Visualisation and Pattern Recognition of Heart Rate Variability

Ben Raymond

Thesis submitted for the degree of
Doctor of Philosophy

March 1999

Departments of Electrical and Electronic Engineering
and Applied Mathematics
The University of Adelaide
Australia

Abstract

Analysis of heart rate variability (HRV) may be used as a noninvasive probe of autonomic nervous system activity. It has found application in many situations where autonomic function is of interest, such as diabetic autonomic neuropathy, hypertension, and recovery following myocardial infarction. This thesis considers various signal processing aspects of HRV analysis; in particular, those of data visualisation and classification.

Frequency domain techniques are popular for HRV processing since the oscillations in the signal are largely separable by frequency. The analysis of spectral HRV data is conventionally performed by considering the power contained in sub-bands within the spectrum. This approach is not ideal: the choice of the limits of the sub-bands can be difficult, and the band-power values do not have a convenient physiological interpretation.

In this thesis, data visualisation techniques - specifically the least-squares scaling and the generative topographic mapping - are applied to the analysis of spectral HRV data. These techniques take a high-dimensional data set and construct a two-dimensional representation in which the relationships between the data are preserved. This allows simple graphical depiction and visual exploration of the data by the user. With spectral data, this approach allows the physiological significance of the data to be ascertained without the need to perform sub-band analysis. The techniques are quite general and may be applied to the analysis of data from other processing methods, such as nonlinear HRV analysis and cardiovascular system modelling. The use of the least-squares scaling and the generative topographic mapping for the visualisation of spectral HRV data is demonstrated on data from two studies: one on sleep apnoea, and the other on the effects on propranolol and posture on heart rate variability.

The fusion of HRV with information from other sources is one method of improving the reliability of HRV analysis. The use of classification techniques for information fusion is investigated in the second part of this thesis. For the discrimination of heart rate data by posture and propranolol, the fusion of HRV information with mean heart rate is shown to yield better discrimination than the use of HRV alone. The use of HRV with other diagnostic information for the detection of sleep apnoea is also explored.

Table of contents

Errata

page xii	The definition of **myocardial infarction** should read: "An area of muscle death in the heart caused by an acute lack of blood supply to the affected area."
page xiii	The definition of the **vagus nerve** should read: "The vagus nerve is the major branch of the parasympathetic nervous system. One of the functions of the vagus nerve is to retard heart rate."
page 3 paragraph 3	The last sentence should read: "These methods suffer from various limitations as described in Section 2.4. The inexpensive, non-invasive nature of the data acquisition process, coupled with the vast array of applications warrants the development of more reliable HRV processing techniques."
page 5 paragraph 3	The last sentence should read: "The combination of heart rate variability with oxyhaemoglobin desaturation information is shown to be of use for the detection of sleep apnoea."
page 42 paragraph 3	Add: "Mappings in three dimensions are possible and may offer a less ambiguous depiction of a data set than a two dimensional mapping. However, special hardware is required to display three dimensional images and so two dimensional mappings are by far the most commonly used."
page 63, opening paragraph	Add: "No examples of the application of the GTM or the LSS to the visualisation of HRV data were found in the literature. The contribution of this section is to demonstrate the use of these techniques for this purpose."
page 69, opening paragraph	Add: "The ECG was sampled at 250Hz."
page 130 paragraph 2	Add: "Here, the LSS and GTM are essentially acting as feature extractors. Other feature extraction algorithms may also be appropriate for this task."
page 161	Equation D1 may be found in Bishop *et al.*, 1998b.

Statement of originality

I hereby declare that this work contains no material which had been accepted for the award of any other degree or diploma in any university or other tertiary institution and that, to the best of my knowledge, it contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent for a copy of my thesis, when deposited in the University Library, being made available for loan and photocopying.

Ben Raymond

Signed:

Date: 20/4/99

Acknowledgements

Many people deserve my gratitude for their assistance with this work. My research supervisors Bob Bogner, Jag Mazumdar and Nanda Nandagopal have each made significant contributions to the direction and content of my research, and to them I am very grateful. Bob in particular has been inestimably helpful in providing critical feedback and evaluation of my work. My research colleagues David Crisp, Michael Lee, Cheng Chew Lim, Daniel McMichael, Kenneth Pope, and Jamie Sherrah have also provided valuable advice and support. I am grateful to Mary Raymond and Tania Streeter for reviewing draft material.

I would like to thank Norm Blockley, David Bowler and Grant Ward for maintaining the computer systems in excellent health: without their efforts, very little work would get done. In my day-to-day activities I have had the pleasure of working with many fantastic people: Andrew Allison, Richard Bowyer, Kathryn Burgess, Anna Buzescu, Hong Gunn Chew, Charlie Green, Colleen Greenwood, Stephen Guest, Kay Leverett, Yadi Parrott, Mary Parry, and Zhishun She. My thanks are also due to Tony Parker and Peter Gill, the heads of my departments, and Henry d'Assumpcao, the director of the Centre for Sensor, Signal and Information Processing, who have always provided support for my work.

The experimental data used in Chapter 4 was provided by Clarice Chian, Anne Tonkin and David Taverner of the Department of Clinical and Experimental Pharmacology at the University of Adelaide. The sleep apnoea data used in Chapter 5 was graciously provided by Michael Hilton and Ruth Cayton of the Birmingham Heartlands Hospital in the United Kingdom.

International conference travel for this work was assisted by grants from the Cooperative Research Centre for Sensor Signal and Information Processing, the Departments of Electrical and Electronic Engineering and Applied Mathematics at the University of Adelaide, The National Heart Foundation of Australia, the Walter and Dorothy Duncan Trust, and ResMed Asia Pacific.

Some of the MATLAB code used in this work was written by other people: I would like to thank Mark Orr of the Centre for Cognitive Science at Edinburgh University for the RBF

toolbox, Markus Svensén of the Neural Computing Research Group at Aston University for the GTM toolbox, and Geoff Jarrad and Daniel McMichael of the Cooperative Research Centre for Sensor, Signal and Information Processing for the shared mixture classifier code.

Finally, I would like to thank my family and friends for their support during my studies.

