Vital Sign Monitoring and Data Fusion in Haemodialysis



Yasmina Borhani

The Queen's College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Supervised by Prof. Lionel Tarassenko

Hilary 2013

To Yosef, Shayesteh and Aryan

Acknowledgements

It has been a great privilege to work under the supervision of Lionel Tarassenko, who has been an invaluable source of inspiration and guidance throughout the course of my work. I am very grateful for his support over the years, which gave me the confidence to think outside the proverbial box, and the freedom to be creative and exploratory in my research. I would also like to thank David Meredith and Clare MacEwan for sharing their wisdom in all aspects of haemodialysis and providing feedback from a medical perspective at key stages of my work, as well as Chris Pugh, whose input set the direction of the project. I am further indebted to Sheera Sutherland for her incredible efforts in collecting the data which forms the basis of this work, as well as to Andy Mosson for his technical support. Thanks also to Gari Clifford and David Clifton, with whom I had many beneficial discussions about my work and to Christina Orphanidou for first introducing me to Lionel back in 2008.

The advice given to me by my maternal grandfather has been pivotal in helping me make key decisions throughout various stages of my life, and I am forever indebted to him for encouraging me to pursue this DPhil to begin with, as well as for his endless support and inspiration from afar. My lovely parents, who by default were the audience of my occasional complaints, have always supported me in my endeavours and encouraged me to follow my interests from a very young age. I cannot thank them enough for this. A special thank you also to my aunt Sheila, who has been an inspirational figure in my life, for all the valuable advice throughout my educational path, for accompanying me up to Oxford on my very first day, and for the many treats over the years that brightened so many of my DPhil days. I am very fortunate to have had my uncle Ali in London throughout my time as a DPhil student, and am very grateful for all the kindness he has bestowed upon me, treating me as he would his own daughter and making me feel at home away from home.

Finally, I would like to thank the many friends and colleagues at the IBME who made my time in lab so much more enjoyable, and who turned this chapter of my education into a chapter of my life.

Vital Sign Monitoring and Data Fusion in Haemodialysis

Yasmina Borhani
The Queen's College

Doctor of Philosophy

Hilary Term 2013

Intra-dialytic hypotension (IDH) is the most common complication in haemodialysis (HD) treatment and has been linked with increased mortality in HD patients. Despite various approaches towards understanding the underlying physiological mechanisms giving rise to IDH, the causes of IDH are poorly understood.

Heart Rate Variability (HRV) has previously been suggested as a predictive measure of IDH. In contrast to conventional spectral HRV measures in which the frequency bands are defined by fixed limits, a new spectral measure of HRV is introduced in which the breathing rate is used to identify and measure the physiologically-relevant peaks of the frequency spectrum. The ratio of peaks leading up to the IDH event was assessed as a possible measure for IDH prediction. Changes in the proposed measure correlate well with the magnitude of abrupt changes in blood pressure in patients with autonomic dysfunction, but there is no such correlation in patients without autonomic dysfunction.

At present, routine clinical vital sign monitoring beyond simple weight and blood pressure measurements at the start and end of each session has not established itself in clinical practice. To investigate the benefits of continuous vital sign monitoring in HD patients with regard to detecting and predicting IDH, different population-based and patient-specific models of normality were devised and tested on data from an observational study at the Oxford Renal Unit in which vital signs were recorded during HD sessions. Patient-specific models of normality performed better in distinguishing between IDH and non-IDH data, primarily due to the wide range of vital sign data included as part of the training data in the population-based models. Further, a patient-specific data fusion model was constructed using Parzen windows to estimate a probability density function from the training data consisting of vital signs from IDH-free sessions. Although the model was constructed using four vital sign inputs, novelty detection was found to be primarily driven by blood pressure decreases.

Contents

List of Acronyms **xiv**

1 Introduction **1**

1.1 Patient Demographics 3
1.2 Haemodialysis as Renal Replacement Therapy 4
1.2.1 Dry Weight Estimation 6
1.3 Current Monitoring during Haemodialysis 8
1.3.1 Electrocardiogram 8
1.3.2 Heart Rate Variability Measures 11
1.3.3 Blood Oxygen Saturation 14
1.3.4 Blood Pressure 15
1.3.5 Breathing Rate 16
1.3.6 Haematocrit 17
1.4 Definitions of IDH 18
1.5 Causes of IDH 20
1.5.1 Cardiac Refilling 20
1.5.2 Autonomic Nervous System 21
1.6 Existing Methods for IDH Prevention 22
1.6.1 Pulse Transit Time 23
1.6.2 Relative Changes in Blood Volume (Haematocrit) 24
1.6.3 Photoplethysmograph-based methods 26
1.6.4 Heart Rate Variability (HRV) 27
1.7 Existing Methods for Predicting Patient Deterioration 35
1.8 Thesis Overview 37

2 Data Collection **40**
2.1 Signals and Derived Vital Signs 42
2.2 The Dataset 43
2.3 Artefact Rejection and Signal Quality 45
2.4 Dataset Statistics 46
2.5 Longitudinal Data Subset 48

3 Heart Rate Variability **51**
3.1 Introduction and Background 51
3.2 Short-Time Fourier Transform 53
3.2.1 The STFT Window 56
3.2.2 Hamming Window 57
3.3 Spectral Components 57
3.4 Defining the LF and HF Bands 59
3.5 A Combined Approach: Non-stationary Peak Measures 62
3.5.1 Pre-Processing 64
3.5.2 Power Spectral Density Calculation 66
3.5.3 Window Specifications 66
3.5.4 Non-Stationary Peak Determination 68
3.6 Application to Haemodialysis Data 72
3.6.1 Comparison of the Fixed-Band and Non-Stationary Peak Measures 73
3.6.2 Stable Sessions 76
3.6.3 Sessions with a Compensatory Response 78
3.6.4 Sessions without a Compensatory Response 79
3.6.5 Effects of Interventions 85
3.7 Patient-Specific Analysis 89
3.8 Discussion 90

4 High Dimensional Visualisation **98**
4.1 Methods for High-Dimensional Visualisation 99
4.1.1 Kohonen Self-Organizing Maps 99
4.1.2 Generative Topographic Mapping 100

4.1.3 Sammon Mapping 100
4.1.4 Neuroscale 101
4.2 Visualisation of Patient Data 104
4.2.1 Extrapolation Properties 104
4.2.2 Initialisation and Training 105
4.2.3 Quantisation 107
4.3 Definition of Normality in Dialysis Patients 107
4.4 Population-Based Model of Normality 110
4.4.1 Effects of Varying Individual Vital Signs 115
4.4.2 Normal and Abnormal Vital Sign Data Projections 120
4.5 Alternative Population-Based Models 126
4.5.1 Relative Changes in Vital Signs 126
4.5.2 Stratification based on Autonomic Function 133
4.5.3 Subset Statistics 135
4.5.4 Training 136
4.5.5 Results 138
4.5.6 Conclusions 142

5 Patient-Specific Models of Normality 146
5.1 Separation of Normalisation and Training Data 147
5.2 Patient-Specific Training Data 150
5.3 Projections of non-IDH Sessions 152
5.4 Projections of IDH Sessions 154
5.5 Vital Sign Trajectories in Time 157
5.5.1 Normal Sessions 159
5.5.2 IDH Session 163
5.5.3 Dialysis Onset 164
5.6 Conclusions 164

6 Patient-Specific Data Fusion Models for HD Patients 166
6.1 A Patient-Specific Model for Detecting Patient Deterioration . . . 167
6.2 Creating the Model of Normality 167
6.2.1 Data Preprocessing 167

6.2.2 Parzen Windows . . . 168
6.2.3 Novelty Calculation . . . 170
6.3 Application to Dialysis Data . . . 171
6.3.1 Persistence Criterion . . . 173
6.3.2 OB 10 . . . 173
6.3.3 OB 19 . . . 176
6.3.4 OB 20 . . . 179
6.3.5 Conclusions . . . 181
6.4 PSI with zero-order held MAP . . . 183
6.4.1 OB 10 . . . 183
6.4.2 OB 19 . . . 184
6.4.3 OB 20 . . . 186
6.5 Performance Evaluation . . . 187

7 Conclusion **193**
7.1 Summary of Results . . . 193
7.2 Future Work . . . 197
7.2.1 Continuous Blood Pressure Monitoring . . . 197
7.2.2 Additional Parameters . . . 199

A Dataset Overview **202**

B Further Neuroscale Examples **207**
B.1 Examples of IDH and non-IDH session Neuroscale Projections . . 207
B.2 Training Data for Patients OB7 and OB30. . . . 207
B.3 Test Data for Patients OB7 and OB30 . . . 214
B.4 Examples of Patient-Specific Neuroscale Projections . . . 214

List of Figures

1.1 A schematic of the blood and dialysate circuits set up during HD. . 5
1.2 A typical ECG waveform. 10
1.3 Components of a typical PPG waveform. 15

2.1 Histogram of number of sessions recorded for each patient. 41
2.2 Sample time plots of the vital sign parameters. 47
2.3 Histograms of vital signs from IDH and no-IDH sessions. 49

3.1 A representative spectrogram with the minute-by minute power spectra for an entire HD session. 55
3.2 Distribution of breathing rates in HD study population B 60
3.3 Exemplar one-minute ECG segment and resulting power spectrum 68
3.4 One-minute median respiratory frequencies alongside the power spectral density. 70
3.5 Power spectrum for a patient whose respiratory frequency is higher than the upper limit of the fixed HF band. 71
3.6 LF and HF power computed from the relevant peaks only. 71
3.7 Example frequency spectrum of a one-minute window of RR intervals with no subsidiary peaks. 75
3.8 Example frequency spectrum with additional peak in LF range. . . 76
3.9 OB41: An example of a patient with stable blood pressure and heart rate throughout the session. 77
3.10 OB32(a): An example of a patient with stable vital signs throughout the session. 78
3.11 OB32(b): An example of a characteristic response to a dialysis-induced decrease in blood pressure. 80

3.12 OB34: An example of a lack of rise in heart rate in response to a decrease in blood pressure. 81
3.13 OB10(a): An example of a decrease in the mean value of the LF measure in conjunction with a decrease in blood pressure. 83
3.14 OB10(b): Example of sharp decreases in the LF measure prior to decreases in mean arterial blood pressure greater than 20%. 84
3.15 Effects of a permanent reduction in the UF rate on the LF measure. 86
3.16 Effects of a temporary reduction in the UF rate on the LF measure. 87
3.17 Average deviation in the LF measure from 100% for sessions from three patients. 91

4.1 Schematic of the Neuroscale algorithm. 103
4.2 Vital sign histograms of the training data compared with the IDH and non-IDH subsets. 112
4.3 Two examples of a heart rate time series, its median-filtered version and a fitted parabola. 114
4.4 Neuroscale projection of the population-based training data. . . . 116
4.5 Varying individual parameters across the population-based Neuroscale model. 117
4.6 Parallel decrease in HR and MAP from the start of the session to the heart rate nadir. 119
4.7 Vital signs from a non-IDH session for patient OB 18 projected into two dimensions using the RBF network trained with a population-based definition of normality. 123
4.8 Vital signs from an IDH session for patient OB 18 projected into two dimensions. 124
4.9 OB18: Vital sign histograms of two sessions 125
4.10 Histograms of reference values at start of sessions. 128
4.11 Neuroscale projections of changes in vital signs in a normal session for OB 18. 131
4.12 Neuroscale projections of changes in vital signs in an IDH session for OB18. 132

4.13 Vital sign histograms of data subsets, separated by autonomic function . . . 137
4.14 Neuroscale projections of a non-IDH session from patient OB 78 by autonomic function. . . . 140
4.15 Neuroscale projections of a non-IDH session from patient OB 70 separated by autonomic function. . . . 141
4.16 Neuroscale projections of an IDH session for patient OB 18, separated by autonomic function. . . . 143
4.17 Neuroscale projections of an IDH session for patient OB 66, separated by autonomic function. . . . 144

5.1 Patient-specific training data for two patients projected onto the population-based Neuroscale map. . . . 151
5.2 Illustration of the seven patient-specific training data sessions on the patient-specific map. . . . 152
5.3 Time series plots of the vital signs from two sessions giving rise to outliers when projected onto the Neuroscale map. . . . 154
5.4 Projection of vital sign data from seven non-IDH sessions onto the patient-specific map for patient OB7. . . . 155
5.5 Projection of vital sign data from seven non-IDH sessions onto the patient-specific map for patient OB30. . . . 156
5.6 Three IDH sessions for patient OB19 projected onto a patient-specific Neuroscale map. . . . 158
5.7 Neuroscale projection of a non-IDH session for patient OB19. . . . 159
5.8 Three IDH sessions for patient OB10 projected onto a patient-specific Neuroscale map. . . . 160
5.9 Neuroscale projection of a non-IDH session for patient OB10. . . . 161
5.10 Three IDH sessions for patient OB20 projected onto a patient-specific Neuroscale map. . . . 162
5.11 Neuroscale projection of a non-IDH session for patient OB20. . . . 163

6.1 OB10: Vital signs and PSI for one non-IDH and three IDH sessions. 174
6.2 OB19: Vital signs and PSI for one non-IDH and three IDH sessions. 177

6.3 OB20: Vital signs and PSI for one non-IDH session and three IDH sessions. 181
6.4 OB10: Vital signs and PSI with zero-order held MAP data. 185
6.5 OB19: Vital signs and PSI with zero-order held MAP data. 187
6.6 OB20: Vital signs and PSI with zero-order held MAP data. 188

B.1 Examples of non-IDH sessions projected onto the population-trained Neuroscale map. 209
B.3 Examples of vital sign data from IDH sessions projected onto the population-trained Neuroscale map. 211
B.4 Training data for patients OB7 and OB30. 214
B.5 Time series plots of the test data for OB7. 217
B.6 Examples of vital sign data from IDH sessions projected onto patient-specific Neuroscale maps. 218
B.7 Examples of vital sign data from non-IDH sessions projected onto patient-specific Neuroscale maps. 220

List of Tables

1.1 Definitions of IDH found in the literature. 19

2.1 B. Braun machine parameters recorded during dialysis 42
2.2 Age distribution for the patients in the recorded dataset. 44
2.3 Patient demographics for the study population. 44
2.4 Artefact rejection limits 45
2.5 Means and standard deviations of IDH and non-IDH sessions. . . 48
2.6 Subset of seven patients for longitudinal analysis. 50

3.1 Bandpass filter specifications. 67

4.1 Means and standard deviations of the population-based training dataset. 111
4.2 Means and standard deviations of the relative changes in each parameter with respect to its 5-minute median value at the start of each respective session. 130
4.3 Means and standard deviations of the data subsets, separated by autonomic function. 135

5.1 Means and standard deviations of the heart rate training data and of the heart rate over entire sessions which include the training data. 148

6.1 Percentiles of the PSI for the training sessions of OB10, OB19 and OB20. 172
6.2 Means and standard deviations of the vital signs for OB10, OB19 and OB20 used for model normalisation. 175
6.3 Timings of alert generation using different persistence criteria. . . 180

6.4 Comparison of alert timings for the PSI calculated with interpolated and zero-held MAP values. 183
6.5 Confusion Matrix 189
6.6 Confusion matrix for alert generation by the PSI based on interpolated MAP values. 190
6.7 Confusion matrix for alert generation by the PSI based on zero-hold MAP values. 190

A.1 Overview of the sessions recorded for each patient in the study. . . 206

List of Acronyms

SBP Systolic Blood Pressure

DBP Diastolic Blood Pressure

PRSA Phase Rectified Signal Average

PDF Probability Density Function

PTT Pulse Transit Time

pmp per million population

RRT Renal Replacement Therapy

GMM Gaussian Mixture Model

MAP Mean Arterial Pressure

IDH Intra-Dialytic Hypotension

HR Heart Rate

HRV Heart Rate Variability

ECG Electrocardiogram

HD Haemodialysis

SA Sino-Atrial

AV Atrio-Ventricular

RSA Respiratory Sinus Arrhythmia

SOM Self-Organizing Map

FFT Fast Fourier Transform

PSI Patient Status Index

GTM Generative Topographic Mapping

STFT Short-Term Fourier Transform

SPWVD Smoothed Pseudo-Wigner-Ville Distribution

CWT Continuous Wavelet Transform

AR Auto-Regressive

LF Low-Frequency

HF High-Frequency

SDNN Standard Deviation of Normal-to-Normal RR intervals

CRF Chronic Renal Failure

ESRD End-Stage Renal Disease

EM Expectation-Maximisation

RBF Radial Basis Function

PPG Photoplethysmograph

UF Ultrafiltration

Chapter 1

Introduction

Chronic Renal Failure (CRF) is a condition characterised by progressive and irreversible loss of kidney function, and is fatal if left untreated. One of the key functions of the kidneys includes blood pressure regulation through water and electrolyte balancing. As such, the kidneys serve as the body's filter for excess fluid. End-Stage Renal Disease (ESRD) describes the final stage of CRF in which the kidneys are no longer able to perform the functions necessary in daily life. This level of failure necessitates Renal Replacement Therapy (RRT), the most common form of which is Haemodialysis (HD).

The main objectives of HD are threefold: to remove excess fluid from the body, to restore electrolyte balance, and to clear the blood of waste products such as urea. It therefore effectively mimics the functions of the kidneys. However, while normal kidney function is continuous, most HD sessions take place over the course of 4 hours, three times per week. As three sessions cannot be evenly spaced within a week, one session will always follow two HD-free days, while the other two sessions will follow one HD-free day. This common schedule for dialysis

dates back to the late 1960s [1] and is a compromise between physiological and practical considerations such as logistics, quality of life and costs.

Each HD session impacts human physiology in an extreme way: the concentrations of ions and solutes are changed significantly, while up to three litres of fluid are drawn from the body. Given that on average 57% of the human body consists of water [2], the removal of 3 litres of fluid in an average person weighing 70 kg represents a change of 7.5% in body fluid. Despite physiological compensation mechanisms in the body to respond to such extreme changes, complications frequently arise in HD patients as a result of these fluctuations in body water and solute content [3]. Most notably, these include an abrupt fall in blood pressure during the HD session, in conjunction with symptoms such as cramps, nausea and vomiting [4]. This is known as Intra-Dialytic Hypotension (IDH) and is the most common complication during HD, occurring on average in 20% to 30% of all dialysis sessions [5]. In comparison, a recent monthly audit at the Oxford Renal Unit showed that of the 1051 sessions recorded in that month, nursing interventions were required in 136 sessions (13%). Of these event sessions, 57% were symptomatic. 53 of the 93 patients receiving dialysis treatment that month experienced at least one event (57%).

In response to an IDH event, fluid is quickly replaced in some patients in the form of saline infusions. In other cases, fluid removal is paused intermittently to allow for the patient to recover, or terminated prematurely. Either of these interventions counteract the main purpose of HD, which is the removal of excess fluid from the body. They can therefore give rise to fluid overload in the long term due to inadequate clearance [6], which in turn increases the risk of mortality [7]. More severe episodes of IDH can result in damage to the heart tissue [8, 9], which

can ultimately lead to heart failure [10, 11]. Compared to the general population, the percentage of sudden cardiovascular deaths is significantly higher in patients treated with dialysis [12].

1.1 Patient Demographics

The UK Renal Registry in 2011 [13] reported 50,965 adult patients receiving RRT in the UK, which encompasses various types of treatment including HD as well as other forms of dialysis and kidney transplants. Approximately 75% of all RRT patients in the UK are receiving HD treatment. An increase in the prevalence of ESRD from 523 per million population (pmp) in 2000 to 832 pmp in 2010, as well as in the incidence rate of new patients from 95 pmp in 2001 to 107 pmp in 2010 was also reported. The median age of HD patients in 2010 was 66.3 years, while in 2000 it had been 62 years. Overall, the prevalence for males is higher than that for females across all age groups. For both males and females, the peak prevalence in 2010 was in the 75- to 79-year age group, but the male prevalence in this age group is almost double that of female prevalence at 2,765 pmp and 1,406 pmp respectively. In comparison, the peak prevalence in 2000 was in the 55- to 64-year age group.

While there has been a general increase in prevalence of ESRD patients over the past 10 years, the increase in prevalence amongst the elderly has been disproportionate with respect to this trend. This increase can be mainly attributed to an ageing population and a general shift in demographics, but also to technological advances which allow for treatment of patients previously deemed too fragile for treatment [14, 15]. Crucially for treatment outcome however, it should be noted

that in elderly patients often other conditions are present in addition to ESRD, which, in conjunction with the more fragile nature of elderly patient physiology, could give rise to more frequent complications during HD. As a result, there is an increasing need for effective therapy which minimizes intradialytical symptoms.

Cardiac death is the single greatest cause of mortality in ESRD patients on dialysis. In the United States, the annual death rate for dialysis patients in 2004 was 230 deaths per 1000 patient years [16, 17]. Of these, 64% were attributed to cardiac arrest [16]. While patients with chronic kidney failure have generally been found to be at higher risk of cardiac arrest as compared to healthy patients, the risk increases with disease progression. In ESRD patients receiving HD, the risk of cardiac arrest after 3 years of treatment was found to be 24% [16].

1.2 Haemodialysis as Renal Replacement Therapy

Dialysis refers to a procedure by which various substances in solution are separated based on their different diffusion properties across a semipermeable membrane [18]. In HD, blood is drawn from the patient's body and passed through a dialyser where it is cleared by a countercurrent of dialysate fluid across a semipermeable membrane. The clean blood is subsequently passed back to the patient's bloodstream. A schematic of the blood circuit set up during HD, as well as of the countercurrent dialysate circuit, is shown in Figure 1.1. The elements of these two circuits are now considered in turn.

The blood circuit starts at the arterial vascular access point, where blood is drawn from the patient. For this purpose, a negative pressure gradient is main-

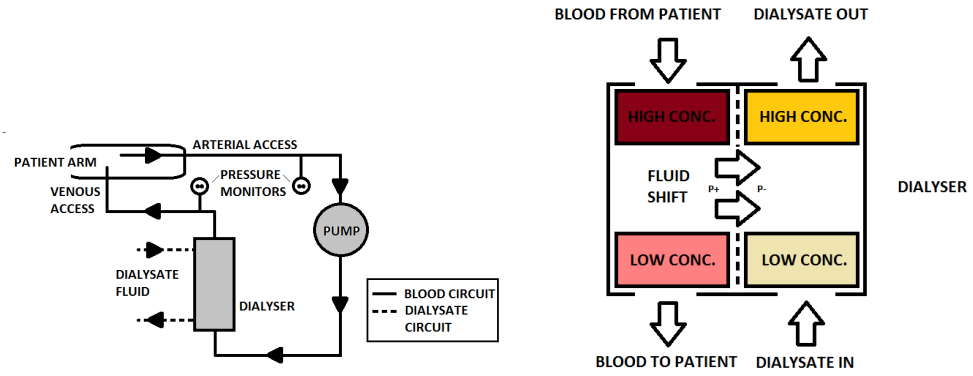


Figure 1.1 – A schematic of the blood and dialysate circuits set up during HD, illustrating the fluid and solute removal process. The countercurrent flow between the blood and dialysate illustrated in (b) ensures that the concentration of designated solutes in the blood is reduced when passing through the dialyser, while their concentration in the dialysate increases.

tained between the access point and the blood pump by a pressure monitor. The blood is then pumped through the dialyser at a variable speed between 200 to 600 mL/min, depending on patient needs. The dialyser consists of thousands of hollow fibres, through which the blood is pumped, and is the site of blood clearance as well as fluid removal. Inside the dialyser, the fibres are surrounded by dialysate fluid which forms part of the dialysate circuit described below. The fibre walls serve as semipermeable membranes, and with this setup the surface of exposure between the blood and the dialysate fluid is maximised. Further, a countercurrent flow is set up between the blood and the dialysate fluid as shown in Figure 1.1(b) in order to maintain as high a concentration gradient as possible across the membrane. This serves to maximise the rate of solute removal.

The dialysate fluid consists of a sterilized solution of different ions which is composed such that urea and other waste products, as well as minerals such as

potassium and phosphate, diffuse from the blood into the dialysate fluid by virtue of the concentration gradient across the membrane. On the other hand, dialysate concentrations of sodium and chloride for instance, which are not meant to be cleared from the blood, are similar to those in the blood to prevent diffusion. Hence the composition of dialysate fluid allows for selective removal of ions from the blood. While one element of the dialysate circuit consists of fresh dialysate fluid which is passed through the dialyser at a variable rate (typically between 250 and 450 ml/min) to effect solute removal from the blood, excess fluid is also removed from the blood through this circuit. For this purpose, a hydrostatic pressure gradient is set up across the membranes within the dialyser (as illustrated in Figure 1.1(b)), the magnitude of which determines the rate of fluid removal. This is known as Ultrafiltration (UF), and the UF rate can be altered over the course of a dialysis session by varying the pressure gradient across the membrane.

Having passed through the dialyser, the cleared blood is passed back into the patient's bloodstream through the venous vascular access point. Again, the pressure is monitored at this point to ensure that its gradient remains positive.

1.2.1 Dry Weight Estimation

As one of the main objectives of HD is to restore correct fluid balance in the patient, the amount of fluid to be removed through UF is estimated at the beginning of the session and the UF rate is set accordingly. For this purpose, the concept of dry weight has established itself in clinical practice as the patient's prescribed target weight, which should ideally be achieved by the end of each session. The volume of fluid that should be removed during a given HD session is then cal-

culated as the difference between the patient's actual weight prior to the session and the patient's prescribed dry weight. By definition, the dry weight reflects the lowest weight a patient can tolerate without symptoms or hypotension [4], and which allows for the patient to maintain blood pressure within normal limits until the next session [19]. Inaccurate estimation of dry weight leads to inappropriate volumes of fluid being removed. Overestimation of dry weight leads to chronic volume overload due to inadequate fluid removal, while underestimation of dry weight results in too much fluid being removed.

Most frequently, dry weight is determined by trial and error and in estimating it, clinicians take into account various aspects including the patient's medical history and blood test results as well as their vital signs. It as been suggested that a patient's dry weight be reviewed at least every two weeks for the following six dialysis sessions [20], but in some UK units it is formally reviewed only once every three months. Instead, ad-hoc adjustments are frequently made at the beginning of each session if deemed necessary [21]. While methods, such as bioimpedance spectroscopy [22], in which the total fluid volume in the body is extrapolated from its impedance to a high-frequency current, have been proposed for determining dry weight, no method has established itself in clinical practice for practical reasons.

Frequent review of dry weight is necessary in order to account for changes in non-water body mass, which can result from a variety of factors including illness or depression [23]. As a further complication to the accurate adjustment of dry weight, patients can remain well above their true dry weight without presenting signs of hypertension until up to 12 hours after dialysis [24]. As patients who are well above their dry weight rarely experience acute intradialytic symptoms,

there is little immediate incentive for dry weight adjustment. At present, the standard approach to dry weight determination can be inadequate in some cases and any supplementary information to support the decision-making process would be advantageous.

1.3 Current Monitoring during Haemodialysis

In routine practice, patients are not currently monitored continuously during HD on a regular basis. In the absence of any irregularities, routine recordings consist of weight and blood pressure measurements at the beginning and end of the session, as well as intermittent measurements of tympanic temperature. Patients who report symptoms in advance of or during a session, as well as patients who are known to frequently experience IDH, are monitored more closely. In these patients, blood pressure is measured every 30 minutes. After serious IDH events, monitoring is further supplemented by Electrocardiogram (ECG) and Photoplethysmograph (PPG) readings. With the exception of the haematocrit (a dialysis-specific parameter defined in Section 1.3.6), no parameter is monitored continuously during routine dialysis. In the following, all relevant parameters are described, including details of their acquisition.

1.3.1 Electrocardiogram

The ECG is a non-invasive recording on the skin surface of the differential voltage that arises from depolarizing and polarizing cardiac cells, which are triggered by specific pacemaker nodes during each heartbeat. The Sino-Atrial (SA) node is the heart's primary pacemaker site, spontaneously generating electrical impulses

that cause local cardiac cells to depolarize. These impulses spread from the SA node, which is located at the top of the right atrium, through the atria to the Atrio-Ventricular (AV) node, which is also located in the right atrium. The impulse continues travelling through the bundle of His, which consists of cardiac muscle fibers and constitutes the only conductive path between the atria and the ventricles. The bundle of His consists of right and left branches, which propagate the depolarisation wave into the right and left ventricles, respectively. Following ventricular depolarisation is a period known as the refractory period, in which the ventricular myocardium repolarises and no further conduction takes place. This cardiac cycle repeats itself at every heartbeat.

Surface electrodes can be placed on the limbs, arms and chest in different configurations to obtain the ECG. In routine analysis, four electrodes are placed on each of the limbs (right and left arms and legs), with one electrode, usually the right leg, being the reference electrode [25]. The voltages between the remaining electrodes give 3 leads, with the most commonly used one being Lead-II from the right arm to the left leg as it provides the most clinically relevant information. A typical Lead-II ECG is shown in Fig. 1.2. With each heartbeat, three wave structures are apparent: the P-wave, the QRS complex and the T-wave.

P-wave The P-wave represents the depolarisation of the atria, originating from the SA node. This depolarisation takes place prior to the main contraction of the ventricles and therefore is represented by only a small change in voltage. The period of zero voltage following the P-wave, also known as the PQ junction, reflects the time it takes for the impulse to travel through the atrioventricular node to the

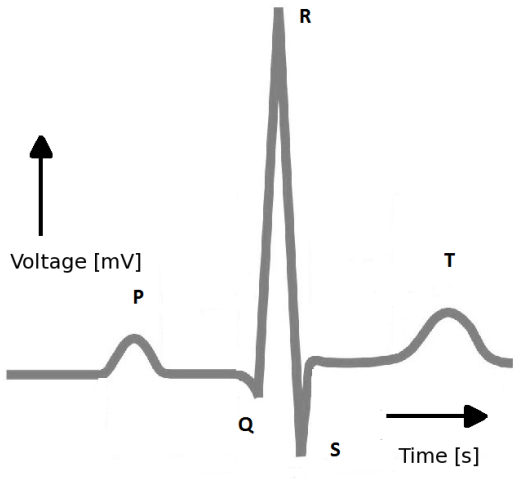


Figure 1.2 – A typical ECG waveform.

ventricles.

QRS complex Ventricular depolarisation is reflected in the QRS complex and results in ventricular contraction. This phase of the cardiac cycle is known as systole. As the ventricles have a much greater muscle mass than the atria, this wave of depolarisation is much larger than the P-wave. Atrial repolarisation occurs at the same time as ventricular depolarisation, but it cannot be observed on the ECG since the repolarisation voltage is much lower than the voltage of ventricular depolarisation.

T-wave Ventricular repolarisation is represented by the T-wave. This corresponds to diastole, the refilling of the heart with blood.

The prevalent ECG morphology shown in Fig. 1.2 can be affected by abnormalities such as ectopic beats. These arise as a result of blocks in the regular

conduction pathway described above. In the absence of impulses from the SA node, the AV node takes over the pacemaker function, but at a lower frequency than the SA node (40-60 beats/minute (bpm) vs. 100-120 bpm). If the AV node fails to act as a pacemaker or the conduction pathway to the ventricles is blocked, the cells of the ventricles have an inherent depolarisation frequency of 15-40 bpm which allows them to act as pacemakers at this much lower frequency as a last resort. These alternate rhythms are known as escape rhythms and exhibit very different ECG morphologies.

The time-averaged Heart Rate (HR) is most commonly measured as the number of contractions over a 60-second period, while the instantaneous HR is computed as $\frac{60}{RR}$, where RR is the time interval between two successive R-peaks in seconds.

1.3.2 Heart Rate Variability Measures

HRV in its most basic form can be defined as the variation in time of the intervals between successive heartbeats as recorded from the ECG. In HD patients, different measures of HRV have been shown by many authors to be significant predictors of cardiac death independently of other known risks in the patient population. Although HRV is known to be decreased in patients with cardiac disorders, it has been shown that HD affects a variety of HRV measures independently of any cardiac conditions [26]. From the ECG, different measures of HRV can be derived, with the two main approaches in this context being time-domain and frequency-domain methods.

Ideally, the fiducial point for each heartbeat should be the P-wave, as it cor-

responds with the activation of the SA node which marks the beginning of the cardiac cycle. However, the P-wave is usually a low-amplitude waveform and therefore not easily detectable. Conventionally, the R-peak is taken as the fiducial point instead as it is easily identifiable by peak detection [27], and because variations in the PR-interval are correlated with RR-interval modulations in the absence of conduction disorders [28]. In establishing the RR-interval sequence from the ECG, only normal heartbeats are used, with those resulting from cardiac arrhythmias having to be excluded.

1.3.2.1 Time-Domain Methods

Time-domain measures of heart rate variability are based on statistical methods that are applied to the RR-interval sequence. Simple statistical measures include the mean RR interval or the difference between the longest and shortest RR interval in a given ECG sequence. More complex statistical measures can be derived from either the sequence of RR intervals or from the time difference between successive RR intervals for a given time period [29]. The standard deviation of the RR intervals (SDNN) for instance, is given by the square root of the variance of RR intervals in the time period under consideration and is the most commonly used time-domain index of HRV. The SDNN is dependent on the length of the RR interval sequence that is being analysed, and is commonly calculated over a 5-minute period.

As a measure of HRV, the SDNN cannot distinguish between variations originating from either periodic or random sources and its computation can also be influenced by the lack of artefact recognition [30]. Variations of the SDNN include the SDANN which is the standard deviation of the average of RR intervals

calculated over short periods, or a measure of the average of 5-minute SDNNs calculated over a longer period of time. Other time-domain measures include the peak-to-trough difference in heart rate per respiratory cycle [31] or the number of adjacent pairs of normal beats that differ by more than 50 ms, denoted as NN50 [32].

Peak-to-trough methods as described in [31] quantify periodic deviations from the baseline HR due to Respiratory Sinus Arrhythmia (RSA) (as defined in Section 1.3.5). At high respiratory frequencies however, the accuracy of these methods is diminished as the RSA amplitude is reduced [33].

1.3.2.2 Frequency-Domain Methods

HRV calculations in the frequency domain are based on power spectral density analysis to determine the distribution of power in the RR interval sequence as a function of frequency. Both parametric methods, such as autoregressive techniques [34], and non-parametric methods based on linear filtering, such as the Fourier transform [35] and the wavelet transform [36], are commonly used.

Both approaches have their advantages and disadvantages. While the Fourier transform does not make any assumption as to the structure of the underlying process, it does require strict periodicity and stationarity of the data as well as an *a priori* selection of the frequency range and the number of frequencies represented within that range. Auto-Regressive (AR) methods, while also requiring periodicity and stationarity of the data, assume a certain distribution for the underlying process, but do not require any *a priori* knowledge of the range of frequencies represented. As with time-domain methods, the results of spectral analysis are partially dependent on the length of recording analysed. Spectral analysis is pri-

marily performed on short-term recordings of no more than 5 minutes [37], but some work on longer-term recordings of up to 24 hours has also been carried out [38]. However, due to the non-stationary nature of the changes in RR interval throughout the day [39], computing the power spectrum over long periods effectively averages the power of the frequency components over the entire recording and hence detailed effects on a shorter timescale cannot be followed.

1.3.3 Blood Oxygen Saturation

98.5% of blood oxygen is bound to haemoglobin in red blood cells. Arterial blood oxygen saturation (SaO_2) gives an indication as to the percentage of haemoglobin binding sites occupied by oxygen in arterial blood:

$$SaO_2 = \frac{HbO_2}{HbO_2 + Hb}$$

where HbO_2 is oxygen-bound haemoglobin and Hb is free haemoglobin. SpO_2 is the measurement of SaO_2 using pulse oximetry, giving rise to the PPG signal. Light of two different wavelengths (red and infrared, 660 nm and 905 nm respectively) from two light-emitting diodes (LEDs) is transmitted through a translucent part of the body, usually the earlobe or fingertip, to a receiving photodiode. Oxygen-bound haemoglobin and its de-oxygenated form have different absorption coefficients at these wavelengths, and therefore the ratio of oxygenated to deoxygenated haemoglobin can be established from the ratios of light attenuation at the two wavelengths. With every heartbeat there is an inflow of arterial blood into the vessels. The DC component of light attenuation is caused by surrounding tissues such as bone and skin. It can be removed to give the cardiac-synchronous

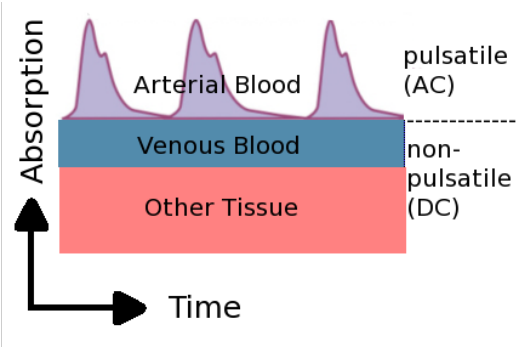


Figure 1.3 – Components of a typical PPG waveform. The pulsatile component results from the inflow of arterial blood with every heartbeat. The non-pulsatile component represents the absorption by venous blood and other tissue such as bone and skin.

AC component which represents the light attenuation due to the arterial component only [40]. The various components of the PPG signal are shown in Fig. 1.3. The pulsatile variation in light attenuation also gives an indication of the heart rate, commonly referred to as the pulse rate.

1.3.4 Blood Pressure

With every heartbeat, pressure is exerted on the blood, and ultimately on the walls of the blood vessels, to circulate it through the body. A maximum pressure is reached when the ventricles contract (Systolic Blood Pressure (SBP)) and a minimum occurs upon relaxation of the heart muscles (Diastolic Blood Pressure (DBP)). Blood pressure is commonly measured non-invasively, using a sphygmomanometer to observe oscillations in blood flow [2]. A cuff is inflated to a pressure above the SBP and subsequently lowered to a pressure below the DBP over a period of 30 seconds. The cuff pressure is constant when either it is above SBP (which occludes blood flow) or when it is below DBP (which allows for free

blood flow). Between these two pressures, when blood flow is restricted, oscillations in the cuff pressure can be measured and used to compute the SBP and DBP values. Oscillations peak when the cuff pressure equals the Mean Arterial Pressure (MAP), which is a weighted average of diastolic and systolic pressures such that

$$MAP = \frac{2}{3}DBP + \frac{1}{3}SBP.$$

1.3.5 Breathing Rate

The breathing rate, also known as the respiratory rate, is a measure of the number of breaths taken per minute and various methods exist to estimate this parameter. The most simple is to visually inspect the motion of the chest wall and count the number of breaths, or alternatively to use a stethoscope to listen for breathing sounds. Given the relatively low frequency of respiration, manual breathing rate measurements should ideally be done over the course of one minute. More sophisticated methods include impedance pneumography [41] and the use of strain gauges to estimate the breathing rate. In impedance pneumography electrodes are placed on the chest and are used to inject a high-frequency current into the body. The resulting voltage across the thorax is then measured, from which the chest's electrical impedance can be calculated. As this varies during each breathing cycle, the resulting signal can be used to obtain the breathing rate. Strain gauges can be integrated into respiratory belts that are worn tightly across the chest by the patient. Distention of the belt due to inspiration allows for the breathing rate to be calculated.

In addition to these direct measurements, indirect methods have been sug-

gested for estimating the breathing rate, most notably based on a phenomenon known as RSA. Respiration in known to cyclically affect the heart rate, by causing it to increase upon inspiration and to decrease upon expiration [42, 43]. This change in heart rate is mediated by the autonomic nervous system through vagal nerve stimulation of the heart. By measuring the time difference between successive R-peaks in the ECG over time, a signal that is synchronous with respiration can be obtained. Based on this principle, various methods have been devised for obtaining the breathing rate from the ECG [42, 44]. Further, it has been suggested that RSA can be used to measure autonomic nervous system function [45]. This is considered in more detail in Chapter 3.

1.3.6 Haematocrit

Haematocrit gives an indication of the fraction of blood volume that is occupied by red blood cells [2]. Reference values for adult males are 42-52%, and 36-48% for adult females [25]. In dialysis patients, haematocrit is measured continuously and non-invasively during a session using the Crit-Line Instrument (HemaMetrics, formerly In-Line Diagnostics). The Crit-Line monitor is based on a similar principle to the pulse oximeter: blood is illuminated by an LED once it has been drawn from the body and based on both the absorption properties of haemoglobin and the scattering prperties of red blood cells passing through the blood chamber, the concentration of red blood cells can be determined. Percentage changes in the blood volume (BV) can then be approximated as

$$\Delta BV = \left(\frac{C_0}{C_t} - 1\right) x100$$

where C_0 and C_t are the concentrations of red blood cells at the start of HD and at a given time t during HD, respectively.

Beyond its application in estimating blood volume, it has been suggested that the slope of the haematocrit with respect to time could also be used to assess dry weight [46], although the application of this method in a recent clinical study resulted in a higher rate of patient hospitalisation [47].

1.4 Definitions of IDH

The definition of IDH varies widely in the literature. While some definitions use absolute value thresholds for either systolic or mean arterial blood pressure to identify episodes of IDH [48, 49, 50], others focus on abrupt changes in blood pressure [51]. Definitions consisting of a combination of absolute thresholds and changes in blood pressure have also been suggested [52, 53, 54]. Other indicators such as the need for intervention or the reporting of symptoms by the patient have also been included in some definitions of IDH [53, 54]. Table 1.1 shows a selection of different definitions commonly found in the literature. The lack of consensus over a global definition for IDH highlights the difficulty in correctly assessing IDH as well as the potential dependency of the definition on the physiological status of a given patient.

For the purpose of this thesis, a definition that combines those of Hoeben et al [52] and Poldermans et al [53] is adopted, as suggested by Meredith et al [57]. This definition accounts for both the abruptness of the fall in systolic blood pressure when considering IDH, as well as including an absolute level for the

Author	Definition	Symptoms/Interventions
Hung et al [48]	systolic BP < 90 mmHg	none
Barnas et al [49]	MAP < 65 mmHg	none
Cai et al [55]	systolic BP < 90 mmHg or fall ≥ 30 mmHg	none
Santoro et al [56]	systolic BP < 100 mmHg	
	fall ≥ 25 mmHg	symptoms and intervention
Schmidt et al [50]	systolic BP < 90 mmHg	none
Mancini et al [54]	systolic BP < 90 mmHg	none
	fall $> 25\%$	symptoms
Mizumaza et al [51]	systolic BP fall > 50 mmHg	symptoms
Hoeben et al [52]	systolic BP fall > 20 mmHg to below 100 mmHg	symptoms
Poldermans et al [53]	systolic BP fall $> 25\%$ or to below 100 mmHg	symptoms

Table 1.1 – Definitions of IDH found in the literature.

blood pressure below which exposure to end-organ damage due to a lack of blood perfusion could occur. For the analysis in subsequent chapters, IDH is defined as a fall in systolic blood pressure of more than 20% with respect to the previous measurement in conjunction with a systolic blood pressure below 100 mmHg.

1.5 Causes of IDH

Multiple causes of IDH have been identified, their common feature being the inadequacy of physiological compensation mechanisms in dealing with the UF-induced decrease in blood volume, also known as hypovolaemia [58]. In this section, the different conditions giving rise to IDH are considered in turn.

1.5.1 Cardiac Refilling

The total blood plasma volume in an average human equals 2.7-3.0 litres. During a typical dialysis session, up to 3 litres of fluid are removed from the blood by UF, which corresponds to the entire human blood plasma volume. Nevertheless, the plasma volume typically only decreases by 10-20% [59, 60]. To maintain plasma volume, fluid is shifted from the extracellular compartment of the body, which comprises the fluid-filled space between the cells of the body, into the blood compartment.

As fluid is removed from the blood compartment by UF, the pressure in the extracellular compartment is increased, while the hydrostatic pressure in the blood compartment is reduced. This causes fluid to move into the bloodstream [61]. This process is termed cardiac refilling and its rate has been reported to be between 0.25 and 0.56 litres/hour [62], which is significantly lower than typical UF rates during

dialysis. The efficacy of this compensatory mechanism is largely dependent on the hydration status of the patient. In this context, an accurate estimation of the patient's dry weight, as described in Section 1.2.1, is key to maintaining a good balance between fluid removal by UF and the patient's refill capacity. By underestimating the dry weight, the capacity for this shifting in fluids is overestimated, and dehydration of the extracellular compartment alongside a low refill results. This has been found to be one of the main mechanisms leading to IDH [63].

1.5.2 Autonomic Nervous System

As previously indicated, the rate of cardiac refilling alone cannot sufficiently counteract the typical rates of fluid removal during dialysis. The resulting volume reduction in the blood compartment triggers other mechanisms to maintain blood pressure, primarily through a series of specialised neurons, known as baroreceptors, located at key locations in the vascular system. These neurons are sensitive to distention of the walls of the vasculature and can therefore monitor changes in blood pressure and relay these to the brain. This so-called baroreceptor reflex is in effect a negative feedback loop designed to maintain blood pressure levels. As blood pressure decreases, the activated reflex triggers various mechanisms governed by the autonomic nervous system. These include vasoconstriction, that is the constriction of the blood vessels, which leads to an increase in the overall resistance to blood flow throughout the vascular system [64], as well as the redistribution of blood volume to more critical regions of the circulation. Further, both the contractility of the heart as well as the heart rate are increased, with the latter mechanism being known as tachycardia. Both serve to increase the blood

pressure.

Impairment of the various mechanisms in this feedback loop has been found to be a cause of IDH [65]. In particular, a reduction in baroreceptor sensitivity, and hence a lack in the subsequent compensatory mechanisms described above, has been identified as a cause [65, 66], although contradicting evidence suggesting preservation of the baroreceptor reflex during IDH has also been presented [67]. This indicates that it may be necessary to distinguish between at least two different types of IDH based on the patient's baroreceptor sensitivity or autonomic function in general.

Regarding the responses activated by the baroreflex, a lack of the expected tachycardic response to fluid removal has been observed in dialysis patients [59], with some studies reporting a paradoxical decrease in heart rate preceding IDH episodes [68, 64]. Finally, as previously mentioned, more severe episodes of IDH can result in damage to the heart tissue [8, 9], which in turn can impact upon the ability of the heart to respond to fluid removal with increased contractility and an increase in heart rate [69].

1.6 Existing Methods for IDH Prevention

Generally, medical intervention and changes in UF rate only take place once symptoms of IDH have been observed, that is once the event has been triggered. Clearly, preventative intervention would be more favourable for patient well-being, especially in the elderly and in patients with existing cardiovascular problems, where the risk of mortality is already elevated. The chief difficulty in pre-empting IDH lies in detecting haemodynamic instabilities as they occur. Based on the

various physiological explanations for IDH presented in Section 1.5, different methods have been developed for detecting and predicting IDH. These are mostly based on the measurement of physiological parameters, including heart rate, blood pressure, SpO_2 and the haematocrit. In this section, the existing methods for IDH prediction are presented. However, despite efforts towards understanding the mechanisms giving rise to IDH, none of the proposed methods in this section have established themselves in clinical practice and no reduction in the incidence rate of IDH has been achieved over the past years.

1.6.1 Pulse Transit Time

The most intuitive approach to detecting hypotension would be to continuously monitor the blood pressure for decreasing trends and various attempts have been made at doing this non-invasively. These mainly revolve around the Pulse Transit Time (PTT), that is the time it takes for the pulse wave to travel from the aortic valve to a peripheral site in the body. The PTT is inversely proportional to blood pressure: a rise in blood pressure causes stiffening of the arterial walls, and therefore reduces the PTT. Conversely, a fall in blood pressure results in a decreased PTT as a result of distention in the arterial walls.

Using the ECG and PPG signals, the PTT can be measured as the time between the R-wave of the ECG (see Fig. 1.2), which corresponds to the opening of the aortic valve, and the peak in the PPG signal which represents the arrival of the pulse wave at the finger. As the speed at which the pulse wave travels is directly proportional to blood pressure, it has been suggested that the PTT could be used as a marker for beat-to-beat blood pressure. In a recent study by Payne et al. [70]

however, the PTT was not found to be reliable enough to be used as a surrogate marker for systolic blood pressure. Instead, variations in systolic blood pressure correlated well with the PTT measure once calibrated intermittently using actual oscillometric blood pressure readings, as suggested by Chen et al. [71]. Ahlstrom et al. [72] performed comparative measurements on dialysis and non-dialysis patients, and found the correlation between systolic blood pressure and the PTT to be $r = 0.8 \pm 0.06$ in the non-dialysis patients. While blood pressure changes were induced in the non-dialysis patients for the purpose of analysis, this was not done for the dialysis patients. However, the natural fluctuations in blood pressure during dialysis were not sufficient to merit statistical analysis. Nevertheless, in the small number of cases where large blood pressure variations did occur, these were reflected in the PTT. As such, no conclusive evidence has been provided in relation to the utility of the PTT in the context of HD.

1.6.2 Relative Changes in Blood Volume (Haematocrit)

In the absence of a continuous measure for blood pressure, changes in blood volume as estimated using the haematocrit (see Section 1.3.6), have been studied as an alternative marker for predicting IDH. In principle, manual adjustments can be made to the UF rate over the course of the session in response to changes in blood volume in an attempt to maintain optimal fluid balance. In particular, while there is no marker for the hydration status of the extracellular compartment, the haematocrit can be used to assess the adequacy of refill during dialysis subject to the assumptions that the red blood cell count of the plasma does not change over the course of dialysis and that there is a uniform mixing of red blood cells in the blood

[73]. While this method for estimating refill capacity is well established in clinical practice, it has recently been argued by Mitra et al [74] that the latter assumption does not hold during HD because of the redistribution of blood that takes place from peripheral regions of the circulation to more central parts, as described in Section 1.5.2. This redistribution causes the mixing of arterial and venous blood, which differ in oxygen content and haematocrit. This results in the reduction in blood volume during HD being underestimated, which could explain why incidence rates of IDH have not been reduced despite efforts in monitoring changes in blood volume. However, the study done by Mitra et al [74] only comprised 10 subjects, and results are yet to be confirmed in a more extensive study.

Despite the relationship between changes in blood volume and IDH being studied extensively since the advent of dialysis treatment [75], the majority of studies have not found any significant correlation between the two [63, 60, 76, 77]. Further, it has been shown by Andrulli et al [78] in a study of 123 HD patients that hypotensive episodes could not be predicted by the overall change in blood volume over the course of the session, although irregularities in the blood volume curve were found to be a useful predictor. Similarly, Mitra et al [76] suggest that a change in the trajectory of decrease in blood volume from exponential to linear is predictive of IDH. Finally, work by Steuer et al [79] suggests that intradialytic symptoms related to hypovolaemia occur at a patient-specific haematocrit, regardless of the UF rate. By changing the UF rate to keep the haematocrit below this patient-specific threshold, they achieved a twofold improvement in patient outcome. In practice, the threshold would have to be established for each patient at the beginning of their treatment. Changes in red blood cell count, resulting from for instance blood transfusions or haemorrhage, would require re-establishing the

patient's threshold.

1.6.3 Photoplethysmograph-based methods

Photoplethysmograph-based methods have also been suggested for predicting hypotension. Mancini et al. [80] argue that variations in blood oxygen saturation can anticipate IDH. It is suggested that a critical threshold exists for the standard deviation of the SpO_2, a transgression of which is predictive of hypotension. Lima et al. [81] make use of the PPG signal, as shown in Fig. 1.3, to estimate peripheral perfusion, which is a measure of blood flow to the peripheral vasculature. Given that the redistribution of blood flow by vasoconstriction is a commonly encountered in response to reductions in blood pressure [82] (as described in Section 1.5.2), a measure for peripheral perfusion would be a useful substitute for assessing decreases in blood pressure. Current methods for assessing peripheral perfusion include Laser Doppler flow [83] and the central-to-toe temperature difference. Lima et al. propose the use of the ratio between the pulsatile and non-pulsatile elements of the PPG waveform (as illustrated in Fig. 1.3) as an index of perfusion. By choosing this ratio over the actual PPG signal, they allow for usage in critically ill patients, where vasoconstriction results in a very low signal that is amplified by the pulse oximeter [40] and therefore requires normalisation. Although the focus of their work was on finding a suitable threshold in this measure for abnormal peripheral perfusion, this method could be useful in the dialysis setting for detecting changes in peripheral perfusion.

1.6.4 HRV

Fukuta et al. [84] found a substantial reduction in multiple long-term HRV measures extracted from 24h ECG recordings in chronic HD patients when compared with age-matched healthy controls. The measures under consideration included time-domain and frequency-domain measures, in particular the mean and standard deviation of the normal-to-normal intervals, as well as the power in the low frequency (LF, 0.03-0.15 Hz) and in the high-frequency (HF, 0.15-0.4 Hz) band of the Fast Fourier Transform (FFT) of a 24h sequence of RR-intervals extracted from the ECG recording (these spectral measures are described in more detail in Section 3.3). 120 HD patients were studied in total (age 61 $\pm$ 11, range 21-79), alongside their age-matched, healthy controls. In the 26-month follow-up period, 21 patients died, with 10 dying of cardiac causes and 11 dying of other causes. The study population was therefore divided into 3 groups, the surviving patients, the group of cardiac-related deaths and the group of non-cardiac related deaths. It was found that all HRV-measures extracted from the 24h ECG recordings were reduced in all three groups of the HD population compared with the control group. Amongst the HD patient groups, no significant differences in HRV measures could be identified between the survivors and the group of non-cardiac related deaths. By regression analysis, a decrease in the ratio of low-frequency to high-frequency power in the frequency spectrum of the RR-intervals was found to be a significant predictor of cardiac death in the HD population, while none of the HRV measures were found to be predictive of non-cardiac related deaths. Further, while it is often argued that diabetes is a possible cause of reduced heart rate variability in chronic HD patients, the analysis done by Fukuta et al [84] suggests that the predictive

value of a decrease in the HRV measures is independent of the effects of diabetes.

With respect to the use of spectral measures in long-term ECG measurements, it should be noted that the underlying physiological mechanisms modulating the heart rate and hence giving rise to the Low-Frequency (LF) and High-Frequency (HF) components are not stationary throughout a 24h period [39]. As already mentioned, extracting the HRV measures from long-term ECG measurements results in averaging of the various instances of heart rate modulation throughout the period under consideration. Given that it has been shown that both physical and mental activity can significantly influence the HRV measures [85], this may affect the results of long-term spectral analysis.

Yoshii et al [86] investigated how spectral HRV parameters changed during HD in association with cardiac output, stroke volume and total peripheral vascular resistance as assessed by a transthoracic electrical admittance cardiograph. Only hemodynamically stable patients were included in the study for this purpose. The patient cohort consisted of 19 patients who had not had any cardiovascular collapse in the two months preceding the study. The blood pressure, which was measured non-invasively every 10 minutes during the session, was maintained throughout all sessions studied for the patient cohort. A reduction in parasympathetic activity and an increase in sympathetic activity was observed during HD treatment, as interpreted from the conventional spectral LF and HF measures. Further, a marked increase in total peripheral vascular resistance was observed, together with a decrease in both cardiac output and stroke volume. These results suggest that the blood pressure is maintained during dialysis primarily through an increase in peripheral vascular resistance rather than an increase in cardiac output and stroke volume.

Vita et al [87] studied 30 chronic renal failure patients (22 men, 8 women) aged between 19 and 85 years (mean 58 $\pm$ 14 years) using both conventional measures of autonomic dysfunction such as the Valsalva test [88] and spectral measures of the power in the LF and HF bands derived from an AR model of a sequence of 500 RR-intervals. The results from the conventional tests for autonomic dysfunction, showed that 63% of the study population had moderate to severe dysfunction in both the sympathetic and parasympathetic branches of the nervous system, while the remainder exhibited normal autonomic function. Based on this distinction, spectral analysis was performed separately on the patients with and those without autonomic dysfunction. The results were compared with an age-matched control group. While there was no significant difference between the spectral measures between the two patient groups, the LF power was significantly lower in both patient groups when compared with the control group. It should be noted however, that the conventional tests used for evaluating autonomic function in ESRD patients have been found to lack reproducibility and sensitivity [89], calling into question the distinction between patients with and patients without autonomic dysfunction and possibly explaining the lack of significant difference in the spectral measures for these two groups.

