to construct the spline) do not carry intuitive physiological meaning. It was found that the 'hyperbolic model with QT offset' could be used instead with little sacrifice in terms of the accuracy of the $QT-\overline{RR}$ relationships produced. This model could be readily received by the cardiology community as a 'modified Bazett formula'. The resemblance to the widely used Bazett formula is an important consideration given that the tests presented in this chapter involved a relatively small number of subjects and a small range of heart rates; the superior performance of the spline model cannot be assumed to extend to other physiological scenarios without more extensive testing.

IV. Using Unipolar Electrograms to Measure Dynamics of Electrophysiology

The previous chapter covered methods developed for the use of the ECG as a non-invasive indicator of the effects of autonomic activity on the heart. The practical convenience of the ECG makes it easily applicable to most psychological experiments. However, as discussed in section II.D, invasive measurements offer much more detailed information regarding the heart's local behaviour and the spatiotemporal distribution of activity. Invasive animal studies have granted valuable insights into the general nature of autonomic control of the heart, but these findings are inevitably subject to some uncertainty in the manner in which they translate to human physiology. Direct measurements from humans are therefore essential in exploring ANS input to the heart and the arrhythmic potential associated with that input. The scarcity of *in situ* human data has been cited as a major reason why pharmacological interventions have achieved only limited success in curing arrhythmia^[Nash et al, 2006]. The availability of such data is of course heavily limited by ethical considerations, but opportunities exist where data collection can be incorporated into surgical procedures without exposing the patient to any additional risk.

Unipolar electrograms (UEGs) are routinely used during a procedure known as 'ablation therapy' to target electrophysiological abnormalities and assess the effectiveness of the intervention. As described in section II.E, page 47, UEGs have been proven as a useful measure of the timing of activation and repolarisation of the myocytes in the vicinity of the exploring electrode. The following sections describe software-based tools developed to enable the practical use of UEGs for assessing dynamic changes in the underlying electrophysiological properties. These dynamics grant insight into autonomic input and other behaviours with relevance to the genesis of arrhythmia. The novel algorithms thus grant access to an abundant source of human *in vivo* data on localised effects of cardiac autonomic input, unlocking new methods of research into arrhythmogenesis and the causes of Sudden Cardiac Death.

The task of developing these analytic tools was broken down into the following objectives. Efforts to achieve these objectives are discussed individually in sections IV.A to IV.E. Section IV.F summarises the new abilities afforded by the algorithms and section IV.G describes the scope for future improvements.

Objectives:

1. Identify the obstacles that have thus far prevented widespread use of UEGs to infer autonomic input from cardiac electrophysiological dynamics.
2. Outline the design requirements for analytic tools to overcome these obstacles.
3. Develop a graphical user interface (GUI) to support the development of those tools.
4. Develop signal-processing algorithms to satisfy the established design requirements.
5. Demonstrate that the algorithms satisfy the requirements.

IV.A. Obstacles for the use of UEGs to measure dynamic behaviour

Any unknown system can be characterised by studying the dynamics of its response to a varying input^[Bendat & Piersol, 1993]. This is the principle underlying the use of heart-rate variability to assess autonomic input to the SA node, and the same principle can be invoked when investigating the role of ANS input to other parts of the heart, particularly the ventricles. Currently, the most common use of UEGs is for mapping the overall spatiotemporal distribution of steady-state behaviour and identifying the sites at which normal activity breaks down in pathological cases^{[Kimber et al, 1996], [Kusumoto, 1999], [de Bakker & Wittkampf, 2010]}. By extending the use of these electrograms to identify the dynamic mechanisms that are active in healthy subjects, it may be possible to illuminate the mechanisms behind otherwise unexplained cases of Sudden Cardiac Death (SCD). It must be noted, however, that the use of UEGs to study the dynamics of electrophysiological behaviour (rather than treating the properties as quasi-static) is subject to several additional considerations, as described in the following sections.

IV.A.1. The need for automation

As is the case for HRV analysis, the study of electrophysiological dynamics requires that recordings be sufficiently long with respect to the timescale of the dynamics of interest. Multiple electrodes must be used in order to differentiate between the behaviour at different sites. As a result, the datasets to be considered in these studies is inevitably large. For example, a relatively short recording of one minute using 20 electrodes, at an average heart rate of 75 beats-per-minute, will involve around $1 \times 75 \times 20 = 1500$ beats to be examined. Clearly the datasets involved in these studies will generally be so large as to be intractable without automated identification of activation and repolarisation times. However, the wide range of UEG morphologies typically encountered (see section II.E) makes automation a difficult feature-recognition task. The algorithms described in section IV.D overcome these difficulties to identify local activation and recovery times in UEG recordings, thereby enabling observations of the dynamics of local electrophysiological behaviour.

IV.A.2. The expected covariance between ARI and APD

The electrophysiological property that is most readily described as being a “local” property is the action potential duration (APD), since its value is largely independent of behaviour in neighbouring regions (although some interdependence does occur through the electrotonic effects described on page 61). The activation-recovery interval (ARI) was introduced on page 47 as a convenient reflection of APD that can be measured from UEGs. Hence the dynamics of ARI measurements are a potentially useful window on regionally-specific autonomic input to the myocardium.

Across large datasets, ARI (using the Wyatt method) and APD have been shown to be well-correlated, but errors for individual beats typically fall in the range ± 10 ms with much larger values occurring occasionally (see Figure 49)^[Haws & Lux, 1990]. After controlling for heart-rate, the subtle variations in APD caused by fluctuating autonomic input are expected to be similar in magnitude to




Figure 49: Reproduced from [Haws & Lux, 1990] with permission. ARI measurements from UEGs plotted against APDs measured from transmembrane potential recordings. The two are shown to be strongly correlated for measurements from 7 sites in 5 pigs during two different conditions (ischemia and normal myocardial perfusion).

these errors. When the object is to characterise the dynamic changes in APD rather than APD itself, these errors are acceptable as long as they present a consistent bias for consecutive beats in a particular recording. That is, ARI and APD must be strongly correlated at each individual site for small changes in their values (roughly ± 10 ms). If that is the case, the inconsistency of the errors seen in the relationship for broader datasets can be explained either by differences between individuals, between recording sites or between gross physiological states (e.g. changes in baseline ANS activity that are large compared to the small oscillations around the baseline). So one must ask: can we expect ARI and APD to be well correlated for short-term variations in APD in a single UEG recording?

Comparisons of ARIs and APDs across many consecutive beats at a single site are not available in the literature. Funding and time constraints prohibited the acquisition of such data in this project. However, the nature of the single-recording ARI-APD correlation can be inferred to some extent from the present understanding of the genesis of UEG morphology.

The strength of this correlation and its susceptibility to corruption will be dependent on the morphology of the electrogram itself. The UEG's 'remote component', as identified through the model of Potse *et al.* [Potse *et al.*, 2009], can be assumed to cause some deviation from a linear correlation. As explained on page 54, this influence causes ARI to slightly underestimate APD for positive T-waves and overestimate APD for negative T-waves. However, it was shown that the polarity of the UEG T-wave is determined by the timing of local activity relative to global activity, which can be assumed to not change substantially during small variations in autonomic input. Hence the induced non-linearity can be assumed to be small.

In general, the ARI-APD correlation will be most easily corrupted when the UEG morphology yields an unstable measurement of activation time or repolarisation time, as illustrated in Figure 50. According to the Wyatt method (page 44), these times are identified in the UEG as the steepest downstroke or upstroke, respectively, in the deflection of interest. These times are local

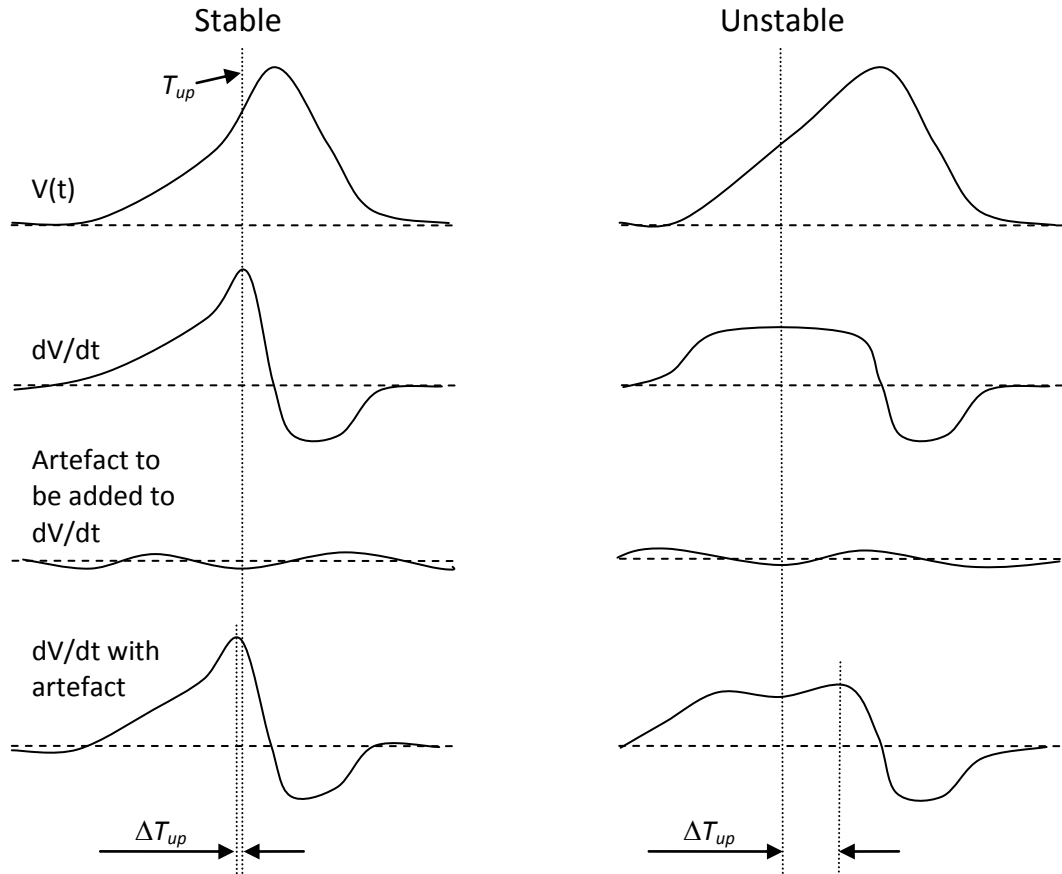


Figure 50: Assessing the stability of the point T_{up} as an indicator of local repolarisation time in the T-wave of the UEG. The signal on the right yields a less stable measurement because the slope of the T-wave does not change sharply in the vicinity of T_{up} . When a small artefact is added the timing of T_{up} changes substantially (ΔT_{up}) for the ‘unstable’ T-wave, whereas the change in the measurement for the ‘stable’ T-wave is negligible.

extrema in the signal’s time derivative, dV/dt . If dV/dt does not change sharply in the vicinity of these extrema, their timing will be more easily altered by the addition of noise, remote effects, or other artefacts.

If T_{up} or the activation time \dot{V}_{min} are unstable in the sense shown in Figure 50, the dynamics of the measured ARI (T_{up} minus \dot{V}_{min}) could contain noise of an amplitude large enough to obscure physiological dynamics of APD. An algorithm was developed to stabilise these measurements by taking into account the signal morphology surrounding the steepest point, as described in section IV.C.

IV.A.3. Dynamic artefacts

Given that, in certain cases, the dynamics of ARI measurements are easily corrupted by artefacts unrelated to changes in local APD, one must consider the expected nature of such artefacts and their manifestation in the ARI measurement. In particular, one must consider the risk that the dynamics associated with these artefacts will be mistaken for electrophysiological dynamics. Equation (6) (page 52) shows that the UEG is a summation of the dot product between the lead field and a dipole current-source distribution. Changes in the source term (or just the local component of it) can generally be treated as the dynamics of interest, while variations in the lead

field can be seen as a distortion of those dynamics. The lead field is dependent on the positioning of the electrodes relative to the myocardium and on the surrounding conductive properties, so variations in the lead field can be caused by variations in either of these factors. The importance of each of these two sources of artefact will be highly dependent on the experimental setup and on the nature of the dynamics of interest, so they should each be considered with specific regard for the experiment at hand.

IV.B. Summary of design requirements

IV.B.1. Fundamental requirements

To formalise the problem of developing tools to enable UEG-based investigation of cardiac electrophysiological dynamics, the following design requirements were identified:

IV.B.1.a. Robustness to unwanted artefacts

The task of automated UEG analysis is, in essence, a feature-recognition task. The electrophysiological artefacts of interest can be recognised based on empirical and theoretical understanding. However, the recordings also contain non-physiological artefacts, such as the deflections caused by a pacing stimulus used to artificially excite the heart during experiments. The presence of such artefacts complicates the task of identifying the features of interest, and it is essential that any automated analysis software should take such complications into account.

IV.B.1.b. Reliable identification of activation waves

All automated ECG analysis begins with the identification of QRS complexes, because these are the most distinct electrophysiological features. Once they have been located, other ECG features can be identified based on physiological expectations of their morphology and timing relative to the QRS. The activation wave can play a similar role in automated UEG analysis. Hence reliable identification of activation waves can be seen as essential to all other UEG processing tasks.

100 percent reliability is not a realistic goal, but a more reasonable minimum requirement can be established. As described in section IV.A.1, the principal objective of automated UEG processing is to make tractable the analysis of the large datasets required to study heterogeneous electrophysiological dynamics. Hence, an algorithm for automated identification of activation waves would be useful as long as the time taken to locate errors and apply manual corrections is significantly less than the time that would be taken to perform a fully manual analysis. Errors in activation wave detection can easily be identified by examining the time-series of intervals between activation waves. In these series, detection errors manifest as sharp discontinuities. As long as the majority of the waves are detected accurately, these discontinuities will be easily identified by brief visual inspection. By this reasoning, a minimum reliability criterion for the identification of

activation waves was established. Any such algorithm should regularly achieve 90 percent accuracy, using the following definition:

$$accuracy = \frac{N - (m + f)}{N} \times 100\% \quad (37)$$

where N is the total number of beats recorded, m is the number of missed activation waves, and f is the number of false-positive detections.

IV.B.1.c. Reliable identification of repolarisation times

As mentioned in section IV.A.2., ARI is a useful surrogate measure of APD, which can be seen as a localised electrophysiological property that is subject to autonomic modulation. UEG-based assessment of such modulation requires reliable identification of a repolarisation time corresponding to each activation time, so that ARI can be calculated as the interval between the two. No single, numerical criterion can be established for the reliability of repolarisation-time detection because the requirement will depend on the dynamics of interest and their sensitivity to detection errors. For example, in section IV.E the adequacy of the developed repolarisation-time detection algorithm is assessed in terms of how faithfully the spectral content of an ARI series can be reproduced, because spectral analysis is central to the experimental applications described in section V.C.

IV.B.2. Secondary requirements

In addition to the need for an accurate reflection of electrophysiological dynamics, the following requirements were considered during the design of the algorithm. A key concern throughout is the need for minimal human intervention, for two reasons. Firstly, any such interventions reduce the tractability of studies involving large datasets. Secondly, any need for intervention implies that the user should be trained to some extent in the use of this tool. Hence minimisation of the need for human intervention is essential to maximise the algorithm's potential as a more widely implemented investigative tool.

IV.B.2.a. Automatic adaptation to different signal magnitudes and morphologies

The QRS detection algorithm described in Appendix IV employed two user-defined parameters that were specific to each ECG signal to account for variations in signal amplitude and morphology. The magnitude of UEGs also varies between signals, depending on electrode size and positioning as well as on pre-processing parameters such as gain. To minimise human intervention and allow the algorithm to be implemented on a variety of hardware configurations (amplifiers, pre-processing filters, etc.), it was necessary that the algorithm adapt automatically to the magnitude and morphology of the signal, preferably in a time-varying manner to allow for any changes in electrode positioning or pre-processing parameters.

IV.C. Development environment

To facilitate the development process, a graphical user-interface named Wave_Inspector was designed using Matlab. Wave_Inspector, shown in Figure 52, incorporates the following key features to accelerate the process of testing and debugging new algorithms for electrogram analysis:

- Easy loading of new signals based on the patient identifier and electrode name.
- Selection of measurements to implement from a drop-down list.
- Display of electrograms and the measurement series ('Waveform' and 'Measurement', respectively, in F). A selection of buttons allows navigation within these windows (zoom in, zoom out, forwards, backwards, fast forwards (larger increments), and fast backwards). A specific time in the signal can be reached by entering a value (in seconds) into the 'Jump to' box. Red lines on the 'Measurement' window indicate the field of view of the 'Waveform' window. A green line indicates the timing of the measurement closest to the centre of the 'Waveform' window.
- Construction lines can optionally be drawn on the 'Waveform' window to validate measurements. In Figure 52, red lines indicate activation times and black lines indicate recovery times.
- Various filters can be quickly applied to the electrogram to examine their effects. The signal can also be differentiated or integrated any number of times.

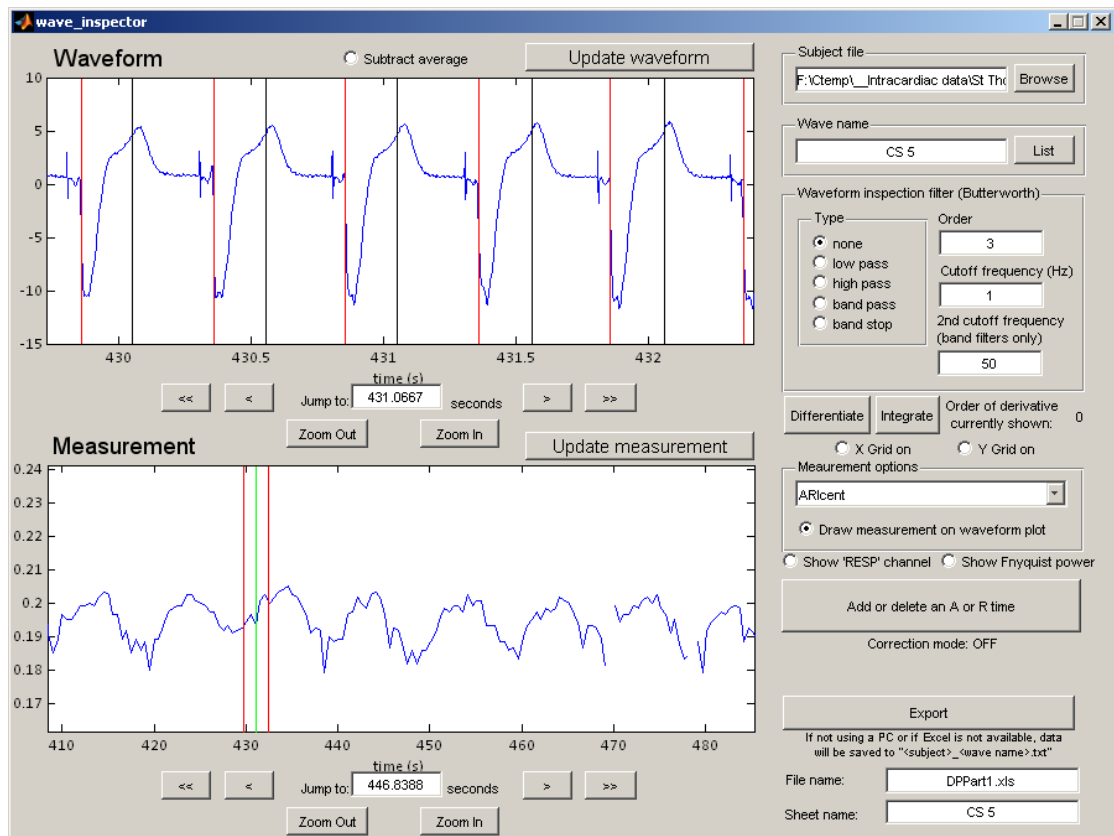


Figure 52: A screenshot of Wave_Inspector, the interface designed to facilitate testing of prototype algorithms for UEG processing.

- Individual measurements of activation time or repolarisation time can be manually added or deleted if necessary.
- Measured values can be exported to Excel or to Matlab's main workspace for separate manipulation or for permanent storage.

IV.D. Description of developed algorithms

Automated mapping of activation times from UEGs has been achieved with reasonable success and is widely implemented [Gepstein et al, 1997], [He et al, 1999], [Shenasa et al, 2009]. At least one prior attempt has been made to automate the measurement of repolarisation times from UEGs, [Witkowski & Penkoske, 1988] but such methods are not in frequent use today. The likely reason for this is that, while the Wyatt method is theoretically well-founded and works well in idealised scenarios, it is prone to large errors in non-ideal scenarios such as those described in section IV.A and in section II.E. Hence extensive manual corrections would typically be required. The following sections describe newly developed algorithms that minimise this problem and allow automated tracking of activation and repolarisation dynamics in the region surrounding an electrode. In doing so, the algorithms serve as a novel tool for the investigation of physiological mechanisms governing cardiac electrophysiology. The previously described requirements are satisfied by the combination of three separate algorithms that perform the following functions: removal of pacing spikes, identification of local activation times, and identification of local recovery times.

IV.D.1. Removal of pacing artefacts

The large artefacts produced by the pacing stimulus posed a problem with regard to the automated detection of activation times. The artefacts typically appear in the vicinity of the activation wave, and contain sharp deflections that are likely to be mistaken for the local activation time, \dot{V}_{\min} , when automated detection is used. Conventional filtering could not be used to attenuate the pacing artefact due to the overlap in frequency content between the artefact and the activation wave. Furthermore, the low-pass filtering used to reduce noise in the signal often broadened the artefact, causing overlap with the activation wave, as shown in Figure 53. Note that in this case a Butterworth filter has been used with a conservatively low order of 3 to minimise the spreading effect. The spread to the left can be easily discerned in the initial flat portion of trace B. It can be assumed that the affected area extends by roughly the same amount to the right of the artefact, but is less easily distinguished due to its overlap with the activation wave. This region is marked on the trace as the Area Of Influence (AOI). The use of Finite Impulse Response (FIR) filters was explored as a means of constraining the AOI, but it was found that a relatively high-order FIR was needed to achieve effective filtering, so the AOI could not be reduced. It was therefore deemed necessary to identify and remove the artefact before any further processing. Replacing it with a




Figure 53: *Top* – A 3.5 s segment of a UEG recording. The dashed box indicates the region shown in detail in traces A-D. *Trace A* – A pacing artefact and activation wave, with no processing applied. *Trace B* – Application of a low-pass Butterworth filter (3rd order, zero-phase, 100 Hz cut-off) does not effectively eliminate the pacing artefact and causes it to broaden and overlap with the activation time, potentially disrupting the detection of this event. The area in which the influence of the filtered artefact is perceptible is marked with a dashed line, which leads up to the original trace for reference. *Trace C* – An unfiltered version of the signal, with the pacing artefact removed by a custom algorithm. The area influence by the artefact (marked out by dotted lines) is noticeably smaller than in the filtered version, leaving a greater portion of the activation down-stroke unaffected. *Trace D* – The filter used to form trace B was applied to trace C. Because the pacing artefact was replaced with low-frequency content, the filtering does not noticeably influence the surrounding signal morphology.

segment of little high-frequency content minimises the effect of any subsequent filtering on the surrounding signal.

The task of identifying pacing artefacts can be simplified if the recording system is configured to also record from the pacing electrode. For each of the UEGs, the signal portion during which the stimulus is applied, plus a trailing segment to account for capacitance effects, can then be identified and replaced with low-frequency content. However, many of the signals used in this project were recorded before the need to remove pacing artefacts was appreciated. Furthermore, the requirement of specialised hardware configurations would limit the simplicity with which the developed algorithms could be adopted by other research groups. It was therefore

deemed necessary to develop an algorithm that would detect a pacing artefact from a single UEG, without additional information. This is achieved in the steps described below.

This algorithm and others described later in this chapter are presented as a sequence of data-manipulations and automated decisions. These processes make use of various numerical parameters, and the reasoning used to select their values is given where possible. However, it should be noted that in other cases the values were chosen by an iterative process of experimentation and refinement until they converged towards values that consistently performed the desired function without adverse side effects. In cases where no such converged parameter-value could be found, the associated decision-criterion was removed and alternative approaches were developed. This approach is referred to as heuristic development.

- I. **Filtering:** Create a heavily low-pass filtered version of the signal (3rd order zero-phase Butterworth filter, 12 Hz cut-off), segments of which will be used to replace pacing artefacts.
- II. **Differentiation:** Differentiate the original signal 8 times. The number 8 was chosen heuristically to give reliable discrimination between pacing artefacts and other deflections. The key point is that the pacing artefact is distinct from the rest of the signal in that it contains very sharp deflections. Each successive differentiation increases the amplitude of the pacing artefact relative to the rest of the signal. This approach was found to limit the spread of the artefact more effectively and more predictably than high-pass filtering.
- III. **Threshold:** A magnitude threshold is established as 2.5 times the 99th percentile of the differentiated signal. Again, these values were determined heuristically.
- IV. **Windowing:** When a point in the differentiated signal is found to have an absolute value greater than the threshold, the point with the maximum absolute value in the subsequent 0.01 s is located. A window of 0.012 s either side of this point is scanned and the first and last points to exceed 1/15 times the maximum are taken as the start and end of the pacing artefact. These boundaries are extended by 0.002 s and 0.004 s, respectively, to ensure that the artefact is well encapsulated.
- V. **Exclusion criterion:** To ensure that a sharp ‘intrinsic deflection’ associated with local activation is not mistaken for a pacing artefact, this window is dismissed if the overall change in the original signal (difference between values at the start and end of the window) is greater than 75% of the range in that window. In other words, the contained signal must not have a substantial net slope.
- VI. **Transplant:** If the window is not dismissed, the corresponding section of the low-pass filtered signal is ‘transplanted’ into the original signal. A linear component is added to the transplanted section to ensure that no discontinuities occur at the boundaries.

The full Matlab code for this algorithm is provided in Appendix IV. Figure 53 demonstrates the effectiveness of the algorithm in replacing the pacing spike while having minimal influence on the surrounding signal. An alternative approach was explored previously, using pattern-matching with an *a priori* assumption of the artefact’s essential biphasic morphology. This approach was found to be less reliable because capacitive effects unpredictably reshaped the tail-end of the artefact. Capacitive effects of body tissues are typically neglected in electrogram analysis, because it has been shown that they are not large for the frequencies associated with the electrophysiological sources^[Plonsey, 1969]. However, the pacing artefact contains substantial energy at higher frequencies, for which capacitive effects are no longer negligible.

A limitation of the multiple-differentiation approach described above is that it is less reliable than the pattern-matching technique for identifying very small pacing artefacts. However, this is not a functional limitation, because any artefacts small enough to avoid detection were found not to be problematic for activation detection. When it is absolutely necessary to locate artefacts (e.g. in order to assess the interval between pacing and local activation), the timing of the artefact can be taken from the electrode in which it is most prominent.

IV.D.2. Identification of activation waves

As mentioned previously, the point \dot{V}_{\min} (the steepest downward slope in the UEG’s activation wave) is widely recognised as an index of the timing of local depolarisation. Several authors^{[Witkowski & Penkoske, 1988], [Anderson et al, 1991], [Ndrepepa et al, 1995], [He et al, 1999], [Klemm et al, 2007]} have described the automated detection of local activation times from UEGs, but all of these approaches use simple detection criteria (usually, the comparison of the signal’s derivative with a threshold) and depend on manual exclusion of unreliable signals or basic automatic criteria such as “rejecting signals that appeared as flat lines or that fluctuated between minimum and maximum saturation values”^[Ndrepepa et al, 1995]. Bhakta & Miller^[Bhakta & Miller, 2008] summarise commercially available activation mapping systems, including the widely used Carto® (Biosense Webster Inc., Diamond View, CA, USA) and EnSite® (St. Jude Medical Inc., St. Paul, MN, USA) systems but, again, all of these rely on the operator to exclude signals and signal portions that are not well suited to analysis. Furthermore, they are intended for spatial mapping of activation times rather than close inspection of beat-to-beat variations, which they neglect or smooth out.

An algorithm was designed to identify local activation times without the need for extensive manual reduction of the dataset or automatic rejection of signals with only brief corrupt sections. This was achieved by incorporating novel criteria to automatically distinguish between activation waves and other large, steep deflections, as explained in the algorithm’s outline below. These characteristics enable the algorithm to quickly produce long sequences of activation time measurements without human intervention. A 500-second UEG recording could typically be processed in less than 15 seconds on a desktop PC (Windows XP, 2.8 GHz CPU, 2.98 GB RAM), thus satisfying the speed criterion established in section IV.B.2.c (page 111). Modification of the

\dot{V}_{\min} definition of local activation time, in order to adapt the measurement to time-varying properties, is explained in section IV.D.4. The full Matlab code for this algorithm is provided in Appendix V.

- I. **Filtering:** The pacing artefacts are removed from the signal, as described in section IV.D.1, then a bandpass filter (4th order zero-phase Butterworth) is applied. The lower cut-off frequency was chosen as 0.5 Hz, as has previously been used by experienced researchers and clinicians^{[He et al, 1999], [Stevenson & Soejima, 2005]}. Several researchers^{[Witkowski & Penkoske, 1988], [Anderson et al, 1991], [Ndrepepa et al, 1995]} have advocated smaller values for the lower cut-off frequency, but for dynamic analyses it is essential to attenuate the respiratory motion artefact (typically 0.15-0.4 Hz). The upper cut-off was chosen as 150 Hz, which was found to achieve an acceptable compromise between the two goals of attenuating high-frequency noise and avoiding excessive distortion of the intrinsic deflection in the activation wave; the intrinsic deflection typically has the highest frequency content of all the UEG features with electrophysiological origins. A second copy of the signal is produced with bandpass cut-off frequencies of 0.5 Hz and 30 Hz. This copy is used to assess the broad morphology of the signal while strongly attenuating any high-frequency artefacts. For clarity, the 0.5-150 Hz signal will be referred to as the main signal and the 0.5-30 Hz signal will be referred to as the low-passed signal. It should be noted that, once the algorithm for removal of pacing artefacts was fully developed, it was found that the main signal could also be used in place of the low-passed signal for most cases with no noticeable effect on the reliability of activation detection.
- II. **Scanning:** The main signal is scanned in overlapping ‘scanning windows’ of 0.15 s width at intervals of 0.01 s. This window size was chosen to be as small as possible while reliably encompassing the full duration of ventricular activation (see Figure 54), even when the use of artificial pacing slowed the wave of activity by not immediately employing the specialised conduction cells. To identify activation waves, these scanning windows are compared against a ‘broad window’ in the low-passed signal (see Figure 54 and Figure 55 (page 119)). Under the assumption that activation waves are the largest features in the local signal, the voltage range of the broad window provides an order-of-magnitude approximation of the peak-to-peak magnitude of an activation wave (including those slightly smaller than their neighbours).
- III. **Detection criteria:** Hence the first criterion for activation detection is that the voltage range of the downward sloping portions of the scanning window must exceed 0.3 times the range of the broad window. This broad-window index is only recalculated once every




Figure 54: A UEG with vertical lines superimposed to mark activation times. The seven activation waves in view were all detected successfully, despite the fact that arrhythmia and intermittent pacing have caused substantial variations in signal morphology. Dashed outlines and grey-filled boxes show two typical examples of the scanning window and broad window, the contents of which are compared during activation detection. The heights of the windows indicate the range of the data contained within them.

0.2 s to avoid excessive computational expense. Ideally, the broad window would contain only the next single activation wave, so that the algorithm would not be at all affected by beat-to-beat changes in signal magnitude. The fact that activation times have not yet been detected makes this impossible. To ensure that the index adapts quickly to changes in signal magnitude, the boundaries of the broad window are defined as follows:

- The start point of the broad window is usually chosen as 0.09 s after the last activation time, at which point that activation wave should have almost completely subsided. Excluding the last activation wave improves the algorithm's speed of adaptation. If no activation time has been detected in the last 1 s, the broad-window start point is chosen as 0.3 s prior to the current point in the scan.
- The end point of the broad window is set as 0.6 s ahead of the current point in the scan.

The scan restarts 0.15 s after each detected activation time, so the minimum possible width of the broad window can be calculated as $0.6 \text{ s} + [0.15 \text{ s} - 0.09 \text{ s}] = 0.66 \text{ s}$. Because some of the experiments involved in this project used high (paced) heart rates, this minimum was sufficient to ensure that the broad window usually contained at least one activation wave. However, the algorithm was also reliable for longer beat-to-beat intervals; additional criteria (described below) successfully eliminated false detections that would otherwise have been caused by an erroneous assessment of the local signal magnitude. In implementations where heart rate is expected to regularly drop below ~ 55 beats-per-minute, it might become necessary to change the minimum broad-window length to approximately 2 s for extra

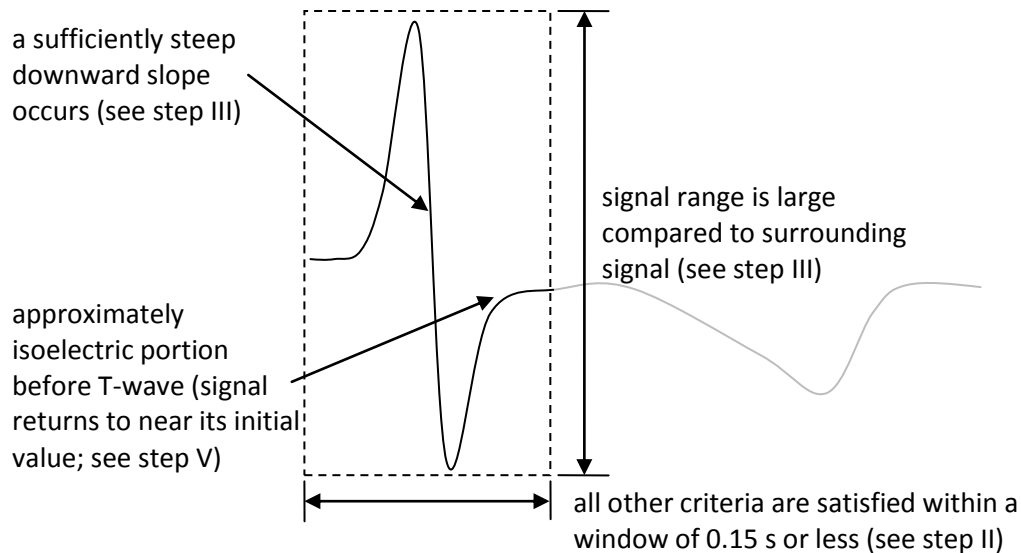


Figure 55: Qualitative descriptions of the various criteria applied in the scanning window to identify activation waves. Note that the criteria are all applicable to the full range known activation wave morphologies (positive, negative, biphasic, fractionated, etc.). Hence the algorithm is robust to dynamic changes in morphology.

insurance against false-positive detections. However, this adjustment would worsen the algorithm's adaptation speed, possibly resulting in missed activation waves just before a sharp increase in signal amplitude.

- IV. **Window refinement:** If the previous criterion is satisfied (range of downward portions of scanning window exceeds 0.3 times the range of the broad window), then the scanning window is refined; the start and end boundaries are pulled in to exclude signal segments whose deflections are negligible compared to the rest of the window (range < 5%). As well as being important in identifying activation waves themselves, well-defined boundaries of the activation are useful when searching for other electrogram features such as recovery times, as described in the next section.
- V. **Further detection criteria:** After refinement of the scanning window, the range of its downward portions must still exceed 0.3 times the broad-window range. Also, the difference in voltage between the first and last points in the window must be less than half the range covered in between. In other words the signal must return to a very rough approximation of its initial value. Finally, the minimum slope in the window must be more negative than 1st percentile of the slopes in the broad window.
- VI. **Further window refinement:** If all criteria have been satisfied, this window is deemed to contain an activation wave. To ensure that the most appropriate choice of scanning window for this beat is found, the scan continues until the window has moved on by 1.5 window lengths, using a finer time-increment (~0.004s). All candidate windows fulfilling the selection criteria are noted so that the one that best encapsulates the activation wave can be chosen. All candidate windows with a range less than 95% of the maximum range are

immediately rejected. Of the remaining windows, the ones with the largest range in the downward sloping segments are selected. If more than one candidate remains, the final choice is the window with the smallest difference between the first and last points.

- VII. **Activation time selection:** Within this chosen window, the local activation time can be selected as \dot{V}_{\min} . However, an improvement upon this index is discussed in section IV.D.4.

IV.D.3. Identification of local recovery times

As described in section II.E, the steepest upward slope on a T-wave (T_{up}) is widely used to indicate the timing of local repolarisation in UEGs. As for activation times, computers are often used with human guidance to locate the steepest slope without introducing operator bias. [Witkowski & Penkoske, 1988] describe an algorithm that incorporates morphological analysis of the surrounding signal to locate T-waves. However, the algorithm implicitly assumes that a clear T-wave will be present in each beat, without substantial corrupting artefacts, which is not always the case.

This section describes the key components of an algorithm to identify local repolarisation times while automatically excluding beats that are unlikely to yield reliable measurements. An earlier version of the algorithm was described in [Western et al, 2010]. The full Matlab code for this algorithm is provided in Appendix VIII.