Glossary and abbreviations

ABP *arterial blood pressure*
(See page 8)

ANS *autonomic nervous system*
The sympathetic and parasympathetic branches of the autonomic nervous system are primarily responsible for the control of heart rate. (See page 8)

AR *autoregressive*
A form of modelling in which each sample of a signal is modelled as a weighted sum of previous sample values. (See page 17)

ECG *electrocardiogram*
(See page 12)

GTM *generative topographic mapping*
The generative topographic mapping is a tool for the visualisation of high dimensional data. (See page 54)

HF *high frequency*
The high frequency component of heart rate variability is caused by respiration and occurs at respiratory frequency. (See pages 10 and 23)

HRV *heart rate variability*
Variations in heart rate. The term "heart rate variability" is also often used to describe heart *interval* variability, although this is not strictly correct. (See page 7)

IPFM *integral pulse frequency modulator*
(See page 14)

LDA *linear discriminant analysis*
A simple approach to classification, in which the data density for each class is modelled using a normal distribution. (See page 131)

LF *low frequency*
The low frequency component of heart rate variability is an oscillation at nominally 0.1 Hz. The precise frequency at which this oscillation occurs is called the LF peak frequency. The LF component is sometimes referred to as the mid-frequency (MF) component. (See pages 10 and 23).

LSS — *least-squares scaling*
The least-squares scaling is a form of multidimensional scaling, which employs a least-squares cost function. (See page 47)

MAR — *multivariate autoregressive model*
An extension of the autoregressive model, which can be used to model the interactions between multiple signals. (See page 26)

MDS — *multidimensional scaling*
A tool for the visualisation of high-dimension data. (See page 45)

MI — *myocardial infarction*
Heart attack. (See page 30)

parasympathetic nerve — see vagus nerve

PCA — *principal components analysis*
A well-understood method of dimension reduction. It operates by finding the projection of the data onto the principal eigenvectors of the covariance matrix.

PSD — *power spectral density*
(See page 14)

relative supervision — A method of training radial basis function networks to produce least-squares scalings. In this method the construction of the mapping and the estimation of the network parameters is carried out simultaneously. (See page 82)

RR interval — The interval between R-peaks on the electrocardiogram, commonly used as a measure of the interval between heart beats. (See page 12)

RRV — *heart period (RR interval) variability*

RSA — *respiratory sinus arrhythmia*
Also known as the high frequency (HF) component of heart rate variability. (See pages 10 and 23)

SA — *sinoatrial*
The sinoatrial node is normally the pacemaker for the heart. Each pulse from the SA node causes the heart to beat. (See page 8)

SAHS — *sleep apnoea/hypopnoea syndrome*
The sleep apnoea/hypopnoea syndrome describes a condition in which respiration stops or is impeded during sleep. (See page 105)

SaO_2 — *arterial haemoglobin saturation*
An indirect but convenient measure of the oxygen content of the arterial blood. (See page 106)

SMC — *shared mixture classifier*
A classifier which uses mixtures-of-Gaussians to model the data density distributions. (See page 130)

SOM — *self organising map*
An unsupervised neural network which can be used for data visualisation. (See page 56)

stress
The value of the error function for a multidimensional scaling. The stress value is essentially a measure of how well the scaling represents the original data, with lower stress values indicating a better representation. (See page 86)

sympathetic nerve
The sympathetic nerve is a branch of the autonomic nervous system. Sympathetic nerve activity acts to increase the heart rate. (See page 8)

sympathovagal balance
A qualitative term used to describe the relative contributions of the sympathetic and vagus nerves to heart rate control. (See page 9)

vagus nerve
Activity in the vagus nerve (also known as the parasympathetic nerve) acts to slow the heart rate. (See page 8)

VLF — *very low frequency*
Activity below about 0.05 Hz is known as the very low frequency component of heart rate variability. (See page 12)

Related publications

Raymond, B., Nandagopal, D., Mazumdar, J., and Taverner, D. (1995). Neural network based feature extraction scheme for heart rate variability. *Proceedings of the SPIE Conference on Applications and Science of Artificial Neural Networks*, volume 2492, pages 893–904, Orlando, USA.

Raymond, B., Mazumdar, J., and Nandagopal, D. (1997a). Modelling the shift of the low frequency component of heart rate variability. *Computers in Cardiology 1997*, pages 407–410, Lund, Sweden. IEEE Computer Society Press, Los Amitos.

Raymond, B., Taverner, D., Nandagopal, D., and Mazumdar, J. (1997b). Classification of heart rate variability in patients with mild hypertension. *Australasian Physical & Engineering Sciences in Medicine*, **20**(4):207–213.

Raymond, B. and Mazumdar, J. (1998). A neural network for heart rate variability analysis. *Engineering Mathematics and Applications Conference*, pages 427–430, Adelaide, Australia.

Raymond, B., Bogner, R., Mazumdar, J., and Hilton, M. (1998). Visualisation of heart rate variability data using topographic mappings. *Computers in Cardiology 1998*, pages 141-144, Cleveland, USA. IEEE Computer Society Press, Los Amitos.

Chapter 1 Introduction

Variations in instantaneous heart rate reflect the beat-to-beat influence of the cardiovascular control systems. The clinical significance of these variations, assessed through simple statistical measures such as the variance of heart rate, was recognised as early as 1965 (Hon and Lee, 1965). However, it was not until spectral techniques were made popular (Sayers, 1973; Womack, 1971) that the potential scope and value of heart rate variability processing became apparent. Typically, the heart rate variability (HRV) signal contains an oscillation which is synchronous with respiration (known as respiratory sinus arrhythmia, and ordinarily found between 0.2 and 0.4 Hz), an oscillation at nominally 0.1Hz, thought to be due at least in part to blood pressure regulation (Baselli *et al.*, 1994), and a variety of oscillations at lower frequencies (Moser *et al.*, 1994). Processing the signal in the spectral domain allows these components to be identified and treated separately.

The control of the heart rate is effected through the sympathetic and parasympathetic branches of the autonomic nervous system and HRV can thus be used as a probe of autonomic function (Pomeranz *et al.*, 1985; Akselrod *et al.*, 1981). The method is non-invasive, since only the time of each heart beat is required and this is easily obtained from the electrocardiogram. Heart rate variability analysis has been shown to be of value in a wide range of applications, including assessing recovery after myocardial infarction (Malik and Camm, 1994; Bigger Jr *et al.*, 1988; Wolf *et al.*, 1978) and investigating the changes in autonomic function which accompany hypertension (Guzzetti *et al.*, 1988; Pagani *et al.*, 1984) and diabetic autonomic neuropathy (Bellavere *et al.*, 1992; van den Akker *et al.*, 1983).

1.1 Research problem

Despite the attractions of non-invasive data acquisition and wide potential application, HRV is subject to various limitations which restrict its practical applicability. The broad motivation of this thesis is to identify deficiencies in current HRV processing techniques and to develop solutions to remedy these deficiencies.