In a similar study by Giordano et al [90], the same spectral measures were evaluated and compared in 10 healthy subjects, 10 type II diabetic patients and 20 ESRD patients of whom nine were type II diabetic and 11 were non-diabetic. For the ESRD patients, the HRV parameters were also computed during HD sessions. In the pre-dialytic period, defined as the hour preceding the dialysis session, diabetic patients showed a higher level of sympathetic activity when compared to the non-dialytic patient group. An increase in cardiac sympathetic activity was

observed in both groups of ESRD patients (diabetic and non-diabetic) when compared with the pre-dialytic period. In the post-dialytic period, which lasted until the morning after the session, the non-diabetic patients showed an increase in parasympathetic activity while the diabetic patients did not show a change in activity when compared with the pre-dialytic period. These results support the hypothesis that diabetic patients undergoing HD experience a more persistent shift towards higher sympathetic activity in response to HD, while in non-diabetic patients the increase in sympathetic activity during HD is reversed after the session by an increase in parasympathetic activity.

In a study by Zitt et al [91] significant differences were also observed between diabetic and non-diabetic patients. The patient cohort consisted of 17 patients, of whom 8 were diabetic and 9 were non-diabetic. All patients included in the study were deemed to be hypotension-resistant. Spectral HRV parameters were measured during dialysis alongside non-invasive measures of stroke volume, cardiac output and total peripheral resistance. Diabetic patients showed an overall reduction in autonomic activity or a lack in autonomic response, while non-diabetic patients showed an increase in sympathetic activity during HD. Further, in agreement with the results of Yoshii et al [86], a decrease in stroke volume and cardiac output as well as an increase in total peripheral resistance were observed in both diabetic and non-diabetic patients.

Rubinger et al [92] also performed a similar study on chronic renal failure patients undergoing haemodialysis. They also divided the HD patient population into subgroups, but based the division on the number and severity of hypotensive events in the three months leading up to the study. More specifically, they defined the group of stable patients to consist of patients who had not had more than one

hypotensive event and had not had any signs of circulatory collapse (as defined below) in the 3 months preceding the study. The unstable patients were defined as those who had had at least three monthly hypotensive episodes in the pre-study period, including one episode of circulatory collapse.

For the purposes of the study, a hypotensive event was defined as a symptomatic drop in mean blood pressure of more than 20% of the pre-dialysis value and requiring an intervention. Circulatory collapse was defined as a fall in systolic blood pressure below 85 mmHg which required temporary cessation of the HD session, accompanied by severe symptoms commonly associated with hypotension, such as severe chest pain and loss of consciousness. 56 patients took part in this study, which consisted of 24h ECG monitoring of midweek HD sessions and the pre- and post-dialysis periods. Both time- and frequency-domain HRV measures were considered, including the mean of the RR intervals, the Standard Deviation of Normal-to-Normal RR intervals (SDNN), and various spectral measures calculated using the AR method and 5-minute sequences of RR interval data. When comparing the stable and unstable subgroups, Rubinger et al [92] found that both time-domain and frequency-domain measures were significantly lower in the group of unstable patients when compared with the stable group. This applied to the pre- and post-dialysis period, as well as during dialysis. Further, during HD, while the frequency-domain measures were not significantly different between the two groups in general, the LF/HF ratio was significantly lower in the group of unstable patients when compared to the stable patients. The heart rate was found to be significantly increased in both groups during HD as well.

Cavalcanti et al [93] investigated the spectral density function of RR-interval sequences from 15 hypotension-prone and 15 hypotension-resistant patients dur-

ing one entire HD session free from hypotensive events. The group of hypotension-resistent patients consisted of patients who had not had a cardiovascular collapse (defined below) in the 12 sessions preceding the study, while the group of hypotension-prone patients consisted of patients who had had at least two episodes of cardiovascular collapse in the 12 sessions preceding the study. Cardiovascular collapse was defined in this case as a fall in systolic blood pressure below 95 mmHg with either a drop in blood pressure of more than 20 mmHg in the previous 10 minutes or the presence of hypotension-related symptoms.

The ECG was measured continuously throughout the session, while the blood pressure was measured every 10 minutes using an inflatable cuff. Power spectra were obtained from 3-minute sequences of RR-intervals using the AR method. Comparing the spectra of the stable and unstable groups, it was concluded that the LF component was more dominant than the HF component in the stable patients, while the HF component was more dominant than the LF component in the unstable patients. Further, the LF/HF ratio was consistently less than 2 in the unstable group, and consistently greater than 2 in the stable group, indicating the dominance of the LF component in the stable group, and the dominance of the HF component in the unstable group.

Pelosi et al [34] performed a similar study on a smaller population of 12 patients, of whom 6 had not had any episode of IDH in the month preceding the study (stable patients), while the other 6 patients had had an episode of IDH in at least two of the three sessions in the week preceding the study (unstable patients). For the purpose of this study, IDH was defined as a reduction in mean blood pressure greater than 20% of the predialytic value, accompanied by symptoms such as dizziness and nausea. The patients were aged between 48 and 78 years ($60\pm$ 10

years) and had an average pre-dialysis blood pressure of 140/90 mmHg. Further, none of the patients suffered from other conditions such as diabetes, arrhythmias or neuropathy that could affect their HRV. In addition to monitoring the ECG during HD, continuous arterial blood pressure measurements were also obtained with a Finapres monitor [94]. This device is based on the volume clamp principle first introduced by Penaz [95]. With each cardiac cycle, variations in blood volume in the finger can be measured using a PPG. An inflatable cuff is applied to the finger such as to maintain a constant PPG signal and thus a constant blood volume in the finger. The cuff pressure is then equivalent to the arterial pressure, and can be used as a real-time estimate. Pelosi et al computed the power spectra for both the RR-interval sequence and the blood pressure variability using the AR technique. Both the LF and HF components were calculated for each 256-point interval of the RR-interval sequence. For the blood pressure spectrum, only the LF component was considered.

As described in Section 1.5, tachycardia (a rise in heart rate) is expected in response to fluid removal in order to maintain blood pressure. In this study, only the stable patients showed a significant rise in heart rate in response to fluid removal, and the highest heart rate was reached at the lowest blood pressure reading in each case. Even during IDH episodes, the unstable patients did not show a significant change in heart rate. Overall the stable population showed a direct and linear relationship between changes in systolic blood pressure and changes in RR interval, while the unstable patients showed either no relation or a poor relation. Based on the spectral measures, sympathetic activity (described in more detail in Section 3.1) was found to be enhanced in all patients in the first 40 minutes of dialysis. Further, in the stable patients the level of sympathetic activity increased

in parallel to the increase in HR throughout the session. In contrast, sympathetic activity decreased in later phases of the session in the unstable patients and suddenly dropped during hypotensive events. While these findings correspond well with the expected physiological responses to fluid removal outlined in Section 1.5, only a small number of patients was studied and in order to establish which spectral measures of heart rate variability were optimal for detecting and predicting IDH, a larger study would be necessary.

Ino-Oka et al [96] investigated the relationship during HD between the LF/HF ratio and systolic blood pressure, which was measured non-invasively at 30-minute intervals throughout each session, in 167 patients. Linear regression analysis was applied for this purpose, and three subgroups emerged within the patient cohort: patients with a positive correlation between the LF/HF ratio and systolic blood pressure, patients with an inverse correlation and patients without any significant correlation, with each group consisting of approximately one third of the patient cohort. No significant differences were found in the age distribution, duration of dialysis, or the amount of fluid removed when comparing the different subgroups. However, the positively-correlated group characteristically showed a hypotension-resistant response, together with a low LF/HF ratio, while the inversely-correlated group exhibited a hypotension-prone response together with a high LF/HF ratio. These results suggest that patients with cardiovascular dysfunction in response to fluid removal during HD may be distinguished from patients with an adequate cardiovascular response based on the correlation between systolic blood pressure and the LF/HF ratio.

1.7 Existing Methods for Predicting Patient Deterioration

Changes in vital signs commonly precede various types of adverse events in patients, whether during dialysis or otherwise [97]. Different alarm systems have been developed in the hospital environment for detecting such deterioration and prompting staff intervention. This approach not only allows for a more efficient distribution of resources such as nursing staff, but also identifies subtle changes in physiology that will potentially result in an adverse event at an early stage such as to allow for preventative interventions.

The most simple such system is a single-channel alarm, which is triggered when a monitored parameter exceeds a predefined threshold. This approach however, is prone to false alarms as a result of the rigid threshold criteria, and extensions to multiple parameters have been developed in response [98, 99]. The advantage of multi-parameter systems is that associations between the vital sign parameters are taken into account. The criteria for triggering the alarm are rule-based and clinical diagnosis by clinicians is commonly used to define the rules. As a result, large volumes of patient data have to be reviewed by clinicians for rule definition and refinement to achieve acceptable performance. These so-called expert systems have only found limited application in the hospital setting.

Machine learning is a promising alternative to such expert systems. A classifier is trained using examples from a training dataset to map a given set of input parameters onto classes, without the need for any explicit rules. A disadvantage of this method is that it requires examples of adverse events in order to be trained (supervised learning). These are usually rare in relation to normal data and it is

unlikely in medical applications that examples of every possible adverse event are available for training. As a result, incorrect classification may occur [100]. Novelty detection overcomes this problem, as it can identify new data that has not been encountered during training. Tarassenko et al [101] recently developed a data fusion based patient monitoring system, which is able to identify abnormalities in a patient's vital signs without having previously encountered such abnormalities during training. This approach differs from supervised learning, in that classification is based on the features of the dataset rather than on labels. The data fusion system is built on a probabilistic model of normality in five dimensions (corresponding to five vital sign recordings), which is learned from a large sample of vital sign data previously collected from a representative group of normal patients [101]. If a single input parameter deviates from the population mean by more than 3 standard deviations, or if several parameters together deviate from their normal values by smaller amounts, an alert is generated. Deviations from normality are thus classified as novel and are quantified by an index of normality.

This index is obtained by assuming that, given an estimated n-dimensional joint probability distribution over the n vital sign vectors, any new vital sign vector can be expressed as a combination of parameters from the underlying distribution. Therefore, the probability of a vital sign vector arising from the underlying distribution can be calculated and used to establish an index of normality. This index assumes higher values for less likely vital sign combinations.

1.8 Thesis Overview

IDH remains the most frequent and arguably least understood complication associated with HD, despite various approaches towards both understanding the underlying physiological mechanisms giving rise to IDH as well as devising methods for preventing IDH that could find practical application in a clinical setting. Given the variety of mechanisms that have been identified as giving rise to IDH, it appears to be a multi-faceted condition and this has been acknowledged in more recent analysis [102]. In particular, IDH episodes arising from dysfunction in the autonomic nervous system should be distinguished from other occurrences of IDH. The common feature of all methods suggested for preventing IDH however, is that they focus on one causative mechanism of IDH only. This may explain why no single method in isolation has distinguished itself as yielding an improvement in IDH incidence. Further, with patients attending three HD sessions per week, HD is by nature a highly repetitive form of treatment which allows for the accumulation of a large amount of longitudinal data if the patient is monitored. However, routine vital sign monitoring beyond simple weight and blood pressure measurements at the start and end of each session, has not established itself in clinical practice. Finally, the lack of consensus over a global definition of IDH points towards the dependence of this condition on each individual patient's physiology.

The main objective of this thesis is to determine whether continuously monitoring vital sign data during haemodialysis can assist in distinguishing between "normal" and "abnormal" patterns of physiology in individual patients, and further to establish whether a data fusion model based on such data is capable of predicting episodes of IDH. In this context different definitions of normality are ex-

plored, including global, condition-specific and patient-specific definitions, with the latter being based upon longitudinal vital sign data from individual patients.

For this purpose, vital signs were continuously monitored during HD in 77 patients over the course of one year. The details of this observational study which took place at the Renal Unit of the Churchill Hospital are given in the following chapter, together with some basic analysis of the dataset and its different subsets.

HRV has been identified as a relevant parameter in relation to IDH in HD patients, and a variety of measures in both the time and frequency domains have been devised in this context. With regard to spectral measures, the two branches of the autonomic nervous system have conventionally been thought to only affect the frequency spectrum of the RR-interval time series of the ECG within fixed frequency bands. However, this approach does not account for the association between RSA and the autonomic nervous system. In Chapter 3, a new spectral measure for estimating HRV is introduced, in which the breathing rate is used to identify physiologically relevant peaks in the frequency spectrum of the RR intervals rather than taking entire frequency bands into account. This allows for more flexibility with respect to the location of the relevant peaks, and the physiologically relevant power within the spectrum to be considered.

In the subsequent chapters, different models of normality are explored using a visualisation technique known as the Neuroscale algorithm. In Chapter 4, the definition of normal physiology in HD patients is first considered in relation to a global model based on vital sign data from all non-IDH sessions that were recorded as part of the study. Two alternative models are subsequently considered. The first model is trained using the relative changes in the vital signs from non-IDH sessions with respect to their value at the start of the session. In the sec-

ond model, the training data is separated based on whether the respective session was recorded from a patient with dysfunction in the autonomic nervous system or not. Based on this distinction, two separate condition-based models are devised and only applied to vital signs from patients with the same conditions. For each of the proposed models, the ability to distinguish between normal vital signs and vital signs stemming from an IDH episode is analysed.

By virtue of the longitudinal nature of the data collection study, patient-specific models of normality are also created using the Neuroscale algorithm. These are trained exclusively on vital sign data from non-IDH sessions for each patient, and are studied in more detail in Chapter 5, particularly in relation to their ability to distinguish between IDH and non-IDH vital signs in lower-dimensional projections.

In the final part of this thesis, a patient-specific data fusion model is introduced based on a probability density distribution being calculated for each patient using vital sign data from normal sessions. This allows for a probability to be estimated for new data, and alerts to be generated below a set probability. In this context, the choice of threshold and of the persistence criterion for alert generation is explained and the model's ability to predict IDH ahead of time is assessed. Finally, there is a discussion of other aspects that can be investigated to further the research presented in this thesis.

Chapter 2

Data Collection

The data on which the work in this thesis is based was collected during an observational study from October 2011 to October 2012 at the Renal Unit of the Churchill Hospital in Oxford (REC reference number 11/SC/2070), following from a pilot study conducted in 2009. The composition of the dataset is described below.

Seventy-nine patients, numbered OB1 to OB79, were recruited after giving informed consent. Two patients (OB16 and OB72) withdrew prior to completion of the study. The study consists of non-invasive vital sign monitoring during regular outpatient haemodialysis sessions. The number of sessions recorded per patient is shown in Fig. 2.1. The time between recorded sessions varies between two days and several months, but where the gap between two sessions was more than two days, sessions were recorded on the same day of the week in order to overcome the inherent asymmetry in the period between dialysis sessions (described in Chapter 1). 83% of sessions were recorded within one month of the previous session and 92% were recorded within two months of the previous session.

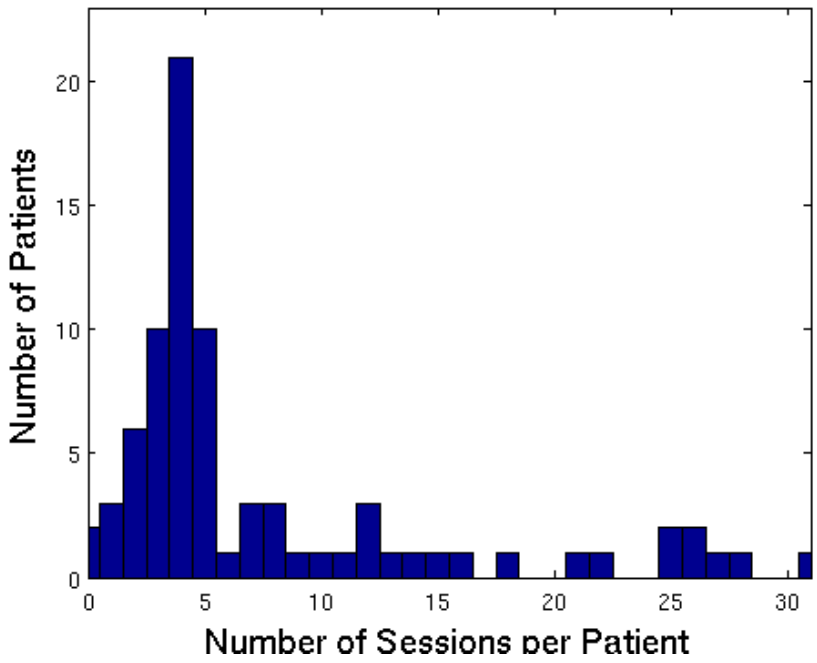


Figure 2.1 – Histogram of number of sessions recorded for each patient. 603 sessions were recorded in total from 77 patients. The maximum number of sessions recorded per patient is 31, and 75% of patients had at least 4 sessions recorded.

Dialysis was performed using Braun Dialog$^+$ machines (B. Braun Melsungen AG, Germany), and the treatment parameters listed in Table 2.1 were recorded in a separate Braun database called Nexadia. The vital signs were recorded using the Equivital™Vital Signs Physiological Monitor (Hidalgo Ltd, Cambridge UK), which consists of a chest belt with skin electrodes and an expansion sensor. The parameters recorded include the ECG and PPG signals (described in Sections 1.3.1 and 1.3.3, respectively), as well as thoracic cavity movement from a strain gauge. From these waveforms, the heart rate, blood oxygen saturation and respiratory rate can be derived, respectively. The algorithms used to obtain the heart rate and the breathing rate are outlined in Section 2.1. The blood oxygen saturation was obtained directly as an output from the pulse oximeter, as described in

Dialysis Machine Parameters	
Venous Pressure	Arterial Pressure
UF Rate	Transmembrane Pressure
Blood Flow	Dialysate Temperature
Conductivity	Dialysate Flow
Blood Volume	Anticoagulation Rate
UF Volume	Anticoagulation Volume

Table 2.1 – B. Braun machine parameters recorded during dialysis.

Section 1.3.3.

2.1 Signals and Derived Vital Signs

The ECG was recorded at a sampling frequency of 256 Hz, and the PPG and strain gauge signals were recorded at 100 Hz and 25.6 Hz, respectively. The beat-to-beat heart rate is derived from the ECG by applying a Pan-Tompkins-based peak detector [27] to the signal to identify the QRS complexes. The reciprocal of the time interval between two successive beats yields the instantaneous heart rate which is commonly given in beats per minute (bpm). However, due to poor signal quality, the peak detector occasionally fails to identify any QRS complexes. In such cases, one minute of ECG following the failed peak detection was discarded, and the peak detector re-applied. On average, 41 minutes of ECG data per 4-hour session had to be discarded on this basis.

Similarly, by applying a peak detector to the strain gauge signal and measuring the time interval between successive peaks, an instantaneous breathing rate signal was derived.

In addition to the parameters recorded with the Equivital vital sign monitor,

systolic and diastolic blood pressures were measured every 15 minutes using a cuff on the resting arm that did not serve as the site for dialysis access. The blood pressures were also recorded in the Nexadia database. The MAP, defined by the weighted sum of the systolic and diastolic blood pressures as $\frac{1}{3}$ systolic BP $+\frac{2}{3}$ diastolic BP, is commonly used in the literature, and was therefore calculated for each of the systolic and diastolic blood pressure readings. Further, the haematocrit was recorded every 20 seconds using a Crit-line monitor (HemaMetrics, formerly In-Line Diagnostics) and stored in a separate file.

2.2 The Dataset

In total, 603 dialysis sessions were recorded from 77 patients. The age of the study population ranges from 21 to 91 years, and the median age is 65, compared to 66.3 for patients across the UK [13]. Overall, 2030 hours and 52 minutes of dialysis sessions were recorded, with an average session length of 3:57:04 hours. The shortest recorded session was 2:59:57 hours and the longest was 5:02:42 hours.

The study population consists of 77 patients, of whom 55 (71%) are male and 22 (29%) are female. Female patients account for 21% of sessions, while male patients provide 79% sessions. The age structure of the study population is shown in Table 2.2. On average, 2.4 litres of fluid ($\pm$1.05 litres) were removed per patient per session.

The patient demographics are presented in Table 2.3. For the study population, the median length of renal replacement therapy is 30 months (ICQ 14-66 months) while the median length of HD treatment is 26 months (ICQ 11-49 months).

Age Range	All	Male	Female	UK Average
18-39	12% (9)	15% (8)	5% (1)	15.1%
40-64	38% (29)	40% (22)	32% (7)	50.5%
65-74	22% (17)	20% (11)	27% (6)	19.3%
75+	28% (22)	25% (14)	36% (8)	15.1%
Total	100% (77)	71% (55)	29% (22)	100%

Table 2.2 – Age distribution for the patients in the recorded dataset. The first column shows the overall age distribution of the study population, while the second and third columns show the relative distributions of the male and female populations separately. For comparison, the overall prevalence of renal replacement therapy patients in the UK is shown in the fourth column. (It should be noted that this includes both haemodialysis and peritoneal dialysis patients.)

Ethnicity		**Diabetes**	40% (31)
White	88% (68)	Type 1	8% (6)
Black	3% (3)	Type 2	32% (25)
South Asian	3% (3)	Insulin	23% (18)
Other	3% (3)	**Medication**	
Primary Diagnosis		ACEi	21% (16)
Glomerulonephritis	31% (24)	A2RB	21% (16)
Diabetic Nephropathy	21% (16)	β-Blocker	38% (29)
Hypertensive/Renovascular	12% (9)	α-Blocker	16% (12)
Pyelonephritis/Reflux	6% (5)	Ca-Channel Blocker	34% (26)
Polycystic Disease	5% (4)	Other	10% (8)
Renal Dysplasia	4% (3)	**Antihypertensives per patient**	
Other	9% (7)	none	30% (23)
Unknown	12% (9)	1	18% (14)
Ischameic Heart Disease	39% (30)	2	40% (31)
PCI	25% (19)	3	7% (5)
known LVH	21% (16)	4	5% (4)
Smoker	19% (15)	**Listed for Transplant**	25% (19)

Table 2.3 – Patient demographics for the study population.

	Lower Limit	Upper Limit
Heart Rate (bpm)	40	180
MAP (mmHg)	20	180
SpO_2 (%)	70	none
Breathing Rate (breaths/min)	4	45

Table 2.4 – Artefact rejection limits. Data outside of these limits was discarded.

2.3 Artefact Rejection and Signal Quality

Based on clinical knowledge, artefact rejection limits were applied to each recorded parameter as shown in Table 2.4. Individual datapoints that fell outside of these limits were discarded. Further, beat-to-beat changes in the heart rate that exceeded 20% were also discarded for being artefactual.

Further to artefact rejection, some sessions were not of adequate signal quality to warrant their use as part of further analysis. These sessions were discarded entirely in the subsequent work. As a simple measure for data quality, the variance of each vital sign was calculated for each session. For sessions that had vital signs with high variance, the data was discarded with a heuristically-set threshold. As a result, data from 141 sessions was excluded from the dataset, leaving 462 sessions for further analysis. An example of a typical time plot of the recorded parameters is shown in Fig. 2.2a, with an example of data from a discarded session shown in Fig. 2.2b.

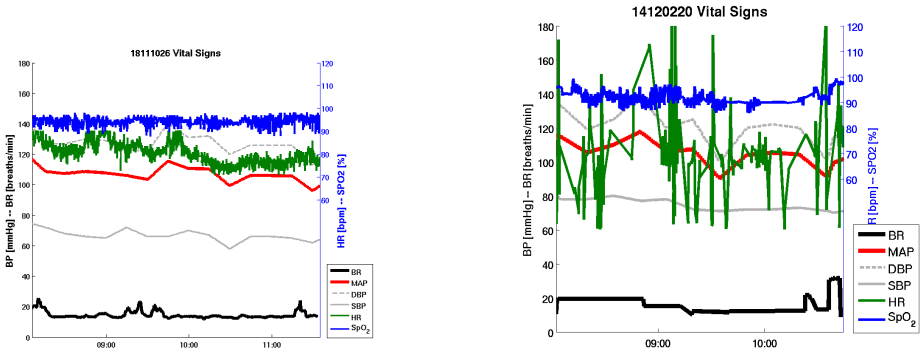
By far the most common problem in relation to signal quality was loss of the HR signal, which points to poor ECG quality. Further, approximately 50% of breathing rates, as derived from the strain gauge signal, are considered to be artefactual based on the clinically derived physiological limits shown in Table 2.4. This is indicative of the fact that the strain gauge signal is not always of good

enough quality to discern independent respiratory cycles as peaks. Regarding both signals, it should be noted that there were three different sizes for the belts using which both the strain gauge signal and ECG were recorded. Although attempts were made to further vary the sizes by inserting adjustable straps at the belt clip, the belts were not always well fitted to the patient. This affected the quality of the ECG signal since the ECG electrodes are embedded in the belt and, if not fitted correctly, are not in full contact with the skin.

An example time plot of the vital sign parameters recorded during a session is shown in Fig. 2.2a. The heart rate is shown in green and the oxygen saturation is shown in blue, with the relevant scale for both parameters being on the right side of the plot. The mean arterial blood pressure is shown in red, with the systolic and diastolic blood pressures in grey. It should be noted that while the blood pressure is only measured every 15 minutes, the plotted time series is interpolated between the measured datapoints, primarily for visualisation purposes. The breathing rate is shown in black. The relevant scale for the blood pressure and breathing rate is on the left side of the plot.

2.4 Dataset Statistics

As the chief objective of this work is to predict IDH in HD patients, the recorded vital sign data was divided into two subsets based on whether or not patients experienced an episode of IDH in each session. A given session was defined as an IDH-session if the systolic blood pressure abruptly fell below 100 mmHg at any time during the session, where an abrupt fall is defined as a change greater than



(a) **(b)**

Figure 2.2 – Sample time plots of the vital sign parameters. The heart rate (HR), breathing rate (BR) and oxygen saturation (SpO_2) as well as the systolic (SBP), diastolic (DBP) and mean arterial (MAP) blood pressures are shown. The left plot shows a typical time series, while the right plot is an example of a discarded session. The time series was discarded due to the high HR variance, as well as the large proportion of missing data.

20% with respect to the previous (actual) measurement. Based on this definition, the choice of which is justified in more detail in Section 1.4, 162 (26.9%) of the 462 recorded sessions were classified as IDH sessions. These 162 IDH sessions were recorded from 48 different patients, that is 62% of the study population. Similar rates have been reported in the literature [5]. A complete overview of the sessions recorded for each patient is given in Appendix A. The means (μ) and standard deviations (σ) of each vital sign parameter are shown separately for the IDH and non-IDH sessions in Table 2.5. Interestingly, the differences in both the means and standard deviations between the two subsets are not very great, particularly with regard to the blood pressure where larger discrepancies between the two subsets might be expected.

Histograms of the different data subsets are shown in Fig. 2.3. While no sig-

	μ_{IDH}	μ_{nonIDH}	σ_{IDH}	σ_{nonIDH}
HR (bpm)	79.0	76.9	14.1	16.3
SpO_2 (%)	94.3	93.6	4.7	5.2
BP (mmHg)	100.9	104.0	19.1	18.6
BR (breaths/min)	21.6	23.1	7.6	9.8

Table 2.5 – Means and standard deviations of IDH and non-IDH sessions.

nificant differences can be observed in the breathing rate and oxygen saturation distributions when comparing data from IDH sessions (shown in red) with data from non-IDH sessions (shown in green), the heart rate and mean arterial blood pressure distributions differ. IDH is defined by an abrupt decrease in blood pressure to low values. It is therefore expected that the blood pressure histogram for the IDH sessions will be more populated at lower values than the distribution of non-IDH sessions. The heart rate histogram for IDH sessions has a peak at 80 bpm, whereas the distribution for non-IDH sessions has a broader peak at lower heart rates between 70 and 80 bpm. The two distributions are similar at very low heart rates between 40 and 50 bpm as well as at heart rate values above 90 bpm.

2.5 Longitudinal Data Subset

Following the overall analysis presented above, the dataset is further partitioned into a subset which consists only of data from patients who had more than 12 sessions recorded as part of the study, of which at least three sessions were sessions in which the patient experienced IDH according to the definition given in Section 1.4. These criteria were chosen empirically to enable longitudinal analysis of patient data. In total, 6 of the 79 patients who participated in the study fell into this

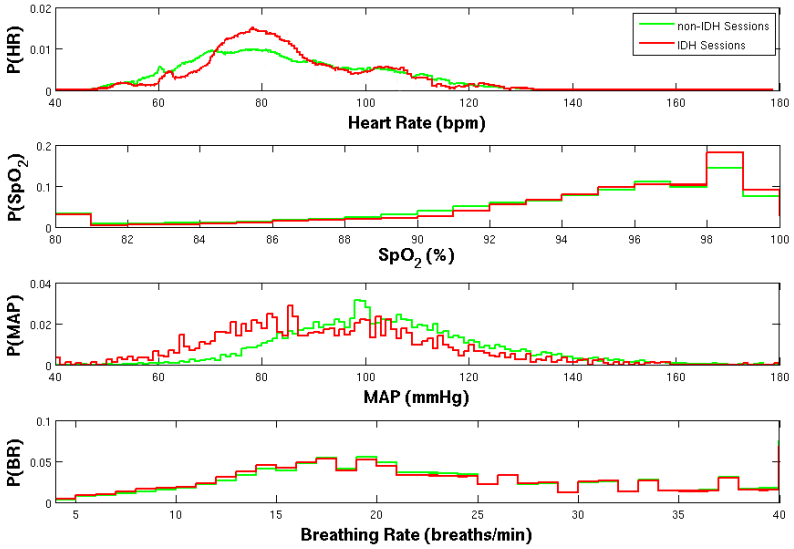


Figure 2.3 – Normalised histograms of vital signs from IDH and non-IDH sessions. No significant differences can be observed in the breathing rate and oxygen saturation vital sign histograms. The mean arterial blood pressure values for the IDH sessions are lower than for the non-IDH sessions, as would be expected.

category after removing sessions of inadequate data quality according to Section 2.3. Details of these patients and of the sessions recorded is shown in Table 2.6.

					HR [bpm]		SpO_2 [%]		MAP [mmHg]		BR [breaths/min]	
	Age	Sex	Sessions	IDH sessions (%)	μ	σ	μ	σ	μ	σ	μ	σ
OB1	68	m	26	12 (46%)	71.2	5.8	92.5	2.4	142.8	17.2	20.9	8.7
OB7	68	m	31	12 (39%)	59.3	5.5	93.8	4.0	117.2	18.0	24.4	5.9
OB10	61	m	21	14 (67%)	77.5	3.0	95.7	4.0	84.5	16.8	17.9	6.2
OB12	73	m	13	5 (38%)	66.2	4.0	97.9	1.4	102.6	5.8	19.5	4.7
OB19	63	m	15	5 (33%)	69.4	5.9	96.2	2.5	117.8	14.1	23.7	9.1
OB20	55	m	28	16 (57%)	70.5	6.0	96.0	3.2	120.8	16.1	21.5	6.8

Table 2.6 – Subset of seven patients for longitudinal analysis. The means and standard deviations of the four vital signs are calculated from the non-IDH sessions only.

Chapter 3

Heart Rate Variability

3.1 Introduction and Background

In order to maintain homeostasis, the human heart rate exhibits quasi-periodic variations, the frequency of which can be quantified either in statistical terms or using spectral estimation methods. Heart Rate Variability (HRV) is formally defined as the variation in time of the period between consecutive heartbeats. The time between two heartbeats and ultimately the heart rate are predominantly regulated by the autonomic nervous system (ANS) in response to various stimuli including central command, respiration and baroreceptor feedback [103]. In particular, it is the balance between the sympathetic and parasympathetic branches of the nervous system that control the heart rate. An increase in sympathetic activity (or a decrease in parasympathetic activity) leads to an accelerated heart rate, whereas a decrease in sympathetic activity (or an increase in parasympathetic activity) causes a deceleration in the heart rate.

The importance of HRV as a clinically relevant parameter was first identified

in fetal heart rate monitoring, where the absence of beat-to-beat HRV was held to be indicative of abnormalities in the central nervous system's modulation of the heart rate [104]. Since then, different measures of HRV have been used to quantify a number of physiological abnormalities, including the risk of cardiac failure [105] and diabetic neuropathy [106].

In normal individuals without cardiac disease, the heart rate has a high beat-to-beat variability: it fluctuates with respiration, increasing with inspiration and decreasing with expiration, a phenomenon also known as RSA (described in Section 1.3.5). Higher HRV is associated with functionally efficient autonomic control and different measures of HRV (as outlined in Section 1.3.2) have previously been used to evaluate the risk of sudden cardiac death in various conditions including cardiac disorders, multiple sclerosis, and diabetes [107, 108].

Physiologically, HRV arises as a result of sympathetic and parasympathetic mediation of the heart. Sympathetic effectors include epinephrine and norepinephrine which act on the sino-atrial and atrio-ventricular nodes to increase the heart rate. The parasympathetic nervous system affects the heart rate via acetylcholine which also acts on the sino-atrial and atrio-ventricular nodes to decrease the heart rate.

Decreased HRV in the time domain has been established as an independent risk factor for mortality in both cardiovascular disease and healthy populations. More specifically, it has been linked with adverse cardiac events in healthy subjects. In the 1998 UK heart study [109], the annual mortality rate in patients with reduced time domain HRV measures from a population of 433 patients was found to be 51.4%, compared with 5.5% in patients with normal HRV values.

The objective of this chapter is to establish whether spectral measures of HRV may have any predictive value with regard to IDH or alternatively as to the sus-

ceptibility of patients to IDH in general. A general overview of both spectral and time-domain methods was given in Chapter 1, and the methods in use for assessing HRV in HD patients are introduced below. Based on the results obtained with these methods and the physiological interpretation of the power spectrum of the RR-interval time series, a new method for assessing HRV in HD patients during dialysis is introduced and applied to the patient population described in Chapter 2.

3.2 Short-Time Fourier Transform

Several methods have been proposed for capturing frequency information during transient cardiac responses over longer periods of time. Most commonly, the short-time Fourier transform (STFT) [110] has been used for this purpose. The STFT is an extension of the standard Fourier transform, which, for a continuous function $f(t)$ in time, is given by:

$$F(\omega) = \int_{-\infty}^{\infty} f(t)e^{-j\omega t}dt \tag{3.1}$$

where ω is the angular frequency $\omega = 2\pi f$. In the case of a discrete-time signal $x[n]$ sampled at N points $n = 0, 1, ...N-1$, this becomes:

$$X(\omega_n) = \sum_{n=0}^{N-1} x[n]e^{-j\omega_k n} \tag{3.2}$$

where $X(\omega_n)$ is the discrete-time Fourier transform (DTFT). The angular frequency ω_n is discretised as well because while X could in principle be evaluated

at any frequency, only a finite number of outputs N will be significant given that only N points are considered as part of the transform. The power spectrum of the Fourier transform is obtained by taking the square of the magnitude of the Fourier transform, $P(\omega) = F(\omega)^2$, and the power in specific frequency ranges can be obtained by integrating $P(\omega)$ over that range of frequencies.

The STFT provides both frequency and time decomposition of a non-stationary signal by calculating consecutive power spectra of short segments of the signal. It is formally defined as:

$$X[n,\lambda) = \sum_{m=-\infty}^{\infty} x[n+m]w[m]e^{-j\lambda m} \tag{3.3}$$

where $w[m]$ is a window sequence, with m indicating the position of the window in time, and λ is the continuous frequency variable of the STFT. As the representation is time-dependent, the one-dimensional sequence $x[n]$ is transformed into a two-dimensional function of the discrete time variable n and the frequency variable λ. In principle, the STFT can be interpreted as the discrete-time Fourier transform of $x[n+m]$ taken in the window $w[m]$. At each value of n, a different section of the signal is analysed through the window. By sequentially plotting the power spectra of the signal segmemts in time, we obtain a spectrogram, as shown in Fig. 3.1.

Other methods have been suggested for computing the power spectra of non-stationary signals over time, amongst them the Smoothed Pseudo-Wigner-Ville Distribution (SPWVD) [111], time-varying AR analysis and the Continuous Wavelet Transform (CWT) [112].

All methods mentioned above have their limitations. For the Short-Term

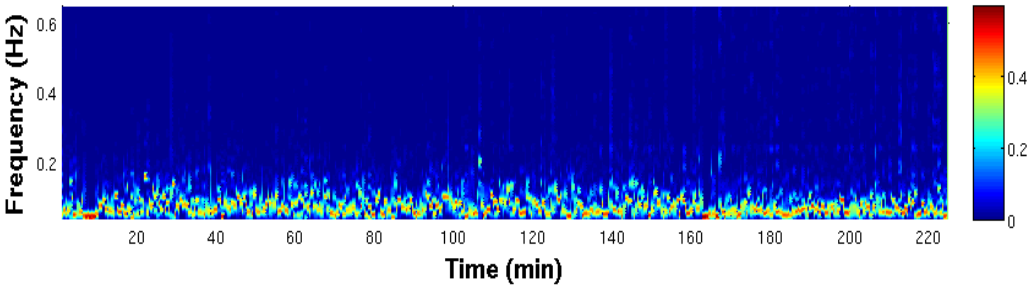


Figure 3.1 – A representative spectrogram, showing the minute-by-minute power spectra for an entire HD session. The colour scheme indicates the normalised density of the power spectrum at given times and frequencies.

Fourier Transform (STFT), there is a trade-off between the time and frequency resolutions, as described in more detail in Section 3.5.3. In the AR models, the time and frequency resolutions can be set independently, but instead it becomes difficult to identify the power of the model poles at a given frequency. For the CWT, the time resolution is dependent on the frequency. This is advantageous when looking at higher frequencies, as it increases the time resolution. However, difficulties arise in computing the spectral power when the frequency resolution of the transform does not match the desired frequency band. The high time-frequency resolution of the SPWVD makes it a strong method for non-stationary analysis, particularly in signals whose spectra change rapidly in time, over the course of only a few samples. For stationary signals however, there is no advantage in using the SPWVD over the STFT, as the spectra correspond [113].

In conclusion, the STFT is a suitable method for the purpose of the analysis in this chapter, as the main focus will be on changes in HRV over the course of a HD session, with the HRV measures computed for windows of one-minute duration, during which the HRV is assumed to be quasi-stationary.

3.2.1 The STFT Window

As mentioned previously, the STFT is equivalent to the DTFT of the sequence $x[n+m]$ taken in the window $w[m]$:

$$X[n,\lambda) = DTFT\{x[n+m]w[m]\} = \frac{1}{2\pi}DTFT\{x[n+m]\} * DTFT\{w[m]\}$$

At a given n, this becomes:

$$X[n,\lambda) = \frac{1}{2\pi}\int_0^{2\pi} e^{j\theta n}X(e^{j\theta})W(e^{j(\lambda-\theta)})d\theta$$

where W is the Fourier transform of the window function w. Since X is periodic with period 2π, variable of integration θ ranges from 0 to 2π. This is effectively the convolution between the Fourier transform of the window, given by $W(e^{j\omega})$, and the Fourier transform of the shifted signal.

The window in the STFT serves to limit the transformed section of the signal such that the frequency characteristics are approximately constant over its entirety (assumption of quasi-stationarity). Signals occurring at more rapidly varying frequencies should thus be analysed with shorter windows, while longer windows are appropriate for slower frequency variations. However, as the window length is shortened, the number of samples included in the Fourier transform is inevitably reduced, thus decreasing its frequency resolution. As a result, improvements in the frequency resolution of the spectrogram can only be achieved at the expense of temporal resolution and optimizing the window size is a key issue in this regard.

3.2.2 Hamming Window

The choice of pre-processing window is also important with respect to spectral leakage to nearby frequencies. This may occur due to the fact that the signal segments contained within each window can have major discontinuities at their edges, giving rise to side-lobes in the periodogram. The Hamming window [114], which is given by

$$W_H(n) = \begin{cases} 0.54 - 0.46\cos(\frac{2\pi n}{N}), & 0 \leq n \leq N, \\ 0 & \text{otherwise.} \end{cases} \tag{3.4}$$

for a window length of $L = N + 1$, has been found to provide the optimal trade-off between spectral leakage, amplitude of side-lobes and width of the central peak (which determines the ability to distinguish frequency components) [115].

3.3 Spectral Components

Based on the widely accepted assumption of two peaks in the frequency spectrum of the RR-interval sequence, the HRV spectrum is conventionally partitioned into two bands, the low frequency (LF) band from 0.04 to 0.15 Hz and the high frequency (HF) band from 0.15 to 0.4 Hz. The power in the different frequency bands is in principle affected by either the sympathetic nervous system (in the LF case) or parasympathetic nervous system (in the HF case), although evidence exists that the parasympathetic nervous system might also affect the LF frequencies, as discussed below.

Results of spectral analysis are commonly given in terms of the ratio of power

in the LF and HF bands, which is believed to be indicative of the sympathetic-parasympathetic balance [116, 37], although controversy still surrounds this metric [117]. The evidence for the associations of each frequency band with the two branches of the autonomic nervous system is given in the following sections.

Low Frequency (0.04-0.15 Hz) The low frequency band was first associated with the sympathetic nervous system [116, 37] through controlled experiments in both humans and dogs, in which sympathetic activity was either functionally increased through orthostatic stimuli (e.g. postural change) or pharmacologically inhibited. Subsequent experiments however, have proven this association to be overly simplistic. Upon both surgical [118] and pharmacological [35] inhibition of the sympathetic nervous system, the LF peak remains prominent, although markedly reduced, implying that other mechanisms also affect LF power.

Other studies suggest a lack of correlation between the sympathetic nervous system and LF power. Hopf et al [119] did not observe any changes in LF power upon pharmacological inhibition of sympathetic activity. Similarly Kingwell et al [120] found no direct correlation between the sympathetic nerve firing rate [121] and LF power. Therefore, while LF band power reflects sympathetic activity, it cannot be treated as an exclusive measure. Rather, the evidence suggests that it is a composite marker, encompassing multiple mechanisms of the autonomic nervous system.

High Frequency (0.15-0.4 Hz) Correlation of the high frequency band with the parasympathetic nervous system is more specific, as changes in HRV due to RSA are captured by the high frequency peak in the HRV spectrum and RSA in turn

is considered to arise as a consequence of parasympathetic control mechanisms [45].

Nevertheless, there are several confounding factors that may affect the relationship between RSA and parasympathetic activity, leading to dissociation under certain conditions [122]. Changes in the HF peak in the absence of changes in parasympathetic activity can for instance be caused by changes in tidal volume [123].

Finally, studies involving heart transplant patients indicate that HRV has a purely mechanical component as well. Given that transplanted hearts have no cardiac nerve innervation, the small changes in RR interval in these patients are attributed to mechanical effects arising from stretching of the atrial walls as a result of changes in cardiac filling and in thoracic pressure during respiration [118].

In the light of these findings, one should keep in mind the confounding factors that can cause dissociation between the autonomic nervous activity and the spectral power when using the LF and HF peaks as specific markers for sympathetic and parasympathetic activity, respectively.

3.4 Defining the LF and HF Bands

The sympathetic and parasympathetic nervous systems affect heart rate at different frequencies, as explained in Section 3.3. Conventionally, the frequency ranges in which the two systems operate (0.04-0.15 Hz for the sympathetic and 0.15-0.4 Hz for the parasympathetic nervous system) have been assumed to be fixed and the LF and HF measures are calculated as the respective integrals of the frequency spectrum of RR intervals over these frequencies. However, it has been argued

that this approach is inadequate under conditions in which the sympathetic and parasympathetic responses are varying in time, such as during stress testing or intensive exercise [124].

This alludes to the main argument against fixed bands for analysing the frequency spectra of RR intervals, namely the nonsteady nature of human physiology: natural variations aimed at maintaining homeostasis can go beyond the fixed boundaries. In particular, the respiratory frequency in humans can vary beyond 0.15-0.4 Hz which corresponds to breathing rates between 9 and 24 breaths/minute. For instance, 19 % of all recorded breathing rates in the study population described in Chapter 2 are above 24 breaths/minute, as shown in Fig. 3.2.

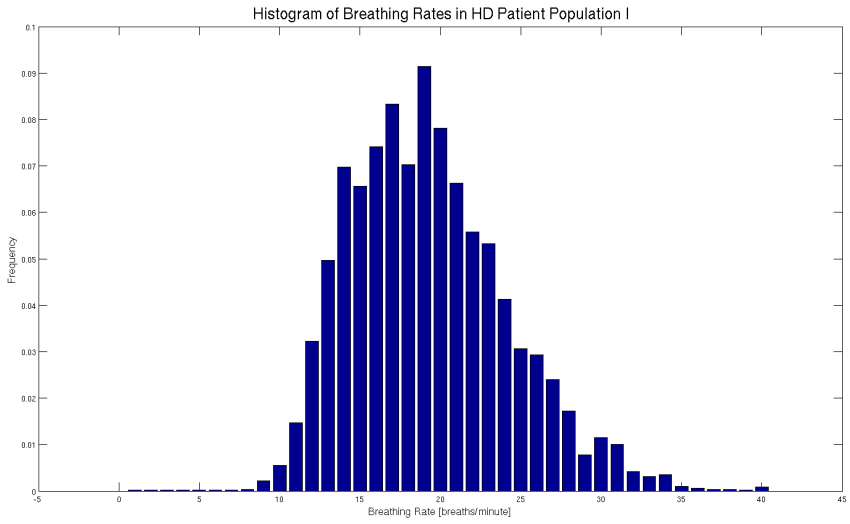


Figure 3.2 – Distribution of breathing rates in HD study population B. 19% of all recorded breathing rates are above 24 breaths/minute, which is the upper limit of the fixed HF band.

Given that the respiratory frequency affects HRV through RSA, adhering to fixed frequency boundaries does not allow for extreme variations to be taken into

consideration. Incidentally, these are the observations of most interest in patients at risk. Further, from a patient-specific perspective, dynamic bands are preferable, as they track more closely transients in the individual's physiology by interpreting the data in a more intelligent fashion.

Different suggestions have been made as to how the dynamic boundaries should be determined. Bailon et al [111] suggest using respiratory information to determine the HF band, while either keeping the LF band stationary at 0.04-0.15 Hz or defining it from 0.04 Hz up to the lower limit of the HF band. The HF band is centered on the respiratory frequency, which can be established either from direct measurements or indirectly from the ECG, and has a bandwidth of 0.25 Hz in accordance with the non-stationary bandwidth. The lower boundary, however, cannot fall below 0.15 Hz in order to avoid overlap between the two bands.

Alternatively, it has been suggested that the upper limit of the HF band be set to half the mean HR in order to include the whole range of potential respiratory frequencies. However, in cases of exceptional stress as for instance during exercise, this approach would include spurious high frequency components that are not considered to result from parasympathetic activity [124].

Goren et al [125] devised an algorithm which looks for the spectral peak within a predefined frequency range (0.02-0.15 Hz for LF and 0.15-0.6 Hz for HF) and then locates the two minima surrounding this peak. The minima define the boundaries of the band under consideration. This algorithm is more adaptive to patient physiology, as it only takes the power in the relevant peaks of the spectrum into consideration, in contrast with both the fixed- and variable-band approaches, which consider all power in the given frequency ranges. However, in the absence of further information as to the location of the peaks, the approach suggested by

Goren requires prior knowledge as to the frequency bands in which the peaks should be located.

3.5 A Combined Approach: Non-stationary Peak Measures

The use of non-stationary HF bands as proposed by Bailon et al [124] is particularly useful in cases where the respiratory rate is outside of the range of 9-24 breaths/min (which corresponds to a respiratory frequency of 0.15 to 0.4 Hz), as illustrated in Fig. 3.5. In such cases, the physiologically relevant high-frequency peak lies outside the fixed HF band, which ranges from 0.15 to 0.4 Hz, and its power would therefore not be included in the conventional HF measure described in Section 3.3. Further, the resulting HF power integral from the fixed band is misleading in such cases, since it measures power in another part of the spectrum, excluding the relevant peak. In principle, therefore, a non-stationary HF band which is centered on the respiratory frequency is advantageous as it can only improve the accuracy of HF power extraction. In cases for which the respiratory frequency is at the centre of the conventional HF band, the stationary and non-stationary measures are equivalent. In cases for which the respiratory frequency is at the edge of the fixed HF band, or even outside of it, the non-stationary band yields a more appropriate HF measure as it focusses on the physiologically-relevant power.

However, as can be seen in the power spectrum in Fig. 3.3, the width of both the dominant LF and HF peaks only constitutes a fraction of the respective bands. In the LF range there is ordinarily only one peak, representative of sympathetic

activity, and as outlined in Section 3.3, the main peak in the HF range of the power spectrum is associated with the respiratory frequency. The physiological origin of the remaining peaks in the spectrum is controversial. Some have argued for the existence of other peaks in addition to the LF and HF peaks: a mid-frequency peak associated with baroreflex activity or a very-low-frequency peak associated with spontaneous oscillations in arterial blood pressure, known as Mayer waves [126]. However, in both cases the precise physiological mechanisms underlying the oscillations are not clear [127, 128]. Further, spurious peaks due to (periodic) noise in the ECG or harmonics of the LF or HF peaks frequently appear in the power spectrum as well. As such, by using the conventional widths of the LF and HF bands (0.11 and 0.25 Hz, respectively), the power at frequencies surrounding the dominant peaks in such subsidiary peaks, is included in the respective measures. As the physiological relevance of subsidiary peaks is unclear, there is no reason to include them in the LF and HF measures. To overcome this problem, the most sensible approach to quantifying the sympathetic and parasympathetic activity using the power spectrum of the RR time-interval time series is to focus on the power in the main LF and HF peaks, so as to include the most physiologically relevant data, analogous to the approach taken by Goren et al [125].

Given that there is commonly only one peak in the LF range, an alternative LF measure can be defined as the integral of the power in the highest peak in the LF range (0.04 - 0.15 Hz), with the limits of integration being the two minima on either side of the dominant peak. The dominant peak in the HF band corresponds to the respiratory frequency, as RSA modulates the heart rate at a frequency representative of the respiratory rate. The HF peak shifts in accordance with changes in respiratory rate [129], whereas the amplitude of the HF peak is not affected by

the respiratory rate over the normal range of respiratory frequencies [130].

In the dataset on which this work is based, the respiratory rate is recorded using a strain gauge, and this information is used to identify the relevant HF peak in the power spectrum. The instantaneous respiratory rate can be calculated as the reciprocal of the time difference between successive peaks in the strain gauge signal. Median averages of the instantaneous respiratory frequency signal can be used to establish where the dominant HF peak should be located in the frequency spectrum of a given window of RR-intervals. The power of this peak is used as the HF measure, with the limits of integration being the two minima surrounding the peak.

The detailed steps taken to calculate these measures are outlined in the following section.

3.5.1 Pre-Processing

The following analysis was performed on the ECG recorded from HD patients in the study as described in Chapter 2.

3.5.1.1 RR-interval sequence

The peaks of the ECG R-waves were detected using a Pan-Tompkins based algorithm [27], and the sequence of RR-intervals was established by taking the time difference between succesive peaks. The timestamp of the interval was chosen as the timestamp of the earlier peak of the interval.

3.5.1.2 Ectopic Beat Removal

Ectopic beats, that is beats that are not initiated by the SA node (see Section 1.3.1), disturb the normal sinus rhythm and result in impulse-like artefacts surrounding the ectopic beat in the RR-interval series. As such, they introduce spurious frequency components in the power spectrum of the RR-intervals. Further, it has been demonstrated that ectopic beats occur frequently in HD patients [131]. As the sympathetic and parasympathetic nervous systems control the heart rate via the SA and AV nodes, beats not originating from the SA node should not be included in the analysis. Therefore, to assess the function of the autonomous nervous systen during HD, ectopic beats should be excluded from HRV spectral analysis [132].

The most common methods for ectopic beat removal consist of either removing entire segments of data from analysis (if the incidence of ectopic beats is very high) or alternatively removing the ectopic beats from the ECG and interpolating the RR interval sequence [132]. Different criteria have been suggested for distinguishing between ectopic and normal beats: Kamath et al [132] suggest that interbeat intervals that differ by more than 20% when compared with the previous valid interval should be rejected as being artefactual. This limit was applied to the RR-interval sequence, as described in [133].

3.5.1.3 Resampling

The RR-interval sequence can conceptually be thought of as a recursive function, since the timestamp of each RR-interval depends on the length of the previous interval. Therefore, the RR-interval sequence is inherently unevenly sampled. In order to apply the Fast Fourier Transform however, the sequence must be evenly

sampled [110]. Commonly either linear or cubic spline interpolation is used for this purpose, and linear interpolation was used in this work. Sampling frequencies between 2 and 10 Hz have been suggested, but it appears that so long as the Nyquist criterion is satisfied, the resampling frequency does not significantly affect the spectral measures [134]. Given that an upper artefact rejection limit of 180 bpm was chosen for the heart rate based on the basis of clinical knowledge (see Section 2.3), a sampling frequency of at least 3.03 Hz is required to satisfy the Nyquist criterion. In accordance with previous work [38], a sampling frequency of $f_s = 4$ Hz was chosen.

3.5.2 Power Spectral Density Calculation

Prior to applying the FFT, each window was multiplied with a Hamming window of the same length, as described in Section 3.2.2. To correct for nonstationarities in the RR-interval sequence, it is commonly detrended. Further, an FIR bandpass filter (0.05-0.6 Hz) is applied to the sequence, in order to remove frequency components that are not of interest with respect to patient physiology. The detailed specifications of the filter are listed in Table 3.5.2. The power spectral density of the RR-interval sequence was calculated using the STFT, as described in Section 3.2. A representative power spectrum, alongside the ECG segment from which it originated, is shown in Fig. 3.3.

3.5.3 Window Specifications

Conventionally, the lowest frequency of interest in HRV analysis is 0.04 Hz. Such a frequency can be identified in the spectrum by using windows of a length of at

Type	FIR
Stopband	DC-0.0346 Hz and 0.866 Hz-fs/2
Passband	0.05-0.6 Hz
Sampling frequency	4Hz
Attenuation	40 dB
Order	581

Table 3.1 – Bandpass filter specifications.

least $T = \frac{1}{0.04Hz} = 25s$.

Further, for an evenly-sampled signal, the Nyquist frequency is $f_c = \frac{1}{2\Delta t} = \frac{N}{2T}$, which must be greater than the highest frequency component in in signal that we wish to resolve. Therefore, as the upper limit of the bandpass filter is at 0.866 Hz, the window must also satisfy $\frac{N}{2T} \geq 0.866$.

The accuracy of the HF peak measure is highly dependent on the accuracy of the breathing rate obtained from the strain gauge. Due to the presence of other peaks in the vicinity of the HF peak of interest, an inaccurate breathing rate can lead to detection of the wrong HF peak and ultimately an incorrect estimation of the HF power. From the strain gauge signal, the instantaneous breathing rate can be established and the choice of window length determines the mean respiratory rate which is used to find the HF peak.

Based on these criteria, a window length of one minute was chosen, primarily in order to support correct HF peak detection by reducing the period over which the respiratory rate is averaged. This results in 240 points per window, and 256 points are used for the FFT. Consecutive windows have an offset of 30 seconds, giving rise to an overlap of 30 seconds of data from one window to the next.

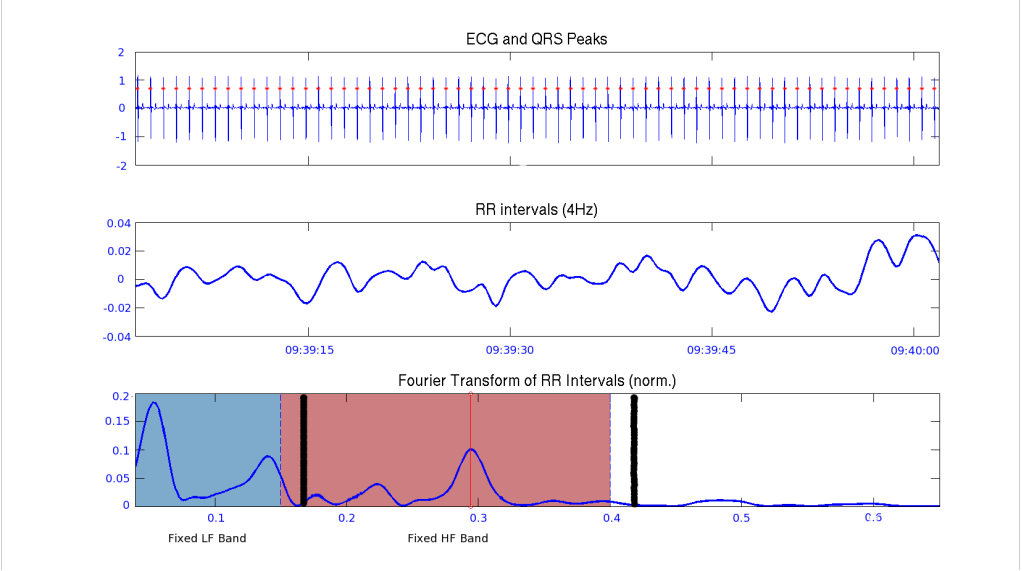


Figure 3.3 – Exemplary one-minute ECG segment and resulting power spectrum. A one-minute segment of ECG and the detected QRS peaks are shown in the first panel. The resulting RR-interval sequence, resampled at 4Hz, is shown in the second panel. The power spectrum in the third panel is obtained by applying the FFT to the RR-interval sequence. The fixed limits of the LF and HF bands are highlighted in blue, while the limits of the non-stationary HF band are highlighted in black on the frequency axis. The respiratory frequency, obtained from the strain gauge, is highlighted in red and denotes the centre of the non-stationary HF band. As the respiratory frequency is within the limits of the fixed HF band, the fixed and non-stationary HF bands overlap and contain approximately the same power.