- I. **Filtering:** The pacing artefacts are removed from the signal and activation times are identified using the previously described algorithms before the recovery-time detection algorithm is implemented. A bandpass filter (3rd order zero-phase Butterworth) is applied with cut-off frequencies of 0.5 Hz and 35 Hz. This upper limit was chosen to reduce noise and the 50 Hz mains artefact as much as possible without being so low as to smooth out key physiological deflections in the T-wave. It was originally set at 30 Hz, but was raised when a single case was found in which excessive smoothing caused erroneous identification of T_{up} .
- II. **Scanning:** As in the activation wave detection algorithm, a ‘broad window’ is established as a representation of the local signal magnitude, against which individual features can be compared. For each beat-to-beat interval, the start point of the broad window is defined as 0.1 seconds before the first activation time. The end point is chosen as either 1 second after the first activation time or the end of the next activation wave, whichever is earlier.
- III. **T-wave windowing; end limit:** Next, the algorithm establishes boundaries within which the recovery wave (T-wave) is expected to occur. The start boundary is taken as the end of the preceding activation wave. The end boundary, T_{end} , is taken as the earliest of four candidates. The first candidate is the start of the subsequent activation wave. The other three candidates are heuristic formulae based on the cycle length of the preceding beat,

$CL(i)$. Where $A(i)$ is the i^{th} activation time, $CL(i)$ is measured in seconds from $A(i-1)$ to $A(i)$.

$$T_{\text{end}}(i) = A(i) + 0.8 \times CL(i) \quad (38)$$

$$T_{\text{end}}(i) = A(i) + 0.5 \times \sqrt{CL(i)} \quad (39)$$

$$T_{\text{end}}(i) = A(i) + 0.6s \quad (40)$$

These formulae, plotted in Figure 56, ensure that the T-wave is only sought in the period in which its appearance is physiologically appropriate. Note that (39) can be loosely interpreted as a limiting value of 0.5 s for the QT_C variable, calculated using Bazett's formula^[Bazett, 1920] ((30), page 96). (38) and (40) serve to correct this relationship for extreme values of CL.

IV. **T-wave windowing; start limit:** As an approximate representation of the ST segment, a single point, t_{ST} , is identified. Within a window of 60 ms from the end of the activation-wave, t_{ST} is defined as the last point at which the difference between the actual slope and the mean slope is less than one hundredth of a standard deviation. This definition ensures that t_{ST} occurs after any trailing features of the activation wave but before any substantial components of the T-wave. t_{ST} is then incrementally shifted later in time until the absolute value of the signal's 1st derivative reaches a local minimum (i.e. the signal appears to be roughly isoelectric).




Figure 56: Heuristic formulae used to establish $T_{\text{end}}(i)$, the end of the window in which the T-wave should be located. Each line is labelled with the number of the equation used to construct it. The shaded area represents the time after $A(i)$ in which the T-wave can occur.

V. **Identification of T_{up} :** The search-window for the time of local repolarisation extends from t_{ST} to T_{end} . The point T_{up} is identified as the time of the maximum slope within this window.

VI. **Rejection criteria:** Several criteria are applied to assess the reliability of the measurement of T_{up} :

- a. The slope at T_{up} must be positive.
- b. The range of voltages covered by the signal within 30 ms either side of T_{up} must be at least 0.2 times the range between t_{ST} and T_{end} .
- c. The range between t_{ST} and T_{end} must be at least 0.015 times the range in the broad window.
- d. T_{up} must not occur immediately at the start of the search-window, since this might not be a true local maximum in the slope.

If any of these criteria are not met, the recovery time for this beat is designated as ‘not found’.

VII. **Secondary rejection criteria:** If the above criteria are all met, a second set of criteria is applied:

- a. T_{up} must not occur at the very end of the search-window, since this might not be a true local maximum in the slope (similar to condition d. above).
- b. At some point between T_{up} and T_{end} , the slope must return to below a threshold equal to the range in the broad window divided by 0.5 s.

Failure to meet either of these criteria suggests that the point T_{end} may not have been suitably defined, hence an attempt is made to shift T_{end} earlier, to the next point at which the slope criterion described in condition b, above, is satisfied. This action sometimes excludes a non-repolarisation/non-local feature with a steep upslope from the search-window so that it is not mistakenly identified as the local-repolarisation deflection. If the newly defined T_{end} does not allow all of the above criteria to be met, the recovery time for this beat is designated as ‘not found’. Otherwise, T_{up} can be taken as the local recovery time. Modification of the T_{up} definition of local recovery time, in order to adapt the measurement to time-varying properties, is explained in section IV.D.4

IV.D.4. Stabilisation of local event times

As explained in Figure 50 (page 108), certain UEG morphologies can yield unstable measurements of activation and recovery times. This instability can obscure subtle repolarisation dynamics when ARIs are used as a surrogate measure of APD. An example of such disruption can be seen in Figure 57. The top panel shows a series of ARI measurements switching sporadically between two very different profiles. In the middle and bottom panels, it can be seen that the UEG morphology is not altering substantially from-beat-to-beat.

Instead, the timing of T_{up} is switching between two local maxima in dV/dt . Hence, it appears that the dynamics of the ARI measurement are not an accurate reflection of the underlying dynamics in APD. In cases of manual measurement, a reviewer might choose only the later of these peaks as the nominal recovery time to impose consistency in the measurement. However,




Figure 57: *Top panel*– In this example of rate adaptation, ARI switches sporadically between two extremes, obscuring the typical pattern of adaptation. *Middle panel*– Two consecutive beats in the UEG signal. \dot{V}_{min} and T_{up} are marked as, respectively, dashed and solid vertical lines. The timing of T_{up} switches between two ends of the T-wave upstroke, suggesting the kind of instability described in Figure 50. *Bottom panel*– Inspection of the signal’s derivative confirms that T_{up} is switching between two ends of a broad peak in dV/dt .

this approach can be problematic if the chosen peak then disappears due to more gradual morphological changes. Hence it is not easily implemented in an automated system.

In section II.E it was shown that UEG T-wave morphology is the combination of an upward stroke produced by local repolarisation behaviour and a downward stroke from more distant behaviour. By that interpretation, the whole of the T-wave's broad positive peak in dV/dt represents the time during which repolarisation in the UEG's local component contributes more strongly than the remote component described in the model of Potse *et al* [Potse et al, 2009]. A more meaningful reflection of repolarisation dynamics would therefore take into account the whole of the peak in dV/dt , rather than picking one of several smaller maxima as the instantaneous index of local recovery. The same principle can be extended to activation time measurements, although in most cases the negative peak in dV/dt during activation is more sharply defined than the T-wave peak.

Modified definitions of the nominal activation and recovery times assigned to a UEG were developed to achieve such a measurement. The resulting indices, named A_S (stabilised activation time) and R_S (stabilised recovery time), are calculated as follows:

- I. The conventional indices, \dot{V}_{\min} and T_{up} , are calculated using the previously described algorithms.
- II. For each beat, a window containing the peak of interest is identified. For activation, this window is the refined search-window established for this beat during activation detection. For recovery, the start boundary is either the point t_{ST} or 60 ms before T_{up} , whichever is later. The end boundary is T_{end} or 100 ms after T_{up} , whichever is earlier. These definitions minimise the chances of non-repolarisation artefacts being included in the window.
- III. A slope threshold is defined, so that the measurement only incorporates signal segments where dV/dt exceeds this threshold. For activation, this threshold is set as 0.7 times the slope at \dot{V}_{\min} . For recovery, the threshold is set as 0.5 times the slope at T_{up} . These factors were chosen heuristically to eliminate incidences of spurious dynamics, such as those shown in Figure 57. For activation, a stricter threshold is used to ensure that only the intrinsic deflection (when one is visible) would be incorporated into the measurements. For recovery, a lower threshold was used to provide extra stability, since the destabilising ripples across the broad peak in dV/dt could be large compared to the peak itself.
- IV. R_S is then calculated from the following formula.

$$R_S = \frac{\sum_{k \in P} t(k) \left[\frac{dV}{dt}(k) - Z \right]}{\sum_{k \in P} \left[\frac{dV}{dt}(k) - Z \right]} \quad (41)$$

where P is the set of samples that exceed the slope threshold Z , $dV/dt(k)$ is the k^{th} point in the signal's derivative, and $t(k)$ is the time corresponding to that point. By this formula, R_S is mathematically equivalent to the centroid of the area enclosed by the slope threshold and the peak of the signal dV/dt . This area is shaded in the bottom panel of Figure 64 (page 133). A_S can be calculated using the same formula, with P and Z altered appropriately.

IV.E. Evaluation of developed algorithms

IV.E.1. Pacing artefact removal and activation wave detection

An example of the benefits of pacing artefact removal has been presented in Figure 53 (page 114), although more extensive evaluation is required. Because the principal purpose of the pacing spike removal algorithm is to improve the reliability of the activation-wave detection algorithm, the evaluation of the latter can also serve as the evaluation of the former.

The performance of the activation-wave detection algorithm was found to be more than sufficient for the purposes of the experiments described later in this thesis. A comprehensive assessment of the algorithm's performance could be achieved during this project, due to the time that would be required to manually annotate an extensive set of electrograms for comparison. One professionally annotated library of unipolar electrograms is known to be commercially available (the Ann Arbor Electrogram Library, www.electrogram.com) but the licensing expenses were beyond the budget of this project. Hence the electrograms recorded in the course of this project were the only ones available for validation of the algorithm. During constant-pacing sections, detection errors were extremely rare. To test the algorithm in a more challenging situation, activation waves were identified manually for a series of 100 beats in an unpaced recording from a subject experiencing persistent arrhythmia (atrial fibrillation), such that the magnitude and morphology of the waves varied substantially between ventricular activations. Additional morphological variations were caused by the fact that catheter positions were still being adjusted before the start of an experiment. The accuracy of detection in each of 20 electrodes was measured according to the definition given by (37) (page 110). The results are summarised in Table 5.

100 percent accuracy was achieved in 14 out of 20 electrodes. Figure 58 presents representative examples of cases in which the algorithm failed. Substantially reduced accuracy occurred at the electrodes closest to the apex of the left ventricle. The apparent cause was that

algorithm validation - accuracy percentages

	1	2	3	4	5	6	7	8	9	10
RV	100	100	100	100	100	100	100	100	100	100
LV	86	94	97	99	100	100	100	100	98	98

Table 5: Accuracy percentages, as defined in (37), for the activation wave detection algorithm applied to 10 electrodes in each ventricle. The recordings consisted of 100 beats, during persistent arrhythmic behaviour. Electrodes 1-10 are distributed across 35mm from the apex to the base of the ventricles. The electrode configuration is depicted in Figure 80, page 168.




Figure 58: Examples of cases in which the activation detection algorithm fails, taken from the electrograms used in the accuracy assessment that is summarised in Table 5. *Top trace* – A small activation wave is not detected due to the influence of the large subsequent wave in the broad window. *2nd trace* – Forces between the electrode and the myocardium at the catheter tip cause MAP like morphologies. The downward spikes associated with local activation are smaller than usual and are superimposed onto upward deflections, but can still be detected in some cases (for example, in the second beat here, but not the first). *3rd trace* – At other times, these forces and the resultant artefact appear only during ventricular contraction. The artefacts' effect on the broad window cause the algorithm to overlook small, fractionated activation waves, and occasionally they are mistaken for activation waves themselves. Comparison with a simultaneous recording from a nearby electrode (bottom trace) reveals which deflections are true activation waves and which are artefact.




Figure 59: A demonstration of the recovery time detection algorithm’s ability to exclude unreliable measurements. This electrogram is from the same recording as the 3rd trace in Figure 58. Dashed lines indicate activation times and solid lines indicate recovery times. *Beat A* – Recovery time ‘not found’ because the signal immediately surrounding T_{up} was too small compared to the rest of the T-wave window. *Beats B and C* – ‘found’. *Beat D* – ‘not found’ because local recovery did not produce a positive slope. *Beat E* – ‘not found’. Two substantial up-strokes appear, but both are too late to be considered local recovery. A small earlier up-stroke appears, but is too small compared to the surrounding signal to be considered reliable. *Beat F* – ‘found’. *Beat G* – ‘not found’ because the only substantial up-stroke occurs too late. *Beat H* – ‘found’.

these electrodes were occasionally forced against the myocardial wall. In some beats (e.g. 2nd trace) this caused MAP-like morphology (explained in section II.D.2.c), suggesting consistent forcing. In other beats, a large artefact emerged around the time at which ventricular contraction would be expected to occur (e.g. 3rd trace), suggesting that the contraction itself provided the force. Signal corruption of this kind was only observed in the distal electrodes, where the force between myocardium and electrode could act in the catheter’s axial direction; the catheter can be assumed to be stiffer in the axial direction than in the transverse direction, preventing the electrodes from moving freely with the myocardial wall.

Because the algorithm does not perfectly match a human’s ability to identify activation waves, it should not be used without supervision to characterise complex arrhythmias. Its advantage lies in providing activation times from long recordings because the user is not required to manually identify the time envelope within which activation occurs. Rapid adaptation to changes in overall electrogram morphology allow the algorithm to be applied to signals involving a combination of paced and unpaced beats, while ensuring that the algorithm is readily adaptable to a real-time implementation in future.

IV.E.2. Recovery-time detection

Figure 59 demonstrates the recovery-time detection algorithm’s ability to reject unreliable measurements on a beat-by-beat basis, referred to henceforth as auto-vetting. The task of meaningfully quantifying the accuracy of auto-vetting is a difficult one. The difference between reliable and unreliable recovery time measurements is more ambiguous than the difference between activation waves and other deflections. Cases can be found in which the decision made by the algorithm does not match that of a human reviewer, but ambiguous cases can also trigger

disagreement among human reviewers. It would therefore be inappropriate to quantify the auto-vetting accuracy rate using a beat-by-beat comparison against a manually annotated dataset (the approach used for the activation wave detection algorithm and summarized in Table 5). In tuning the algorithm's behaviour for ambiguous cases, it was deemed that unnecessary rejection of a few beats would be preferable to the inclusion of occasional erroneous measurements, in order to better serve the aim of exposing the dynamics of a variable. Erroneous measurements can substantially obscure the dynamics of a variable because their values are likely to differ widely from the surrounding values, whereas omitted measurements can be interpolated across with only a small effect on the apparent dynamics. This principle is illustrated in Figure 61.

To provide a more general impression of the effectiveness and importance of auto-vetting in preventing erroneous distortion of measurement dynamics, the following analysis was carried out. The comparison seen in Figure 61 was implemented for all twenty of the endocardial ventricular electrograms that were simultaneously recorded from this subject during the experiments described in section V.C. The outcomes of these comparisons were divided into four categories, as shown in Table 6. Note that the nature of the first category highlights a previously unmentioned benefit of auto-vetting: signals with consistently ambiguous morphologies can be automatically rejected by imposing a threshold on the percentage of beats rejected through auto-vetting. In multi-electrode recordings, it is typically not possible to achieve useful recordings from every electrode, so poor quality recordings must be excluded before post-processing. Automating this task further reduces the time-cost of human intervention during analysis.

Outcome	Number of electrodes (out of 20)
Signal section rejected (>10% of beats rejected)	5
No corrections necessary	7
Some manual corrections made. No automatic rejections made.	3 (see Table 7)
Automatic rejections made.	5 (see Table 7)

Table 6: Summary of outcomes when the analysis of Figure 61 was extended across all 20 simultaneous electrogram recordings from the same subject.




Figure 60: A sample of the electrogram used to produce the ARI measurements in Figure 61. Measured activation and recovery times are marked with dashed and solid vertical lines, respectively. Auto-vetting was switched off. In the first of the two beats shown, the boundary between the activation wave and T-wave upstroke is unclear. This ambiguous morphology ultimately leads to an erroneous identification of the recovery time.

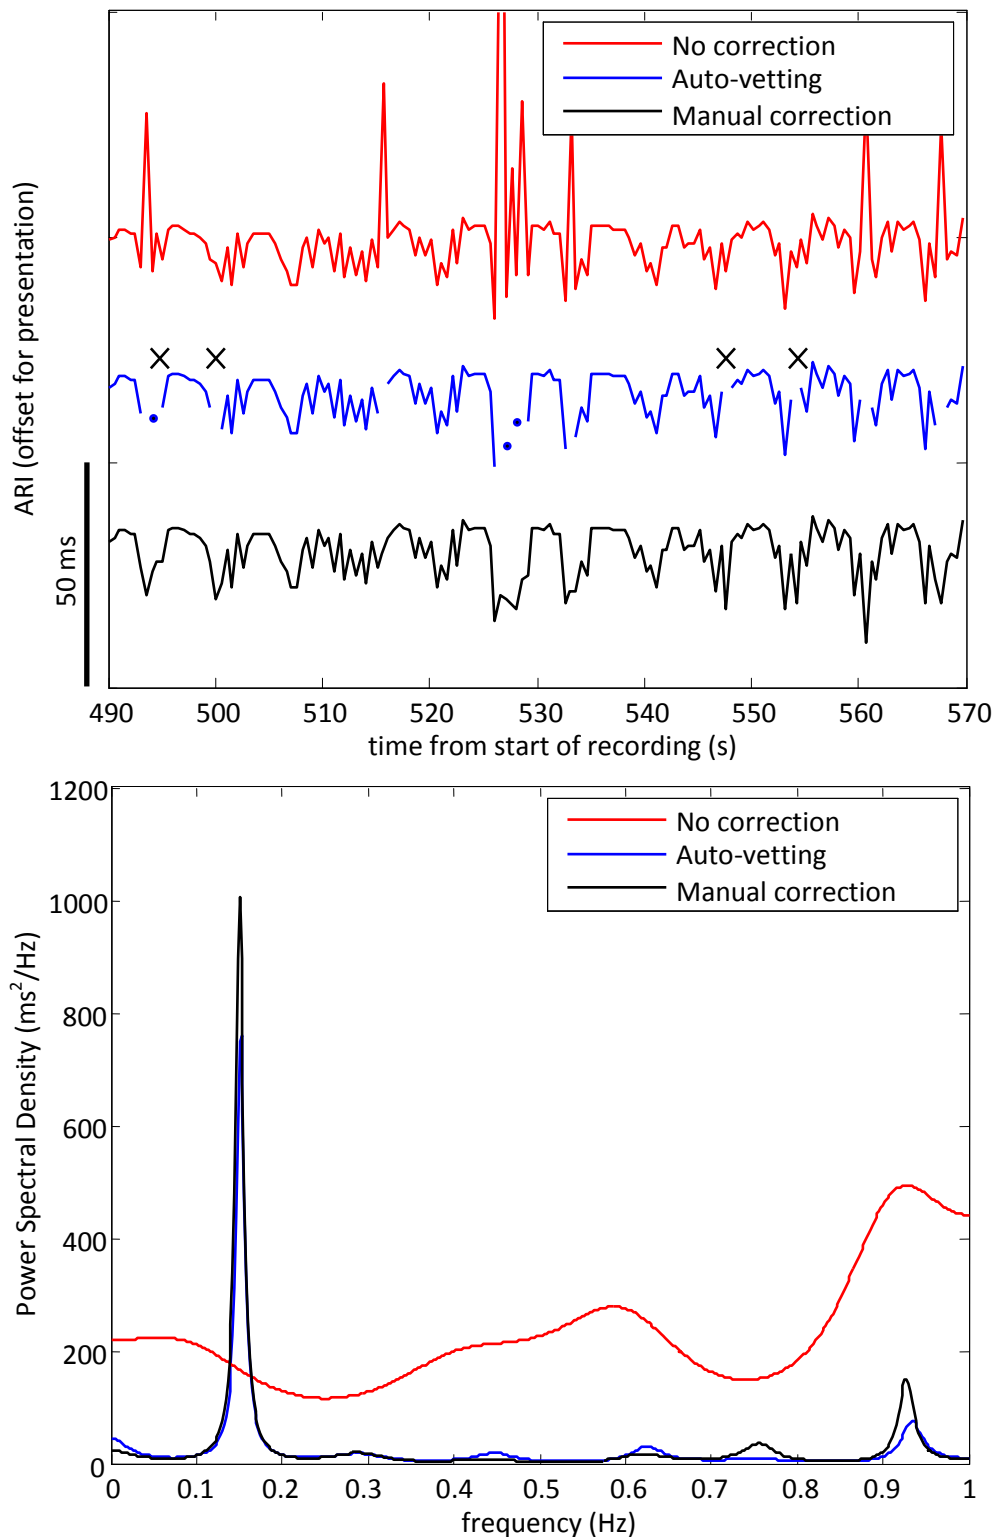


Figure 61: *Top panel*– Three versions of the same ARI series, taken from electrogram signal of which an example is shown in Figure 60. The first was achieved by switching off the algorithm’s auto-vetting features. The second was achieved by reinstating auto-vetting, so that ambiguous morphologies such as that seen in Figure 60 would not substantially distort the overall dynamics in ARI. Beats that were unnecessarily rejected are marked “X”. The third series was produced by manually correcting the first series. *Bottom Panel*– Power spectra calculated from these 80-second ARI series using the autoregressive method. The errors associated with ambiguous morphologies lead to substantial distortion of the series’ measured frequency content. In contrast, rejecting ambiguous beats and interpolating across the resultant gaps in the ARI series preserves the key features of the frequency spectrum, even though more beats than necessary were rejected.

For the eight cases in which corrections were necessary, the accuracy of the spectral calculations could be compared with and without auto-vetting. Two quantitative measures of accuracy were established for this purpose:

$$\text{Mean error: } ME = \frac{\sum_{f=f_1}^{f_N} |PSD_a(f) - PSD_m(f)|}{N} \quad (42)$$

$$\text{Percent total error: } \%TE = \frac{\sum_{f=f_1}^{f_N} PSD_a(f) - \sum_{f=f_1}^{f_N} PSD_m(f)}{\sum_{f=f_1}^{f_N} PSD_m(f)} \quad (43)$$

$PSD_a(f)$ and $PSD_m(f)$ are the power spectral densities of the automatic and manually corrected ARI series, respectively, at frequency f . f ranges across $N=500$ values from $f_1 = 0$ to $f_N = 1\text{Hz}$, the Nyquist Frequency of the series. Table 7 compares values of these performance indicators with and without the use of auto-vetting to find PSD_a .

The first three rows of Table 7 show that, although the auto-vetting algorithm failed to reject beats for which manual corrections were necessary in these signals, the resultant distortion of the spectral content was relatively small. The ‘percent total error’ values may seem large until one considers that in Figure 61 the key features are well-preserved by the use of auto-vetting, and the percent total error measured in this case was 24 %. This index is prone to large values when the denominator is small, i.e. when the total power of the manually-corrected spectrum is low. The low ‘mean error’ values confirm that the absolute distortion of spectral content resulting from uncorrected errors was low in all three of these cases.

	Electrode name	Mean error (ms ² /Hz)		Percent total error	
		No correction	Auto-vetting	No correction	Auto-vetting
No automatic rejections made.	RV4	4.0	4.0	53 %	53 %
	RV6	2.8	2.8	36 %	36 %
	RV10	3.0	3.0	31 %	31 %
Automatic rejections made	RV2	219.2	3.5	1459 %	24 %
	RV3	17.4	0.3	116 %	2 %
	RV5	4.3	0.9	15 %	3 %
	RV7	783.8	2.0	11419 %	29 %
	RV8	850.0	1.1	16023 %	21 %

Table 7: Quantitative comparison (see equations (42) and (43)) of the distortion of ARI dynamics (spectral content) with and without auto-vetting switched on. Electrode RV2 is the signal used in Figure 60 and Figure 61.




Figure 62: Adapted from [Franz et al, 1988] with permission (by payment). APD changes in response to step changes in the cycle length during artificial pacing of an *in vivo* human heart. APD (measured from MAP recordings) responds with a step change followed by a more gradual convergence. Black arrows indicate the times at which the step change in pacing rate occurs.

The remaining five rows of Table 7 show that, where the auto-vetting process did reject beats, it yielded a more accurate representation of the series' spectral content in every case. The mean error in the auto-vetted spectrum never exceeded that presented in Figure 61 (electrode RV2). In contrast, for two cases the uncorrected spectrum yielded mean errors and percent total errors that were substantially greater than for the case in Figure 61.

This study confirms that auto-vetting can be used to substantially reduce the extent to which occasional ambiguous T-wave morphologies distort calculations of repolarisation dynamics. The approach thus enables the time- and labour-savings of automated recovery time detection while minimising the extent to which the dataset needs to be reduced to ensure that subsequent dynamic analyses are meaningful.

To illustrate a useful application of automated detection of recovery times from UEGs, a simple example of repolarisation dynamics is considered. [Franz et al, 1988] manually analysed MAP recordings to show that a step change in the cardiac cycle length causes a step change in APD followed by roughly 2-minutes of gradual convergence towards a steady-state value of APD. This behaviour is shown in Figure 62. Figure 63 illustrates that this 'rate adaptation' behaviour can be exposed in measurements of ARIs (activation-recovery intervals) from unipolar electrograms (UEGs). Characterising this behaviour at a single site requires roughly 100 measurements or more to be taken. Hence manual analysis can be very laborious, limiting the size of the data-sets taken on. Automatic calculation of ARIs with automatic rejection of unreliable measurements allows this rate adaptation to be characterised across multiple sites simultaneously, with only minimal corrections from a human reviewer. The developed algorithm was recently employed by Dr. Alfonso Bueno-Orovio, Dr. Blanca Rodriguez and colleagues (University of Oxford, Department of Computer




Figure 63: The behaviour described by in Figure 62 is evoked at the onset of pacing and exposed using automated detection of activation-recovery intervals from unipolar electrograms. The top trace shows that the cycle length changes suddenly from its natural value of roughly 1230 ms to the paced cycle length (500 ms). The bottom trace shows that the activation-recovery interval responds with a sudden initial change, followed by a more gradual convergence towards a settled value.

Science) to expose regional differences in rate adaptation using 20 intracardiac UEGs from 18 patients. Such differences are expected to have important implications for the genesis of arrhythmia during sudden changes in heart rate. This study is the first to expose these heterogeneities in *in vivo* human ventricles, which illustrates the importance of enabling the analysis of extensive datasets from UEGs. The work is summarised in a paper accepted for presentation at the *The 7th International Workshop on Biosignal Interpretation*^[Bueno-Orovio et al, 2012] and a related paper soon to be submitted to *The Journal of Physiology*.

IV.E.3. Stabilisation of local event times

The stabilised measurements of activation and repolarisation time – A_S and R_S – will sometimes differ by a few milliseconds from the conventional indices, which have been verified against more accurate measurements of event timings. However, the dynamics of the stabilised measurements can be expected to more faithfully reflect the dynamics in the underlying action potential morphologies. The bottom panel of Figure 64 shows that, when ARIs are calculated using these new definitions of activation and recovery time for the same signal as in Figure 57 (page 123), the spurious dynamics are attenuated to reveal a clear rate-adaptation profile. Respiratory




Figure 64: Compare with Figure 57. *Top panel*– When ARI is calculated from the stabilised indices A_s and R_s , the switching behaviour is eradicated, leaving a clearer impression of the rate adaptation profile. *Middle panel*– Two consecutive beats in the UEG signal. The shaded box indicates the portion of the signal used to calculate R_s , which sits stably near the centre of the T-wave upstroke. *Bottom panel*– Inspection of the signal’s derivative confirms that the timing of R_s is stable relative to the broad peak in dV/dt . R_s is the centroid of the shaded region between dV/dt and a threshold of 0.5 times the maximum in dV/dt .

oscillations in ARI are also revealed. The nature of these respiratory oscillations is discussed extensively in section V.C.

Another demonstration of the importance of stabilised measurements of activation and repolarisation times is presented in Figure 65. ARIs calculated using the conventional indices show bursts of oscillations at a rate faster than the respiratory frequency. This appears to be an




Figure 65: When the conventional (unstabilised) indices of activation and recovery time (\dot{V}_{\min} and T_{up}) are used, the sensitivity of the ARI measurement to alternans is extremely inconsistent. Use of the stabilised indices on the same signal reveals that alternans is in fact present throughout this section.

intermittent manifestation of the phenomenon known as T-wave alternans. Alternans is typically defined as a consistently alternating pattern in electrophysiological behaviour, with a period of twice the cycle length. The presence of T-wave alternans has been shown to frequently precede ventricular arrhythmias, and is an effective indicator of increased risk of Sudden Cardiac Death [Rosenbaum et al, 1994], [Narayan & Smith, 1999], [Gehi et al, 2005]. Hence it is currently a subject of substantial research interest. Studies in canine [Konta et al, 1990], [Gelzer et al, 2008] and guinea-pig [Pastore et al, 1999] hearts have shown that discordant alternans (where spatial heterogeneities exist with respect to the phase at which alternans is occurring at any instant) may be particularly potent in triggering arrhythmia. Hence the ability to use UEGs to investigate heterogeneous alternans with *in situ* human hearts constitutes an important new tool for researchers.

Close inspection of the UEG used to produce Figure 65 indicated that the larger alternations in the dashed ARI trace arose as a ‘switching’ behaviour, similar to that shown in Figure 57. The regularity of the alternation suggests that the switching is driven by some genuine form of alternans, but it is likely that the magnitude of the switching substantially exaggerates the changes in APD. Furthermore, no dramatic change in UEG morphology could be found to coincide with the points at which the alternans suddenly emerges and disappears from the ARI measurement. This observation confirms that the sensitivity of conventional ARI measurements to subtle electrophysiological dynamics is extremely variable, even within a single recording. Hence these measurements are not appropriate for assessing the magnitude of alternans or the timing of its onset. The solid line in Figure 65 shows that when the stabilised indices \mathcal{A}_S and R_S are used, a smooth development of the onset of alternans can be observed. A team at London’s Heart Hospital (including Dr. Pier Lambiase, Dr. Xiao Jie, Dr. Malcolm Finlay, and Dr. Justine Bhar-Amato) is




Figure 66: An example of switching behaviour in the conventional activation time index, \dot{V}_{\min} , marked by solid arrows for each beat. The extent to which this switching corrupts the dynamics of the index can be seen by plotting the intervals between the pacing artefact and the index (dashed line, bottom panel). When the stabilised index A_S is used (solid line, bottom panel), it is revealed that the changes in activation wave morphology are driven by consistent alternans.

currently extending the use of these stabilised ARI measurements to real-time investigation of alternans, enabling novel research protocols of the kind described in Figure 51, page 111.

The preceding examples illustrate the benefits of stabilising the recovery time index. As stated previously, activation time measurements are typically more stable, but Figure 66 presents a case in which stabilisation of the activation time is also important. The ill-defined down-stroke of the activation wave gives rise to switching in the index \dot{V}_{\min} . This behaviour would corrupt the dynamics of any ARI measurement calculated using these activation times. Use of the stabilised

index, \mathcal{A}_s , eliminates the erratic switching and reveals a consistent alternans influence on the activation wave.

IV.F. Summary of achievements for chapter IV

It has been demonstrated in the preceding sections that the developed algorithms enable the detection of activation and recovery times from large datasets with minimal human intervention. Furthermore, new definitions for the nominal activation and recovery times were shown to stabilise the measurements, allowing more meaningful characterisation of the dynamic changes in UEG morphology. These analytic tools enable the use of the UEG as a means of observing heterogeneous autonomic input to the heart. This approach has the potential to yield abundant human *in situ* data because UEGs are already used extensively in routine, minimally invasive therapeutic procedures such as ablation therapy. Such data will be crucial in advancing the present understanding of the role played by the ANS in the genesis of fatal arrhythmias. The individual accomplishments that combine to enable this approach can be summarised as follows.

Automatic removal of pacing artefacts with minimal distortion:

The use of pacing to control heart rate is an important tool in the study of nervous modulation of cardiac behaviour because many of the electrophysiological properties of interest are dependent on heart rate variations as well as direct ANS input. Artificial electrical stimulation of the myocardium unavoidably imposes a distinct artefact on the UEG, and can corrupt subsequent signal processing if not properly accounted for.

The algorithm described in IV.D.1 identifies pacing artefacts based on their substantial high-frequency content. The method used to remove the artefact avoids conventional filtering techniques, instead replacing the artefact with only a short section of heavily low-pass filtered signal. Distortion of the surrounding signal is thus minimised to prevent the corruption of subsequent signal analyses.

Activation wave recognition:

Any beat-by-beat analysis of a cardiac electrical signal requires accurate identification of separate beats. For UEG analysis, this is most easily achieved by recognising activation waves, the most distinct regular feature of the signal. An algorithm was developed to automate this process. The decision criteria employed in the algorithm use only general assumptions about the morphology of the activation waves, thereby accounting for the wide range of individual morphologies that can occur (i.e. positive, negative, and polyphasic activation waves). The algorithm was tested using a challenging set of test recordings, in which persistent arrhythmia gave rise to a variety of abnormal activation patterns and, hence, a variety of activation wave morphologies. The algorithm was found

to identify activation waves with at least 98 percent accuracy, except in electrodes that showed evidence of inappropriate positioning.

Recovery time recognition:

Another algorithm was developed to implement automated recognition of recovery times in UEG signals. Again, this algorithm used only generalised assumptions about signal morphology to account for the full spectrum of T-waves configurations that can occur. Included are criteria to identify and exclude cases in which the estimated recovery time is unreliable. Such occurrences are inevitable in UEG recordings, so this ‘auto-vetting’ ability is essential for immediate implementation of subsequent calculations. It was demonstrated that, the frequency content of ARI time-series calculated by this algorithm closely match those taken manually, which is not the case when ‘auto-vetting’ is not implemented. Hence this feature is crucial in extending the use of UEG signals to real-time measurement of electrophysiological dynamics. Additionally, the algorithms have been formulated to be suitable for real-time implementations, which would enable novel experimental procedures based on feedback of dynamic properties, as described in Figure 51 (page 111).

Stabilised indices of activation time and recovery time:

It was demonstrated that the conventional indices used to identify activation and recovery times in UEGs (\dot{V}_{\min} and T_{up}) are not well suited to the characterisation of electrophysiological dynamics; the small errors to which the indices are prone are of similar magnitude to the dynamic fluctuations of interest. For that reason, new stabilised indices of activation and recovery time were proposed (A_S and R_S). These new indices do not improve the accuracy of the measurement but, by construction, they ensure that the errors are consistent from beat to beat, assuming no drastic changes in signal morphology occur. This characteristic prevents distortion of any dynamic calculations that are based on the indices.

IV.G. Scope for improvement

Although the developed tools have been shown to be effective for their intended purpose, there are several ways in which they might be improved in future efforts.

Multi-electrode decision making:

Cases of ambiguous UEG morphology can sometimes be resolved by comparing simultaneous deflections across multiple electrodes, as demonstrated in the lowest two panels of Figure 58. Furthermore, the use of bipolar electrograms (constructed as the difference between two UEGs) can often be used distinguish local activity from more remote behaviour, at least for activation waves. These approaches are often used during manual identification of activation and recovery times, and could conceivably be incorporated into an automated detection algorithm.

Fractionated electrograms:

As described in section II.E.3, certain situations produce electrograms for which it could be misleading to assign only one activation and recovery time to each beat, because the timing of these events is temporally distinct for two or more clusters of cells within the field-of-view of the exploring electrode. The developed algorithms do not contain special provisions for dealing with fractionated electrograms. It would be relatively easy to develop a criterion to automatically exclude such morphologies – for example, by identifying T-waves containing two or more substantial upstrokes separated by a region of reduced dV/dt . This approach was not adopted in this project because, as stated in the justification for the stabilised measurements, the dynamics of changes in these event timings were of greater interest than the timings themselves; the stabilised measurements typically incorporate all major deflections in fractionated electrograms, so that dynamics in the timing of any of those deflections will be captured to some extent, even though the nominal event times will not accurately represent any of the individual cell clusters.

A caveat on this approach is that dynamic changes in the lead field may change the relative influence of the different clusters on the UEG morphology. False dynamics would then be introduced to the stabilised measurements as the relative influence shifts between clusters. In fact, this caveat is not specific to fractionated electrograms, nor is it specific to the newly developed, stabilised measurements.

Variable specificity in measurements:

The preceding discussion of fractionated electrograms highlights the fact that any index of activation or recovery time is a simplification of a complex spatiotemporal distribution of activity. The steepest slopes in UEG deflections are used as the conventional indices of local behaviour because the underlying current sources are changing quickly enough for their influence on the local potential field to be distinct at these times. However, these changes are not instantaneous, and are not perfectly synchronised among all cells in a region of myocardium. As a result, their manifestation in the local potential field can sometimes be spread across a relatively long signal segment. The stabilised measurements are designed to provide a more appropriate simplification of the spatiotemporal distribution of activity, in the form of a weighted-average of the event time, but the activity distribution is still reduced to an instantaneous event. The confidence that can be attributed to this measurement as a meaningful index of underlying behaviour will be greater in some cases than in others. Hence one might consider using the temporal spread of UEG features (e.g. the width of the shaded region in Figure 64, page 133) as a complementary measurement, indicating the specificity associated with the stabilised event times for each beat. This additional index could then be used to decide whether to include particular signals in further dynamic calculations, or to estimate the confidence attributable to the results of those calculations.