In order to facilitate interpretation and comparison of results, the HRV signal is usually reduced to a few interesting features. In frequency domain, spectra are most commonly assessed in terms of the power contained within sub-bands of the spectrum. In the time domain, the signal is characterised in terms of variance or other statistical measures, such as the proportion of intervals in which the difference between successive intervals is greater than 50 milliseconds (Kleiger *et al.*, 1995). Nonlinear methods may lead to characterisation in terms of system complexity (Signorini *et al.*, 1994) or entropy (Zebrowski *et al.*, 1994).

In the frequency-band approach, the high-frequency band encompasses respiratory frequencies, and is accepted to be a marker of parasympathetic modulation of heart rate. Bands at lower frequencies are also commonly used and both sympathetic and parasympathetic activity influence the oscillations in these bands. There is no accepted marker which relates purely to sympathetic modulation of heart rate. The ratio of low-frequency to high-frequency power has been suggested for this purpose (Pagani *et al.*, 1986), but in the general case, "will not provide an accurate measure, neither of the sympathetic tone, nor of the ratio of sympathetic to parasympathetic activity" (Akselrod, 1995, p. 156). The use of sub-bands for analysis also requires the choice of the limits of these bands for which different researchers argue different values (e.g. Takalo *et al.*, 1994; Malliani *et al.*, 1994b). Time domain analysis similarly lacks a specific marker of sympathetic activity. Methods of nonlinear analysis are in a relatively early stage of development and require further work in application and interpretation (Task Force, 1996).

Thus, the problem of finding suitable descriptors of the HRV signal is one which is not well addressed by existing techniques. The identification of a few salient features which provide a useful characterisation of a data set is not unique to HRV analysis. A variety of general data visualisation algorithms exist for this purpose, such as principal components analysis (Hotelling, 1933) and multidimensional scaling (Torgerson, 1958; Young and Householder, 1938). These methods seek to construct a two dimensional representation of the data set, allowing graphical depiction and visual interpretation by the user. Visualisation is most effectively carried out interactively, allowing the user to direct and fine-tune the process. In this work, it is shown that data visualisation algorithms, in particular multidimensional scaling and the generative topographic mapping, can be extremely useful for the analysis of HRV information.

Despite its widespread use as a probe of autonomic function, HRV analysis does not directly provide information on autonomic tone. Since variations in heart rate are caused by

variations in autonomic activity, the amplitude of HRV offers information relating only to modulation of autonomic tone and not to tone itself. It is common to assume that modulation depth is proportional to tone (Akselrod, 1995), but this assumption is clearly invalid if autonomic tone begins to saturate (Malik and Camm, 1993). Additionally, noise (for example, measurement noise in the beat times or spectral estimation error due to transient signals) can further obscure the results of HRV analysis.

In this work, we investigate the combination of HRV information from other signals and sources, as a method of improving the reliability and utility of HRV analysis. By considering several information sources in parallel, the detrimental effects of noise in one source can be reduced. A convenient means of accomplishing this information "fusion" is through classification. In this approach, the available information is combined in such a fashion as to maximise the correct classification of the data into predefined categories. We show that the combination of HRV with mean heart rate and other diagnostic information can improve the discrimination of heart rate data, compared to the use of HRV alone.

1.2 Justification for the research

A number of commercial instruments exist which offer facilities for HRV analysis. Most employ relatively simple processing techniques, such as time-domain measures and Fourier transform-based spectral analysis (Kennedy, 1995). The inexpensive, non-invasive nature of the data acquisition process, coupled with the vast array of applications warrants the development of more reliable HRV processing techniques.

Data visualisation techniques are widely used for other information processing applications, particularly in the field of biomedicine (van Gils *et al.*, 1997; Lehtinen *et al.*, 1997). Visualisation techniques do not attempt to make an explicit classification of the data but simply present a clear, graphical depiction of the relationships within the data set to the user, often in an interactive fashion. This approach explicitly includes the human operator as a part of the information processing chain, exploiting their excellent visual pattern recognition and cluster analysis skills as well as their expert knowledge of the significance of the data. Such interactivity has been identified as extremely beneficial for data exploration (Siedlecki *et al.*, 1988b).

Graphical techniques are well-established as tools for HRV processing. Their robustness to outliers (Malik, 1995) makes these methods particularly attractive for long-term data analysis, where arrhythmias and other noise are difficult to avoid (Malik *et al.*, 1993). Existing geometrical methods, however, operate only on heart interval series data. The use of more general data visualisation techniques will retain the advantages of graphical analysis, while extending the range of data types which may be processed.

The success of visualisation methods in many areas of biomedical data processing, along with the established nature of other graphical methods for HRV analysis, suggest that visualisation algorithms have much to offer for HRV analysis.

Classification as a means of accomplishing information fusion has also found prior application in HRV analysis. Depressed HRV is known to be an indicator of risk of mortality following myocardial infarction (Wolf *et al.*, 1978). The predictive accuracy of HRV in this regard has been shown to improve when combined with other indicators of risk such as left ventricular ejection fraction (Bigger Jr *et al.*, 1992b). However, the use of classification techniques in this manner has not been widespread in the field of HRV processing.

1.3 Thesis outline

The thesis proper begins with a concise review of existing HRV processing methods and an exposition of research issues in this area. Brief reviews of the background physiology and of clinical applications of HRV analysis are also given.

The development of data visualisation techniques as tools for HRV analysis begins in Chapter 3. Some desirable attributes of visualisation algorithms for biomedical data analysis are first identified. It is suggested that data visualisation algorithms should be inductive (able to project new data onto the mapping efficiently), interactive, flexible (able to operate on a variety of data types), quantifiable (the final result can be objectively assessed), nonlinear and computationally tractable. Three well-known visualisation algorithms are reviewed in the light of these suggestions. Multidimensional scaling (in particular the least-squares scaling) and the generative topographic mapping are identified as the most suitable algorithms. These algorithms are subsequently applied to HRV data in Chapter 4. Suitable preprocessing for spectral data is discussed, and data from a simple physiological experiment is used to provide a demonstration.

The ability to project new data, not initially used during map construction, is a particularly valuable asset for data visualisation. This allows the comparison of unknown data against a database of known samples, thus offering the potential for use as a diagnostic aid. While the least-squares scaling has a number of favourable properties, the ability to generalise is not amongst them. A popular means of redressing this disadvantage is through the use of neural networks. Once a network is trained to produce a least-squares scaling, new data may be projected with low computational overhead. Chapter 5 explores aspects of the training of radial basis function networks for this purpose. Particular attention is given to regularisation, since appropriate regularisation is essential for good generalisation performance.