3.5.4 Non-Stationary Peak Determination

As explained in Section 3.5, the non-stationary LF measure is defined as the integral of the power in the highest peak in the LF range (0.04 - 0.15 Hz), with the limits of integration being the two minima on either side of the dominant peak. The non-stationary HF measure is based on the respiratory frequency, which is obtained from strain gauge measurements. Once the peaks of the strain gauge signal are detected, the reciprocal of the time difference between the peaks gives the instantaneous respiratory frequency. As with the RR-interval signal, the times-

tamp of the interval is chosen as the timestamp of the earlier peak. The 5-minute mean of one-minute medians of the respiratory frequency signal is used to establish the location of the non-stationary HF peak in each 5-minute window. Since both the respiration rate and the power spectrum are averaged over five minutes, it is possible that the relevant HF peak is not exactly aligned with the respiratory rate. Therefore, the two peaks in the frequency spectrum closest to the respiratory frequency are identified using a peak detection algorithm. To decide which of the two peaks is the relevant HF peak, both the amplitudes of the peaks as well as their distance to the respiratory frequency are taken into account. The ratio of the amplitude of the peak to its distance from the respiratory frequency is used to find the relevant HF peak; the peak with the highest value of ratio is selected as the HF peak. The limits of integration are the two minima surrounding the chosen peak.

Figure 3.4 shows an example of how the respiratory frequency medians correspond to the high-frequency peaks of the power spectra. To compensate for sudden transients, a 9-point median filter is applied to the time series.

3.5.4.1 Normalisation

The conventional LF and HF measures based on fixed frequency bands are commonly given both in absolute terms and as a proportion of total power, where the total power is equivalent to the sum of the LF and HF powers [135]. By using non-stationary definitions of LF and HF that are limited to the dominant peak of each band, two options arise for normalising the measures: the LF and HF measures could either be normalised by the entire power of the spectrum, or by the sum of the LF and HF powers only. Given that the physiological meaning of the

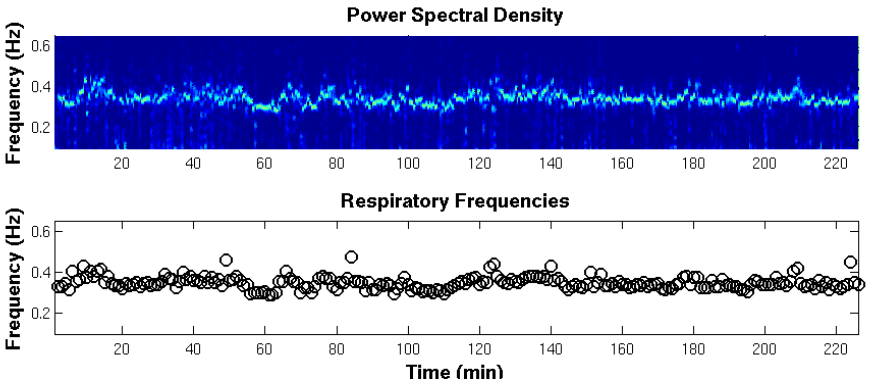


Figure 3.4 – One-minute median respiratory frequencies alongside the power spectral density. The high-frequency peaks of the power spectrum, obtained by spectral analysis of the ECG and shown in the first panel, correspond well with the one-minute respiratory frequency medians, obtained from the strain gauge signal and shown in the second panel. This confirms that the high-frequency peaks arise as a consequence of respiratory sinus arrhythmia, which in turn corresponds with the respiration frequency.

spectrum beyond the designated LF and HF peaks is unknown, it appears more sensible to exclude it entirely from the LF and HF measures and hence from the normalisation factor. Further, including the total spectral power in the normalisation factor can give rise to misleading changes in the LF and HF measures due to changes in the total power, when in fact the LF and HF power may not have changed. By normalising using the sum of the LF and HF power, each measure is effectively expressed in relation to the other, with the sum of the two equalling 1:

$$\frac{LF}{LF+HF}+\frac{HF}{LF+HF}=1.$$

In the following plots, the normalised LF measure is given as a percentage of $LF/(LF+HF)$. A representative power spectrum illustrating the need to exclude the power of subsidiary peaks is shown in Fig. 3.6.

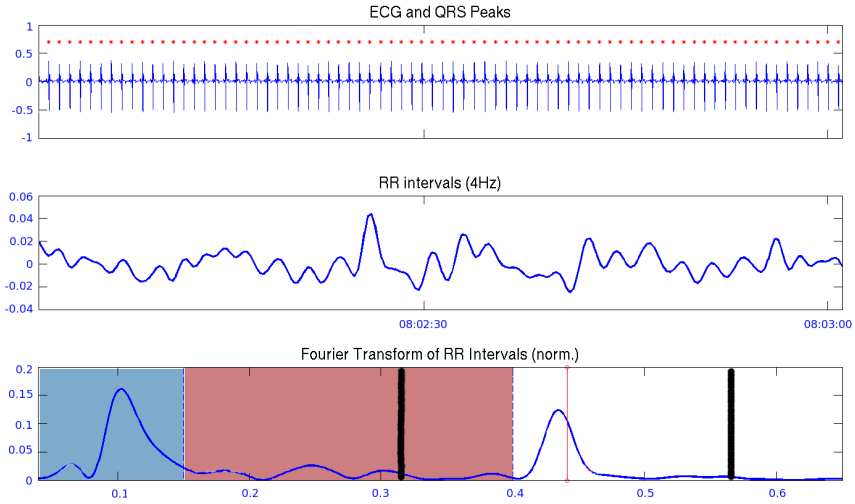


Figure 3.5 – Power spectrum for a patient whose respiratory frequency is higher than the upper limit of the fixed HF band. The power spectrum in the third panel shows a dominant peak at the respiratory frequency, which is higher than the upper limit of the fixed HF band. The limits of the fixed LF and HF bands are highlighted in blue, while the limits of the non-stationary HF band are highlighted in black on the frequency axis. The respiration frequency at the centre of the non-stationary HF band is highlighted in red. The non-stationary HF band includes the RSA peak in the spectrum, while the fixed HF band does not. As a result, the power measured by the non-stationary HF band is very different to that contained in the fixed band.

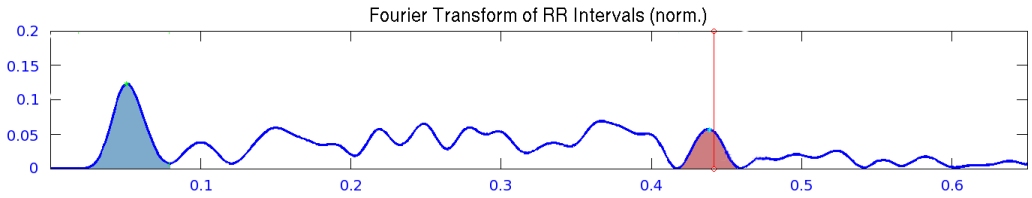


Figure 3.6 – LF and HF power computed from the relevant peaks only. The power spectrum shows two dominant peaks, one in the low-frequency range and another at the respiratory frequency. There are other peaks in the spectrum, whose physiological origin and meaning is unknown. The suggested LF and HF measures therefore include the relevant peaks only. The relevant LF peak is defined as the highest peak in the LF range (0.04 - 0.15 Hz), while the HF peak is defined as the peak closest to the respiratory frequency. The two minima surrounding each peak are used as the limits of integration. By taking this approach, only the peaks whose physiological relevance is known are included in the measure.

3.6 Application to Haemodialysis Data

The minute-by-minute LF power measure, as described in Section 3.5, is considered over entire haemodialysis sessions from the database described in Chapter 2 and compared with the conventional fixed-band measure. Figures 3.9 to 3.14 contrast the evolution of these measures in time with the changes in mean arterial blood pressure (MAP). In principle, three types of responses to HD can be distinguished based on this analysis:

- ◇ Sessions in which the patient is not adversely affected by fluid removal;
- ◇ Sessions in which the patient experiences a decrease in blood pressure and the patient's cardiovascular system is able to compensate for the fall in blood pressure by a rise in heart rate;
- ◇ Sessions in which the patient experiences a decrease in blood pressure and the patient's cardiovascular system is not able to compensate for the fall in blood pressure by a rise in heart rate.

In the following sections, the results from the non-stationary peak measure are first compared with those from the conventional fixed-band measure and the disadvantages of the fixed-band approach are highlighted. Subsequently, the three types of responses mentioned above are presented in more detail and the value of the normalised LF measure based on non-stationary peak power for distinguishing between the different types of responses is explored. With the exception of the sessions presented in Section 3.6.5, none of the sessions presented had nursing interventions. The plots presented each consist of four panels, showing the

overall one-minute frequency spectra in time, and the normalised and median-filtered LF measures, as well as the heart rate and blood pressure for the entire session, respectively.

3.6.1 Comparison of the Fixed-Band and Non-Stationary Peak Measures

The results presented in this section indicate that both the fixed-band and the non-stationary peak LF measure generally follow the same trends throughout the dialysis session, as can be seen in the second panel of Figs. 3.9-3.14. At times both measures are in full agreement, but in the majority of cases the normalised LF measure based on fixed frequency bands gives lower percentages of total power for the LF component than the non-stationary peak measure. Less frequently, the fixed-band measure yields higher percentages of LF power than the non-stationary peak measure. As described in Section 3.5, peaks other than the physiologically relevant LF and HF peaks are often present in the frequency spectrum due to noise or as harmonics of the main peaks or for other yet unknown physiological reasons and it is the presence of these peaks that determines the relationship between the fixed-band and non-stationary peak measure. In particular, it should be noted that for the purpose of comparison, the results are given in terms of normalised LF power in both cases, where the normalisation factor is the sum of LF and HF power. In the case of the fixed-band measure, this is effectively the LF power normalised by the total power within the spectrum, while for the non-stationary peak measure this is the power in the LF peak normalised by the sum of power in the designated LF and HF peaks.

In the absence of subsidiary peaks in a given one-minute window of RR intervals, as is the case in Fig. 3.7, both measures are in disagreement since the inclusion of additional frequencies in the fixed-band does not add power to the measure. More frequently however, there are subsidiary peaks in the spectrum which give rise to discrepancy between the two measures depending on their location and height. For instance, if additional prominent peaks appear in the fixed LF band, this can give rise to an increase in the percentage of power in the fixed-band LF measure relative to the non-stationary peak measure which would not be affected by such an additional peak. This is illustrated in Fig. 3.8. Equally, additional peaks in the fixed HF band result in a reduction in the percentage of power in the fixed-band LF measure.

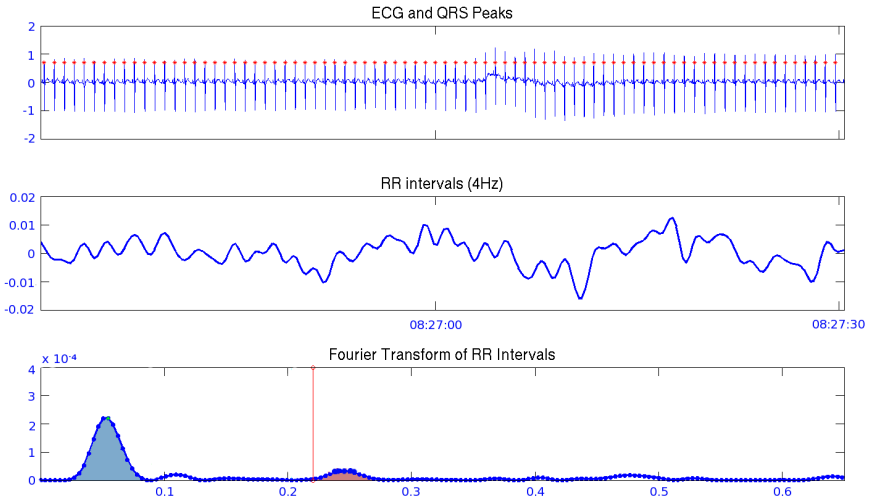


Figure 3.7 – Example frequency spectrum of a one-minute window of RR intervals with no subsidiary peaks. The first panel shows a one-minute segment of ECG and the QRS complexes used to establish the RR interval sequence, which is shown in the second panel. The third panel shows the Fourier transform of the RR interval sequence. The power is primarily distributed between the relevant LF and HF peaks. In such cases, the fixed-band and the non-stationary peak LF measures fully correspond.

In this regard, it should be noted that the fixed LF and HF bands are of different widths, namely 0.11 and 0.25 Hz respectively. Assuming that the location and height of subsidiary peaks is random, the difference between the widths of the fixed bands will on average result in a lower normalised fixed-band LF power in comparison with the normalised LF power from the non-stationary peak measure because there will, on average, be more additional peaks in the fixed HF band than in the LF band.

Therefore, while both measures follow the same trends when measuring the normalised LF power during dialysis, the fixed-band measure gives lower proportions of LF power when compared with the non-stationary normalised LF mea-

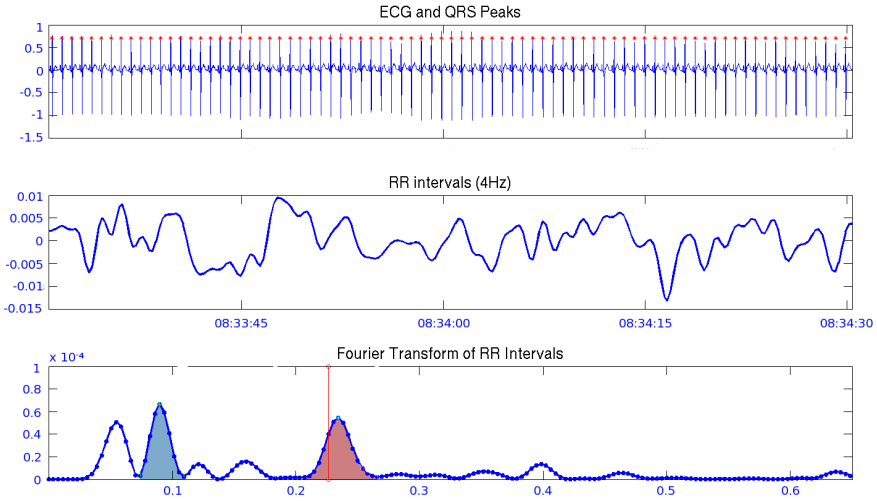


Figure 3.8 – Example frequency spectrum with additional peak in LF range. The additional peak, which can arise, for instance due to noise, increases the proportion of LF power in the absence of physiological changes.

sure. This is due to the inherent difference in size of the fixed bands, which results in a shift in the power distribution in favour of the larger HF band. As a result, only the non-stationary peak LF measure is considered in the following.

3.6.2 Stable Sessions

Stable sessions are characterised by stable vital signs, most importantly with regard to blood pressure. In both examples of stable sessions presented below, both the heart rate and the MAP remain constant throughout the session despite fluid removal. In the first example (patient OB41, Fig. 3.9), the MAP fluctuates between 90 mmHg and 112 mmHg while the heart rate ranges between 90 and 147 bpm. In the second example (patient OB32, Fig. 3.10), the MAP varies between 112 and 132 mmHg while the heart rate fluctuates between 72 and 113 bpm. In

both examples, the normalised LF peak-based measure remains close to 100% throughout the entire session, with the lowest value being 85% for OB41 and 86% for OB32. This indicates that in each window the LF peak constitutes the majority of the power in comparison to the HF peak, and as the LF power is taken to be indicative of sympathetic activity, this indicates that the main activity of the autonomic nervous system is in the sympathetic branch.

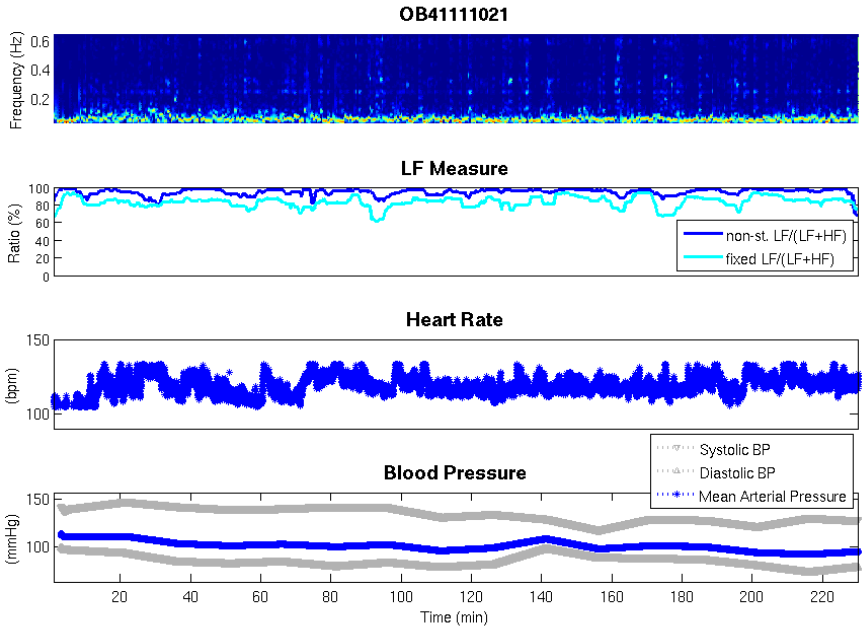


Figure 3.9 – OB41: An example of a patient with stable blood pressure and heart rate throughout the session. At the onset of the session, there is a significant rise in heart rate from 90 bpm to 144 bpm within the first 30 minutes, as can be seen in the third panel. This indicates an immediate physiological response to the start of fluid removal and is an indication of stability. Further, the blood pressure remains stable throughout the session. The normalised LF measure remains close to 100% for the duration of the session, indicating a high level of sympathetic activity. This agrees with the rapid response at the onset of fluid removal and the maintenance of both the blood pressure and the heart rate at stable levels.

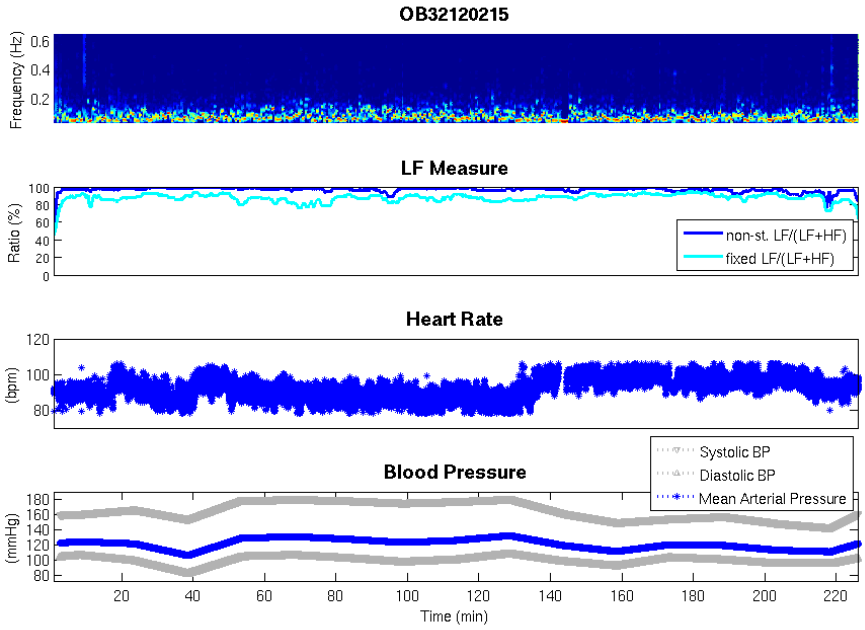


Figure 3.10 – OB32(a): An example of a patient with stable vital signs throughout the session. The blood pressure remains within a range of 20 mmHg for the duration of the session, as shown in the fourth panel. The heart rate also does not show large long-term variations, but fluctuates in the short term between 80 and 100 bpm. In the only instance where the blood pressure falls by more than 20% within 30 minutes, the heart rate increases from an average value of 90 bpm to 105 bpm, which is the desired physiological response. Throughout the session, the LF measure remains close to 100%, indicating a high level of sympathetic activity.

3.6.3 Sessions with a Compensatory Response

The second type of sessions are those in which there is a decrease in blood pressure upon fluid removal, which is compensated for by a rise in heart rate, also known as tachycardia, (see Section 1.5 for more details). This is a favourable physiological response from a clinical point of view, and does not necessitate an intervention from the nursing staff. In identifying whether a patient who is experiencing a decrease in blood pressure is at risk of having a hypotensive event, it

is crucial to distinguish these patients whose cardiovascular system is capable of compensating for the fluid loss and resulting blood pressure decrease, from those whose cardiovascular system cannot compensate.

Patient OB32 also has a characteristic tachycardic response to fluid removal and subsequent blood pressure decrease, as shown in Fig. 3.11. In the first 80 minutes of the session, the mean arterial blood pressure is constant at 108 to 118 mmHg, while the heart rate shows short-term fluctuations around an average value of 90 bpm. In the following 140 minutes until the end of the session, the blood pressure monotonically decreases from 118 mmHg to 87 mmHg. Simultaneously, the heart rate monotonically increases from an average value of 90 bpm to nearly 120 bpm towards the end of the session. As in the previous example, the LF measure remains at almost 100% throughout, reaching its lowest value of 82% 180 minutes into the session. While the blood pressure further decreases after 180 minutes, the heart rate no longer increases. This coincides with the decrease of the LF measure.

3.6.4 Sessions without a Compensatory Response

As explained in Section 1.5, IDH typically arises as a consequence of inadequate compensatory mechanisms in response to fluid removal during HD. Most commonly, this manifests itself in the lack of a tachycardic response to a decrease in blood pressure. The sessions considered in this section have this feature in common. In the first example (patient OB34, Fig. 3.12), the patient shows the opposite response to what would be desirable from a clinical point of view: as the blood pressure decreases, the heart rate decreases in parallel. At the beginning

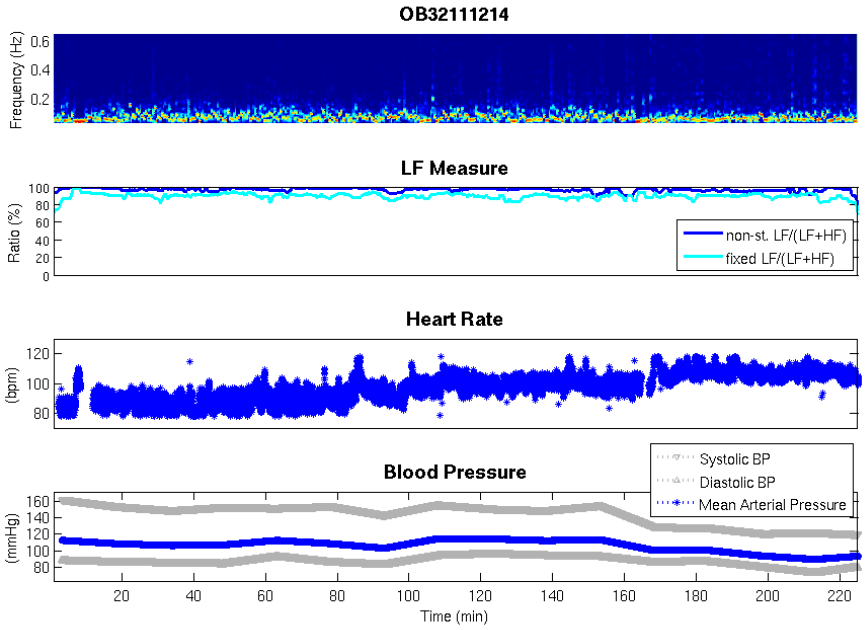


Figure 3.11 – OB32(b): An example of a characteristic response to a dialysis-induced decrease in blood pressure. While both the blood pressure and the heart rate are constant in the first 80 minutes of the session, the mean arterial blood pressure subsequently decreases from 110 to 87 mmHg while the heart rate rises correspondingly from an average value of 90 bpm to 110 bpm. In the last 50 minutes of the session, the blood pressure further decreases while the heart rate does not rise. The LF measure is very close to 100% throughout the session.

of the session, the patient's MAP is 80 mmHg but it decreases from the onset of the session and reaches its lowest value at 47 mmHg after 100 minutes. During this time, the heart rate which is between 85 and 90 bpm at the beginning of the session starts to decrease after 30 minutes and is at around 80 bpm when the blood pressure is lowest. It further decreases, reaching values as low as 75 bpm towards the end of the session. The blood pressure slightly recovers in the thrid hour of the session and remains constant for the remainder of the session at around 60 mmHg.

In contrast with the previous sessions, the LF measure varies significantly

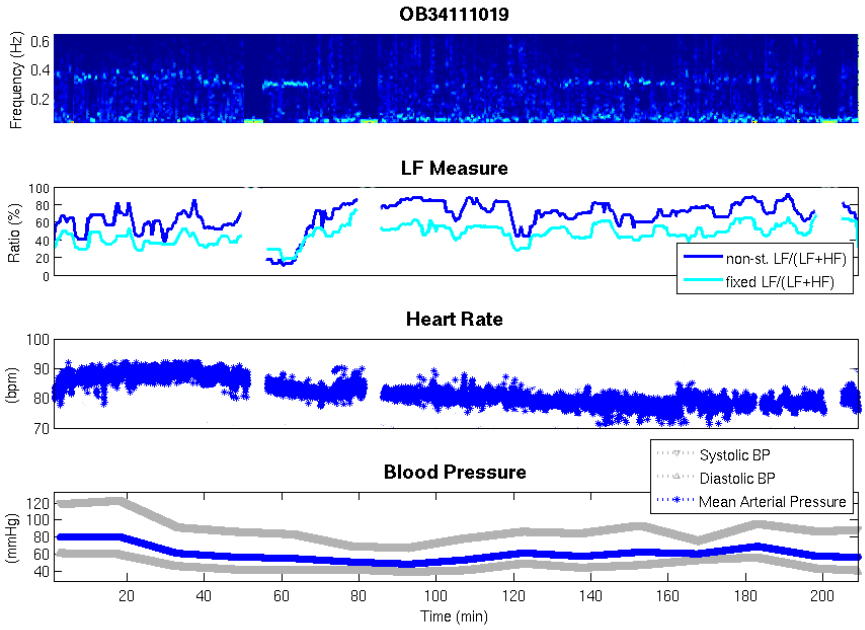


Figure 3.12 – OB34: An example of a lack of rise in heart rate in response to a decrease in blood pressure. While the blood pressure decreases monotonically in the first two hours of the session, the heart rate also decreases from an average value of 85 bpm to 76 bpm. In the second half of the session, the heart rate remains constant and the MAP recovers slightly to 60 mmHg. In contrast to previous examples, the LF measure varies significantly throughout the session. It reaches its lowest value one hour into the session at 10%, having previously fluctuated between 25% and 85%. It subsequently rises to values between 80% and 100% for an hour, but rapidly decreases again to 38% and maintains an average value of 70%, although fluctuating, until the end of the session.

throughout the session, most notably 60 minutes into the session when it reaches its lowest point at 10%. The LF measure is lowest when the heart rate is decreasing most in conjunction with a decrease in blood pressure. In the first half hour of the session, while the heart rate is constant, the blood pressure is decreasing and the LF measure fluctuates between 60% and 95% prior to the sudden decrease after 60 minutes. After this rapid decrease in the LF measure, the heart rate slightly recovers from 80 bpm to 85 bpm but then decreases further, while the LF measure

increases again and fluctuates between 55% and 95% until the end of the session.

Patient OB10 does not have a typical tachycardic response to a decrease in blood pressure either, as the two examples in Figs. 3.13 and 3.14 illustrate. In the first example (Fig. 3.13), the MAP is constant in the first 90 minutes of the session at around 85 mmHg, having decreased from its initial value of 106 mmHg. It then further decreases to 44 mmHg within the following 30 minutes and only slightly recovers in the final hour of the session to 50 mmHg. In the first 90 minutes of the session, while the blood pressure is constant, the heart rate increases from 70 bpm to 80 bpm, but does not rise much further when the blood pressure decreases, apart from individual brief peaks at 85 bpm. Towards the end of the session, the heart rate decreases to 75 bpm, and the heart rate thus remains between 70 and 80 bpm throughout the session, despite the 60% change in mean arterial blood pressure from 106 mmHg to 44 mmHg during the session. The LF measure fluctuates between 40% and 90%, and reaches its lowest value of 26% at two instances at the end of the session, at 195 and 215 minutes. This coincides with the lowest mean arterial pressure values. The LF measure fluctuates around 70% for the first 90 minutes of the session, varying between 50% and 80%. It then takes on lower values, varying between 40% and 75% in the following hour until if further decreases in the final hour. The gradual decrease in the mean value around which the LF measure fluctuates corresponds with the gradual decrease in blood pressure in the absence of a compensatory heart rate response.

In the second example (Fig. 3.14), the blood pressure remains constant at just above 80 mmHg in the first 90 minutes, with one decrease to 63 mmHg after 50 minutes. The heart rate does not increase significantly in the first 50 minutes of

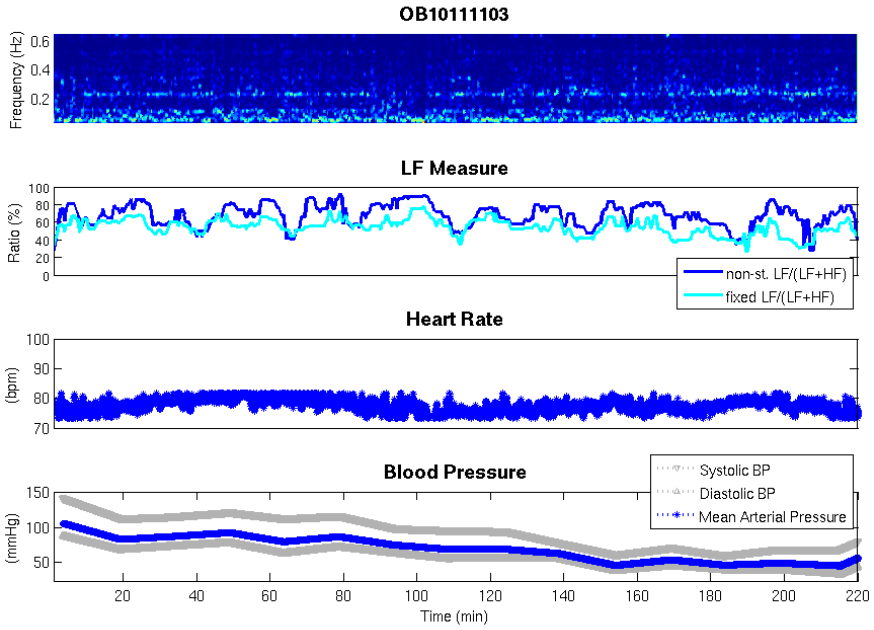


Figure 3.13 – OB10(a): An example of a decrease in the mean value of the LF measure in conjunction with a decrease in blood pressure. The heart rate increases at first in response to an initial decrease in blood pressure at the onset of the session, and the blood pressure remains constant for the first 90 minutes of the session. During this time, the LF measure fluctuates around a mean value of 70%. Subsequently, the heart rate decreases back to its initial value at the start of the session, and the blood pressure begins to decrease monotonically until the end of the session. Together with the decrease in blood pressure, the mean value around which the LF measure fluctuates decreases as well to around 55%, decreasing further in the final 30 minutes of the session.

the session, although it has a slightly higher rate of variability. The LF measure fluctuates between 50% and 70% in the first 40 minutes of the session, but rapidly decreases to 10% after 40 minutes and just prior to the 20% decrease in MAP to 63 mmHg. The LF measure then rises to values between 60% and 95% for the following 40 minutes, but again falls sharply to 25% once the blood pressure begins to decrease at 90 minutes. During the remainder of the session, the heart rate remains between 75 bpm and 80 bpm, with no significant increase despite

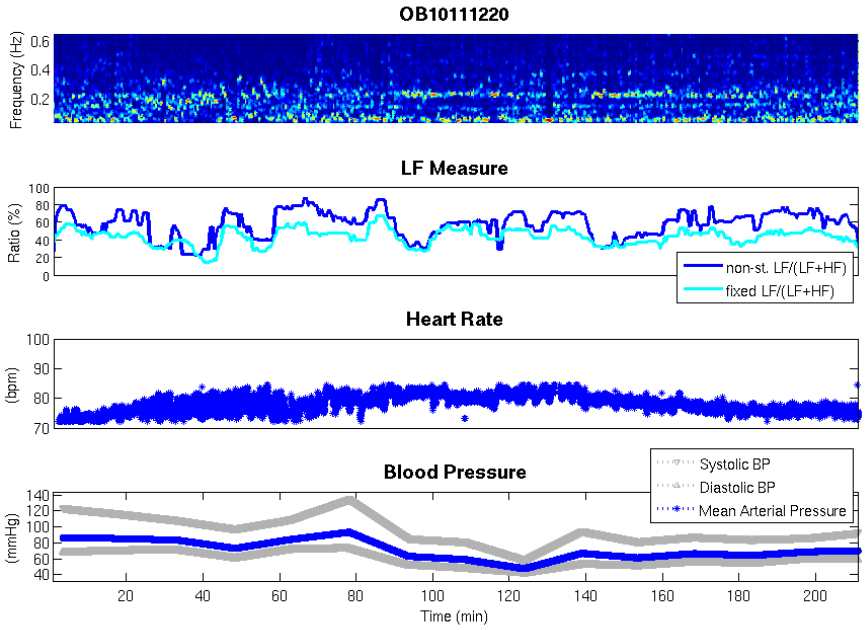


Figure 3.14 – OB10(b): Example of sharp decreases in the LF measure prior to decreases in mean arterial blood pressure. In the absence of an increase in heart rate in response to a dialysis-induced decrease in blood pressure, the LF measure decreases sharply at three instances, each preceding a decrease in mean arterial blood pressure greater than 20% (at 40, 90 and 120 minutes). Apart from these three troughs, the LF measure fluctuates around a mean value of 60% throughout the majority of the session, indicating a lower level of sympathetic activity than in the stable sessions and the sessions with a compensatory response (for other patients) presented in Sections 3.6.2 and 3.6.3.

the MAP decreasing from 79 mmHg to 42 mmHg during this time. The LF measure fluctuates around a mean value of 60% for the remainder of the session, after reaching a trough at 24% just prior to the mean arterial blood pressure reaching its lowest value of 42 mmHg. As is the case with the previous examples in this section, the overall fluctuations in the LF measure and hence in sympathetic activity are indicative of the cardiovascular system's inability to compensate for the decrease in blood pressure in light of fluid removal.

3.6.5 Effects of Interventions

The immediate response to IDH in patients undergoing haemodialysis is to place the patient in Trendelenburg's position, in which the body is laid flat on the back and the legs are elevated at an angle of 15 to 30 degrees, in order to maintain blood perfusion to the vital organs. Further nursing interventions include infusion of isotonic saline and reduction of the UF rate in order to raise the circulating blood volume [136]. Little work has been done on investigating how such interventions affect HRV, and whether changes in HRV can be used to assess the impact of the intervention on the physiology of the patient. Cavalcanti et al [93], whose analysis is most similar to that presented in this chapter and includes a time series of spectral power calculations in the fixed LF and HF bands, only considered sessions that were IDH and intervention free. In work done by Barnas et al [49], although 22 of the 48 sessions involved hypotensive events and interventions, no analysis was performed on the effects of interventions on the LF and HF measures. In this section, two examples of interventions in response to IDH are presented in conjunction with the LF measure, as presented in Section 3.5.

In the first example (Fig. 3.15), the patient experiences a first decrease in MAP from 100 mmHg to 72 mmHg within the first 20 minutes of the session, which is met by a corresponding rise in heart rate from 75 bpm to 90 bpm. The heart rate subsequently decreases slightly to 85 bpm, while the blood pressure remains constant around 70 mmHg. After 120 minutes, the patient experiences another decrease in MAP to just above 40 mmHg. Twenty minutes after this decrease, with the blood pressure not recovering, an intervention takes place in the form of a decrease in the UF rate to its minimal value and the UF rate is held at this

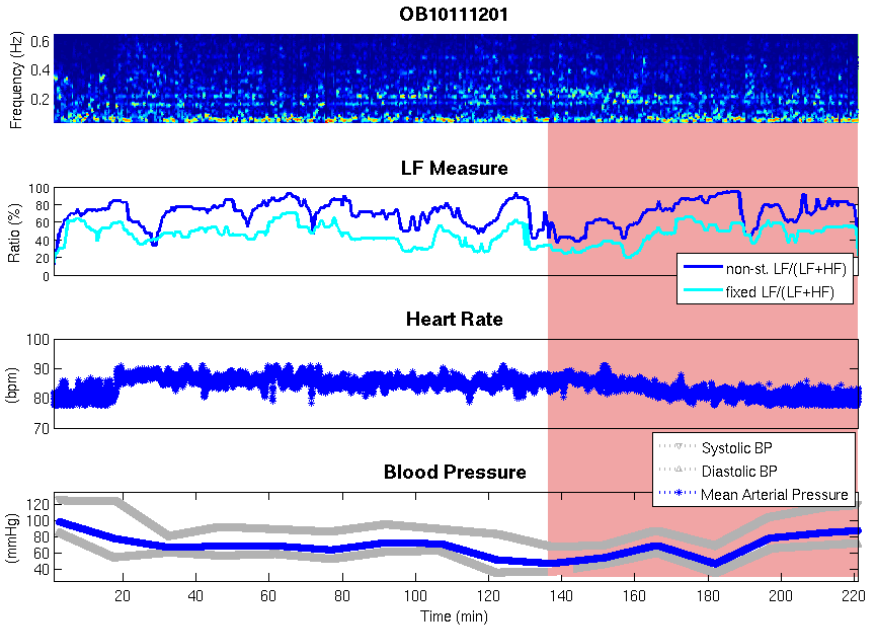


Figure 3.15 – Effects of a permanent reduction in the UF rate on the LF measure. In this session, the UF rate is reduced to and kept at its minimum value in the final hour of the session following a sudden decrease in blood pressure in the third hour of the session. The background of the plot is shaded in red to indicate the intervention. Following the intervention, the LF measure rapidly increases to above 90%.

value until the end of the session (shaded in red on Fig. 3.15). A peak in the heart rate, reaching values of 95 bpm can be observed just before the intervention takes place.

In the first 120 minutes of the session, the LF measure varies between 70% and 95%, but decreases markedly to just below 50% once the MAP falls to around 40 mmHg. Just after the intervention, at which time the LF measure has reached its lowest value, the LF measure rapidly rises to above 90%. Following the intervention and the rise in the LF measure, the MAP rises to just above 70 mmHg, but subsequently falls back to 46 mmHg as the heart rate slowly decreases following

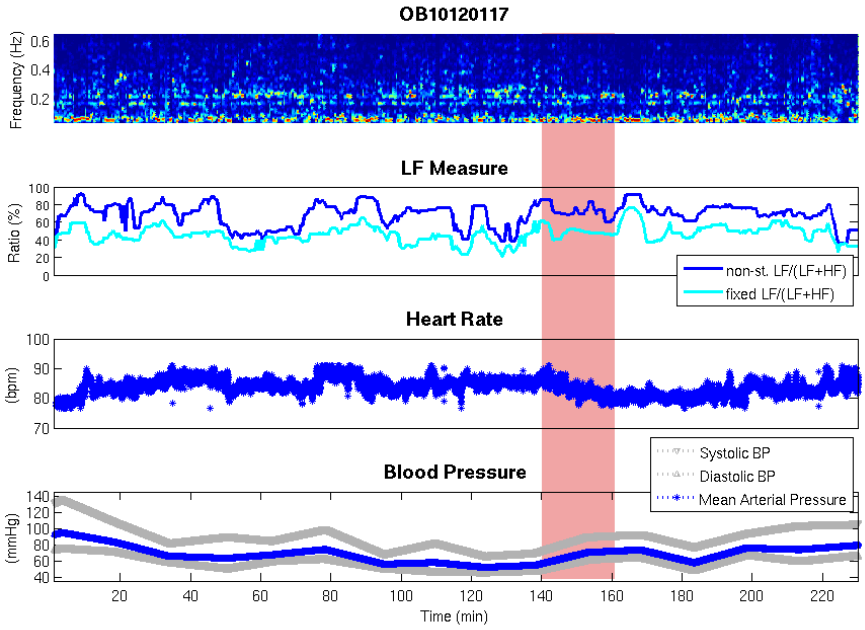


Figure 3.16 – Effects of a temporary reduction in the UF rate on the LF measure. Following a rapid decrease in blood pressure after two hours of dialysis, the UF rate is reduced to its minimum value for 20 minutes (from minute 140 to 160, shaded in red). During this time, MAP recovers from 50 mmHg to almost 80 mmHg and the heart rate decreases slightly from 85 bpm to 80 bpm. Following the intervention, the LF measure rapidly increases from 40% to values above 80%. However, once the intervention is over and the UF rate is restored, blood pressure again decreases while the heart rate gradually increases, possibly as a compensatory response. In contrast to the large fluctuations in the LF measure prior to the intervention, the LF measure remains relatively constant after the intervention.

the intervention. This fall in blood pressure is accompanied by a fall in the LF measure to 76%. MAP then recovers to 81 mmHg at the end of the session, and the LF measure remains between 85% and 100% throughout the recovery.

In the second example (Fig. 3.16), the blood pressure is initially just below 100 mmHg at the beginning of the session, but falls to 50 mmHg within the first two hours of the session. During this time, the heart rate initially rises from 80

bpm to 90 bpm but subsequently decreases again to 85 bpm. In the second hour, despite the blood pressure continuously decreasing, the heart rate does not change apart from a temporary rise to 90 bpm for approximately 10 minutes. Towards the end of the session, the heart rate decreases back to its initial value of 80 bpm and remains at this value until the final 30 minutes in which a slight increase in the heart rate is observed. Over the course of the session, the LF measure fluctuates between 50% and 97%, but reaches its lowest value just below 40% when the blood pressure is also at its lowest value. At this point, an intervention consisting of a reduction in the ultrafiltration rate to its minimum occurs and a marked increase in the LF measure to almost 100% is observed, alongside a recovery in MAP to 78 mmHg.

In both examples, both the blood pressure and the LF measure are at their lowest value just before the intervention takes place. Following the interventions, the LF measure rises significantly to very high percentage values in both cases. Also, in both cases subsequent decreases in blood pressure are preceded by decreases in the LF measure. It appears therefore, that the LF measure is sensitive to interventions and that reductions in the ultrafiltration rate serve to restore sympathetic activity to raise the blood pressure when the body faces volume depletion. However, it should be noted that these preliminary conclusions are based on a small number of cases in which interventions took place, and a larger study would be required to draw more detailed conclusions on the physiological effects of interventions and the possible value that the LF measure presented in this chapter could add to the analysis.

3.7 Patient-Specific Analysis

In light of the patient-specific approach adopted later in this thesis, it is worth investigating whether the non-stationary spectral measure is capable of distinguishing between the different physiological responses in individual patients. For this purpose, the LF measure is calculated for normal and abnormal sessions for three patients from the subset described in Section 2.5.

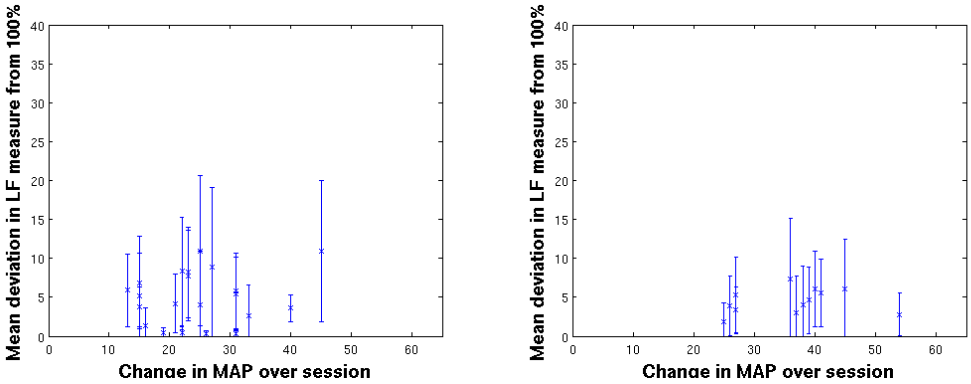
The results of this chapter suggest that lower values for the LF measure are indicative of a lack of the usual physiological compensation mechanisms, particularly of a rise in heart rate in response to a fall in blood pressure. Therefore, in order to quantify the efficacy of the LF measure in distinguishing between stable sessions with adequate physiological compensation for fluid removal and unstable sessions, the average change of the LF measure from 100% is calculated for each session considered. The resulting mean changes together with the standard deviation of the change are plotted for each session and sorted by the difference between the greatest and smallest MAP value for the respective session. The results for patients OB 7, OB 17 and OB 19 are shown in Figure 3.17.

Based on an orthostatic test for autonomic dysfunction (described in more detail in Section 4.5.2), patient OB 7 is classified as a patient with autonomic dysfunction, while OB 17 and OB 19 are not. This is reflected in the degree of variation in the LF measure from 100% against the change in MAP over the course of each session. The plots for OB 17 and OB 19, shown in Figures 3.17a and 3.17b suggest that there is no correlation between the degree of change in the LF measure from 100% when plotted against the change in mean arterial blood pressure over the course of the session in patients without autonomic dysfunction. This

could be as a result of the fact that in these patients, the compensatory response of the body to fluid removal is adequate in light of the change in blood pressure. In contrast, the results for OB 7 shown in Figure 3.17c, show a division in the deviation in the LF measure from 100% between sessions in which the change in blood pressure is less than 50 mmHg over the course of the session, and those in which the change is greater. While the value of blood pressure 50 mmHg is most likely specific to this patient's physiology, it is interesting to note that such a threshold exists for the change in MAP, beyond which the degree of change in the LF measure from 100% over the course of the session is significantly greater. Assuming that the LF measure is reflective of autonomic function, this agrees with the observation that the physiological compensatory mechanisms in patients with autonomic dysfunction are less pronounced or not present.

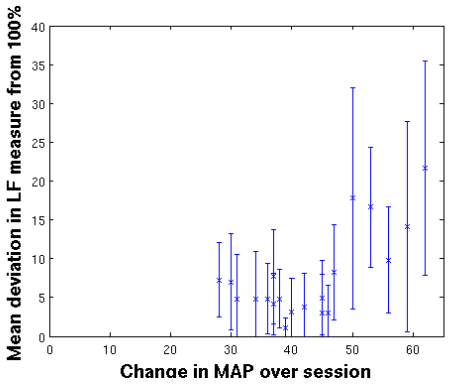
3.8 Discussion

Spectral measures of HRV have been established as a variable for independent risk assessment of mortality in a variety of conditions, including for instance diabetic neuropathy [106]. In HD patients, HD treatment has been shown to affect HRV independently of cardiac conditions [26]. In particular, fluid removal during HD has been shown to be a potent stimulator of the sympathetic nervous system by causing hypovolaemia which in turn leads to an unloading of cardiopulmonary receptors as part of the baroreflex feedback loop [137]. Previous work related to spectral measures of HRV in HD patients has exclusively considered fixed bands for calculating the LF and HF measures from the power spectrum of RR-intervals



(a) Patient OB17

(b) Patient OB19



(c) Patient OB7

Figure 3.17 – Average deviation in the LF measure from 100% against the change in mean arterial blood pressure over the course of the session for all sessions from three patients. The standard deviations are shown as bars around the mean values. Patient OB7 is one of the patients who matched the criteria outlined in Section 4.5.2 for autonomic dysfunction, while patients OB17 and OB19 are not known to have any such conditions. These results suggest that in patients in whom there is an adequate physiological response to fluid removal, the deviation in the LF measure from 100% does not appear to be correlated with the change in MAP over the session, a clear separation can be seen in patients with autonomic dysfunction.

either throughout the entire session using shorter sequences of the RR-interval signal to calculate a time series of spectral measures, or more commonly from a short sequence of data at the beginning and end of a session.

The approach taken in the work done to date on spectral HRV measures in HD patients is inadequate in a number of aspects, with respect to assessing the function of the autonomic nervous system in its sympathetic and parasympathetic branches. Firstly, the conventional measure used in all previous HD-related work is based on a fixed definition of bands for the LF and HF measures, namely 0.04-0.15 Hz for the LF band and 0.15-0.4 Hz for the HF band. The spectral power integral within these limits is used as the respective measure. However, this approach has two key disadvantages. Firstly, by fixing the limits of integration in the power spectrum, variations in the spectral power beyond the limits of integration are not taken account for. In particular, given that the dominant peak in the HF range represents the respiratory frequency and that the respiratory frequency can exceed 0.4 Hz (which corresponds to 24 breaths/min), fixing the upper limit of the HF range to 0.4 Hz can potentially result in the dominant peak not being included in the HF measure. This effectively renders the fixed-band HF measure unreliable, as, depending on the variations in the patient's respiratory frequency, it remains unclear whether the physiologically-relevant respiratory peak has been included in the measure or not. As such, the fixed-band HF measure is unsuitable for drawing inferences with respect to autonomic nervous system activity.

Further, it should be noted that the origin of other peaks in the frequency spectrum, with the exception of the dominant LF and HF peaks, is unclear [128]. Some have argued for the existence of a mid-frequency peak associated with baroreflex activity or a very-low-frequency peak associated with spontaneous oscillations in

arterial blood pressure [126]. Additional uncertainty as to the physiological meaning of these subsidiary peaks arises from spurious peaks due to periodic noise in the ECG or due to harmonics of the LF or HF peaks. Their inclusion in the LF and HF measures, therefore, cannot be justified.

Given that there is only one dominant peak in each frequency band that is of physiological relevance, a new method for measuring the activity in the sympathetic and parasympathetic branches of the autonomous nervous system has been introduced in this chapter. This method combines the approaches developed by Goren et al [125] and Bailon et al [124] outside of the HD context, in that it finds the relevant peaks, and computes only the power contained in these peaks (as defined by their surrounding minima) for the LF and HF measure, thereby only including the phyioslogically-relevant sections of the power spectrum.

With the exception of work done by Cavalcanti et al [93] and Pelosi et al [34], most previous work done on HRV in dialysis patients has either revolved around time-domain measures or has compared the fixed-band LF and HF measures for short periods at the beginning and end of the dialysis session. Although Cavalcanti et al [93] calculate the spectral measures in 3-minute windows throughout the dialysis session, their analysis does not include any hypotensive events. While the study performed by Pelosi et al [34] does include hypotensive events, it consists of a study population of only 12 patients. The method presented in this chapter represents the first application of a non-stationary spectral measure of HRV to HD data from an entire session. As the dataset includes hypotensive and non-hypotensive sessions, as well as sessions in which interventions took place, the predictive value of the method with regard to the incidence of IDH can be assessed in a wide variety of cases.

The standard interventions for patients experiencing IDH consist of a change in position, combined with either a saline infusion or a reduction in the UF rate, or both. The chief aim of HD however, is to remove the fluid accumulated in the patient's body in the interdialytic period, and the interventions described above run counter to this objective. In particular, if the dry weight target is not met repeatedly, this predisposes the patient to fluid accumulation in the body and chronic hypertension in the long term [6]. Equally, the incidence of IDH should be avoided as much as possible as repeated occurrences of IDH have been linked with an increased risk of mortality [7]. Therefore, from a clinical perspective, it would be advantageous to know when the patient's cardiovascular system will not be able to compensate for a decrease in blood pressure in response to fluid removal, thus requiring an intervention, and when the physiological compensation mechanisms are sufficient to avoid IDH, as this would optimize treatment outcome.

The results from this chapter indicate that the non-stationary spectral measure is a useful measure for distinguishing between patient sessions in which an intervention is required due to inadequate compensation mechanisms, and those in which no such intervention is necessary. In sessions in which no change in blood pressure is observed, the power in the LF peak constitutes almost 100% of the power throughout the entire session. In sessions in which the patient experiences a decrease in blood pressure, the LF measure allows us to distinguish between patients whose heart rate gradually rises in response to the decrease in blood pressure and those in whom the heart rate either does not rise adequately or even decreases together with the blood pressure. In the former, the LF measure takes on similar values to those in patients without a decrease in blood pressure, with the power in the LF peak representing more than 80% of the power in both the LF

and HF peaks. In the latter however, the LF measure commonly fluctuates heavily throughout the session, varying by more than 70% in the course of the session and taking on values as low as 10%. These fluctuations as such distinguish these sessions from those in which no IDH occurs. More specifically however, troughs in the LF measure appear to precede changes in the patient's physiology that are contrary to the expected compensatory response. For instance, the LF measure reaches its lowest value when the heart rate is decreasing most rapidly in conjunction with a decrease in blood pressure, preceding the lowest blood pressure reading in time. However, a more extensive study as to the timing of troughs in the LF measure with respect to the IDH events would be required to draw a more definitive conclusion in this regard.

In terms of autonomic nervous system activity, the results indicate that in patients in whom the blood pressure either does not decrease or decreases in conjunction with an adequate compensatory response, the vast majority of activity is in the sympathetic nervous system throughout the session. Although derived using different LF measures, our results agree with those of Cavalcanti et al [93], Rubinger et al [92] and Pelosi et al [34]. Cavalcanti et al found that the sympathetic activity was more dominant in stable patients compared with unstable patients, while Pelosi et al found the sympathetic activity to be more dominant in all patients at the beginning of the session, yet decreasing later in the session in unstable patients and dropping suddenly during hypotensive events. Similarly, Rubinger et al found the conventional fixed-band LF/HF ratio to be significantly lower in unstable patients during HD sessions when compared to that of the stable patients. While in all three cases patients were classified as stable or unstable based on their history of previous IDH episodes, our results indicate that the pa-

tient's response to fluid removal can be assessed using the LF measure introduced in this chapter. In patients who will encounter hypotensive episodes in the course of a session, fluctuations in the LF measure can be observed from the onset of the session. In contrast, in patients in whom the blood pressure remains stable throughout the session, or in cases where the patient is able to compensate for a decrease in blood pressure through a rise in heart rate, the LF measure is stable and high at close to 100% throughout the session.

Looking at the changes in the LF measure in the light of interventions, two conclusions can be drawn. Firstly, sharp decreases in the LF measure appear to be indicative of the need for an immediate intervention, as interventions by the nursing staff (who did not have the LF measure at hand at the time of data collection) occurred shortly after sharp decreases in the LF measure in response to a lack of recovery in the MAP from low values. Further, the effectiveness of the intervention can be assessed by the recovery of the LF measure. In practice, this can be used to reinstate the full UF rate as soon as possible after an intervention, namely once the patient's sympathetic activity has been restored to a level that can sustain further fluid removal. However, as previously noted, the above conclusions are based on a small number of interventions and a more extensive investigation of the effects and timings of interventions would be required to establish how interventions affect the LF measure.

As such, the LF measure introduced in this chapter adds predictive value to the assessment of the patient's physiology during the session, indicating both whether the patient's cardiovascular system will be able to compensate for the effects of fluid removal through dialysis and preceding instances in which the patient's physiology most deteriorates. However, while the results of the patient-

specific analysis indicate that the LF measure is useful for detecting a lack in the necessary compensatory mechanisms in patients with autonomic dysfunction, it appears that in patients without autonomic dysfunction the measure cannot distinguish between sessions with different degrees of change in MAP over the course of the session. It should be noted that this could alternatively be due to the fact that in the patients considered the compensatory mechanisms are adequate despite large changes in MAP, and therefore warrants further investigation.