Further validation needed:

It should be noted that the described algorithms were developed by a process of gradual refinement to maximise performance when evaluated on the available patient data-sets. These recordings contained some natural heart beats, but predominantly consisted of artificially paced beats, which can produce abnormal activation patterns. The morphology of UEGs is dependent on electrode positioning and on the spatiotemporal spread of activity. Hence it is possible that the performance of the algorithms will be altered when using different electrode configurations or different pacing scenarios. As mentioned previously, the algorithms are currently in use by a team at The Heart Hospital in London; this group has informally reported successful implementation of the algorithms in ventricular epicardial electrograms and in atrial intracardiac electrograms. However, a comprehensive and systematic assessment of the algorithms' performance across a wide range of scenarios has not yet been performed. Furthermore, the stabilised measurements have yet to be validated against a more accurate measure of local electrophysiological behaviour, such as MAP recordings, to confirm whether they are more accurate than the conventional indices for measuring action potential durations. Nevertheless, the superiority of the stabilised indices has already been demonstrated for cases where an accurate representation of dynamic changes is more important than the accuracy of individual measurements.

V. Experimental Applications

Sections III and IV covered methodological considerations and novel methods developed for the purpose of inferring the nature of autonomic control of the heart based on ECG and UEG recordings. This section will describe the implementation of these methods in several different experiments conducted over the course of this project. These experiments were all designed to explore questions of psychological, psychiatric, or physiological interest. For this thesis, however, the primary concern is the efficacy of the methods and the extent to which they enable novel physiological insights to be gained. The experiments can be divided into three categories, each described in a separate section:

- V.A. Can the ECG reveal evidence of emotionally generated asymmetry in autonomic outflow to the heart? – A collaboration with psychologists at the University of Southampton.
- V.B. Can HRV analysis be used to expose cardiac effects of transient emotional responses? – A collaboration with psychiatrists at the University of Oxford.
- V.C. Can intracardiac UEG recordings expose autonomic modulation of ventricular electrophysiology, driven by respiratory behaviour? – A collaboration with cardiologists at St. Thomas' Hospital, London.

A key collaborator in all of the experiments mentioned above was Peter Taggart, Professor of Cardiac Electrophysiology at The Heart Hospital, London.

V.A. Can the ECG reveal evidence of emotionally generated asymmetry in autonomic outflow to the heart?

V.A.1. Objectives

A recent study involving 10 cardiology patients^[Critchley et al, 2005] found that those subjects who presented predominantly right-sided midbrain activity during mental and physical stress challenges also presented changes in ECG features that were significantly different from those in other subjects. This observation gives support to the Brain-Heart Laterality (BHL) hypothesis. However, the number of subjects studied was small, as necessitated by the expense and time-requirements of brain-scanning studies used to identify subjects with right-sided midbrain activity. This section describes an experiment to test a more practical means of investigating the BHL hypothesis in large cohorts.

As described on page 69, it has been suggested that it may be possible to prove the Brain-Heart Laterality hypothesis by providing an asymmetric stimulus to the brain. This stimulus may yield a measurable change in ECG parameters that can be explained by the asymmetry of ANS

innervation of the heart (Figure 29, page 66). If so, this approach would provide a useful technique for researching the clinical significance of the BHL hypothesis, including the true implications for cardiac electrical stability, in large cohorts. Separate modulation of heart rate and ventricular repolarisation properties, dependent on which side the stimulus is presented to, would be one example of such a measurable change in ECG parameters. An objective of particular relevance to this thesis was to explore whether the variable QT_C , calculated by the methods described in section III.B.2, could be used to heart-rate-independent autonomic control of ventricular repolarisation.

V.A.2. Methods

V.A.2.a. Stimulus

At the University of Southampton's Department of Psychology, Dr. Richard Gramzow conducted an experiment to investigate lateralised emotional processing induced by presenting images from the International Affective Picture System (IAPS). IAPS is an image library that has been validated as a standardised tool for researchers to induce certain emotional responses^[Lang et al, 1997]. Images were presented in pairs with a neutral image on one side and a distressing ("negative") image on the other, so that the optic nerves would transmit the images to different sides of the visual cortex (see Figure 67). Each pair of images was displayed for only 170 milliseconds, with intervals of 2.3 seconds, so that the subjects did not have time to redirect their gaze. If the asymmetry of the stimulus was retained throughout cerebral processing, and substantially preserved in the subsequent ANS activity, one would expect to observe differences in the ECG according to which side of the peripheral vision the distressing images were presented on.

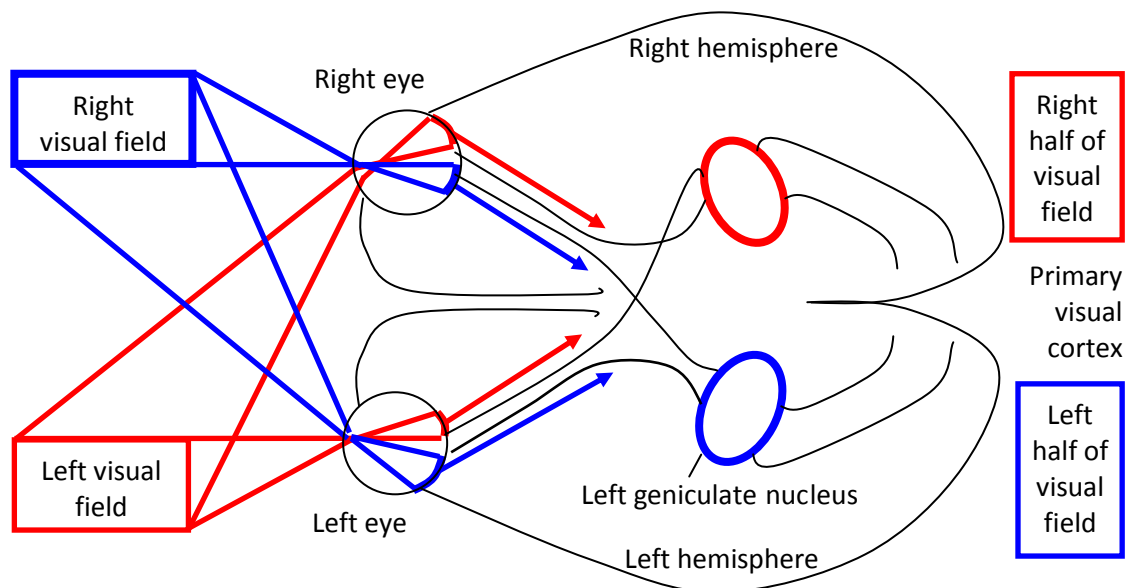


Figure 67: The ‘wiring’ of the optic nerve causes visual information from different sides of the visual field to be transmitted to different sides of the brain. Early processing in the lateral geniculate nuclei is also carried out separately for the two visual fields. Anatomy from [Deutsch & Deutsch, 1993].




Figure 68: Timeline illustrating the protocol used for half the subjects in the Brain-Heart Laterality experiments. For the other half, the left-right order was reversed. The start time for each image series was controlled by the patient and recording automatically by the computer program used to present the images. Half of the subjects were tested with the negative (distressing) image on the left side for the first series, and half were tested with the negative image on the right side for the first series.

81 pairs of images were presented with random order in two epochs, with a break of approximately one minute between them, as shown in V.A.3.a. The negative image was always presented on one side during the first epoch and on the other side during the second epoch. The ordering of the left- and right-sided epochs was determined pseudo-randomly for each of the 54 subjects tested. Half the subjects were tested with the negative image on the left for the first epoch, while the other half were tested with the negative image on the right first. This balance ensures that in statistical analyses the effects of laterality and of the order in which the two sides were tested could be separated.

V.A.2.b. ECG Measurements

Heart rate and QT interval measurements were probed for statistically significant effects of the stimuli. The length of the recordings (roughly 1000 beats/subject \times 54 subjects) necessitated the use of the automated QRS and QT measurement systems described in Appendix IV and section III, respectively. QT_C was also calculated, using the entire recorded ECG to establish the QT-RR relationship for each subject. This proved to be a poor choice of training data, as will be described in the discussion section. However, the time and expense associated with more rigorous characterisation of the QT-RR relationship would have defied the purpose of developing a convenient means of testing the BHL hypothesis in large cohorts.

It was hypothesized that QT_C intervals might be more sensitive than standard QT intervals as an indicator of asymmetric ANS activity. This hypothesis is based on the observation that heart rate is predominantly controlled by the right branch of the ANS and ventricular repolarisation is predominantly controlled by the left branch (Figure 29, page 66). QT_C intervals could therefore be interpreted as reflecting left-sided ANS activity while controlling for the effects of right-sided activity.

To ensure that QT_C calculations are based on an accurate model of the QT-RR relationship, a QT-RR training set should be acquired for each subject. These training data should be independent from the experiment recordings, and should cover a large range of values. Thus in this case, the training set should be established over an extended period of changes in body-position between supine and standing^[Malik et al, 2008]; the resultant changes in blood pressure induce a change in heart rate, with little influence from emotional state or, if a tilt-table is used, exertion. Because the potential use of QT_C calculations was not considered when the experiments were conducted, no

separate training dataset was acquired. Instead, the full recordings taken during the experiments and the preceding few minutes rest were used for the training set. The resting period before the start of the first image series was considered as an alternative training-set, but this approach was found to yield poor correlation between QT and RR, due to the low variability of the RR intervals in this period.

V.A.2.c. Collective analysis and statistical comparisons

Once the sequences of RR intervals, QT intervals, and QT_C intervals had been calculated for all subjects, it was necessary to distinguish between changes that occurred in response to the stimuli and those that were unrelated. It was assumed that each subject's overall response was a superposition of these two components, and that the response-to-stimuli component was similar in each subject. This component was therefore identified by finding the mean time-series across subjects for the measurement in question, as follows:

- I. The time-series (RR intervals, QT intervals, or QT_C intervals) for each of the 54 subjects was normalised to have a mean value of 1 second, then interpolated with cubic splines and resampled at 1Hz.
- II. The time axis for each resampled series was aligned so that equivalent events would occur at the same time for each subject.
- III. The mean of these new series was then found as the 'average response' for a particular type of epoch.

The averaged-response plots from the brain-heart laterality experiments are shown in Figure 69, Figure 70, Figure 71, and Figure 72, comparing the effects of left- and right-side presentation of images, and comparing the effects of the first and second epochs, regardless of which side images were presented on. To test the statistical significance of any observed differences in these comparisons, Welch's t-test was applied to each point in time. For example, to compare the effects of the first and second epochs on QT intervals 10 seconds after the start of the epoch, the two datasets used in the t-test were (a) the QT intervals at 10 seconds from the start of the first epoch (one value for each subject), and (b) the QT intervals at 10 seconds from the start of the second epoch. Examples of the p-values produced by these tests can be seen plotted against time in the lower panels of Figure 69 and Figure 70.

V.A.3. Results

V.A.3.a. QT intervals

Figure 69 and Figure 70 show the averaged response to the IAPS image epochs of the brain-laterality experiments in terms of QT intervals (found by Berger's method), with an identifiable response to the protocol. QT intervals increase by approximately 1.5 percent when the images are first displayed, and then gradually return to their previous levels over the next two minutes, presumably due to habituation. The most substantial common response is produced at the




Figure 69: *Upper panel* – Average response of QT intervals to the IAPS images. The plot was formed by averaging all the epochs in which the negative image was displayed on the left in one series (blue), and all the epochs in which the negative image was displayed on the right in the other (red). *Lower panel* – P-values calculated using Welch’s t-test, indicating the statistical significance of any difference between the ‘left’ and ‘right’ groups of measurements at each point in time. The dashed line marks the conventional significance threshold, $p = 0.05$. At no time does the measurement cross fall below this threshold to achieve statistical significance.

end of the epoch, where RR intervals shorten by approximately 2 percent. Also, QT intervals shorten by approximately 1 percent before the start of the images. These observations suggest that the observed fluctuations are more strongly controlled by the change in activity between resting and focussing than by the images themselves.

The lower panel of Figure 69 reveals that at no point did the differences between left- and right-sided presentation of images become statistically significant. However, the lower panel of Figure 70 shows that a significant difference ($p < 0.05$) exists between the first and second epochs during much of the first two minutes.

V.A.3.b. RR intervals

Figure 71 presents the ‘average response’ in terms of RR intervals, which follows a similar time-course to that of QT intervals. RR intervals increased by approximately 10 percent when the




Figure 70: Same as Figure 69, with the exception that these plots compare the 1st and 2nd IAPS epochs for each subject, regardless which side images were presented on. P-values notably decrease for the first two minutes of images, achieving statistical significance ($p < 0.05$) for much of this period.

images are first displayed, then gradually return to their previous levels over the next two minutes seconds. At the end of the epoch, RR intervals shorten by approximately 15 percent. A 10 percent shortening of RR intervals occurs in anticipation of the start of the images. t-tests were applied in the same manner as for QT intervals, but revealed no statistically significant differences associated with these plots.

V.A.3.c. QT_C intervals

Figure 72 shows the averaged time-course of the QT_C intervals. The heart-rate correction has eliminated the most prominent deflections seen in Figure 69 and Figure 70. A period of increased variability in QT_C appears approximately 60s from the start of the IAPS epoch, corresponding approximately with the time at which the RR interval response has returned to its baseline. Again, t-tests revealed no significant differences associated with these plots.




Figure 71: Average response of RR intervals to the IAPS images. The plot in the top panel was formed by averaging all the epochs in which the negative image was displayed on the left in one series (blue), and all the epochs in which the negative image was displayed on the right in the other (red). The plot in the lower panel was formed by averaging the first (blue) and second (red) IAPS epochs, regardless of which sides the images were presented.

V.A.4. Discussion

V.A.4.a. Summary of observations

Neither RR intervals nor QT intervals presented any clear differences attributable to the presentation-side of the negative images. The use of an ‘average response’ to separate the effects of the stimuli from unrelated fluctuations was successful in that RR and QT intervals presented clear deflections related to experiment events. The sudden decrease in heart-rate observed at the start of each IAPS epoch agrees with the results of previous studies^{[Winton et al, 1984], [Palomba et al, 1997], [Palomba et al, 2000]}. The observation of a habituation effect in QT intervals, with marginal statistical significance (Figure 70), indicates that the psychological effects of the stimuli are indeed manifested in the ECG parameters, even though no left-right differences were observed. Based on the averaged QT_C response depicted in Figure 72, there appears to be very little heart-rate-independent modulation of repolarisation properties during these experiments.

V.A.4.b. Methodological Critique

Implementation of the heart-rate correction algorithm to calculate QT_C demonstrated the importance of the choice of training dataset to fully characterise the QT-RR relationship. The range




Figure 72: Average heart-rate corrected repolarisation response to IAPS images, measured as QT_C .

of the training sets used in these experiments was limited; heart rate does not vary substantially during the resting period, so the most extreme values in each subject's training set are determined predominantly by the changes that occur in response to the experiment. Any heart-rate-independent modulation of repolarisation, simultaneous with the direct influence of heart rate, would therefore be incorporated into the QT -RR relationship and eliminated from QT_C . Hence, the relatively flat response seen in Figure 72 does not decisively rule out the presence of heart-rate-independent changes in repolarisation properties. Nevertheless, the inconvenience that would be associated with acquiring a satisfactory, independent QT -RR training set for each subject represents a substantial limitation of QT_C measurements in general.

The lack of a separate training set to establish the QT -RR relationship is not the only possible explanation for the fact that no left-right distinction was found in this experiment. It may be that, although asymmetry is preserved in early visual processing, the subsequent emotional response involves enough cross-communication between hemispheres of the brain to eliminate that asymmetry. Another possibility is that the duration for which each pair of pictures was displayed (170 ms) was too short to induce a substantial emotional response, with clear differentiation

between the negative and neutral images. Furthermore, it is unlikely that presenting two images simultaneously will produce emotional effects similar to when the images are presented singly, considering that the subject's gaze is not allowed to settle on either picture. As described in section V.B, IAPS images were used again in subsequent experiments, with each image displayed alone for five seconds to allow its psychological effect to take hold.

V.A.5. Conclusions

The absence of a statistically significant effect, combined with the uncertainty regarding the effectiveness of the stimulus, made it impossible to conclude whether lateralised emotional stimuli could reliably give rise to asymmetric ANS activity and associated changes in ECG parameters. In any future exploration of the proposed method, the specific stimulus to be used should first be tested for its ability to reliably evoke a lateralised cerebral response using fMRI or PET scans of brain activity in a small set of subjects.

The variable QT_C is unlikely to be useful as part of a practical approach to investigating the BHL hypothesis using the general population; the need to acquire an extensive training dataset of RR and QT intervals for each subject would largely cancel out the benefits of not having to use brain-scanning techniques to identify the minority of subjects who naturally lateralise emotional processing.

V.B. Can HRV analysis with spectral averaging be used to expose cardiac effects of transient emotional responses? – ‘The Oxford Experiments’

V.B.1. Objectives

This section covers a series of experiments conducted in collaboration with researchers from the Department of Psychiatry at the University of Oxford. These experiments all followed a similar procedure based on a previous experiment^[Phan et al, 2005] that was successful in inducing measurable differences in brain activity according to whether the subject voluntarily suppressed or ‘maintained’ their natural response to an emotional stimulus. That stimulus came in the form of images from the IAPS collection selected for being highly “aversive”. Examples of the content of these images include burn victims, dead animals, and funeral scenes. For the ‘maintain’ condition, subjects were instructed to attend to and experience naturally (without trying to alter) the emotional state elicited by the pictures. For the ‘suppress’ condition, they were asked to reduce the negative affect from the images using the cognitive strategy of reappraisal. Each subject was instructed in the use of this technique, which involves reinterpreting the images in the context of an imagined scenario in which the negative implications are less severe.

Two of the experiments at Oxford (covered in sections V.B.3 and V.B.4) sought to determine whether the procedure used by [Phan et al, 2005] evokes different responses from ‘normal’ subjects and from those with psychiatric disorders (in particular, panic disorders and neuroticism). A third experiment (covered in section V.B.5) examined whether treatment with an antidepressant drug altered the behaviour of neurotic subjects during this experiment. ECG recordings were gathered in each experiment to examine whether the procedure also induced measurable changes in ANS activity, as indicated by heart rate variability (HRV). Substantial habituation to the stimuli was expected to occur in less than a minute, so the stimulus periods were limited to 45 seconds or less, shorter than the conventionally recommended minimum for HRV analyses. Hence the question of principal interest for the purposes of this thesis was whether the spectral averaging method proposed in III.B.1 could provide a useful indication of the short-lived autonomic response to the emotional stimuli.

Although measurements of cardiac autonomic effects are the primary concern, the results presented in the following sections include measurements of other effects (brain activity and self-reported emotion) because they help to characterise the overall physiological response to the stimuli. HRV results can be more richly interpreted and assessed when the broader physiological context is taken into account.

V.B.2. Methods

The general protocol followed in all of these experiments is illustrated in Figure 73. Variations specific to each experiment are presented separately in the sections that follow, as are the results. Section V.B.6 then covers the discussion of all the results collectively.

V.B.2.a. Experiment protocol

The experimental procedure was:

- 30 seconds of fixation on a crosshair on-screen to establish a baseline physiological state.
- 4 seconds of on-screen instructions, telling the subject whether they are to suppress or maintain their emotional reaction to the subsequent picture block.
- 30 seconds of pictures (5 pictures displayed for 5 seconds each, with 1 second between).
- 4 seconds in which the subject is to indicate their self-assessment of their emotional response on a scale from 1 to 4 (neutral to most negative).
- Repeat the above procedure until it has been carried out eight times, alternating between 'suppress' and 'maintain' for the picture block. Successive patients alternated as to whether the first instruction was to maintain or suppress, in order to avoid inducing a bias related to the order in which the tasks were executed.

V.B.2.b. ECG analysis

The automated QRS-detection procedure described in Appendix IV was implemented to examine the experiments' effects on heart rate and heart rate variability (HRV). These measurements were analysed separately for each of the crosshair-fixation and image-display blocks.

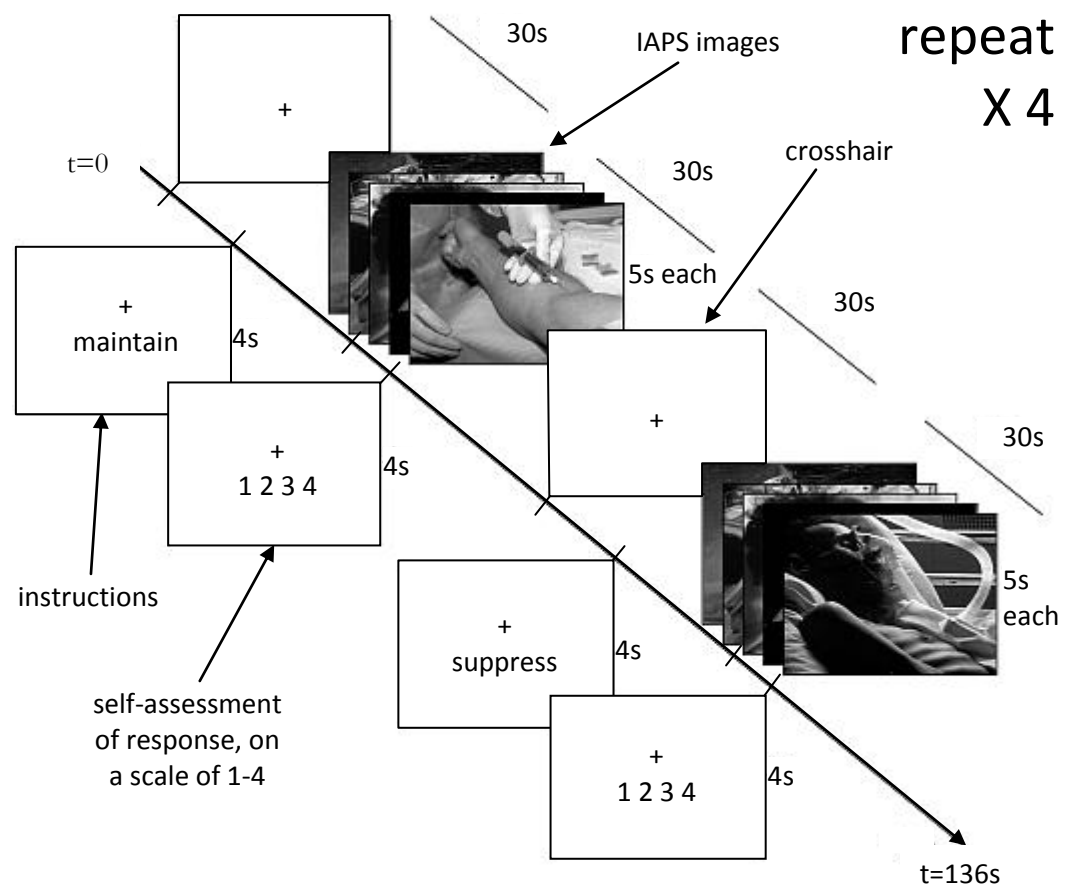


Figure 73: A graphical presentation of the IAPS protocol used in the emotion regulation experiments, with time running from top-left to bottom-right.

After differentiating as to whether each of these blocks corresponded to a ‘suppress’ or ‘maintain’ instruction, the blocks fell into four categories: ‘suppress’, ‘maintain’, ‘fixation before suppress’, and ‘fixation before maintain’. Each recording contained 4 instances from each category.

Because the length of these blocks is shorter than the recommended 2-minute minimum for HRV analysis^[TFESCNASPE, 1996], the ‘spectral averaging’ technique was used, as justified in section III.B.1; the mean of the spectra from the 4 instances within a particular category was calculated to provide HRV measurements. The spectra were calculated from auto-regressive (AR) models, as per section II.G.2. The AR models themselves were fitted after resampling the signal in each block to a length of $L=128$ points using cubic-spline interpolation, to reduce quantisation errors.

Statistical analyses of these measurements were performed using ANOVAs (Analyses of Variance) and Welch’s t -test^[Bain & Engelhardt, 2000] to examine whether any changes in a measurement – for example, the difference between HF power during ‘suppress’ and ‘maintain’ blocks – were statistically significant across the cohort or within a particular cohort.

V.B.3. First experiment – comparing panic sufferers with a control group

V.B.3.a. Methods

19 panic sufferers and 18 subjects from the general population were recruited for this experiment, conducted by Dr. Andrea Reinecke at the University of Oxford. The protocol in Figure 73 was used while subjects lay supine in an MRI scanner, so that fMRI could be used to monitor the experiment’s effects in terms of localised brain activity. The distortion of the ECG signal by the MRI scanner necessitated the use of the segment-mean signal processing technique, described in section III.A. One of the panic sufferers was excluded from the experiment when the MRI scan revealed a brain tumour. Also, 6 subjects were excluded from ECG analysis because the electrodes did not remain attached throughout the procedure, and another 5 were excluded because QRS complexes could not be reliably identified. The remaining cohort subjected to full analysis consisted of 13 panic patients and 12 controls.

V.B.3.b. Results

Psychiatric questionnaires and subjective affect ratings:

A series of questionnaire-based psychiatric tests confirmed that the two groups were psychologically distinct from one another. As shown in Figure 74, the subjective affect ratings given by the subjects after each image block show a significant difference according to whether they were asked to ‘suppress’ or ‘maintain’ their response, in that more negative emotions were evoked during the maintain blocks. No clear difference was found between the panic and control groups. The subjectivity of these ratings must be acknowledged as a caveat on their interpretation.




Figure 74: Average subjective affect ratings after each of the image blocks, with error bars representing the standard errors. A significant difference exists between the ‘maintain’ and ‘suppress’ conditions, but not between the two groups. Subjective affect ratings collected and compiled by Dr. Andrea Reinecke.

Brain activity:

A more objective measure of the brain’s response comes in the form of the fMRI data. These were processed by Dr. Andrea Reinecke, but are summarised here to allow comparison with HRV results. These showed a significant difference between the ‘suppress’ and ‘maintain’ blocks in line with what one might expect, given the nature of the task. Suppressing, compared with maintaining, was found to be associated with increased activity in brain areas associated with strategic processes, self-awareness, and conscious modulation of attention and emotions (Brodmann area 10, $p < 0.05$; superior frontal gyrus, $p < 0.01$). Suppression also presented reduced activity in areas associated with subconscious emotional processing (thalamus/hippocampus, $p = 0.001$; insular cortex, $p < 0.001$; precuneus, $p < 0.001$).

A significant distinction between the two groups in terms of the difference between the ‘suppress’ and ‘maintain’ responses was found in the right amygdala ($p < 0.05$). This region is associated with emotional processing and activation of the sympathetic nervous system. In the panic group, the right amygdala was much more active while maintaining than while suppressing, whereas in the control group the activity in this region was similar for both tasks.

HRV analysis:

For the purposes of this thesis, the results of greatest interest from this experiment are the incidences in which the HRV variables were successful in exposing a statistically significant difference between two experiment conditions. To explore the specific nature of HRV parameters calculated with the use of spectral averaging, section V.C.4 compares the different variables in terms of their ability to expose experiment effects. The significant differences exposed in this particular experiment can be summarised as follows:

Statistical significance of differences in HRV parameters

Comparison		<i>VLF</i> (0-0.04 Hz)	<i>LF</i> (0.04- 0.15 Hz)	<i>HF</i> (0.15- 0.4 Hz)	<i>LF_n</i>	<i>HF_n</i>	<i>LF/HF</i>
Fixation vs. Images, full cohort:		0.45	0.001 -	0.075 -	<0.001 -	0.001 +	0.025 -
Maintain vs. Suppress	Panic group:	0.99	0.65	0.31	0.039 -	0.044 +	0.091 -
	Control group:	0.52	0.31	0.63	0.099 +	0.11 -	0.22
Panic vs. Control	(Maintain) minus (Suppress):	0.62	0.31	0.28	0.008 -	0.011 +	0.036 -
	(Fix. before Maint.) Minus (Fix. before Supp.):	0.08 -	0.24	0.19	0.12 -	0.09 +	0.32

Table 8: A summary of p-values found using Welch's t-test to describe the distinction between various experiment conditions in terms of six HRV based measurements. Statistically significant differences ($p \leq 0.05$) are shown in bold. These and any near-significant differences ($p \leq 0.12$) are accompanied by a '+' or '-' indicating the sense of the difference. For example, a '+' indicates that the mean measurement was greater for the condition named after "vs."

- *LF*, *HF_n*, *LF_n*, and *LF/HF* showed that the subjects' HRV response to the image blocks was distinct from their response to the crosshair-fixation blocks.
- *HF_n* and *LF_n* exposed a difference between the 'maintain' and 'suppress' conditions in panic sufferers.
- *HF_n*, *LF_n*, and *LF/HF* exposed a distinction between the panic and control groups in terms of how their response to the 'maintain' condition differed from their response to the 'suppress' condition.

Table 8 summarises the results of the HRV analysis in terms of the p-values associated with any differences between experiment conditions. Statistically significant differences ($p \leq 0.05$) are shown in bold. These and any near-significant differences ($p \leq 0.12$) are accompanied by a '+' or '-' indicating the sense of the difference. For example, it can be seen from the first cell of the *LF* column that, across all subjects, the average LF power was significantly lower while images were displayed than during crosshair-fixation. The normalised indices, *LF_n* and *HF_n*, were also found to discriminate between the image and fixation blocks with strong statistical significance, as was the ratio *LF/HF*, albeit to a lesser extent. HF power discriminated with only marginal significance. Considering all of the results in the first row together, it appears that spectral power tends to decrease across the frequency range of interest (LF and HF) when the images are presented, but normalisation exposes a relative increase in the HF band (*HF_n*). This shift in relative magnitude between the LF and HF peaks is also manifested in the decrease in the *LF/HF* ratio during the image blocks.

p-values for the *VLF* power are also presented in Table 8. Although this frequency range is not indicative of autonomic activity, it is an important consideration in this case since the spectra




Figure 75: Mean values of LF/HF, with error bars representing the standard error of the samples. The two groups show opposite tendencies in the difference between maintaining and suppressing. Taken individually, this tendency is only significant for the panic group ($p=0.044$), but the difference between the two groups' tendencies is more significant ($p=0.036$). Graph compiled by Dr. Andrea Reinecke.

were calculated from very short signals. The spectra can be expected to be less stable at lower frequencies, for which fewer cycles were captured. More leakage between frequencies can occur at this end of the spectrum, and a measured variation in *LF* power might in fact be explained by *VLF* oscillations that have leaked into the *LF* range. Such behaviour can be ruled out when there is no significant variation in the *VLF* power itself, as is the case for all but one of the *VLF* comparisons in Table 8.

When comparing the effects of maintaining and suppressing, only the normalised variables capture a statistically significant difference, and only for the panic group (second row in Table 8), although a near significant difference in the opposite sense is found for the control group.

A group effect can be found by subtracting the measurement for the 'suppress' blocks from that for the 'maintain' blocks. As shown in the penultimate row of Table 8, this approach exposes a significant difference between the two groups in the variables LF_n , HF_n , and LF/HF . All of these variables indicate that, for panic sufferers compared with controls, the difference between 'suppress' and 'maintain' is more of a shift towards sympathetic dominance over parasympathetic activity in the 'maintain' blocks. As shown in Figure 75, the balance shifts in the opposite sense for

the control group. Curiously, applying a similar approach to the crosshair-fixation blocks (last row in Table 8) reveals a near significant group difference in LF_n and HF_n that is of the same sense as the effect in the image blocks. If this effect is genuine it must be caused either by anticipation of the coming block or by a residual/recovery from the preceding image block.

A caveat to be noted when interpreting these p-values is that, when many separate t-tests are applied, it becomes likely that in some cases the null hypothesis will be wrongly rejected. For example, out of 100 independent t-tests that all yield $p=0.05$, one can expect the null hypothesis to be true in roughly 5 of these comparisons, even though their p-values suggest statistical significance.

The results presented in this section are discussed fully in section V.B.6 in combination with the results of two related experiments, which are presented in the intervening sections.

V.B.4. Second experiment – comparing subjects with different levels of neuroticism

V.B.4.a. Methods

Neuroticism is a personality trait associated with “a maladaptive emotion regulation and... alterations in ANS function”^[Di Simplicio et al, 2011]. This experiment, conducted in collaboration with Dr. Martina Di Simplicio at the University of Oxford, involved 30 subjects from the general population, selected for having particularly low or high neuroticism (scores of ≤ 6 or ≥ 16) according to the Eysenck Personality Questionnaire^[Eysenck & Eysenck, 1975]; 10 subjects fell into the high neuroticism group. The protocol was as described in Figure 73, this time without concurrent MRI scanning. An additional 1-minute ‘resting’ ECG recording was gathered while the subjects sat still in front of the computer before the start of the task.

V.B.4.b. Results

Subjective affect ratings:

As in the previous experiments, the subjective affective ratings were consistently higher during the ‘maintain’ blocks than during the ‘suppress’ blocks ($p < 0.001$), while no significant distinction existed between the groups when comparing the ‘maintain’ and ‘suppress’ affect ratings ($p = 0.5$).

HRV analysis:

After HRV parameters were calculated for all subjects, these values were provided to Dr. Di Simplicio, who carried out statistical analyses using repeated-measures ANOVAs. A summary of the most salient results is presented in the following paragraph, allowing the ‘spectral averaging’ technique to be assessed in terms of the HRV variables’ abilities to discern effects of the stimulus.

Statistical significance of differences in HRV parameters

Comparison	<i>LF</i>	<i>HF</i>	<i>LF_n</i>	<i>HF_n</i>	<i>LF/HF</i>
Group	0.99	0.67	0.20	0.35	0.20
Task x Group	0.14	0.016	0.62	0.54	0.48

Table 9: p-values from a repeated measures ANOVA to expose group effects and interactions between group and task ('Maintain' vs. 'Suppress'). Resting values of each HRV variable were used as covariates in the ANOVA to control for their effect.




Figure 76: In low-neurotics, the power in the HF band of the HRV spectrum was found to be higher when suppressing one's emotional response to an aversive stimulus. In high-neurotics, a non-significant tendency of the opposite sign was found.

No differences between the two groups were found in the 'resting' portion of the ECG. However, these 'resting' measurements proved useful in controlling for the effect of individual differences in HRV. Table 9 summarises the p-values produced from an ANOVA applied to various HRV parameters to expose any significant differences between the two groups of subjects, with 'resting' values of each parameter entered as covariates. The first row examines the statistical significance of any group effect (difference between high- and low-neurotic subjects) across all image blocks. The second row examines the distinction between the groups in terms of how each subject's response to the 'Maintain' task differs from their response to the 'Suppress' task. Such effects are referred to as a 'task x group interaction'. Note that the only statistically significant effect revealed by this ANOVA is the task x group interaction found using the variable *HF*. When the same analysis was run without using 'resting' HRV parameters as covariates, the *HF* task x group interaction was only marginally significant ($p=0.05$), which demonstrates the effectiveness of controlling for individual differences.

As a supplement to the above-mentioned analyses carried out by Dr. Di Simplicio, individual t-tests were applied to explore the *HF* task x group interaction in greater detail. The *HF* variable revealed a significant effect among low neurotics, with higher values for the 'Suppress' task than 'Maintain' ($p=0.027$). An opposite tendency was indicated in high neurotics, but without statistical significance ($p=0.155$) (see Figure 76). It should be noted that the normalised HRV variables have not revealed any significant differences between subjects with high neuroticism and

those with low neuroticism, whereas in the previous experiment these were the most discriminative of the variables tested.

V.B.5. Third experiment – comparing the effects of an antidepressant and a placebo in subjects with high neuroticism

V.B.5.a. Methods

After the previous experiment revealed a significant difference between subjects with high levels of neuroticism and those with low levels, a follow-up experiment was designed to examine whether this characteristic was responsive to pharmaceutical intervention. Forty-two subjects with high levels of neuroticism (scores of ≥ 16 on the Eysenck Personality Questionnaire) were recruited from the general population. At the screening session, a five minute ‘resting’ ECG was recorded while the subjects lay supine in a quiet environment. Over seven days, each subject took a course of either a placebo or citalopram, which is one of a class of antidepressants known as selective serotonin reuptake inhibitors (SSRIs). At the end of the seven days, each subject underwent the emotion regulation procedure used in the previous experiments, with the modifications described below:

- The previous experiments exposed no significant difference between groups in terms of subjective affect ratings, even though objective measures (fMRI and HRV) did expose group differences. It was hypothesised that a difference in subjective affect ratings might be achieved with a stronger stimulus. Hence the IAPS pictures used for this experiment were selected for having more intense content, as measured by the rating system that is incorporated into the IAPS database.
- As in the second experiment, a 1-minute ‘resting’ ECG was recorded immediately before the start of the experiment to expose individual differences in HRV and to allow these differences to be accounted for in the statistical analyses.
- To prevent the orienting response described in section III.B.1 from corrupting HRV analysis, the length of the image and crosshair-fixation blocks was extended to 45 seconds, so that a signal of at least 30 seconds could be extracted after the orienting response subsided. Upon later inspection of the signals, it was determined that it would be sufficient to exclude the first 10 seconds of each leaving a 35-second section for HRV analysis.
- As mentioned in section III.B.1, the orienting response itself can be treated as an additional measurable ECG parameter. The indices used to describe the orienting response in this experiment were the minimum heart rate in the first 2 seconds after stimulus onset (HR_{min}) and the maximum heart rate in the next 2 seconds (HR_{max}), as proposed by [Hodes et al, 1985]. The mean heart rate from the 15 seconds preceding the stimulus onset was subtracted from these values to control for long-term fluctuations in heart rate.