The use of prior information during map construction is also addressed in Chapter 5. While the basic least-squares scaling and generative topographic mapping algorithms are unsupervised, the incorporation of prior data knowledge can often improve cluster separation or otherwise improve the interpretability of the mapping. Links between the use of prior information and the network regularisation are drawn, and the findings of the chapter are demonstrated on data from a sleep apnoea study.

The discussion of regularisation and generalisation leads naturally to the use of visualisation techniques for classification tasks. The extraction of features upon which classification decisions can be made is a similar task to that of data visualisation, and so it is plausible to use visualisation techniques for this feature extraction. The use of classification techniques as a means of combining heart rate variability with other information is broached in Chapter 6. Features are extracted from the heart rate variability signal using the visualisation techniques developed in the preceding sections, and then classification is performed. The classification with the combination of mean heart rate and heart rate variability is shown to outperform the classification with heart rate variability alone, for the tasks of detecting the changes in heart rate with posture and propranolol. The combination of heart rate variability with other diagnostic information is also considered and shown to be of use for the detection of sleep apnoea.

Finally, the thesis is concluded with a summary and suggestions for further work.

1.4 Scope

The scope of this thesis is delimited to the investigation of the technical aspects of the processing of the heart rate variability signal, focusing primarily on the principles and algorithms needed to accomplish this task. The identification of new clinical applications and the development of models of the cardiovascular system are outside the scope of this work.

Chapter 2 Heart rate variability

The term *heart rate variability* (HRV) refers to changes in the heart rate about its mean level. In a normal subject, even at rest, the heart rate varies from one beat to the next. Variations in blood pressure linked to respiration were first documented by Hales in 1733 (cited in Malliani *et al.*, 1991a). Arterial blood pressure variations with a period of approximately 10 seconds (Mayer waves) were noted in 1876 (Mayer, cited in Malliani *et al.*, 1991a).

The first indication that heart rate variations were linked to nerve activity came in 1932, when activity in the sympathetic nerve, a branch of the autonomic nervous system, was seen to change in rhythm with cardiac cycle and respiration (Adrian *et al.*, 1932). Similar activity but simultaneous to arterial Mayer waves was documented in 1965 (de Molina and Perl, 1965). The first clinical application of HRV analysis was presented in 1965, when a decreased variance of the foetal heart rate was found to be an indicator of foetal distress (Hon and Lee, 1965).

Early analysis of HRV was limited to simple time domain methods. The introduction of frequency-domain techniques (Sayers, 1973; Luczak and Laurig, 1973; Womack, 1971). allowed the various frequency components of the signal to be identified for the first time. Applications were found almost immediately (e.g. Chess *et al.*, 1975).

Animal experiments (Akselrod *et al.*, 1981) demonstrated the use of HRV for investigating variations in autonomic tone. Using pharmacological blockade of the sympathetic and parasympathetic nervous system, the components present in the HRV spectrum were matched with the activity in these nerves. This work was the first illustration of the utility of HRV analysis as a probe of autonomic function. The work was repeated on human subjects (Pomeranz *et al.*, 1985; Pagani *et al.*, 1984; Selman *et al.*, 1982) with similar results.

The non-invasive nature and simple data acquisition were appreciated by medical researchers, and the method quickly found application as a probe of autonomic function. As the popularity of HRV studies increased, new methods of processing the signal also appeared.

The modern field of HRV processing is extremely diverse, involving spectral estimation, system modelling, and nonlinear and chaotic analysis. The literature is very active - a combined search of the Medline and INSPEC databases for 1998 alone revealed over 400 journal and conference articles.

2.1 Physiology of heart rate variability

The function of the mammalian respiratory and cardiovascular systems is to supply oxygen to, and remove carbon dioxide from the tissues of the body. Blood is used to transport oxygen and carbon dioxide, as well as other substances including nutrients and hormones (Scott, 1986). The exchange of carbon dioxide in the blood for oxygen occurs by simple diffusion within the lungs. In order to fulfil the body's oxygen requirements over a range of conditions, the flow of blood through the body must be variable. The heart is the pump responsible for propelling blood through the body. Variation of either stroke volume (the volume of blood ejected with each heart beat) or heart rate, or both, will result in a variation of the rate of blood flow from the heart.

In normal circumstances, the pacemaker for the heart is the sinoatrial (SA) node. This is a group of cells within the heart which exhibit spontaneous depolarisation, allowing the node to produce autonomous, rhythmic electrical impulses. Each of these impulses causes a wave of depolarisation to propagate through the myocardium, causing contraction of the cardiac muscles and pumping of blood.

The discharge rate of the SA node is controlled primarily by the dual innervation of the sympathetic and vagal (or parasympathetic) branches of the autonomic nervous system. Activity in the sympathetic nerve increases the firing rate, while activity in the vagus has the opposite effect. A schematic diagram of nerve activity and heart rate is shown in Figure 2.1.

In order to maintain adequate blood flow to the body tissue, there exist a number of feedback systems which contribute to the control of heart rate. These cardiovascular reflexes principally provide regulation of arterial blood pressure and gases, although other reflexes exist, such as those regulating blood flow to the brain (Vogel, 1992).

The carotid sinus and aortic arch baroreceptors, located in the carotid sinus in the neck and in the arch of the aorta respectively, provide pressure-sensitive feedback to the cardiac control

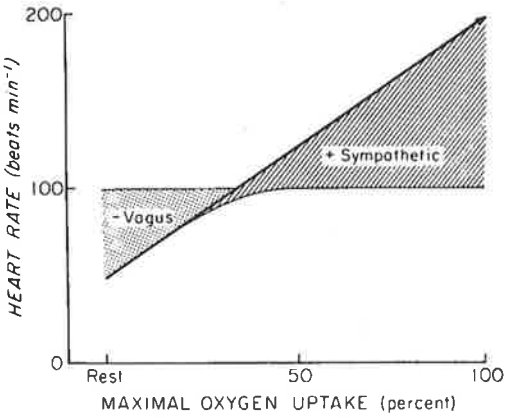

Figure 2.1: Changes in the sympathovagal balance with exercise. Sympathetic and vagal activity tends to be reciprocal: sympathetic stimulation is accompanied by vagal withdrawal. Reproduced from Warner and Cox, 1962.

centres. The feedback signal is a train of discharge pulses, and the rate of discharge is increased by the arterial wall stretch induced by a pressure increase. The mean rate of discharge of carotid sinus is directly proportional to the mean arterial pressure (Vogel, 1992).