Finally, two limiting factors of the study should be highlighted. Firstly, as the blood pressure is only measured every 15 minutes, it is difficult to establish whether, and if so by how long, the LF measure precedes the recorded decrease in blood pressure. In order to draw more meaningful conclusions as to the predictive value of the proposed measure, more frequent or possibly even continuous blood pressure recordings would be advantageous. Further, the accuracy of the LF measure is highly depended on the correct locatisation of the relevant peaks in the power spectrum. As the respiratory frequency is used for this purpose, the accuracy of the measure in effect depends on the quality of the respiratory signal obtained from a strain gauge or other sensor (for example impedance pneumography).

Chapter 4

High Dimensional Visualisation

Understanding high-dimensional datasets consisting for instance of multiple vital sign signals in time, is a difficult task. Visualisation techniques are a useful tool for mapping such data into a lower dimensional space in order to better understand the structure of the data. Further, as dialysis is a repetitive form of treatment, comparative projections of vital sign data can be employed to identify significant changes over time in a patient's physiology that may not be immediately obvious from univariate analysis. Typically, two or three dimensions are chosen to enable plotting. A common feature of all such projection techniques is that they aim to preserve as much as possible a key aspect of the data, such as for instance the interpoint distances in the original high-dimensional space. Several algorithms exist for this purpose, and are introduced below. In this chapter, the Neuroscale algorithm is applied to HD vital sign data such that the centre of the two-dimensional projection space represents normal physiology. This is achieved by first defining normal physiology in dialysis patients, and subsequently using this data to normalise and train the Neuroscale algorithm. Both a population-based approach, in

which vital sign data from the entire patient population is used as training data, as well as a patient-specific approach, in which a patient's own vital sign data is used to train a Neuroscale map specific to that patient, can be adopted. In this chapter, different approaches to training a population-based model of normality are introduced, and the advantages and disadvantages of each in identifying abnormalities in vital signs are discussed.

4.1 Methods for High-Dimensional Visualisation

Our high-dimensional dataset consists of four-dimensional vital sign data in time, as described in Chapter 2, with each dimension representing one vital sign parameter.

In this section, the main high-dimensional visualisation methods are presented, and motivation for the choice of method used in the remainder of the chapter is given.

4.1.1 Kohonen Self-Organizing Maps

The Kohonen Self-Organizing Map (SOM) is a topology-preserving method for high-dimensional data visualisation. Data vectors are typically fed into a multi-layer feedforward network, where each node represents a topographical location thus constituting a grid of discrete points in space [138, 139]. The weights of the nodes are adjusted in the training phase to group data on the grid such that data that is close in high-dimensional space remains close in the lower-dimensional representation. In this sense the SOM resembles the k-means algorithm [140], in that the assignment of data vectors to points on the grid is continuously updated,

i.e. ”self-organizing”. In contrast to k-means however, SOMs do not optimise a well-defined cost function [141] and as a result neither convergence nor the parameters of the model can be assessed objectively.

4.1.2 Generative Topographic Mapping

The Generative Topographic Mapping (GTM) algorithm can be interpreted as a probabilistic version of the SOM, as it also maps high-dimensional data onto a grid of discrete points in a lower-dimensional space. A low-dimensional (usually 2D) nonlinear manifold is fitted to the data in high-dimensional space using for instance a feed-forward neural network. The high-dimensional data is modelled by a mixture of radially symmetric Gaussians centred on the datapoints. The Expectation-Maximisation (EM) algorithm is then used to learn the parameters of the mapping and the variance of the Gaussians. Using these parameters, high-dimensional data can be projected into the lower-dimensional space for visualisation.

4.1.3 Sammon Mapping

Sammon mapping [142] is a data visualisation technique for mapping high-dimensional data onto two dimensions. The underlying criterion for the transformation is that the Euclidean interpoint distances in the 2D visualisation space should be as close as possible to the corresponding distances in the higher dimensional space. This is equivalent to minimising the so-called Sammon stress objective function E:

$$E = \frac{1}{\sum_{i=1}^{N}\sum_{j>1}^{N} d_{ij}^{*}} \sum_{i=1}^{N}\sum_{j>1}^{N} \frac{(d_{ij} - d_{ij}^{*})^2}{d_{ij}^{*}} \tag{4.1}$$

where d_{ij} denotes the Euclidean distance between two of N points $\mathbf{x_i}$ and $\mathbf{x_j}$ in an m-dimensional space, $d_{ij} = \sqrt{(\mathbf{x_i} - \mathbf{x_j})^T(\mathbf{x_j} - \mathbf{x_i})}$, and d^*_{ij} represents the distance between the corresponding two points $\mathbf{y_i}$ and $\mathbf{y_j}$ in the lower, d-dimensional visualisation space, $d^*_{ij} = \sqrt{(\mathbf{y_i} - \mathbf{y_j})^T(\mathbf{y_j} - \mathbf{y_i})}$ with $i, j = 1...N$. The stress measure E quantifies the difference between the configuration of points in the m-dimensional space and their configuration in the lower dimensional space. To preserve as much as possible the relationship between the high-dimensional distribution of points in the lower-dimensional space, the difference between the two configurations is minimised. The stress function gives a measure of distortion in the mapping.

The algorithm starts from an initial (e.g. random) configuration of points in the lower dimensional space and calculates the stress measure E, given by Eqn. 4.1. Then the configuration is optimised iteratively by a gradient descent method until a minimal E is found. Each calculation of E involves $O(n^2)$ computations, making the optimisation problem a computationally intensive operation with an increasing number of points. Another drawback of the algorithm is its lack of generalisation. Once an optimal mapping has been established for a number of points, it is not possible to add more points onto the map without retraining it due to the lack of an explicit mapping function [143].

4.1.4 Neuroscale

The Neuroscale algorithm [144] trains a Radial Basis Function (RBF) neural network to map high-dimensional data into lower dimensional space. The given high-dimensional data points $\mathbf{x_i}$ are the inputs to the neural network, the output being the points in the visualisation space, $\mathbf{y_i}$, and the objective function is the Sammon

stress measure E (Eqn. 4.1). Each hidden unit implements a radial-basis function ϕ. The outputs of the hidden units are linearly combined using weights W to produce an overall transformation given by $\mathbf{y_i} = \mathbf{f}(\mathbf{x_i}, \phi, \mathbf{W})$. The distances in visualisation space are given by $d_{ij} = ||\mathbf{f}(\mathbf{x_i}) - \mathbf{f}(\mathbf{x_j})||$, which can also be written as:

$$d_{ij}^2 = \sum_{l=1}^{n} \left(\sum_k w_{lk} \left[\phi_k \left(||\mathbf{x_i} - \mu_k|| \right) - \phi_k \left(||\mathbf{x_j} - \mu_k|| \right) \right] \right)^2 \tag{4.2}$$

where ϕ_k are the basis functions, μ_k are the centres of the functions and w_{lk} are the weights of the basis functions. This assumes Gaussian radial-basis functions of the form $\phi_k(||\mathbf{x_i} - \mu_k||) = e^{-\frac{||\mathbf{x_i} - \mu_k||^2}{2\sigma^2}}$, where σ denotes the width of the function. To minimize the objective function E, we combine Eqns. 4.1 and 4.2 and take the partial derivative $\frac{\partial E}{\partial w_{lk}}$ of the stress measure with respect to the weights. The weights can then be adjusted so as to minimize the stress measure. This is done using a gradient descent method.

New points can subsequently be added onto the visualisation map by simply feeding them through the trained RBF network. A schematic of the Neuroscale algorithm is shown in Figure 4.1.

One advantage of the Neuroscale algorithm over both the SOM and GTM algorithms is that in the Neuroscale algorithm the location of points in the lower-dimension projection is not confined to a discrete grid. Further, while the cost function for SOMs is not well-defined, thus making it difficult for convergence of the model to be assessed, the stress measure E and its partial derivative with respect to the weights of the basis functions of the neural network are well-defined.

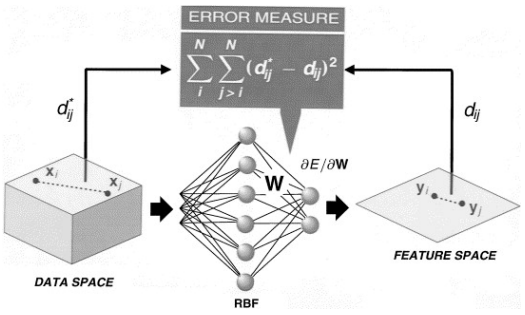


Figure 4.1 – Schematic of the Neuroscale algorithm. Reprinted from [145] with permission from Elsevier.

The chief advantage of the Neuroscale algorithm over Sammon mapping is its ability to generalise from a set of training data points to other points. Once the error measure E has been optimised for a given set of high-dimensional points, the resulting Sammon projection into lower-dimensional space cannot be extended to other points without re-calculating an optimised E. As previously mentioned, each calculation of E involves $O(n^2)$ computations for n points thus making it impractical to re-calculate E with an increasing number of points. Further, the lack of generalisation in Sammon mapping means that different subsets of high-dimensional data cannot be visually compared in lower-dimensional space. Our main motivation for applying a topology-preserving high-dimensional visualisation technique to HD vital sign data is to compare vital sign data from a given session with data considered to be normal. The generalisation property of the Neuroscale algorithm is crucial in this regard, and therefore the Neuroscale algorithm is chosen for the analysis in this chapter.

4.2 Visualisation of Patient Data

Four-dimensional vital sign data consisting of heart rate, oxygen saturation, mean arterial blood pressure and breathing rate measurements, based on the data collected from patients in the study described in Chapter 2, is visualised using the Neuroscale algorithm. Every vital sign measurement is normalised by the zero-mean, unit-variance transformation:

$$x_n = \frac{x - \mu_p}{\sigma_p}$$

where x_n is the normalised parameter value, μ_p is the mean of the parameter in the training data and σ_p is the standard deviation of the parameter in the training data. This removes scaling differences in the data and sets the mean of each transformed parameter to zero. Upon normalisation, $(0\,0\,0\,0)$ represents the "most" normal datapoint and the Euclidean distance of a vital sign projection to this point gives a measure of the abnormality of the data.

Conventionally, for the purpose of identifying abnormal changes in the data, the neural network is trained on normal data only [146]. As a result, abnormal vital sign data, that is combinations of vital sign parameters previously unseen during training, is mapped outside of the region defined by the normal data.

4.2.1 Extrapolation Properties

When presented with new data, the trained network projects data that is similar to the training data by interpolating between the training examples. However, when the data is significantly different to that used during training, the RBF network

extrapolates to outputs which are based on known activations of the hidden units. If only normal data is used to train the RBF network, (the definition of normal being discussed in Section 4.3), abnormal data will be projected outside the region defined by the training data. This allows us to distinguish between normal and abnormal data, but not to draw further conclusions regarding the abnormal data or the transition from normal to abnormal data.

4.2.2 Initialisation and Training

Each hidden unit in the RBF network implements a radial basis function characterised by a centre μ_k and a width σ. In the second layer, the outputs of the l hidden units are weighted by weights w_{lk}. When initialising the algorithm, the centres of the hidden units are chosen as a random subset of the training data to ensure correct localisation. The widths of the radial basis function are all set to $\sigma = 1$ and the initial weights of the output layer are drawn randomly from a standard normal distribution centred at zero with a standard deviation of $\frac{1}{30}$.

Depending on the number of sessions recorded for each patient, the training data can consist of thousands of training patterns. Training the neural network with this volume of data is a computationally intensive task and likely to be hampered by poor local minima. In order to reduce the number of training patterns with minimal loss of information, the data is clustered into a subset of prototype patterns using the k-means clustering algorithm [140, 147]. The objective of this iterative algorithm is to minimise the error function given by:

$$\varepsilon = \sum_{i=1}^{k} \sum_{x_j \in S_i} \left| x_j - m_i \right|^2 \tag{4.3}$$

where k is the number of clusters, S_i represents the i^{th} cluster and m_i is the centroid of the distribution of points in S_i. The algorithm is initialised by randomly setting k clusters in the data space and assigning each point to its nearest cluster centre. The positions of the cluster centres are recalculated by making the centroid of the data points assigned to the cluster the new cluster centre. The data points are then assigned to their nearest cluster centre and the process is repeated. This continues until convergence or until the change in ε between iterations is lower than a predefined value.

A choice of $k = 500$ for the number of prototype patterns is empirical, but has been shown to be appropriate in related models of normality [148]. Of these 500 prototypes, the 10% (i.e. 50) least populated clusters are removed on the basis that they are most likely to be outliers. Generally, for a single-layer RBF network the number of hidden units is chosen to be one order of magnitude greater than the dimensionality of the input data [149]. Further, the number of hidden units is typically chosen to be one order of magnitude less than the number of input data points [150], or 20% of the number of input data points [151] in order to ensure that the data space is adequately populated. Here, the dimensionality of the input data is 4, while the number of input data points is 450. The number of hidden units in the RBF network is therefore chosen to be 100. The stopping criterion for training the network is set such that the change in the error E between iterations is lower than a predefined value.

4.2.3 Quantisation

Oxygen saturation readings are quantised at integer values in the recorded dataset. To avoid streaks in the Neuroscale projections, normally-distributed noise centred at $\mu = 0$ with a standard deviation of $\sigma = \frac{1}{6}$ is added to the SpO_2 values.

The MAP is only measured every 15 minutes, while the three other parameters are sampled at least once per minute. For display purposes, the MAP values are linearly interpolated between readings. The interpolated signal is then sampled once per minute. As this is still a lower sampling frequency than that for the three other vital signs, a zero-order hold is applied between the interpolated MAP values.

4.3 Definition of Normality in Dialysis Patients

Haemodialysis is a necessary treatment modality for patients with Chronic Renal Failure, that serves to restore normal fluid balance and blood composition. In addition to the intradialytic complications described in Chapter 1, HD treatment can give rise to interdialytic complications as well. While the former is caused by fluid removal from the blood at a rate faster than can be compensated for by refill through interstitial fluid, the latter is caused by insufficient fluid removal, leading to long-term accumulation of fluid in the body (see Section 1.2.1). Therefore, depending on the volume of excess fluid in the patient's body at the onset of each session, dialysis can either restore normal fluid levels in a fluid-overloaded patient, or can reduce body fluid to levels that are not physiologically tolerable, leading to symptoms such as nausea and ultimately IDH. Achieving a balance between these two states is key to successful treatment of ESRD patients.

Given the repeated accumulation and removal of fluid and the associated complications described above, it is difficult to define normal physiology in HD patients. In fluid-overloaded patients, normal physiology will be achieved towards the end of the session, while in other patients the physiology at the onset of a dialysis session is most likely to be representative of normal physiology. In reality, normal physiology is probably achieved at some point in the course of the session, when normal fluid balance is reached. However, since dialysis sessions are scheduled two to three days apart, the objective is to remove as much fluid as physiologically tolerable, which is reflected in the concept of dry weight. Dry weight however, does not necessarily reflect a state of optimal fluid balance and more often corresponds to less than the optimal amount of fluid in the body. Further, as described in Section 1.2.1, a patient's dry weight is primarily established by trial and error in clinical practice, and the concept as such is not rooted in patient physiology. Given that the dialysis parameters, such as the UF rate, are adjusted at the start of each session to ensure that this weight is achieved by the end of the session, the dry weight, although a vague concept, constitutes the "physiological target" of every dialysis session. It is commonly estimated up to three months in advance for the following sessions, and although often re-adjusted prior to the start of each session, no adjustments are made during the session in response to the patient's physiological reaction to fluid removal, unless an adverse event occurs.

An alternative approach would be to make normal physiology the primary target of each dialysis session, instead of the estimated dry weight. Recent work by Javed et al [152] suggests that a state of normal physiology is achieved at the trough of the heart rate time series during dialysis, therefore implying that vital signs can potentially be used as target parameters in dialysis treatment instead of

a predefined weight. In the work described by Javed et al, an initial decrease and subsequent rise in heart rate was identified in patients in whom more than two litres of fluid were removed during dialysis. The initial decrease in heart rate can be interpreted as a recovery from the fluid-overloaded state at the start of dialysis, in which the heart rate is elevated by mechanoreceptors near the right and left atria in response to the high blood pressure from fluid accumulation. This mechanism is also known as the Bainbridge reflex [153]. The subsequent rise in heart rate most likely constitutes the characteristic compensatory response to fluid removal observed in HD patients. This was investigated in more detail in Section 3.6.3. As such, the time at which the patient's heart rate is at its lowest value during dialysis represents the best approximation to normal physiology.

The respective sessions from which the data constituting normal physiology is obtained must also be considered. Clearly, data from sessions in which IDH episodes occurred in accordance with the definition of IDH given in Section 1.4 should be excluded, even if the episode took place after the deemed point of normal physiology. The subset of normal data could be narrowed down further by only including data from patients who never experienced IDH, as opposed to including all non-IDH sessions. This is effectively a question of whether IDH is seen as a condition inherent to a given patient's physiology or as a result of incorrect dialysis prescription such as poor dry weight estimation or excessive fluid removal. In order to allow for comparison between the population-based and patient-specific models of normality that are introduced in the subsequent chapters, normal data was obtained from all IDH-free sessions regardless of whether or not the patient experienced IDH in other sessions.

4.4 Population-Based Model of Normality

In this chapter, a population-based model of normality is trained using the Neuroscale algorithm. In the population-based approach, normal data from the entire dialysis patient population is included in the training dataset. For this purpose, the 603 recorded sessions were divided into two subsets of IDH and non-IDH sessions, as described in more detail in Section 2.4. The distinction between IDH and non-IDH sessions is based on the definition given in Section 1.4. Only vital sign data from the non-IDH sessions was used to train the population-based model. Vital sign data from the 15 minutes of the session preceding the point of the lowest heart rate is chosen to be representative of normal physiology, as described in Section 4.3. The subset of non-IDH sessions consists of 300 sessions. Of these non-IDH sessions, 12 sessions were selected at random to constitute the test dataset. Therefore, in total, 15 minutes of vital signs from 288 sessions constitute the training dataset, giving rise to 2.5 x 10^5 training patterns. The means and standard deviations of this subset are given in Table 4.1. In comparison, the means and standard deviations of all IDH and non-IDH sessions are given in Table 2.5. As the training data by definition consists of the 15 minutes of data preceding the heart rate nadir, the mean heart rate of the training dataset is lower than both the mean heart rate of the IDH as well as that of the non-IDH sessions. There are no significant differences in the mean breathing rate and mean oxygen saturation values for the three data subsets. The mean value for the mean arterial blood pressure in the training data is significantly higher compared to both the IDH and non-IDH sessions. This is expected as the training data is supposed to represent normal physiology and is commonly followed by lower blood pressure values as

	μ	σ
HR (bpm)	74.41	15.13
SpO_2 (%)	94.2	4.72
BP (mmHg)	115.48	21.23
BR (breaths/min)	21.68	7.43

Table 4.1 – Means and standard deviations of the population-based training dataset.

further fluid is removed past this point of physiological normality.

Histograms showing the distribution of vital signs amongst the training data (shown in blue) in comparison with the entire IDH (shown in red) and non-IDH (shown in green) sessions are shown in Figure 4.2. The histograms are normalised such that the area of each distribution is 1. As the training data by definition consists of the 15 minutes of data preceding the lowest heart rate in non-IDH sessions, it is expected that the distribution of heart rates would be skewed towards low heart rates, as shown in the first histogram. With respect to the breathing rate and oxygen saturation histograms, no significant differences can be observed across the three subsets. With regard to the mean arterial blood pressure, the training dataset contains far fewer MAP values below 100 mmHg than both the IDH and non-IDH sessions, while it contains more readings at higher values, particularly above 120 mmHg.

The parameters given in Table 4.1 are used to normalise each vital sign measurement in the training dataset by the zero-mean, unit-variance transformation given by:

$$x_n = \frac{x - \mu_p}{\sigma_p}, \tag{4.4}$$

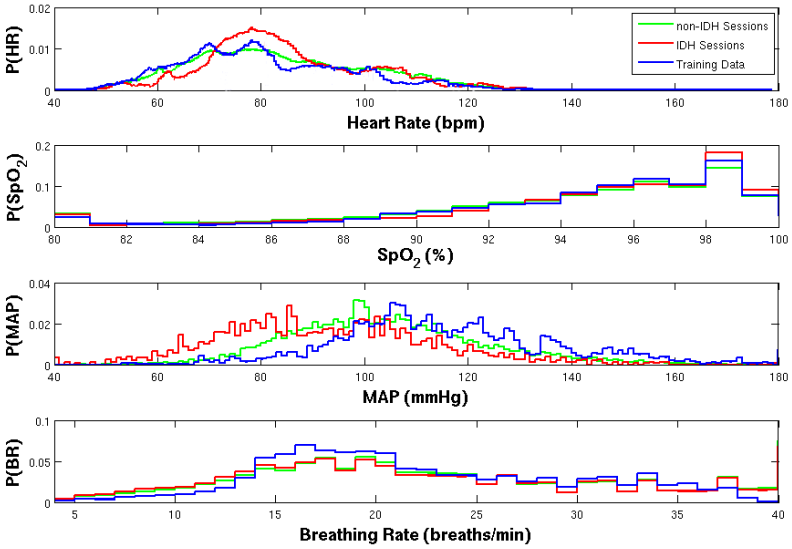


Figure 4.2 – Vital sign histograms of the training data (blue), compared with the IDH (red) and non-IDH (green) subsets. By definition, the training data is skewed towards low heart rate values, given that it is based on the 15 minutes of data preceding the heart rate nadir. No significant differences can be observed in the breathing rate and oxygen saturation distributions. The training data contains far fewer MAP values below 100 mmHg than both the entire IDH and non-IDH sessions. Instead, higher blood pressure values, particularly above 120 mmHg, are much more prevalent in the training data compared with the IDH and non-IDH sessions.

where x_n is the normalised parameter value, μ_p is the mean of the parameter in the training data and σ_p is the standard deviation of the parameter in the training data. This removes scaling differences in the data and sets the mean of each transformed parameter to zero. Thus, upon normalisation, $(0\,0\,0\,0)$ represents the most "normal" datapoint and the distance from the origin gives a measure of deviation of the data from normality.

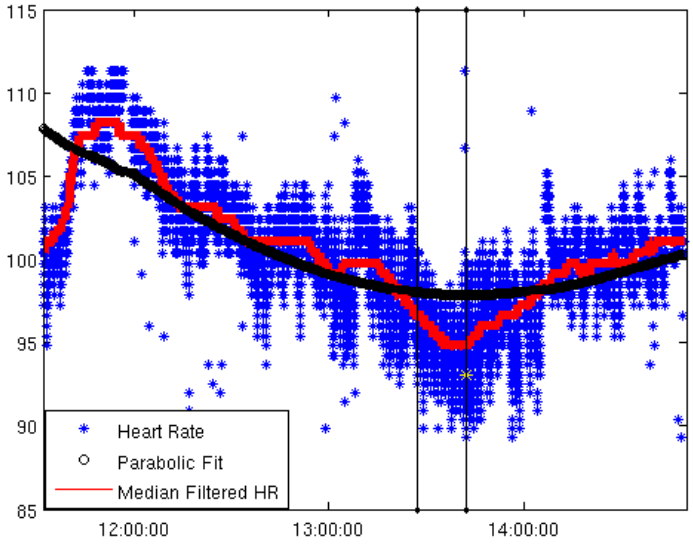
To determine the lowest heart rate during a given dialysis session, the RR-intervals are derived from the ECG using a Pan-Tompkins based peak detection algorithm [27], as described in Section 2.1. The scaled reciprocal of the RR inter-

vals gives the heart rate. Heart rate values outside of the artefact rejection limits given in Table 2.4 were discarded. To find the nadir of the heart rate time series, the heart rate is median filtered over 5-minute windows shifted every minute, and a parabola of the form

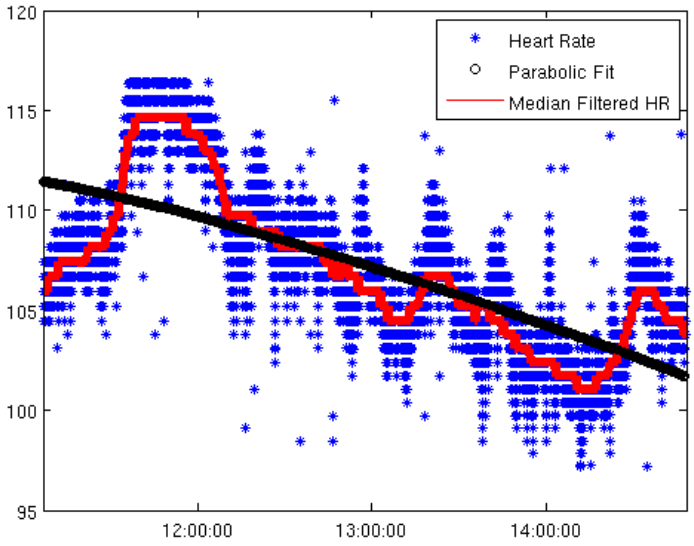
$$p(x) = a_2x^2 + a_1x + a_0 \tag{4.5}$$

is fitted to the median-filtered heart rate data. The minimum of the fitted parabola, given by $x_0 = -\frac{a_1}{2a_2}$, is taken to represent the time of the lowest heart rate. By taking this approach, RSA-associated fluctuations as well as sudden transients and outlier data are discarded when looking for the heart rate nadir. An example in which the heart rate nadir is detected using a parabola fitted to the median-filtered heart rate signal is showed in Fig. 4.3a.

In sessions for which the rate at which the heart rate decreases at the start of the session is very different to the rate at which the heart rate subsequently increases, the parabola is fitted such that its minimum lies beyond the range of the heart rate time series. Alternatively, the heart rate time series can be such that only an inverse parabola can be fitted, as illustrated in Fig. 4.3b. In such cases in which no minimum can be found for the fitted parabola within the range of the heart rate time series, the global minimum of the median-filtered heart rate series is used as the heart rate nadir instead of the minimum of the parabola. This minimum is detected using a trough detector.



(a)



(b)

Figure 4.3 – Two examples of a heart rate time series (in blue), its median-filtered version (in red) and the fitted parabola (in black). In the first example, the vertical black lines indicate the 15-minute segment that is considered to represent normal physiology. In the second example, an inverse parabola is fitted to the median-filtered heart rate time series and the parabola has no minimum. In this case, the global minimum of the median filtered heart rate time series is used to define the end of the window of vital signs representing normal physiology.

4.4.1 Effects of Varying Individual Vital Signs

Once the Neuroscale algorithm is trained with the 450 four-dimensional prototype patterns as explained in Section 4.2.2, the region of normality in two dimensions can be defined by applying all four-dimensional datapoints that constitute the training data at the input of the trained RBF network. A two-dimensional plot showing this region defined by the training data is shown in Figure 4.4. This region forms the reference area for all subsequent vital sign projections in this section. Vital sign projections onto this region can be interpreted as being indicative of normal physiology, in line with the population-based definition of normality. Equally, vital sign projections outside of this region are indicative of abnormal physiology, when compared with the population-based normal data.

Since the Neuroscale algorithm reduces the dimensionality of a given dataset such that the interpoint distances in the higher dimensional space are preserved as much as possible in the lower dimensional space, the axes of the Neuroscale plot do not represent physiological parameters as such, and are dimensionless. In order to establish the physiological significance of the two-dimensional projection space, individual, normalised vital sign parameters are varied individually across a range of -3σ and 3σ from the mean, while the remaining vital signs are kept at their mean value (zero following the zero-mean unit-variance transformation described above). The blood oxygen saturation for the population has a mean value of 94.2% with a standard deviation of 4.72%. However, it cannot exceed 100% and therefore is not varied beyond 1.2σ. The resulting four-dimensional datapoints are projected onto the two-dimensional Neuroscale map trained as described above. The results for each vital sign parameter (heart rate, oxygen satu-

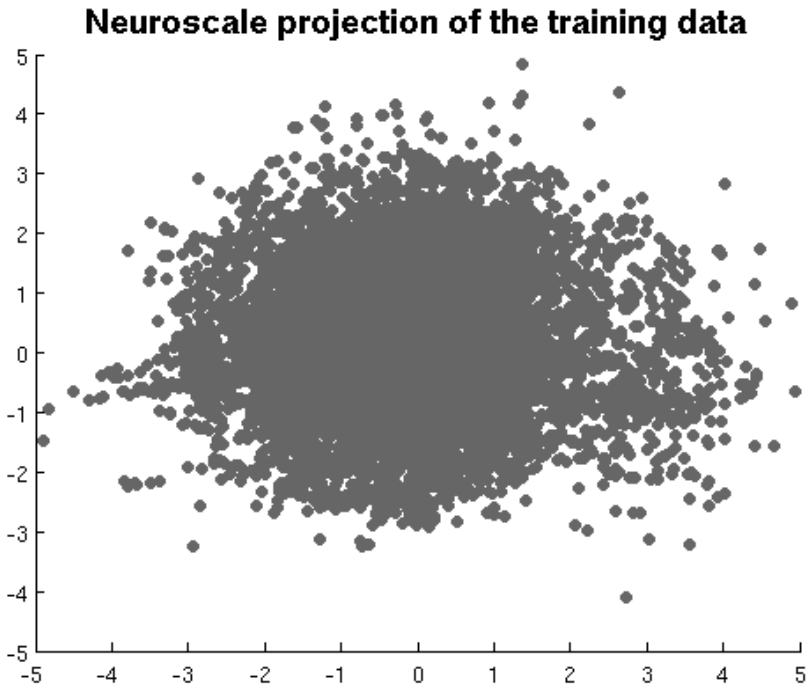


Figure 4.4 – Neuroscale projection of the population-based training data. The training data, originally consisting of 2.5x10^5 datapoints is passed through the RBF network originally trained on 450 representative data clusters. The centre of the map, at $(0\,0)$ represents the most normal set of vital signs, but it should be noted that the axes of the map are otherwise dimensionless.

ration, mean arterial blood pressure and breathing rate) are shown in Figure 4.5, projected onto the region of normality (as previously shown in Fig. 4.4).

Variations in the heart rate and oxygen saturation are projected onto the same axis, with variations below the respective means being projected onto the right of the map and variations above the mean being projected onto the left half of the map for both parameters. The spacing of the datapoints on the map is indicative of the distribution of the vital signs in the training data. For instance, the datapoints on the heart rate trajectory are more widely spaced for values below the

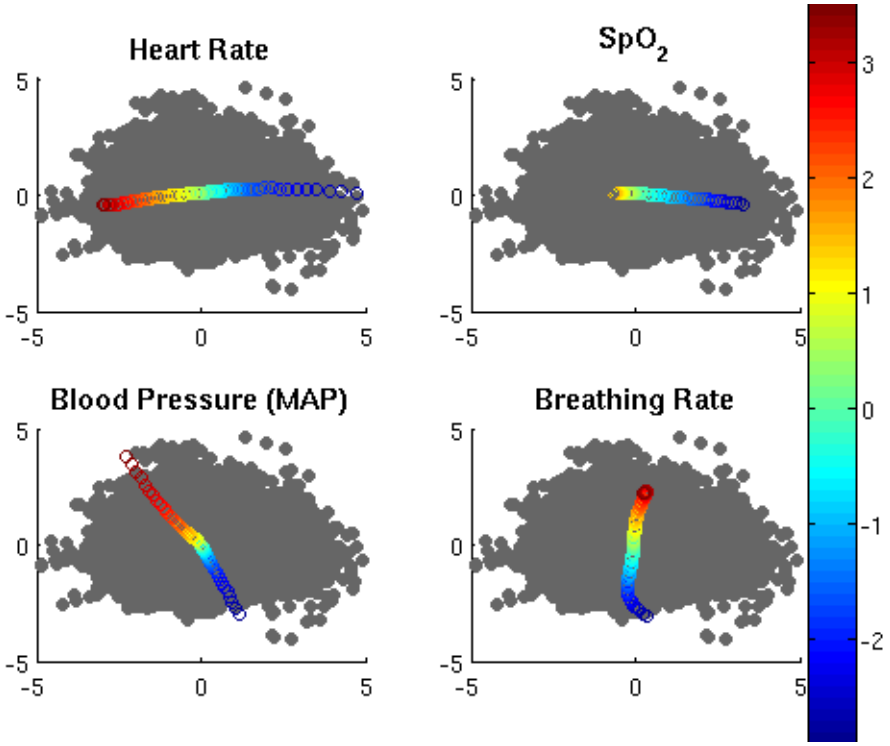


Figure 4.5 – Varying individual parameters across the population-based Neuroscale model. Each vital sign is varied from -3σ (blue) to 3σ (red), while all other parameters are kept at the mean of the training population. As the SpO_2 cannot exceed 100% has a mean value of 94.2 with a standard deviation of 4.72, it is varied between -3σ and 1.2σ. The projections of variations in heart rate, SpO_2, mean arterial blood pressure and breathing rate are shown, with the colourbar representing the variations in terms of σ. While the heart rate and SpO_2 vary along the horizontal axis, the MAP varies diagonally and the breathing rate varies along the vertical axis. The upper left quadrant appears to be representative of heart rate and blood pressure values above the mean, while the lower right quadrant of the map represents heart rate and blood pressure values below the mean. As such, the Neuroscale algorithm appears to capture dependencies between the parameters.

mean (shown in blue), than for values above the mean (shown in red). This corresponds to the distribution of heart rates shown in Fig. 4.2, in which there are significantly fewer heart rate values more than two standard deviations below the mean than there are values that are more than two standard deviations above the mean. The breathing rate varies vertically on the Neuroscale map, with values below the mean breathing rate being projected onto the lower half of the map and values above the mean being projected onto the upper half of the map. The mean arterial blood pressure varies diagonally across the map, with values above the mean being projected onto the top left quadrant and values below the mean being projected onto the lower right quadrant. A shift to the left of the centre of the map represents a simultaneous rise in heart rate and blood pressure, whereas a shift to the right represents a simultaneous fall in heart rate and blood pressure.

In principle, as demonstrated in Section 3.6, the heart rate rises in response to a decrease in blood pressure during fluid removal, and the projection of the normal vital sign parameters appears to contradict this fundamental relationship. However, as observed by Javed et al [152] and described in Section 4.3, dialysis patients commonly experience an initial decrease in heart rate in parallel with a decrease in blood pressure as excess fluid is removed from the body at a rate not exceeding the body's rate of refill. Once the body has recovered from the fluid-overloaded state, a rise in heart rate is commonly observed in dialysis patients as a compensatory response to further fluid removal. As explained in Section 4.4, we define normal physiology using the 15 minutes of vital signs prior to this change in physiology from a parallel decrease in heart rate and blood pressure to a compensatory increase in heart rate in response to a further decrease in blood pressure. Therefore, the vital sign data used to train the Neuroscale algorithm consists ex-

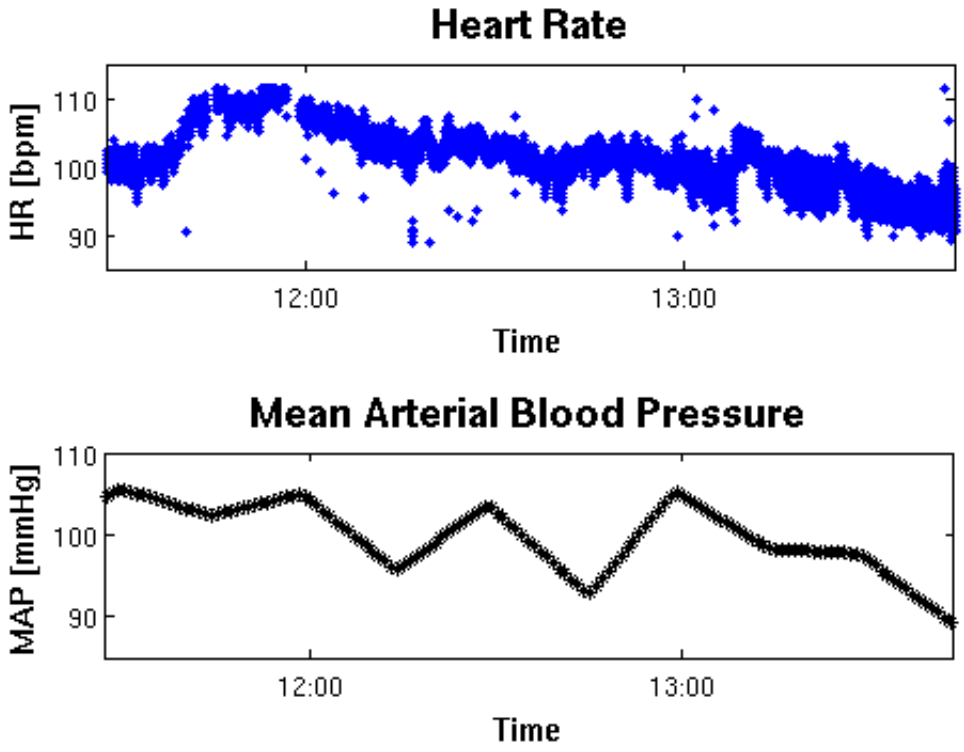


Figure 4.6 – Parallel decrease in HR and MAP from the start of the session to the heart rate nadir. As only data from the 15 minutes preceding the heart rate nadir is used for training the Neuroscale algorithm, the training data is biased towards proportional changes in HR and MAP.

clusively of data in which the blood pressure and heart rate change in parallel, and does not extend to the compensatory phase in which an inverse relationship is observed. This is illustrated in Figure 4.6, which shows the heart rate previously shown in Figure 4.3a from the start of the session until the estimated nadir, alongside the MAP during this time. Although there are fluctuations in the measured MAP, there is an overall decrease over time in both parameters.

4.4.2 Normal and Abnormal Vital Sign Data Projections

In the following section, the population-based Neuroscale algorithm, trained as described above, is applied to data from both IDH and non-IDH sessions which constitute the test dataset. For this purpose, 12 sessions were chosen at random from the 300 non-IDH sessions of adequate quality. Further, all 162 IDH sessions were included in the test dataset, giving 174 test sessions in total. An example of Neuroscale visualisation of two sessions from patient OB 18 is shown in Figures 4.7 and 4.8. The patient is a 60-year old male who had been on dialysis for two years at the time of the study. The vital signs in the first session (Fig. 4.7a) are stable throughout the session, with the mean arterial blood pressure (shown in red) decreasing gradually from its initial value of 117 mmHg to 104 mmHg at the end of the session. (It should be noted at this point that while the systolic and diastolic blood pressures are shown on all vital sign plots (in grey), they are not included in any of the models in this thesis.) The heart rate (shown in green) initially starts at 85 bpm and varies between 80 and 90 bpm throughout the session until it decreases at the end of the session to 75 bpm. The oxygen saturation (shown in blue) and breathing rate (shown in black) remain stable at 90-98% and 10-24 breaths/minute, respectively. The Neuroscale algorithm projects these vital signs across the centre of the region defined by the population-based training data, as shown in Figure 4.7b. The change in colour of the projected vital signs is representative of their progression in time. The start of the session is shown in dark blue, with subsequent hours following the colourscheme shown on the right in Fig. 4.7b.

In the following session (Fig. 4.8a), while the patient's mean arterial blood

pressure is higher at the onset of the session at 128 mmHg, it decreases sharply within the first 45 minutes of the session to 68 mmHg. The heart rate decreases during this time as well, from 86 to 70 bpm, and remains at around 70 bpm for the remainder of the session. The oxygen saturation is initially between 90 and 98 %, but decreases during the session to very low values.

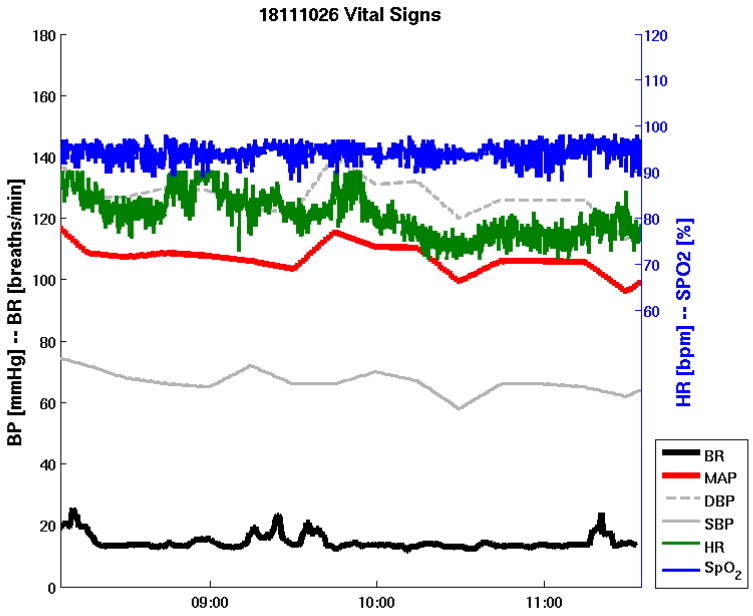
However, neither the IDH episode at the start of the session, nor the decrease in blood oxygen saturation are reflected as abnormal in the Neuroscale projection. While some data is mapped close to the edges of the region of normality, no data is projected outside this region, thus indicating that the data is similar to that of the training set. Otherwise the Neuroscale algorithm would not be able to interpolate the two-dimensional projection using the radial basis function network trained using normal data. Comparing the location of the projections on the map with the analysis in Figure 4.5, the decrease in blood pressure (shown in a medium shade of blue) is projected onto the lower right part of the map, which corresponds with the region in Figure 4.5 onto which low blood pressure values were projected when the remaining vital signs were kept at their repective means. Equally, the vital signs in the final hour (shown in red) in which the patient desaturates are projected onto the right edge of the map, which again corresponds with the region onto which low oxygen saturation values are projected as shown in Figure 4.5. Finally, while the vital sign data at the start of this session (shown in dark blue) is projected onto almost exactly the same region of the population-trained map as the previous session (see Fig. 4.7b), the projection of the vital signs from the abnormal session are spread over the entire right half of the map, indicating a larger degree of variation in the vital signs than in the normal session. More examples of Neuroscale projections of both IDH and non-IDH sessions are shown

in Section B.1 of Appendix B.

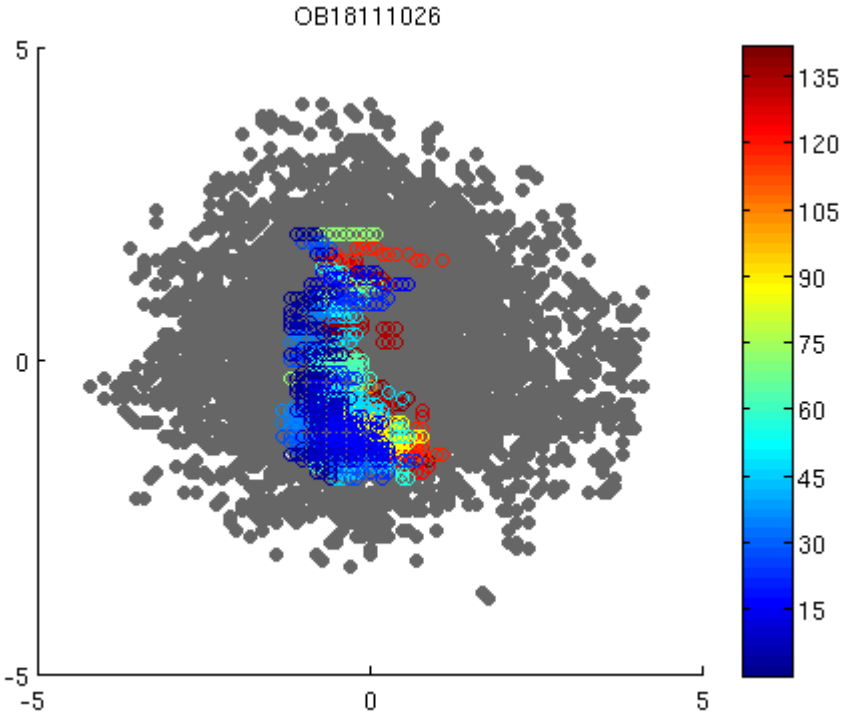
Therefore, although this patient clearly experiences an episode of IDH as defined in Section 1.4 at the beginning of the second session and has abnormally low oxygen saturation values towards the end of this session, this is not adequately reflected in the Neuroscale projections when using a population-based training dataset that encompasses vital signs from 288 different sessions.

The histograms of the normal and IDH session for OB18 are shown in Fig. 4.9 in comparison with the vital sign distributions for the population-based training data. To enable comparison between the distributions, the histograms are normalised such that for each vital sign the frequency of the value with the highest frequency is 1. For the heart rate, the histogram of the training population ranges from 45 to 120 bpm, while those of the two sessions for OB18 cover 70 to 90 bpm and 70 to 86 bpm, respectively. Similarly, the distribution of blood pressure values for the training population ranges from 60 to 170 mmHg, while those of the two sessions for OB18 range from 95 to 118 mmHg and 70 to 128 mmHg, respectively. The differences in the distributions for oxygen saturation and breathing rate are less pronounced when comparing the IDH session with the training data. However, both the oxygen saturation and breathing rate distributions for the normal session are more concentrated around a small range of values than is the case in the IDH session. Overall the histograms demonstrate that the training data covers a much wider range of vital sign data than data from an individual patient, including data from an IDH session.

In conclusion, a population-based model of normality is unsuitable for distin-

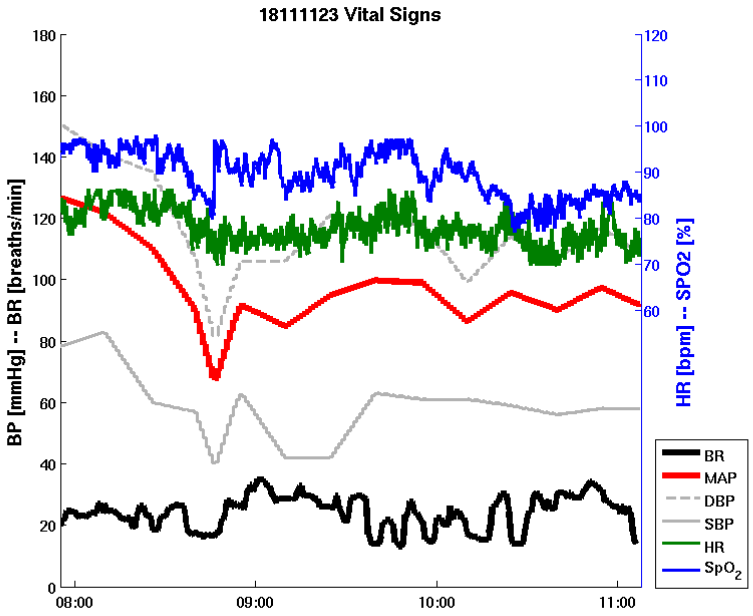


(a)

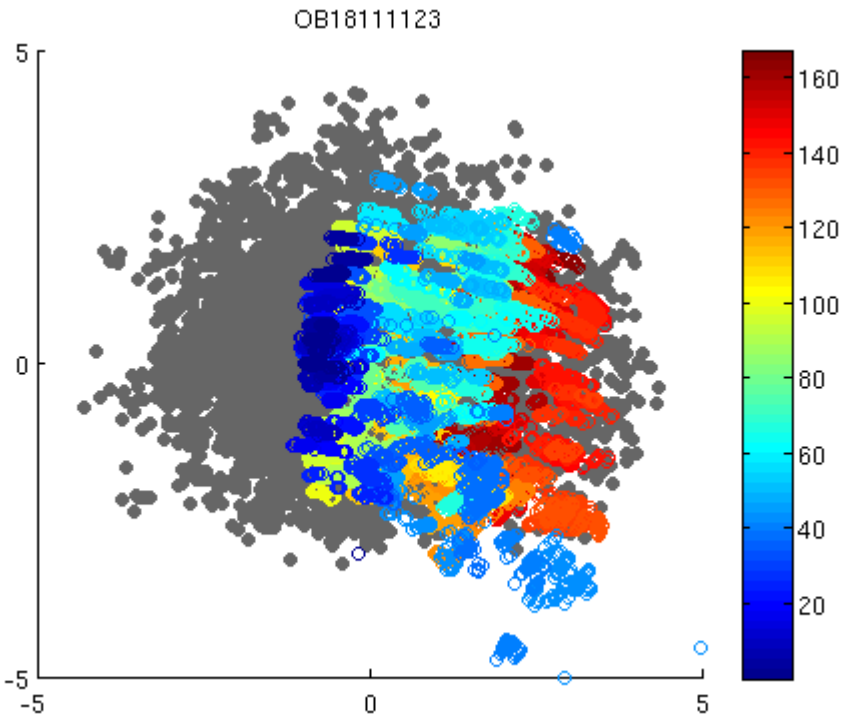


(b)

Figure 4.7 – Vital signs from a non-IDH session for patient OB 18 projected into two dimensions using the RBF network trained with a population-based definition of normality. The region defined by the training data is shown in grey, while the vital sign data is projected in colour, reflecting the progression of the session in time from start (blue) to end (red). The colourbar is displayed in minutes from the start of the session. The vital signs are plotted onto the centre of the Neuroscale map.



(a)



(b)

Figure 4.8 – Vital signs from an IDH session projected into two dimensions. Despite the abnormalities in both the mean arterial blood pressure and the oxygen saturation in this session, the vital signs are projected onto the region of the Neuroscale map defined by the normal population data. The vital signs are projected across the right half of the map.

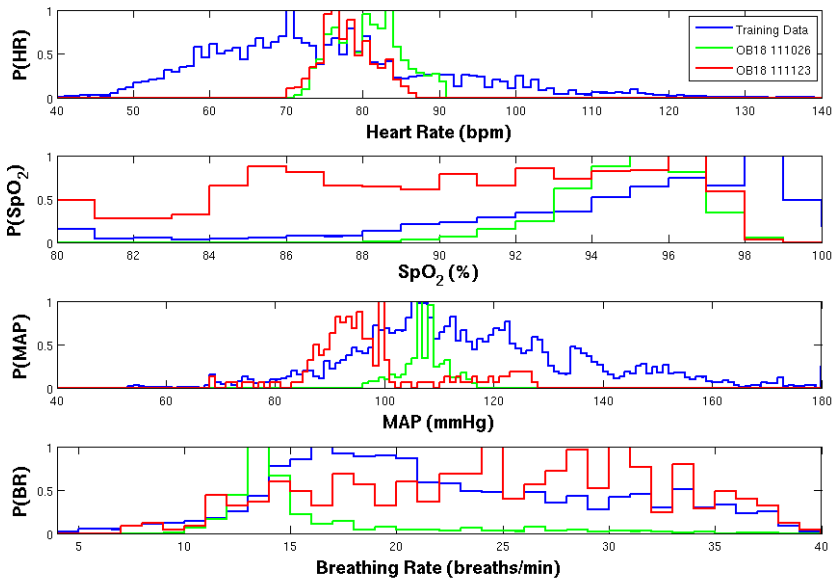


Figure 4.9 – OB18: Vital sign histograms of two sessions in comparison with histograms of the training population. The histograms of the training population (shown in blue) cover a much wider range of values across all vital signs when compared with the histograms of the two sessions for OB18 (shown in green and red). This is particularly evident in the heart rate and mean arterial blood pressure distributions. With regard to the breathing rate and oxygen saturation distributions, no significant difference can be observed between the histograms of the IDH session (in red) and the training data (in blue). However, the distribution for the normal session (in green) covers a much smaller range of values for these two parameters. It should be noted that the training dataset consists only of vital sign data from the 15 minutes preceding the heart rate nadir, while the histograms of the two sessions for OB18 are calculated over the entire session. The distributions are normalised such that for each vital sign the frequency of the value with the highest frequency is set to 1.

guishing between normal and abnormal patient vital signs. It should also be noted that regardless of the Neuroscale projections, vital sign combinations that reflect normal physiology in one given patient may constitute abnormal physiology in another patient.

4.5 Alternative Population-Based Models

As demonstrated in the previous section, a population-based Neuroscale model of normality is unable to distinguish between vital signs representing normal physiology and those representing an episode of IDH for an individual patient, based on the wide range of vital sign data that is included in such a model. In the following section, alternative population-based models are suggested instead. In the first approach, relative changes in vital signs with respect to the start of the session are used to train the RBF network with the Neuroscale algorithm, rather than the absolute values. In a second model, the population is further divided based on physiological characteristics and different maps are trained for the two subsets.

4.5.1 Relative Changes in Vital Signs

As can be seen from the histograms of individual sessions shown in the previous section, it appears that normal and IDH sessions can be distinguished in part on the basis of the amount of variation in the vital signs. By calculating the difference in the vital sign values with respect to those at the start of the session, variations in vital signs could be quantified regardless of the absolute values. By considering relative changes, the effect of associating normality with vital sign values which, may represent normal physiology in one patient, but abnormal physiology

in another is removed.

Finally, given that the definition of IDH used in this thesis is based on a relative change in blood pressure (albeit with respect to the previous reading and not the start of the session), it can be argued that the model of normality should be based on relative changes in the vital signs.

As described in Section 4.4.2, the training dataset in this model is composed of the vital signs from the same 288 non-IDH sessions used previously. For these normal sessions, a reference value is first calculated for each vital sign by taking its 5-minute median from the start of each session. The relative change in each vital sign in time with respect to this value constitutes the training data. Given that the vital signs were recorded at different sampling frequencies, the four vectors representing the relative changes in each of the vital signs are aligned such as to obtain equally sampled vectors by applying a zero-hold on the less frequently sampled data where necessary (as explained in Section 4.2.3).

Histograms of the reference values for each parameter are shown in Fig. 4.10. While the reference values of the non-IDH sessions are evenly distributed along a range of possible values for each vital sign, those of the IDH sessions are more concentrated at certain values, in particular with respect to the heart rate and blood pressure. This is most likely a statistical manifestation of the unbalanced nature of the two distributions. Firstly, data from 162 IDH sessions is being compared with data from 288 non-IDH sessions. Further, within the two datasets, it should be noted that different patients contributed variable amounts of vital sign data depending on the respective number of IDH and non-IDH sessions recorded. For instance, of the 162 IDH sessions considered, 16 sessions (10%) were recorded from patient OB20. As outlined in more detail in Chapter 5, vital signs recorded

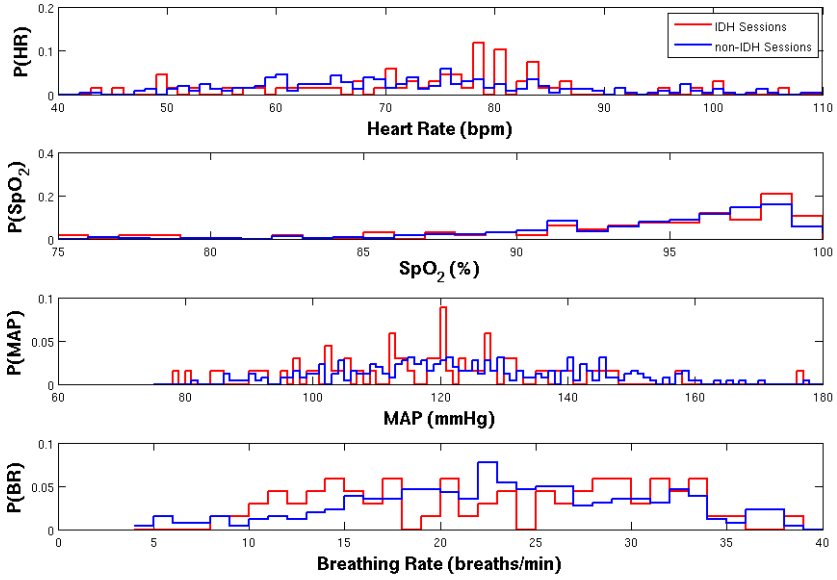


Figure 4.10 – Histograms of reference values at the start of each session. The reference values are calculated as the 5-minute median of each vital sign from the start of each session. The histograms are separated by IDH and non-IDH sessions. While the reference values for the non-IDH sessions are more evenly distributed, those of the IDH sessions appear more concentrated at certain values. This could be a reflection of the unbalanced number of sessions considered, and more particularly of the unbalanced number of sessions recorded from different patients.

across different sessions in the same patient often take on similar values, particularly at the start of dialysis.