7 subjects were excluded from HRV analysis on the basis of the quality of their ECG recordings. Another 3 were excluded as outliers in terms of the primary variable of interest, HF ,

Significance of differences between screening and pre-task recording

Comparison	LF	HF	LF _n	HF _n	LF/HF
Time (screening vs. pre-task recording)	0.20	0.009 +	0.12	< 0.001 -	0.001 +
Time x Group	0.82	0.69	0.67	0.50	0.85

Table 10: p-values from a repeated measures ANOVA examining the differences in ‘resting’ HRV between the initial screening and the pre-task ECG recording (after the seven day course of SSRI or placebo). In the first row, a ‘+’ or ‘-’ indicates the sign of the difference. For example, a ‘+’ indicates that the mean measurement was greater for the condition to the right of “vs.”.

Significance of differences between ‘maintain’ and ‘suppress’ conditions

Comparison	LF	HF	LF _n	HF _n	LF/HF
Task (Maintain vs. Suppress)	0.52	0.72	0.059 +	0.018 -	0.091 +
Task × Group	0.18	0.043	0.85	0.96	0.44

Table 11: p-values from a repeated measures ANOVA examining the differences in HRV between the ‘maintain’ and ‘suppress’ image blocks. Pre-task ‘resting’ values of the HRV parameters were entered as covariates to control for individual differences.

because they were found to have resting values that differed by more than three standard deviations from the group mean. The remaining cohort consisted of 33 subjects. Of these 17 were given placebo and 16 were given SSRI. Each group contained of 9 male subjects.

V.B.5.b. Results

As in the previous experiment, statistical analyses of the HRV parameters and other variables were carried out by Dr. Di Simplicio using repeated-measures ANOVAs, and a summary of the most relevant results is presented in the following paragraphs to enable a comparison of the performances of the HRV variables.

No significant difference was found between the subjects given placebo and those given the antidepressant in terms of the subjective affect ratings or resting HRV parameters.

Table 10 presents the p-values produced by an ANOVA used to expose any differences in resting HRV between the initial screening session and the 1-minute recording that immediately preceded the emotion-regulation protocol to see if the drugs affected baseline HRV. HF , HF_n , and LF/HF exposed significant differences between the two recordings (across all subjects). These effects should be interpreted with caution, because the subjects’ posture changed from supine in the screening session to sitting upright in the pre-task recording. This change can be expected to produce an increase in LF power and a decrease in HF power^[TFESCNASPE, 1996]. The HF_n and LF/HF changes described in the first row of Table 10 are in agreement with these expectations, but the increase in HF cannot be explained by the change in posture. No significant time x group interaction was found for any of the variables, so the increase in HF is presumably caused by other environmental factors, perhaps related to anticipation of the imminent task.




Figure 77: A significant group \times task interaction was found in HF. Although neither within-group difference is statistically significant (p-values shown from t-tests), the difference between these tendencies is significant ($p=0.043$).




Figure 78: Gender differences in HF during the image blocks. The drug SSRI appears to reduce HF when suppressing, relative to when maintaining ($p=0.043$). This effect is most significant in women.

Table 11 presents the p-values from another ANOVA, examining the differences between responses to the ‘maintain’ and ‘suppress’ image blocks. The variable HF_n was found to be significantly lower during suppression blocks than during maintain blocks, while LF_n and LF/HF showed a near significant tendency to be higher during suppression. A statistically significant group \times task interaction was found for the HF variable. This effect is presented in Figure 77. Subjects given a placebo showed a tendency to have higher values of HF in suppress than maintain, while drug takers presented an opposite tendency. Although neither tendency itself is statistically significant, the difference between them is.

The significance of this group \times task interaction was found to improve when controlling for gender differences; entering gender as a covariate in the ANOVA yielded $p=0.03$ for the HF group \times task interaction. The role of gender is explored in greater detail in Figure 78, which shows that this effect is rooted in the fact that females given placebo present significantly higher HF when

p-values, orienting response

Comparison	HR_{min}	HR_{max}
Task (Maintain vs. Suppress)	0.67	0.718
Task x Group	0.31	0.018

Table 12: p-values from ANOVAs used to expose statistically significant effects in the orienting response.

suppressing. Men in this group present no such difference. The antidepressant appears to reduce the HF power when suppressing compared to when maintaining, as seen in Figure 77. In Figure 78, this effect can be seen in both men and women, but it is much more significant in women. The gender differences in the drug's effects might be attributable to overall gender differences; a significant gender x task interaction was found ($p=0.028$), in that women presented significantly higher HF when suppressing than when maintaining, but men did not. Furthermore, one of the questionnaire-based psychiatric tests (the Emotion Regulation Questionnaire^[Gross & John, 2003]) revealed that the women were more disposed towards the use of reappraisal strategies to regulate their emotional responses in daily life ($p=0.045$).

Table 12 presents the results of a statistical analysis of the subjects orienting responses to the image blocks. HR_{min} , the minimum heart-rate in the window 0-2 seconds after the start of the picture block, did not expose any significant effects. HR_{max} , the maximum heart rate from the subsequent 2-second window, exposed a task x group interaction; in the placebo group, HR_{max} was significantly higher when maintaining than when suppressing. No such distinction emerged in the antidepressant group. This result might be seen as being in agreement with those presented in the previous paragraph, in the sense that the antidepressant appears to abolish the difference in autonomic response between suppressing and maintaining emotions.

V.B.6. Discussion

V.B.6.a. Methodological evaluation

A question of primary concern for this thesis is 'Did the HRV parameters, calculated using the spectral averaging method outlined in section III.B.1, provide meaningful indices of the autonomic response to the experiment?' Statistically significant differences in HRV parameters were found between psychologically distinct groups, and as an effect of pharmaceutical intervention. However, one cannot immediately assume that these results represent effects on the autonomic response.

Can the observed effects in LF and HF be explained as artefacts of VLF effects?:

As illustrated in Figure 42 (page 90), occasional large peaks centred in the VLF band appeared, and the power from these peaks leaked into the LF band and, to a lesser extent, the HF band. This contribution to the LF and HF bands can be assumed to be unrelated to the ANS mechanisms ordinarily associated with these frequencies. It was therefore necessary to ascertain whether this 'leakage' from the VLF band might account for the statistically significant responses in other HRV parameters.

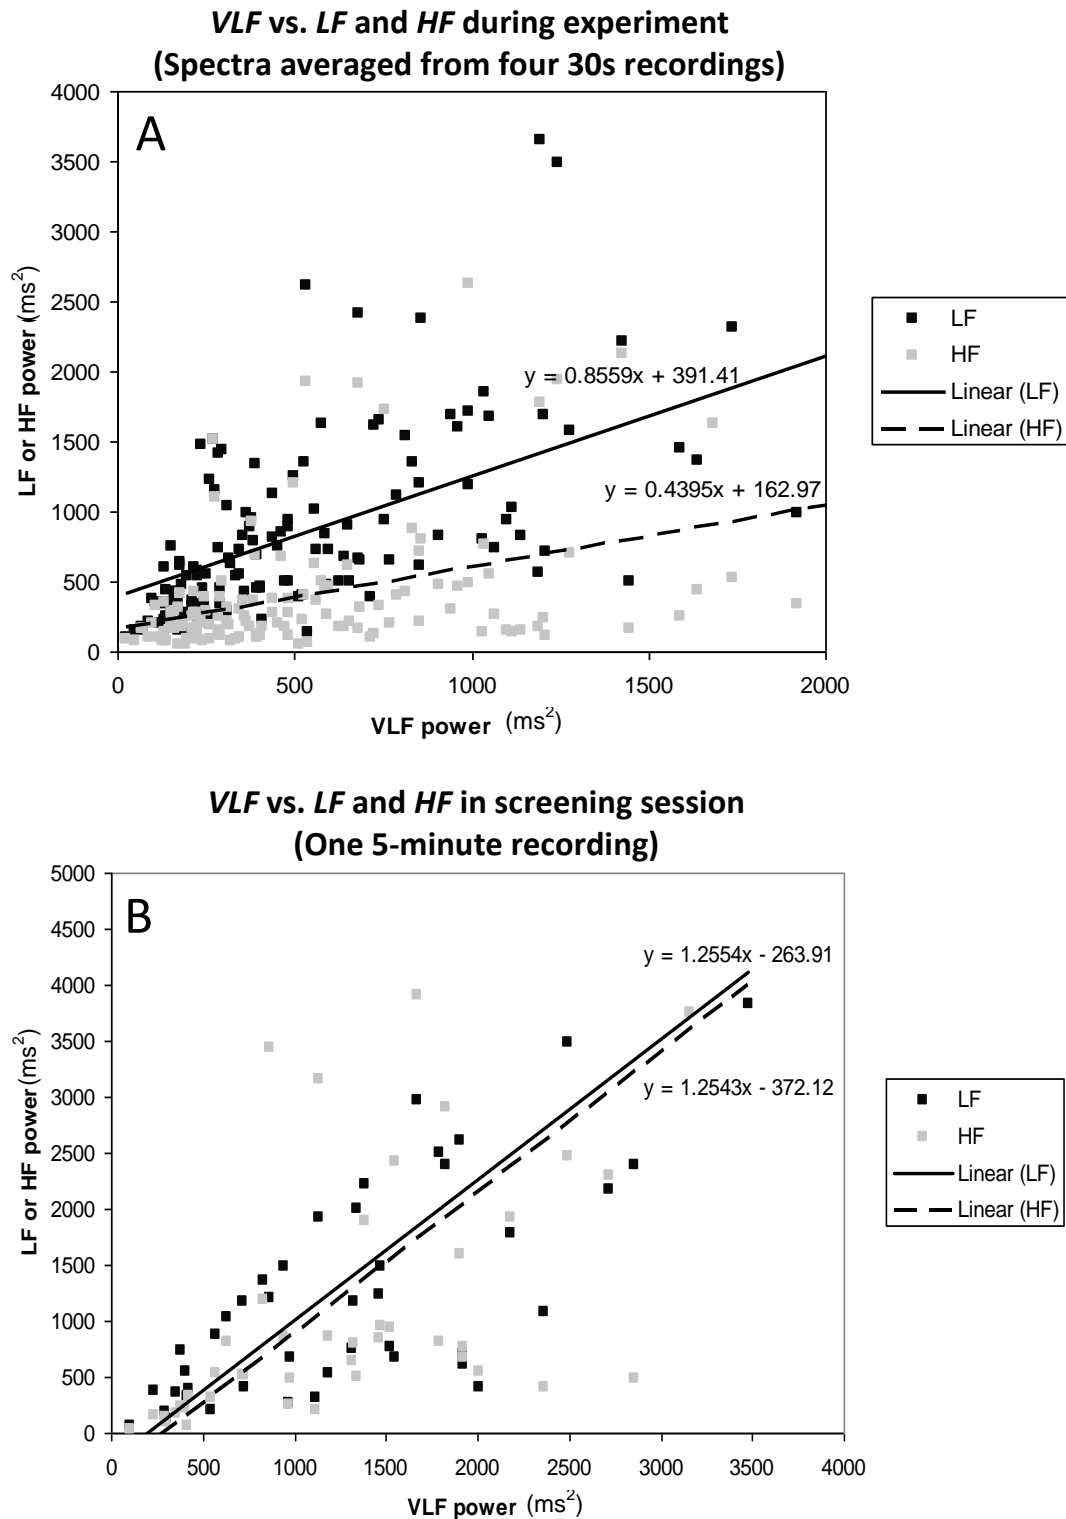


Figure 79: Correlations between *VLF* and the two basic indices of autonomic activity, *LF* and *HF*. Although positive correlations exist, they cannot be wholly explained by low-frequency leakage due to the use of short signals; the correlations still exist when longer recordings are used.

Figure 79 explores this question using as an example the results of the second neuroticism experiment, in which a significant $\text{Group} \times \text{Task}$ interaction was found for *HF*. Figure 79.A. shows that, during the experiment (all image and fixation blocks), *LF* and *HF* both presented a positive correlation with *VLF*. One might expect such a relationship to be induced by leakage from large *VLF* peaks, which would be exacerbated by the use of excessively short recordings. However, a

similar correlation was observed in the 5-minute recordings from the screening session, as shown in Figure 79.B. Hence the relationship cannot be wholly explained as an artefact from the use of short signals. Still, the fact that a much more positive correlation appears in the *LF-VLF* relationship than in the *HF-VLF* relationship when short signals are used (Figure 79.A) suggests that leakage does indeed play a significant role, since leakage from VLF into the adjacent LF bandwidth will be more substantial than leakage into the HF bandwidth. As suspected, the meaningfulness of the variable *LF* declines more substantially than that of *HF* when spectra are calculated from signals shorter than the recommended 2 minutes, due to VLF oscillations contributing more significantly to the *LF* values.

Taken alone, this insight does not answer the question of whether the statistically significant results from the Oxford experiments are meaningful. The veracity of individual results presented in the preceding sections must be explored on an individual basis.

In the experiments with neurotic subjects, the most important results came from the variable *HF* (see Table 9, Table 10, Table 11; pages 156 and 158), which has been shown to be less susceptible than the variable *LF* to VLF leakage. The notion that these *HF* results might in fact be driven by leakage from VLF changes is contested by the fact that they were not accompanied by statistically significant results in the intervening LF band. If *VLF* leakage was responsible for the effect, one would expect its influence to be more pronounced in the LF band.

In the experiment comparing panic-attack sufferers with control subjects, differences between the two groups were found in the variables LF_n , HF_n , and LF/HF . These variables are all sensitive to changes across both LF and HF frequencies. The absence of any significant effect in the VLF band implies that these results cannot be attributed to leakage from outside the LF-HF frequency range. The fact that these variables presented significant results where *LF* and *HF* did not suggests that the salient difference in *HF* between groups was accompanied by an opposite change in *LF*.

Did the non-stationary orienting response corrupt *LF* and *HF* results?

An important limitation in the first two Oxford experiments is that the signals used for HRV analysis did not exclude the orienting response. In the third experiment, changes in the variable HR_{max} indicated that neurotic subjects (on placebo) experience a more substantial short-lived acceleration in heart rate when maintaining than when suppressing their emotional response.

If a similar tendency occurred in the first two experiments, it is likely to have induced a bias towards increased measurements in the LF and/or HF band for the ‘maintain’ condition, since the time-scale on which the orienting response occurs is roughly matched to these frequency ranges. As explained in the previous section, the results of the experiment with panic sufferers indicate opposite changes in the LF and HF bands, so these changes are unlikely to be caused purely by a bias induced by the orienting response. In the experiment comparing subjects with high and low neuroticism, high neurotics did present a non-significant tendency towards increased *HF* while maintaining, but low neurotics presented a significant tendency in the opposite sense. This result could only be explained as a bias from the orienting response if the task dependency of HR_{max} is

opposite in low-neurotics to that in high-neurotics, a possibility which may be explored in future analysis.

In the experiment examining the effects of SSRI antidepressants, longer task blocks were used so that the first 10 seconds of each block could be excluded from HRV analysis, thereby avoiding the introduction of bias from the orienting response. In this case, HF still proved to be sensitive to changes in heart rate variability related to emotional processing. However, the change in HF between suppressing and maintaining for high-neurotics was very different in the third experiment (Figure 78) compared with the second (Figure 76). This discrepancy may be caused by corruption from the orienting response or by a genuine change in the HF spectral content, perhaps due to the use of a more intense set of images in the third experiment.

Comparing the performances of different HRV parameters:

The first of the Oxford experiments revealed significant differences between panic-attack sufferers and controls in three variables: LF_n , HF_n , and LF/HF . Of these, LF/HF has arguably the most straightforward interpretation. As described in section II.G.2, it can be interpreted as a reflection of the balance between sympathetic and parasympathetic outflow to the heart. In section III.B.1.c, it was shown that LF_n and HF_n can also be interpreted as measures of sympathetic-parasympathetic balance, with mathematical properties that differ substantially from those of LF/HF .

The variable that most often exposed a statistically significant effect was HF_n . This might be explained by the fact that HF_n is more sensitive to changes that occur when LF is small. It was noted in section III.B.1.b that the application of HRV analysis to ECG recordings shorter than one minute can be expected to yield highly variable estimations of power at low frequencies. Figure 43 (page 93) shows that the magnitude of $\partial HF_n / \partial LF$ is much greater for smaller values of LF , and $\partial HF_n / \partial HF$ is greater for smaller values of LF as long as $LF > HF$, which is usually the case^{Bigger et al., 1995}. As explained in the previous section, LF measurements in these experiments were often corrupted by 'leakage' from large peaks in the VLF range. The normalised variables may have given better statistical significance than LF/HF by being less sensitive to cases with spuriously large LF values, where VLF leakage will act as a source of noise. In contrast, the sensitivity of LF/HF to changes in the HF band increases monotonically as LF increases (Figure 44, lower panel), so that the measure is more sensitive to the cases with greater low-frequency corruption. the sensitivity of LF/HF to changes in the LF band does not vary with LF (upper panel).

To summarise, the apparent effectiveness of HF_n in distinguishing experiment effects is likely to be caused by its selectivity; during statistical calculations, this variable effectively gives greater weighting to cases where LF is small (and where, implicitly, the corrupting effect of using short-time recordings is lower). In these experiments, that characteristic was useful in that it exposed statistically significant effects on sympathetic-parasympathetic balance that were not detected by LF/HF . However, because HF_n introduces uneven weighting to the statistical calculations, it is possible that the effect is confined to a subset of patients (those who were given greater weighting).

V.B.6.b. Physiological insights

Comparing panic sufferers and control subjects:

The results in LF_n , HF_n , and LF/HF from the first Oxford experiment indicate that the autonomic balance in panic sufferers, when suppressing one's emotional response, is similar to that of 'normal' subjects. When maintaining their response, however, panic sufferers present a significant shift towards sympathetic dominance, whereas controls do not (Figure 75). This phenomenon might be expected from the general nature of panic disorders, in that the swing towards sympathetic dominance could be interpreted as an exaggerated defence response.

This result has three important implications. Firstly, it demonstrates that HRV can serve as a useful objective measure of a subject's response to psychological stimuli of this nature, revealing significant differences between groups while self-reported subjective measures (affect ratings) did not. Secondly, it provides evidence of psychosomatic consequences of panic disorders. As summarised by [Hemingway et al, 2001], numerous studies have found a link between panic disorder and the occurrence of sudden cardiac death (SCD). This new evidence suggests a mechanism by which such deaths occur. Increased sympathetic dominance in response to a threat allows increased cardiac output at the expense of electrophysiological stability. When such a response is generated more frequently than necessary, there are more chances for a fatal arrhythmia to develop. A third important implication of the findings is that they point to a treatment method; the fact that group differences in autonomic behaviour were abolished during strategic suppression of the emotional response suggests that psychiatric coaching would be useful in attenuating the cardiac risks associated with panic disorders.

Substantial overlap occurs between the HRV parameters of the panic and control groups, so psychiatric tests alone would not be useful in any hypothetical screening for vulnerability to SCD. However, it may be that the statistical differences are driven by a few 'extreme responders' who are also particularly vulnerable to SCD. That possibility might be explored in future by running tests similar to this experiment in combination with long-term tracking of patient outcomes, in order to expose any correlation between the subjects' response to the experiment and the incidence of SCD. If these 'extreme responders' do present increased vulnerability, a similar procedure might be used to screen for this risk among panic patients to improve the assessment of treatment priorities.

Neuroticism and pharmaceutical intervention:

In the second Oxford experiment, a distinction between subjects with high and low neuroticism was found in that low neurotics showed a more distinct difference in HF according to whether they suppressed or maintained their emotional response to the images, with higher HF during suppression. The physiological insights from this experiment are described in a recently published paper^[Di Simplicio et al, 2011]. As described in the previous section, it cannot be confirmed whether the HRV results of this experiment were corrupted by changes in the orienting response. However, the third Oxford experiment used an improved measurement to exclude the effects of the orienting response. In this experiment, it was found that the administration of SSRI (antidepressant) reduced HF power when neurotic subjects attempt to suppress their emotional response. If HF is taken as a measure of the mean level of parasympathetic outflow to the heart,

the effect of the drug can be interpreted as a reduction of the subject's ability to consciously suppress the defence response and the associated decline in cardiac stability.

V.B.7. Conclusions

The fact that statistically significant results were discovered in all three of the Oxford experiments indicates that spectral averaging, the modified approach to HRV analysis proposed in section III.B.1, does indeed yield useful, objective measures of cardiac regulation related to psychological activity. In particular, the method has revealed psychosomatic differences between psychologically distinct groups, as well as changes in psychosomatic behaviour caused by pharmaceutical intervention. In general, this approach extends the applicability of HRV analysis by allowing the characterisation of responses that are attenuated by habituation on a time-scale shorter than the normal recommended 2-minute minimum signal length. The spectra produced showed instability and leakage among low frequencies, which may explain why statistically significant group differences were not discovered in the variable LF . However the unique properties of the normalised measurements, LF_n and HF_n , exposed significant changes in autonomic balance that were not captured by the other HRV parameters. These benefits can be expected in any experiment employing short-time recordings.

V.C. Exploring respiratory effects on ventricular repolarisation

V.C.1. Introduction

V.C.1.a. Objectives

Breathing is known to play a role in cardiac autonomic regulation; the most familiar manifestation is respiratory sinus arrhythmia (RSA, described on page 68), which arises from autonomic innervation of the SA node. These heart rate oscillations and the associated oscillations in blood pressure and other physiological measurements are widely observed and are often used as surrogate measures of autonomic behaviour^[TFESCNASPE, 1996]. However, the underlying mechanisms remain the subject of substantial research interest and debate^{[Eckberg, 2009], [Karemaker, 2009], [Larsen et al, 2010]}. The autonomic nerves are known to be present in myocardial regions other than the SA node^[Kowallik & Meesmann, 1995], but the nature and implications of respiratory involvement in myocardial electrophysiology have not been studied extensively.

As explained in section II.B, action potential duration is known to vary with heart rate due to the rate sensitivity of the myocytes. It is not clear to what extent direct input from autonomic nerves might have a complementary or antagonistic effect on this behaviour. Such effects may have arrhythmic potential, especially if they emerge heterogeneously (varying across the myocardium). Their exposition could give rise to new diagnostic and therapeutic approaches for the known link between respiratory disorders and sudden cardiac death^{[Hanly & Zuberi-Khokhar, 1996], [Lanfranchi et al, 1999], [Javaheri et al, 2007]}. This section covers a series of experiments designed to explore the effects of respiration on ventricular repolarisation properties.

The UEG analysis algorithms described in section IV facilitate the use of human *in situ* recordings for this investigation, providing unique data on the behavior of the intact human nervous system. Hence the experiments were also used to explore the effectiveness and utility of these novel algorithms, as will be discussed in section V.C.4.

V.C.1.b. Overview

Conducted in collaboration with Dr. Jaswinder Gill at St. Thomas' Hospital, London, the experiments all involved the use of unipolar electrogram (UEG) recordings inside the left and right ventricles of human patients, immediately before they underwent ablation therapy to treat atrial arrhythmia. The patients' hearts were paced at a constant rate using an electrical stimulus to control for the effects of heart rate, thereby exposing rate-independent effects of respiration on local electrophysiology.

Initially, these experiments focused on the hypothesis that changing the frequency of respiration causes heart-rate-independent changes in action potential duration (measured as activation-recovery intervals, ARIs). During preliminary investigations, it was noted that ARI oscillated at the respiratory frequency. Further investigations explored the characteristics of this phenomenon in greater detail – specifically, the magnitude and distribution of the oscillations

across the ventricles, their persistence during breath-holds, the effects of drug-induced autonomic blockade, and the phase relationships between respiration, blood pressure, and ARI. These analyses demonstrate for the first time that respiration drives oscillatory dynamics of repolarisation properties via direct autonomic input to the ventricles.

Furthermore, it is shown that, in the context of the abnormal autonomic behavior associated with Central Sleep Apnea (CSA), this mechanism could generate arrhythmias. During recording, five subjects spontaneously exhibited breathing patterns indicative of Central Sleep Apnea (CSA). This syndrome is characterised by a regular waxing and waning of respiratory volume, as seen in Figure 87 (page 179). As the name suggests, the symptoms usually arise in sleep, but are occasionally observed in conscious subjects. CSA often appears in patients with congestive heart failure, and several studies have found it to be linked with increased mortality in this cohort^{[Hanly & Zuberi-Khokhar, 1996], [Lanfranchi et al, 1999], [Javaheri et al, 2007]}. The syndrome is associated with various autonomic abnormalities, including a shift towards sympathetic dominance of the autonomic balance^[Lanfranchi et al, 1999]. Furthermore, patients with CSA have been found to experience more arrhythmic events during sleep than patients without CSA^[Javaheri, 2006]. This evidence suggests that autonomic control of the heart may predispose CSA patients to unstable cardiac behaviour. The results presented in section V.C.3 and discussed in section V.C.4 indicate that CSA alters the normal cardiovascular rhythms associated with respiration, providing a candidate mechanism for the known link between CSA and arrhythmia. These results are published in [Western *et al*, 2011].

V.C.2. Methods

V.C.2.a. General method

Although various stimulus protocols were used for the experiments described in this section, the electrode configuration was the same for each patient. Two catheters, each with 10 electrodes distributed from the tip along a 35 mm length, were inserted in the femoral vein and fed through the inferior cava to reach the heart's right atrium. One catheter was then fed through the tricuspid valve and positioned in the right ventricle, against the septal wall. The other catheter was pierced through the atrial wall to the left atrium, then fed through the mitral valve and positioned against the posterior wall of the left ventricle. A third catheter was inserted in the right ventricle and used to pace the ventricles (to control the timing of activation by delivering a brief electrical stimulus) near the apex. The final positions of the catheters is shown in Figure 80. The pacing period varied between subjects, but was chosen in the range 350 ms – 500 ms as the longest period at which the pacing stimulus was found to consistently pre-empt the heart's natural rhythm. The reference electrode was a 100 mm x 150 mm rectangular patch positioned on the patient's skin at the navel. Activation and recovery times (A_S and R_S) were determined for each beat in each electrode using the algorithms described in section IV.D.4.

Patients were asked to control their breathing by following an on-screen animation developed by Dr. Ben Hanson, allowing the respiratory frequency to be controlled. For confirmation, breathing behaviour was continuously recorded using a tension sensor attached to an




Figure 80: *Left* – Diagram showing the positioning of the electrode catheters within the heart. RA and LA are the right and left atria. RV and LV are the right and left ventricles. *Right* – X-ray image of the catheters in place.

elasticated band around the patient's abdomen. The tension measured in this band was assumed to be approximately proportional to the circumference of the band, and thus gave an un-calibrated representation of lung volume. Arterial blood pressure (ABP) was recorded from the femoral artery since blood pressure has a feedback role in autonomic cardiovascular regulation. ABP, abdominal circumference, and electrograms were recorded in synchrony at a sample rate of 1200 Hz.

Ethical approval for the studies was granted by the Guy's and St. Thomas' Hospitals Ethics Committee and conformed to the standards set by the Declaration of Helsinki (latest revision: 59th WMA General Assembly). Informed consent was obtained in writing from all the subjects. Subjects were not sedated, and the use of any anti-arrhythmic drugs was discontinued from 5 days prior to the study.

V.C.2.b. Testing for effects of breathing rate on steady-state ARI

As explained in section V.C.1.b, the first question explored in the St. Thomas' experiments was whether respiratory rate caused any increase or decrease in the action potential duration of ventricular myocytes. A relationship of this kind could potentially arise through various mechanisms, such as a hypothetical change in mean autonomic outflow when the respiratory gating effect (page 68) is driven at different frequencies.

16 subjects were asked to breathe at a rate of 6 BrPM (breaths-per-minute), 9 BrPM, 15 BrPM, and 30 BrPM for 2 minutes at each frequency. The respiration signal recorded from the chest-band was used to identify the interval within which the patient successfully followed the intended breathing pattern, which was typically longer than 90 seconds. For each electrode in the recording, the median value of the ARIs in this interval was taken as the representative 'steady-

state' value at that breathing rate. The median was preferred to the mean as a representative measure because the median is less sensitive to outlier measurements.

The results of this analysis are illustrated in Figure 81 (page 172).

V.C.2.c. Assessing the magnitude and distribution of oscillations in ARI

Oscillations in ARI were observed at some ventricular measurement sites for every subject at every breathing frequency tested. The following method was designed to quantify the magnitude and significance of these oscillations during respiration at 15 BrPM (0.25 Hz) across all electrodes. This breathing rate was chosen because it lies near the centre of the HF band (0.15-0.4 Hz), normally associated with spontaneous breathing.

- I. Any voids left in the ARI series by the auto-vetting of recovery time measurements were replaced using cubic-spline interpolation.
- II. An AR model was fitted to the each ARI series, and the power spectral density of the series was calculated, as per section II.G.2.
- III. The 0.25 Hz component of this power spectrum was compared with a reference 'noise' bandwidth, which covered the natural range of respiratory frequencies, 0.15 Hz – 0.4 Hz, while excluding the range 0.2 Hz – 0.3 Hz to avoid overlapping the true respiratory frequency. The test statistic Z was defined as

$$Z = \frac{M_R - \bar{M}_N}{\sigma_N} \quad (44)$$

where M_R is the magnitude of the 0.25 Hz spectral component, \bar{M}_N is the mean magnitude of the noise band, and σ_N is the standard deviation of the noise band. Respiratory ARI oscillations in a given electrode were said to be statistically significant when $Z \geq 1.96$, equivalent to a 97.5% confidence limit if the noise band is assumed to have a Gaussian distribution. This approach was inspired by a method advocated by [Martinez & Olmos, 2005] for the automatic detection of the phenomenon known as T-wave alternans, typically identified as a beat-to-beat oscillation in T-wave magnitude.

- IV. Electrodes were excluded from analysis if more than 10% of beats were rejected by the recovery-time detection algorithm or if no AR model could be found to adequately describe the spectral content of the series, based on the prediction-error whiteness test described on page 75.

V.C.2.d. Assessing phase relationships.

The phase relationships between oscillatory signals are often used to grant insight into the causality between them. As pointed out in the recent debate on the origins of respiratory sinus arrhythmia, respiratory oscillations in cardiovascular behaviour are non-sinusoidal and their

morphology can vary. Conventional mathematical definitions of phase (approximating the signal as a sinusoid) can therefore produce misleading results^[Karemaker, 2009]. However, phase relationships can still be examined by comparing the timings of particular events in the signals. As shown in Figure 83, the timings of the onset of inspiration and exhalation for each breath were determined by visual inspection for each breath. These events were chosen as timing indices because they were found to produce prominent, short-lived artefacts in the respiratory recordings for each patient, thereby producing a reliable, temporally well-defined marker to be compared against the extrema of the ARI and blood pressure waveforms.

During cases of CSA, an interesting pattern was observed in the morphological relationship between respiration and respiratory oscillations in ABP. It was noted that the latency between onset of inspiration and the subsequent peak in systolic ABP gradually lengthened during the development of each hyperpnea (crescendo in respiratory volume), then shortened (see Figure 88, page 180). The following method was developed to quantify this effect:

- I. The timings of the onset of inspiration and the maximum systolic blood pressure were determined for each breath using a simple custom Matlab algorithm. Any erroneous estimations of these timing indices were corrected manually.
- II. The delay between onset of inspiration and peak systolic blood pressure was calculated for each breath.
- III. The start, middle (peak), and end of each hyperpnea were identified from visual inspection of the respiratory trace. A representative value of the delay at each of these events was calculated as the mean of the three values nearest the event.
- IV. For each subject, the mean of these representative values was calculated for start-, middle-, and end-hyperpnea across all CSA cycles.
- V. Welch's t-test was applied to assess the statistical significance of the difference between start- and middle-hyperpnea and between middle- and end-hyperpnea, as summarised in Figure 89 (page 181).

V.C.2.e. Inducing autonomic blockade

To investigate the role of autonomic nerve activity in these oscillations, 6 subjects had drugs administered intravenously to alter sympathetic and parasympathetic input to the heart. After performing the original protocol at four different breathing frequencies, metoprolol-tartrate (10-15 mg) was administered to partially block sympathetic activity and the protocol was repeated. Atropine sulphate (≈ 1.2 mg) was then administered to partially block parasympathetic activity, and the protocol was repeated a third time. Pacing was briefly halted while the drugs were administered so that the subject-specific dosage could be controlled to achieve a change in heart rate of approximately 10 BPM from each drug.

To provide a single statistic describing the change in the magnitude of ARI oscillations, the electrode with the clearest oscillations across all frequencies (before drugs) was identified for each subject. The power spectrum of the ARI series across the full breathing protocol was calculated using the Fourier Transform, and the power spectral density was averaged across the frequency range 0.08 Hz – 0.6 Hz to give a single value representing the magnitude of ARI oscillations for that drug condition in that subject.

To allow a more detailed portrayal of the drugs' effects on individual patients, a wavelet transform (see section II.G.2) was applied to the respiratory signals and ARI series to show their frequency content as it varies throughout the protocol. The wavelet family chosen for this analysis was the Gaussian-5 wavelet, as defined in the Matlab Wavelet Toolbox™. A Gaussian-type wavelet was chosen because its frequency spectrum (Fourier Transform) has a single smooth peak, thereby allowing a meaningful mean frequency to be assigned to each scale (s in equation (21), page 77). The Gaussian-5 wavelet was chosen in particular because it was found to offer an appropriate compromise between temporal resolution and frequency resolution for these experiments, allowing experiment events to be clearly distinguished along both axes. Example results are provided in Figure 86 (page 178).

V.C.3. Results

V.C.3.a. Respiratory rate does not affect steady-state ARI

As shown in Figure 81, no substantial change in median ARI occurs as an effect of breathing rate. However, the small differences in median ARI between 15 BrPM and 6 BrPM and between 30 BrPM and 6 BrPM appear to be statistically significant when Welch's t-test is applied to the data ($p \leq 0.05$ in both cases). This effect might be related to the increased effort required to breathe at the more unnatural rates (6BrPM and 30 BrPM), as will be discussed in section V.C.4. However, the results should be interpreted with caution; the data used to construct each plot are not truly independent readings since each datum belongs to a set of up to 20 that correspond to simultaneous recordings from a particular patient.

V.C.3.b. Respiratory oscillations in repolarisation properties

Although respiratory frequency was found to have no substantial effect on median values of ARI, ARI was observed to oscillate at the respiratory frequency (as exemplified in Figure 82, page 173). The phenomenon may have important consequences with regard to the development of arrhythmia, especially when considering that it may be modulated heterogeneously by a non-uniform distribution of autonomic connections to the myocardium. This section describes further analysis of these respiratory oscillations in ARI, which have not been previously described in the literature.




Figure 81: Histograms showing the difference in median ARI between 6 breaths-per-minute and three other breathing rates. These calculations were performed for 20 electrodes in 16 different subjects. Electrodes were excluded if fewer than 90 percent of beats produced successful recovery time identification. Also shown are the mean difference and associated p-value (from Welch's t-test) for each graph. It must be noted that, as explained in the text, these are not true t-tests because the data are not truly independent readings.