An increase in arterial blood pressure (ABP) will elicit an increased firing rate in the baroreceptor. This in turn alters the sympathovagal balance, decreasing the activity of the sympathetic nerve and increasing that of the vagus. The decreased sympathetic activity will also activate other mechanisms including decreased myocardial contractility, leading to an overall decrease in cardiac output, and therefore blood pressure.

The information supplied by the baroreceptors to the brain is essentially concerned with short-term changes in blood pressure. If the blood pressure levels deviate from their normal levels for an extended period of time (more than a few days), the baroreceptors will adapt to the new pressure levels.

If the change in blood pressure is large, then the usual reciprocal changes in sympathetic and vagal activity become distorted. For large increases in ABP (in excess of 20 mmHg) sympathetic activity becomes completely suppressed, and further increases in pressure will be counteracted by increased vagal activity alone. The opposite holds for a large decrease in blood pressure (Vogel, 1992).

In addition to the two principal baroreceptive sites mentioned above, other sites exist in large systemic veins and the walls of the heart itself.

The peripheral chemoreceptors are located in the carotid and aortic bodies and are stimulated by changes in blood oxygen, carbon dioxide, and acidity. Chemoreceptors are principally concerned with respiratory control, and so will have a secondary effect on heart rate through the mechanisms responsible for respiratory sinus arrhythmia. Under normal conditions, however, the effect of the chemoreceptors on the cardiovascular system is minimal (McDonald, 1980).

2.1.1 Sources of heart rate variability

The major oscillations in the heart rate, along with their causes, are listed below.

2.1.1.1 Respiratory sinus arrhythmia

Respiratory sinus arrhythmia (RSA) is a variation in heart rate which is coupled to respiratory activity. At a breathing rate of 0.2 Hz (one breath every five seconds), heart rate typically accelerates during inspiration and decelerates during exhalation.

The SA node response to changes in sympathetic activity is slow compared to that of vagal activity. For frequencies above about 0.15 Hz, the SA node is mediated primarily by the vagus nerve.

In addition to direct neural mechanisms, mechanical factors contribute to RSA. The inflation of the lungs during inspiration induces a drop in the pressure within the thoracic cavity. This accelerates the return of blood from the veins to the right side of the heart, causing an increase in heart rate through a mechanism known as the Bainbridge reflex. Once the increased venous return reaches the left side of the heart, it causes an increase in the left ventricular output and hence a rise in arterial blood pressure. This rise will in turn cause a drop in heart rate through the baroreflex control of the heart. The change in intrathoracic pressure generated by lung inflation is dependent on posture.

2.1.1.2 The ten-second rhythm

The oscillations first noted by Mayer in arterial blood pressure also appear in heart rate and are commonly referred to as the low-frequency (LF) component of HRV. The pioneering work by Hyndman explained the variations as a spontaneous oscillation, a consequence of the nonlinearity in the short-term blood pressure regulation loop (Hyndman *et al.*, 1971). A

similar opinion was adopted by others (Rosenblum and Kurths, 1995; Madwed *et al.*, 1989; Kitney *et al.*, 1985).

In 1974, evidence was presented which suggested that these variations were generated by a central oscillator (Preiss and Polosa, 1974). Through the use of a pressure stabilising device, the oscillations in cat arterial blood pressure were removed. Predictably, this resulted in a lack of detectable ten-second oscillations in the feedback nerve activity from the chemo- and baro-receptors. However, in the sympathetic nerve activity (responsible for control of the SA node and thus heart rate) the oscillations remained, although the modulation was "not as strong" (Preiss and Polosa, 1974). A similar experiment using pharmacological pressure stabilisation indicated that feedback was not necessary for these oscillations to occur in the sympathetic nerve activity, and that there must be a central oscillator responsible for their production (Preiss and Polosa, 1974).

More recently, multivariate modelling of the cardiovascular system (Baselli *et al.*, 1994; Baselli *et al.*, 1988b; de Boer *et al.*, 1987; Wesseling *et al.*, 1982) has reconciled these two theories. Oscillations at these frequencies from an external source (which may be a central oscillator or broadband noise) are conveyed directly to the SA node. The slow response of the sympathetic nerve causes a delay in the blood pressure regulation loop, and thus a resonance, which acts to amplify these oscillations (Baselli *et al.*, 1994). The frequency of the oscillation (the frequency of the LF peak in the HRV spectrum) is set by the resonant frequency of the loop.

Physical exercise causes a decrease in the LF peak frequency (Baselli *et al.*, 1991; Kamath *et al.*, 1991) with a corresponding increase in frequency during post-exercise recovery (Kamath *et al.*, 1991). Mental stress similarly causes a decrease in peak frequency (Singh *et al.*, 1993; Nandagopal *et al.*, 1985), as does tilt (Raymond *et al.*, 1997; Weise *et al.*, 1989; Weise and Heydenreich, 1989). Physical and mental exercise and an upright posture are all conditions under which sympathetic control of the heart dominates. Small doses of atropine, thought to cause augmented vagal cardiac inhibition, cause an increase in LF peak frequency; large doses, which paradoxically act to depress parasympathetic and enhance sympathetic action, cause a decrease in peak frequency (Weise *et al.*, 1989).

2.1.1.3 Slow variations

In addition to the two components already described, a multitude of slower oscillations (with period in the range of minutes or hours) also contribute to HRV. These include respiratory tidal volume variations, oscillations in angiotensin levels in circulating blood, circadian rhythms and submultiples thereof, and rhythms associated with REM sleep periodicity (Moser *et al.*, 1994). The physical correlates of slow variations are both poorly understood and an area of continuing research (Task Force, 1996).

2.2 Processing methods

The basic data required for HRV analysis is the occurrence time of each heart beat, which may be conveniently and non-invasively acquired from the electrocardiogram (ECG). Strictly speaking, the P-wave of the ECG corresponds to the onset of depolarisation and hence the time of beat initiation (Figure 2.2).