There are two characteristic physiological responses to fluid removal during dialysis: a compensatory increase in heart rate in response to the fall in blood pressure, as well as a decrease in heart rate as a recovery from the fluid-overloaded state at the start of dialysis. In order to include the entire range of physiological responses to fluid removal that typically occur during haemodialysis, the relative changes throughout entire non-IDH sessions are included in the training data. The

normalisation parameters consist of the means and standard deviations of the relative changes in each parameter. These are shown in Table 4.2. In comparing the means of the relative changes in each vital sign parameter over all IDH and non-IDH sessions, the only significant difference is observed in the mean arterial blood pressure. The average change in MAP over all non-IDH sessions is a 8.91 mmHg decrease, compared to a 17.63 mmHg decrease in the IDH sessions. The average initial MAP in the IDH sessions is significantly lower than for the non-IDH sessions, at 117.84 mmHg compared with 123.46 mmHg. This is expected, given that by definition an episode of IDH involves low blood pressures.

It should be noted that the shortcoming of such univariate analysis in this context is that no conclusions can be drawn as to the concurrent changes in vital signs during dialysis, most notably of heart rate in conjunction with blood pressure, that are part of maintaining normal physiology. By training an RBF network with the Neuroscale algorithm using data from non-IDH sessions in which compensatory changes in physiology occur in response to a decrease in blood pressure during fluid removal however, concurrent four-dimensional changes in vital signs can be projected onto two dimensions so as to identify the relevant compensatory changes as normal physiology. Further, in contrast to the analysis in Section 4.4.2, this is achieved without regard to the absolute values of the vital signs at any given time.

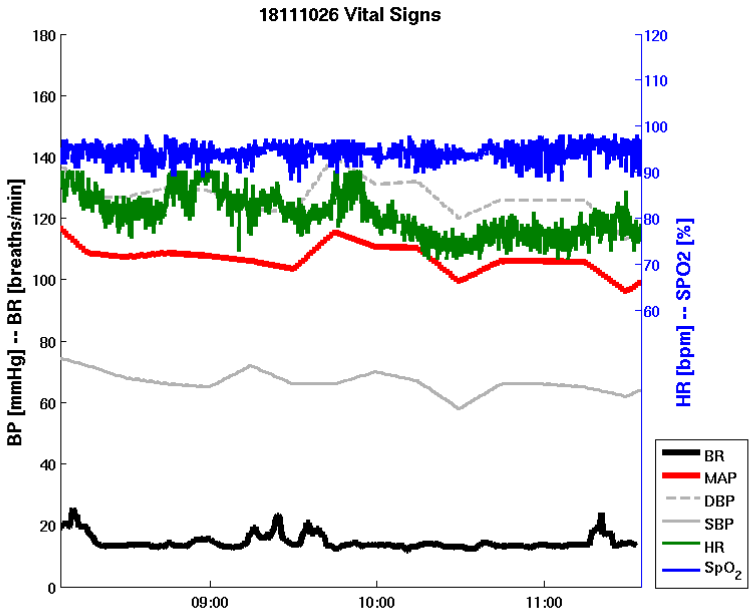
The Neuroscale projections of the relative changes in vital signs for the two sessions considered previously for patient OB18 are shown in Figure 4.11b and 4.12b. As illustrated in Section 4.4.2, the Neuroscale map trained on absolute vital sign values representing normal physiology was unsuitable for distinguishing between normal and abnormal physiology for that patient during haemodialysis.

	non-IDH Sessions			IDH Sessions		
	Ref. Values	μ	σ	Ref. Values	μ	σ
HR (bpm)	72.17	2.20	9.04	73.92	2.44	8.18
SpO_2 (%)	94.23	-0.21	3.41	93.88	-0.28	4.48
MAP (mmHg)	123.46	-8.91	16.24	117.84	-17.63	18.54
BR (breaths/min)	22.76	-0.80	8.74	22.4	-0.58	9.41

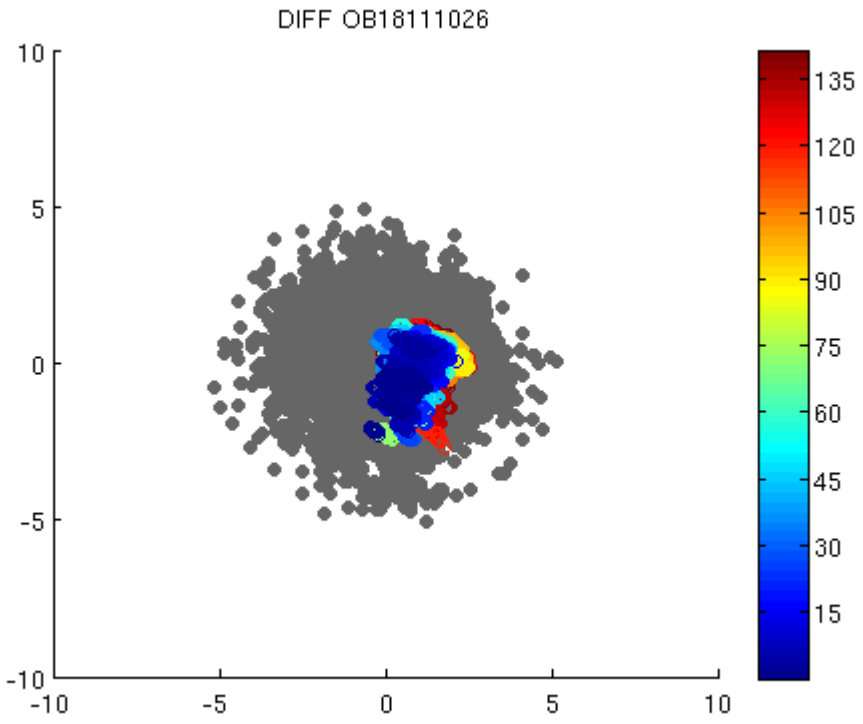
Table 4.2 – Means and standard deviations of the relative changes in each parameter with respect to its 5-minute median value at the start of each respective session. The statistics are separated for IDH and non-IDH sessions. The mean initial value for each vital sign over all sessions is also given for the IDH and non-IDH sessions. Only the mean initial values for the MAP are significantly different between the IDH and non-IDH sessions, at 123.46 and 117.84 mmHg, respectively. Similarly, the average change with respect to the initial value of each session is significantly greater in the IDH sessions than in the non-IDH sessions, at -8.91 and -17.63, respectively.

In the Neuroscale map trained on relative changes in vital signs over the course of entire non-IDH sessions, the vital sign changes from the first session, in which no adverse event occurs, are projected onto the centre of the map. In contrast, the relative changes in vital signs in the second session are projected outside of the region of normality at the specific times during the session in which there is an abnormal decrease in oxygen saturation as well as a sharp fall in blood pressure which is classified as an episode of IDH. The episode of IDH occurs within the first hour of the session, which is represented in light blue in the Neuroscale map, while the abnormally low oxygen saturation levels occur in the last hour of the session, which is projected in red on the Neuroscale map. The IDH episode is projected beyond the upper edge of the Neuroscale map, while the desaturation at the end of the session is projected beyond the bottom left edge of the map.

While in the previous map the region of normality as defined by the training data of the Neuroscale map represents combinations of absolute vital sign values

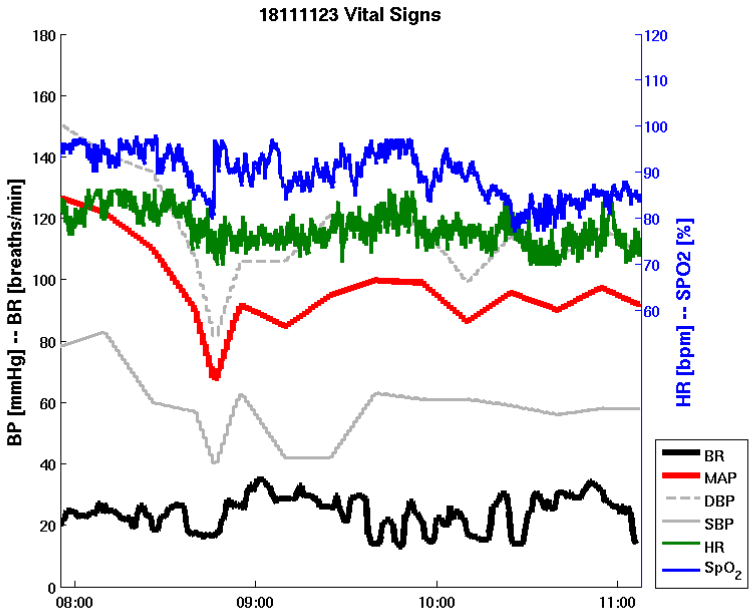


(a)

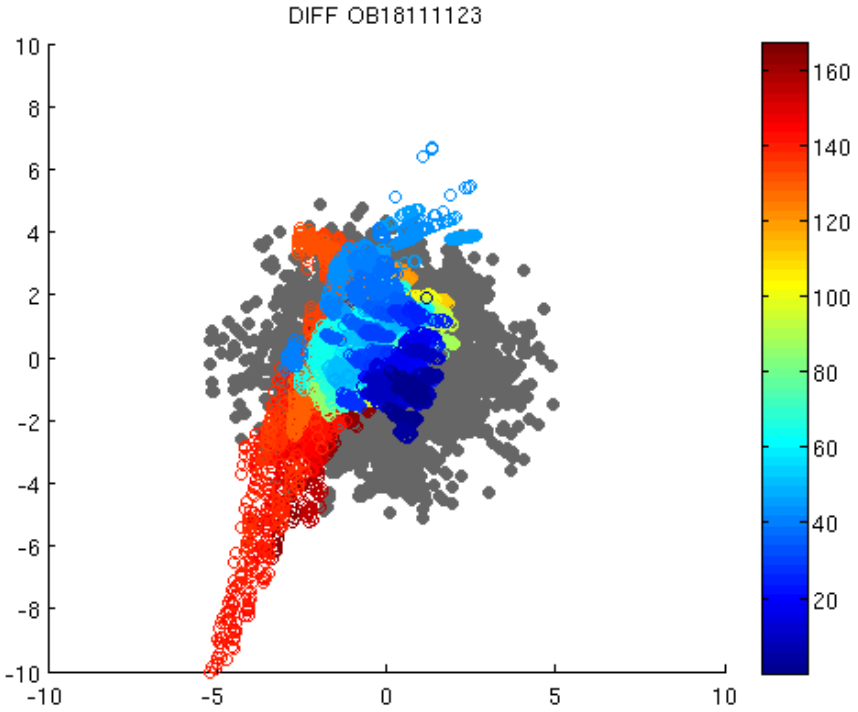


(b)

Figure 4.11 – Neuroscale projections of changes in vital signs in a normal session for OB 18. All vital sign data from the normal session is projected onto the centre of the map.



(a)



(b)

Figure 4.12 – Neuroscale projections of changes in vital signs in an IDH session for OB18. While the data from the normal session is projected onto the centre of the map (see Fig. 4.11b), the abnormal changes in both the oxygen saturation and the blood pressure in this session are projected outside the region of normality. The IDH episode at the start of the session is shown in light blue at the top edge of the map, while the desaturation towards the end of the session is shown in red in the lower left corner of the map.

that are considered to be normal physiology, the map introduced in this section represents relative changes in vital signs that are considered to be normal. As such it appears to be more suitable for distinguishing between adverse events and normal physiology during dialysis. The reasons for this are twofold. Firstly, it overcomes the chief disadvantage of the absolute model, by which despite the fact that the training data is confined to a 15-minute window during each non-IDH session representing normal physiology, the range of vital signs included in the training data is very wide. The model trained on relative changes however, classifies vital signs as abnormal based on their degree of variation over the course of the session from the values at the start of the session. Further, by virtue of including data from the entire session, this map is more likely to identify as normal such changes in the vital signs that represent physiological compensation mechanisms throughout the session.

4.5.2 Stratification based on Autonomic Function

Another alternative to the population-based model introduced in Section 4.4.2 is to partition the population based on conditions known to affect patient physiology during dialysis treatment, and to train separate algorithms for the different cohorts. One such condition that is known to affect the heart rate response to blood volume depletion is autonomic nervous system dysfunction [154]. Specifically, patients with autonomic dysfunction do not show a characteristic rise in heart rate in response to a fall in blood pressure (a more extensive review of autonomic dysfunction in relation to HRV is given in Section 1.5). In the absence of this compensatory mechanism, patients with autonomic dysfunction are not only

more prone to IDH but generally exhibit a different physiological response to fluid removal than normal patients do. As such, these differences justify the separation of the dataset in order to train different Neuroscale maps. By separating the training data according to different physiological responses, it is more likely that the Neuroscale algorithm can distinguish between normal and abnormal physiology, the definition of which varies between the two subgroups.

In order to distinguish patients with autonomic dysfunction from those who do not have this condition, a measure for autonomic dysfunction was devised based on work done by Kaufmann et al [155]. This involved assessing each patient's blood pressure and heart rate upon a change of posture from seated to standing. The normal physiological response of the body upon standing, which results in the sudden accumulation of up to 1 litre of blood in the lower extremities, is an increase in heart rate of 10 to 25 bpm. This autonomic response is mediated by the baroreceptor reflex (described in more detail in Section 1.5) and is activated by a decrease in the cardiac output of the heart and a resulting drop in blood pressure. As such, it limits the drop in systolic blood pressure to 5 to 10 mmHg.

In patients with autonomic dysfunction, no such response is observed. As a result, upon standing, a larger drop in systolic blood pressure is commonly observed in conjunction with either no increase in heart rate, or an increase of less than 10 bpm. Therefore, patients with autonomic dysfunction were defined as those who exhibited a drop in systolic blood pressure of more than 20 mmHg, together with an increase in heart rate of less than 10 bpm. The change in blood pressure and heart rate was measured prior to each session to exclude as much as possible the effects of volume depletion after dialysis. Patients were classified as having autonomic dysfunction if they met the above criteria in at least one session. Using this

	AD		AD, IDH		Normal		Normal, IDH	
	μ	σ	μ	σ	μ	σ	μ	σ
HR (bpm)	65.69	9.52	74.19	8.72	78.32	16.08	70.43	10.88
SpO_2 (%)	93.68	5.15	94.14	5.74	94.71	4.25	92.31	5.97
MAP (mmHg)	118.66	23.27	105.9	18.74	113.68	20.05	107.35	12.7
BR (breaths/min)	22.62	7.66	20.69	7.38	21.25	7.22	20.8	7.74

Table 4.3 – Means and standard deviations of the data subsets, separated by autonomic function. The statistics are calculated separately for IDH and non-IDH sessions from patients with autonomic dysfunction (AD, first and second columns), as well as for IDH and non-IDH sessions from normal (non-AD) patients.

definition, 21 of the total of 77 patients in the study were identified as having autonomic dysfunction. In total, 232 sessions were recorded for these 21 patients, and episodes of IDH occurred in 84 of these sessions (36%). In comparison, the subset of patients without autonomic dysfunction consists of 371 sessions, of which 78 had IDH episodes (21%). This confirms the presumption that patients with autonomic dysfunction are more prone to IDH due to the lack in compensatory heart rate increase in response to a blood pressure decrease.

4.5.3 Subset Statistics

The means and standard deviations of the sessions from patients with autonomic dysfunction (AD), separated for IDH and non-IDH sessions, are shown in Table 4.3, in comparison with the means and standard deviations of both IDH and non-IDH sessions from patients without autonomic dysfunction. The histograms of the underlying data are shown in Fig. 4.13.

The histograms show significant differences in the heart rate distributions when comparing the two patient subsets, while the other three vital signs are similar. In particular, while in patients without autonomic dysfunction the heart rate

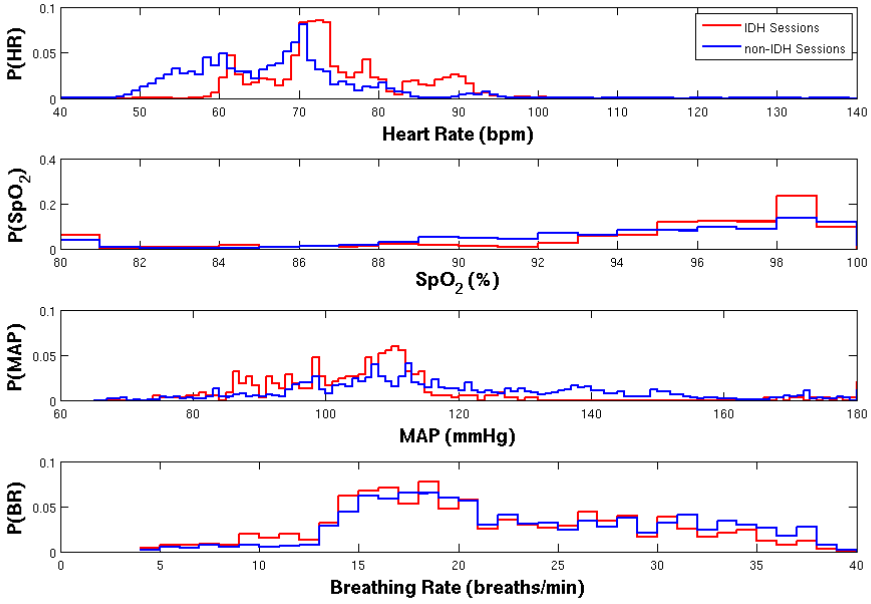
distribution is lower in sessions with IDH than in those without, this relationship is reversed in patients with autonomic dysfunction. In these patients, the heart rate distribution of non-IDH sessions is lower than that of sessions with IDH.

These histograms illustrate two key points. Firstly, the two distributions of non-IDH sessions confirm that patients with autonomic dysfunction do not show the characteristic increase in heart rate during dialysis, since the non-IDH heart rate distribution is lower for the patients with autonomic dysfunction than for the patients without. This is further supported by the mean values for the heart rate given in Table 4.3. Not only are there significant differences in the heart rate means of non-IDH sessions between patients with and without autonomic dysfunction (65.69 vs. 78.32 bpm for patients with and without autonomic dysfunction, respectively).

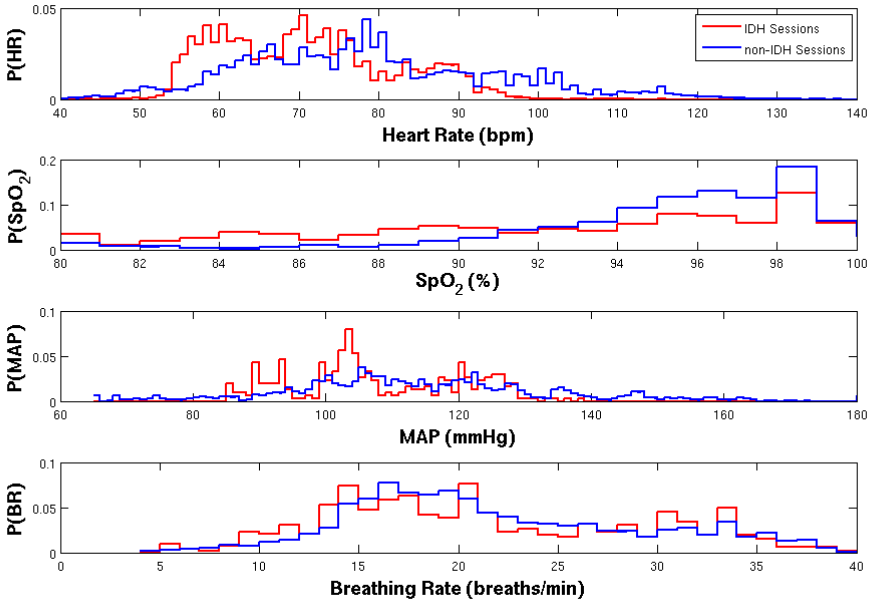
Further, comparing the heart rate distributions of IDH sessions for the two patient subgroups, the IDH distribution is shifted to the left for patients without autonomic dysfunction than for those with autonomic dysfunction. This is indicative of differences between the two patient subgroups in the physiology leading to and during IDH episodes. In this regard it should be noted that while there are significant differences in the heart rate distributions of the two patient subsets, no such differences can be observed in the blood pressure distributions.

4.5.4 Training

Using this data, two separate RBF networks were trained using the Neuroscale algorithm. For the first map, only data from patients who were deemed not to have autonomic dysfunction based on the definition given in Section 4.5.2 was included



(a)



(b)

Figure 4.13 – Vital sign histograms of data subsets, separated by autonomic function. The histograms on the top (a) are for data from patients with autonomic dysfunction, while the histograms on the bottom (b) are for data from patients without autonomic dysfunction. In each subgroup of data, data from IDH and non-IDH sessions is plotted separately (in red and blue). While the histograms for oxygen saturation, mean arterial blood pressure and breathing rate are similar, there are significant differences between the heart rate histograms between patients with and those without autonomic dysfunction.

in the training subset. Of these patients' sessions, only the non-IDH sessions were used. For the second map, the training subset consisted of all non-IDH sessions from patients deemed to have autonomic dysfunction. For both maps, the same procedure for training the Neuroscale algorithm was followed as described in Section 4.3. In short, the heart rate nadir was determined for each session, and the data in the 15 minutes preceding this point was used as the training data. For both maps, the training data was reduced to 450 four-dimensional prototype patterns by using the k-means algorithm to find 500 clusters and removing the 50 (10%) least populated clusters. These prototype patterns were used to train the Neuroscale algorithm, as previously described in Section 4.4.

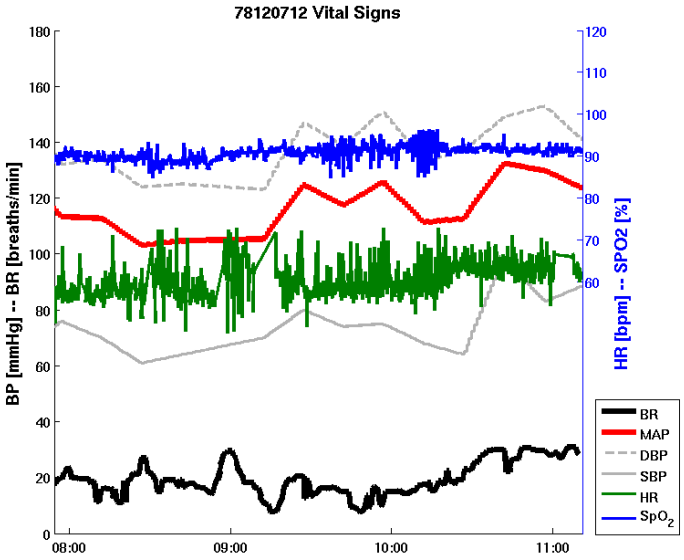
4.5.5 Results

After training the two separate maps, only data from patients with autonomic dysfunction was tested on the map trained using data from this group of patients. Equally, only data from patients deemed not to have autonomic dysfunction was tested on the other map. For the purpose of testing the two maps, 10% of the normal sessions as well as all IDH sessions were retained and not included in the training data. Figures 4.14 and 4.15 show the vital sign time series and Neuroscale projections of two non-IDH sessions from a patient with autonomic dysfunction (OB70, Figs. 4.15a to 4.15c) and a patient without autonomic dysfunction (OB78, Figs. 4.14a to 4.14c). For comparison, the vital sign projections onto both the population-based model introduced in Section 4.4 (in the third column) as well as onto the condition-specific maps (in the second column) are shown. In both sessions, the vital signs are stable and no adverse events occur. Comparing the

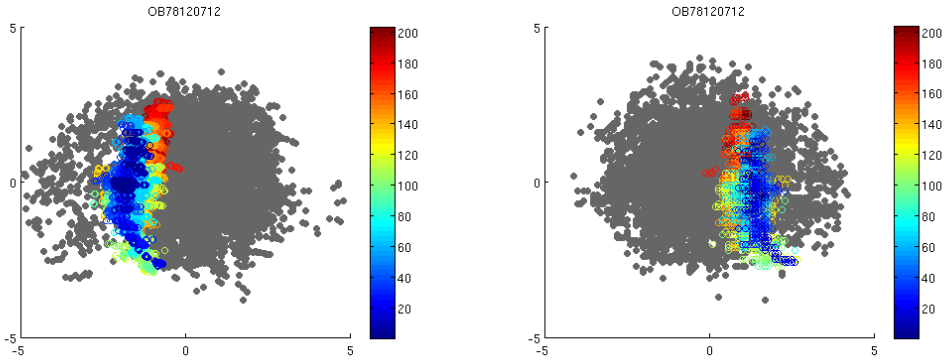
projections onto the two different maps, the data from both normal sessions is projected onto the centre of the respective maps which represent the regions of normality for each patient subgroup. Similarly, the projections of the two sessions onto the population-based Neuroscale map, as shown in Figs. 4.14c and 4.15c, are onto the region of normality as defined by that map.

Figures 4.16 and 4.17 show the Neuroscale projections of two sessions with IDH, one from a patient with autonomic dysfunction (OB66, Figs. 4.17b and 4.17c) and another from the patient without autonomic dysfunction previously studied (OB18, Figs. 4.16b and 4.16c). The time series of vital signs for both sessions are shown in Figs. 4.16a and 4.17a. The Neuroscale maps in the second column are trained on the condition-based subgroups of data, while the map used for the projections in the third column is trained on data from the entire population. As each of the sessions contains an episode of IDH, it is expected that the vital signs at the time of the episode are projected outside of any region defined as normal. As previously established, the population-based Neuroscale map does not display this behaviour, as illustrated in Figs. 4.16c and 4.17c, in none of which the vital signs during the IDH episode are projected outside of the region of normality.

In the session visualised for OB66 (Fig. 4.17a), the IDH episode occurs at around 10:30, two and a half hours into the session, when the mean arterial blood pressure falls abruptly to 60 mmHg. In the Neuroscale plots, this is represented by the orange and red points. In the projection in which the region of normality is defined by the subset of autonomic dysfunction patients (Fig. 4.17b), the IDH episode is clearly projected outside the region of normality. In contrast, the IDH

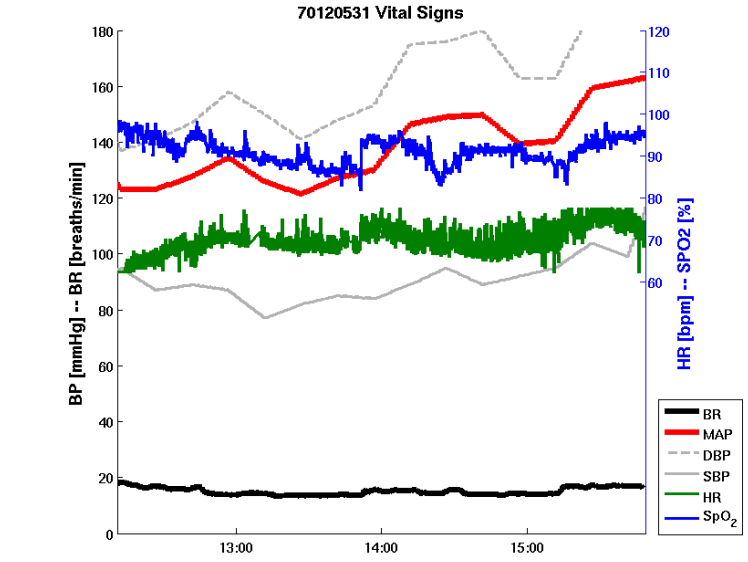


(a)

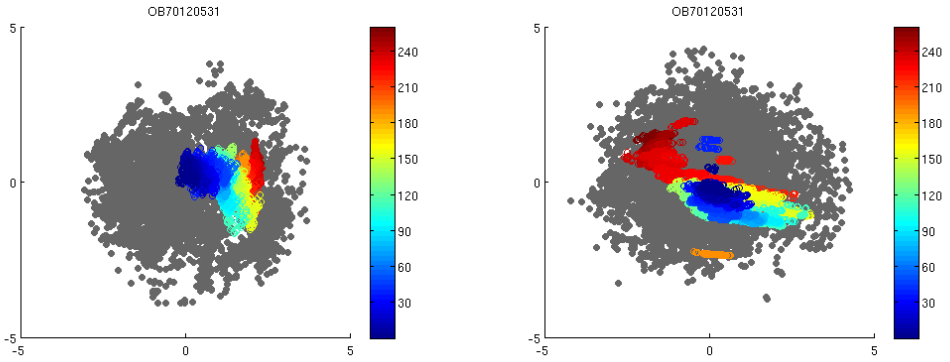


(b) (c)

Figure 4.14 – Neuroscale projections of a non-IDH session from patient OB 78 by autonomic function. Patient OB 78 does not have autonomic dysfunction as defined in Section 4.5.2. The two Neuroscale plots shown are trained using the subgroup of data from patients without autonomic dysfunction (left) and using data from the entire population (right), as previously described in Section 4.4.2. In both cases, the vital sign data is projected onto the region of normality.



(a)



(b) (c)

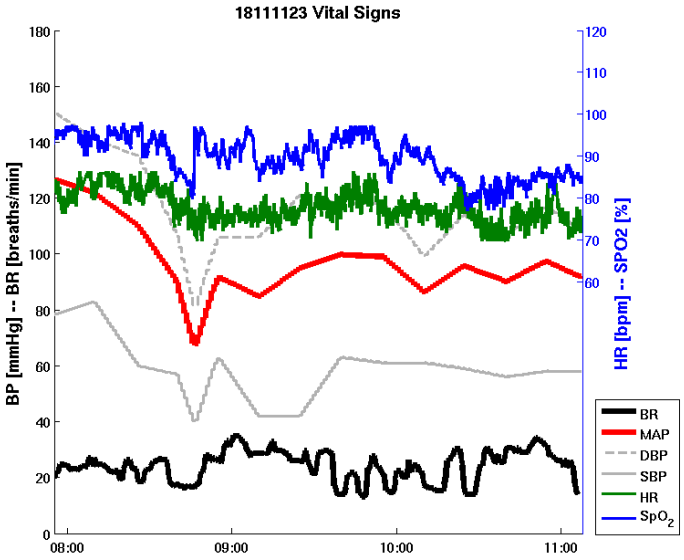
Figure 4.15 – Neuroscale projections of a non-IDH session from patient OB 70 separated by autonomic function. Patient OB 70 satisfies the definition of a patient with autonomic dysfunction. The Neuroscale map on the left is trained using the condition-based subgroup of data from patients with autonomic dysfunction, while the map on the right is trained using data from the entire population. Both the condition-based and population-based models project the non-IDH vital sign data onto the region of normality.

episode in patient OB18, which occurs within the first hour of the session, is not projected outside the region of normality when using the map defined by the vital sign data from patients without autonomic dysfunction, as shown in Fig. 4.16b.

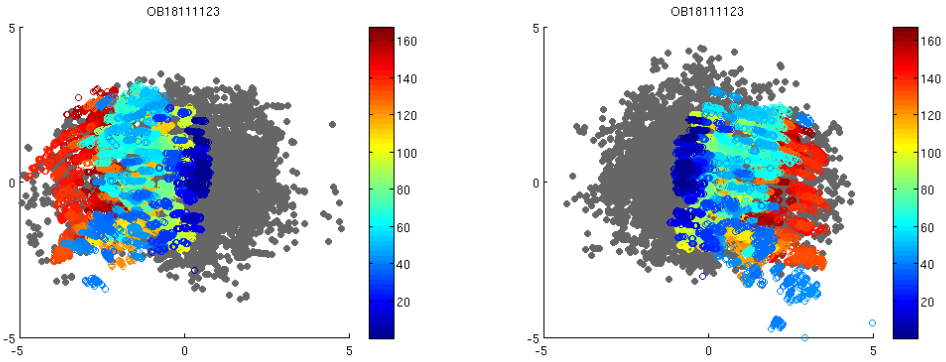
4.5.6 Conclusions

In conclusion, these results demonstrate that a condition-specific model of normality based on autonomic dysfunction can perform better at identifying abnormalities in the vital signs than a population-based model does. It is known that autonomic dysfunction affects the body's physiological response to fluid removal during dialysis. In particular, patients with autonomic dysfunction do not exhibit a characteristic increase in heart rate in response to a fall in blood pressure. Given this inherent difference, a model of normality trained exclusively on normal data from this group of patients performs better in identifying abnormal physiology because it is more specific to the particular characteristics of these patients' physiology.

However, despite improvements in identifying abnormalities in patients with autonomic dysfunction, a model trained only on patients without autonomic dysfunction shows no improvement in comparison with a model based on data from the entire population. When comparing the means and standard deviations of the training data for the cohort of patients without autonomic dysfunction with those of the training data for the population-based model given in Table 4.1, no significant differences can be observed as a result of the removal of the subset of patients with autonomic dysfunction. It should be noted that the subset of patients with au-

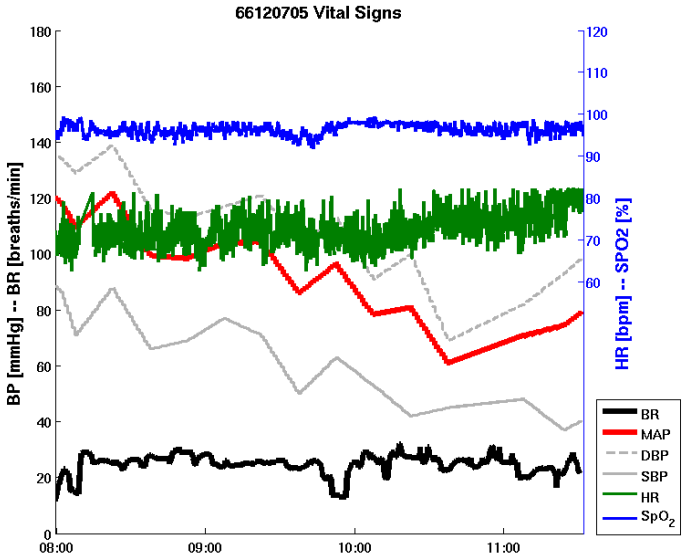


(a)

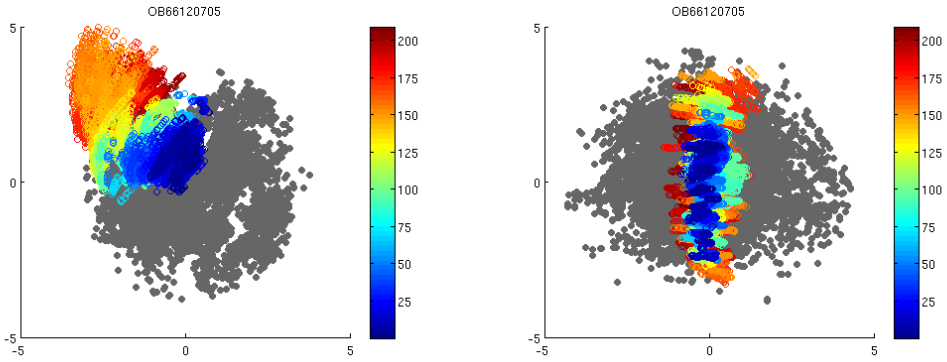


(b) (c)

Figure 4.16 – Neuroscale projections of an IDH session for patient OB18, separated by autonomic function. Patient OB18 does not satisfy the definition for autonomic dysfunction. The vital signs are projected onto the region of normality for both the condition-based and the population-based Neuroscale maps.



(a)



(b) **(c)**

Figure 4.17 – Neuroscale projections of an IDH session for patient OB66, separated by autonomic function. Patient OB66 satisfies the definition of a patient with autonomic dysfunction. The condition-specific map performs better than the population-based map in identifying abnormal physiology in the subgroup of autonomic dysfunction patients.

tonomic dysfunction constitutes less than one quarter of the entire dataset, both in terms of the number of sessions as well as total number of hours.

The main distinguishing factor between patients who suffer from IDH and those who do not lies in the relevant physiological compensatory mechanisms described in Section 1.5. By separating the data based on factors known to affect these mechanisms, the ability of a model of normality to distinguish between normal and abnormal data is improved. To further enhance the specificity of the model, each patient's individual response to fluid removal could be used to train a model of normality that is particular to the physiology of that patient. This approach is further investigated in the next chapter.

Chapter 5

Patient-Specific Models of Normality

Although Haemodialysis is a highly repetitive and controlled form of treatment which lends itself to long-term, patient-specific data collection and analysis, little work has been done in this respect. As can be seen from the results in Section 4.4, a Neuroscale map trained on vital sign data from patients whose physiology is considered to be normal cannot necessarily map the abnormal vital sign data from a particular patient outside the normal vital sign data from the entire population. The results from Section 4.5.2 indicate that the performance of a model of normality for dialysis patients can be improved by making it more specific to the patient's condition, if the latter is known to affect physiology. In this chapter, we consider whether it might be possible to identify changes in a patient's vital signs which are abnormal with respect to that patient's own model of normality constructed from previous vital sign data.

Population-based vital sign models have a number of disadvantages in the context of HD. Firstly, each dialysis session is governed by a number of dialysis-specific parameters such as the UF rate and dialysate sodium concentration, which

can affect the vital signs. In part, a patient's vital signs during dialysis reflect the ability of the cardiovascular system to respond to fluid removal, as demonstrated in more detail in Chapter 3. In this context it is difficult to identify absolute boundaries of normality that are valid for all patients, as is illustrated by the lack of consensus over a global definition of IDH in terms of absolute blood pressure boundaries (discussed in Section 1.4). While each four-dimensional datapoint represents a set of simultaneously recorded vital signs and therefore captures the interrelations between the parameters (such as for instance an elevated heart rate to compensate for a low blood pressure), this is not sufficient to determine absolute limits of normality. Other patient-specific factors, such as the refill capacity, affect a patient's ability to respond to fluid removal.

Long-term medications such as anti-hypertensive drugs introduce a further confounding factor with an effect on vital signs. Finally, as explained in more detail in Section 1.2.1, the length of time that the patient has been receiving HD treatment can also affect vital signs.

5.1 Separation of Normalisation and Training Data

A patient-specific Neuroscale algorithm, which is trained using normal vital sign data for that patient, is considered in this section. As described previously, the 15 minutes of vital signs preceding the point of lowest heart rate in IDH-free sessions are considered to represent the "most normal" physiology. This point is found either by fitting a parabola to the heart rate time series, or if no such parabola can be fitted, by using the global minimum of the time series.

For our investigations in this chapter, we used the subset of 6 patients origi-

	μ_{train}	μ_{all}	σ_{train}	σ_{all}
OB19 HR (bpm)	71.7	74.8	5.9	6.7
OB66 HR (bpm)	68.1	71.6	5.7	10.6

Table 5.1 – Means and standard deviations of the heart rate training data and of the heart rate over entire sessions which include the training data. The values demonstrate the need to expand the training data to include data from the entire session, in order for the Neuroscale algorithm not to project normal data outside of the region of normality, which is defined by the training data.

nally described in Section 2.5 who had more than 12 sessions recorded of which at least three were IDH sessions. To train patient-specific Neuroscale maps, the first 7 non-IDH sessions were used to obtain the training data for each patient, and the remaining sessions which in each case include at least one further normal and three IDH sessions, were used as test data.

Table 5.1 show the means and standard deviations of the heart rates in the 15 minutes preceding the heart rate nadir in the seven sessions that constitute the training data for two patients compared with the means and standard deviations over the entirety of these sessions. In both examples, the mean heart rate over the entire sessions is higher than during the 15 minutes preceding the heart rate nadir. This is expected, given that the 15 minutes preceding the heart rate nadir by definition consist of the lowest heart rate values of the session. Further, the standard deviation over the entire sessions (1680 minutes) is up to twice the standard deviation of the training data (105 minutes). Given that the standard deviation of the heart rates preceding the heart rate nadir is significantly lower than that over the entire sessions, a Neuroscale map trained using the 105 minutes of data only is likely to have the data from the remainder of the non-IDH sessions projected outside the region of normality.

Further, as the dialysis session progresses in time, the vital signs change in response to fluid removal. By adopting a definition of normality that only takes into account a subset of vital sign data in a relatively short time (15 minutes), the remaining vital sign data are deemed to be abnormal, especially as the heart rate nadir by definition represents the physiological transition from a state of recovery from excess fluid accumulation to a state of compensation for further fluid removal. This argument is presented using heart rate data only, but is equally applicable to the other vital signs, particularly the blood pressure which typically decreases monotonically over the course of the dialysis session.

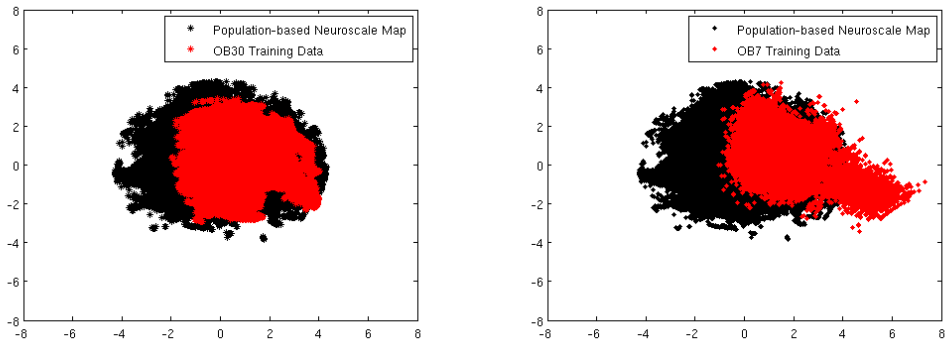
To include both the recovery and the compensatory stages of dialysis in the model of normality, vital sign data from entire non-IDH dialysis sessions are used for the training data. By training the Neuroscale algorithm using vital sign data from the entire session, the mean of the training data no longer reflects the most normal combination of vital signs.

The patient-specific Neuroscale maps in this section are trained using all vital signs from seven non-IDH sessions from each patient. The zero-mean, unit-variance normalisation however, uses the means and standard deviations of the 15 minutes of data preceding the heart rate nadir of each of these sessions. Thus, the training data consists of approximately 28 hours of data, while the normalisation parameters are derived from 105 minutes of data. The (0,0) centre of the Neuroscale map still represents normal physiology, but the region of normality is defined by the training data. Given that no IDH episodes are included in the training data, it is expected that any abnormal data would be projected outside of the region of normality.

5.2 Patient-Specific Training Data

To compare an example of patient-specific training data with the population-based training data from Section 4.4, the training data for theNeuroscale map for patient OB30 is projected onto the population-based Neuroscale map trained in Section 4.4, shown in Figure 5.1a. The patient-specific data (shown in red) overlaps with only a subset of the population-based Neuroscale map (shown in black), with the left edge of the population-based map being excluded. It should be noted however, that the training data for the patient-specific map consists of vital sign data from the entire session, as explained in Section 5.1, whereas the data used to train the population-based map consists only of the 15 minutes preceding the heart rate nadir in the 288 non-IDH sessions recorded. Therefore, in this plot, the vital sign data from seven entire non-IDH sessions of one patient is being compared with the vital sign data from the 15 minutes preceding the heart rate nadir in 288 non-IDH sessions from all 76 patients.

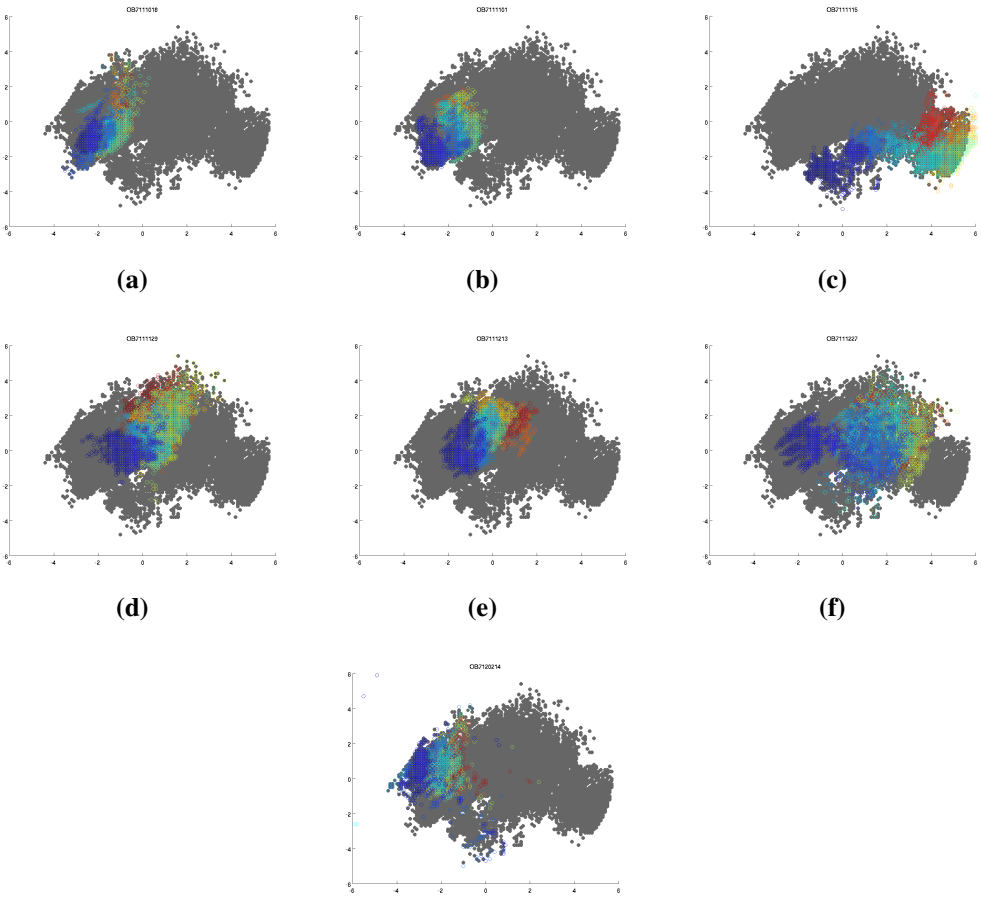
While the training data of patient OB30 is fully contained in the region of normality described by the population-based map, that of patient OB7 shown in Figure 5.1b is in part projected outside the region of normality of the population-based map. This is most likely as a result of the difference in the data included in the training set, as described in the previous paragraph. Interestingly, both sets of training data are projected onto the right half of the population-based map only. Referring back to Figure 4.5, in which the different regions of the population-based Neuroscale map are explored, the left half of the map corresponds to high heart rates. The upper left quadrant is representative of high blood pressures up to three standard deviations away from the mean of the population training data.



(a) **(b)**

Figure 5.1 – Patient-specific training data for two patients projected onto the population-based Neuroscale map. While the majority of patient-specific training data is contained in the region of normality of the population-based map, it is in some parts projected outside the region of normality, as is the case for OB7 (b).

For the purpose of illustration, the data from the seven sessions that constitute the training data for patient OB7 are highlighted individually on the patient-specific Neuroscale map for this patient, as shown in Figure 5.2. The corresponding time series plots are shown in Appendix B. These plots highlight several points. Firstly, although all seven sessions included in the training data are classified as normal non-IDH sessions, the vital sign data from different sessions when projected onto two dimensions with the Neuroscale algorithm do not necessarily overlap. The overall map is representative of the different vital sign combinations that occur in the training data. The projections also illustrate the changes in physiology in each session with time.



(a) **(b)** **(c)**

(d) **(e)** **(f)**

(g)

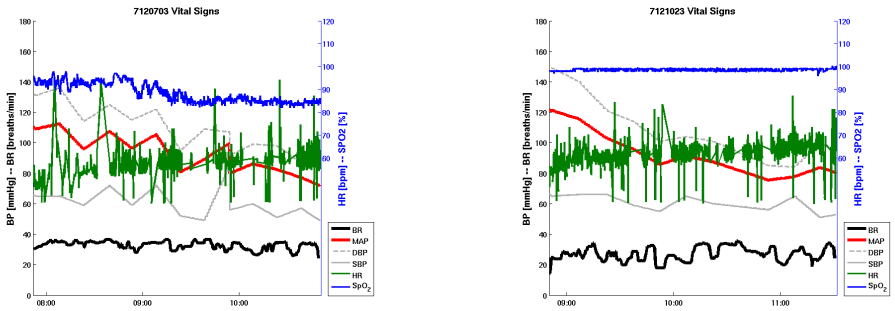
Figure 5.2 – Illustration of the seven patient-specific training data sessions on the patient-specific map trained using the data from these sessions.

5.3 Projections of non-IDH Sessions

In this section, seven other normal sessions are projected onto the patient-specific Neuroscale maps of the two patients considered in the previous section (OB7 and OB30). The results of the projection of normal data are shown in Figures 5.4 and 5.5, while the corresponding time series are shown in Appendix B. Looking at these examples it can be seen that the data from non-IDH sessions is projected

onto the region defined by the training data in the patient-specific Neuroscale maps. However, in some cases such as in Figs. 5.4 (d) and 5.4 (g), outliers do occur. The time series plots of these two sessions are shown in Figure 5.3. It can be seen that the outliers in the Neuroscale projection of the first session arise due to the poor quality of the HR data in the first 30 minutes of the session, in which the signal occasionally drops out or otherwise consists of outlier values. Similarly, in the second session, the HR signal quality decreases towards the end of the first hour of the session, with several gaps in the signal as well as outlier values. While such data of poor quality is projected outside of the region of normality and thus could be mistaken for a deterioration in the patient's physiology, it should be noted that such outliers can be distinguished from abnormalities in the vital signs that arise from changes in physiology in that they are not sustained and deviate from the trajectory of the vital sign projections over time. As can be seen in the following section, abnormal data from IDH sessions is not only projected outside of the region of normality, but the progression of the vital sign projections can be followed across different regions of the map in time. This is discussed in more detail in Section 5.5.

Further examples of non-IDH sessions projected onto patient-specific Neuroscale maps are shown in Appendix B. Based on these results, it can be concluded that patient-specific Neuroscale algorithms can construct a map which is representative of normal physiology.



(a) **(b)**

Figure 5.3 – Time series plots of the vital signs from two sessions giving rise to outliers when projected onto the Neuroscale map. Poor data quality, particularly of the heart rate signal, results in the data being projected outside of the region of normality.

5.4 Projections of IDH Sessions

As described in Section 2.5, the subset of patients chosen for the patient-specific analysis in this thesis consists only of patients who have had more than 12 sessions recorded, of which at least three were IDH sessions. To establish whether a patient-specific Neuroscale map is appropriate for distinguishing between IDH and non-IDH data, a map was trained for each patient in the subset with seven sessions of non-IDH data. Three IDH sessions were then projected onto the patient's map. These sessions are shown in Figures 5.6, 5.8 and 5.10. For reference, vital sign data from one non-IDH session that was not included in the training data is also projected onto the patient-specific Neuroscale map in each case, as shown in Figures 5.7, 5.9 and 5.11. Further examples of IDH sessions from patients who had fewer than 3 IDH sessions recorded yet who had enough normal sessions to train a patient-specific model of normality are shown in Appendix B.

In all three examples shown in this section, the entire vital sign data from the

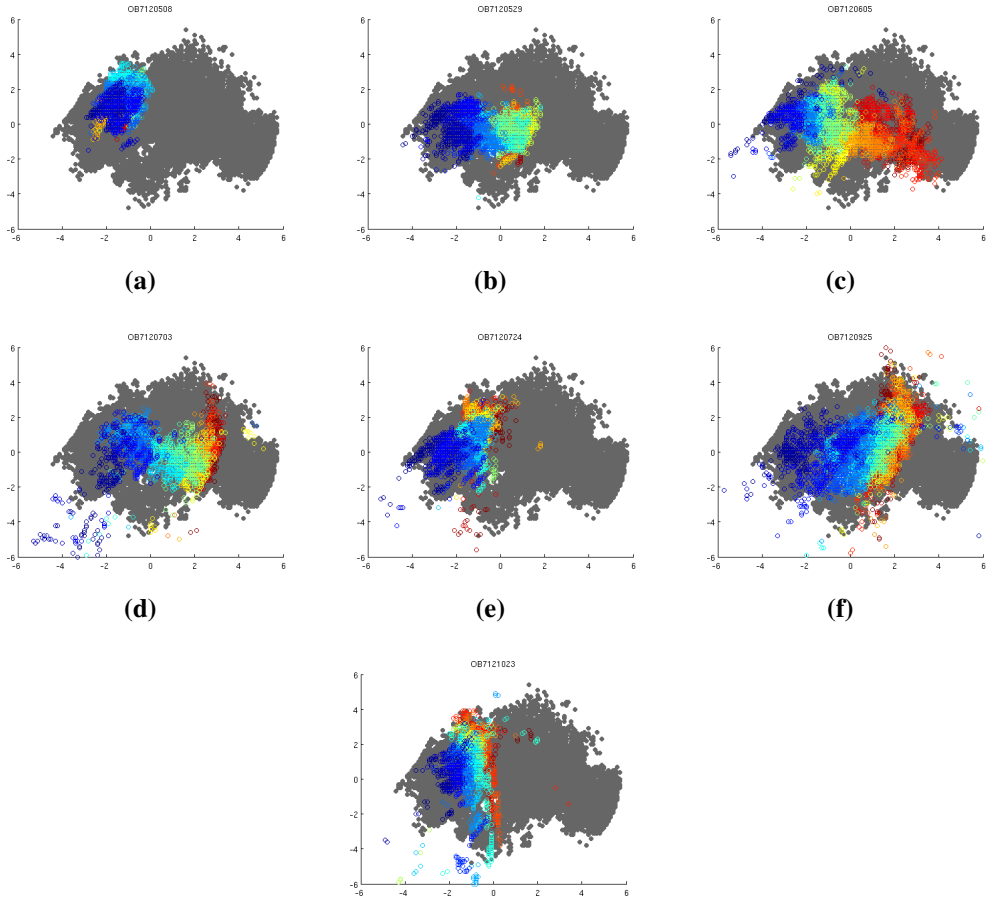


Figure 5.4 – Projection of vital sign data from seven non-IDH sessions onto the patient-specific map for patient OB7. With the exception of outlier data, the vital signs from normal sessions are projected onto the region of normality for the patient.

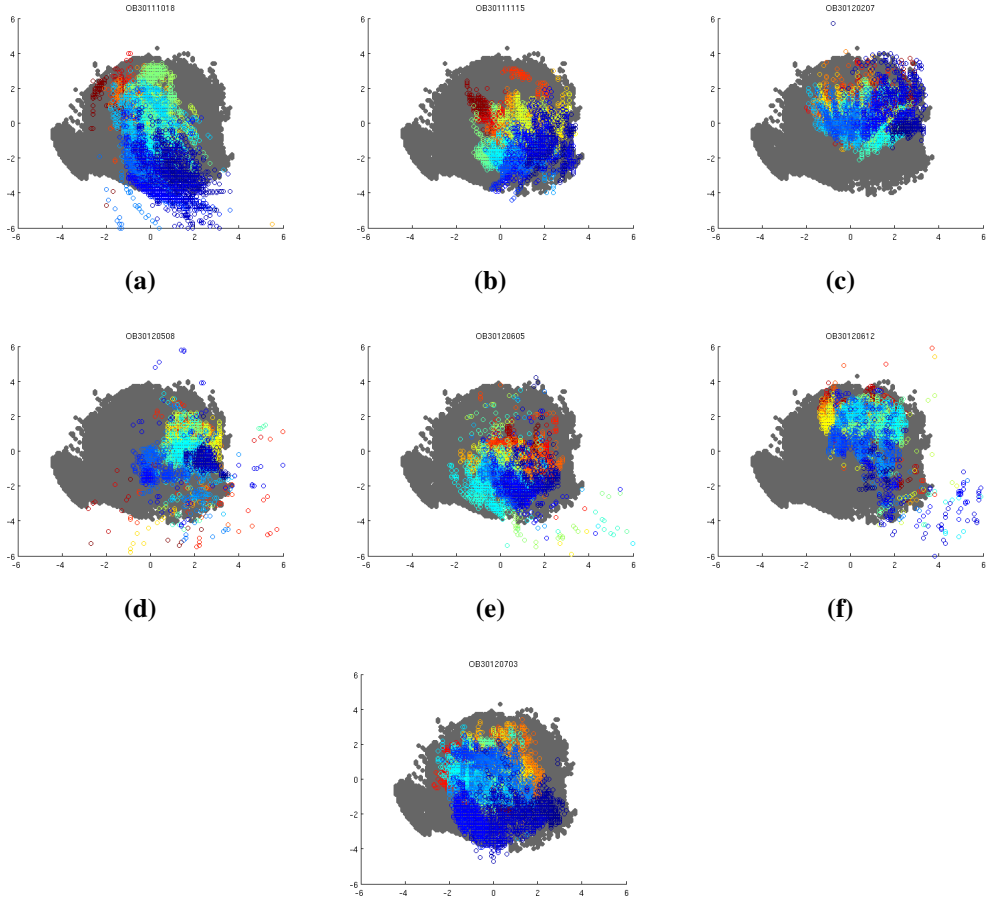


Figure 5.5 – Projection of vital sign data from seven non-IDH sessions onto the patient-specific map for patient OB30.

non-IDH sessions is projected onto the region defined by the training data, as was the case for the data for patient OB7 previously shown in Figure 5.4.