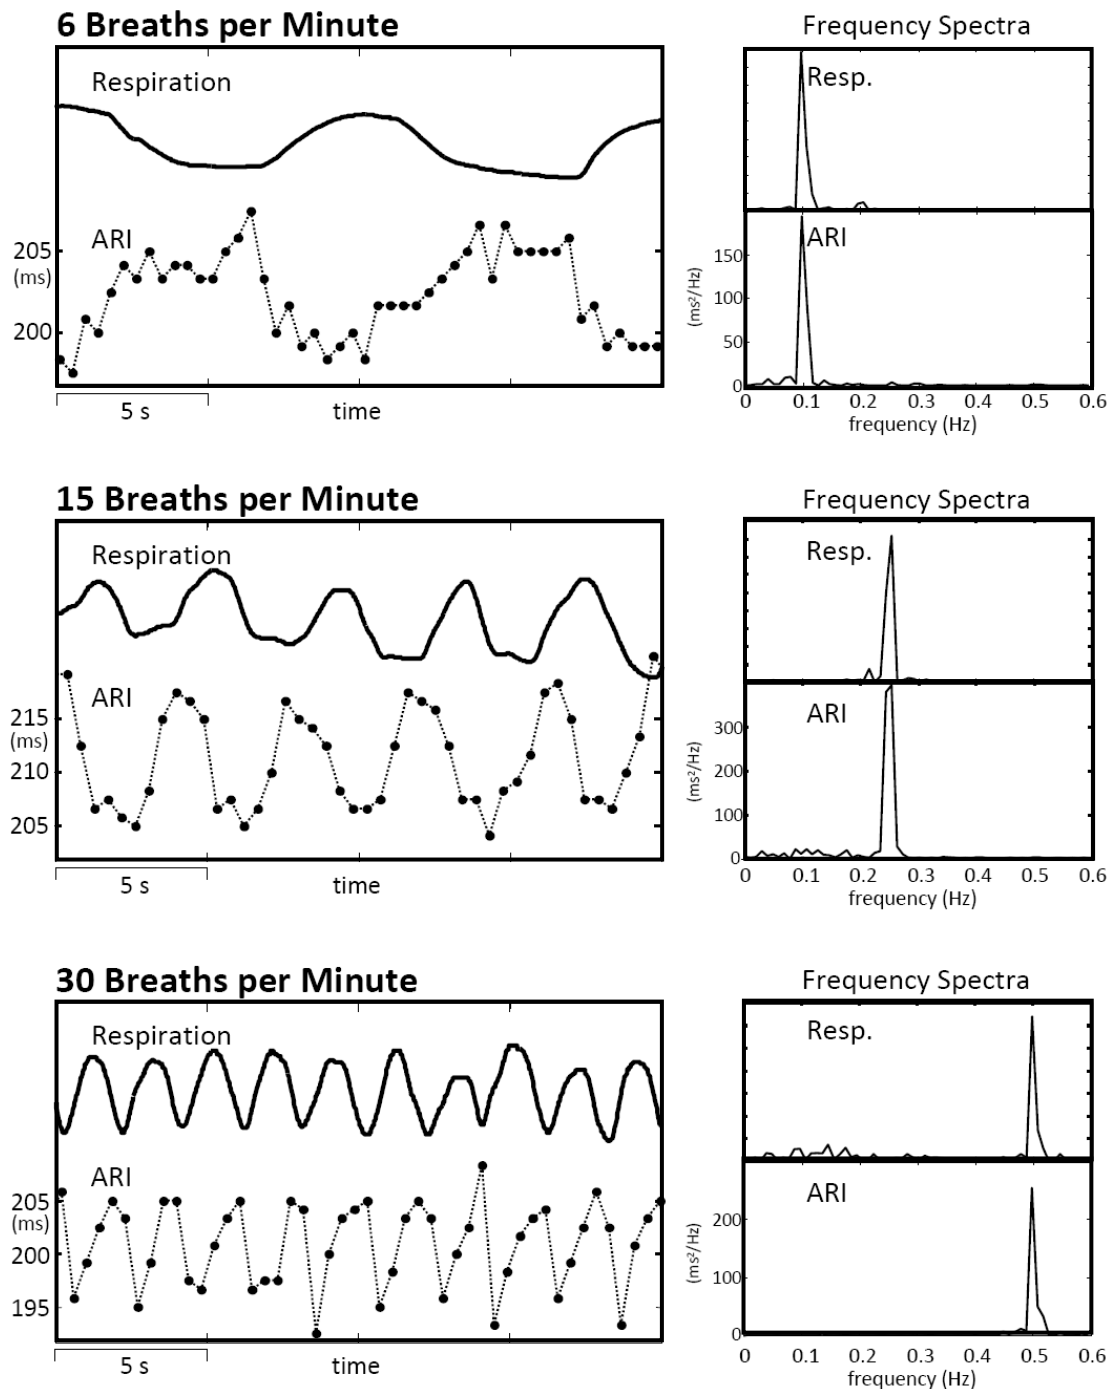


Figure 82: Examples of oscillations in ARI at three different respiratory frequencies, along with spectra calculated from the respiratory trace (abdomen circumference) and ARI series using the Fourier Transform. The spectra confirm that the dominant frequency of the ARI oscillations matches the respiratory frequency precisely.

Prevalence and magnitude:

In the cohort of 16 subjects, the mean number of electrodes (out of 20) excluded from analysis was 4.7 (max. 11, min. 0). The mean number of electrodes in which statistically significant respiratory oscillations in ARI were detected was 11 (max. 20, min. 6). These results are summarised in Table 13. Statistically significant respiratory oscillations were also observed in systolic and diastolic arterial blood pressure (ABP), measured as the maximum and minimum values, respectively, of ABP in each beat. The amplitude of oscillations in systolic ABP was greater than in

Subject	# of electrodes rejected	# of electrodes with significant oscillations	Peak-to-peak amplitude of largest ARI oscillations (ms)	Location of largest ARI oscillations	Amplitude of oscillations in systolic ABP (mmHg)	Amplitude of oscillations in diastolic ABP (mmHg)
1	2	12	14	RV 2	21	10
2	0	20	8	RV 1 and LV 6	7	2
3	3	12	26	RV 1	10	8
4	11	6	19	RV 2	16	15
5	6	9	25	LV 9	18	7
6	4	11	7	LV 8	8	4
7	5	14	20	RV 5	16	6
8	9	7	24	RV 8	10	8
9	6	11	16	RV 10	15	7
10	2	13	18	RV 6	24	8
11	4	16	20	LV 5	22	11
12	0	11	7	RV 6	10	5
13	5	10	7	RV 1	9	4
14	8	7	11	LV 10	20	17
15	9	8	6	LV 2	11	6
16	1	6	12	RV 10	12	6

Table 13: Summary of oscillations in activation-recovery intervals (ARI) and arterial blood pressure (ABP) across 16 subjects. Measurements of the peak-to-peak amplitude of oscillations were conducted by visual inspection to avoid corruption by noise. The locations found to present the largest ARI oscillations for a particular individual are labelled LV or RV (left or right ventricle) and numbered 1 to 10, indicating the sequential position of the electrode on the catheter. Electrode 1 is that closest to the apex of the ventricles (catheter tip), and electrode 10 is closest to the base of the ventricles.

diastolic ABP in every case. These blood pressure oscillations may be mechanistically related to the ARI oscillations via baroreflex feedback, as will be discussed in section V.C.4.

Phase relationships:

In Figure 83, vertical lines have been added manually to mark the onset of inspiration (dashed line) and onset of exhalation (solid line), omitting any cases where no such point can be clearly identified.

Where significant respiratory ARI oscillations were found, inspiration was found to coincide with a shortening in ARI. The coarse time-resolution of the ARI measurement (one heart-beat interval) makes it difficult to ascertain whether this shortening of ARI begins before, after, or in synchrony with the onset of inspiration. However, in a few cases, such as the 12 BrPM pane in Figure 83, the shortening of ARI consistently began more than one full beat before the onset of inspiration. The lengthening of ARI typically began before the onset of exhalation, often by several beats during slow breathing rates.

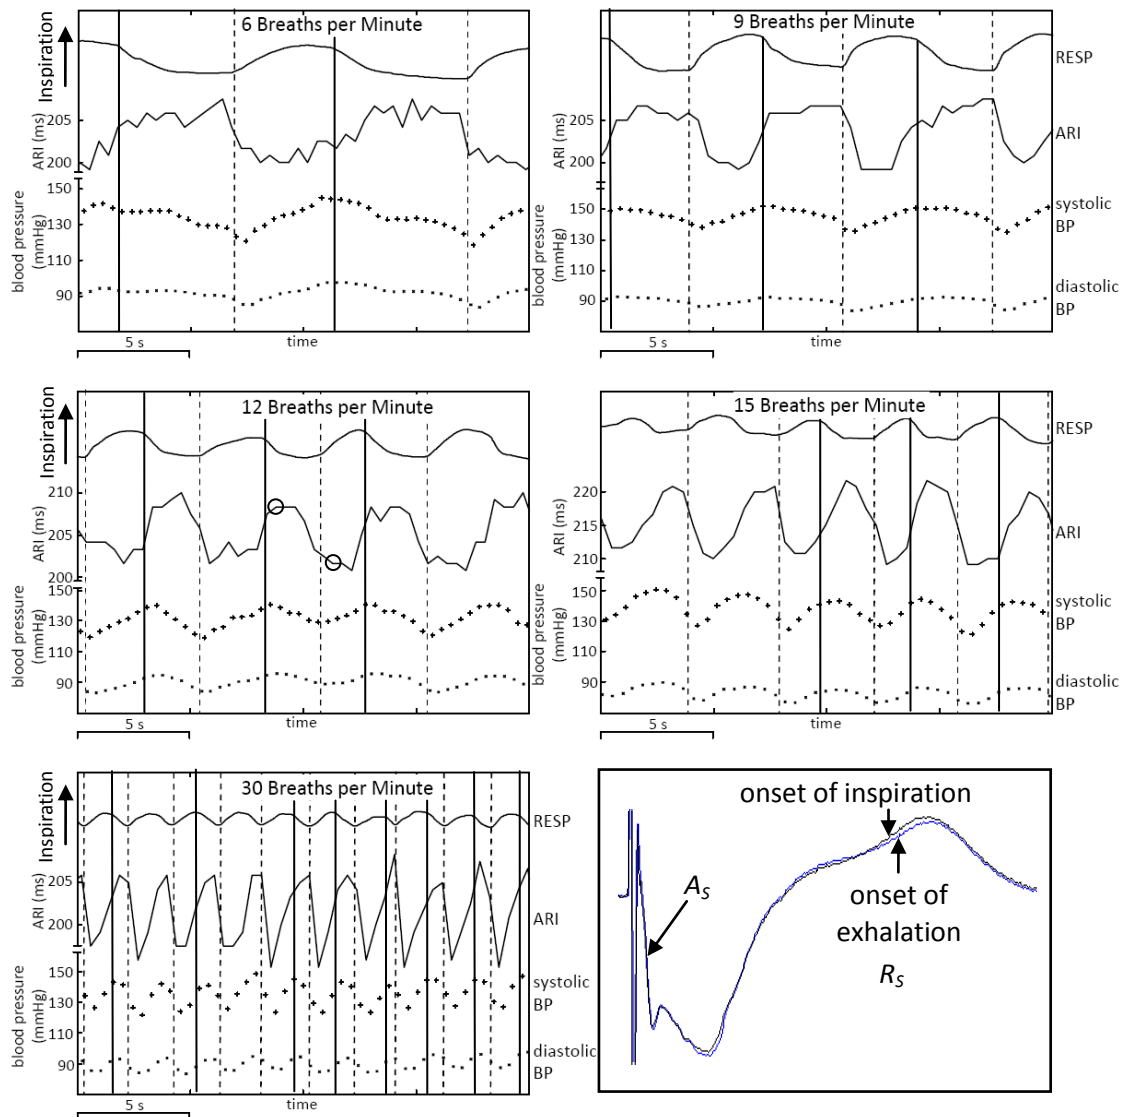


Figure 83: Examining the phase relationships between respiration, ARI, and blood pressure at five different respiratory frequencies for a single subject. All ARI series are calculated from the same electrode (RV 5). Dashed lines mark the onset of inspiration, and solid vertical lines mark the onset of exhalation. In the bottom-right pane, the electrograms corresponding to the two circled beats in the 12 BrPM pane are superimposed for comparison, temporally aligned so that the pacing artefacts (the first deflection seen) are coincident. The measured activation times (A_s) are also coincident, whereas the measured recovery time (R_s , marked by vertical black arrows) occurs earlier on inspiration than on exhalation.

The timing of the oscillations in systolic and diastolic ABP was such that the minimum pressures occurred at or just after (within 2 beats of) the onset of inspiration. The timings of the maxima were more variable, but coincided with the beginning of exhalation to within one quarter of a breathing cycle.

Oscillations during breath-holds:

14 subjects were also asked to perform breath-holds during the protocol. ARI oscillations at frequencies in or near the respiratory range were observed in 5 of these 14 patients during the breath-hold. The recording of respiratory movements confirmed that no corresponding abdomen motion occurred at this time. The full range of frequencies observed was 0.27 Hz – 0.7 Hz. The amplitude of the oscillations varied through the breath hold, but never exceeded that of the




Figure 84: Oscillations in ARI and ABP during a breath-hold.

oscillations observed during regular breathing in the same electrode. Blood pressure generally oscillated at the same frequency as ARI during breath-holds, but with reduced amplitude compared to during regular breathing, as shown in Figure 84.

Effects of autonomic blockade:

One of the patients was excluded from this analysis because the drugs caused an unexpectedly large change in mean systolic blood pressure, from 120 mmHg to 80 mmHg. In the other five subjects to whom drugs were administered, the change in blood pressure was much smaller, from $127 \pm 7 / 80 \pm 3$ mmHg (systolic/diastolic, mean \pm standard error of the mean) before drugs to $125 \pm 9 / 82 \pm 3$ mmHg after sympathetic blockade, then $124 \pm 12 / 81 \pm 4$ after parasympathetic blockade. Note that the duration of the effects of metoprolol tartrate is sufficiently long that sympathetic blockade can be assumed to have persisted during parasympathetic blockade. Mean values of ARI were found to not change significantly after the administration of drugs, with one exception in that ARIs in the right ventricle lengthened from 181 ± 5 ms to 185 ± 6 ms in the right ventricle after sympathetic blockade ($p < 0.05$ according to Welch's t-test).

Sympathetic blockade was generally found to reduce the magnitude of respiratory oscillations in ARI. In each subject, the effect was typically found to vary in magnitude between the different respiratory frequencies, but no clear pattern emerged across subjects in terms of a particular frequency being more sensitive to the intervention. As shown in Figure 85, the power spectral density of ARI decreased in all subjects after sympathetic blockade ($p < 0.01$). After subsequent parasympathetic blockade, the nominal magnitude of oscillations decreased or stayed the same for all but one subject, in whom the power spectral density increased substantially. No statistically significant difference was found between the 'sympathetic blockade' and 'sympathetic and parasympathetic blockade' conditions, but the latter condition remained significantly different




Figure 85: Averaged power spectral density (PSD) in ARI (left) and ABP (right) in the range 0.08 Hz – 0.6 Hz, at the site with the clearest oscillations, shown for five subjects and all three drug conditions. The solid black squares indicate the mean value for each condition. One subject (dashed line) differed from the others in that parasympathetic blockade led to an increase in PSD for ARI and BP. This subject was excluded from the calculation of the mean, but not from statistical comparisons, which showed significant differences between the drugged and undrugged states, as indicated by the given p-values.

from the control condition ($p < 0.05$), despite the variable effects of parasympathetic blockade. Changes in the power spectral density of ABP generally resembled those of ARI, although the statistical significance of the difference between the control and ‘sympathetic blockade’ conditions was reduced to $p < 0.05$.

Figure 86 presents the results of applying a wavelet transform to the respiratory signals and ARI series, as described in section V.C.2.e, page 171. The wavelet coefficients, presented as the vertical and colour axes of the 3-dimensional plots, indicate the magnitude of ARI oscillations at each frequency. At each point in time, the dominant components in ARI closely match those of the respiratory signal, but are attenuated with the administration of drugs. The magnitude and morphology of the respiratory signal do not change noticeably as can be seen in the middle trace on the back wall of panels B-D.

Oscillations during spontaneous cases of Central Sleep Apnea:

Levels of consciousness varied significantly among the five subjects who presented CSA. In the most severe case, the subject had recently been awoken from deep but unседated sleep, and in two other cases the subject was fully awake with no drugs administered. In one case, the subject

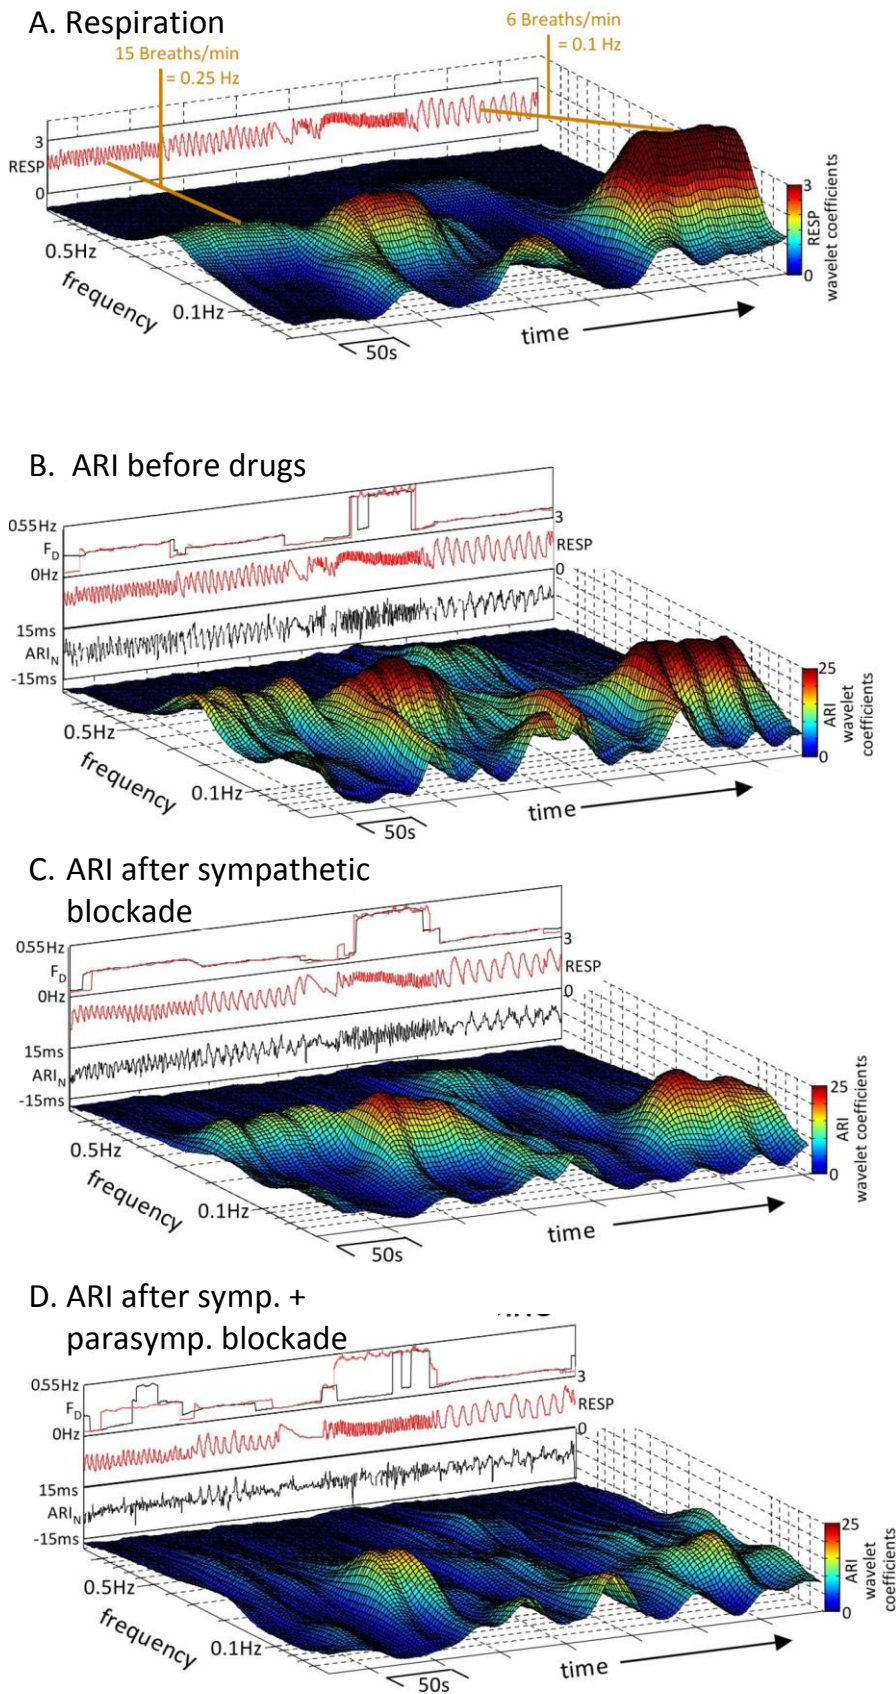


Figure 86: Time-frequency representation of ARI oscillations, using the wavelet transform. *Panel A* – Wavelet transform applied to the respiratory signal. Clear ridges coincide with the different respiratory frequencies (15 BrPM = 0.15 Hz, 30 BrPM = 0.5 Hz, 6 BrPM = 0.1 Hz). The respiratory to which the transform was applied is presented on the back wall. *Panel B* – Wavelet transform applied to an ARI series, with clear ridges occurring in the same places as in panel A. On the back wall, F_D is the dominant frequency in the ARI series, calculated as the frequency with the largest wavelet coefficient at each point in time (F_D for the RESP signal is calculated from the wavelet transform in panel A.). This trace confirms that the oscillations in ARI track the respiratory frequency closely. The RESP signal is presented again on the back wall for reference, along with the ARI series to which the wavelet transform was applied. *Panel C* – Same as panel B, except that the ARI and RESP signals used are from after the administration of metoprolol to block sympathetic activity. From the colour map, it can be seen that the ridges corresponding to ARI oscillations are, on average, reduced in magnitude. *Panel D* – Same as panels B and C, but using RESP and ARI signals from after parasympathetic blockade. In this subject, parasympathetic blockade has caused further reduction in the magnitude of ARI oscillations, to the extent that the dominant frequency F_D for ARI no longer tracks the respiratory frequency consistently.




Figure 87: Oscillations in ARI and ABP observed during Cheyne-Stokes respiration, a form of CSA.

The sites at which ARIs were recorded are labelled as “RV” or “LV” (left or right ventricle) and numbered 1-10 in reference to the sequential position of the electrode along the catheter, 1 being the electrode at the catheter tip, near the apex of the ventricle. ARI and ABP oscillate at the respiratory frequency (~ 0.3 Hz here) and the slower CSA frequency (~ 0.015 Hz). The degree to which each of the two frequencies is expressed in ARI differs substantially between electrodes. Differences can also be observed in the phase of the slower CSA oscillation, e.g. between LV 10 and LV 1.

was conscious but drowsy due to the administration of drugs, and in one case the drugs rendered the subject unconscious.

In most cases, the frequency of respiration increased with increasing volume of respiration. In all cases, ABP was found to oscillate at the respiratory frequency, as well as at the lower CSA frequency (approximately 0.3 Hz and 0.015 Hz, respectively, in Figure 87). Clear oscillations in ARI were observed at the respiratory frequency in all 5 subjects and at the CSA frequency in 4 subjects. In each case there was substantial heterogeneity in the degree to which these oscillations were expressed, as exemplified in Figure 87. At some sites, only one frequency could be detected. Some sites would express the superposition of both oscillations and others would show no clear pattern.

The latency between onset of inspiration and the subsequent peak in systolic ABP gradually lengthened during the development of each hyperpnea and shortened after the peak of the hyperpnea, as seen in Figure 88. This pattern differs widely from the changes in respiratory interval (dashed line), indicating that changes in respiratory cycle length alone are not sufficient to explain the phase shift in ABP. This effect was observed in all subjects, as demonstrated in Figure 89. Welch’s t-test confirmed statistical significance for the difference between stages of hyperpnea (middle vs. start: $p < 0.02$; middle vs. end: $p < 0.01$).




Figure 88: During each hyperpnea, as respiratory volume increased an increase was also seen in the time delay between each onset of inspiration and the subsequent peak in systolic blood pressure. The two events are marked for each beat with white dots in the top two panels. In the bottom trace, the solid line shows the varying delay between the two events. The relationship cannot be explained by changes in the duration of each breath, as indicated by the dashed line in the bottom panel. On the left side of the figure, construction lines illustrate how the traces in the bottom panel are constructed from the timings of events in the respiration and ABP traces.

V.C.4. Discussion

V.C.4.a. The effect of respiratory rate on steady-state ARI

Given that the mean differences between steady-state ARI at 15 BrPM and at 6 and 30 BrPM are small (see Figure 81, page 172), these results are not deemed to be clinically important. Nevertheless, one plausible explanation for such an effect exists, considering that 6 BrPM and 30 BrPM are extremely slow and fast breathing rates, respectively. The increased attention and mental effort required to consciously follow these less natural rates could conceivably evoke increased sympathetic drive – and, thus, shorter action potential durations – compared with the more natural rate of 15 BrPM. Although the effect seen here is small, future studies of this kind should take into account the possibility that consciously following extreme breathing rates might alter ANS activity via a stress response.

V.C.4.b. Are the respiratory oscillations in ARI real or artefact?

In Table 13 it was shown that significant oscillations in ARI arose in all subjects. However, this does not necessarily prove the existence of oscillations in the underlying electrophysiological properties. One must consider the possibility that they are solely attributable to the influence of respiratory motion on electrogram morphology. Various mechanisms by which such effects might occur are considered in the following paragraphs.




Figure 89: The shifting temporal relationship between ABP and respiration, as introduced in Figure 88, was observed in all instances of CSA. From plots similar to Figure 88, representative values of the delay between onset of inspiration and peak systolic pressure were chosen for the beginning, middle, and end of each hyperpnea cycle. These values were then averaged across all cycles observed in a particular subject (min. 2 clear cycles, max. 10) to produce the values presented in this figure. T-tests indicated that the values for the start and end of each hyperpnea differed from the values in the middle with statistical significance (p-values shown on graph).

As described on page 108, any changes in electrogram morphology that are not caused by changes in the electrophysiological source distribution can be described as a change in the lead field within that source distribution. The lead field can be altered by changing the geometric relationship between the electrodes and the surrounding anatomy, or by changing the effective conductivity of different regions within that anatomy. Influences of both kinds undoubtedly occur, as listed below:

- As the chest and abdomen expand during inspiration, the distance between the heart and the reference electrode (in this case, on the skin surface at the abdomen) is likely to increase slightly.
- As the lungs fill with air, their effective conductivity decreases substantially^{[Malmivuo & Plonsey, 1995], [Plonsey, 1965]}, which can be assumed to alter the lead field in the surrounding region, including the heart.
- The stroke volume of the different chambers of the heart is known to vary with respiration (page 68) ^[Innes et al, 1993]. This can be assumed to yield a change in the magnitude of the Brody effect, described on page 52, another lead field alteration.

- The exploring electrodes may have moved with respect to the myocardial wall. Movement perpendicular to the wall is unlikely to have been significant, since good contact was maintained during the more substantial wall movements that occur with each contraction. Movement in the axial direction (i.e. with the catheter sliding in and out of the ventricle) cannot be ruled out. The positioning of the catheters was occasionally monitored using an X-ray video feed, and no such effect was noticed. Small axial movements may have gone unnoticed, but they would not be expected to regularly cause ARI oscillations of the magnitude observed. [Lessick et al, 2001] investigated the stability of a similar catheter design in the context of respiratory movements and cardiac contraction, although it should be noted that this work was funded by the catheter's manufacturers, presenting a possible conflict of interest. Lessick *et al* used a modified catheter with a retractable needle that could be used to anchor the tip to the myocardium. They tracked the movement of the catheter tip in porcine hearts, comparing trajectories in 3D-space throughout the cardiac cycle with and without anchoring. They concluded that catheter trajectories "taken several beats apart from one another are separated by not more than 2mm". In our study, differences in ARI between adjacent electrodes (spaced 2 mm – 5 mm apart) were typically smaller than the amplitude of oscillations in ARI, so electrode movement parallel to the surface of the myocardium is unlikely to have contributed substantially to the observed oscillations. The influence of any electrode movement perpendicular to the myocardial wall would be highly dependent on the transmural variation on repolarisation time, which could not be characterised with the catheter electrodes used.

Respiratory changes in the lead field, such as those described above, cannot be ruled out completely from inducing some respiratory oscillation in ARI. In fact clear evidence of their effects exists in the form of baseline wander, observable when comparing isoelectric segments of the unfiltered electrogram recordings. However, two additional observations imply that these mechanisms cannot be wholly responsible for the ARI oscillations. Firstly, in several cases oscillations at typical respiratory frequencies were found to persist during breath-holds, with no evidence of corresponding chest movement (see Figure 84, page 176). Although the waistband sensor used to measure abdominal movements provides an incomplete representation of respiratory motion, it is unlikely that significant motion would have occurred without any detectable effect on abdominal circumference. Secondly, autonomic blockade was found to have a substantial effect on the magnitude of ARI oscillations without noticeably changing respiratory movements. This implies that the oscillations are related to autonomic input to the heart.

A caveat on the interpretation of the autonomic blockade test results is that only five subjects could be included in this analysis. Hence, although the differences between the blockade and control conditions were found to be statistically significant, the statistical test is not robust in

this case. Furthermore, the observed effects do not constitute a thorough characterisation of the effects of autonomic blockade on respiratory oscillations in ARI, especially when considering that the effects of parasympathetic blockade were not consistent among all subjects.

In summary, it is not possible to determine with confidence whether any particular instance of respiratory oscillations in ARI is caused by electrophysiological changes, lead field changes, or both. However, it is highly likely that some oscillations in ARI are genuinely electrophysiological in origin. The existence of such respiratory modulation of ventricular electrophysiology, independent of heart rate, is a novel observation. This discovery is described in a paper recently submitted for publication in the journal *Frontiers in Cardiac Electrophysiology*.

V.C.4.c. Can the oscillations in ARI be taken as an accurate representation of oscillations in APD?

The respiratory oscillations in ARI have been found to be at least partly electrophysiological in origin. But can the underlying changes be truthfully interpreted as a change in action potential duration (APD), or is some other electrophysiological change responsible for the changes in electrogram morphology?

In many electrodes, the delay between the pacing stimulus and the time of local activation was found to vary with respiration. An example of this behaviour is shown in Figure 90. In this example, the morphology of the activation wave changes so that an extra positive deflection appears before the local activation down-stroke. As discussed on page 59, this indicates that the activation travelled towards this electrode parallel to the myocyte fibres at ‘end exhalation’ and across the fibres at ‘end inspiration’. It appears that the spatio-temporal distribution of activation is changing, so the change in activation time is not representative of behaviour local to this electrode. These variations in activation time were more commonly found in the left ventricle, presumably because the hearts were always paced from the right ventricle, so any changes in conduction velocity had a longer distance across which to influence activation times in the left ventricle.

In Figure 90, the variations in activation time are more substantial than the variations in recovery time, and the phase relationship between respiration and the measured ARI oscillations is opposite to the typical relationship described on page 173. This observation calls into question




Figure 90: Two unipolar electrograms from the same electrode in the left ventricle (LV 6) at different stages of the breathing cycle, aligned according to the pacing artefact. The measured times of local activation (A_s) and recovery (R_s) are marked on each trace with vertical lines. The delay between pacing stimulus and activation time was often found to vary with respiration.

whether cases such as this can be considered a reliable representation of changes in APD. As discussed in section II.E, the spatial gradients in membrane potential are typically smaller during repolarisation than during activation, yielding smaller membrane currents. However, the relatively slow development of repolarisation means that these small currents occur across a broad region of tissue at the same time. The morphology of the T-wave is therefore representative of a broader distribution of behaviour, and may be less sensitive than activation-wave morphology to localised changes in timing.

However, as shown in Figure 24, page 62, electrotonic effects are substantial in inducing synchronisation of repolarisation between neighbouring myocytes. In cases such as Figure 90 where the timing of local activation has been altered relative to neighbouring regions, electrotonic synchronisation might be expected to prevent the repolarisation time from changing in the same way. By this interpretation, the measured change in ARI is still representative of the local change in APD, but that change in APD is caused by electrotonus in combination with more remote changes in conduction velocity, rather than by direct autonomic modulation of local repolarisation behaviour.

Hence, when relating measurements of local behaviour to local autonomic input, it is preferable to use measurements from early activating sites, thereby minimising the confounding influence of changes in conduction velocity. It should be noted, however, that ARI measurements at these sites may also be influenced by changes in conduction velocity since these changes would be manifested to some extent in the remote component of the electrogram (as described on page 49). An additional caveat in relating measured changes in ARI to changes in local APD is that, even when the stabilised indices are used, the sensitivity of this relationship will be dependent on UEG morphology. Although the magnitude of oscillations does not necessarily reflect the magnitude of changes in action potential duration, it should be remembered that APD is only a useful description of electrophysiological behavior, rather than a fundamental electrophysiological property. Changes in ARI might occasionally reflect changes in action potential morphology that do not constitute a change in APD but do alter the cell's effective refractory period, such as when the action potential morphology becomes more triangular in shape.

V.C.4.d. Underlying mechanisms

Before the observed respiratory oscillations in ARI can be interpreted as evidence of autonomic input, one must consider whether other mechanisms may have induced respiratory oscillations in action potential morphology. The mechanical effects of respiration might influence myocyte membrane behaviour by the mechanism known as mechanoelectric feedback^{[Taggart & Sutton, 2011], [Kohl et al, 2011]}. The respiratory changes in ventricular stroke volume described by Innes *et al*^[Innes et al, 1993] imply varying levels of strain in the myocytes before contraction. This 'pre-load' is known to influence action potential morphology. Although this mechanism may be involved in some cases of

ARI oscillations, it cannot fully account for the phenomenon; between the two ventricles, the changes in ARI were found to be approximately in phase with each other (excluding the confounding cases in which activation time varied significantly), whereas the respiratory changes in pre-load act oppositely on the two ventricles. Furthermore, autonomic blockade using drugs was found to strongly influence the magnitude of the oscillations without a significant change in respiratory behaviour.

It is therefore likely that the observed respiratory modulation of action potential morphology is driven by autonomic input. In determining the precise nature of that input, one is faced with the two candidate mechanisms considered in the recent debate on the nature of respiratory sinus arrhythmia^{[Eckberg, 2009], [Karemaker, 2009]}: the baroreflex (feedback control) versus centrally mediated respiratory gating (feedforward control).

Baroreflex feedback does not offer a consistent explanation of the observed oscillations. The time required for a change in blood pressure to be transmitted from baroreceptors to the central nervous system, processed, and acted upon constitutes a delay of at least one heart beat and up to 2.5 seconds^[Julien, 2006]. This constant delay would cause the phase of ARI oscillations (as a fraction of the cycle length) to lag behind respiration increasingly at higher frequencies. No such frequency dependence could be discerned in the phase relationship (see, for example, Figure 83).

The results presented are more consistent with the alternative candidate mechanism, central gating of autonomic activity by respiratory networks. This gating may arise in any of three locations: the brainstem^[Spyer & Gilbey, 1988], the cerebral cortex (specific to consciously controlled breathing)^[Evans, 2010], and the cardiac plexus (a network of neurons around the outside of the heart)^[Armour & Hopkins, 1990]. As described on page 68, respiratory gating manifests as an increase in sympathetic activity and a decrease in parasympathetic activity during inspiration. Sympathetic stimulation increases the L-type Ca⁺⁺ current and the slow delayed rectifier K⁺ current, resulting in an increase in contractility and a decrease in action potential duration^{[Trautwein & Hescheler, 1990], [Volders et al, 2003]}. Parasympathetic activity is not generally thought to have a direct effect on the ventricles. Acetylcholine (the neurotransmitter associated with parasympathetic activity) has been shown to activate a particular class of K⁺ current which, contrary to our observations, would yield longer action potentials during inspiration, when parasympathetic activity is decreased^[Yang et al, 1996]. However, the indirect effects of parasympathetic stimulation may well be more important in the ventricles, inhibiting sympathetic input and its effects^[Stramba-Badiale et al, 1991]. By this mechanism, the inspiratory decrease in parasympathetic activity would enhance the effects of the increase in sympathetic activity. The observation that the effects of parasympathetic blockade varied between subjects might be explained by the existence of these two mechanisms that act in opposite senses. Furthermore, the observation that changes induced by parasympathetic blockade were smaller than those induced by sympathetic blockade (Figure 85, page 177) is consistent with the notion that

parasympathetic nerves impose a variable gain on the effects of sympathetic activity, which had already been reduced when parasympathetic blockade was induced.

An important caveat for the physiological interpretation of the results presented is that the autonomic outflow generated by the central nervous system during consciously controlled breathing may differ substantially from that produced during spontaneous breathing^[Pagani et al, 1986]. This possibility is currently being explored in follow-up experiments in which the same measurements as used in the original experiments are recorded during spontaneous breathing, then the recorded breathing pattern is presented on-screen for patients to consciously reproduce.

V.C.4.e. Implications for cardiac stability

Respiratory Sinus Arrhythmia causes heart rate to increase during inspiration and decrease during exhalation. APD is known to change in direct response to changes in heart rate, as evidenced by the shortening of APD that occurs as soon as a heart is paced (Figure 63, page 132). However, this mechanism acting alone would cause changes in APD to lag behind heart rate, leaving the heart unstable during fast breathing or sharp breaths. The observed heart-rate-independent respiratory oscillations in APD (ARI) are therefore likely to improve cardiac stability by ensuring that APD shortens in close synchrony with the shortening cardiac cycle, thereby reducing the chance that the next activation wave will be disrupted by refractory tissue.

While heart-rate-independent control may improve cardiac stability on the whole, it is worth noting that there were substantial heterogeneities in the magnitude of these oscillations. Heterogeneous patterns of repolarisation are believed to increase the likelihood of arrhythmic events^{[Noble & Rudy, 2001], [Hunter et al, 2001], [Clayton & Taggart, 2005]}. Heterogeneous respiratory oscillations in APD might arise from a non-uniform distribution of nerve endings in the myocardium. However, it is also possible that the regional differences in ARI oscillations were caused only by the fact that the sensitivity of the ARI measurement to changes in APD is highly dependent on the overall electrogram morphology, which varied between electrodes.

Heart-rate-independent modulation of APD may also detract from cardiac stability when considering occasional extreme autonomic fluctuations. Evolution may have fine-tuned this control system to closely track heart rate variations at the magnitude and frequencies experienced in everyday life, but at the expense of being less well suited to occasional extreme inputs, such as those induced by an extreme threat/stressor. The fact that APD and heart-rate can be controlled independent of one another provides a mechanism by which the two can become mismatched in unusual scenarios.

These unusual scenarios would not necessarily involve a sudden, extreme stimulus. Cases of abnormal respiratory behaviour may also fall into this category. In section V.C.3 it was shown that, in cases of Central Sleep Apnea (CSA) during cardiac pacing, an apparent desynchronisation occurs in the phase relationship between respiration and arterial blood pressure (ABP) oscillations.