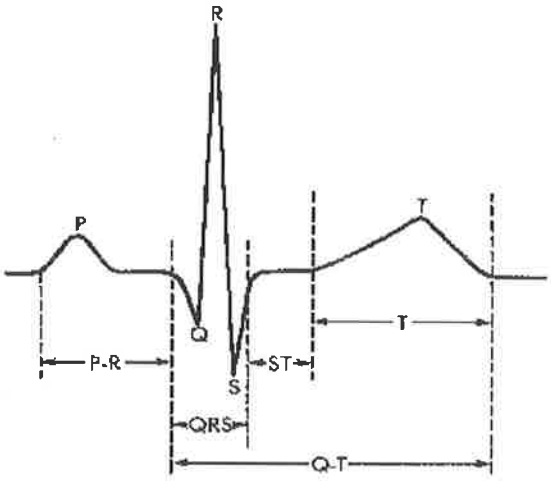

Figure 2.2: Typical electrocardiogram trace, showing the P wave, QRS complex and T wave, and the P-R, Q-T and S-T intervals. Reproduced from Berne and Levy, 1986.

However, the P-wave is quite difficult to detect accurately in comparison with the QRS-complex. The P-R time interval corresponds to the conduction of the impulse through the atrio-ventricular node. Variations in this interval are typically less than 5ms (Koeleman *et al.*, 1985), comparable to the accuracy of most QRS detectors (Laguna *et al.*, 1991; Rompelman *et al.*, 1982). Hence, the R-wave is most commonly used to detect the beat times.

Investigations have been made into the use of the P-wave as a marker of the beat timing with promising results (Laguna *et al.*, 1991).

Errors in the detection of beat times naturally lead to errors in subsequent analysis (see Figure 2.3). For adequate performance, ECG sampling rates of 250-500 Hz are recommended

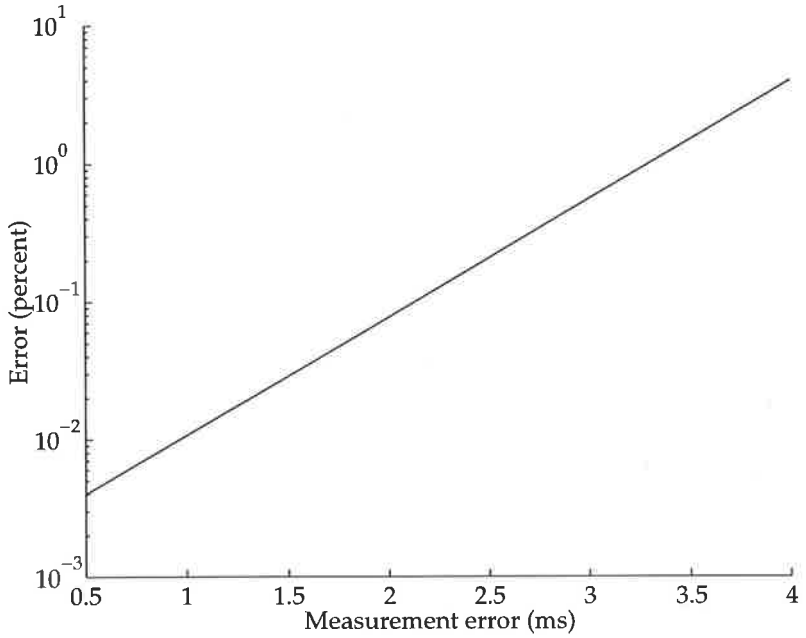

Figure 2.3: Relative error in the power spectrum as a function of the sampling error. (Adapted from Rompelman *et al.*, 1982).

(Hilton *et al.*, 1997; Task Force, 1996), although simple interpolation of the ECG can provide more accurate estimates of the fiducial point of the R-wave with low ECG sampling rates (Bianchi *et al.*, 1993).

Once the beat-times have been acquired, various processing methods may be used to extract information regarding the autonomic control of the heart. Time domain and graphical methods are computationally simple but lack the ability to discriminate between sympathetic and parasympathetic contributions to HRV. Frequency domain techniques offer some improvement in this regard, and a variety of other techniques including chaotic analysis and multivariate modelling may also be used, depending on the application.

2.2.1 Time domain methods

Time domain methods generally work from either the RR interval tachogram (the sequence of time intervals between successive R-waves) or from differences between successive RR intervals. The variance of the RR intervals reflects the action of all modulatory influences on the heart rate, and may be evaluated over short (less than five minutes) or long (12-24 hours) records. The difference between successive intervals is generally due to the fast action of the parasympathetic nerve and may be assessed by (for example) RMSSD (square root of the mean-square difference between successive intervals) or pNN50 (proportion of intervals which exceed previous interval length by 50ms or more).

2.2.2 Graphical methods

The RR interval data may also be assessed by geometric methods, in which a geometric construction is made from the data and then assessed. The geometric construction may take the form of the sample density distribution of intervals or interval differences, or the Poincaré plot (Woo *et al.*, 1992; Anan *et al.*, 1990). These are usually assessed on the basis of size and shape, by measuring various dimensions or proportions (Farrell *et al.*, 1991; Malik *et al.*, 1989). Lorenz or Poincaré diagrams plot each interval length against the next, producing patterns with distinctive shapes, such as fan or comet shapes (Kamen and Tonkin, 1995).

Geometric methods may be applied to data which may contain errors in the recognition of the RR intervals from the ECG (Malik *et al.*, 1993). This is their main advantage (Task Force, 1996). This insensitivity stems at least in part from the ability of the researcher to ignore outliers and concentrate on the main peak of the density curve (Malik, 1995).

Short-term recordings are generally unsuitable for this type of analysis since a large number of data points are needed to get a good representation of the shape.

2.2.3 Frequency domain methods

Frequency domain methods of HRV processing gained wide popularity following the initial work of Sayers (1973) and are now firmly entrenched as a standard tool for HRV processing.

The simplest frequency domain approach is to estimate the spectrum of either the RR interval tachogram or the instantaneous heart rate (obtained by taking the reciprocal of each RR interval). The most immediate difficulty with this approach lies in the fact that these signals are inherently nonuniformly sampled: observations of heart rate are available only each time the heart beats. The processing of nonuniformly sampled signals is generally much more difficult than for their evenly-sampled counterparts. A common first step is to fit a smooth model to the data and resample the signal at a uniform rate. A method of processing the nonuniform signal without such resampling is discussed in Section 2.2.3.3. The RR interval tachogram and the instantaneous heart rate (HR) series are both used in HRV studies. In some cases, the choice between the two may be clear (for example, from model assumptions (Castiglioni, 1995) but often, there is little to guide this choice. Few studies have considered the differences between the two - some notable exceptions are (Castiglioni, 1995; Janssen *et al.*, 1993; Sapoznikov *et al.*, 1992b; Courtemanche *et al.*, 1992). It is also worth noting that the normalisation scheme which is used may reduce the differences between the results of the RR and HR signals (see Section 2.2.3.6).