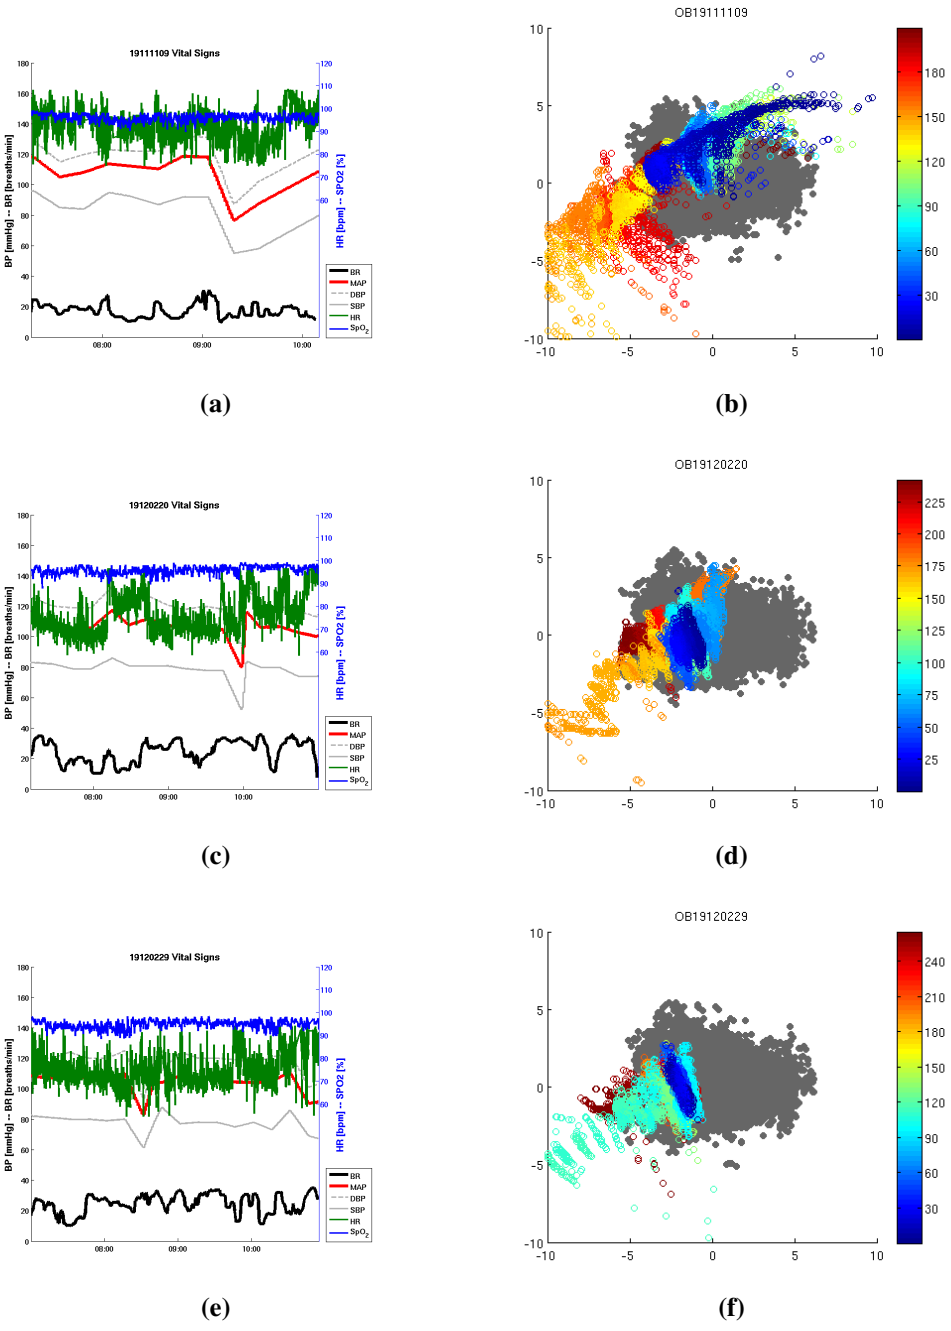
The IDH events are highlighted in each of the time plots and in every case the corresponding points in the Neuroscale projection are clearly outside the region defined by the training data. While vital signs from seven normal sessions were chosen empirically as a suitable amount of data to constitute the patient-specific training data for each patient, the results of this section confirm that this is a suit-

able choice. With a patient-specific Neuroscale map trained using non-IDH data from seven dialysis sessions for that patient, abnormal vital sign data is projected outside the region defined by the training data. Thus normal and abnormal vital sign data can clearly be distinguished in such a map, a significant improvement over the population-based models introduced in Chapter 4.

A further advantage of patient-specific maps is that they allow for comparisons to be drawn between vital sign data from different sessions, particularly their progression in time. This may yield additional insight into the patient's fluid status and general physiological status. These further aspects of the Neuroscale map are discussed in the following sections.

5.5 Vital Sign Trajectories in Time

As previously explained, the continuous removal of fluid during dialysis elicits various physiological responses in the body throughout the session. Further, despite adverse events, dialysis ideally should be seen as a restorative treatment in which the patient's physiology is brought from a state of fluid overload to a state of correct fluid balance. The body's response to this shift is mirrored in the vital signs. While the overall effect on patient physiology is difficult to assess by univariate analysis due to the interrelations between the different vital signs, patient-specific Neuroscale projections allow to compare the changes in vital signs in response to dialysis across multiple sessions.

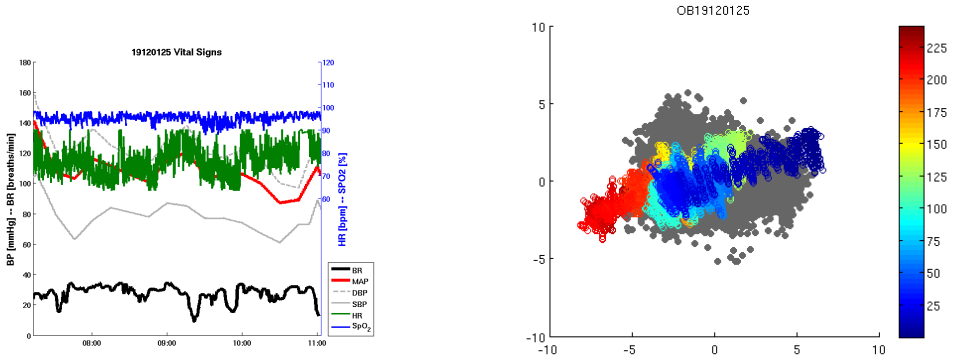


(a) **(b)**

(c) **(d)**

(e) **(f)**

Figure 5.6 – Three IDH sessions for patient OB19 projected onto a patient-specific Neuroscale map.



(a) **(b)**

Figure 5.7 – Neuroscale projection of a non-IDH session for patient OB19.

In the case of a normal session without adverse events, as for instance in the non-IDH session for OB19 shown in Figure 5.7, a shift is observed with time in the projected points across the region of normality, giving rise to a vital sign trajectory across the map. Ideally, each dialysis session would progress along the same trajectory, if the patient presented with the same amount of excess fluid and no other weight gain at the start of each session. However, given that patients' food and fluid intake as well as other factors can affect both the amount of fluid and weight gain in the interim period between dialysis sessions, the patient's fluid status at the start of each dialysis session can vary greatly. Correct assessment of a patient's fluid status at the start of a dialysis session is one of the key clinical considerations that gives rise to the treatment parameters (such as the UF rate) and the target amount of fluid to take off.

5.5.1 Normal Sessions

It is unlikely that a given patient will have the same vital sign trajectory on the Neuroscale map for every normal session, as evidenced by the fact that the vital

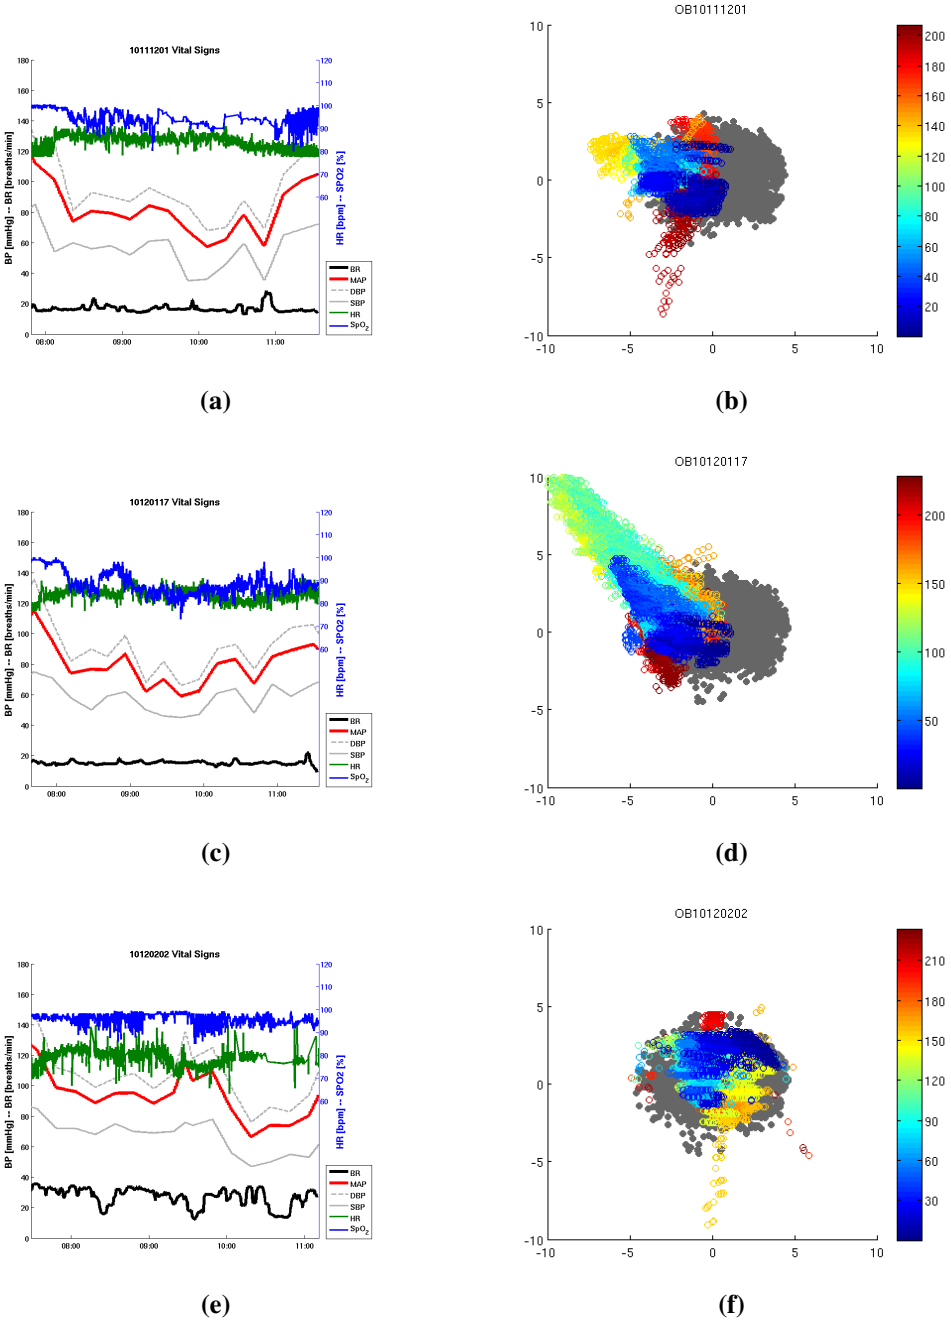


Figure 5.8 – Three IDH sessions for patient OB10 projected onto a patient-specific Neuroscale map.

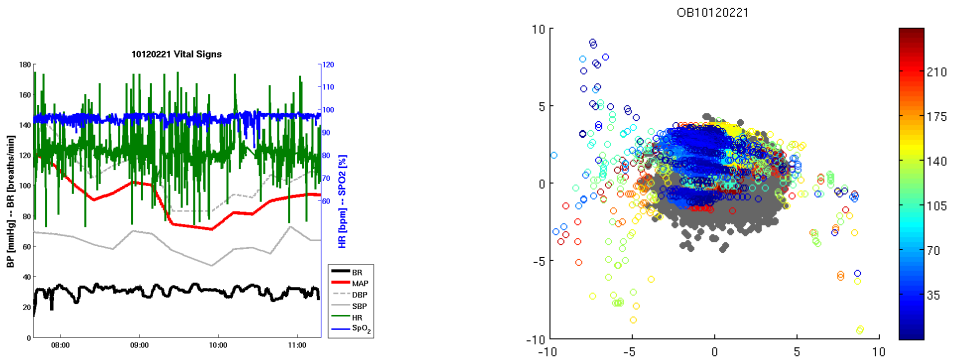
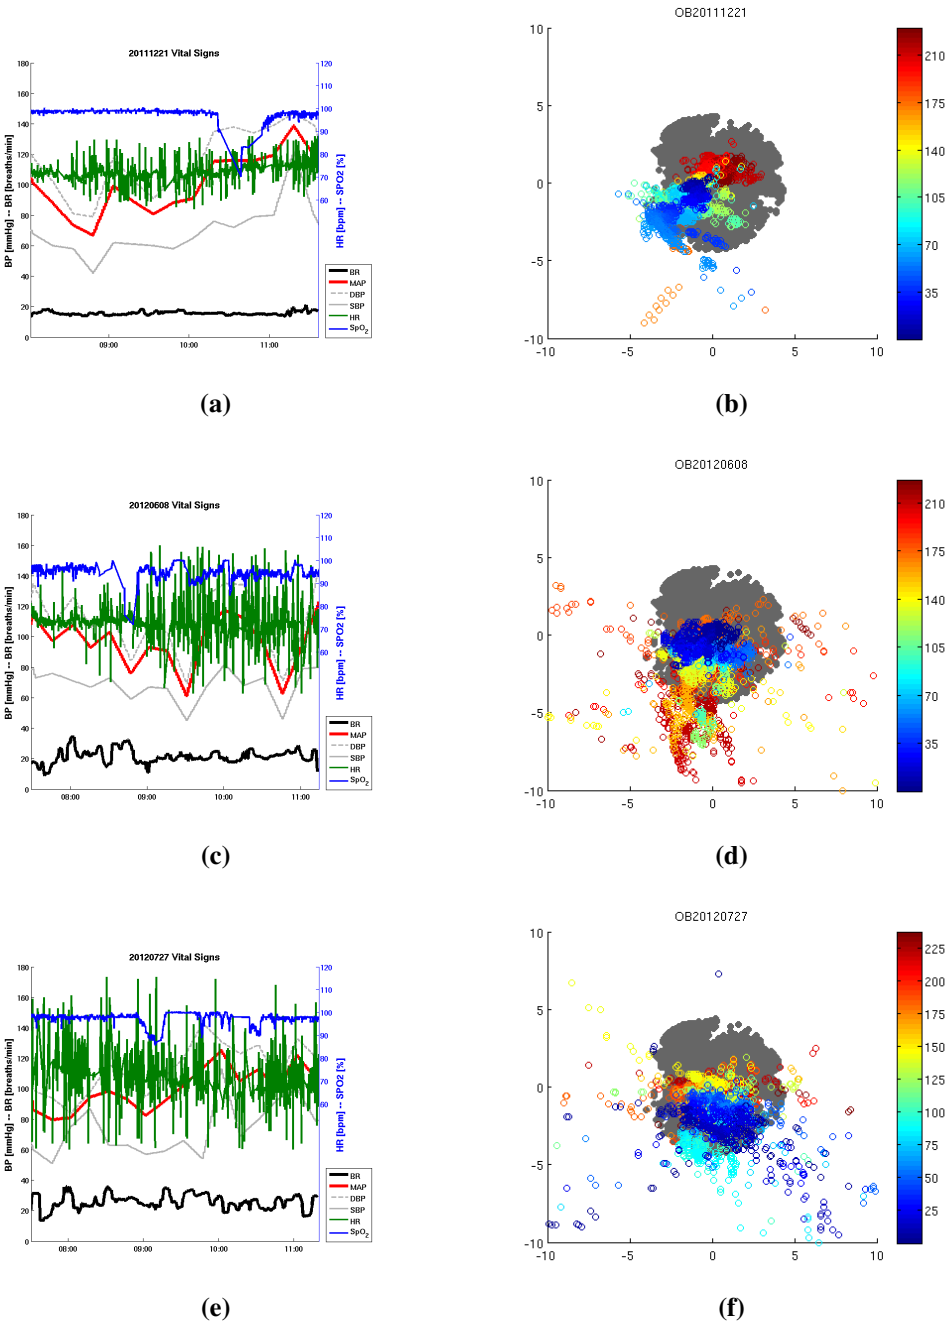


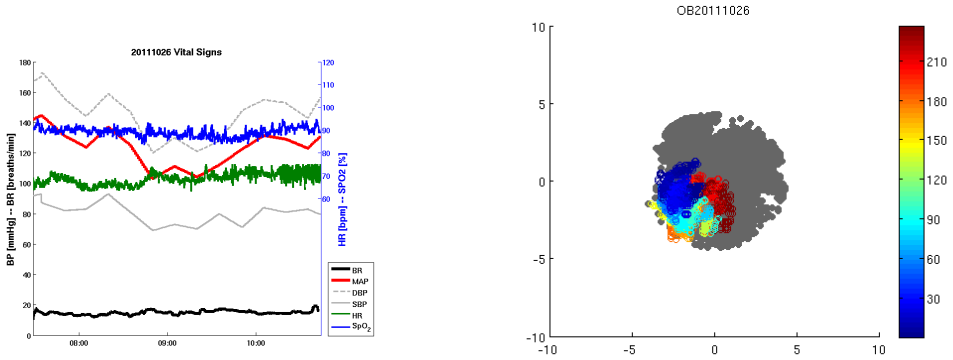
Figure 5.9 – Neuroscale projection of a non-IDH session for patient OB10.

signs from the normal sessions that form the training data for a patient give rise to a broad region of normality on the Neuroscale map. This is illustrated by the projections of the vital signs forming the training data for patient OB7 shown in Figure 5.2. As previously observed, the non-IDH sessions for patient OB7 shown in Figure 5.4 are projected onto slightly different parts of the map despite all consisting of normal data. Nevertheless, similarities can be identified in the projections of the vital signs from non-IDH sessions. At the start of each session, the vital signs, shown in dark blue, are projected onto the left half of the region of normality, and then move over time across the Neuroscale map from left to right. In all sessions, the vital signs at the middle of the session are plotted onto the centre of the map, which reflects the most normal physiology for that patient. The vital signs at the end of each of the sessions are projected onto the right half of the region of normality. Such a progression in the vital signs appears to be representative of a normal session for this patient.



(a) **(b)** **(c)** **(d)** **(e)** **(f)**

Figure 5.10 – Three IDH sessions for patient OB20 projected onto a patient-specific Neuroscale map.



(a) **(b)**

Figure 5.11 – Neuroscale projection of a non-IDH session for patient OB20.

5.5.2 IDH Session

As previously mentioned in Section 4.4.1, the axes of the Neuroscale map are dimensionless. As such, the map is rotation invariant. With the given set of weights, the normal trajectory for patient OB 7 goes from left to right during dialysis. With a different set of (random) initialisation weights, the trajectory could equally well go from right to left. Indeed this is the case for the Neuroscale map for patient OB 19, where the trajectory for a normal dialysis session (Fig. 5.7) is from right to left.

Comparing the non-IDH session for patient OB19 from Figure 5.7 with the IDH session shown in Figure 5.6, the abnormalities in patient physiology in the latter appear in the lower left quadrant, beyond the boundary of the region of normality. It is interesting to note that the projection of the vital signs at the start of the last IDH session (Fig. 5.6 (e) and (f)) corresponds with the projection of the third hour of data in the normal session shown in Figure 5.7 in terms of its location on the Neuroscale map. The patient's weight at the start of this IDH session was

below the prescribed dry weight for the session, and this appears to be reflected in the patient's physiology as projected onto the Neuroscale map.

5.5.3 Dialysis Onset

As the patient's vital signs (particularly the heart rate and blood pressure) are indicative of the patient's fluid status, the projection of the vital signs at the start of a dialysis session could potentially give insight into the degree of excess fluid in the patient's body. While this concept is beyond the scope of this thesis, the projection of the vital signs from non-IDH sessions for patient OB7 (shown in Fig. 5.4) illustrates that the vital signs at the start of each dialysis session, shown in dark blue in each case, are projected onto the same part of the map in every case. This can be contrasted with the results for patient OB19 described in Section 5.5.2, in which the vital signs at the onset of the third IDH session are projected onto a different region of the map than the vital signs at the start of the non-IDH session.

5.6 Conclusions

In this chapter the concept of a patient-specific Neuroscale map for dialysis patients is introduced. Each map is trained using vital sign data from seven normal sessions for that patient. The results of projecting vital sign data from both non-IDH and IDH sessions onto the Neuroscale map specific to that patient provide evidence for the benefits of defining vital sign normality in patient-specific terms in the context of dialysis. In particular, vital sign data from IDH events is projected outside the boundaries of normality defined by the patient-specific map.

These results form the basis of the work in the following chapter, in which a patient-specific data fusion model is proposed for the purpose of identifying and possibly predicting IDH in dialysis patients.

Chapter 6

Patient-Specific Data Fusion Models for HD Patients

The ultimate aim of this work is to develop a patient-specific data fusion model to detect and eventually predict patient deterioration in dialysis. This involves selecting appropriate input parameters for the model and creating a probabilistic model of normality using data from non-IDH sessions. The evidence from the visualisation studies in the previous chapters suggests that a patient-specific model trained using the vital signs from seven non-IDH sessions can distinguish between normal and abnormal vital signs. In this chapter the steps required to construct a data fusion model of normality are outlined. A patient-specific probability density distribution is constructed for each patient from normal data, which allows a probability to be estimated for new data. A resulting index of normality is calculated for both IDH and non-IDH test sessions in an attempt to establish the suitability of such a model in predicting IDH ahead of time.

6.1 A Patient-Specific Model for Detecting Patient Deterioration

While the previous work on novelty detection for patient deterioration is based on data fusion models derived from representative patient populations, the use of patient-specific models is suggested in this chapter based on the results from the patient-specific Neuroscale analysis in Chapter 5. In Section 5.4, it was demonstrated that an RBF network trained using patient-specific data from seven non-IDH sessions was capable of visualising the differences between IDH and non-IDH data. The same training dataset can be used to train a patient-specific data fusion model for detecting and possibly predicting IDH in a dialysis patient. In the following, a detailed outline of the construction of such a patient-specific data fusion model is given.

6.2 Creating the Model of Normality

6.2.1 Data Preprocessing

As the different vital signs in the dataset were recorded at different frequencies, the first step in data preprocessing consists of aligning the individual parameters such that a set of simultaneous input vectors is obtained. This is done by applying a zero-order hold between measurements and resampling the signals at 5-second intervals. Since the blood pressure readings were recorded at the lowest frequency (once every 15 minutes), the signal is interpolated at one-minute intervals prior to applying the zero-order hold. This is an important consideration, to which we

will return later. As previously, the artefact rejection limits given in Table 2.4 are applied.

The training and normalisation data used to create the model of normality are the same as for the patient-specific Neuroscale algorithm, as described in Section 5.2. For each patient, the vital signs from 7 non-IDH sessions form the training data. The vital sign data is normalised using the zero-mean unit-variance transform (given in Section 4.2), with the patient-specific means and standard deviations derived from the normalisation data, which consists of the 15 minutes of vital sign data preceding the heart rate nadir in non-IDH sessions, as described in Section 5.1. While the normalisation parameters are derived from 105 minutes of data, the training dataset consists of the vital signs from approximately 28 hours of data (7 dialysis sessions, each of 4 hours). In order to reduce the number of training patterns with minimal loss of information, the data is clustered into a subset of prototype patterns using the k-means clustering algorithm [140]. As with the Neuroscale algorithm, 500 clusters were chosen initially and the 50 least populated clusters were removed for being at the edges of the region of normality.

6.2.2 Parzen Windows

To construct a model of normality, the assumption is made that the vital sign data is drawn from an underlying Probability Density Function (PDF). Using the 450 prototype patterns generated by the k-means clustering algorithm, a PDF of the vital signs, which are treated as the random variable, is estimated such that the probability that a new data pattern $\mathbf{x}$ comes from the training data set of d-dimensional patterns, where $d = 4$ in this case. This is done with Parzen windows

[156], a method which places a kernel function on each point of the training data. The PDF is then estimated as a linear combination of the kernel functions.

The advantage of a Parzen windows model over other PDF estimation algorithms such as the Gaussian Mixture Model (GMM) for example, is that, as a non-parametric method, it makes no prior assumption as to the shape of the underlying distribution. The density function and the parameters of the model are derived from the training data.

Considering a region of volume V_n containing k_n of the N training patterns, the probability density p_n in the region containing $\mathbf{x}$ can be estimated by the fraction of samples that fall into the region normalised by the volume of the region:

$$p_n(\mathbf{x}) = \frac{k_n/N}{V_n} \tag{6.1}$$

In order to evaluate k_n, a window function θ is used such that

$$\int \theta(x)dx = 1. \tag{6.2}$$

The window function is usually chosen to be a PDF, so that the estimated probability density function p_n also satisfies the properties of a PDF. A popular choice of window function is a Gaussian PDF, as it is infinitely differentiable, making p_n infinitely differentiable as well:

$$p_n(\mathbf{x}) = \frac{1}{N(2\pi)^{d/2}} \sum_{i=1}^{N} \frac{1}{h_n^d} \exp\Big(-\frac{|\mathbf{x}-\mathbf{x_i}|^2}{2h_n^2}\Big) \tag{6.3}$$

where d is the dimension of the Gaussian, $\mathbf{x_i}$ denotes its centre, and h_n is the window width, which in the case of Gaussian distributions, can be considered as

the variance σ of the Gaussian:

$$p(\mathbf{x}) = \frac{1}{N(2\pi)^{d/2}\sigma^d} \sum_{i=1}^{N} \exp\Big(-\frac{|\mathbf{x}-\mathbf{x_i}|^2}{2\sigma^2}\Big) \tag{6.4}$$

The probability density p can thus be expressed as a sum of spherical Gaussians centred at every training pattern $\mathbf{x_i}$, normalised by the number of training patterns, N. The variance σ is typically taken to be the average over all training patterns of the average distance between each training pattern and its 10 nearest neighbours [157].

6.2.3 Novelty Calculation

Having established an estimate of the underlying probability density of the training data, we can compute the novelty Q of a new test pattern $\mathbf{x}$ as follows:

$$Q = \log\Big(\frac{1}{p(\mathbf{x})}\Big) - \log\Big(\frac{1}{p_{max}(\mathbf{x})}\Big) = \log\Big(\frac{p_{max}(\mathbf{x})}{p(\mathbf{x})}\Big) \tag{6.5}$$

Here, novelty is defined to be the logarithm of the reciprocal of the probability density function, in order to emphasize data with a low probability density. p_{max} is the value of the PDF evaluated at the normalised mean (0000), which represents the most probable and therefore highest value of $p(\mathbf{x})$. This is subtracted in order to ensure that the novelty of data is zero when all parameters are normal. The novelty Q evaluated according to Equation 6.5 provides a measure of abnormality; the higher Q is, the more abnormal the data is. Novelty can therefore be used as a Patient Status Index (PSI), with low values of PSI indicating normality.

6.3 Application to Dialysis Data

For each patient, a model of normality is trained as described above and for each session the PSI is calculated retrospectively from the four-dimensional vital sign data, which is normalised using the mean and standard deviation of the training data for each particular patient. In the absence of data from one of the four channels, the missing parameter is held at the last available measurement for one minute. If the data is missing for longer, the PSI is calculated using the mean value of the missing parameter from the training data. The mean is chosen as it reflects the ”most normal” value for the parameter and therefore will not affect the value of the PSI. Possible deteriorations occurring while one parameter is missing may still be detected by increases in the PSI as a result of changes in the other inputs. At least three of the four parameters are needed to calculate the PSI. However, it should be noted that in calculating the PSI of the training data, missing data is held for one minute only. Subsequently, data from all other channels is also discarded until data from all channels is available again.

To distinguish between normal and abnormal vital signs, a threshold is set on the PSI calculated from the training data. Previous approaches to determining an optimal threshold in population-based models include the use of fixed empirical thresholds [158] and more complex statistical methods [159]. Given the patient-specific nature of our model, the threshold is set to be the 99th percentile of the PSI values generated for the seven non-IDH sessions for each patient. This ensures that 99% of the training data is classified as normal. The percentile values for the three exemplar patients considered in the previous chapter are given in Table 6.1.

For each of the three patients, OB10, OB19 and OB20, considered in the pre-

Percentiles	OB10	OB19	OB20
90^{th}	1.80	1.98	2.65
95^{th}	2.32	2.47	3.20
99^{th}	3.51	3.54	4.05

Table 6.1 – Percentiles of the PSI for the training sessions of OB10, OB19 and OB20.

vious chapter, the PSI is calculated for the same three IDH sessions and one non-IDH session. The PSI at the values of the vital signs for each session are shown in Figures 6.1, 6.2 and 6.3. The vital signs are shown in the upper panel, while the PSI is shown in the lower panel. The PSI threshold for each patient is shown in red on the PSI plot. In each figure, the non-IDH session in the test set is shown in (a), while the IDH sessions are shown in (b) to (d). The means and standard deviations of the training data for each patient are shown in Table 6.2.

The means and standard deviations of the heart rates and oxygen saturations are similar across all three patients with the exception of the mean heart rate for OB10 which is slightly higher than that for OB19 and OB20. The mean MAP for OB10 is lower than that of OB19 and OB20. The standard deviation of the MAP for patients OB10 and OB20 are between 16 and 17 mmHg, while that of OB19 is at 14.1 mmHg.

Given the importance of the blood pressure in identifying episodes of IDH in dialysis patients, two approaches are taken for calculating the PSI. In the first instance, the interpolated MAP values are used to compute the PSI in order to establish whether the patient-specific PSI is capable of identifying abnormalities in the vital signs and predicting IDH. Following from this analysis, the PSI is then re-calculated using the zero-order hold MAP values. As the MAP is measured every 15 minutes, this in effect serves to determine whether an IDH event can be

predicted 15 minutes in advance of its occurrence, using the definition of Section 1.4. In both cases, only interpolated MAP values are used to train the model of normality.

6.3.1 Persistence Criterion

In order to avoid generating alerts for transient changes in physiology for which the body can compensate, it is common for a persistence criterion to be introduced such that the PSI is required to exceed the designated threshold for a certain length of time within a given time window [158]. For the following sessions, the timing of alert generation with respect to the IDH event is considered for three different persistence criteria: the PSI is required to exceed the threshold for either two minutes in a three-minute window, three minutes in a four-minute window, or four minutes in a five-minute window (referred to as the minimum, medium and maximum persistence criteria, respectively). While a longer persistence criterion may result in a reduction in the number of false alerts arising from transient changes, it may also result in the removal of true positives, and as such it needs to be optimised. The timings of the first alert generated using the different persistence criteria are shown in Table 6.3 for each session.

6.3.2 OB 10

In the non-IDH session considered for patient OB10, the PSI is lower than the 3.51 threshold (see Table 6.1) throughout the session, except for three instances in which the threshold is exceeded. However, as can be seen in Table 6.3, the PSI does not exceed the threshold for a long enough period to generate an alert using

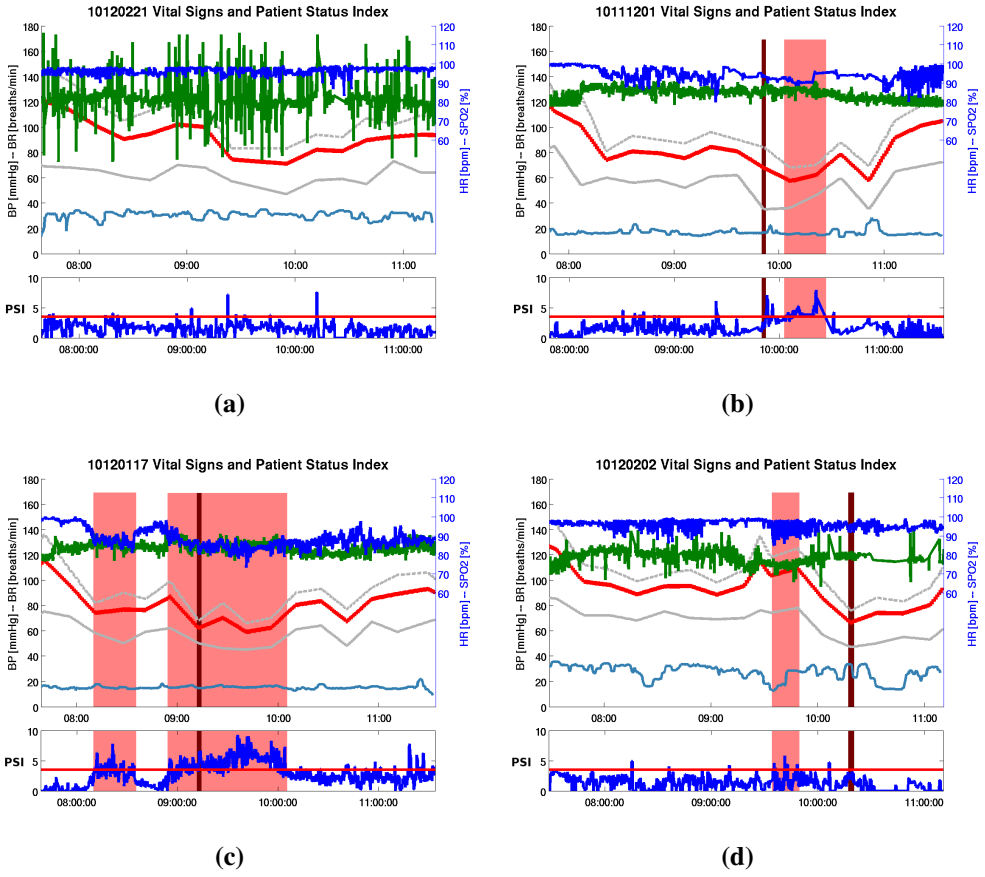


Figure 6.1 – OB10: Vital signs and PSI for one non-IDH and three IDH sessions. IDH events are highlighted in dark red, while elevations in the PSI above the threshold are highlighted in light red.

any of the three persistence criteria.

In the first IDH session considered for patient OB10, the threshold is exceeded by the PSI for nearly 30 minutes starting just after 10:00 (see Fig. 6.1 (b), highlighted in light red). This overlaps with a decrease in the MAP from 80 mmHg to 62 mmHg, which occurs after 10:00. While the MAP recovers in the following 30 minutes, it falls again just before 11:00, again to 62 mmHg. However, at this second decrease no abnormality is detected by the PSI. Although the IDH event

	OB10		OB19		OB20	
	μ	σ	μ	σ	μ	σ
HR (bpm)	77.5	3.0	69.4	5.9	70.5	6.0
SpO_2 (%)	95.7	4.0	96.2	2.5	96.0	3.2
MAP (mmHg)	84.5	16.8	117.8	14.1	120.8	16.1
BR (breaths/min)	17.9	6.2	23.7	9.1	21.5	6.8

Table 6.2 – Means and standard deviations of the vital signs for OB10, OB19 and OB20 used for model normalisation.

occurs at 9:51:15 (highlighted in dark red on the plot), an alert is not generated until 10:09:45 with the minimum persistence criterion (see Table 6.3).

In the second IDH session (shown in Fig. 6.1(c)), the MAP decreases from 113 mmHg at the start of the session to 57 mmHg within two hours. Following the first decrease to 76 mmHg within the first 30 minutes, the oxygen saturation falls below 90% and fluctuates in the second hour of the session between 85% and 97%. These variations in the vital signs result in the PSI exceeding this patient's threshold at two instances during the session, the first being at 8:12:00 and the second starting just before 9:00:00. The change in blood pressure at the start of the session does not satisfy the definition of IDH, while the second decrease in MAP gives rise to an event at 09:12:31. All three persistence criteria give rise to an alert prior to the event, although the minimum criterion generates an alert more than 45 minutes earlier than the medium and maximum criteria, which generate an alert 7 minutes (medium criterion) and 26 seconds (maximum criterion) in advance of the event, respectively.

In the final IDH session (shown in Figure 6.1 (d)) there is a sudden fall in blood pressure from 114 mmHg to 64 mmHg which is preceded by a desaturation. However, this is not detected as an abnormal change in vital signs by the PSI. The PSI exceeds its threshold three times while the blood pressure is decreasing (as

highlighted in light red), but there is no prolonged elevation in the PSI. This does not result in an alert for any of the three persistence criteria. As can be seen in Table 6.2, the mean MAP for this patient is 84.5 mmHg, with a standard deviation of 16.8 mmHg. Thus a change in MAP from 114 mmHg, which is within two standard deviations from the mean, to 64 mmHg which again is within two standard deviations from the mean, is unlikely to be detected as an abnormality in the absence of changes in the other parameters. However in this session the other vital signs remain within two standard deviations from their mean, and therefore the decrease in MAP is not highlighted as an abnormality by the PSI in this patient.

6.3.3 OB 19

The non-IDH session for patient OB19 is shown in Figure 6.2 (a). Although the blood pressure in this session does not satisfy the definition of IDH, the PSI exceeds the patient-specific threshold of 3.54 (see Table 6.1) towards the end of the session, as highlighted in light red. At this time, the MAP falls to 83 mmHg while the heart rate is between 75 and 90 bpm. The mean heart rate from the training data is 69.4 bpm with a standard deviation of 5.9 bpm, while the mean of the MAP is 117.8 with a standard deviation of 14.1 mmHg. Therefore, both the MAP and the heart rate are more than two standard deviations from the mean in the final 30 minutes of the session. At its lowest point, the MAP is 2.5 standard deviations from the mean value of the training data (the other two vital signs are within two standard deviations of the mean throughout the session). The vital signs in the final 30 minutes of the session therefore give rise to a high PSI which exceeds the threshold. All three criteria give rise to alerts in advance of the lowest blood

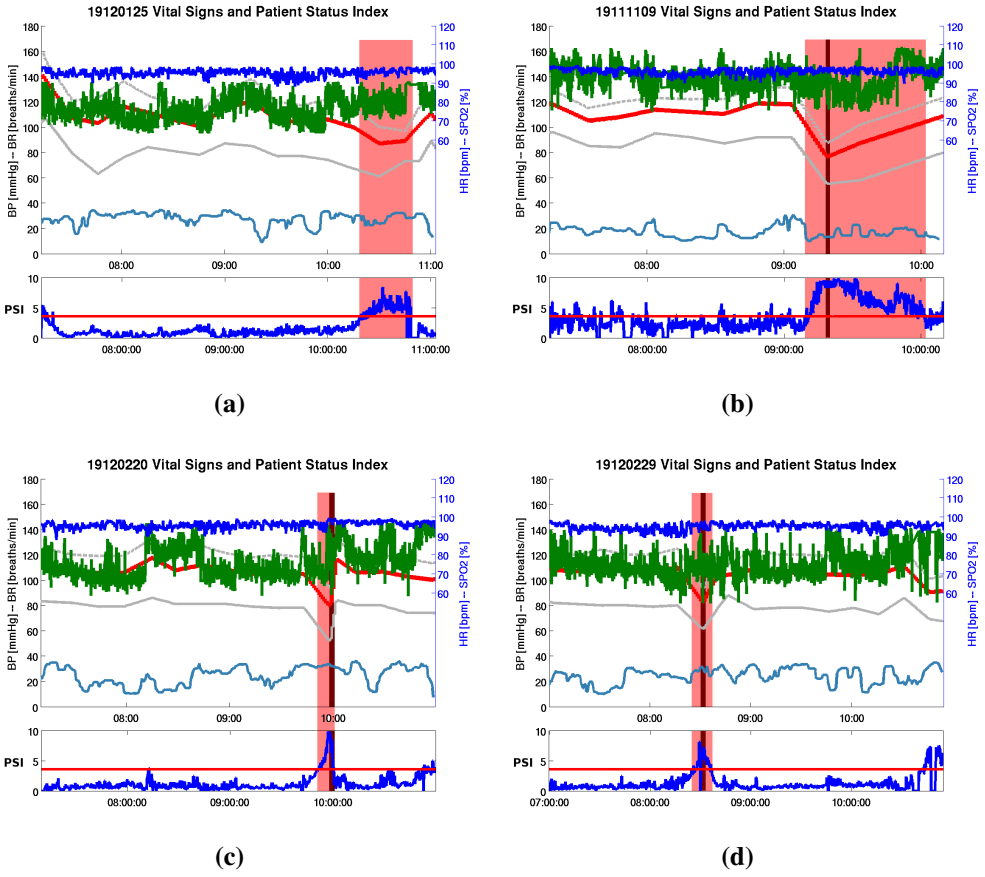


Figure 6.2 – OB19: Vital signs and PSI for one non-IDH and three IDH sessions. All three IDH events are detected by the PSI ahead of their occurrence by at least one of the three persistence criteria. While the first session does not satisfy the definition of IDH, the PSI is elevated above the threshold towards the end of the session.

pressure at 10:29:46, despite the fact that this blood pressure does not constitute an IDH event by definition. In contrast, while the heart rate varies between 75 and 90 bpm in the 20 minutes after 9:00 as well, the MAP is between 110 and 120 mmHg at that time, which is within one standard deviation from the MAP mean. This session highlights the advantages of a patient-specific model of normality in identifying abnormalities in vital signs which would not be considered abnormal using a global definition of IDH.

In the first IDH session for OB19, the MAP is relatively constant in the first two hours of the session, varying between 103 mmHg and 120 mmHg. While this does not result in persistently high values in the PSI, fluctuations in the heart rate between 86 and 110 bpm give rise to temporary spikes in the PSI that exceed the alert threshold. Just after 9:00 however, the MAP suddenly decreases to 77 mmHg which is 2.9 standard deviations from the mean MAP for this patient. An IDH event occurs at 9:18:46, although no significant changes can be observed in the other vital signs at that time. As the MAP begins to decrease, the PSI exceeds the threshold and remains above it until the end of the session. The PSI first exceeds the alert threshold between 9:13:09 and 9:14:46 depending on the criterion used, which is up to 5 minutes prior to the event.

Similar to the first IDH session, the MAP is constant in the first three hours of the second IDH session shown in Figure 6.2 (c), and then suddenly decreases with an event occurring at 9:57:30. This is accompanied by a sudden increase in heart rate. The rapid fall in the MAP causes a sharp rise in the PSI, which is otherwise far below the threshold in this session. All three persistence criteria generate an alert prior to the event, at 9:53:00, 9:53:50 and 9:54:41, respectively.

In the final IDH session, there is a sudden fall in MAP just after 8:30, which is preceded by a gradual decrease in heart rate from the start of the session. During the IDH episode, which occurs at 8:31:34, the heart rate suddenly increases by 10 bpm. The PSI exceeds the threshold during the decrease in MAP, and alerts are generated at 8:28:52, 8:31:07 and 8:32:19 by the different persistence criteria. Therefore, only the minimum and medium criteria generate an alert ahead of the IDH event.

6.3.4 OB 20

In the non-IDH session considered for OB20, there are no sudden decreases in the MAP, nor are there any other abnormalities in the vital signs. At the start of the session however, the patient appears to be fluid overloaded given that the MAP is 1.6 standard deviations above the mean value for this patient. As a result, the PSI exceeds its threshold immediately at the start of the session. However, given that patients are frequently fluid overloaded at the start of dialysis sessions, any alerts generated within the first 15 minutes of a session are not considered for the purpose of our analysis. It should be noted however that the PSI introduced in this chapter could equally find application for identifying abnormalities in physiology at the onset of dialysis.

In the first IDH session for this patient, shown in Fig. 6.3 (b), there are two sharp decreases in MAP which constitute IDH events, one at 9:30:23 and another at 10:48:35. In the first instance, the MAP decreases within 15 minutes from 88 mmHg, which is two standard deviations below the mean, to 61 mmHg, which is 3.7 standard deviations below the mean. In the second instance, the MAP decreases from 74 mmHg to 62 mmHg over the course of 30 minutes. None of the persistence criteria generate an alert in advance of the first decrease in MAP, while all three criteria generate an alert prior to the second decrease in MAP.

In the second IDH session (shown in Fig. 6.3 (c)), the MAP is low from the onset of the session, and an IDH event occurs at 9:02:49. The threshold is exceeded at numerous instances within the first half hour of the session, and each of the persistence criteria generate an alert at 7:36:03 (minimum), 7:36:53 (medium) and 7:37:44 (maximum), respectively. Thus all three criteria generate an alert

Session	Event	min.	med.	max.
OB10 120221	no event	no alert	no alert	no alert
OB10 111201	09:51:15	+00:18:30	+00:19:21	+00:20:12
OB10 120117	09:12:31	-00:50:34	-00:07:18	-00:00:26
OB10 120202	10:18:25	no alert	no alert	no alert
OB19 120125	10:29:46	-00:06:00	-00:05:08	-00:04:10
OB19 111109	09:18:46	-00:05:37	-00:04:00	no alert
OB19 120220	09:57:30	-00:04:30	-00:03:40	-00:02:49
OB19 120229	08:31:34	-00:02:42	-00:00:27	+00:00:45
OB20 111026	no event	no alert	no alert	no alert
OB20 120608	09:30:23	+00:02:22	+00:06:07	+01:17:27
OB20 120727	09:02:49	-01:12:50	-01:12:33	-01:11:25
OB20 111221	08:49:33	-00:16:45	-00:16:37	-00:16:28

Table 6.3 – Timings of alert generation using different persistence criteria. Alert timings are given with reference to In cases in which all three criteria generate an alert prior to the event, the earliest alert is always generated by the minimum criterion, as expected. In the following work, the minimum criterion is used for alert generation.

ahead of the IDH event in this session.

For the final IDH session, all three persistence criteria give rise to an alert between 8:32:48 and 8:33:05, which is up to 16 minutes prior to the IDH event. Within the first hour of the session, the MAP decreases from 101 mmHg to 67 mmHg while all other vital signs remain stable. This gives rise to a gradual increase in the PSI, eventually leading to an alert. The MAP values are low for this particular patient, whose mean MAP from the training data is 120.8 mmHg with a standard deviation of 16.1. Thus once the MAP falls below 88 mmHg, it is more than two standard deviations below the mean.

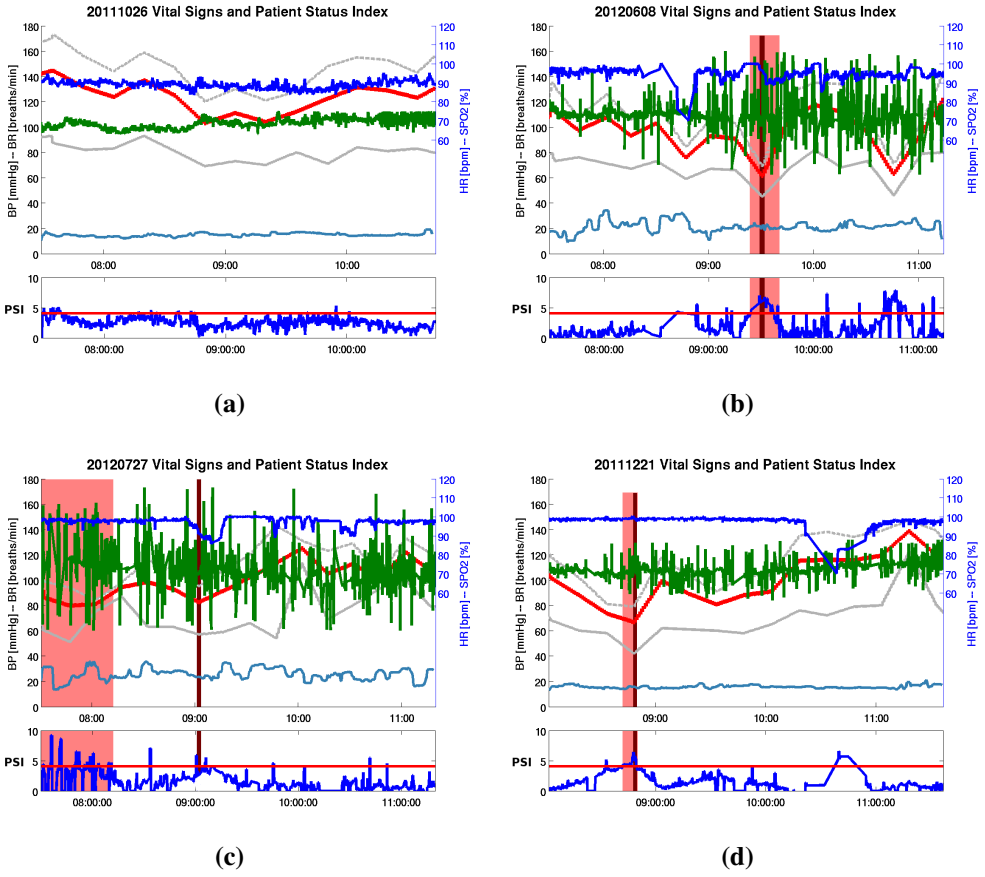


Figure 6.3 – OB20: Vital signs and PSI for one non-IDH session and three IDH sessions. In the non-IDH session, the PSI exceeds the threshold of 4.05 at the start of the session, due to the patient being fluid overloaded.

6.3.5 Conclusions

The results from this section indicate that a patient-specific PSI with a threshold based on the 99^{th} percentile of the PSI values calculated using the training data and a persistence criterion is capable of distinguishing abnormal vital signs from normal ones. In particular, of the nine IDH sessions considered, six IDH events were detected ahead of time by a prolonged elevation of the PSI above the threshold for at least two minutes in a three minute window. In two of the remaining

three sessions, alerts were also generated, but not in advance of the IDH event. Only in one of the nine IDH sessions no alert was generated at all.

The minimum persistence criterion, for which the PSI must exceed its threshold for at least two minutes in a three-minute window, performed best in the sessions considered. In the cases for which alerts are generated using any of the three criteria, the alert generated by the minimum criterion precedes those of the medium and maximum criteria. Therefore, in the following analysis, the first persistence criterion will be used to determine the timing of the generated alerts.

Further, the results illustrate the occasional shortcomings of a global definition of IDH which only takes changes in systolic blood pressure into account. More specifically, the results in Figure 6.2 (a) demonstrate that although the blood pressure values may not satisfy the definition of IDH in a particular instance, the overall combination of vital signs may nevertheless indicate abnormality in that patient.

Finally, as previously noted, the results of this section are based on the interpolated MAP values which would not be available in real time. Given that we have established that the PSI can distinguish between normal and abnormal vital signs when the trends in the MAP are indicated by interpolation, we now investigate whether IDH can be predicted using the zero-hold MAP values. In the majority of cases in this section, while the PSI exceeds the threshold in advance of the lowest MAP value, it is only in two instances that this occurs more than 15 minutes (the interval between blood pressure recordings) in advance of the lowest MAP value. It should be noted in this context however that while we assume that the IDH event occurs at the lowest MAP value, the timing of the event is also subject to the low sampling frequency of the MAP and cannot be determined accurately.

Session	Event	Interpolated	Zero-order hold
OB10 120221	no event	no alert	no alert
OB10 111201	09:51:15	+00:18:30	+00:18:27
OB10 120117	09:12:31	-00:50:34	-00:50:38
OB10 120202	10:18:25	no alert	-02:46:25
OB19 120125	10:29:46	-00:06:00	-03:03:28
OB19 111109	09:18:46	-00:10:37	+00:16:57
OB19 120220	09:57:30	-00:04:30	+00:02:54
OB19 120229	08:31:34	-00:02:42	+00:01:52
OB20 111026	no event	no alert	no alert
OB20 120608	09:30:23	+00:02:22	-00:47:45
OB20 120727	09:02:49	-01:12:50	-01:12:48
OB20 111221	08:49:33	-01:13:30	-01:13:13

Table 6.4 – Comparison of alert timings for the PSI calculated with interpolated and zero-order held MAP values. Alert timings are given with reference to the timing of the event.

6.4 PSI with zero-order held MAP

In this section, the PSI is calculated for the same sessions as in the previous section using the zero-hold MAP values instead of the interpolated ones. As the training data is still based on the interpolated MAP values, the thresholds for alert generation do not change (see Table 6.1). Alert generation is now based on the minimum persistence criterion. The resulting timings of alert generation in comparison with those from the previous section (using interpolated MAP values) are shown in Table 6.4.

6.4.1 OB 10

Comparing the timings of alert generation in the normal session for OB10, the PSI remains below the threshold for the entire session using both the zero-held

and interpolated MAP values. In both cases, no alert is generated for this session. In the first IDH session (shown in Fig. 6.4 (b)), the alert for the PSI using zero-held MAP values is generated at the same time as that of the interpolated MAP values. However, neither of the alerts precede the IDH event, which occurs at 09:51:15.

In the second IDH session (Fig. 6.4 (c)), the alert generated by the PSI based on the zero-order held MAP is again generated at the same time as that of the PSI with interpolated MAP values. In this case, it precedes the event by 50 minutes. In the final IDH session (Fig. 6.4 (d)), while the PSI with interpolated MAP values does not generate an alert, the PSI based on zero-order held MAP values generates an alert ahead of the event which is at 10:18:25. The alert is generated at the start of the session as a result of a drop in MAP from 125 mmHg to 98 mmHg. This example illustrates that for a model trained using interpolated MAP values, sudden changes in the zero-order held MAP values of the test data may generate alerts where interpolated values do not. On this occasion, the model performs better on the zero-order held MAP data than on the interpolated data.

6.4.2 OB 19

While the alert timings for patient OB10 were similar regardless of the kind of MAP data used, the results for the PSI using zero-order hold MAP values are significantly different from those using interpolated MAP values for patient OB19. In particular, while alerts are generated ahead of the IDH event in all three IDH sessions when using interpolated MAP values, no alerts are generated ahead of the IDH events when using the zero-order held MAP values.