The observation of such changes during CSA demands attention when considering the association between CSA and sudden cardiac death. A similar phenomenon, described as ‘bidirectional phase walk’, was recently observed in another system of cardio-respiratory autonomic interactions, between heart-rate and cardiac sympathetic nerve (CSN) discharge during respiratory effort [Gebber et al, 2005]. Unidirectional phase-walk is conventionally interpreted as evidence of weak coupling between two oscillators. In [Gebber et al, 2005], the bidirectional nature of the changes in phase relationship was taken as evidence of strong coupling of the CSN discharge to a central oscillator whose phase is modulated by respiratory inputs. While this explanation is plausible for our system, we also consider another interpretation: it may be that the apparent change in phase relationship is caused by a shift in influence between two or more follower oscillators, all strongly coupled to the same forcing oscillation but at different phases. This concept can be explained by considering a simple sinusoidal model in which multiple oscillators are driven at the same frequency but with different amplitudes and phases. The observed output y is taken as the superposition of these N oscillations.

$$y = \sum_{n=1}^N A_n \sin(\omega t + \theta_n) \quad (45)$$

The observed output can be expressed as a single sinusoid at the same frequency:

$$y = B \sin(\omega t + \phi) \quad (46)$$

with the magnitude and phase as described below.

$$B = \left(\left(\sum_{n=1}^N A_n \cos \theta_n \right)^2 + \left(\sum_{n=1}^N A_n \sin \theta_n \right)^2 \right)^{1/2} \quad (47)$$

$$\phi = \arctan \left(\frac{\sum_{n=1}^N A_n \sin \theta_n}{\sum_{n=1}^N A_n \cos \theta_n} \right) \quad (48)$$

Note that the phase of the output, ϕ , can be modulated by changing the amplitude of any of the driven oscillators, rather than just their phase. Respiratory oscillations in blood pressure are known to be driven by various mechanisms (cardiac, autonomic, mechanical), with different phase relationships between them. Although the oscillations involved are non-sinusoidal and do not combine as a simple linear sum, this model is useful in illustrating a basic concept: that a shifting phase relationship between the input and output of a system involving multiple parallel branches can be explained by a simple change in the gain of any of those branches. Each hyperpnea is defined by a crescendo in respiratory effort, and respiratory gating is known to modulate the gain of the baroreflex. It is therefore plausible that the observed change in the phase relationship is caused by the fact that the level of respiratory gating varies during hyperpnea.

The fact that this desynchronisation has not been previously described suggests that it may only be noticeable when pacing is used to control heart rate, since heart rate variability strongly influences ABP oscillations [Tan & Taylor, 2010]. However, the physiological insight yielded by this technique may still be an important one. Given that the role of respiratory sinus arrhythmia – whether it be harmful, benevolent or neutral – is not currently well understood, it is difficult to judge the clinical implications of any shifting phase relationship. However, given that ABP provides a feedback loop to multiple aspects of the cardiovascular control system (including ventricular APD, it seems), the desynchronisation between ABP and respiration complicates the range of input scenarios that may be presented to cardiac electrophysiology during CSA. The resultant behaviour therefore becomes less predictable, possibly increasing the chances that multiple inputs will briefly combine to disrupt the normal progression of cardiac electrical excitation and give rise to a fatal arrhythmia.

V.D. Summary of findings

This chapter has described experiments exploring several different approaches to the inference of autonomic cardiac control mechanisms from the analysis of electrical recordings.

Attempts to measure asymmetric nervous input to the heart:

Section V.A covered the use of beat-to-beat ECG measurements such as RR intervals and QT_C intervals in an attempt to test the Brain-Heart Laterality hypothesis. The results revealed no significant difference in cardiac function according to whether aversive images were presented in the left or right visual field, but close inspection of the results suggested that this could be due to the weakness of the stimulus rather than the absence of the hypothesised mechanisms. Hence the experiment was inconclusive.

Nevertheless, it provided useful insight as to whether the variable QT_C can be used to discern heart-rate independent modulation of repolarisation properties. This variable has long been used to normalise the influence of heart rate, allowing more reliable diagnosis of repolarisation abnormalities. In our experiments, it did not prove useful in discerning subtle, emotionally driven changes in autonomic control. This may be due to the absence of any such changes, but the results also point to an error in the choice of training data-set used to identify the QT -RR relationship. The fact that the experiment section in which QT_C showed the largest common response was the section after habituation took effect indicates that the most relevant variations in QT and RR dominated the fitted QT -RR relationship, and were thus eliminated from QT_C . A better approach would be to extract the training data-set ECG recording during various changes in posture and exertion. The fitted QT -RR relationship would then reflect the range of values associated with maintaining

homeostasis during physical challenges, and the variable QT_C might be more sensitive to psychological stimuli.

Measuring the HRV response to habituation-prone stimuli:

Section V.B discussed the use of heart-rate variability (HRV) to monitor autonomic cardiac control in psychological experiments that necessitated recordings shorter than the recommended duration in order to avoid habituation. In particular, we explored whether averaging spectra calculated from short sections, recorded under similar conditions, would yield spectral measurements sufficiently stable to distinguish the effects of psychological differences or pharmaceutical interventions. This approach was found to yield spectra with broad, ill-defined peaks at low frequencies, whose energy would leak into adjacent frequencies. Hence the measurement LF was generally unstable. It is therefore unsurprising that this particular measure of autonomic control did not yield a distinction between groups in any of the experiments.

In contrast, HF exposed a difference between subjects with high and low levels of neuroticism, and also exposed an effect of the drug SSRI on autonomic control in high-neurotic subjects. It was also shown that LF/HF , LF_n , and especially HF_n were sensitive to these experiments as indices of autonomic balance. Based on the number of statistically significant effects exposed, HF_n appeared to be particularly robust to the low-frequency corruption caused by averaging spectra from short signals. This performance can be explained as a consequence of the variable's increased sensitivity to cases in which low-frequency leakage is small, hence it should be noted that the observed effects may have been concentrated within that subset of patients. Collectively, these results indicate that, with the 'spectral averaging' technique, useful measurements of HRV can indeed be achieved in psychological experiments subject to rapid habituation.

Novel insights into cardiac autonomic control, inferred from the dynamics of UEG measurements:

Section V.C presented results from experiments in which the algorithms described in section IV were used for automatic measurement of activation-recovery intervals (ARIs) from unipolar electrograms. In particular, these experiments focussed on the physiological mechanisms linking respiratory behaviour with cardiac electrophysiology. It was shown that steady-state (median) values of ARI are not sensitive to respiratory rate, but the value of ARI often oscillates with respiration. It was not usually possible to precisely characterise the local electrophysiological changes underlying observed ARI oscillations, due to the many forms of artefact associated with respiratory movements and the variable sensitivity of the ARI measurement to changes in APD. Still, the use of pharmaceutical interventions and breath-hold tests indicated that these oscillations could not be wholly attributed to movement artefact. In general, our results suggest that ventricular APD is sensitive to the same respiratory gating of autonomic activity as is thought to contribute to respiratory sinus arrhythmia. This mechanism may have important consequences for cardiac stability. Although the precise nature of those consequences is currently unclear, it is hoped that dissemination of the discovery will provoke further investigation of the mechanisms involved.

An additional novel observation was the desynchronisation between blood pressure oscillations and respiration during episodes of Central Sleep Apnea (CSA). Given that different frequencies of ARI oscillation were found to manifest heterogeneously across the ventricles during CSA, the desynchronisation of an important feedback parameter could give rise to chaotic patterns of electrophysiological behaviour. A mechanism of this nature might help to explain the increased incidence of Sudden Cardiac Death among CSA patients.

VI. Conclusion

VI.A. Summary of achievements

This thesis presents various novel extensions to existing methods for inferring nervous input to the human heart using surface ECG recordings and invasive measurements of cardiac electrical activity. These methods were implemented in various experiments, yielding new insights into physiological and psychosomatic phenomena.

The identification of subject-specific relationships between QT intervals and RR intervals measured from ECG recordings was explored as a means of exposing any imbalance between the left and right branches of the autonomic nervous system due to psychological processes. Such an imbalance, the existence of which is proposed in the Brain-Heart Laterality hypothesis, could be expected to impose separate modulation of heart rate and ventricular repolarisation properties. These changes would be revealed in the variable QT_C , if calculated appropriately. Our experiments did not confirm the Brain-Heart Laterality hypothesis. However, the hypothesis could not be confidently refuted; the data cast doubts on the effectiveness of the psychological stimuli used, and it was noted that the performance of the QT_C variable is highly dependent on the training data used to establish the subject-specific QT-RR relationship. A separate contribution to the field of QT_C measurements was made in section III.B.2, where it was shown that a three-point cubic spline consistently models the QT-RR relationship more accurately than the wide range of biparametric models employed by previous authors.

Another group of experiments explored the possibility of using unusually short ECG recordings to characterise the autonomic response (in terms of heart-rate variability (HRV)) to psychological stimuli. This approach was based on the hypothesis that any loss of spectral accuracy would be justified by the enabled reduction of the effects of habituation. Habituation, being a relatively fast process, ordinarily prevents investigations of this kind. In our modified approach, spectra calculated from multiple short recordings, under similar conditions, were averaged together to provide a representation of the subject's heart-rate variability response to those conditions. The approach exposed a statistically significant distinction between psychological groups (panic sufferers versus control subjects, and high-neuroticism versus low-neuroticism) in terms of their response to an emotion-suppression task. It was also found that an antidepressant (SSRI) had a statistically significant effect on the response of high-neuroticism subjects to the task. Hence it seems that spectral averaging can indeed yield useful measurements of HRV, enabling further experiments into the cardiac autonomic effects of psychological stimuli. The performances of the various HRV indices in this specialised application were compared. In particular, it was noted that the normalised high-frequency index, HF_N , can be interpreted as a measure of autonomic balance that is more sensitive than LF/HF to stochastic HRV responses when very short signals are used.

A novel filtering technique was developed to reduce the artefact introduced to ECG recordings by the simultaneous use of fMRI brain-imaging. In a subset of the experiments mentioned above, this filtering made it possible to compare autonomic outflow to the heart with the

brain activity involved in regulating that outflow, confirming a difference in activity in a brain region associated with emotional processing and autonomic activation (right amygdala).

It is difficult to ascertain whether an observed task-dependent difference in autonomic activity would translate to detrimental or beneficial effects on an individual's general health. However, the confirmation of differences in psychosomatic response between psychologically distinct groups provides a mechanistic argument to support the numerous phenomenological studies that have linked psychological pathologies with poorer physical well-being.

Contributions to the field of invasive electrocardiographic recordings have also been explained. Algorithms were developed to allow automated detection of activation and recovery times from unipolar electrograms (UEGs) with only minimal human intervention. These algorithms enable the handling of large data-sets, such as those required to study heterogeneous electrophysiological dynamics. Because the conventional indices used to measure activation and recovery times from UEGs were found to be unsuitable for dynamic analyses, new definitions of these indices were developed and shown to provide a more meaningful representation of changes in electrophysiological properties. These stabilised measurements, combined with the automated detection algorithms, enable investigations into the mechanisms of cardiac control based on observed dynamics.

The stabilised activation and repolarisation indices were used to explore the involvement of respiratory behaviour in autonomic input to the heart. Heart-rate independent oscillations in ARI indicated direct modulation of ventricular action potential morphology by the autonomic nervous system. This novel observation extends the present understanding of respiratory involvement in autonomic control of the heart. It also constitutes a contribution to the persistent debate on the nature of respiratory sinus arrhythmia; the phase relationships between respiration and ARI were not found to be frequency-dependent, which suggests that respiratory modulation of ANS input to the heart is dominated by central mechanisms such as 'gating', rather than by the baroreflex response.

It was noted that, for UEG recordings, separating respiratory motion artefacts from related electrophysiological effects is difficult. Pharmaceutical manipulation of the ANS was used in this case to partially overcome the problem. Nevertheless, the developed measurements are best suited to measuring electrophysiological dynamics that are not accompanied by motion artefacts (or other lead field changes) with similar dynamics.

In a subset of patients, spontaneous incidences of Central Sleep Apnea were found to produce heterogeneous oscillations in ARI across the ventricles. Additionally, the phase relationship between respiration and blood pressure oscillations was found to shift during each episode of hyperpnea. This combination of ventricular electrophysiological (ARI) effects and the desynchronisation of an important feedback mechanism (ABP) were hypothesised to be pro-arrhythmic, potentially contributing to the previously observed link between CSA and Sudden Cardiac Death.

VI.B. Plans and suggestions for further work

The work described in this thesis might be extended in terms of further validation of the developed methods of analysis, new applications for those methods, and further investigation of the novel physiological insights described. The common purpose in all of the works described in this thesis was to improve the accessibility of *in vivo* human data on interactions between the heart and the autonomic nervous system. The validation of several novel tools for this purpose enables a ‘systems identification’ approach to this field, and it is hoped that this will lead to increased contributions from engineers experienced in the characterisation of unknown systems. An improved understanding of the mechanisms that lead to fatal arrhythmias could ultimately prevent thousands of unnecessary deaths in the UK alone each year. More specific suggestions for further work are detailed in the paragraphs below.

The benefits of the short-recording HRV technique, used to overcome the obstacle of habituation in psychological experiments, might be further confirmed by comparing measurements from spectra averaged from four 30-second stimulation periods with spectra from 2-minute continuous recordings. The spectral content is likely to be different due to the effects of habituation, preventing direct comparison. However, if the ‘spectral averaging’ approach is found to provide better discrimination between, for example, different psychological groups, this would confirm that the sacrifice in terms of the spectral accuracy achieved from shorter recordings is justified by the benefits of reducing the overall effect of habituation. The precise nature of this trade-off is expected to be highly dependent on the psychological task used.

The new method has already exposed differences between psychologically distinct groups in terms of their psychosomatic response to an emotion-suppression task. Hence the approach is of use to groups investigating other psychopathologies and stimuli. The clinical importance of the distinctions observed in this thesis are currently being further explored by our collaborators at the University of Oxford’s Department of Psychiatry, who are designing follow-up experiments to elucidate the nature of the psychosomatic differences.

As mentioned previously, the tools developed for UEG analysis have been employed by researchers at the University of Oxford, headed by Dr. Blanca Rodriguez, to investigate spatial heterogeneity of ARI rate-adaptation profiles. A team at the Heart Hospital in London, headed by Dr. Pier Lambiase, is developing a real-time implementation of the algorithms to allow *in situ* assessment of tissue properties during heart surgery. The observed behaviours will be compared with computational models of alternans in development at the University of Manchester, under Dr. Mark Boyett, to provide a better understanding of the mechanisms driving this phenomenon.

The stabilised indices A_S and R_S have been shown to provide more meaningful representations of electrophysiological dynamics. However, the sensitivity of the resulting ARI measurements to the underlying changes in action potential duration has not yet been thoroughly characterised. As described in the recommendations at the end of section IV, this can be achieved in future by a direct comparison with measurements from MAP recordings.

References

- Adenauer, H. et al., 2010. Is freezing an adaptive reaction to threat? Evidence from heart rate reactivity to emotional pictures in victims of war and torture. *Psychophysiology*, 47(2), pp.315-322.
- Akselrod, S. et al., 1981. Power spectrum analysis of heart rate fluctuation: a quantitative probe of beat-to-beat cardiovascular control. *Science*, 213, pp.220-222.
- Appelhans, B.M. & Luecken, L.J., 2006. Heart rate variability as an index of regulated emotional responding. *Review of General Psychology*, 10(3), pp.229-240.
- Anderson, K. et al., 1991. Determination of local myocardial electrical activation for activation sequence mapping. A statistical approach. *Circulation Research*, 69(4), pp.898-917.
- Appel, M.L. et al., 1989. Beat to beat variability in cardiovascular variables: Noise or music? *Journal of the American College of Cardiology*, 14(5), pp.1139-1148.
- Armour, J.A. & Hopkins, D.A., 1990. Activity of in vivo canine ventricular neurons. *American Journal of Physiology - Heart and Circulatory Physiology*, 258(2), pp.H326 -H336.
- Bain, L.J. & Engelhardt, M., 2000. *Introduction to Probability and Mathematical Statistics* 2nd ed., Pacific Grove: Duxbury Press.
- de Bakker, J.M.T. & Opthof, T., 2002. Is the apico-basal gradient larger than the transmural gradient? *Journal of Cardiovascular Pharmacology*, 39(3), pp.328-331.
- de Bakker, J.M. & Wittkampf, F.H., 2010. The pathophysiologic basis of fractionated and complex electrograms and the impact of recording techniques on their detection and interpretation. *Circulation: Arrhythmia and Electrophysiology*, 3(2), pp.204 -213.
- Barbieri, R. & Brown, E., 2006. Analysis of heartbeat dynamics by point process adaptive filtering. *Biomedical Engineering, IEEE Transactions on*, 53(1), pp.4-12.
- Barbieri, R. et al., 2005. A point-process model of human heartbeat intervals: new definitions of heart rate and heart rate variability. *American Journal of Physiology: Heart and Circulatory Physiology*, 288(1), pp.H424-435.
- Batchvarov, V.N. et al., 2002. QT-RR relationship in healthy subjects exhibits substantial intersubject variability and high intrasubject stability. *American Journal of Physiology - Heart and Circulatory Physiology*, 282(6), pp.H2356 -H2363.

- Bazett, H.C., 1920. An analysis of the time-relations of electrocardiograms. *Annals of Noninvasive Electrocardiology*, 2(2), pp.177-194.
- Bendat, J.S. & Piersol, A.G., 1993. *Engineering Applications of Correlation and Spectral Analysis*. 2nd ed. New York: Wiley-Blackwell.
- Bennett, D.H., 2002. *Cardiac Arrhythmias: Practical Notes on Interpretation and Treatment*. 6th ed. London: Hodder Arnold.
- Berger, R.D. et al., 1997. Beat-to-beat QT interval variability: novel evidence for repolarization lability in ischemic and nonischemic dilated cardiomyopathy. *Circulation*, 96(5), pp.1557-1565.
- Bernardi, L. et al., 1989. Respiratory sinus arrhythmia in the denervated human heart. *Journal of Applied Physiology*, 67(4), pp.1447-1455.
- Bhakta, D. & Miller, J.M., 2008. Principles of electroanatomic mapping. *Indian Pacing and Electrophysiology Journal*. 8(1), pp.32-50.
- Bigger, J.T. et al., 1995. RR variability in healthy, middle-aged persons compared with patients with chronic coronary heart disease or recent acute myocardial infarction. *Circulation*, 91(7), pp.1936-1943.
- Bolter, C.P., Wallace, D.J. & Hirst, G.D., 2001. Failure of Ba²⁺ and Cs⁺ to block the effects of vagal nerve stimulation in sinoatrial node cells of the guinea-pig heart. *Autonomic Neuroscience: Basic & Clinical*, 94(1-2), pp.93-101.
- Borst, C. & Karemaker, J.M., 1983. Time delays in the human baroreceptor reflex. *Journal of the Autonomic Nervous System*, 9(2-3), pp.399-409.
- Bowker, T. et al., 2003. Sudden, unexpected cardiac or unexplained death in England: a national survey. *QJM*, 96(4), pp.269-279.
- Box, G.E.P., Jenkins, G.M. & Reinsel, G.C., 2008. *Time Series Analysis: Forecasting and Control* 4th ed., New York: Wiley Blackwell.
- Bradley, M.M., Lang, P.J. & Cuthbert, B.N., 1993. Emotion, novelty, and the startle reflex: habituation in humans. *Behavioral Neuroscience*, 107(6), pp.970-980.
- Bueno-Orovio, A. et al., 2012. Slow adaptation of ventricular repolarization as a cause of arrhythmia? Abstract to be published in the proceedings of *The 7th International Workshop on Biosignal Interpretation*. Como, Italy, 2-4 July 2012.

- Chen, P. et al., 1991. Epicardial activation and repolarization patterns in patients with right ventricular hypertrophy. *Circulation*, 83(1), pp.104-118.
- Chen, Z., Brown, E. & Barbieri, R., 2009. Assessment of autonomic control and respiratory sinus arrhythmia using point process models of human heart beat dynamics. *IEEE Transactions on Biomedical Engineering*, 56(7), pp.1791-1802.
- Clayton, R. & Taggart, P., 2005. Regional differences in APD restitution can initiate wavebreak and re-entry in cardiac tissue: A computational study. *BioMedical Engineering OnLine*, 4(1), p.54.
- Clayton, R.H. & Holden, A.V., 2004. Propagation of normal beats and re-entry in a computational model of ventricular cardiac tissue with regional differences in action potential shape and duration. *Progress in Biophysics and Molecular Biology*, 85(2-3), pp.473-499.
- Clerc, L., 1976. Directional differences of impulse spread in trabecular muscle from mammalian heart. *The Journal of Physiology*, 255(2), pp.335 -346.
- Cohen, L.B. et al., 1974. Changes in axon fluorescence during activity: Molecular probes of membrane potential. *The Journal of Membrane Biology*, 19(1), pp.1-36.
- Cohen, M.A. & Taylor, J.A., 2002. Short-term cardiovascular oscillations in man: measuring and modelling the physiologies. *The Journal of Physiology*, 542(3), pp.669 -683.
- Colli Franzone, P. et al., 1998. Spread of excitation in 3-D models of the anisotropic cardiac tissue. III. Effects of ventricular geometry and fiber structure on the potential distribution. *Mathematical Biosciences*, 151(1), pp.51-98.
- Colli Franzone, P. et al., 2007a. Monophasic action potentials generated by bidomain modeling as a tool for detecting cardiac repolarization times. *American Journal of Physiology - Heart and Circulatory Physiology*, 293(5), pp.H2771 -H2785.
- Colli Franzone, P. et al., 2007b. Determining recovery times from transmembrane action potentials and unipolar electrograms in normal heart tissue. In *Functional Imaging and Modeling of the Heart*. pp. 139-149.
- Conrath, C.E. & Opthof, T., 2006. Ventricular repolarization: an overview of (patho)physiology, sympathetic effects and genetic aspects. *Progress in Biophysics and Molecular Biology*, 92(3), pp.269-307.
- Coronel, R. et al., 2006. Monophasic action potentials and activation recovery intervals as measures of ventricular action potential duration: Experimental evidence to resolve some controversies. *Heart Rhythm*, 3(9), pp.1043-1050.

- Critchley, H.D. et al., 2003. Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain*, 126(10), pp.2139 -2152.
- Critchley, H.D. et al., 2005. Mental stress and sudden cardiac death: asymmetric midbrain activity as a linking mechanism. *Brain*, 128(1), pp.75-85.
- Debbas, N.M. et al., 1999. Human atrial repolarization: effects of sinus rate, pacing and drugs on the surface electrocardiogram. *Journal of the American College of Cardiology*, 33(2), pp.358-365.
- deBoer, R.W., Karemaker, J.M. & Strackee, J., 1987. Hemodynamic fluctuations and baroreflex sensitivity in humans: a beat-to-beat model. *Am J Physiol Heart Circ Physiol*, 253(3), pp.H680-689.
- Deutsch, S. & Deutsch, A., 1993. *Understanding the Nervous System: An Engineering Perspective*, New York: Wiley-Blackwell.
- Di Simplicio, M. et al., 2010. Cardiovascular autonomic nervous system alterations during emotion regulation in subjects at risk for psychopathology. In proceedings of the *65th Annual Meeting of the Society of Biological Psychiatry*. New Orleans, USA 20-22 May 2010.
- Di Simplicio M. et al, 2011. Decreased heart rate variability during emotion regulation in subjects at risk for psychopathology. *Psychological Medicine*, Available through Cambridge Journals Online [Accessed 7 June 2012].
- Dimick, R.N. et al., 1987. Optimizing electrocardiograph electrode placement for cardiac-gated magnetic resonance imaging. *Investigative Radiology*, 22(1), pp.17-22.
- Draper, M.H. & Weidmann, S., 1951. Cardiac resting and action potentials recorded with an intracellular electrode. *The Journal of Physiology*, 115(1), pp.74-94.
- Durrer, D. et al., 1961. The electrocardiogram in normal and some abnormal conditions : In revived human fetal heart and in acute and chronic coronary occlusion. *American Heart Journal*, 61(3), pp.303-314.
- Eckberg, D.L., Nerhed, C. & Wallin, B.G., 1985. Respiratory modulation of muscle sympathetic and vagal cardiac outflow in man. *The Journal of Physiology*, 365(1), pp.181-196.
- Eckberg, D.L., 2003. The human respiratory gate. *The Journal of Physiology*, 548(2), pp.339 -352.
- Eckberg, D.L., 2009. Point:Counterpoint: Respiratory sinus arrhythmia is due to a central mechanism vs. respiratory sinus arrhythmia is due to the baroreflex mechanism. *J Appl Physiol*, 106(5), pp.1740-1742.

- Euler, D.E., 1999. Cardiac alternans: mechanisms and pathophysiological significance. *Cardiovascular Research*, 42(3), pp.583-590.
- Evans, K.C., 2010. Cortico-limbic circuitry and the airways: Insights from functional neuroimaging of respiratory afferents and efferents. *Biological psychology*, 84(1), pp.13-25.
- Eysenck, S.B.G. & Eysenck, H.J., 1975. *Manual of the EPQ (Eysenck Personality Questionnaire)*, London: University of London Press.
- Fleming, J., 1980. *Interpreting the Electrocardiogram* 1st ed., New York: Springer.
- Frank, E., 1956. An accurate, clinically practical system for spatial vectorcardiography. *Circulation*, 13(5), pp.737-749.
- Franz, M.R., 1983. Long-term recording of monophasic action potentials from human endocardium. *The American Journal of Cardiology*, 51(10), pp.1629-1634.
- Franz, M.R., 1991. Method and theory of monophasic action potential recording. *Progress in Cardiovascular Diseases*, 33(6), pp.347-368.
- Franz, M.R., 1999. Current status of monophasic action potential recording: theories, measurements and interpretations. *Cardiovascular Research*, 41(1), pp.25 -40.
- Franz, M.R., 2005. What is a monophasic action potential recorded by the Franz contact electrode? *Cardiovascular Research*, 65(4), pp.940 -941.
- Franz, M.R. et al., 1986. In vitro validation of a new cardiac catheter technique for recording monophasic action potentials. *European Heart Journal*, 7(1), pp.34 -41.
- Franz, M.R. et al., 1987. Monophasic action potential mapping in human subjects with normal electrocardiograms: direct evidence for the genesis of the T wave. *Circulation*, 75(2), pp.379-386.
- Franz, M.R. et al., 1988. Cycle length dependence of human action potential duration in vivo. Effects of single extrastimuli, sudden sustained rate acceleration and deceleration, and different steady-state frequencies. *Journal of Clinical Investigation*, 82(3), pp.972-979.
- Franzone, P.C., Guerri, L. & Tentoni, S., 1990. Mathematical modeling of the excitation process in myocardial tissue: influence of fiber rotation on wavefront propagation and potential field. *Mathematical Biosciences*, 101(2), pp.155-235.

- Fridericia, L.S., 1921. Die systolendauer im elektrokardiogramm bei normalen menschen und bei herzkranken. *Acta Medica Scandinavica*, 54(1), pp.17-50.
- Gebber, G.L., Das, M. & Barman, S.M., 2005. Dynamic changes in baroreceptor-sympathetic coupling during the respiratory cycle. *Brain Research*, 1046(1-2), pp.216-223.
- Geddes, L.A. & Baker, L.E., 1967. The specific resistance of biological material—A compendium of data for the biomedical engineer and physiologist. *Medical & Biological Engineering*, 5(3), pp.271-293.
- Gehi, A.K. et al., 2005. Microvolt t-wave alternans for the risk stratification of ventricular tachyarrhythmic events: a meta-analysis. *Journal of the American College of Cardiology*, 46(1), pp.75-82.
- Gelzer, A.R. et al., 2008. Dynamic mechanism for initiation of ventricular fibrillation in vivo. *Circulation*, 118(11), pp.1123-1129.
- Gepstein, L., Hayam, G. & Ben-Haim, S.A., 1997. A novel method for nonfluoroscopic catheter-based electroanatomical mapping of the heart. In vitro and in vivo accuracy results. *Circulation*, 95(6), pp.1611-1622.
- Gepstein, L., Hayam, G. & Ben-Haim, S.A., 1997. Activation-repolarization coupling in the normal swine endocardium. *Circulation*, 96(11), pp.4036-4043.
- Geselowitz, D., 1989. On the theory of the electrocardiogram. *Proceedings of the IEEE*, 77(6), pp.857-876.
- Geselowitz, D.B., 1992. Description of cardiac sources in anisotropic cardiac muscle : Application of bidomain model. *Journal of Electrocardiology*, 25(Supplement 1), pp.65-67.
- Geselowitz, D.B., 1967. On bioelectric potentials in an inhomogeneous volume conductor. *Biophysical Journal*, 7(1), pp.1-11.
- Geselowitz, D.B. & Miller, W.T., 1983. A bidomain model for anisotropic cardiac muscle. *Annals of Biomedical Engineering*, 11(3-4), pp.191-206.
- Goldman, L., 1995. Sodium channel inactivation from closed states: evidence for an intrinsic voltage dependency. *Biophysical Journal*, 69(6), pp.2369-2377.
- Greenbaum, R.A. et al., 1981. Left ventricular fibre architecture in man. *Heart*, 45(3), pp.248-263.

- Gross, J.J. & John, O.P., 2003. Individual differences in two emotion regulation processes: Implications for affect, relationships, and well-being. *Journal of Personality and Social Psychology*, 85(2), pp.348-362.
- Hanly, P. & Zuberi-Khokhar, N., 1996. Increased mortality associated with Cheyne-Stokes respiration in patients with congestive heart failure. *Am. J. Respir. Crit. Care Med.*, 153(1), pp.272-276.
- Hanson, B. et al., 2009. Interaction of activation–repolarization coupling and restitution properties in human. *Circulation: Arrhythmia and Electrophysiology*, 2(2), pp.162 -170.
- Haws, C. & Lux, R., 1990. Correlation between in vivo transmembrane action potential durations and activation-recovery intervals from electrograms. Effects of interventions that alter repolarization time. *Circulation*, 81(1), pp.281-288.
- He, Y.H. et al., 1999. An interactive graphical system for automated mapping and display of cardiac rhythms. *Journal of Electrocardiology*, 32(3), pp.225-241.
- Hemingway, H., Malik, M. & Marmot, M., 2001. Social and psychosocial influences on sudden cardiac death, ventricular arrhythmia and cardiac autonomic function. *European Heart Journal*, 22(13), pp.1082-1101.
- Henriquez, C.S., 1993. Simulating the electrical behavior of cardiac tissue using the bidomain model. *Critical Reviews in Biomedical Engineering*, 21(1), pp.1-77.
- Hodes, R.L., Cook, E.W. & Lang, P.J., 1985. Individual differences in autonomic response: conditioned association or conditioned fear? *Psychophysiology*, 22(5), pp.545-560.
- Hodgkin, A.L. & Huxley, A.F., 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), pp.500-544.
- Hunter, P.J., Kohl, P. & Noble, D., 2001. Integrative models of the heart: achievements and limitations. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 359(1783), pp.1049 -1054.
- Hutter, O.F. & Trautwein, W., 1956. Vagal and sympathetic effects on the pacemaker fibers in the sinus venosus of the heart. *The Journal of General Physiology*, 39(5), pp.715-733.
- Innes, J.A. et al., 1993. Within-breath modulation of left ventricular function during normal breathing and positive-pressure ventilation in man. *The Journal of Physiology*, 460, pp.487-502.

- Ino, T. et al., 1988. Relation of monophasic action potential recorded with contact electrode to underlying transmembrane action potential properties in isolated cardiac tissues: a systematic microelectrode validation study. *Cardiovascular Research*, 22(4), pp.255 -264.
- Irnich, W., 1985. Intracardiac electrograms and sensing test signals: electrophysiological, physical, and technical considerations. *Pacing and Clinical Electrophysiology*, 8(6), pp.870-888.
- Jackson, J.D., 1998. *Classical Electrodynamics* 3rd ed., Hoboken: John Wiley & Sons.
- Jalife, J., 2000. Ventricular Fibrillation: Mechanisms of Initiation and Maintenance. *Annual Review of Physiology*, 62(1), pp.25-50.
- Javaheri, S., 2006. Sleep disorders in systolic heart failure: a prospective study of 100 male patients. The final report. *International Journal of Cardiology*, 106(1), pp.21-28.
- Javaheri, S. et al., 2007. Central sleep apnea, right ventricular dysfunction, and low diastolic blood pressure are predictors of mortality in systolic heart failure. *Journal of the American College of Cardiology*, 49(20), pp.2028-2034.
- Josephson, M.E., 2008. *Clinical cardiac electrophysiology: techniques and interpretations*, Philadelphia: Lippincott Williams & Wilkins.
- Julien, C., 2006. The enigma of Mayer waves: facts and models. *Cardiovascular Research*, 70(1), pp. 12-21.
- Kadish, A., 2004. What is a monophasic action potential? *Cardiovascular Research*, 63(4), pp.580 -581.
- Kamath, M.V. & Fallen, E.L., 1993. Power spectral analysis of heart rate variability: a noninvasive signature of cardiac autonomic function. *Critical Reviews in Biomedical Engineering*, 21(3), pp.245-311.
- Karemaker, J.M., 2009. Counterpoint: Respiratory sinus arrhythmia is due to the baroreflex mechanism. *Journal of Applied Physiology*, 106(5), pp.1742-1743.
- Katona, P.G. et al., 1970. Cardiac vagal efferent activity and heart period in the carotid sinus reflex. *The American Journal of Physiology*, 218(4), pp.1030-1037.
- Katona, P.G. & Jih, F., 1975. Respiratory sinus arrhythmia: noninvasive measure of parasympathetic cardiac control. *Journal of Applied Physiology*, 39(5), pp.801-805.
- Kautzner, J., 2002. QT Interval Measurements. *Cardiac Electrophysiology Review*, 6(3), pp.273-277.

-
- Kay, S.M., 1999. *Modern Spectral Estimation: Theory and Application*, Upper Saddle River: Prentice Hall PTR.
- Kimber, S. et al., 1996. A comparison of unipolar and bipolar electrodes during cardiac mapping studies. *Pacing and Clinical Electrophysiology*, 19(8), pp.1196-1204.
- Klemm, H.U. et al., 2007. Catheter based simultaneous mapping of cardiac activation and motion: a review. , 7(3), pp.148-159.
- Koch, H., 2004. Recent advances in magnetocardiography. *Journal of Electrocardiology*, 37, Supplement(0), pp.117-122.
- Kohl, P., Bollensdorff, C. & Garny, A., 2006. Effects of mechanosensitive ion channels on ventricular electrophysiology: experimental and theoretical models. *Exp Physiol*, 91(2), pp.307-321.
- Kohl, P., Sachs, F. & Franz, M.R., 2011. *Cardiac Mechano-Electric Coupling and Arrhythmias* 2nd ed., Oxford: Oxford University Press.
- Köhler, B., Hennig, C. & Orglmeister, R., 2002. The principles of software QRS detection. *IEEE Engineering in Medicine and Biology Magazine: The Quarterly Magazine of the Engineering in Medicine & Biology Society*, 21(1), pp.42-57.
- Kondo, M., Nesterenko, V. & Antzevitch, C., 2004. Cellular basis for the monophasic action potential. Which electrode is the recording electrode? *Cardiovascular Research*, 63(4), pp.635-644.
- Kowallik, P. & Meesmann, M., 1995. Independent autonomic modulation of the human sinus and av nodes: evidence from beat-to-beat measurements of pr and pp intervals during sleep. *Journal of Cardiovascular Electrophysiology*, 6(11), pp.993-1003.
- Konta, T. et al., 1990. Significance of discordant ST alternans in ventricular fibrillation. *Circulation*, 82(6), pp.2185-2189.
- Kupersmith, J., Krongrad, E. & Waldo, A.L., 1973. Conduction intervals and conduction velocity in the human cardiac conduction system. Studies during open-heart surgery. *Circulation*, 47(4), pp.776-785.
- Kusumoto, F.M., 1999. Unipolar recording in cardiac electrophysiologic studies. *Journal of Interventional Cardiac Electrophysiology*, 3(2), pp.121-127.

- Lane, R. & Schwartz, G., 1987. Induction of lateralized sympathetic input to the heart by the CNS during emotional arousal: a possible neurophysiologic trigger of sudden cardiac death. *Psychosomatic Medicine*, 49(3), pp.274 -284.
- Lane, R.D. & Jennings, J.R., 1996. Hemispheric asymmetry, autonomic asymmetry, and the problem of sudden cardiac death. *Brain Asymmetry*. Boston: MIT Press.
- Lanfranchi, P.A. et al., 1999. Prognostic value of nocturnal cheyne-stokes respiration in chronic heart failure. *Circulation*, 99(11), pp.1435-1440.
- Lang, P., Bradley, M. & Cuthbert, B., 1997. International affective picture system (IAPS): technical manual and affective ratings. Available at: <http://www.unifesp.br/dpsicobio/adap/instructions.pdf> [Accessed July 15, 2010].
- Larsen, P.D. et al., 2010. Respiratory sinus arrhythmia in conscious humans during spontaneous respiration. *Respiratory Physiology & Neurobiology*, 174(1-2), pp.111-118.
- LeGrice, I.J. et al., 1995. Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *The American Journal of Physiology*, 269(2 Pt 2), pp.H571-582.
- Lessick, J. et al., 2001. Assessment of NOGA catheter stability during the entire cardiac cycle by means of a special needle-tipped catheter. *Catheterization and Cardiovascular Interventions: Official Journal of the Society for Cardiac Angiography & Interventions*, 52(3), pp.400-406.
- Levick, J., 2003. *An Introduction to Cardiovascular Physiology* 4th ed., London: Hodder Arnold.
- Levy, M.N. & Martin, P.J., 1979. Neural control of the heart. In *Handbook of Physiology. Section 2: The Cardiovascular System, vol. 1, ed. by R. M. Berne*. Bethesda, MD: American Physiological Society, pp. 581-620.
- Levy, M.N. & Pappano, A.J., 2006. *Cardiovascular Physiology: Mosby Physiology Monograph Series* 9th ed., Maryland Heights: Mosby.
- Lewis, J. et al., 1987. Complex dynamics resulting from repeated stimulation of nonlinear oscillators at a fixed phase. *Physics Letters A*, 125(2-3), pp.119-122.
- Liu, D. & Antzelevitch, C., 1995. Characteristics of the delayed rectifier current (i_{kr} and i_{ks}) in canine ventricular epicardial, midmyocardial, and endocardial myocytes : a weaker i_{ks} contributes to the longer action potential of the M cell. *Circulation Research*, 76(3), pp.351-365.