An alternative approach is to consider the sequence of beat times as the output of an integral pulse frequency modulator (IPFM) (Hyndman and Mohn, 1975). This models the action of the autonomic nervous system (represented by a continuous signal $m(t)$) on heart rate. The sequence of beat times may be written as a sum of delta functions:

$$r(t) = \sum_k \delta(t - t_k) , \tag{2.1}$$

where each t_k is a beat time. The IPFM integrates the input signal $m(t)$ until some threshold is reached, at which time an output spike (beat) is generated and the model reset. The model is characterised by the equation:

$$T = \int_{t_{k-1}}^{t_k} (1 + m(t))\, dt , \tag{2.2}$$

where T is the model threshold and $m(t)$ is the input signal to the model. The spectral structure of $m(t)$ is of interest here. Early work (Bayly, 1968) showed that for a sinusoidal input $m(t)$ it is possible to recover the spectral structure of $m(t)$ by simply low-pass filtering the impulse train signal $r(t)$. However, this is only possible if the modulation depth (the

amplitude of $m(t)$) is small and the modulation frequency (the frequency of $m(t)$) is small compared to the average pulse repetition rate.

Since that work, many methods of estimating $m(t)$ have been proposed, such as the spectrum of counts (de Boer *et al.*, 1985b; de Boer *et al.*, 1984), the spectrum of intervals (Rompelman *et al.*, 1982) and Berger's method (Berger *et al.*, 1986).

The spectrum of counts is obtained by low-pass filtering the event sequence $r(t)$ in order to recover $m(t)$. Unfortunately, for general $m(t)$, this is not necessarily possible. Harmonics will be present in the estimated spectrum, in the form of interference terms between the input frequencies and the free-running frequency of the model (the pulse repetition frequency for $m(t) = 0$) (Nakao *et al.*, 1997). It has been shown that it is not possible to recover $m(t)$ unless the frequencies in the modulating signal are much less than the mean inter-beat frequency (Rompelman *et al.*, 1982; Bayly, 1968). This can be a particular problem if respiration is fast. The spectrum of intervals additionally shows sum and difference terms between the input signals (de Boer *et al.*, 1985b).

Berger's method also employs low-pass filtering, but on the stepwise continuous instantaneous heart rate signal rather than the pulse series itself (Figure 2.4). This can be viewed explicitly as the reconstruction of a signal which is observed at non-uniform sampling intervals. In this case, reconstruction is only possible if the signal is band-limited to half the average sample rate (Strohmer, 1991). Again, this requirement is threatened when respiration is fast.

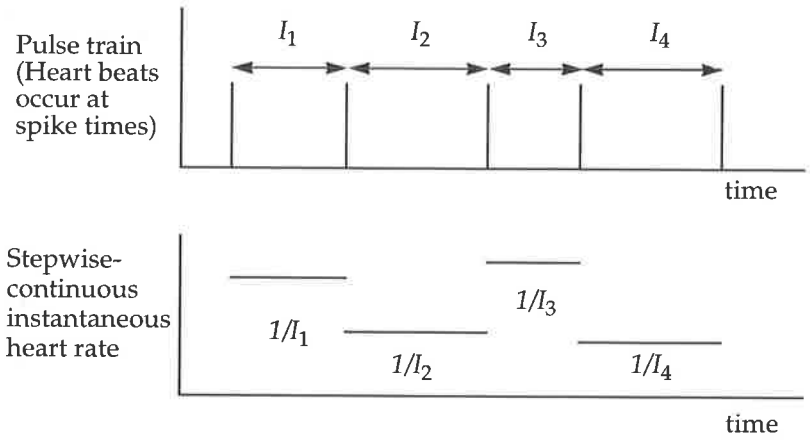

Figure 2.4: The stepwise-continuous instantaneous heart rate is constructed from the reciprocal of the inter-pulse intervals. Berger's algorithm low-pass filters and then uniformly resamples this signal.

The inability to unambiguously recover $m(t)$ is one of the basic limitations of HRV analysis. Work is continuing, and different models of the SA node behaviour may offer some help in this regard (Kitney and Seydnejad, 1996; Sato *et al.*, 1994).

Obtaining a reliable spectral estimate is crucial to any (spectral) HRV analysis procedure. Consequently, the HRV field (as is probably true of many fields which depend heavily on spectral information) is quick to reflect the current fashion in spectral estimation techniques.

Early studies of HRV employed the periodogram or fast Fourier transform (Pomeranz *et al.*, 1985; Akselrod *et al.*, 1981) for power spectral density (PSD) estimation. The requirement for stationary data suggests the use of relatively short data segments; however, the resolution of such methods is inversely proportional to data record length.

2.2.3.1 Autoregressive modelling

Autoregressive (AR) spectral estimation (Lacoss, 1971) has gained wide acceptance in HRV studies (Kamath *et al.*, 1988; Nandagopal *et al.*, 1985; Bartoli *et al.*, 1983). The AR method offers high frequency resolution which is independent of record length. The resulting spectral estimates also tend to be smooth, which is appealing for subsequent manual interpretation (Figure 2.5).

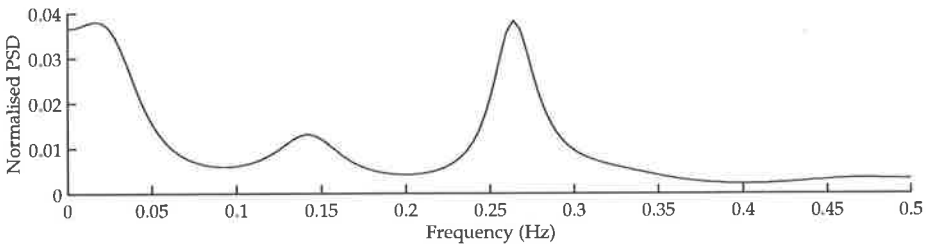

Figure 2.5: Typical autoregressive HRV spectral estimate. A model order of 10 was used here.

The method provides high resolution through an implicit extrapolation of the autocorrelation function, and the resolution (in cycles/sample) has been shown to be (Marple Jr, 1976):

$$\Delta f_{AR} = \frac{1.03}{p\,[\eta\,(p+1)\,]^{0.31}}, \tag{2.3}$$

for $p\eta > 10$, where p is the model order and η is the signal to noise ratio (SNR).

Thus, while the resolution is independent of record length, it is adversely affected by observation noise. In low SNR situations, AR methods may not offer any better resolution than classical methods. Additionally, as record length decreases (with respect to the model order p), the variance of the PSD estimate increases and the spectral estimate begins to show spurious peaks (Kay, 1988).