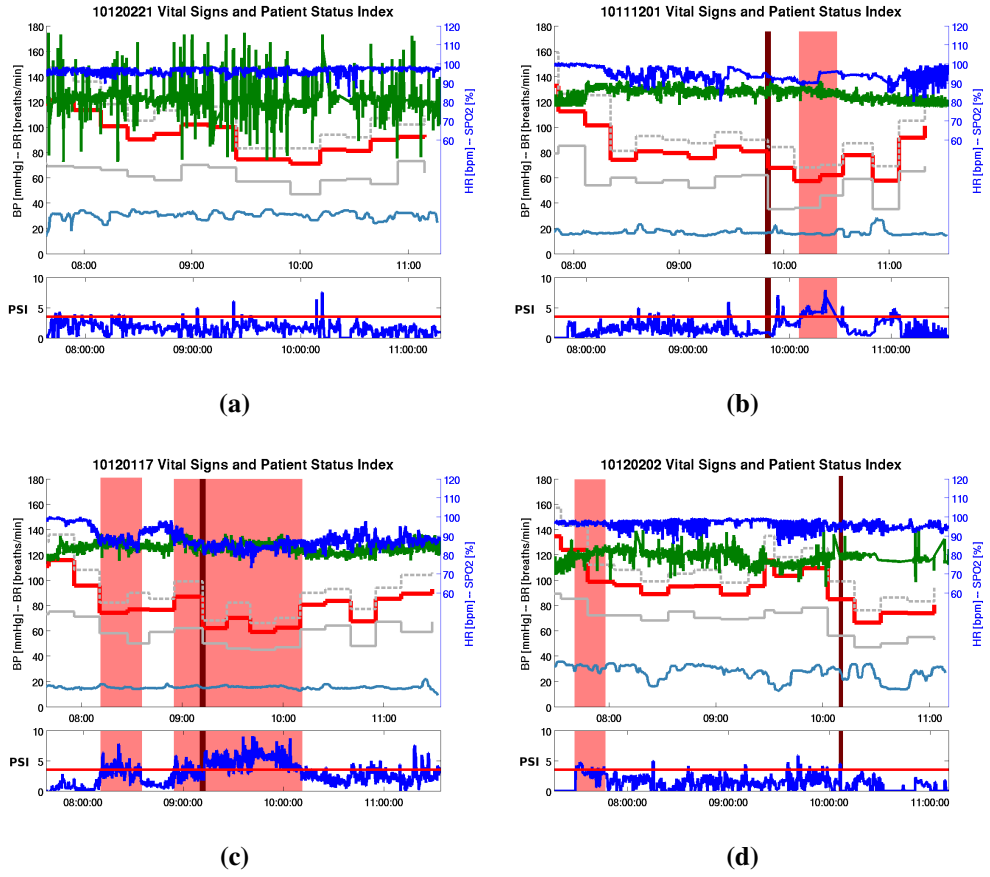


Figure 6.4 – OB10: Vital signs and PSI with zero-order held MAP data.

In the first session, which as previously mentioned is considered to be a session with abnormal physiology despite not satisfying the general definition for IDH, the PSI with zero-held MAP also exceeds the threshold. In contrast to the PSI with interpolated MAP however, this occurs both at the start of the session, as well as towards the end.

In the first IDH session (shown in Fig. 6.5 (b)), while the PSI exceeds the threshold on several occasions prior to the IDH event at 9:18:46, none of these occurrences is long enough for an alert to be generated. In this session, while the PSI based on interpolated MAP values generates an alert 5 minutes prior to the

event, the PSI is unable to generate an alarm prior to the event without having an indication of the trend of the MAP between actual measurements. The same argument applies to the second IDH session (shown in Fig. 6.5 (c)), in which the PSI based on interpolated MAP values generates an alert ahead of the event at 9:57:30, while the PSI based on zero-order held MAP values does not generate an alert ahead of the event. Similarly, in the third IDH session (Fig. 6.5 (d)) the PSI based on zero-order held MAP values does not generate an alert ahead of the event, while the PSI with interpolated MAP values does.

6.4.3 OB 20

As described in the previous section, the PSI exceeds the threshold at the onset of the non-IDH session for patient OB20 (Fig. 6.6 (a)) due to the fluid overloaded status of the patient, but this alert is discarded for the purpose of our analysis. With the exception of this alert, no further alerts are generated over the course of the non-IDH session. In the first IDH session (Fig. 6.6 (b)), the PSI based on zero-hold MAP values generates an alert at 8:42:38. In contrast to the PSI with interpolated MAP values, this alert is generated ahead of the IDH event at 9:30:23. This is primarily due to the abrupt fall in MAP at 8:40:00 which causes a step change in the PSI and makes it exceed the threshold.

In the second IDH session (Fig. 6.6 (c)), the PSI exceeds the threshold on several occasions in the first hour of the session, thus generating an alert ahead of the IDH event at 9:02:49. In the final IDH session, an alert is generated prior to the IDH event at 8:49:33, although only less than two minutes in advance of the event. In comparison, the PSI with interpolated MAP values generates an alert

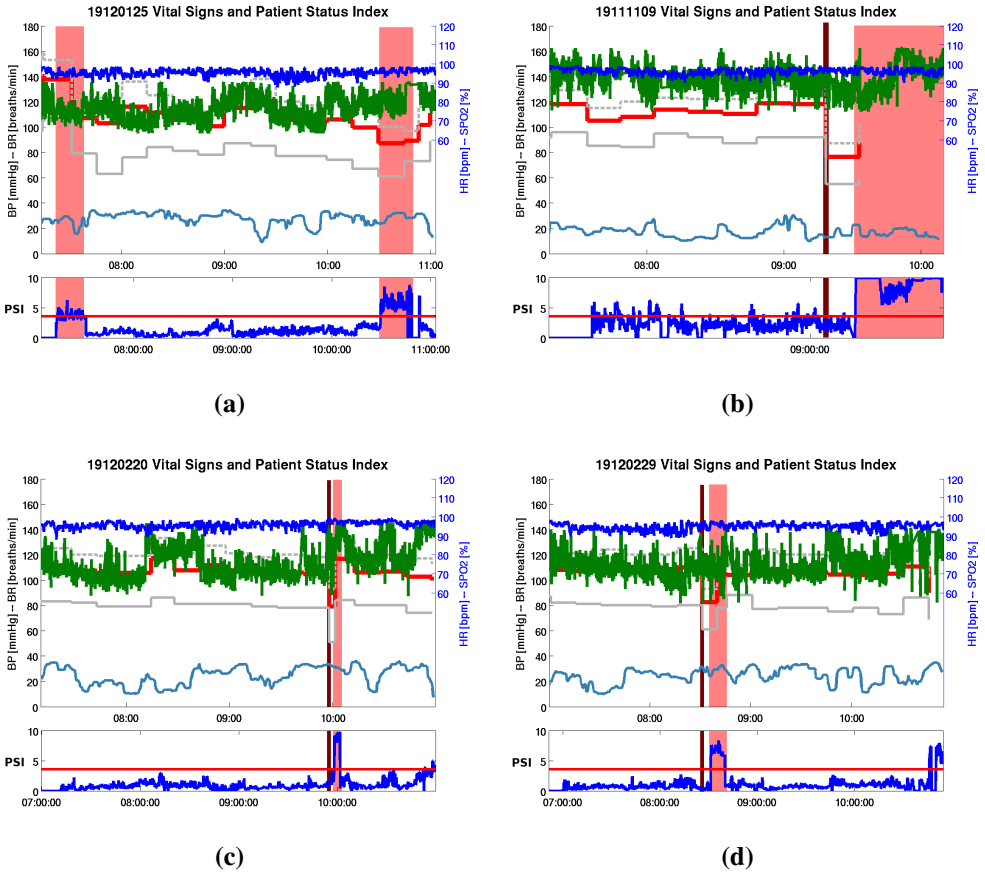


Figure 6.5 – OB19: Vital signs and PSI with zero-order held MAP data. None of the IDH events are detected ahead of time when using the zero-order held MAP values, while all IDH events are detected prior to their occurrence using the interpolated MAP values.

more than 15 minutes ahead of the IDH event.

6.5 Performance Evaluation

To evaluate the ability of the PSI to discriminate between abnormal vital signs that lead to IDH events and normal physiology, both the PSI based on interpolated MAP values as well as that based on zero-hold MAP values are calculated for 34

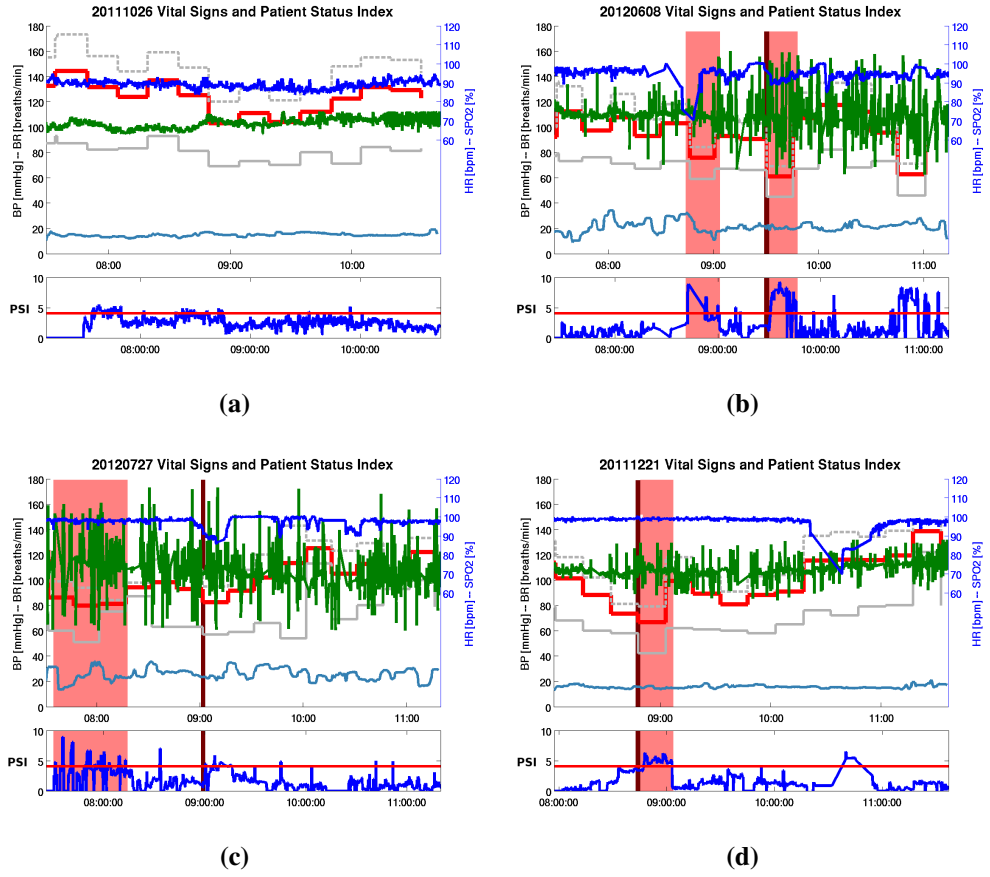


Figure 6.6 – OB20: Vital signs and PSI with zero-order held MAP data.

sessions from the subset of 6 patients mentioned in Section 2.5. All these patients have at least 7 non-IDH sessions to constitute the training data, as well as at least one IDH session and a further non-IDH session on which to test the PSI.

The primary objective of the PSI is to alert the nursing staff to changes in the vital signs which are considered abnormal and which could potentially lead to a potential IDH event during dialysis being diagnosed in advance of the event. In order to do this most effectively, the PSI must on the one hand correctly identify genuine deteriorations in patient physiology and flag them to the nursing staff as soon as possible, but on the other hand is must also correctly recognise normal

	IDH session	non-IDH session
Alert	True Positive (TP)	False Positive (FP)
No alert	False Negative (FN)	True Negative (TN)

Table 6.5 – Confusion Matrix

physiology which does not warrant an intervention. These two outcomes, which are known as true positives and true negatives respectively, should ideally be maximised. Equally, both the number of alerts generated in the absence of genuine events (known as false positives), as well as the number of genuine events which do not trigger an alert (known as false negatives) should ideally be minimised. This is to avoid wasting resources and desensitising nurses to the alerts in the case of false positives, and to avoid delays in potential preventative interventions which could lead to further deterioration in the case of false negatives. These outcomes can be combined to form a confusion matrix [160], as shown in Table 6.5. Further, the sensitivity, which is a measure of the ability of the PSI to correctly identify positive events, as well as the specificity, which is a measure of the ability of the PSI to correctly identify normal physiology, can be established using the above outcomes as follows:

$$\text{Sensitivity} = \frac{\text{True Positives}}{\text{True Positives}+\text{False Negatives}}, \quad \text{Specificity} = \frac{\text{True Negatives}}{\text{True Negatives}+\text{False Positives}}$$

Of the 34 sessions considered, 19 sessions have IDH events while the remaining 15 sessions are non-IDH sessions. For each IDH session, the timing of the first alert, if applicable, is established. If the alert is generated prior to the event, this is counted as a true positive, while alerts that are generated at or after the event are classified as false negatives. Similarly, each non-IDH session in which no alert is generated is considered to be a true negative, while those in which alerts are

	IDH session	non-IDH session
Alert	TP = 11	FP = 5
No Alert	FN = 9	TN = 10

Table 6.6 – Confusion matrix for alert generation by the PSI based on interpolated MAP values.

	IDH session	non-IDH session
Alert	TP = 10	FP = 6
No Alert	FN = 10	TN = 9

Table 6.7 – Confusion matrix for alert generation by the PSI based on zero-hold MAP values.

generated are considered to be false positives. The resulting confusion matrices for both the PSI based on interpolated MAP values and that based on zero-hold MAP values are shown in Tables 6.6 and 6.7.

The resulting sensitivities are given by 0.55 for the PSI based on interpolated MAP values, and 0.5 for the PSI based on zero-order held MAP values. The specificity is 0.67 for interpolated MAP values and 0.4 for zero-order held MAP values. For the 10 true positive alerts obtained using zero-hold MAP values, the lead times of the alerts ahead of each event are more than 15 minutes in all cases, with half of the true positive alerts occurring more than one hour in advance of the event. For the true positive alerts obtained with interpolated MAP values, the lead times of the alerts are below 15 minutes in 4 of the 11 cases. In the other cases, the lead times are between 45 minutes and 1.5 hours ahead of the event.

The results for the PSI with interpolated MAP values are better than those of the PSI with zero-hold MAP values, indicating the need for either more frequent blood pressure measurements or alternatively more sophisticated methods for es-

timating trends in the blood pressure between actual measurements. In similar applications, Gaussian processes have been suggested for this purpose [161]. In dialysis in particular however, an alternative method for estimating blood pressure trends has been suggested [162] and will be considered in more detail as part of the Future Work presented in Section 7.2.

This should not only improve the timing of alert generation by the PSI such that events are detected ahead of time, but could also impact upon the timing of events. At present, the definition of IDH as given in Section 1.4 is based upon an abrupt change in the blood pressure over the course of a fixed period of time. Given that in our dataset the blood pressure is measured once every 15 minutes and the definition of an IDH event is based exclusively on changes in blood pressure, an IDH event can only occur at these times during the session which correspond with actual blood pressure measurements. As such, unless abnormalities in the other vital signs occur at least 15 minutes prior to a new blood pressure measurement which satisfies the definition of IDH, it is not possible for the PSI based on zero-hold MAP values to generate an alert in advance of the time at which the IDH is observed. This is illustrated by the results in Table 6.4, where events which were preceded by an alert using the PSI based on interpolated MAP values by less than 15 minutes did not result in alerts being generated ahead of the event using the PSI based on zero-hold data in some cases. In contrast, when the PSI with interpolated MAP values predicted the event more than 15 minutes prior to its occurrence, the PSI with zero-hold MAP values generated an alert ahead of the event as well.

A further point in which the suggested approach could be improved is by optimising the threshold for alert generation. At present, the threshold is empirically

set to be the 99^{th} percentile of the PSI values from the training data for each patient. Given a sufficient amount of test data for each patient, this could be further optimised on a patient-specific basis by minimising the cost function C given by:

$C = \alpha(1 - \text{sensitivity}) + (1 - \alpha)(\text{specificity})$ where α represents the cost of a false negative.

Finally, it should be noted that the data fusion model is based on MAP values, while the definition of IDH is based on an abrupt fall in systolic blood pressure. Therefore, the definition of IDH is not incorporated into the data fusion model and novelty detection relies on a departure from the envelope of normality (as defined by the threshold) in four dimensions. Although MAP constitutes only one of the four vital signs in the model, it appears that the multi-variate novelty is still mostly driven by decreases in blood pressure.

Chapter 7

Conclusion

7.1 Summary of Results

While IDH is the most common complication in HD treatment [5] and has been linked with increased mortality in HD patients [7], the mechanisms giving rise to IDH events are multi-faceted and poorly understood, with no consensus as to a universally applicable definition of IDH in terms of patient physiology. While continuous monitoring of vital signs is not currently a part of routine care during HD, the benefits of continuously monitoring patients during HD were investigated in this thesis with regard to different methods for detecting and predicting IDH. Beyond the general advantages of continuous vital sign monitoring in predicting patient deterioration [101], HD treatment in particular lends itself to accumulating vital sign data from individual patients over time with a view to creating patient-specific models of normality for each patient, given the repetitive nature of this treatment modality.

Chapter 2 introduced the observational data collection study which was car-

ried out at the Oxford Renal Unit in 2011/2012 and which forms the basis for the analysis in this thesis. The data was partitioned into subsets based on whether the vital signs in each session satisfied the definition of IDH adopted in this thesis. Differences were observed particularly in the heart rate and blood pressure distributions of these two subsets. While the occurrence of blood pressures in the range of 60 to 90 mmHg was only found to be slightly higher in the IDH sessions, as would be expected, the occurrence of heart rates between 50 and 70 bpm was lower in the IDH sessions. In contrast, the occurence of heart rates between 70 and 90 bpm was higher in the IDH sessions when compared with data from non-IDH sessions. This agrees with the general observation of a tachycardic response to fluid removal in HD patients as a compensatory mechanism for blood pressure maintenance.

Although a number of measures have been introduced for assessing HRV in HD patients [93], no conclusive evidence has yet been provided for the predictive value of these measures for IDH. Chapter 3 introduced a new spectral measure for HRV which departs from the conventional definition of frequency bands in the spectrum of RR intervals and moves towards a dynamic approach in which breathing rate information is used to identify the physiologically relevant peaks. The ratio of peaks leading up to the IDH event was assessed as a possible measure for prediction. This measure was shown to be capable of distinguishing between sessions in which patients were able to compensate for decreases in blood pressure by an increase in heart rate and sessions in which such a response was absent. However, while changes in the proposed measure correlate well with the magnitude of change in blood pressure in a patient with autonomic dysfunction, such correlation cannot be observed in patients without autonomic dysfunction.

Different population-based models of normality were considered in Chapter 4 in terms of their ability to distinguish vital signs recorded during IDH episodes from normal data. A model trained on normal data from the entire patient population was shown not to be able to distinguish IDH data, primarily due to the wide range of vital sign data included as part of the training data. Two alternative models of normality were investigated, the first being trained on relative changes in vital signs with respect to the start of the session and the second being based on a condition-specific stratification of patients with and without autonomic dysfunction. These performed better than the global population model, yet each was shown to have disadvantages. In the model based on relative changes with respect to the start of the session, large changes in the vital signs, even if occurring gradually over the course of the session are likely to be identified as being abnormal. While the condition-specific model performed better in distinguishing between normal and abnormal data from patients with autonomic dysfunction, no improvement was observed in the model of the remaining patients without autonomic dysfunction.

Chapter 5 introduced the concept of a patient-specific model of normality, in which the model is trained and tested with data from a particular patient only. When applied to both IDH and non-IDH sessions, patient-specific models were shown to correctly distinguish between normal and IDH data. In particular, vital sign data from IDH events is projected outside of the boundaries of normality as defined by the patient-specific map, while data free of IDH events is projected within the boundaries.

The main objective of this thesis was to determine whether continuously monitoring vital sign data during haemodialysis could assist in distinguishing between

”normal” and ”abnormal” patterns of physiology in individual patients, with a view to predicting episodes of IDH. As mentioned above, patient-specific models of normality can correctly distinguish between vital signs from an IDH episode and those recorded during regular haemodialysis, and based on these results, a patient-specific data fusion model was constructed in Chapter 6 with the objective of predicting IDH events. Given that the frequency of blood pressure measurements during the study was low (every 15 minutes), the model was trained and tested with both interpolated and zero-held blood pressure values. Nevertheless, the low sampling frequency of the blood pressure introduced several limitations to the analysis. Firstly, with regard to the timing of the event against which the performance of the data fusion model is being measured, it should be noted that the adopted definition of IDH relies solely on changes in blood pressure with respect to the previous reading, and therefore by definition imposes that the IDH event can only occur at times when the blood pressure is measured. Similarly, in the absence of abnormalities in the other vital signs which precede the IDH event by more than the sampling frequency of the blood pressure, it is not possible for an alert to be generated ahead of the event using the zero-hold blood pressure data. Thus both the timing of the event as well as the timing of the alert are heavily influenced by the low sampling frequency of the blood pressure. While the sensitivity of the model in predicting IDH using interpolated blood pressure values was better than for the model using zero-hold blood pressure values, blood pressure interpolation is non-causal and alternative methods for estimating the blood pressure more regularly should be sought to further enhance the performance of the data fusion model. However, regardless of the type of blood pressure values used, the multi-variate novelty in this model is mostly driven by decreases in blood

pressure despite blood pressure being one of four vital sign parameters included in the model.

7.2 Future Work

Two main areas can be identified to further the work discussed in this thesis with a view to improving the specificity of the patient-specific data fusion model. Firstly, alternatives to the current blood pressure measurements should be considered in order to either obtain a continuous or semi-continuous measure of blood pressure. Secondly, the performance of the data fusion model could possibly be improved by adding more parameters that have been shown to be relevant for IDH and potential parameters are considered below. These two points are considered in more detail in the following sections.

7.2.1 Continuous Blood Pressure Monitoring

While the patient-specific data fusion model presented in this thesis can identify IDH events, its ability to predict them is limited, particularly in the absence of information on the trends in blood pressure between actual measurements. This is partly due to the low sampling frequency of the blood pressure in our study. In this regard, it would be advantageous to use either more frequent, or continuous blood pressure measurements. Blood pressure can be measured continuously by invasive means, using an intra-arterial catheter. However, these are normally only used for critically ill patients and are not suitable for application to dialysis patients for practical reasons. Firstly, this would require routine insertion of a catheter in a patient's artery, possibly on the contralateral arm to the dialysis access point,

which would restrict movement in both of the patient's arms for the entire session. More importantly however, there are risks associated with catheter insertion that outweigh the benefits of having a continuous blood pressure signal.

Alternative non-invasive devices for continuous blood pressure measurements such as the Finometer (Finapres Medical Systems, Amsterdam, The Netherlands) exist. This monitor is based on the principle of vascular unloading, which was first described by Penaz [95]. An inflatable cuff is mounted on the finger to occlude blood flow through the arteries, such that the arterial and cuff pressures are equal at all times [163]. The diameter of the cuff is constantly adjusted for this purpose using a plethysmograph incorporated in the cuff, which is used to determine the total finger volume under the cuff. The accuracy of this monitor has been assessed in numerous studies [164, 165, 166], and given its non-invasive nature it would be a suitable device for our purposes. However, this would require the acquisition of further equipment and, given the current pricing of the monitor at several thousands of pounds, would be unsuitable for large-scale application at this point.

The PTT, that is the time it takes for a pulse wave to travel from the aortic valve to a peripheral site, could also be used as a continuous estimate of blood pressure, as described in more detail in Section 1.6.1. Using the ECG and PPG signals which have both been measured as part of our study, the PTT can be measured as the time between the R wave of the ECG, which corresponds to the opening of the aortic valve, and the peak in the PPG signal which represents the arrival of the pulse wave at the finger. While Payne et al [70] report that the PTT is not reliable enough for use as a blood pressure surrogate, they suggest that it can be used as a measure of variations in blood pressure, and it would be useful to test this with our data.

7.2.2 Additional Parameters

In an attempt to improve the specificity of the patient-specific data fusion model introduced in this thesis, further parameters could be added. In the following, some of the candidate parameters are introduced and their relevance discussed.

7.2.2.1 Peripheral Perfusion

While blood pressure is the primary indicator for IDH, it is more generally used to assess the overall stability of the blood circulation. However, it does not account for changes in microcirculation at peripheral parts of the body [167]. As previously described in Section 1.6.3, blood flow is redistributed from the peripheral to the core circulation which supplies the vital organs as a compensatory response to decreases in blood volume [82]. This results in a decrease in peripheral perfusion. Given that a reduction in blood volume due to UF is known to be a causative factor for IDH [168], peripheral perfusion, as an index for blood volume change, would be a useful parameter for inclusion in the data fusion model. During HD, blood volume changes due to UF have been shown to reduce the blood supply to the microcirculation below the tongue [169]. Several measures have been suggested for assessing peripheral perfusion, including the central-to-toe temperature gradient [170], Laser Doppler flow [83] and the ratio between the pulsatile and non-pulsatile components of the PPG waveform [81]. In the latter computation, note should be taken of the automatic gain compensation for the PPG waveform that is commonly included as part of commercially available photoplethysmographs. As the amplitude of the PPG signal decreases, for instance as a result of low peripheral perfusion, adjustments are made to the signal by gain compensation [40].

Therefore, accurate measurement of peripheral perfusion using this method will only be possible if this effect is accounted for.

7.2.2.2 Phase Rectified Signal Averaging

While the frequency spectrum of RR intervals gives an indication as to the frequencies at which the timings between successive beats are modulated, it does not contain any information as to whether the RR interval is increased of decreased at any given modulation. Phase Rectified Signal Average (PRSA) is a method for separately identifying increases and decreases in the RR interval, also termed accelerations and decelerations of the heart rate in this context [171]. By measuring accelerations and decelerations separately, the PRSA allows the effects of the sympathetic and parasympathetic branches of the autonomic nervous system to be assessed separately. The key advantage of this technique is that it does not take into account non-periodic components in the signal such as artefacts and noise, thus making it particularly robust. Although it has not yet been assessed in HD patients to the best of our knowledge, evidence from assessments in general patient populations indicates that the measure has predictive value with respect to myocardial infarction [171], and based on the relevance of other HRV methods, it is likely that PRSA would be a suitable parameter for inclusion in the data fusion model.

7.2.2.3 Other Measures

Finally, both the Haematocrit and the HRV measure introduced in Chapter 3, could be included in the model, although the limitations of these parameters should be acknowledged at this point. As previously demonstrated, the HRV measure intro-

duced in this thesis appears to be useful for predicting IDH only in patients with autonomic dysfunction. This would have to be accounted for in the model. With regard to the Haematocrit, evidence for a significant correlation with IDH events is lacking [63, 60], although this parameter is routinely used by nursing staff to estimate blood volume and to achieve an optimal balance between the patient's refill capacity and the UF rate in an attempt to avoid IDH. This is consistent with evidence that changes in the Haematocrit trajectory are predictive of IDH [78, 76], and it is therefore suggested that instead of absolute haematocrit values, a model of the haematocrit trajectory in conjunction with the corresponding UF rates might be used.

Appendix A

Dataset Overview

An overview of the sex and age of each of the 77 patients of the study is given in Table A.1, alongside the total number of sessions recorded and the proportion of IDH sessions. A session was classified as an IDH session if the systolic blood pressure satisfied the definition given in Section 1.4.

OB	Age	Sex	Number of Sessions	Number (%) of IDH sessions
1	68	m	26	12 (46%)
2	81	m	22	3 (14%)
3	62	m	8	3 (38%)
4	82	f	3	3 (100%)
5	73	m	3	3 (100%)
6	49	m	8	1 (13%)
7	68	m	31	12 (39%)
8	77	m	25	2 (8%)
9	75	m	5	1 (20%)

Continued on next page

Continued from previous page

OB	Age	Sex	Number of Sessions	Number (%) of IDH sessions
10	61	m	21	14 (67%)
11	56	m	4	0 (0%)
12	73	m	13	5 (38%)
13	79	f	5	0 (0%)
14	45	m	16	3 (19%)
15	66	f	7	3 (43%)
17	80	f	27	1 (4%)
18	59	m	5	1 (20%)
19	63	m	15	5 (33%)
20	55	m	28	16 (57%)
21	67	m	18	4 (22%)
22	72	m	11	1 (9%)
23	54	m	25	2 (8%)
24	72	m	3	2 (66%)
25	77	m	3	2 (66%)
26	43	m	5	2 (40%)
27	21	m	6	1 (17%)
28	62	m	8	2 (25%)
29	70	m	4	0 (0%)
30	35	m	26	0 (0%)
31	30	m	2	1 (50%)
32	41	f	4	0 (0%)

Continued on next page

Continued from previous page

OB	Age	Sex	Number of Sessions	Number (%) of IDH sessions
33	84	m	4	0 (0%)
34	90	m	4	4 (100%)
35	33	m	12	11 (92%)
36	63	m	10	3 (30%)
37	77	f	5	2 (40%)
38	70	f	4	4 (100%)
39	79	m	5	1 (20%)
40	79	m	4	2 (50%)
41	29	f	4	0 (0%)
42	77	f	1	1 (100%)
43	38	m	4	1 (25%)
44	49	m	3	0 (0%)
45	68	f	5	1 (20%)
46	37	m	5	0 (0%)
47	70	m	2	0 (0%)
48	51	m	3	0 (0%)
49	63	m	4	1 (25%)
50	56	f	4	0 (0%)
51	24	m	3	0 (0%)
52	37	m	3	0 (0%)
53	47	m	4	0 (0%)
54	46	m	4	0 (0%)

Continued on next page

Continued from previous page

OB	Age	Sex	Number of Sessions	Number (%) of IDH sessions
55	74	f	2	0 (0%)
56	52	m	3	2 (66%)
57	85	f	7	1 (14%)
58	77	m	1	0 (0%)
59	79	f	4	1 (25%)
60	60	f	5	2 (40%)
61	75	m	4	0 (0%)
62	73	m	4	1 (25%)
63	50	m	12	1 (8%)
64	62	m	7	0 (0%)
65	74	f	1	0 (0%)
66	45	f	9	3 (33%)
67	65	f	4	0 (0%)
68	55	m	14	8 (57%)
69	76	m	4	0 (0%)
70	49	f	2	0 (0%)
71	79	m	2	2 (100%)
73	64	f	4	2 (50%)
74	77	m	2	0 (0%)
75	53	m	5	0 (0%)
76	85	f	3	0 (0%)
77	51	f	4	0 (0%)

Continued on next page

Continued from previous page

OB	Age	Sex	Number of Sessions	Number (%) of IDH sessions
78	75	m	12	0 (0%)
79	69	m	4	2 (50%)

Table A.1 – Overview of the number of IDH sessions and total number of sessions recorded for each patient in the study. Further, the age and sex of each patient is given.

Appendix B

Further Neuroscale Examples

B.1 Examples of IDH and non-IDH session Neuroscale Projections

Figure B.1 show six further examples of vital sign data from non-IDH sessions projected onto the population-trained Neuroscale map (described in detail in Chapter 4).

B.2 Training Data for Patients OB7 and OB30.

Figure B.4 shows the time series plots of the training data for the patient-specific Neuroscale maps of patients OB7 and OB30, as discussed in Section 5.2.

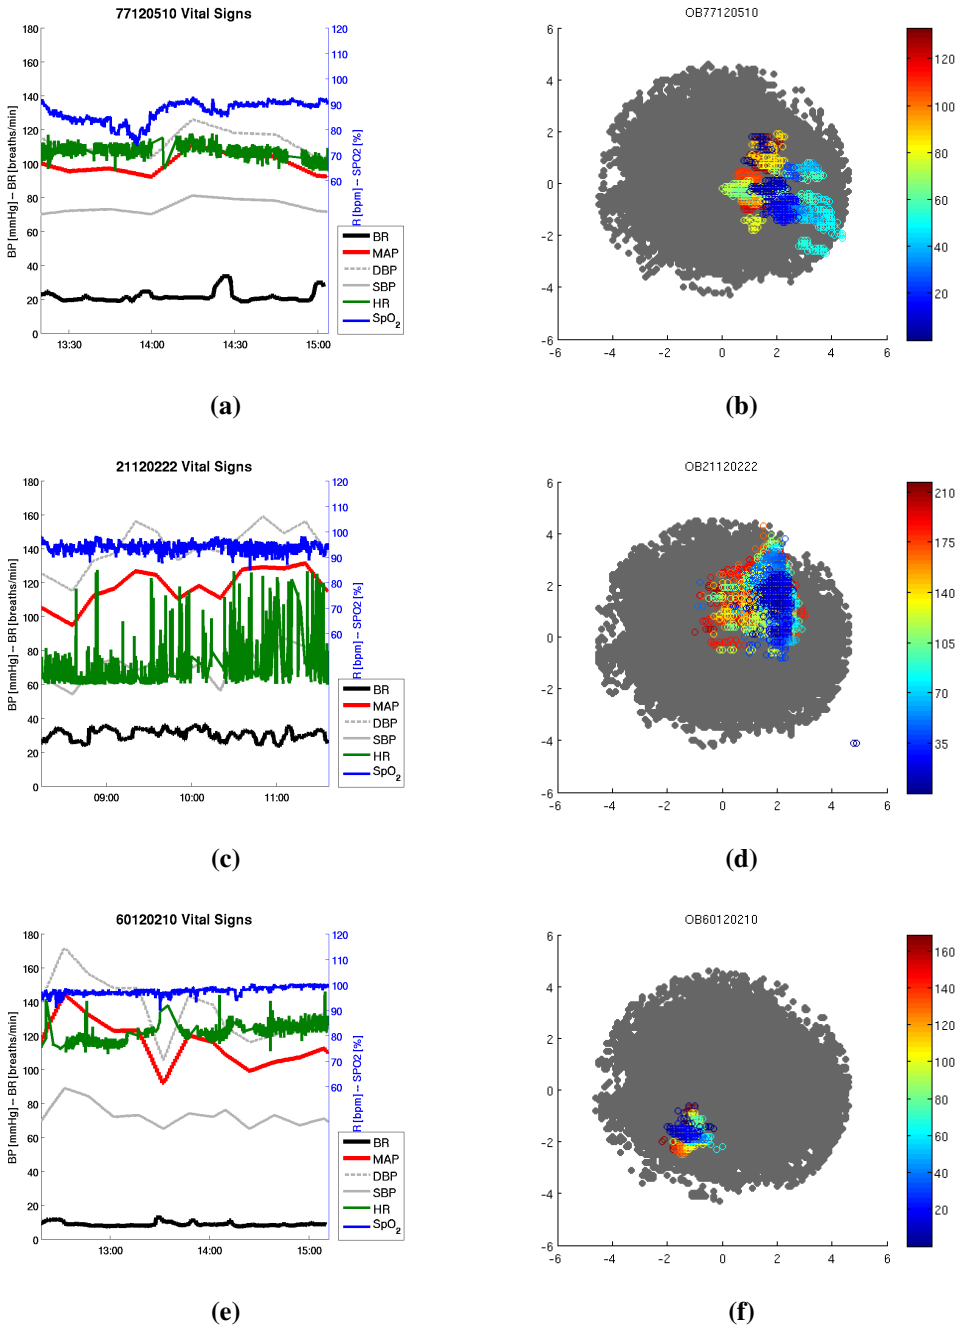


Figure B.1 – *Continued on next page*

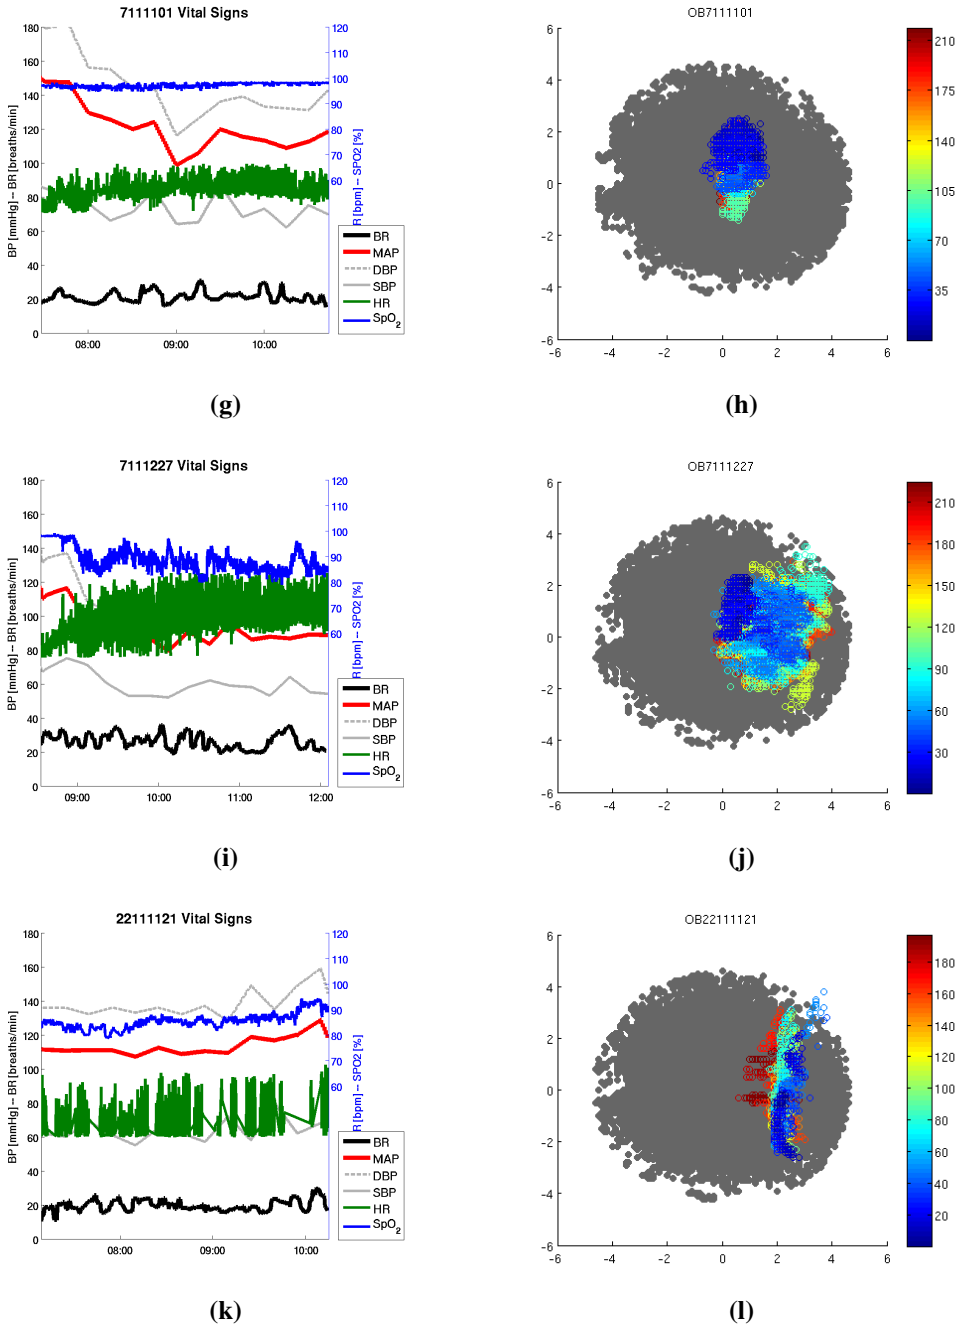
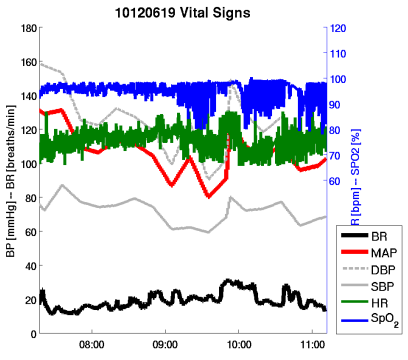
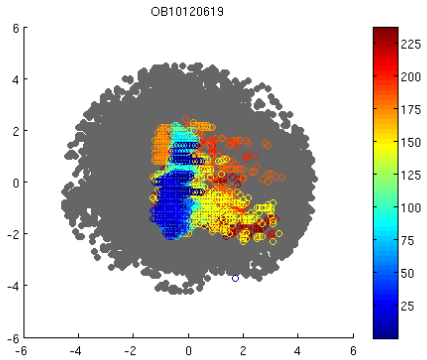


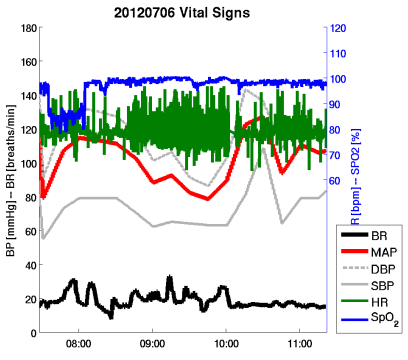
Figure B.1 – Examples of non-IDH sessions projected onto the population-trained Neuroscale map. All vital sign data from these sessions is projected within the boundaries of the region defined by the training data.



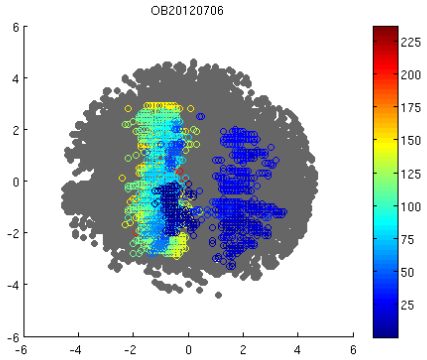
(a)



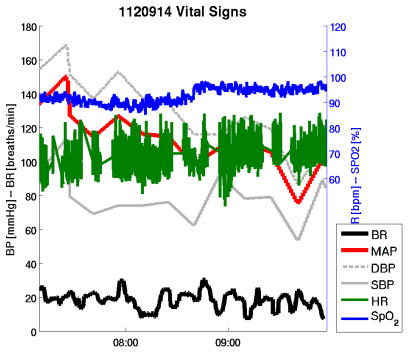
(b)



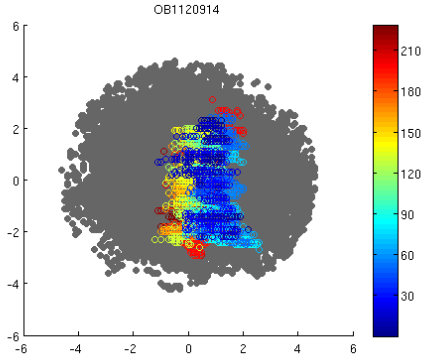
(c)



(d)



(e)



(f)

Figure B.2 – *Continued on next page*

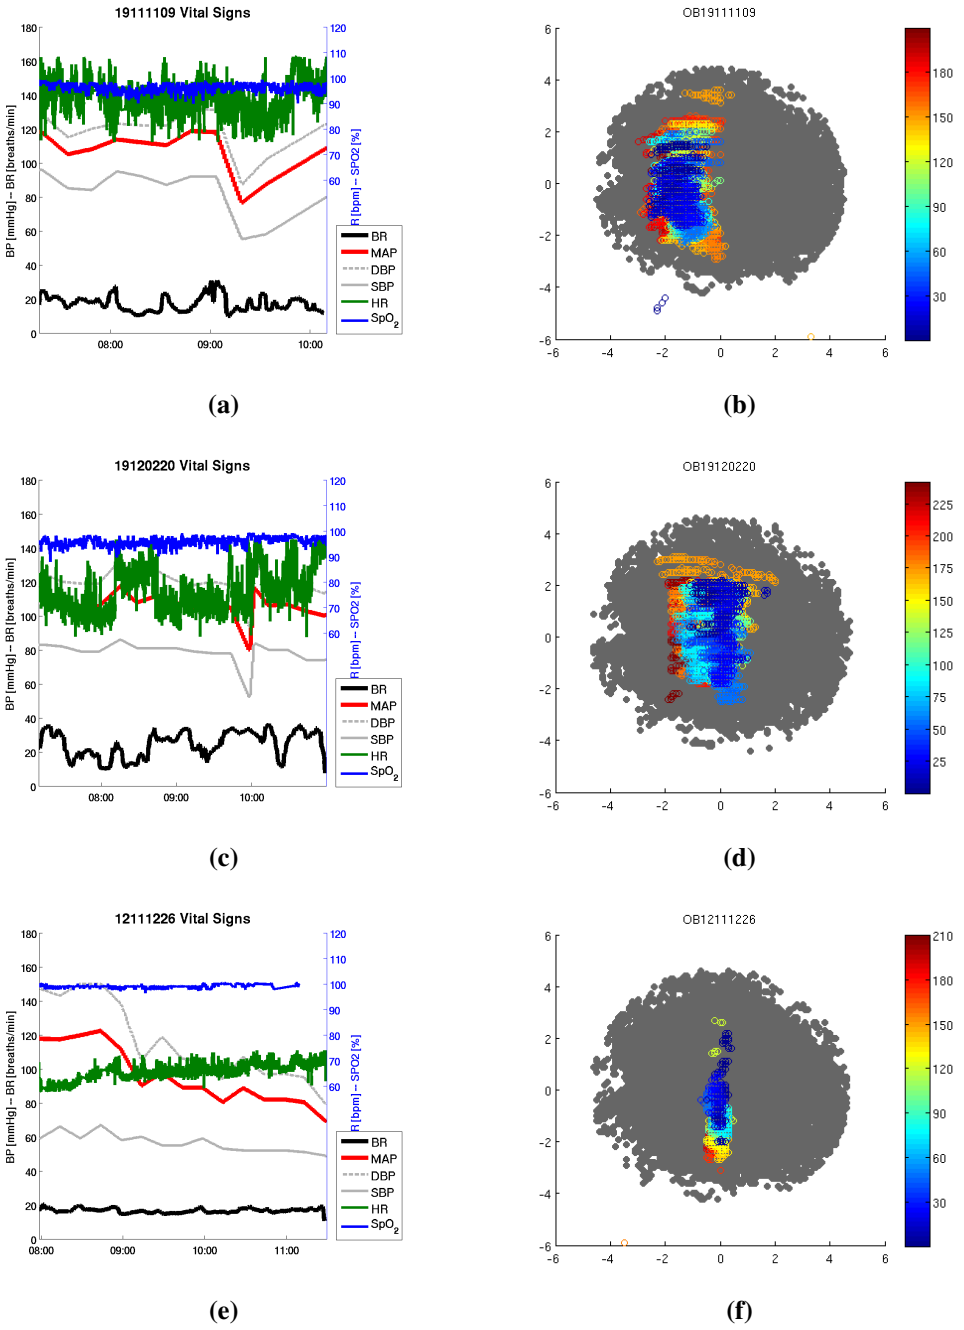


Figure B.3 – Examples of vital sign data from IDH sessions projected onto the population-trained Neuroscale map. Although these sessions contain IDH episodes, the vital signs are projected within the boundaries of the map.

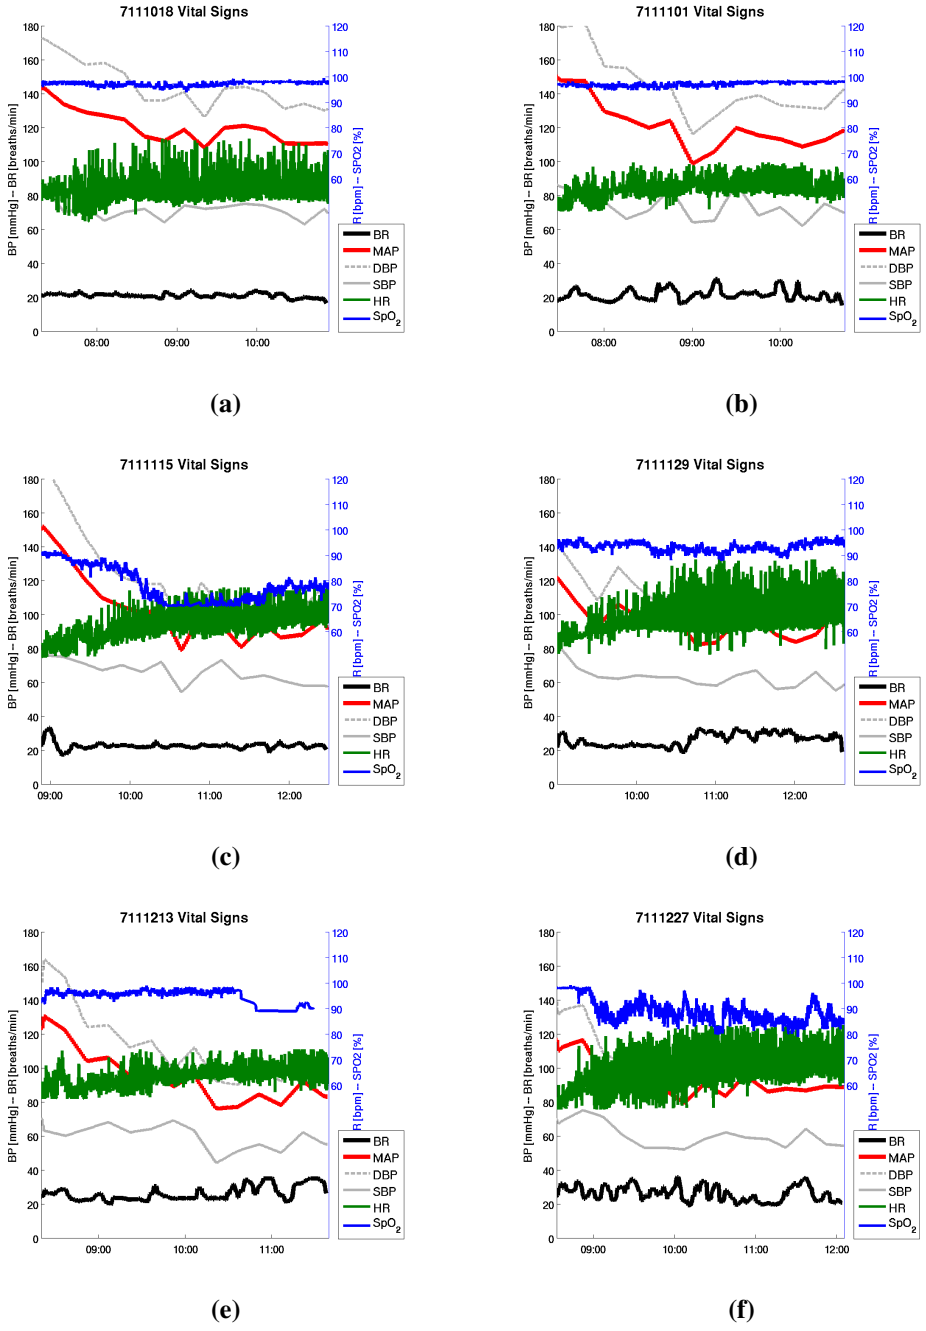
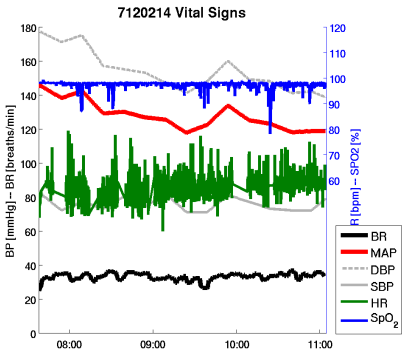
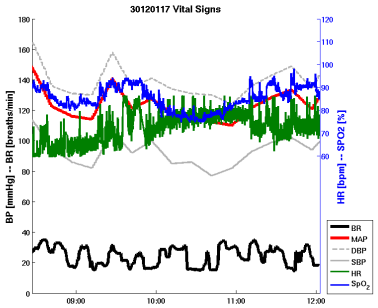


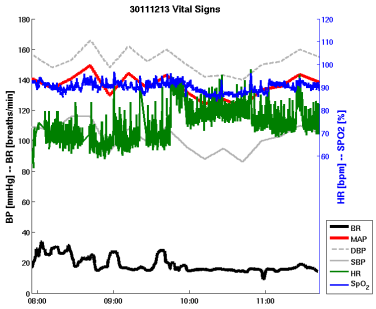
Figure B.4 – *Continued on next page*



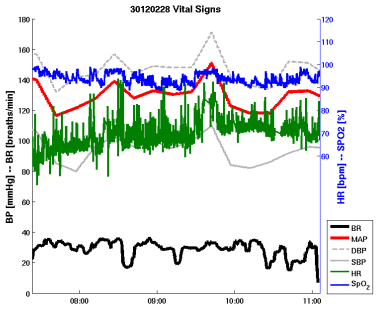
(g)



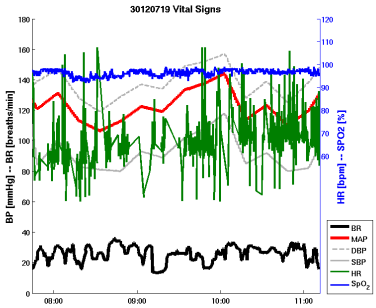
(h)



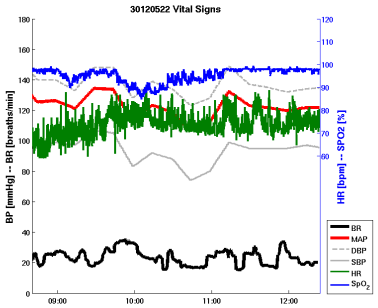
(i)



(j)

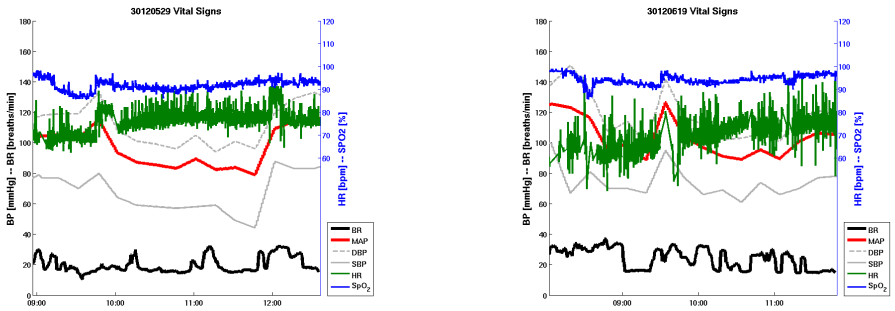


(k)



(l)

Figure B.4 – *Continued on next page*



(m) **(n)**

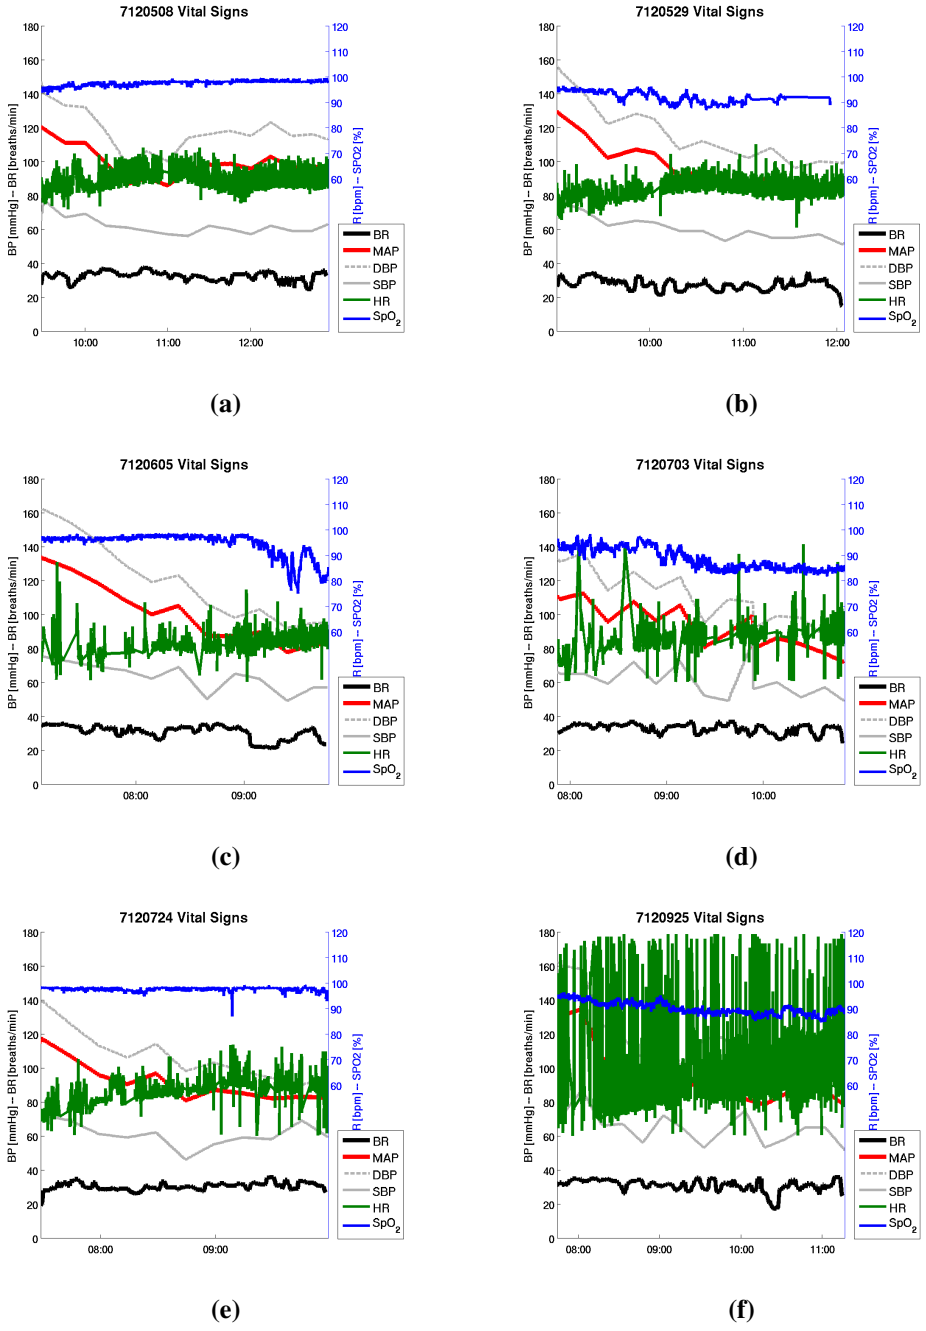
Figure B.4 – Time series plots of the vital sign data used to train the patient-specific Neuroscale maps for patients OB7 (a-g) and OB30 (h-n).