- Louch, W.E., Ferrier, G.R. & Howlett, S.E., 2002. Changes in excitation-contraction coupling in an isolated ventricular myocyte model of cardiac stunning. *American Journal of Physiology - Heart and Circulatory Physiology*, 283(2), pp.H800 -H810.
- Luo, C.H. & Rudy, Y., 1991. A model of the ventricular cardiac action potential. Depolarization, repolarization, and their interaction. *Circulation Research*, 68(6), pp.1501-1526.
- Lyons, R.G., 2004. *Understanding Digital Signal Processing* 2nd ed., Hoboken: Prentice Hall.
- Malik, M. et al., 2002. Relation between QT and RR intervals is highly individual among healthy subjects: implications for heart rate correction of the QT interval. *Heart*, 87(3), pp.220-228.
- Malik, M. & Camm, A.J., 1993. Components of heart rate variability -- what they really mean and what we really measure. *The American Journal of Cardiology*, 72(11), pp.821-822.
- Malik, M. et al., 2008. Subject-specific profiles of QT/RR hysteresis. *Am J Physiol Heart Circ Physiol*, 295(6), pp.H2356-2363.
- Mallat, S., 1999. *A Wavelet Tour of Signal Processing, Second Edition* 2nd ed., Salt Lake City: Academic Press.
- Malliani, A. et al., 1991. Cardiovascular neural regulation explored in the frequency domain. *Circulation*, 84(2), pp.482-492.
- Malmivuo, J. & Plonsey, R., 1995. *Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields*, New York: Oxford University Press USA.
- Mangoni, M.E. & Nargeot, J., 2008. Genesis and Regulation of the Heart Automaticity. *Physiological Reviews*, 88(3), pp.919-982.
- Martinez, J. & Olmos, S., 2005. Methodological principles of T wave alternans analysis: a unified framework. *IEEE Transactions on Biomedical Engineering*, 52(4), pp.599-613.
- McFee, R. & Johnston, F.D., 1953. Electrocardiographic leads: I. Introduction. *Circulation*, 8(4), pp.554 -568.
- Messinger-Rapport, B.J. & Rudy, Y., 1990. Noninvasive recovery of epicardial potentials in a realistic heart- torso geometry. Normal sinus rhythm. *Circulation Research*, 66(4), pp.1023-1039.

- Millar, C., Kralios, F. & Lux, R., 1985. Correlation between refractory periods and activation-recovery intervals from electrograms: effects of rate and adrenergic interventions. *Circulation*, 72(6), pp.1372-1379.
- Miller, W. & Geselowitz, D., 1978. Simulation studies of the electrocardiogram. I. The normal heart. *Circulation Research*, 43(2), pp.301 -315.
- Montano, N. et al., 1994. Power spectrum analysis of heart rate variability to assess the changes in sympathovagal balance during graded orthostatic tilt. *Circulation*, 90(4), pp.1826-1831.
- Nakazawa, K. et al., 2003. Autonomic imbalance as a property of symptomatic Brugada syndrome. *Circulation Journal: Official Journal of the Japanese Circulation Society*, 67(6), pp.511-514.
- Napadow, V. et al., 2008. Brain correlates of autonomic modulation: Combining heart rate variability with fMRI. *NeuroImage*, 42(1), pp.169-177.
- Narayan, S.M. & Smith, J.M., 1999. Spectral analysis of periodic fluctuations in electrocardiographic repolarization. *IEEE Transactions on Bio-Medical Engineering*, 46(2), pp.203-212.
- Nash, M.P. et al., 2006. Whole heart action potential duration restitution properties in cardiac patients: a combined clinical and modelling study. *Experimental Physiology*, 91(2), pp.339-354.
- Ndrepepa, G. et al., 1995. Activation time determination by high-resolution unipolar and bipolar extracellular electrograms in the canine heart. *Journal of Cardiovascular Electrophysiology*, 6(3), pp.174-188.
- Nearing, B.D. & Verrier, R.L., 2002. Progressive increases in complexity of T-wave oscillations herald ischemia-induced ventricular fibrillation. *Circulation Research*, 91(8), pp.727-732.
- Nesterenko, V.V., Kondo, M. & Antzelevitch, C., 2005. Biophysical basis for monophasic action potential. *Cardiovascular Research*, 65(4), pp.942 -944.
- NHS Choices, 2011. Heart rhythm problems - Live Well - NHS Choices. Available at: <http://www.nhs.uk/Livewell/Healthyhearts/Pages/Arrhythmias.aspx> [Accessed May 27, 2011].
- Nielsen, P.M. et al., 1991. Mathematical model of geometry and fibrous structure of the heart. *The American Journal of Physiology*, 260(4 Pt 2), pp.H1365-1378.
- Noble, A. et al., 2005. *The Cardiovascular System: Systems of the Body Series*, Philadelphia: Churchill Livingstone.

-
- Noble, D. & Rudy, Y., 2001. Models of cardiac ventricular action potentials: iterative interaction between experiment and simulation. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 359(1783), pp.1127 -1142.
- Okin, P.M. et al., 2000. Electrocardiographic quantitation of heterogeneity of ventricular repolarization. *Annals of Noninvasive Electrocardiology*, 5(1), pp.79-87.
- Pagani, M. et al., 1986. Power spectral analysis of heart rate and arterial pressure variabilities as a marker of sympatho-vagal interaction in man and conscious dog. *Circulation Research*, 59(2), pp.178-193.
- Palomba, D., Angrilli, A. & Mini, A., 1997. Visual evoked potentials, heart rate responses and memory to emotional pictorial stimuli. *International Journal of Psychophysiology*, 27(1), pp.55-67.
- Palomba, D. et al., 2000. Cardiac responses associated with affective processing of unpleasant film stimuli. *International Journal of Psychophysiology*, 36(1), pp.45-57.
- Parati, G. et al., 1995. Spectral analysis of blood pressure and heart rate variability in evaluating cardiovascular regulation. A critical appraisal. *Hypertension*, 25(6), pp.1276-86.
- Pastore, J.M. et al., 1999. Mechanism Linking T-Wave Alternans to the Genesis of Cardiac Fibrillation. *Circulation*, 99(10), pp.1385-1394.
- Pauza, D.H. et al., 2000. Morphology, distribution, and variability of the epicardiac neural ganglionated subplexuses in the human heart. *The Anatomical Record*, 259(4), pp.353-382.
- Phan, K.L. et al., 2005. Neural substrates for voluntary suppression of negative affect: A functional magnetic resonance imaging study. *Biological Psychiatry*, 57(3), pp.210-219.
- Philipson, K.D., 1997. Myocardial Ion Transporters. In *The Myocardium*. Salt Lake City: Academic Press
- Pichot, V. et al., 1999. Wavelet transform to quantify heart rate variability and to assess its instantaneous changes. *J Appl Physiol*, 86(3), pp.1081-1091.
- Plank, G. et al., 2008. From mitochondrial ion channels to arrhythmias in the heart: computational techniques to bridge the spatio-temporal scales. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1879), pp.3381 -3409.
- Plonsey, R., 1969. *Bioelectric Phenomena*, New York: McGraw-Hill Education.

- Plonsey, R., 1963. Current Dipole Images and Reference Potentials. *Bio-medical Electronics, IEEE Transactions on*, 10(1), pp.3-8.
- Plonsey, R. & Barr, R.C., 2007. *Bioelectricity: A Quantitative Approach* 3rd ed., New York: Springer.
- Potse, M. et al., 2009. Validation of a simple model for the morphology of the T wave in unipolar electrograms. *American Journal of Physiology - Heart and Circulatory Physiology*, pp. H792-H801.
- Pueyo, E. et al., 2004. Characterization of QT interval adaptation to RR interval changes and its use as a risk-stratifier of arrhythmic mortality in amiodarone-treated survivors of acute myocardial infarction. *IEEE Transactions on Biomedical Engineering*, 51(9), pp.1511-1520.
- Rankin, C.H. et al., 2009. Habituation revisited: an updated and revised description of the behavioral characteristics of habituation. *Neurobiology of learning and memory*, 92(2), pp.135-138.
- Rimoldi, O. et al., 1990. Analysis of short-term oscillations of R-R and arterial pressure in conscious dogs. *American Journal of Physiology - Heart and Circulatory Physiology*, 258(4), pp.H967-976.
- Roberts, D., Hersh, L. & Scher, A., 1979. Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog. *Circulation Research*, 44(5), pp.701 -712.
- Roberts, D. & Scher, A., 1982. Effect of tissue anisotropy on extracellular potential fields in canine myocardium in situ. *Circulation Research*, 50(3), pp.342 -351.
- Roden, D.M. et al., 2002. Cardiac ion channels. *Annual Review of Physiology*, 64(1), pp.431-475.
- Rosenbaum, D.S. et al., 1994. Electrical alternans and vulnerability to ventricular arrhythmias. *New England Journal of Medicine*, 330(4), pp.235-241.
- Sagie, A. et al., 1992. An improved method for adjusting the QT interval for heart rate (the Framingham Heart Study). *The American Journal of Cardiology*, 70(7), pp.797-801.
- Salama, G., Lombardi, R. & Elson, J., 1987. Maps of optical action potentials and NADH fluorescence in intact working hearts. *American Journal of Physiology - Heart and Circulatory Physiology*, 252(2), pp.H384 -H394.
- Samson, W.E. & Scher, A.M., 1960. Mechanism of S-T segment alteration during acute myocardial injury. *Circulation Research*, 8(4), pp.780-787.

- Sánchez-Navarro, J.P., Martínez-Selva, J.M. & Román, F., 2006. Uncovering the relationship between defence and orienting in emotion: cardiac reactivity to unpleasant pictures. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 61(1), pp.34-46.
- Sarma, J.S. et al., 1984. An exponential formula for heart rate dependence of QT interval during exercise and cardiac pacing in humans: Reevaluation of Bazett's formula. *The American Journal of Cardiology*, 54(1), pp.103-108.
- Scacchi, S. et al., 2009. A reliability analysis of cardiac repolarization time markers. *Mathematical Biosciences*, 219(2), pp.113-128.
- Schmitt, F.O. & Erlanger, J., 1928. Directional differences in the conduction of the impulse through heart muscle and their possible relation to extrasystolic and fibrillary contractions. *American Journal of Physiology*, 87(2), pp.326-347.
- Scollan, D.F. et al., 2000. Reconstruction of cardiac ventricular geometry and fiber orientation using magnetic resonance imaging. *Annals of Biomedical Engineering*, 28(8), pp.934-944.
- Shaw, R.M. & Rudy, Y., 1997. Electrophysiologic effects of acute myocardial ischemia: a theoretical study of altered cell excitability and action potential duration. *Cardiovascular Research*, 35(2), pp.256-272.
- Shenasa, M., Hindricks, G. & Borggreffe, M., 2009. *Cardiac Mapping*, Hoboken: John Wiley & Sons.
- Sicouri, S. & Antzelevitch, C., 1991. A subpopulation of cells with unique electrophysiological properties in the deep subepicardium of the canine ventricle. The M cell. *Circulation Research*, 68(6), pp.1729-1741.
- Smith, J. et al., 1988. Electrical alternans and cardiac electrical instability. *Circulation*, 77(1), pp.110-121.
- Spach, M.S. et al., 1979. Extracellular potentials related to intracellular action potentials during impulse conduction in anisotropic canine cardiac muscle. *Circulation Research*, 45(2), pp.188-204.
- Spyer, K.M., 1994. Annual review prize lecture. Central nervous mechanisms contributing to cardiovascular control. *The Journal of Physiology*, 474(1), pp.1-19.
- Spyer, K.M. & Gilbey, M.P., 1988. Cardiorespiratory interactions in heart-rate control. *Annals of the New York Academy of Sciences*, 533, pp.350-357.

- Steinhaus, B., 1989. Estimating cardiac transmembrane activation and recovery times from unipolar and bipolar extracellular electrograms: a simulation study. *Circulation Research*, 64(3), pp.449-462.
- Stevenson, W.G. & Soejima, K., 2005. Recording techniques for clinical electrophysiology. *Journal of Cardiovascular Electrophysiology*, 16(9), pp.1017-1022.
- Stramba-Badiale, M. et al., 1991. Sympathetic-parasympathetic interaction and accentuated antagonism in conscious dogs. *American Journal of Physiology - Heart and Circulatory Physiology*, 260(2), pp.H335 -H340.
- Sutton, P.M.I. et al., 1991. Alternans of epicardial repolarization as a localized phenomenon in man. *European Heart Journal*, 12(1), pp.70-78.
- Taccardi, B. et al., 2005. Intramural activation and repolarization sequences in canine ventricles. Experimental and simulation studies. *Journal of Electrocardiology*, 38(4 Suppl), pp.131-137.
- Taccardi, B. et al., 1998. Multiple components in the unipolar electrogram: a simulation study in a three-dimensional model of ventricular myocardium. *Journal of Cardiovascular Electrophysiology*, 9(10), pp.1062-1084.
- Taggart, P. et al., 2000. Inhomogeneous transmural conduction during early ischaemia in patients with coronary artery disease. *Journal of Molecular and Cellular Cardiology*, 32(4), pp.621-630.
- Taggart, P., Critchley, H. & Lambiase, P.D., 2011. Heart-brain interactions in cardiac arrhythmia. *Heart*. Available at: <http://heart.bmj.com/content/early/2011/03/01/hrt.2010.209304.extract> [Accessed March 7, 2011].
- Taggart, P. & Sutton, P., 2011. Load dependence of ventricular repolarization. In *Cardiac Mechano-Electric Coupling and Arrhythmias*. Oxford: Oxford University Press.
- Takalo, R., Hytti, H. & Ihalainen, H., 2005. Tutorial on univariate autoregressive spectral analysis. *Journal of Clinical Monitoring and Computing*, 19(6), pp.401-410.
- Taylor, J.A. et al., 2001. Sympathetic restraint of respiratory sinus arrhythmia: implications for vagal-cardiac tone assessment in humans. *Am J Physiol Heart Circ Physiol*, 280(6), pp.H2804-2814.
- TFESCNASPE (Task Force of the European Society of Cardiology and the North American Society of Pacing Electrophysiology), 1996. Heart rate variability : standards of measurement, physiological interpretation, and clinical use. *Circulation*, 93(5), pp.1043-1065.

- Thayer, J.F. & Lane, R.D., 2000. A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders*, 61(3), pp.201-216.
- Thayer, J.F. & Brosschot, J.F., 2005. Psychosomatics and psychopathology: looking up and down from the brain. *Psychoneuroendocrinology*, 30(10), pp.1050-1058.
- Trautwein, W. & Hescheler, J., 1990. Regulation of Cardiac L-type calcium current by phosphorylation and G proteins. *Annual Review of Physiology*, 52(1), pp.257-274.
- Valens, C., 2004. A really friendly guide to wavelets. Available at: http://pagesperso-orange.fr/polyvalens/clemens/download/arfgtw_26022004.pdf [Accessed January 26, 2009].
- Vaseghi, S.V., 2000. *Advanced Digital Signal Processing and Noise Reduction, 2nd Edition* 2nd ed., Hoboken: John Wiley and Sons Ltd.
- Vigmond, E. et al., 2008. Solvers for the cardiac bidomain equations. *Progress in biophysics and molecular biology*, 96(1-3), pp.3-18.
- Vigmond, E.J., 2005. The electrophysiological basis of MAP recordings. *Cardiovascular Research*, 68(3), pp.502 -503.
- Visser, K.R., 1992. Electric conductivity of stationary and flowing human blood at low frequencies. *Medical & Biological Engineering & Computing*, 30(6), pp.636-640.
- Volders, P.G. et al., 2003. Probing the contribution of IKs to canine ventricular repolarization. *Circulation*, 107(21), pp.2753 -2760.
- Wei, D., 1997. Whole-heart modeling: Progress, principles and applications. *Progress in Biophysics and Molecular Biology*, 67(1), pp.17-51.
- Weiss, J.N., 1997. Ion Channels in Cardiac Muscle. In *The Myocardium*. Salt Lake City: Academic Press
- Weissenburger, J., Nesterenko, V.V. & Antzelevitch, C., 2000. Transmural heterogeneity of ventricular repolarization under baseline and long QT conditions in the canine heart in vivo: torsades de pointes develops with halothane but not pentobarbital anesthesia. *Journal of Cardiovascular Electrophysiology*, 11(3), pp.290-304.
- Western, D., Taggart, P. & Hanson, B., 2010. Real-time feedback of dynamic cardiac repolarization properties. Proceedings of *The 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 114-117, Buenos Aires, Argentina, 1-4 Sep 2010.

- Western, D. et al., 2011. Oscillatory patterns of respiration: consequences for the stability and control of cardiac electrophysiology. Proceedings of *The 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1949-1952, Boston, USA, 30 Aug-3 Sep 2011.
- Wilson, F.N. & Bayley, R.H., 1950. The electric field of an eccentric dipole in a homogeneous spherical conducting medium. *Circulation*, 1(1), pp.84-92.
- Winton, W.M., Putnam, L.E. & Krauss, R.M., 1984. Facial and autonomic manifestations of the dimensional structure of emotion. *Journal of Experimental Social Psychology*, 20(3), pp.195-216.
- Witkowski, F.X. & Penkoske, P.A., 1988. A completely automated activation-repolarization interval algorithm for directly coupled unipolar electrograms and its three-dimensional correlation with refractory periods. *Journal of Electrocardiology*, 21(3), pp.273-282.
- Woo, M.A. et al., 1992. Patterns of beat-to-beat heart rate variability in advanced heart failure. *American Heart Journal*, 123(3), pp.704-710.
- Wyatt, R.F., 1980. Comparison of estimates of activation and recovery times from bipolar and unipolar electrograms to in vivo transmembrane APDs. *Proc IEEE/Eng Med Biol Soc*, 2nd annual conference, Washington, DC, pp.22-25.
- Yan, G.X. & Antzelevitch, C., 1998. Cellular basis for the normal T wave and the electrocardiographic manifestations of the long-QT syndrome. *Circulation*, 98(18), pp.1928-1936.
- Yang, Z.K. et al., 1996. Regional differences in the negative inotropic effect of acetylcholine within the canine ventricle. *The Journal of Physiology*, 492(Pt 3), pp.789-806.
- Yue, A.M. et al., 2004. Determination of human ventricular repolarization by noncontact mapping: validation with monophasic action potential recordings. *Circulation*, 110(11), pp.1343-1350.
- Zhang, R. et al., 2002. Mechanism of blood pressure and R-R variability: insights from ganglion blockade in humans. *Journal of Physiology*, 543.1, pp.337 - 348.
- Zipes, D.P. & Jalife, J., 1999. *Cardiac Electrophysiology: From Cell to Bedside* 3rd ed., Philadelphia: Saunders (W.B.) Co Ltd.

Appendix I. Deriving an Expression for Extracellular Potential, in Terms of Transmembrane Potential, Using the Core-Conductor Model

By definition, the transmembrane potential v_m is

$$v_m = \phi_i - \phi_e \quad (49)$$

where ϕ_i and ϕ_e are the intracellular and extracellular potential. Examining the partial derivative with respect to x , position along the length of the myocyte, and applying Ohm's law, we achieve

$$\frac{\partial v_m}{\partial x} = \frac{\partial \phi_i}{\partial x} - \frac{\partial \phi_e}{\partial x} = -I_i \bar{R}_i + I_e \bar{R}_e \quad (50)$$

where I_i and I_e are the longitudinal currents and \bar{R}_i and \bar{R}_e are the resistance per unit length for the intracellular and extracellular domains. Note that this is the continuous version of the discrete model shown in Figure 17 (page 50). If we define the total axial current, I , as

$$I = I_i + I_e \quad (51)$$

then (50) can be written as

$$\frac{\partial v_m}{\partial x} = -I_i (\bar{R}_i + \bar{R}_e) + I \bar{R}_e. \quad (52)$$

By invoking Ohm's law again to replace $-I_i$, this becomes

$$\frac{\partial v_m}{\partial x} = \frac{\bar{R}_i + \bar{R}_e}{\bar{R}_i} \frac{\partial \phi_i}{\partial x} + I \bar{R}_e. \quad (53)$$

The defining characteristic of the core-conductor model is that all extracellular currents are assumed to be in the longitudinal (axial) direction. In the absence of any external current source, Kirchoff's law dictates that

$$I = I_i + I_e = 0. \quad (54)$$

Rearranging (53) then gives

$$\frac{\partial \phi_i}{\partial x} = \frac{\bar{R}_i}{\bar{R}_i + \bar{R}_e} \frac{\partial v_m}{\partial x}. \quad (55)$$

Integrating with respect to x and ignoring the integration constants as an arbitrary offset voltage, we have

$$\phi_i = \frac{\bar{R}_i}{\bar{R}_i + \bar{R}_e} v_m.$$

(56)

Using (49) we get

$$\phi_e = \frac{\bar{R}_i}{\bar{R}_i + \bar{R}_e} v_m - v_m$$

(57)

or

$$\phi_e = -\frac{\bar{R}_e}{\bar{R}_i + \bar{R}_e} v_m.$$

(58)

Appendix II. Deriving an Expression for the Unipolar Electrogram, Incorporating The Lead Field Concept

For completeness, the full derivation for (2) (page 51), taken from [Geselowitz, 1989] and [Geselowitz, 1992], is as follows. Green's theorem, also known as Green's second identity, can be written for any two scalar fields ψ and ϕ as^[Plonsey & Barr, 2007]

$$V \int_H (\phi \nabla^2 \psi - \psi \nabla^2 \phi) dV = \int_S (\phi \nabla \psi - \psi \nabla \phi) \cdot d\vec{S} \quad (59)$$

Let S be the surface of the body, so that V is the volume contained by the body, then choose ϕ as the extracellular potential field, ϕ_e , and let ψ be σZ , where σ is the effective conductivity at each point in the body. If ∇Z is the lead field for a particular electrode configuration, then Z is the potential field that results when a unit current is applied to that lead. If the body is assumed to be perfectly insulated, then the right-hand side of (59) is zero because no current will flow across the boundary. That is, $\nabla \phi \cdot d\vec{S} = 0$ and $\nabla \psi \cdot d\vec{S} = \sigma \nabla Z \cdot d\vec{S} = 0$. After rearranging, this leaves

$$\int Z \nabla \cdot \sigma \nabla \phi_e dV = \int \phi_e \nabla \cdot \sigma \nabla Z dV \quad (60)$$

Now, let \vec{J}_i be the impressed current density that results from membrane activity and let \vec{J} be the net current density. Applying Ohm's law,

$$\vec{J} = -\sigma \nabla \phi_e + \vec{J}_i \quad (61)$$

In the absence of any sources or sinks, we have

$$\nabla \cdot \vec{J} = 0 \quad (62)$$

Taking the gradient of (61), we can then write

$$\nabla \cdot \sigma \nabla \phi_e = \nabla \cdot \vec{J}_i \quad (63)$$

(63) can be substituted into the left side of (60). Now let p and q be the positions of the exploring and reference electrode. In the right side, Z can be assumed to belong to the class of functions known as Green functions,^[Jackson, 1998] in which case $\nabla^2 Z$ is 0 except for singularities at p and q , where its volume integral is $-1/\sigma$ and $1/\sigma$ respectively. Hence...

$$\int Z \nabla \cdot \vec{J}_i dV = -\phi_e(p) + \phi_e(q) \quad (64)$$

The right side of this equation is the negative of the UEG voltage. For later convenience, the left side can be expressed using \vec{J}_i instead of $\nabla \cdot \vec{J}_i$. This form is achieved as follows. From the divergence theorem (also known as Gauss's theorem) and the chain rule...

$$\int \nabla \cdot (\vec{J}_i Z) dV = \int \vec{J}_i Z \cdot d\vec{S} = \int (Z \nabla \cdot \vec{J}_i + \vec{J}_i \cdot \nabla Z) dV \quad (65)$$

Because \vec{J}_i is zero at the surface – that is, the surface wholly contains the myocardium – the middle expression is zero. Hence the terms on the right are equal and opposite.

$$-\int Z \nabla \cdot \vec{J}_i dV = \int \vec{J}_i \cdot \nabla Z dV \quad (66)$$

This is the expression that allows us to rewrite (64) as...

$$UEG = \int \vec{J}_i \cdot \nabla Z dV \quad (67)$$

Without changing the value of this expression, we can redefine that boundary as the outer surface of the heart, which still contains all non-zero values of \vec{J}_i .

The interpretability of this expression can be improved if we replace \vec{J}_i with a term involving the familiar transmembrane potential, V_m . As mentioned previously, a fundamental principle of the bidomain model is that the membrane current per-unit-volume, i_m , is the only exchange of current between the two domains. This principle can be expressed by the following two equations, which state that the membrane current per-unit-volume is equal to the change in current within each domain^[Miller & Geselowitz, 1978].

$$\begin{aligned} \nabla \cdot \sigma_i \nabla \phi_i &= i_m \\ \nabla \cdot \sigma_e \nabla \phi_e &= -i_m \end{aligned} \quad (68)$$

The impressed current density \vec{J}_i must be related to the source i_m by

$$\nabla \cdot \vec{J}_i = i_m \quad (69)$$

Comparing (68) and (69) we see that \vec{J}_i is also the intracellular current density, $\vec{J}_i = \sigma_i \nabla \phi_i$. \vec{J}_i can alternatively be interpreted as a current dipole moment per-unit-volume (units: amp-meters/meter³, rather than amps/meter²), an interpretation that is more readily reconciled with our model of the two domains as continuous volumes.

Because the extracellular space is expansive, and therefore highly conductive, relative to the intracellular space, it can be assumed that spatial differences in transmembrane potential are dominated by changes in intracellular potential rather than extracellular potential. That is,

$$\nabla V_m = \nabla \phi_i - \nabla \phi_e \approx \nabla \phi_i \quad (70)$$

The source can therefore be written in terms of the transmembrane potential as

$$\vec{J}_i = -\sigma_i \nabla V_m \quad (71)$$

Substituting this expression into (67) yields

$$UEG = -\int \sigma_i \nabla v_m \cdot \nabla Z \, dV. \quad (72)$$

By using a similar manipulation to that used in (65), this can be rewritten as

$$UEG = -\int \sigma_i v_m \nabla Z \cdot d\vec{S} + \int v_m \nabla \cdot \sigma_i \nabla Z \cdot dV. \quad (73)$$

The second term contains another volume integral of $\nabla^2 Z$, but for the newly defined surface, only the exploring electrode is contained within the volume, so the function has just one singularity at p . The integral of that singularity is $-1/\sigma_B$, where $\sigma_B = \sigma_i + \sigma_e$ is the bulk conductivity, the combined effective conductivity of the intracellular and extracellular domains.

Hence we can write

$$UEG = -\int \sigma_i v_m \nabla Z \cdot d\vec{S} - \frac{\sigma_i}{\sigma_i + \sigma_e} v_m(p) \quad (74)$$

In the second term, $v_m(p)$ is the membrane potential at the electrode in the myocardium. Note that this term matches the expression for local extracellular potential derived using the core-conductor model, hence it is used as the UEG's "local component" in the model of [Potse et al, 2009]. The first term is the "remote component".

Appendix III. Deriving an Expression for the Extracellular Potential Induced by a Single Myocyte in an Infinite, Homogeneously Conductive Medium

First, consider a point current source with magnitude i_s . Assuming uniform extracellular conductivity σ_e , the current density \vec{J} at some distance r from the source will be evenly distributed about the surface of a hypothetical sphere with surface area $4\pi r^2$. Hence the magnitude of \vec{J} will be

$$|\vec{J}| = \frac{i_s}{4\pi r^2}. \quad (75)$$

The extracellular potential gradient at a distance r is then given by

$$\frac{d\phi_e}{dr} = -\frac{|\vec{J}|}{\sigma_e} = -\frac{i_s}{4\pi\sigma_e r^2} \quad (76)$$

Relative to some infinitely distant reference potential, the induced potential at some arbitrary point p can then be found by integrating with respect to r :

$$\phi_e = \frac{i_s}{4\pi\sigma_e r_p} \quad (77)$$

The membrane-current-per-unit-length \bar{i}_m of a roughly cylindrical myocyte can be modelled as a distribution of point sources along the axis of the cell. Each infinitesimal section, dx , of the cell's length then makes a contribution $d\phi_e$ to the potential at p .

$$d\phi_e = \frac{\bar{i}_m}{4\pi\sigma_e r} dx \quad (78)$$

The total contribution of this cell to the potential at p can then be found as a function of time by integrating along the cell's length L .

$$\phi_e(p, t) = \frac{1}{4\pi\sigma_e} \int_L \frac{\bar{i}_m(x, t)}{r(x, p)} dx \quad (79)$$

As before, it is convenient to replace $\bar{i}_m(x, t)$ with the transmembrane voltage $v_m(x, t)$, the time-course of which is more easily observable and thus more familiar. An expression relating the two can be achieved by considering the intracellular current and potential, i_i and ϕ_i . From Ohm's law, these quantities are related by the intracellular resistance-per-unit-length, \bar{R}_i , as

$$i_i = -\frac{1}{\bar{R}_i} \frac{\partial \phi_i}{\partial x}. \quad (80)$$

Because the extracellular domain is expansive, $\bar{R}_e \ll \bar{R}_i$, which implies that $\partial \phi_e / \partial x \ll \partial \phi_i / \partial x$. Applying this relationship to the definition of transmembrane voltage ((49), page 212) gives

$$\frac{\partial v_m}{\partial x} = \frac{\partial \phi_i}{\partial x} - \frac{\partial \phi_e}{\partial x} \approx \frac{\partial \phi_i}{\partial x} \quad (81)$$

(80) then becomes

$$i_i = -\frac{1}{\bar{R}_i} \frac{\partial v_m}{\partial x} \quad (82)$$

Consider that the only possible external contribution to the intracellular current is the membrane current-per-unit-length \bar{i}_m . Kirchoff's law then yields

$$\bar{i}_m = -\frac{\partial i_i}{\partial x}. \quad (83)$$

Combining (82) and (83) we get

$$\bar{i}_m = \frac{1}{\bar{R}_i} \frac{\partial^2 v_m}{\partial x^2} \quad (84)$$

To widen the applicability of this expression, \bar{R}_i can be written in terms of a measurable property that is not dependent on cell geometry, the intracellular conductivity σ_i . For a cell radius a , σ_i is related to \bar{R}_i by

$$\bar{R}_i = \frac{1}{\pi a^2 \sigma_i} \quad (85)$$

Combining (84) and (85) into (79), we reach

$$\phi_e(p, t) = \frac{a^2 \sigma_i}{4 \sigma_e} \int_L \frac{\partial^2 v_m}{\partial x^2}(x, t) \frac{1}{r(x, p)} dx \quad (86)$$

The $1/r$ term in the integrand of can be thought of as a transfer function H that is convolved with local membrane behaviour to determine its influence at p . If p is defined by the coordinates x' , y' , and z' , then the convolution form of is

$$\phi_e(p, t) = k \int_L H(x - x') \frac{\partial^2 v_m}{\partial x^2}(x, t) dx, \quad (87)$$

where k combines the constant factors and $H(x - x')$ is

$$H(x - x') = \frac{1}{\sqrt{(x - x')^2 + (y')^2 + (z')^2}} \tag{88}$$

Note that the coordinate system has been defined such that $y = z = 0$ along the axis of the myocyte. This expression allows to examine the relationship between electrode positioning and the spatial specificity of the UEG. Figure 23 (page 61) shows the transfer function $H(x - x')$ for $p = (0, 0, 0.2mm)$ and $p = (0, 0, 2mm)$.

Appendix IV. Custom Software for Automated QRS Detection

AppIV.A. Detection algorithm

QRS complexes, produced by ventricular depolarisation, are typically the most prominent feature of a surface ECG, and are therefore the most reliable indicator of beat-to-beat timing. It is for this reason that, as mentioned previously, most heart rate variability (HRV) measurements are calculated from a series of RR intervals, the intervals between consecutive R waves. Because these measurements are typically taken from lengthy recordings, automated QRS detection algorithms substantially decrease the time-cost of the calculations. Numerous detection algorithms have been developed, with widely varying degrees of complexity, but even very simple algorithms can achieve accuracies of better than 99.5 percent^[Köhler et al, 2002].

The total length of signal to be tested for the experiments discussed in this thesis was relatively small when compared with the 24-hour recordings used for diagnostic purposes, which necessitate extremely high reliability. Hence it was acceptable to implement a relatively simple QRS detection system incorporating a small amount of guidance from a human operator to avoid errors. Section AppIV.B describes a procedure developed to minimise the time-cost associated with this ‘guidance’.

A more challenging requirement was that the algorithm needed to be robust to the large ECG distortions induced by the simultaneous use of MRI, as described in the previous section. It was shown that the higher frequency content of the MRI artefact could be satisfactorily attenuated, leaving only the blood-flow artefact and small distortions caused by processing. An algorithm was developed that successfully identified QRS complexes in the presence of these distortions using two key principles: 1) By incorporating only the simplest assumptions about the morphology of the QRS complexes, the algorithm remains robust to artefacts that distort the signal in unpredictable ways. 2) The use of the previously described segment-mean technique for MRI artefact attenuation ensures that the remaining artefacts can be distinguished from QRS complexes based on their frequency content. The frequency content of blood-flow artefacts is typically much lower than that of the QRS except during ventricular contraction. Ventricular contraction is unlikely to occur at the same time as the QRS; instead, it generally coincides with ventricular repolarisation (see **Figure 2**, page 24). Hence the artefact that coincides with the QRS can be expected to occupy lower frequencies.

It should be noted that, in the presence of certain arrhythmias or extremely high heart rates, some overlap may occur between the QRS and ventricular contraction. In these cases, the blood flow artefact caused by the concurrent use of MRI and ECG may make it difficult to reliably identify QRS complexes. This is an important limitation for experiments of this nature, rather than a limitation of the proposed QRS detection procedure.




Figure 91: An illustration of the R wave detection method described in steps II-IV above, applied to an ECG signal containing blood-flow artefacts caused by the concurrent use of MRI.

The developed algorithm is described below and in Figure 91:

- I. Two criteria are chosen for identifying the R wave: minimum magnitude, M , and maximum duration, T . Because the R waves appear distinctly as short sharp spikes, when chosen carefully these two criteria are usually sufficient for identifying all R waves in a signal.
- II. The signal is scanned using a moving time-window of length T that stretches ahead of the initial point, t .
- III. The difference between the signal at t and $t+T$ is taken as a measure of the local overall trend in the data. This trend is effectively subtracted out of the signal so that we deal only with the local spike, if there is one. This step is effectively a high-pass filter to exploit the difference in frequency content between the QRS and the blood-flow artefact, as explained previously. The key difference between this approach and conventional filtering is that the use of a finite window amounts to a precisely defined cut-off frequency, while the removal of a linear trend avoids unpredictable distortions of the QRS.
- IV. If the changes in the signal from t to $t+T/2$ and from $t+T/2$ to $t+T$ are both greater than M and are of opposite sign, the deflection qualifies as an R wave (which can be assumed to be part of a QRS complex). Although the implicit assumption that the peak of

the R-wave will occur halfway through the time window is not strictly valid, if the parameter T is chosen appropriately the method works well.

- V. The time at which the signal reaches its maximum absolute value within the window is taken as the R-wave time, the representative index of the timing of that heart beat.

Two drawbacks of the proposed QRS detection algorithm should be noted at this point. Firstly, by using a single set of detection parameters for each signal, the algorithm assumes that the basic morphology of the ECG trace does not change significantly throughout the signal. This assumption is valid only when the heart is allowed to follow its natural rhythm, with no extreme stimuli and no ectopic beats occurring. Hence in many experimental settings the assumption does not hold, particularly when intermittent pacing is used. The activation wavefront of a beat initiated naturally in the sinus node progresses very differently from one initiated by a pacing electrode at some other location, so the resultant ECG morphology is very different. It is therefore unlikely that a single set of detection parameters will be well suited to the whole signal in those circumstances. Secondly, the use of user-defined signal-specific parameters slows the overall speed of data processing due to the additional time required to select those parameters. Steps taken to ameliorate this problem are described in the following section.

AppIV.B. An efficient procedure for choosing detection parameters

Automatic detection of QRS complexes greatly increases the tractability of HRV measurements, which require long recordings for which manual annotation would be impractical. However, even the most advanced algorithms require manual verification and correction when novel results are being presented. With this in mind, a graphical user interface (GUI) called RRvarGUI (RR interval variability GUI) was developed to minimise the time taken to perform manual verification of detected QRS complexes. As seen in Figure 92, the GUI provides input boxes for adjustments to the detection parameters, and also displays the ECG itself and the time series of RR intervals. Markers can optionally be superimposed onto these plots to indicate the timing of events in the experiment.

As a tool to assist in the evaluation and optimisation of the QRS detection parameters, the interface also incorporates a Poincaré plot of the RR intervals. A Poincaré plot presents a series with each value plotted against the preceding value. Hence it offers a useful qualitative measure of beat-to-beat variability. With a small amount of practice, it can be interpreted very quickly and is richly informative. Poincaré plots of RR intervals have proven diagnostic value^[Woo et al, 1992] but their implementation as a tool for parameter refinement is novel.

The key reason for employing a Poincaré plot in the optimisation of QRS detection parameters is that they are only sensitive to RR-variability on a time-scale of one heart beat, and large variations on this scale are characteristic of erroneous QRS detection. If the detection

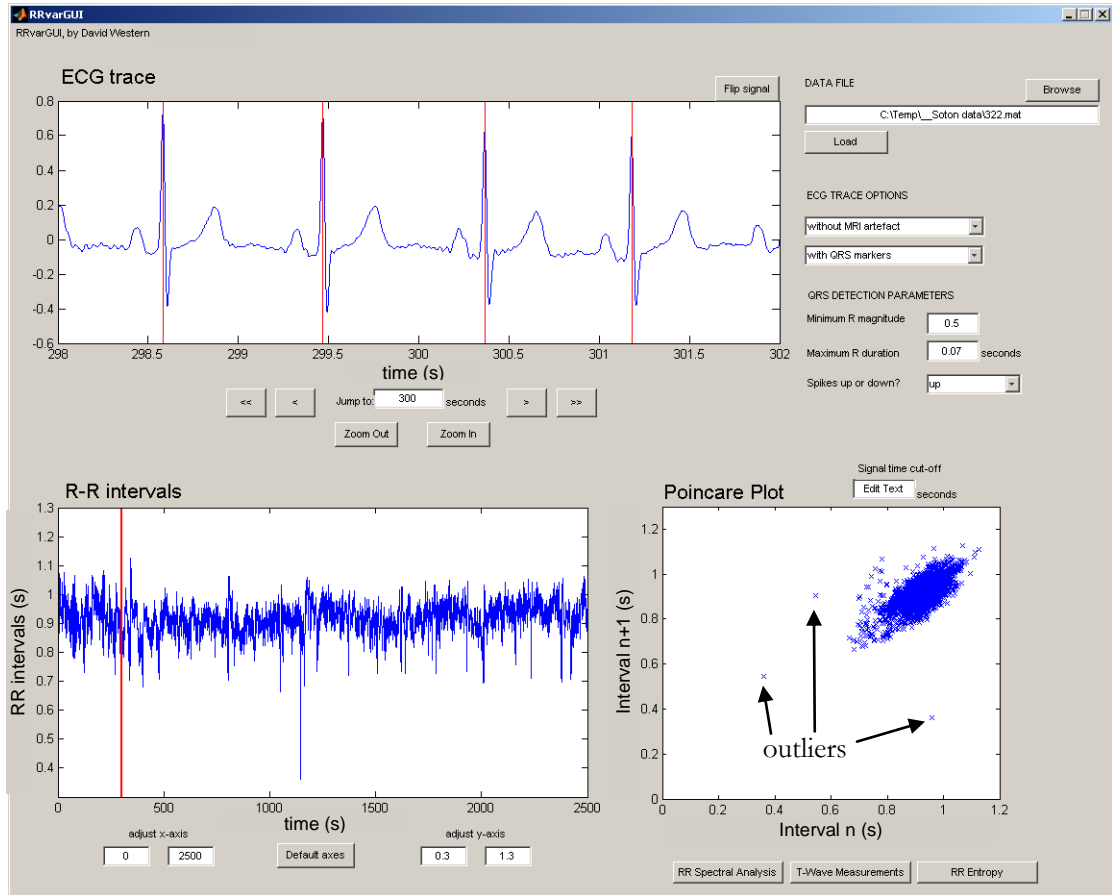


Figure 92: RRvarGUI, developed for ECG processing. Detected QRS complexes are marked by red vertical lines on the ECG trace. On the RR interval plot, two red lines (too close to distinguish from each other in this screenshot) indicate the current viewing window for the ECG trace. In the Poincaré plot, outliers highlight the presence of a detection error. Three buttons in the bottom-right corner can be used to begin optional subsequent processing: “RR Spectral Analysis” (HRV), “T-Wave Measurements”, and “RR Entropy” (HRV).

algorithm has missed a beat or made a false-positive detection, an RR interval will be found that is abnormally long or short, respectively. This incorrect measurement will produce two or three outliers on the Poincaré plot. Smoother, physiological variations in RR intervals will not produce such outliers, even if they amount to a large deviation from the mean RR interval. If only the time-series window is used for verification then these natural variations may be confused with the brief variations associated with detection errors, especially when viewing on a coarse time-scale. Hence the Poincaré plot offers a much faster means of determining whether any detection errors are present, while the RR interval time-series window can then be used to assist in locating those errors.

The intended procedure for selecting QRS detection parameters using RRvarGUI is as follows:

- I. The user makes an initial estimate of an appropriate value for the minimum magnitude criterion, M , by viewing a small section of the signal in the ECG trace window. In Figure 92, $M = 0.5$ is used. The maximum duration criterion rarely needs to be changed from $T = 0.07s$.

- II. The Poincaré plot is examined. If the parameters have been chosen appropriately, one can expect to see a closely grouped cloud of points.
- III. If the Poincaré plot contains outliers, the user examines the RR interval plot to locate the anomaly. In Figure 92, we can see that an abnormally short RR interval of less than 0.4s has been found roughly 1150 seconds from the start of recording.
- IV. The user adjusts the viewing window for the ECG trace to find the suspicious measurement, then assesses the cause of any error and adjusts the detection parameters accordingly. The detection algorithm executes automatically when these parameters are adjusted.
- V. The user repeats steps II-IV until the detection parameters are chosen satisfactorily.

AppIV.C. An interface for manual corrections

For some signals, no single set of parameters exists that achieves error-free QRS detection with the developed algorithm. Rejecting all such signals could substantially reduce the size of the data set, jeopardising the statistical significance of any results from the study. Instead, it is preferable to use manual correction of the measurements to bridge the gap in reliability between human detection and automated detection. Another custom GUI was developed to facilitate the correction process, as seen in Figure 93. This GUI is executed after QRS detection parameters have been chosen and before subsequent analyses such as HRV measurements.

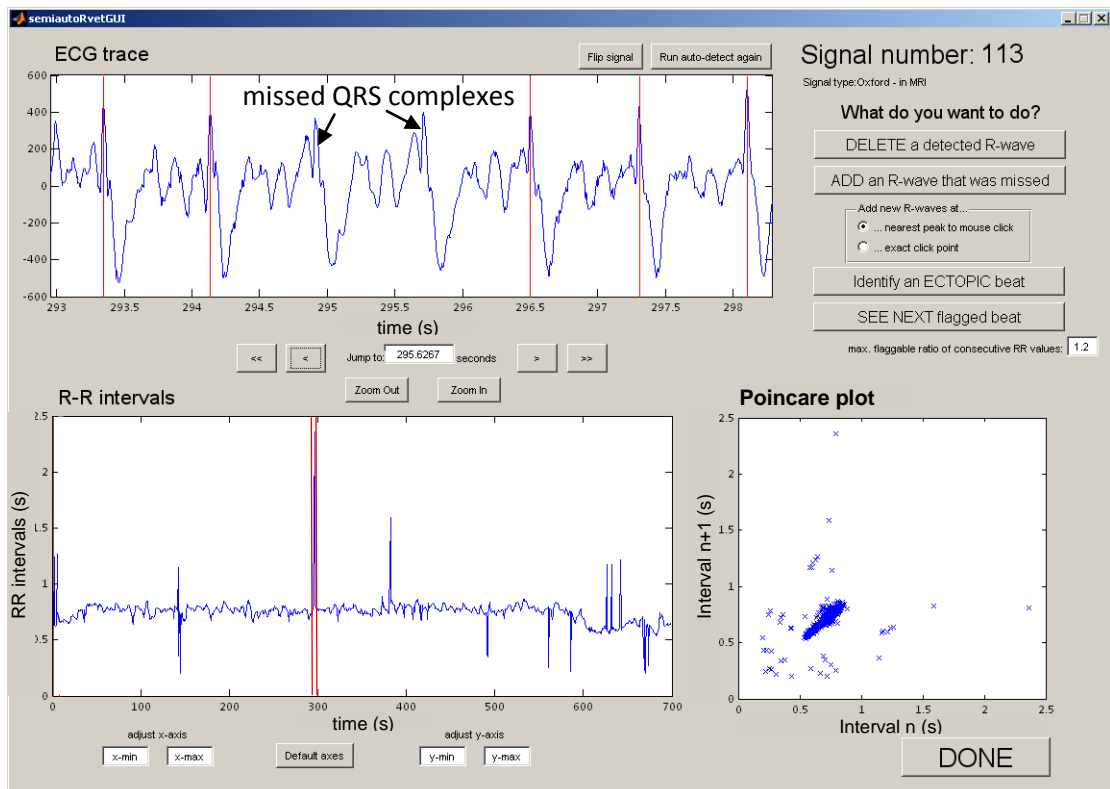


Figure 93: A GUI developed to allow quick and easy correction of detected QRS complexes. By point-and-click interaction, R-waves can be added, deleted, or discounted as ectopics. In this example of a heavily distorted signal, the ECG trace shows two complexes that have not been detected because they are slightly smaller than the others.

The GUI was designed to allow R-wave times to be added, deleted, or discounted by clicking within the ECG trace window. When new R-waves are being added, the user has the option of inserting at the exact point where they click or at the peak nearest to the point where they click. Thus the user can choose between full human control and increased speed for each correction made. The ability to discount certain QRS complexes by labelling them as ‘ectopics’ is important for HRV analysis because it allows them to be replaced by interpolation in the RR interval series, minimising their distortion of the HRV measurement as an indicator of ANS input to the sino-atrial node. In other analyses, such as T-wave measurements, any beats that are expected to corrupt the measurement can be excluded by treating the preceding QRS as an ectopic.

Appendix V. The Berger Method for Calculating QT Intervals

The QT interval can broadly be interpreted as the time between a QRS complex and the subsequent T-wave in an ECG recording. A more precise definition can be achieved by defining fiducial points, the points that mark the beginning and end of the interval. Given the nature of the QT interval as a global approximation of the timing of cardiac behaviour, there is some degree of arbitrariness in the selection of instantaneous fiducial points. When QT intervals are measured manually, they are typically taken from the onset of the QRS wave to the end of the T-wave. There is scope for uncertainty in these reference points since either of them may be temporally indistinct. They are identified as the initial deviation from or return to the isoelectric baseline, but that baseline cannot be defined with certainty and any divergence from it may be gradual. The isoelectric line is typically identified using a low-pass filter or by interpolating between PR intervals (the gaps between each P wave and the onset of the next QRS complex). For the QRS fiducial point, alternatives that are more easily distinguished include the nadir of the Q-wave or the peak of the R-wave. Similarly, the peak of the T-wave, rather than the end of it, is sometimes used to mark the end of the QT interval. When the end of the T-wave is used, ambiguities are normally avoided by extrapolating along the steepest downslope of the T-wave to find the intersection between this construction line and the isoelectric baseline (Figure 94).




Figure 94: Where the end of a T-wave is indistinct the intersection of the steepest downslope with an isoelectric baseline may be used. In this particular case the region of uncertainty is approximately 30 ms long.

A common problem with all of the conventions described above is that each fiducial point is vulnerable to localised variations that are not representative of global behaviour. They produce QT interval measurements that are useful for diagnostic purposes, where extreme deviations from the norm can be identified without the need for great precision; they are not sufficiently precise for tracking the time-course of repolarisation properties in response to subtle stimuli. Berger and colleagues propose a method that provides a more stable measurement, at least on an intra-individual basis, by taking into account the whole shape of the T-wave^[Berger et al, 1997]. Their method is as follows (see also Figure 95):

- I. Let $x(n)$ represent the n^{th} sample of the ECG signal.
- II. R-wave peak for each i^{th} beat is identified at $n = R_i$, using any QRS detection method.

- III. The user defines the beginning and end of a QT interval for a single beat, which will serve as a template and is referred to as beat k . These points are n_0 and n_1 , respectively.
- IV. Each i^{th} beat is stretched temporally about the R-wave peak, R_i , by a factor α . The squared error between the T-wave in question and the template T-wave is defined as...

$$\varepsilon_i(\alpha) = \sum_{j=s}^{n_1-R_k} [x(R_k + j) - x(R_i + \alpha j)]^2 \quad (89)$$

... where s is a blanking period, normally equivalent to 50ms, used to exclude the tail-end of the QRS complex from the error calculation. The value of α that minimises this error is $\hat{\alpha}_i$. Note that α is rounded to obtain an integer sample number.

- V. A nominal QT interval for any i^{th} beat is taken as...

$$QT_i = \hat{\alpha}_i \frac{(n_1 - n_0)}{f_s} \quad (90)$$

... where f_s is the sampling frequency used.

This method does not improve on conventional QT interval measurements in terms of diagnostic value or the reliability of comparisons between individuals, since the definition of the onset and end of the template interval is still somewhat arbitrary. However, the Berger method can be expected to offer reduced sensitivity to small noise deflections in critical areas of the ECG. It is more appropriate than conventional measurements for tracking the time-course of repolarisation behaviour in a single subject because it takes the shape of the whole T-wave into account. For this reason, it was implemented as an extension to RRvarGUI, along with the conventional means of measuring QT interval. The ability to measure T-wave amplitude and area were also incorporated into the program.




Figure 95: Illustration of the Berger method. Note that the template and fitted T-waves have been elevated slightly for clarity.

Appendix VI. Matlab Code to Find and Remove Pacing Artefacts

```
% Locates pacing spikes in a unipolar electrogram.
% Created by David Western,
% Department of Mechanical Engineering, University College London.
function [spikes, masterwave] = replace_pacing_spikes2(spiky, handles)

h = waitbar(0, 'Replacing pacing spikes...');

% Remove baseline wander:
[b, a] = MYbutter(3, 0.5/(handles.Fs/2), 'high');
masterwave = MYfiltfilt(b, a, handles.rawwave);

% Differentiate 'spiky' once for the sake of identifying a fiducial point
% (later):
spikydiff1 = diff(spiky);
temp = spiky;
% Differentiate 'spiky' 8 times!:
spiky = diff(spiky, 8);
% Compensate for the reduction in vector length:
spiky = [spiky(1); spiky(1); spiky(1); spiky(1); spiky(1); spiky(1); spiky(1); spiky(1)];
spiky = [spiky; spiky(end); spiky(end); spiky(end); spiky(end); spiky(end)];

waitbar(0.1, h)

% Create low-pass-filtered wave to replace spikes with
[b, a] = MYbutter(3, [0.5 12]/(handles.Fs/2), 'bandpass');
lpfwave = MYfiltfilt(b, a, handles.rawwave);

% Find the spikes and blank them out
thresh = 2.5*prctile(spiky, 99);
k = round(0.03*handles.Fs);
c = 0;
while k < length(spiky) - round(0.03*handles.Fs)
    k = k + 1;
    if abs(spiky(k)) > thresh/6
        % Seek beginning and end of spiky activity:
        [tip, tipind] = max(abs(spiky(k:k+round(0.01*handles.Fs)))));
        if tip < thresh
            continue
        end
        k = k + tipind - 1;
        ind = find(abs(spiky(k-round(0.012*handles.Fs):
            k+round(0.012*handles.Fs))) > tip/15);
        if ~isempty(ind)
            ind = ind + (k - round(0.012*handles.Fs) - 1);
            startspike = ind(1) - round(0.0017*handles.Fs);
            endspike = ind(end) + round(0.004*handles.Fs);

            if abs(masterwave(endspike) - masterwave(startspike)) >
                0.5*range(masterwave(startspike:endspike))
                % Net-shift criterion
                continue
            end
            halfsec = round(0.5*handles.Fs); % Number of samples in half a
                second (for convenience).
            if k > 2*halfsec & k < length(masterwave) - 2*halfsec
                if abs(mean(masterwave(startspike:endspike))) >
                    0.3*range(masterwave(startspike-
                    halfsec:endspike+halfsec))
                    % The spike is big and roughly monophasic. Ignore it.
                    continue
                end
            end
        end
    end

    % Choose a fiducial point for the spike - the first big
```

```
% upstroke.
maxup = max(spikydiff1(startspike:endspike));
ups = find(spikydiff1(startspike:endspike)>0.25*maxup);
if isempty(ups)
    continue
end

c = c+1;
spikes(c) = ups(1)+startspike;
waitbar(k/length(spiky),h)

% Blank out that region:
implant = lpfwave(startspike:endspike);
% Retrend implant to avoid discontinuities:
% Shift implant so that first point meets masterwave:
implant = implant+(masterwave(startspike)-implant(1));
% Add a trend to implant so that last point meets masterwave:
endlevel = masterwave(endspike);
for d = 2:length(implant)
    implant(d) = implant(d)+(d-1)/(length(implant)-1)*(endlevel-
        implant(end));
end
% Insert implant:
masterwave(startspike:endspike) = implant;

k = k+round(0.03*handles.Fs);
end
end
close(h)
```

Appendix VII. Matlab Code to Find Stabilised Activation Times (\mathcal{A}_f)

```
% Find STABILISED activation times in a unipolar electrogram.
% Created by David Western, 21-7-2010
% Department of Mechanical Engineering, University College London.
function [A,Astart,Aend] = getAcent(handles)

params.window = 0.15; % Window size (in s) in which an activation can occur.
params.Fc = [0.5 150]; % Bandpass filter cut-off frequencies in Hz.
params.R_to_broadRNG = 0.3; % Minimum value of DownRNG/broadRNG.
params.shift_factor = 0.6;
params.WinRNG_to_broadRNG = 0.00; % Minimum value of WinRNG/broadRNG.

[b,a] = MYbutter(4,params.Fc/(handles.Fs/2),'bandpass'); % Filter coefficients.
fwave = MYfiltfilt(b,a,handles.nswave);
[b,a] = MYbutter(3,[48 52]/(handles.Fs/2),'stop'); % Filter coefficients.
fwave = MYfiltfilt(b,a,handles.nswave);

% Miscellaneous preliminaries
N = length(fwave);
broadRNG_ref_per = 0.2*handles.Fs; % Refresh period for broadRNG
broadRNG = range(fwave(1:round(4*broadRNG_ref_per)));
broadSlPt1 = prctile(diff(fwave(1:round(4*broadRNG_ref_per))),1);
[b,a] = MYbutter(3,[0.5 30]/(handles.Fs/2),'bandpass');
fwave2 = MYfiltfilt(b,a,handles.nswave);

nsampwin = round(handles.Fs*params.window); % Number of samples in the size
% of window in which an
% activation spike can occur.

p = 1; % Search position marker.
c = 0; % Count of activation times found.

getAWait = waitbar(0,'Finding activation times...');

jump = round(0.01*handles.Fs);
next_broadRNG_refresh = 4*broadRNG_ref_per;
while p < N-2*nsampwin

    if p>next_broadRNG_refresh && c>0

        if p-A(c)<1*handles.Fs
            ind = min(A(c)+round(0.09*handles.Fs),p-round(
                1.5*broadRNG_ref_per)):p+round(5*broadRNG_ref_p
                er);
        else
            ind = p-round(1*handles.Fs):p+round(3*broadRNG_ref_per);
        end
        if ind(end)>length(fwave2)
            ind = p-4*broadRNG_ref_per:p;
        end
        broadRNG = range(fwave2(ind));
        broadSlPt1 = prctile(diff(fwave2(ind)),1);
        next_broadRNG_refresh = next_broadRNG_refresh+broadRNG_ref_per;
    end
    p = p+jump;

    if rem(p,500)==0
        waitbar(p/N,getAWait)
    end

    SigInWin = fwave(p:p+nsampwin-1);
    slopes = SigInWin(3:end)-SigInWin(1:end-2);
    slopes = [NaN; slopes];
```

```

down = find(slopes<0);
if isempty(down)
    continue
end

DownRNG = range(SigInWin(down));

front_cliplim = round(0.35*nsampwin);
back_cliplim = round(0.65*nsampwin);
if DownRNG > params.R_to_broadRNG*broadRNG
    % Narrow down window:
    % If 1/20 of the window contains less than 1/100th of the range,
    % clip it.
    cutoff = 1;
    sigsec = SigInWin(cutoff:cutoff+round(0.05*nsampwin));
    srrange = max(sigsec)-min(sigsec);
    while srrange<0.02*DownRNG && cutoff<front_cliplim
        cutoff = cutoff+round(0.05*nsampwin);
        sigsec = SigInWin(cutoff:cutoff+round(0.05*nsampwin));
        srrange = max(sigsec)-min(sigsec);
    end
    cutoff = cutoff-round(0.05*nsampwin);
    cutoff = max(1,cutoff);
    SigInWin = SigInWin(cutoff:end);
    ptemp = p+cutoff-1;
    front_cliplim = front_cliplim-cutoff+1;
    back_cliplim = back_cliplim-cutoff+1;
    % Same backwards.
    cutoff = length(SigInWin);
    sigsec = SigInWin(cutoff-round(0.05*nsampwin):cutoff);
    srrange = max(sigsec)-min(sigsec);
    while srrange<0.02*DownRNG && cutoff>back_cliplim+0.1*nsampwin
        cutoff = cutoff-round(0.05*nsampwin);
        sigsec = SigInWin(cutoff-round(0.05*nsampwin):cutoff);
        srrange = max(sigsec)-min(sigsec);
    end
    SigInWin = SigInWin(1:cutoff);
    back_cliplim = min(back_cliplim,cutoff);
    net_shift = abs(SigInWin(end)-SigInWin(1));

    slopes = SigInWin(3:end)-SigInWin(1:end-2);
    down = find(slopes<0);
    DownRNG = range(SigInWin(down));
    if isempty(net_shift) || isempty(DownRNG)
        disp('too much clipping?')
        continue
    end
    if net_shift<params.shift_factor*DownRNG && DownRNG >
        params.R_to_broadRNG*broadRNG
        WinRNG = range(SigInWin);
        PreRNG = range(fwave(max(1,ptemp-
            round(1*length(SigInWin))):ptemp+round(0.01*len
            gth(SigInWin))));
        if nanmin(slopes)<broadSlPt1 && WinRNG>0.85*PreRNG % &&
            WinRNG>params.WinRNG_to_broadRNG*broadRNG
            % This last criterion (WinRNG>0.85*PreRNG) was only
            % necessary before effective removal of pacing artefacts
            % was developed.
            %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
            % All criteria have been met -> we have an activation
            % time.
            c = c+1;

            cand = 1; % Number of candidate windows that may be used to

```

```

% represent this activation.

clear DownRNG net_shift WinRNG candwin candlpwin dropB dropE

% There will be a range of positions around here in
% which we could place the window and still meet the
% criteria. Let's choose the best one.

% Preserve these in case the window refinement fails.
origSigInWin = SigInWin;
try
p2 = p;
while p2 < min(p+round(1.5*nsampwin), length(fwave2)-nsampwin)
    SigInWin = fwave(p2:p2+nsampwin-1);
    slopes = SigInWin(3:end)-SigInWin(1:end-2);
    slopes = [NaN; slopes];

    down = find(slopes<0);
    if isempty(down)
        p2 = p2+ceil(0.004*handles.Fs);
        continue
    end
    DownRNG(cand) = range(SigInWin(down));

    front_cliplim = round(0.35*nsampwin);
    back_cliplim = round(0.65*nsampwin);
    if DownRNG(cand) > params.R_to_broadRNG*broadRNG
        % Narrow down window:
        % If 1/20 of the window contains less than 1/100th of the
        % range,
        % clip it.
        cutoff = 1;
        sigsec = SigInWin(cutoff:cutoff+round(0.05*nsampwin));
        ssrange = max(sigsec)-min(sigsec);
        while ssrange<0.02*DownRNG(cand) && cutoff<front_cliplim
            cutoff = cutoff+round(0.05*nsampwin);
            sigsec =
                SigInWin(cutoff:cutoff+round(0.05*nsampwin));
            ssrange = max(sigsec)-min(sigsec);
        end
        cutoff = cutoff-round(0.05*nsampwin);
        cutoff = max([1 cutoff]);
        SigInWin = SigInWin(cutoff:end);
        p2 = p2+cutoff-1;
        front_cliplim = front_cliplim-cutoff+1;
        back_cliplim = back_cliplim-cutoff+1;
        % Same backwards.
        cutoff = length(SigInWin);
        sigsec = SigInWin(cutoff-round(0.05*nsampwin):cutoff);
        ssrange = max(sigsec)-min(sigsec);
        while ssrange<0.02*DownRNG(cand) &&
            cutoff>back_cliplim+0.1*nsampwin
            cutoff = cutoff-round(0.05*nsampwin);
            sigsec = SigInWin(cutoff-
                round(0.05*nsampwin):cutoff);
            ssrange = max(sigsec)-min(sigsec);
        end
        SigInWin = SigInWin(1:cutoff);
        back_cliplim = min(back_cliplim,cutoff);

    net_shift(cand) = abs(SigInWin(end)-SigInWin(1));

    slopes = SigInWin(3:end)-SigInWin(1:end-2);
    down = find(slopes<0);
    if isempty(down)

```



```

        p2 = p2+ceil(0.005*handles.Fs);
        continue
    end
    DownRNG(cand) = range(SigInWin(down));

    if net_shift(cand)<params.shift_factor*DownRNG(cand) &&
        DownRNG(cand) > params.R_to_broadRNG*broadRNG
        WinRNG(cand) = range(SigInWin);
        PreRNG = range(fwave(max(1,p2-
            round(1*length(SigInWin))):p2+round(0.01*length
                (SigInWin))));
        if WinRNG(cand)>0.85*PreRNG %&&
            WinRNG(cand)>params.WinRNG_to_broadRNG*broadRNG
            % This is a valid candidate. Don't
            % overwrite it.
            candp(cand) = p2;
            candwin{cand} = SigInWin;
            cand = cand+1;
        end
    end
end
    p2 = p2+ceil(0.004*handles.Fs);
end

% The final value of cand never represents a valid
% candidate:
DownRNG = DownRNG(1:cand-1);

net_shift = net_shift(1:cand-1);
WinRNG = WinRNG(1:cand-1);
cut1 = find(WinRNG > 0.95*max(WinRNG));
    % cut 1 may contain multiple values, hence cut 2.
cut2 = find(DownRNG(cut1) == max(DownRNG(cut1)));
net_shift = net_shift(cut1);
net_shift = net_shift(cut2);
    % cut 2 may also contain multiple values, hence cut 3.
[ignore,cut3] = min(net_shift);

candp = candp(cut1);
candp = candp(cut2);
candp = candp(cut3);

for k = 1:numel(cut1)
    temp1{k} = candwin{cut1(k)};
end
candwin = temp1;
clear temp1
for k = 1:numel(cut2)
    temp1{k} = candwin{cut2(k)};
end
candwin = temp1{cut3};
clear temp1
if isempty(candp) || isnan(candp)
    t = p/handles.Fs
end
p = candp; % That's the one!
SigInWin = candwin;
if strcmp(class(SigInWin),'cell')
    SigInWin = SigInWin{1};
end
SigInWin = fwave(p:p+length(SigInWin)-1);
catch
disp('Activation window refinement failed.')
time_of_failure = (p-1)/handles.Fs

```

```

    SigInWin = origSigInWin;
    p = ptemp;
end

% Now find nominal activation time:
slopes = SigInWin(3:end)-SigInWin(1:end-2);
slopes = [NaN; slopes]; % So that the indices of slopes
                    % match those of SigInWin.
slpThresh = 0.6*nanmin(slopes); % Slope threshold.
thresh = slpThresh;
Atrough = find(slopes<thresh);

% The nominal activation time is taken as the centroid of
% the signal at the points it Atrough.
A(c) = sum(Atrough.*(slopes(Atrough)-
                    thresh))/sum(slopes(Atrough)-thresh);
A(c) = round(A(c))+(p-1); % Translate A(c) to global indices.

%     [~,a] = nanmin(slopes); % Replace the previous two lines
%     A(c) = a+(p-1);         % with these two to use the
%                             % conventional index i.e. no
%                             % stabilisation.

Astart(c) = p; % Start of activation wave
Aend(c) = p+length(SigInWin)-1; % End of activation wave

p = Aend(c)+round(0.1*handles.Fs);
end
end
end
end
close(getAWait)

```

Appendix VIII. Matlab Code to Find Stabilised Recovery Times (R_s)

```

% Find STABILISED recovery times in a unipolar electrogram, given
% activation times.
% Created by David Western, 21-7-2010
% Department of Mechanical Engineering, University College London.
function [R,ARI] = getRcent(handles)

fwave = handles.nswave;

[b,a] = MYbutter(4,[0.5 35]/(handles.Fs/2),'bandpass');
% [b,a] = cheby1(3,0.5,[0.5 35]/(handles.Fs/2),'bandpass');
fwave = MYfiltfilt(b,a,fwave);

LocalMag = std(fwave);
N = length(handles.A);
Nfw = length(fwave);
if strcmp(waitbar_opt,'on')
    getRWait = waitbar(0,'Finding recovery times...');
end

for j = 1:length(handles.A)-1 % For each beat...
    ind = max(1,handles.A(j)-round(0.1*handles.Fs)):min([Nfw
        handles.Aend(j+1)
        handles.A(j)+ceil(1*handles.Fs)]);
    LocalMag = range(fwave(ind)); % LocalMag (like "broadRNG" in getAcent)
    % is frequently updated to give a rough estimate of the local
    % magnitude of the signal. This allows us to distinguish between
    % significant and insignificant features.

    if ~any(handles.IgnoreA == handles.A(j)) &...
        ~isnan(handles.AAint(j+1))
        % Based on previous activation detection, this is assumed to be a
        % legitimate beat.

        % Update the waitbar.
        if strcmp(waitbar_opt,'on')
            if rem(j,200)==0
                waitbar(j/N,getRWait)
            end
        end

        % Establish window in which to search for recovery time.
        % Tstart = handles.A(j)+round(0.09*handles.Fs); % Jump ahead of the
        % activation time to avoid the rest of the
        % activation complex.
        Tstart = max(handles.A(j)+round(0.07*handles.Fs),handles.Aend(j));
        Tstop = handles.Astart(j+1)-round(0.02*handles.Fs);
        Tstop = min(Tstop,handles.A(j)+round(sqrt(handles.AAint(j)/
            1000)/2*handles.Fs)); % This empirical
        % formula determines the upper limit of the
        % length of the T-wave window.
        Tstop = min(Tstop,handles.A(j)+round(handles.AAint(j)/
            1000*handles.Fs*0.8)); % A second
        % empirical formula, which does the same
        % thing.
        Tstop = min(Tstop,handles.A(j)+round(0.6*handles.Fs)); % ... and a
        % third. Tstop is now the index
        % representing the latest time that
        % satisfies all three limits.
        Trange = Tstart:Tstop; % The indices of the points in the signal
        % that contain the T-wave.
        if length(Trange)<0.07*handles.Fs
            Tup = NaN;
        else

```

```

slopes = fwave(Trange(3:end))-fwave(Trange(1:end-2));
slopes = [NaN; slopes]; % Corrects the indices, so that 'slopes'
                        % has the same length as fwave(Trange)
d2 = slopes(3:end)-slopes(1:end-2);
d2 = [d2; NaN]; % Corrects the indices, as above.

% Find beginning of Twave (STpoint, an approximately flat point
% around the beginning of the T-wave):
STdomain = 1:min(Tstop-Tstart,ceil(0.06*handles.Fs)); % Region
                                                    % in which STpoint can occur (0-0.06 seconds
                                                    % after Tstart).
sdslope = nanstd(slopes(STdomain)); % How much variation is
                                     % there in the slope here.
mnslope = nanmean(slopes(STdomain));
STcand = find(slopes(STdomain)-mnslope<0.01*sdslope); %
                                                    % Candidates for STpoint.
if isempty(STcand)
    [ignore,STpoint] = nanmin(abs(slopes(STdomain)));
else
    STpoint = STcand(end); % Use the latest of the candidates.
end
while STpoint<length(slopes) &&
    abs(slopes(STpoint+1))<abs(slopes(STpoint))
    && STpoint<round(0.45*handles.Fs)
    STpoint = STpoint+1;
end
Tstart = Tstart+STpoint-1;

% Chop off unnecessary front end of slopes and d2:
slopes = slopes(STpoint:end);
d2 = d2(STpoint:end);

% Recovery time (Tup) is where slope is at a maximum. Cap it at
% 0.45s
TupLim = round(0.45*handles.Fs)-(Tstart-handles.A(j))+1; % Upper
                                                    % limit for Tup.
TupLim = min(length(slopes),TupLim);
[maxslope,Tup] = nanmax(slopes(1:TupLim));
% [maxslope,Tup] = nanmax(slopes); % uncapped.

% Disqualification process:
vicinity = fwave(Tup+Tstart-
                round(0.03*handles.Fs):Tup+Tstart+round(0.
                03*handles.Fs)); % The signal immediately
                % surrounding Tup.
rangelim = 0.2*range(fwave(Tstart:Tstop)); % The minimum
                % acceptable range covered by a T-wave.
if maxslope<0 ||...
    Tup<=2 ||...
    range(vicinity)<rangelim ||...
    rangelim < 0.003*LocalMag
    Tup = NaN; % Disqualified!
elseif Tup>=length(slopes)-1 ||...
    Tup == TupLim ||...
    isempty(find(d2(Tup:end)<0)) ||... % This line is
    % probably redundant with the above.
    isempty(find(slopes(Tup:end)<2*LocalMag/(1*handles.Fs)))
    % Something's not quite right. Let's try adjusting the
    % window we examine to see if that helps.

    if ~isempty(slopes(1:end)<2*LocalMag/(1*handles.Fs))
        % Rethink end of window
        newTstop = find(slopes(TupLim:-
            1:1)<2*LocalMag/(1*handles.Fs),1);
        newTstop = TupLim-newTstop+1;
        if newTstop>0.02*handles.Fs

```

```

        slopes = slopes(1:newTstop);
        [maxslope,Tup] = nanmax(slopes);
        d2 = slopes(3:newTstop)-slopes(1:newTstop-2);
        d2 = [NaN; d2];
        vicinity = fwave(Tup+Tstart-
            round(0.03*handles.Fs):Tup+Tstart+round(0.
                02*handles.Fs));
        rangelim = 0.2*range(fwave(Tstart:Tstart+Tup));
        if Tup>=length(slopes)-1 ||...
            maxslope<0 ||...
            Tup==2 ||...
            Tup == newTstop ||...
            range(vicinity)<rangelim ||...
            rangelim < 0.01*LocalMag
            %
            Tup = NaN;
        else
            Tstop = newTstop+Tstart-1;
        end
    else
        Tup = NaN;
    end
else
    Tup = NaN;
end
end
clear rangelim

end

if ~isnan(Tup)
    Tp = find(slopes(Tup:end)<0,1)-1+(Tup-1);
    if isempty(Tp)
        % No downward slope after T-wave. Use first straight point
        % instead.
        [ignore,TpLim] = min(d2(Tup:min([Tup+round(0.1*handles.Fs)
            length(d2)-1])));
        TpLim = TpLim+Tup-1; % Earliest possible point for Tp.
        d3 = d2(TpLim+1:end)-d2(TpLim:end-1);
        Tp = find(d3<0.02*range(d3),1);
        if isempty(Tp)
            [ignore,Tp] = min(abs(d3));
        end
        Tp = Tp+TpLim-1;
    end
    [ignore,Td2] = nanmin(d2(Tup:min([Tp+round(0.03*handles.Fs)
        length(d2)])));
    Td2 = Td2+Tup-1+1;
    if isempty(Td2)
        R(j) = NaN;
        disp('no d2')
        A = (handles.A(j)-1)/handles.Fs;
    else
        R(j) = Tup+(Tstart-1);
    end
end
else
    R(j) = NaN;
end
end
R(j) = NaN;
end
if ~isnan(R(j))

    % This next bit reassigns R(j) as the hump in dV/dt around R(j).
    candstart = max([R(j)-round(0.06*handles.Fs) Tstart]);
    candstop = min([R(j)+round(0.1*handles.Fs) Tstop]);
    cands = candstart:candstop;
    slopes = fwave(cands(3:end))-fwave(cands(1:end-2));

```

```
slopes = [NaN; slopes];
maxslope = nanmax(slopes);
thresh = 0.5*maxslope;
hump = find(slopes>thresh);
newRj = sum((hump-hump(1)+1).*(slopes(hump)-
    thresh))/sum(slopes(hump)-thresh)+hump(1)-
    1;
R(j) = round(newRj)+(cands(1)-1); % Translate newRj to global
    indices.

ARI(j) = (R(j)-handles.A(j))/handles.Fs;

else
    ARI(j) = NaN;
    R(j) = handles.A(j);
end
end
```