B.3 Test Data for Patients OB7 and OB30

The time series plots for the seven non-IDH sessions projected onto the patient-specific Neuroscale map for patient OB7 (as shown in Section 5.3) are shown in Figure B.5, together with the time series plots of the seven non-IDH sessions for patient OB30.

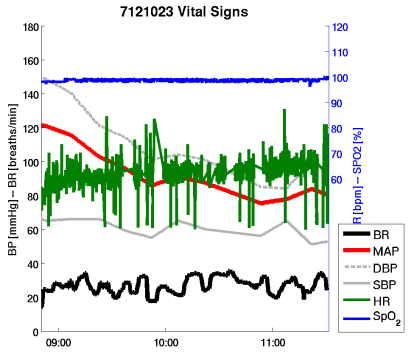
B.4 Examples of Patient-Specific Neuroscale Projections

Further examples of projections of vital sign data from IDH and non-IDH sessions onto patient-specific Neuroscale maps (as described in Chapter 5) are shown in Figures B.6 and B.7. Figure B.6 shows two further examples of IDH sessions from patients OB1 and OB14, respectively, while Fig. B.7 shows four further

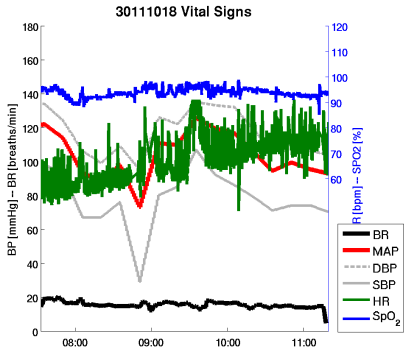


(a) **(b)** **(c)** **(d)** **(e)** **(f)**

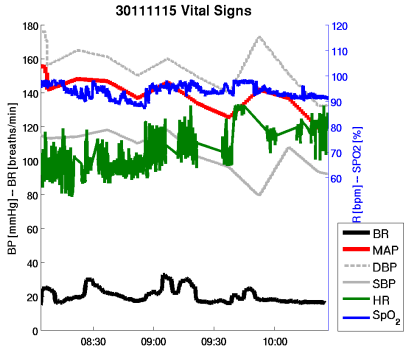
Figure B.5 – *Continued on next page*



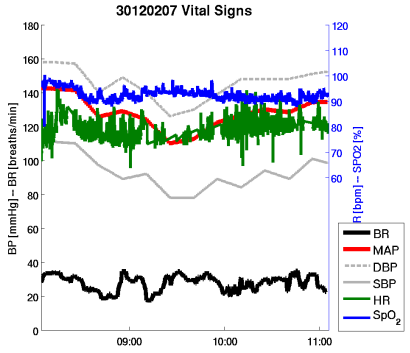
(g)



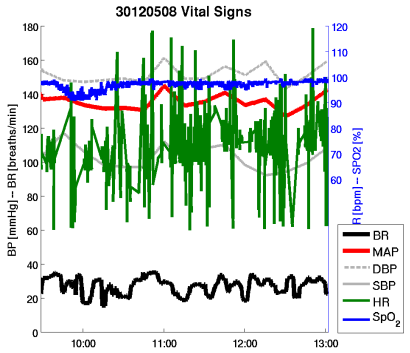
(h)



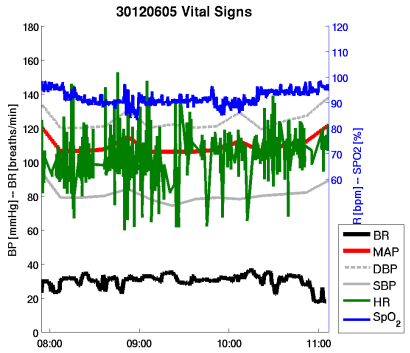
(i)



(j)

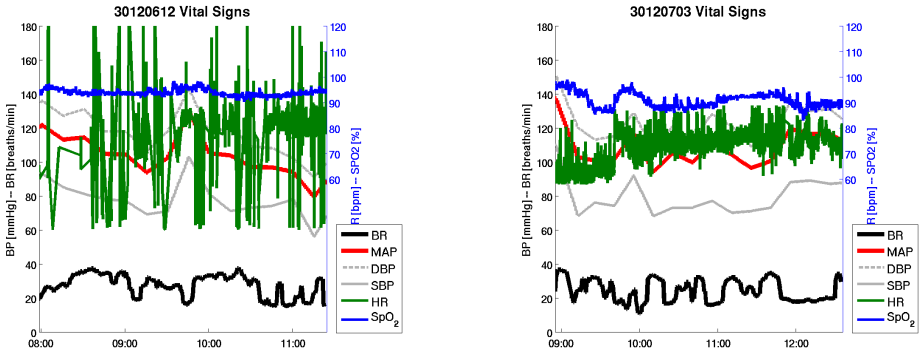


(k)



(l)

Figure B.5 – *Continued on next page*



(m) **(n)**

Figure B.5 – Time series plots of the vital sign data from seven non-IDH sessions each for patients OB7 (a-g) and OB30 (h-n).

examples of non-IDH sessions from patients OB1, OB12 and OB17.

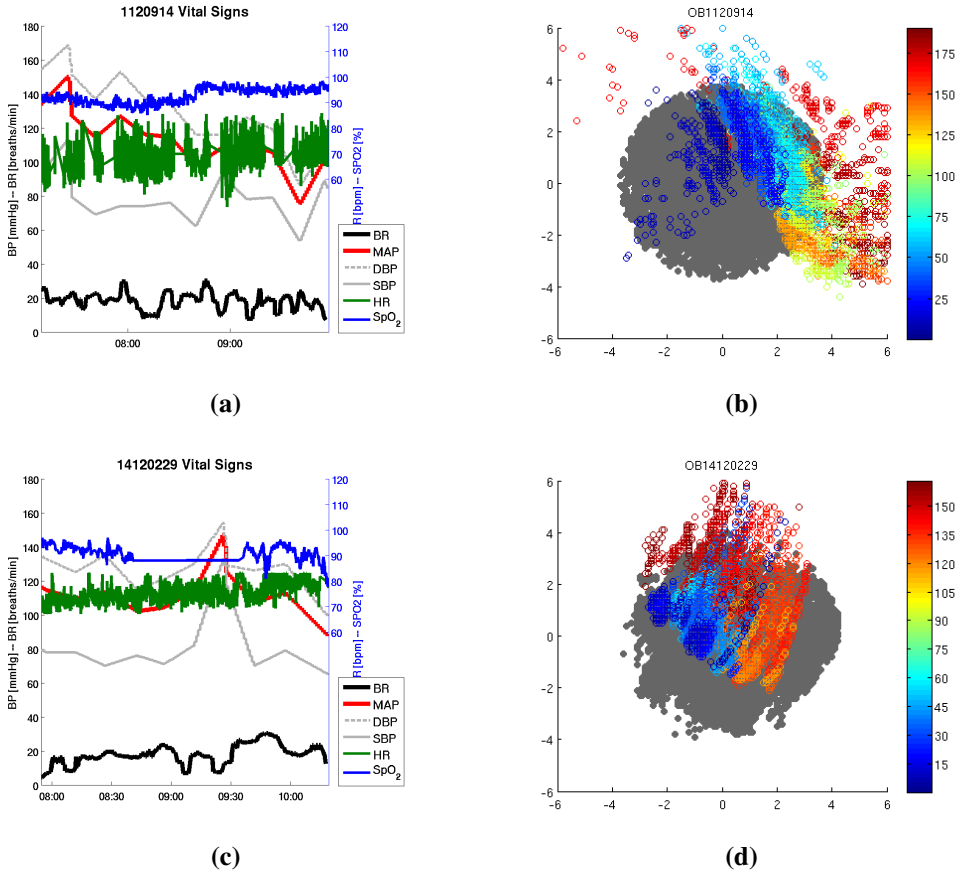


Figure B.6 – Examples of vital sign data from IDH sessions projected onto patient-specific Neuroscale maps for patients OB1 and OB14. The vital signs from these sessions are projected outside of the region defined as normal on the Neuroscale map.

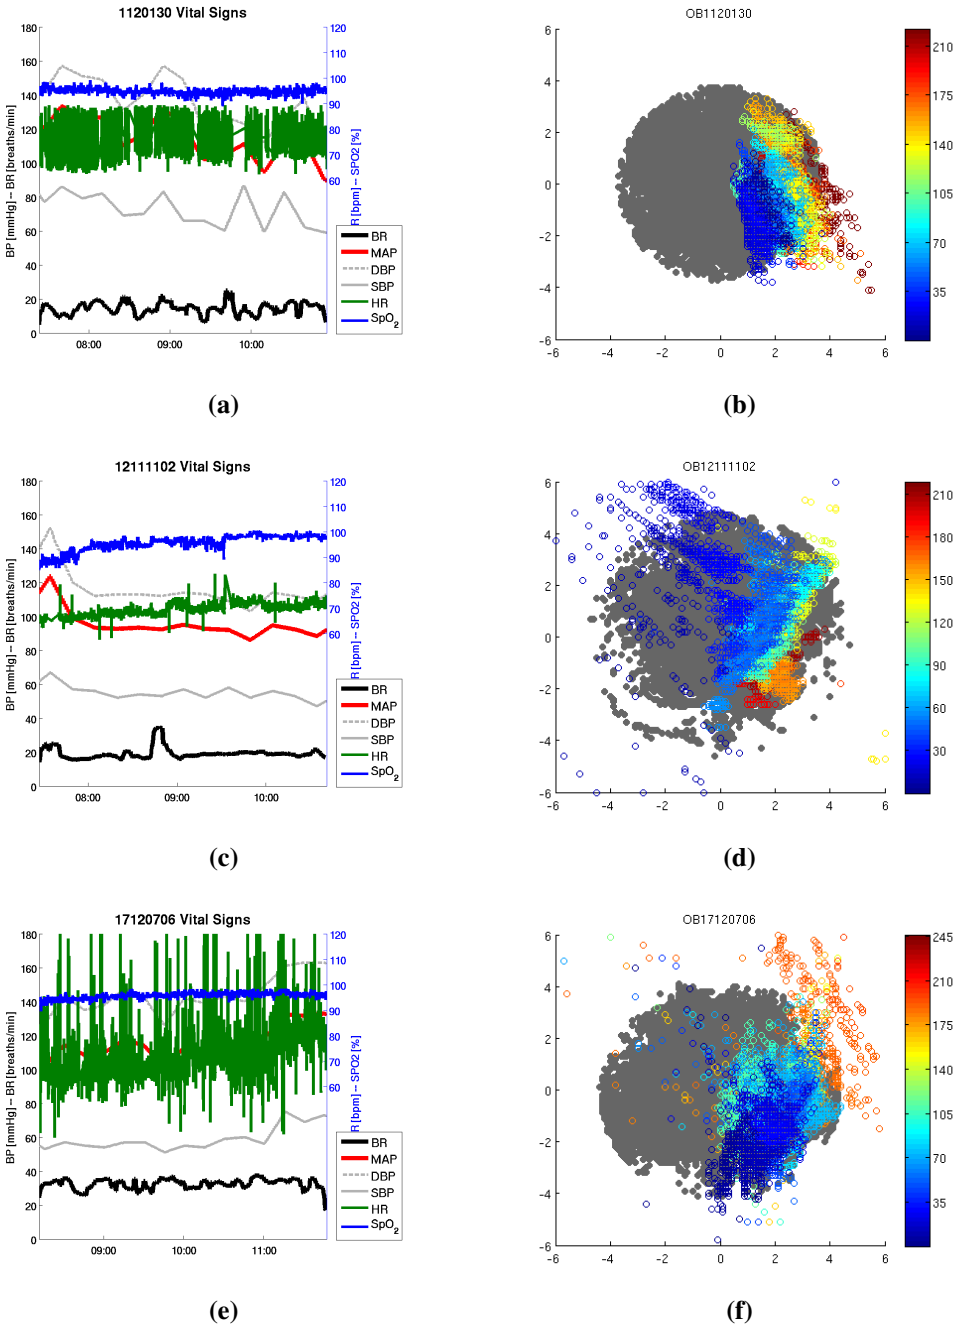
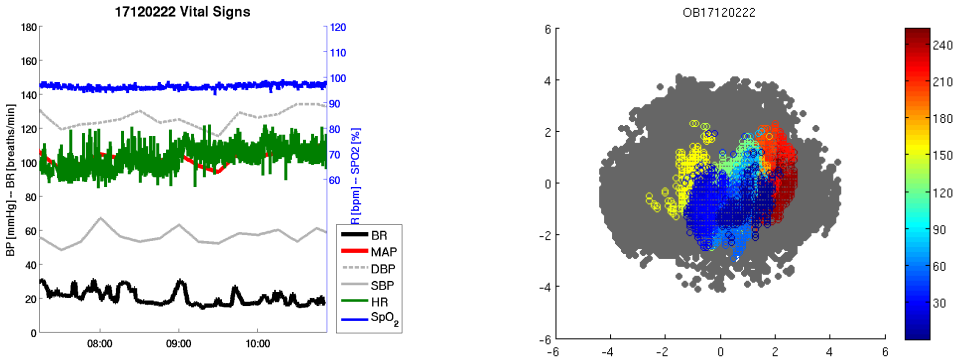


Figure B.7 – *Continued on next page*



(g) **(h)**

Figure B.7 – Examples of vital sign data from non-IDH sessions projected onto patient-specific Neuroscale maps. With the exception of outliers, the vital signs are projected onto the normal region of each patient-specific map.

Bibliography

[1] L. W. Henderson, "Symptomatic intradialytic hypotension and mortality: An opinionated review," *Seminars in Dialysis*, vol. 25, no. 3, pp. 320–25, 2012.

[2] A. Guyton and J. Hall, *Textbook of Medical Physiology, 10th Ed.* W.B. Saunders Company, 2000.

[3] J. Daugirdas, "Dialysis hypotension. a hemodynamic analysis," *Kidney Int*, vol. 39, pp. 233–46, 1991.

[4] L. Henderson, "Symptomatic hypotension during haemodialysis," *Kidney Int*, vol. 17, pp. 571–576, 1980.

[5] W. Bos, S. Bruin, R. W. van Olden, I. Keur, K. H. Wesseling, N. Westerhof, R. T. Krediet, and L. A. Arisz, "Cardiac and hemodynamic effects of hemodialysis and ultrafiltration," *Am J Kidney Dis*, vol. 35, pp. 819–826, 2000.

[6] B. F. Palmer and W. L. Henrich, "Recent advances in the prevention and management of intradialytic hypotension," *J Am Soc Nephr*, vol. 19, no. 1, pp. 8–11, 2008.

[7] P. Zager, J. Nikolic, R. H. Brown, M. A. Campbell, W. C. Hunt, D. Peterson, J. V. Stone, A. Levey, K. B. Meyer, M. J. Klag, H. K. Johnson, E. Clark, J. H. Sadler, and P. Teredesai, "'u' curve association of blood pressure and mortality in hemodialysis patients," *Kidney Int*, vol. 54, pp. 561–569, 1998.

[8] P. Owen, W. Priestman, M. Sigrist, S. Lambie, S. John, L. Chesterton, and C. McIntyre, "Myocardial contractile function and intradialytic hypotension," *Hemodial Int*, vol. 13, pp. 293–300, 2009.

[9] J. Dasselaar, R. Slart, M. Knip, J. Prium, R. Tio, C. McIntyre, P. de Jong, and C. Franssen, "Haemodialysis is associated with a pronounced fall in myocardial perfusion," *Nephrol Dial Transplant*, vol. 24, no. 2, pp. 604–610, 2009.

[10] J. Burton, H. Jefferies, N. Selby, and C. McIntyre, "Hemodialysis-induced repetitive myocardial injury results in global and segmental reduction in systolic cardiac function," *Clin J Am Soc Nephrol*, vol. 4, no. 12, pp. 1925–1931, 2009.

[11] T. Breidthardt and C. McIntyre, "Dialysis-induced myocardial stunning: the other side of the cardiorenal syndrome," *Rev Cardiocasc Med*, vol. 12, pp. 13–20, 2011.

[12] G. Manjunath, A. Levey, and M. Sarnak, "How can the cardiac heath rate be reduced in dialysis patients?," *Semin Dialysis*, vol. 1, no. 15, pp. 18–29, 2002.

[13] www.renalreg.com/Reports/2011.html, "Uk renal registry - the fourteenth annual report," 2011.

[14] F. Locatelli, L. D. Vecchio, and C. Manzoni, "Morbidity and Mortality on Maintenance Haemodialysis.," *Nephron*, vol. 80, pp. 380–400, 1998.

[15] F. Brunner and N. Selwood, "Profile of patients on RRT in Europe and death rates due to major causes of death groups," *Kidney Int*, vol. 42, no. 38, pp. S4–S15, 1992.

[16] U. R. D. System, *USRDS 2006 annual data report: Atlas of End-Stage Renal Disease in the United States*. Bethesda, MD: National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Disease, 2006.

[17] C. Herzog, "Can we prevent sudden cardiac death in dialysis patients?," *Clin J Am Soc Nephrol*, no. 2, pp. 410–411, 2007.

[18] M. Cogan and P. Schoenfeld, *Introduction to Dialysis, 2nd Ed.* Churchill Livingstone, 1991.

[19] L. Mailloux and W. E. Haley, "Hypertension in the esrd patient: pathophysiology, therapy, outcomes, and future directions," *Am J Kid Dis*, vol. 32, no. 5, pp. 705–719, 1998.

[20] W. Purcell, E. Manias, A. Williams, and R. Walker, "Accurate dry weight assessment: Reducing the incidence of hypertension and cardiac disease in patients on haemodialysis," *Nephrology Nursing Journal*, vol. 31, no. 6, pp. 631–645, 2004.

[21] A. Thornley, "A prospective study exploring the use of the crit-line in acute dialysis therapies," *School of Health and Social Care, Oxford Brookes University*, 2006.

[22] J. Jaeger and R. Mehta, "Assessment of dry weight in hemodialysis: an overview," *J Am Soc Nephrol*, vol. 10, pp. 392–403, 1999.

[23] B. Charra, G. Laurent, C. Chazot, E. Calemard, J.-C. Terrat, T. Vanel, G. Jean, and M. Ruffet, "Clinical assessment of dry weight," *Nephrol Dial Transplant*, vol. 11, pp. 16–19, 1996.

[24] J. Cheigh, C. Milite, J. Sullivan, A. Rubin, and K. Stenzel, "Hypertension is not adequately controlled in hemodialysis patients," *Am J Kidney Dis*, vol. 19, pp. 453–459, 1992.

[25] G. Tortora and S. Grabowski, *Principles of Anatomy and Physiology*. John Wiley Sons, 2003.

[26] D. Rubinger, D. Sapoznikov, A. Pollak, M. Popovtzer, and M. Luria, "Heart rate variability during chronic hemodialysis and after renal transplantation: studies in patients without and with systemic amyliodosis," *J Am Soc Nephrol*, vol. 10, pp. 1972–1981, 1999.

[27] P. S. Hamilton and W. J. Tompkins, "Quantitative Investigation of QRS Detection Rules using the MIT/BIH Arrhythmia Database," *IEEE Transactions on Biomedical Engineering*, vol. BME-33, no. 12, pp. 1157–1165, 1986.

[28] R. Shouldice, C. Heneghan, P. Nolan, P. Nolan, and W. McNicholas, "Modulating effect of respiration on atrioventricular conduction time assessed using pr interval variation," *Med Biol Eng Comp*, vol. 40, no. 6, pp. 609–617, 2002.

[29] R. Kleiger, J. P. Miller, J. T. Bigger, A. J. Moss, and M. P.-I. R. Group, "Decreased heart rate variability and its association with increased mortality after acute myocardial infarction," *Am J Cardiol*, vol. 28, pp. 256–62, 1987.

[30] M. Malik, R. Xia, O. Odemuyiwa, A. Staunton, J. Poloniecki, and A. Camm, "Influence of the recognition artefact in automatic analysis of long-term electrocardiograms on a time-domain measurement of heart rate variability," *Med & Biol Eng & Comput*, vol. 31, pp. 539–44, 1993.

[31] F. M. Fouad, R. C. Tarazi, C. M. Ferrario, S. Fighaly, and S. Alicandri, "Assessment of parasympatheticcontrol of heart rate by a noninvasive method," *Am J Physiol Heart Circ Physiol*, vol. 246, pp. H838–H842, 1984.

[32] D. J. Ewing, J. M. M. Neilson, and P. Travis, "New method for assessing cardiac parasympathetic activity using 24 hour electrocardiograms," *Br Heart J*, vol. 52, pp. 396–402, 1984.

[33] P. Grossmann, J. van Beeck, and C. Wientjes, "A comparison of three quantification methods for estimation of respiratory sinus arrhythmia," *Psychophysiology*, vol. 27, pp. 702–714, 1990.

[34] G. Pelosi, M. Emdin, C. Carpeggiani, M. Morales, M. Piacenti, P. Dattolo, T. Cerrai, A. Macerata, A. L'Abbate, and Q. Maggiore, "Impaired sympathetic response before intradialytic hypotension: a study based on spectral analysis of heart rate and pressure variability," *Clinical Science*, vol. 96, pp. 23–31, 1999.

[35] S. Akselrod, D. Gordon, F. Ubei, D. Shannon, A. Barger, and R. Cohen, "Power spectrum analysis of heart rate fluctuation: A quantitative probe of beat-to-beat cardiovascular control," *Science*, vol. 213, pp. 220–222, 1981.

[36] E. Toledo, O. Gurevitz, H. Hod, M. Eldar, and S. Akselrod, "Wavelet analysis of instantaneous heart rate: a study of autonomic control during thrombolysis," *Am J Physiol Regul Integr Comp Physiol*, vol. 284, pp. R1079–R1091, 2003.

[37] A. Malliani, M. Pagani, F. Lombardi, and S. Cerutti, "Cardiovascular neural regulation explored in the frequency domain," *Circulation*, vol. 84, pp. 1482–1492, 1991.

[38] R. Berger, S. Akselrod, D. Gordon, and R. Cohen, "An efficient algorithm for spectral analysis of heart rate variability," *IEEE Trans Biomed Eng*, vol. 9, pp. 900–904, 1986.

[39] R. Furlan, S. Guzetti, W. Crivellaro, S. Dassi, M. Tinelli, G. Baselli, S. Cerutti, F. Lombardi, M. Pagiani, and A. Malliani, "Continuous 24-hour assessment of the neural regulation of systemic arterial pressure and rr variabilities in ambulant subjects," *Circulation*, vol. 81, pp. 537–547, 1990.

[40] J. Sinex, "Pulse oximetry: Principles and limitations.," *Am J Emerg Med*, vol. 17, pp. 59–67, 1999.

[41] F. Allan and S. B. Pacela, "Impedance pneumography - a survey of intrumentation techniques," *Medical and biological engineering*, vol. 4, no. 1, pp. 1–15, 1996.

[42] R. Bailon, L. Sornmo, and P. Laguna, "Ecg-derived respiratory frequency estimation," in *Advanced methods and tools for ECG data analysis* (G. D. Clifford, F. Azuaje, and P. McSharry, eds.), pp. 215–244, Engineering in Medicine and Biology, 2006.

[43] V. Novak, P. Novak, J. De-Champlain, A. Le-Blanc, R. Martin, and R. Nadeau, "Influence of respiration on heart rate and blood pressure fluctuations," *Journal of Applied Physiology*, vol. 74, no. 2, pp. 617–626, 1993.

[44] C. Orphanidou, O. Brain, J. Feldmar, S. Khan, J. Price, and L. Tarassenko, "Spectral fusion for estimating respiratory rate from the ecg," *Proc ITAB*, 2009.

[45] G. G. Berntson, J. T. Cacioppo, and K. S. Quigley, "Respiratory sinus arrhythmia: Autonomic origins, physiological mechanisms, and psychophysiological implications," *Psychophysiology*, vol. 30, pp. 183–196, 1993.

[46] A. Sinha, R. Light, and R. Agarwal, "Relative plasma volume monitoring during hemodialysis aids the assessment of dry weight," *Hypertension*, vol. 55, pp. 305–311, 2010.

[47] M. J. Damasiewicz and K. R. Polkinghorne, "Intra-dialytic hypotension and blood volume and blood temperature monitoring," *Nephrology*, vol. 16, pp. 13–18, 2011.

[48] S. Hung, Y. Hung, H. Fang, J. Jeh, G. Hung, C. Wu, K. Chou, and H. Chung, "Cardiac troponin I and creatine kinase isoenzyme MB in patients with intradialytic hypotension," *Blood Purif*, vol. 22, no. 4, pp. 338–343, 2004.

[49] M. G. W. Barnas, W. H. Boer, and H. A. Koomans, "Hemodynamic patterns and spectral analysis of heart rate variability during dialysis hypotension," *J Am Soc Nephrol*, vol. 10, pp. 2577–84, 1999.

[50] R. Schmidt, O. Roeher, H. Hickstein, and S. Korth, "Prevention of haemodialysis-induced hypotension by biofeedback control of ultrafiltration and infusion," *Nephrol Dial Transplant*, vol. 16, no. 3, pp. 595–603, 2001.

[51] T. Mizumasa, H. Hirakata, T. Yoshimitsu, E. Hirakata, M. Kubo, M. Kashiwagi, H. Tanaka, H. Kanai, S. Fujimi, and M. Iida, "Dialysis-related hypotension as a cause of progressive frontal lobe atropy in chronic hemodialysis patients: A 3-year prospective study," *Nephron Clin Pract*, vol. 97, pp. c23–c30, 2004.

[52] H. Hoeben, A. Abu-Alfa, and R. Mahnensmith, "Hemodynamics in patients with intradialytic hypotension treated with cool dialysate or midodrine," *Am J Kid Dis*, vol. 39, no. 1, pp. 102–107, 2002.

[53] D. Poldermans, A. M. I. 't Veld, R. Rambaldi, A. van den Meiracker, M. can den Dorpel, G. Rocchi, E. Boersma, J. Bax, W. Weimar, J. Roelandt, and E. Zietse, "Cardiac evaluation in hypotension-prone and hypotension-resistant hemodialysis patients.," *Kidney Int.*, vol. 56, pp. 1905–1911, 1999.

[54] E. Mancini, E. Mambelli, M. Irpinia, D. Gabrielli, C. Cascone, F. Conte, G. Meneghel, F. Cavatorta, A. Antonelli, G. Villa, G. Triolo, V. Nuzzo, and A. Santoro, "Prevention of dialysis hypotension episodes using fuzzy logic control system," *Nephrol Dial Transplant*, vol. 22, no. 5, pp. 1420–1427, 2007.

[55] Y. Cai, A. Zimmerman, S. Ladefoged, and N. Secher, "Can haemodialysis-induced hypotension be predicted?," *Nephron*, vol. 92, no. 3, pp. 582–588, 2002.

[56] A. Santoro, E. Mancini, C. Basile, L. Amoroso, S. D. Giulio, M. Usberti, G. Colasanti, G. Verzetti, A. Rocco, E. Imbasciati, G. Panzetta, R. Bolzani, F. Grandi, and M. Polacchini, "Blood volume controlled hemodialysis in hypotension-prone patients: A randomized, multicenter controlled trial," *Kidney Int*, vol. 62, pp. 1034–1045, 2002.

[57] D. J. Meredith, C. W. Pugh, S. Sutherland, L. Tarassenko, and J. Birks, "The relationship between symptoms and blood pressure during maintenance hemodialysis." In review.

[58] B. Straver, P. D. Vries, B. T. Voorde, M. Roggenkamp, A. Donker, and P. T. Wee, "Intradialytic hypotension in relation to pre-existent autonomic dysfunction in hemodialysis patients," *Int J Artif Organs*, vol. 21, pp. 794–801, 1998.

[59] J. Daugirdas, "Pathophysiology of dialysis: an update," *Am J Kidney Dis*, vol. 38[4 Suppl 4], pp. S11–S17, 2001.

[60] H. Krepel, R. W. Nette, E. Akcahuseyin, W. Weimar, and R. Zietse, "Variability of relative blood volume during hemodialysis," *Nephrol Dial Transplant*, vol. 15, pp. 673–679, 2000.

[61] M. Rodriguez, J. Pederson, and F. Llach, "Effect of dialysis and ultrafiltration on osmolarity, colloid osmotic pressure and vascular refilling rate," *Kidney Int*, vol. 28, pp. 808–13, 1985.

[62] J. Rouby, J. Rottembourg, J. P. Durande, J. Y. Basset, and M. Legrain, "Importance of the plasma refilling rate in the genesis of hypovolaemic hypotension during regular dialysis and controlled sequential ultrafiltration-haemodialysis," *Proc Eur Dial Transplant*, vol. 15, pp. 239–44, 1978.

[63] K. Maeda, H. Morita, T. Shinzato, B. V. Vega, H. Kobayakawa, T. Ishihara, H. Inagaki, I. Igarashi, and T. Kitano, "Role of hypovolaemia in dialysis-induced hypotension," *Artif Organs*, vol. 12, pp. 116–21, 1988.

[64] R. L. C. Jr, T. Jacobsen, C. Jost, R. D. Toto, P. Grayburn, T. M. Obregon, R. Fouad-Tarazi, and R. Victor, "Paradoxical withdrawal of reflex vasoconstriction as a cause of hemodialysis-induced hypotension," *J Clin Invest*, no. 90, pp. 1657–1665, 1992.

[65] M. Heber, A. Lahiri, D. Thompson, and E. Raftery, "Baroreceptor, not left vetricular, dysfunction is the cause of hemodialysis hypotension.," *Clin Nephrol.*, vol. 32, no. 2, 1989.

[66] L. J. Chesterton, N. M. Selby, J. O. Burton, J. Fialova, C. Chan, and C. W. McIntyre, "Categorization of the hemodynamic response to hemodialysis: The importance of baroreflex sensitivity," *Hemodialysis International*, vol. 14, no. 1, pp. 18–28, 2010.

[67] D. Sapoznikov, R. Backenroth, and D. Rubinger, "Baroreflex sensitivity and sympatho-vagal balance during intradialytic hypotensive episodes," *Journal of Hypertension*, vol. 28, no. 2, pp. 314–324, 2010.

[68] A. Santoro, E. Mancini, M. Spogano, M. Rossi, F. Paolini, and P. Zuchelli, "A haemodynamic study of hypotension during hemodialysis-induced hypotension using electrical bioimpedance cardiography," *Nephrol Dial Transplant*, no. 5, pp. 147–153, 1990.

[69] C. W. McIntyre, "Recurrent circulatory stress: The dark side of dialysis," *Semin Dial*, vol. 23, no. 5, pp. 449–51, 2010.

[70] R. A. Payne, C. Symeonides, D. Webb, and S. Maxwell, "Pulse transit time measured from the ECG: an unreliable marker of beat-to-beat blood pressure.," *J Appl Physiol.*, vol. 8, pp. 136–141, 2006.

[71] W. Chen, T. Kobayashi, and S. Ichikawa, "Continuous estimation of systolic blood pressure using the pulse arrival time and intermittent calibration," *Med Biol Eng Comput*, vol. 38, no. 5, pp. 569–574, 2000.

[72] C. Ahlstrom, A. Johansson, F. Uhlin, T. Lanne, and P. Ask, "Noninvasive investigation of blood pressure changes using the pulse wave transit time: a novel approach in the monitoring of hemodialysis patients.," *J Artif Organs*, vol. 8, pp. 192–197, 2005.

[73] J. J. Dasselaar, R. M. Huisman, P. E. de Jong, and C. F. Franssen, "Measurement of relative blood volume changes during haemodialysis: merits and limitations," *Nephrol Dial Transplant*, vol. 20, pp. 2043–9, 2005.

[74] S. Mitra, P. Chamney, R. Greenwood, and K. Farrington, "The relationship between systemic and whole-body hematocrit is not constant during ultrafiltration on hemodialysis," *J Am Soc Nephrol*, vol. 15, pp. 463–9, 2004.

[75] K. E. Kim, M. Neff, B. Cohen, M. Somerstein, J. Chinitz, G. Onesti, and C. Swartz, "Blood volume changes and hypotension during hemodialysis," *Trans Am Soc Artif Intern Organs*, vol. 16, pp. 508–514, 1970.

[76] S. Mitra, P. Chamney, R. Greenwood, and K. Farrington, "Linear Decay of Relative Blood Volume During Ultrafiltration Predicts Hemodynamic Instability.," *Am J Kidney Dis*, vol. 40, no. 3, pp. 556–565, 2002.

[77] M. Tonelli, P. Astephen, P. Andreou, S. Beed, P. Lundrigan, and K. Jindal, "Blood volume monitoring in intermittent hemodialysis for acute renal failure," *Kidney Int*, vol. 62, pp. 1075–80, 2002.

[78] S. Andrulli, S. Colzani, F. Mascian, L. ucchi, L. Stipo, M. Bigi, M. Crepaldi, B. Redaelli, A. Albertazzi, and F. Locatelli, "The role of blood volume reduction in the genesis of intradialytic hypotension," *Am J Kid Dis*, vol. 40, no. 6, pp. 1244–54, 2002.

[79] R. Steuer, J. Leypoldt, A. Cheung, H. Senekjian, and J. Conis, "Reducing Symptoms During Hemodialysis by Continuously Monitoring the Hematocrit.," *Am J Kidney Dis*, vol. 27, no. 4, pp. 525–532, 1996.

[80] E. Mancini, L. Corazza, D. Cannarile, M. Soverini, S. Cavalvanti, S. Cavani, A. Fiorenzi, and A. Santoro, "Short Term Variability of Oxygen Saturation during Hemodialysis Is a Warning Parameter for Hypotension Appearance.," *Computers in Cardiology*, vol. 35, pp. 881–883, 2008.

[81] A. Lima, P. Beelen, and J. Bakker, "Use of a peripheral perfusion index derived from the pulse oximetry signal as a noninvasive indicator of perfusion," *Crit Care Med.*, vol. 30, no. 6, pp. 1210–1213, 2002.

[82] R. Bond, "A review of the skin and muscle hemodynamics during hemorrhagic hypotension and shock.," *Adv Shock Res*, vol. 8, pp. 53–70, 1982.

[83] P. Carpentier, "New techniques for clinical assessment of the peripheral microcirculation.," *Drugs*, vol. 59, pp. 17–22, 1999.

[84] H. Fukuta, J. Hayano, S. Ishihara, S. Sakata, S. Mukai, N. Ohte, K. Ojika, K. Yagi, H. Matsumoto, S. Sohmiya, and G. Kimura, "Prognostic value of heart rate variability in patients with end-stage renal disease on chronic haemodialysis," *Nephrol Dial Transplant*, vol. 18, no. 2, pp. 318–325, 2003.

[85] G. Clifford and L. Tarassenko, "Segmenting cardiac-related data using sleep stages increases separation between normal subjects and apnoeic patients," *IPO Physiol Meas*, vol. 25, pp. N27–N35, 2004.

[86] M. Yoshii, J. Minami, T. Ishimitsu, K. Yamakoshi, and H. Matsuoka, "Noninvasice monitoring of hemodynamic changes during hemodialysis by the use of a newly developed admittance cardiograph," *Therapeutic Apheresis and Dialysis*, vol. 9, no. 2, pp. 154–160, 2005.

[87] G. Vita, G. Bellinghieri, A. Trusso, G. Costantino, D. Santoro, F. Monteleone, C. Messina, and V. Savica, "Uremic autonomic neuropathy studied by spectral analysis of heart rate," *Kidney Int*, vol. 56, pp. 232–237, 1999.

[88] G. Vita, P. Princi, R. Calabro, A. Toscano, L. Manna, and C. Messina, "Cardiovascular reflex tests: Assessment of age-adjusted normal range," *J Neurol Sci*, vol. 75, no. 3, pp. 263–74, 1986.

[89] H. Gunderson and B. Neubauer, "Long term diabetic autonomic nerve abnormality," *Diabetologica*, vol. 13, pp. 137–140, 1977.

[90] M. Giordano, D. Manzella, G. Paolisso, A. Caliendo, M. Varricchio, and C. Giordano, "Differences in heart rate variability parameters during the post-dialytic period in type ii diabetic and non-diabetic esrd patients," *Nephrol Dial Transplant*, vol. 16, no. 3, pp. 566–573, 2001.

[91] E. Zitt, U. Neyer, E. Meusburger, M. Tiefenthaler, P. Kotanko, G. Mayer, and A. R. Rosenkranz, "Effect of dialysate temperature and diabetes on autonomic cardiovascular regulation during hemodialysis," *Kidney Blood Press Res*, vol. 31, pp. 217–225, 2008.

[92] D. Rubinger, N. Revis, A. Pollak, M. Luria, and D. Sapoznikov, "Predictors of haemodynamic instability and heart rate variability during haemodialysis," *Nephrol Dial Transplant*, vol. 19, no. 8, pp. 2053–2060, 2004.

[93] S. Cavalcanti, S. Severi, and G. Enzmann, "Analysis of oscillatory components of short-term heart rate variability in hemodynamically stable and unstable patients during hemodialysis.," *Artificial Organs*, vol. 22, no. 2, pp. 98–106, 1998.

[94] K. Wesseling, B. De Wit, G. Van der Hoeven, J. Van Goudoever, and J. Settels, "Physiocal, calibrating finger vascular physiology for finapres," *Homeostasis*, vol. 36, no. 2-3, pp. 67–82, 1995.

[95] J. Penaz, "Photo-electric measurement of blood pressure, volume and flow in the finger," *Digest of the 10th International Conference on Medical and Biological Engineering*, p. 104, 1973.

[96] E. Ino-Oka, J. Urae, M. Sekino, M. Satoh, K. Takeuchi, H. Sekino, Y. Ohtaki, and H. Inooka, "Blood pressure and sympathetic nerve tone relation during hemodialysis may reflect cardiovascular dysfunction," *Intern Med*, vol. 50, pp. 2741–2748, 2011.

[97] J. Kause, G. Smith, D. Prytherch, M. Parr, A. Flabouris, and K. Hillman, "A comparison of antecedents to cardiac arrests, deaths and emergency intensive care admissions in australia and new zealand, and the united kingdom - the academia study," *Resuscitation*, no. 62, pp. 275–282, 2004.

[98] C. Oberli, J. Urzua, C. Saez, M. Guarini, A. Ciprianio, B. Garayar, G. Lema, R. Canessa, C. Sacco, and M. Irarrazaval, "An expert system for monitor alarm integration," *J Clin Monit Comput*, no. 15, pp. 29–35, 1999.

[99] N. Kannathal, U. Acharya, E. Ng, S. M. Krishnan, L. C. Min, and S. Laxminarayan, "Cardiac health diagnosis using data fusion of cardiovascular and haemodynamic signals," *Comput Methods Programs Biomed*, vol. 82, pp. 87–96, 2006.

[100] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[101] L. Tarassenko, A. Hann, and D. Young, "Integrated monitoring and analysis for early warning of patient deterioration.," *British Journal of Anaesthesia*, vol. 97, no. 1, pp. 64–68, 2006.

[102] L. W. Henderson, "The hypo/hypertension quandary," *Artificial Organs*, vol. 35, no. 6, pp. 559–564, 2011.

[103] J. P. Saul, "Beat-to-beat variations of heart rate reflect modulation of cardiac autonomic outflow," *NIPS*, vol. 5, no. 1, pp. 32–37, 1990.

[104] E. Hon and S. T. Lee, "Electronic evaluations of the fetal heart rate patterns preceding fetal death: Further observations," *Am J Obstet Gynecol*, vol. 87, pp. 814–26, 1965.

[105] G. Casolo, E. Balli, T. Taddei, J. Amuhasi, and C. Gori, "Decreased spontaneous heart rate variability on congestive heart failure," *Am J Cardiol*, vol. 64, pp. 1162–67, 1989.

[106] M. Pagani, G. Malfatto, S. Pierini, R. Casati, A. Masu, M. Poli, S. Guzzetti, F. Lombardi, S. Cerutti, and A. Malliani, "Spectral analysis of heart rate variability in the assessment of autonomic diabetic neuropathy," *J Auton Nerv Syst*, vol. 23, pp. 143–53, 1988.

[107] M. T. L. Rovere, G. D. Pinna, R. Maestri, A. Mortara, S. Capomolla, O. Febo, R. Ferrari, M. Franchini, M. Gnemmi, C. Opasich, P. G. Riccardi, E. Traversi, and F. Cobelli, "Short-term heart rate variability strongly predicts sudden cardiac death in chronic heart failure patients," *Circ*, vol. 107, pp. 565–570, 2003.

[108] R. E. Kleiger, P. Stein, and T. B. Jr, "Heart rate variability: Measurement and clinical utility," *Annals of Noninvasive Electrocardiology*, vol. 10, pp. 88–101, 2005.

[109] J. Nolan, P. Batin, and R. A. et al., "Prospective study of heart rate variability and mortality in chronic heart failure: results of the united kingdom heart failure evaluation and assessment of risk trial (u.k.-heart)," *Circ*, vol. 98, pp. 1510–1516, 1998.

[110] A. Oppenheim and R. Schafer, *Digital-time Signal Processing*. Pearson, 2007.

[111] R. Bailon, L. Mainardi, M. Orini, L. Sornmo, and P. Laguna, "Analysis of heart rate variability during exercise stress testing using respiratory informaion," *Biomed Sig Proc Contr*, vol. 4, pp. 299–310, 2010.

[112] L. Mainardi, "On the quantification of heart rate variability spectral parameters using time-frequency and time-varying methods," *Phil Trans R Soc A*, vol. 367, pp. 255–275, 2009.

[113] S. Jasson, C. M. edigue, P. Maison-Blanche, N. Montano, L. Meyer, C. Vermeiren, P. Mansier, P. Coumel, A. Malliani, and B. Swynghedauw, "Instant power spectrum analysis of heart rate variability during orthostatic tilt using a time-/frequency-domain method," *Circulation*, vol. 96, no. 10, pp. 3521–26, 1997.

[114] R. Blackman and J. Tuckey, *The Measurement of Power Spectra, from the point of view of communications engineering*. New York: Dover Publications, 1959.

[115] F. Harris, "On the use of windows for harmonic analysis with the discrete fourier transform," *Proc IEEE*, vol. 66, no. 1, pp. 51–83, 1978.

[116] M. Pagani, F. Lombardi, S. Guzzetti, G. Sandrone, O. Rimoldi, G. Malfatto, S. Cerutti, and A. Malliani, "Power spectral density of heart rate variability as an index of sympatho-vagal interactions in normal and hypertensive subjects," *J Hypertens Suppl*, vol. 2, pp. 383–385, 1984.

[117] D. L. Eckberg, "Point:counterpoint: Respiratory sinus arrhythmia is due to a central mechanism vs. respiratory sinus arrhythmia is due to the baroreflex mechanism," *J Appl Physiol*, vol. 106, no. 5, pp. 1740–1742, 2008.

[118] L. Bernardi, F. Keller, M. Sanders, P. Reddy, B. Griffith, F. Meno, and M. Pinsky, "Respiratory sinus arrhythmia in the denervated human heart," *J Appl Physiol*, vol. 67, pp. 1447–1455, 1989.

[119] H. B. Hopf, A. Skyschally, G. Heusch, and J. Peters, "Low-frequency spectral power of heart rate variability is not a specific marker of cardiac sympathetic modulation," *Anesthesiology*, vol. 82, no. 9, pp. 609–619, 1995.

[120] B. A. Kingwell, J. M. Thompson, D. M. Kaye, G. A. McPherson, G. Jennings, and M. Esler, "Heart rate spectral analysis, cardiac norepinephrine spillover and muscle sympathetic nerve activity during human sympathetic nervous activation and failure," *Circulation*, vol. 90, pp. 234–240, 1994.

[121] M. Esler, G. Jennings, P. Korner, P. Blombery, N. Sacharias, and P. Leonard, "Measurement of total and organ-specific norepinephrine kinetics in humans," *Am J Physiol*, no. 247, pp. E21–E28, 1984.

[122] P. Grossman and E. W. Taylor, "Toward understanding respiratory sinus arrhythmia: Relations to cardiac vagal tone, evolution and biobehavioral functions," *Biological Psychology*, vol. 74, no. 2, pp. 263–285, 2007.

[123] J. A. Hirsch and B. Bishop, "Respiratory sinus arrhythmia in humans: how breathing pattern modulated heart rate," *Am J Physiol Heart Circ Physiol*, vol. 241, pp. H620–H629, 1981.

[124] R. Bailon, P. Laguna, L. Mainardi, and L. Sornmo, "Analysis of heart rate variability using time-varying frequency bands based on respiratory frequency," *Proceesings of the 29th Annual International Congerence of the IEEE EMBS*, no. 2007, pp. 6674–6677, 2007.

[125] Y. Goren, L. Davrath, I. Pinhas, E. Toledo, and S. Akselrod, "Individual time-dependent spectral boundaries for improved accuracy in time-frequency analysis of heart rate variability," *IEEE Trans Biomed Eng*, vol. 53, no. 1, pp. 35–42, 2006.

[126] G. Parati, J. Saul, M. D. Rienzo, and G. Mancia, "Spectral analysis or blood pressure and heart rate variability in evaluating cardiovascular regulation. a critical appraisal," *Hypertension*, vol. 25, pp. 1276–1286, 1995.

[127] M. Malik and A. Camm, *Heart Rate Variability*. Futura Publishing Company, 1995.

[128] C. Julien, "The enigma of mayer waves: Facts and models," *Cardiovascular Research*, vol. 70, pp. 12–21, 2006.

[129] T. Brown, L. Beightol, J. Koh, and D. Eckberg, "Important influence of respiration on human rr interval power spectra is largely ignored," *J Appl Physiol*, vol. 75, no. 5, pp. 2310–2317, 1993.

[130] D. Bloomfield, A. Magnano, J. Bigger, H. Rivadeneira, M. Parides, and R. Steinman, "Comparison of spontaneous vs. metronome-guided breathing on assessment of vagal modulation using rr variability," *Am J Physiol Heart Circ Physiol*, vol. 280, pp. H1145–H1150, 2001.

[131] O. Shapira and Y. Bar-Khayim, "Ecg changes and cardiac arrhythmias in chronic renal failure patients on hemodialysis," *J Electrocardiol*, vol. 25, no. 4, pp. 273–279, 1992.

[132] M. V. Kamath and E. L. Fallen, "Correction of the heart rate variability signal for ectopics and missing beats," in *Heart Rate Variability* (M. Malik and A. J. Camms, eds.), pp. 75–86, Armonk: Futura Publ, 1995.

[133] G. Clifford, P. McSharry, and L. Tarassenko, "Characterizing artefact in the normal human 24-hour rr time series to aid identification and artificial replication of circadian variations in human beat to beat heart rate using a simple threshold," *Computers in Cardiology*, vol. 29, pp. 129–132, 2002.

[134] G. D. Clifford and L. Tarassenko, "Quantifying errors in spectral estimates of hrv due to beat replacement and resampling," *IEEE Trans Biomed Eng*, vol. 52, no. 4, pp. 630–638, 2005.

[135] T. F. of the European Society of Cardiology, the North American Society of Pacing, and Electrophysiology, "Heart rate variability: Standards of measurement, physiological interpretation and clinical use," *Circulation*, vol. 93, no. 5, pp. 1043–65, 1996.

[136] O. Liangos, B. Pereira, and B. Jaber, "Acute complications associated with hemodialysis," in *Chronic Kidney Disease, Dialysis, and Transplantation* (B. Pereira, M. Sayegh, and P. Blake, eds.), pp. 451–471, Elsevier Saunders, 2005.

[137] J. Schadt and J. Ludbrook, "Hemodynamic and neurohumoral responses to acute hypovolemia in conscious mammals," *AJP Heart*, vol. 260, no. 2, pp. H305–H318, 1991.

[138] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol. 43, pp. 59–69, 1982.

[139] T. Kohonen, *Self-Organizing Maps*. Springer, 1995.

[140] J. MacQueen and W. M. S. I. UCLA, "Some Methods for Classification and Analysis of Multivariate Observations," 1966.

[141] E. Erwin, K. Obermayer, and K. Schulten, "Self-organizing maps: ordering, convergence properties and energy functions," *Biological Cybernetics*, vol. 67, pp. 47–55, 1992.

[142] J. J.W. Sammon, "A Nonlinear Mapping for Data Structure Analysis.," *IEE Transactions on Computers*, vol. C18, no. 5, pp. 401–409, 1969.

[143] J. Mao and A. Jain, "Artificial neural networks for feature extraction and multivariate data projection," *IEEE Trans. on Neural Networks*, vol. 6, pp. 296–317, 1995.

[144] D. Lowe and M. Tipping, "Neuroscale: Novel Topographic Feature Extraction with Radial Basis Function Networks.," *Advances in Neural Information Processing Systems*, vol. 9, pp. 543–549, 1997.

[145] M. E. Tipping and D. Lowe, "Shadow targets: A novel algorithm for topographic projections by radial basis functions," *NeuroComputing*, vol. 19, pp. 211–222, 1997.

[146] L. Tarassenko, D. Clifton, P. Bannister, S. King, and D. King, "Novelty detection," in *Encyclopaedia of Structural Health Monitoring* (K. W. et al., ed.), John Wiley and Sons, 2008.

[147] I. T. Nabney, *NETLAB: Algorithms for Pattern Recognition*. Springer, 2002.

[148] D. Wong, "Optimisation of a multi-parameter monitor for early warning of patient deterioration," *Technical Report*, vol. Universtiy of Oxford, 2007.

[149] L. Tarassenko, *A Guide to Neural Computing Applications*. Butterworth-Heinemann Ltd, 1998.

[150] D. A. Clifton, L. A. Clifton, P. R. Bannister, and L. Tarassenko, "Automated novelty detection in industrial systems," in *Advances of Computational Intelligence in Industrial Systems* (Y. Liu, A. Sun, H. T. Loh, W. F. Lu, and E. Lim, eds.), Springer Berlin Heidelberg, 2008.

[151] S. Sundaram, I. G. D. Strachan, D. A. Clifton, L. Tarassenko, and S. King, "Aircraft engine health monitoring using density modelling and extreme value statistics," *Proc 6th Intern Conference on Condition Monitoring and Machinery Failure Prevention Technologies*, pp. 919–930, 2009.

[152] F. Javed, A. Savkin, G. S. H. Chan, P. Middleton, P. Malouf, E. Steel, J. Mackie, and N. Lovell, "Assessing the blood volume and heart rate responses during haemodialysis in fluid overloaded patients using support vector regression," *Physiol Meas*, vol. 30, pp. 1251–1266, 2009.

[153] M. Hakumaki, "Seventy years of the Bainbridge reflex," *Acta Physiologica Scandiavica*, vol. 130, no. 2, pp. 177–185, 1987.

[154] V. M. Campese, M. Romoff, D. Levitan, K. Lane, and S. Massry, "Mechanisms of autonomic nervous system dysfunction in uremia," *Kidney Int*, vol. 220, pp. 246–253, 1981.

[155] H. Kaufmann, N. Kaplan, and R. Freeman, "Mechanisms, causes, and evaluation of orthostatic and postprandial hypotension," in *UpToDate* (D. S. Basow, ed.), Waltham, MA.: UpToDate, 2012.

[156] G. A. Babich and O. Camps, "Weighted parzen windows for pattern classification," *IEE trans. Pattern Analysis and Machine Intelligence*, pp. 567–570, 1996.

[157] C. M. Bishop, "Novelty Detection and Neural-Network Validation," *IEE Proceedings-Vision Image and Signal Processing*, vol. 141, no. 4, pp. 217–222, 1994.

[158] L. Tarassenko, A. Hann, A. Patterson, E. Braithwaite, K. Davidson, V. Barber, and D. Young, "Biosign®: multi-parameter monitoring for early warning of patient deterioration," *Proc 3rd IEE International Seminar on Medical Applications of Signal Processing,*, pp. 71–6, 2005.

[159] D. A. Clifton, S. Hugueny, and L. Tarassenko, "Novelty detection with multivariate extreme value statistics," *Journal of Signal Processing Systems for Signal Image and Video Technology*, vol. 65, no. 3, pp. 371–389, 2011.

[160] B. Kirkwood and J. A. C. Sterne, *Essential Medical Statistics*. Wiley-Blackwell, 2003.

[161] L. Clifton, D. A. Clifton, M. A. F. Pimentel, P. J. Watkinson, and L. Tarassenko, "Gaussian processes for personalized e-health monitoring with wearable sensors," *IEEE Trans Biomed Eng*, vol. 60, no. 1, pp. 193–197, 2013.

[162] K. Nakane, T. Shinzato, and K. Maeda, "Continuous measurement of pressure in internal arteriovenous fistulae and prevention of symptomatic hypotension during hemodialysis," *ASAIO Transactions*, vol. 37, no. 3, pp. M232–4, 1991.

[163] G. Langewouters, J. Settels, R. Roelandt, and K. Wesseling, "Why use finapres or portapres rather than intra-arterial or intermittent non-invasive techniques of blood pressure measurement?," *J Med Eng Technol*, vol. 22, pp. 37–43, 1998.

[164] I. Guelen, B. E. Westerhof, G. L. V. D. Sar, G. A. V. Montfrans, F. Kiemeneij, K. H. Wesseling, and W. J. Bos, "Finometer, finger pressure measurements with the possibility to reconstruct brachial pressure," *Blood Press Monit*, vol. 8, pp. 27–30, 2003.

[165] A. Elvan-Taspinar, L. A. Uiterkamp, J. M. Sikkema, M. L. Bots, H. A. Koomans, H. W. Bruimse, and A. Franx, "Validation and use of the finometer for blood pressure measurement in normal, hypertensive and preeclamptic pregnancy," *J Hypertens*, vol. 21, pp. 2053–60, 2003.

[166] A. E. Schutte, H. W. Huisman, J. M. van Rooyen, N. T. Malan, and R. Schutte, "Validation of the finometer device for measurement of blood pressure in black women," *J Hum Hypertens*, vol. 18, pp. 79–84, 2004.

[167] E. Mistrik, S. D. Sulkova, V. Blaha, M. Kalousova, J. Knizek, P. Moucka, V. Herout, M. Kadlec, R. Stilec, and L. Sobotka, "Evaluation of skin microcirculation during hemodialysis," *Renal Failure*, vol. 32, no. 1, pp. 21–26, 2010.

[168] G. Ligtenberg, M. Barnas, and H. Koomans, "Intradialytic hypotension: New insights into the mechanism of vasovagal syncope," *Nephrol Dial Transplant*, no. 13, pp. 2745–2747, 1998.

[169] R. Bemelmans, E. C. Boerma, J. Barendregt, C. Ince, J. Rommes, and P. Spronk, "Changes in the volume status of haemodialysis patients are reflected in sublingual microvascular perfusion," *Nephrol Dial Transplant*, vol. 24, pp. 3487–92, 2009.

[170] A. Kholoussy, S. Sufian, C. Pavlides, and T. Matsumoto, "Central peripheral temperature gradient: Its value and limitations in the management of critically ill surgical patients.," *Am J Surg*, vol. 140, no. 5, pp. 609–612, 1980.

[171] J. W. Kantelhardt, A. Bauer, A. Y. Schumann, P. Barthel, R. Schneider, M. Malik, and G. Schmidt, "Phase-rectified signal averaging for the detection of quasi-periodicities and the prediction of cardiovascular risk," *Chaos*, vol. 17, no. 1, 2007.