The method also shows other disadvantages: the PSD peak height is not proportional to sinusoid power; and in the case of multiple sinusoids in white noise, the PSD peaks do not appear at the sinusoid frequencies (Kay, 1988). The statistical properties of the method are only known for true AR processes. Application to non-AR processes results in the loss of spectral detail. The model order must be specified, and an inappropriate choice has a detrimental impact on the resulting PSD estimate. Too low an order results in an overly smoothed estimate, while too high a model order (particularly for short or noisy data) causes the estimate to show spurious peaks. Although various criteria (such as the Akaike information criterion and the final prediction error) have been proposed to assist this choice, the correct model order is not obvious. The correct model order may also change with time, imposing further difficulties on long-term spectral analysis. These difficulties prompted a cautionary note in the Task Force report:

> "The basic disadvantage of parametric methods is the need of verification of the suitability of the chosen model and of ... the order of the model." - Task Force, 1996, p. 1046.

Some effort has been made to determine an empirical method of selecting the appropriate AR model order, specifically in the context of HRV analysis (Pinna *et al.*, 1996).

The saving grace of the AR method is that the model structure allows an extremely convenient decomposition of the PSD, which may be exploited for HRV applications. The AR model can be considered as an all-pole filter driven by white noise, with the Kth order AR PSD estimated from the following expression:

$$X(\omega) = \frac{\sigma^2}{\left|1 + \sum_{k=1}^{K} a(k)\, e^{-j\omega k}\right|^2}, \qquad (2.4)$$

where the $a(k)$ are the model parameters and σ^2 is the variance of the driving noise.

The contribution of each pole or complex pole pair to the spectrum may be isolated and used to evaluate the conventional low- and high-frequency indices of the HRV spectrum (Cerutti, 1995). One must exercise some caution in interpreting the model results, however, since the AR model is "an incomplete and possibly inappropriate model of nonlinear heart rate dynamics" (Christini *et al.*, 1995).

Various comparisons between the FFT and AR methods have been made for HRV studies - for example, (Clayton *et al.*, 1997; Yan and Zheng, 1995).

2.2.3.2 Time-frequency methods

Time-frequency methods describe a signal as a function of both time and frequency, lifting the constraint of stationarity and allowing application to transient signals. These techniques have found widespread application in HRV analysis (Yang and Liao, 1997; Wiklund *et al.*, 1997; Vila *et al.*, 1997; Figliola and Serrano, 1997; van Steenis and Tulen, 1996; Pola *et al.*, 1996; Akay and Mulder, 1996; Tazebay *et al.*, 1995; Chan *et al.*, 1994; Novak and Novak, 1993; Venturi *et al.*, 1990) as well as other biomedical signal processing areas (Bentley *et al.*, 1998; Clayton and Murray, 1998; Sava *et al.*, 1998; Williams, 1996; Bentley *et al.*, 1995; Guo *et al.*, 1994a; Guo *et al.*, 1994b; Wood *et al.*, 1992).

The Gabor transform (Gabor, 1946) was an early attempt to improve the time resolution of the FFT:

$$G(\omega) = \int_{-\infty}^{\infty} x(t)\, g_{\alpha}(t-b)\, e^{-j\omega t} dt, \tag{2.5}$$

where $g_{\alpha}(t)$ is a time-localisation windowing function (originally chosen by Gabor to be a Gaussian function) with width parameter α. The short-time Fourier transform (STFT) is a generalised version of the Gabor transform, with an arbitrary windowing function. The time resolution of the STFT is approximately the length of $g_{\alpha}(t)$, and the well-known trade-off between time- and frequency resolution applies. Successive spectra may also be averaged to reduce variance, with a further degradation in time resolution.

A similar approach may be made with the AR spectral estimate, and this is known as the AR block method. A recursive method may also be used, in which the data are not explicitly windowed, but an exponentially weighted error function is used to emphasise recent changes in the signal (Mainardi *et al.*, 1994; Bianchi *et al.*, 1993):

$$E = \sum_{k=1}^{n} \lambda^{n-k} |e(k)|^2 . \tag{2.6}$$

The "forgetting factor" λ controls the time resolution, which is approximately $-1/\ln(\lambda)$ samples (Pola *et al.*, 1996). Small values of λ may cause instability in the recursive least-squares algorithm, and the method is best suited to slowly varying signals.

The Wigner distribution is perhaps the best-known example of a nonparametric time-frequency distribution. It is a member of Cohen's class (Cohen, 1989) which in the continuous case has the representation:

$$P(t, \omega) = \frac{1}{2\pi} \iiint e^{-j\theta t - j\tau\omega + j\theta u} \varphi(\theta, \tau)\, x\left(u + \frac{\tau}{2}\right) x^*\left(u - \frac{\tau}{2}\right) du\, d\tau\, d\theta , \tag{2.7}$$

where $x(t)$ is the time series, $x^*(t)$ is its complex conjugate, and $\varphi(\theta, \tau)$ is known as the kernel of the distribution (Claasen and Mecklenbrauker, 1980).

The Wigner distribution employs a kernel $\varphi = 1$. It yields excellent time resolution but suffers from several drawbacks. It cannot be interpreted strictly as an energy distribution, as it may take negative values (Cohen, 1989). The Wigner distribution also shows cross-terms in the case of multi-component signals. These cross-terms occur midway between components and may obscure other features of the distribution. Averaging in time may be used to reduce both the cross-terms and negative regions (Cohen, 1989), although this obviously sacrifices time resolution.

The exponential distribution (ED) (Choi and Williams, 1989) was an early attempt at reducing cross-terms and employs a kernel $\varphi(\theta, \tau) = e^{-\theta^2\tau^2/\sigma}$. The kernel gives a relatively large weight to $x\left(u + \frac{\tau}{2}\right) x^*\left(u - \frac{\tau}{2}\right)$ when u is close to t, which emphasizes events near time t (Cohen, 1989). While the ED is effective in reducing the amplitude of cross-terms, the choice of σ is a trade-off between cross-term suppression and time-frequency resolution.

The reduced interference distribution (Williams and Jeong, 1992) describes a class of distributions which embrace and generalise the cross-term suppression of the ED. Reduced interference distributions are generated first by choice of kernel, which is designed to reduce cross-terms and enhance the desirable properties of the signal. The details of this design process are given in (Jeong and Williams, 1992). A discrete realisation of the distribution is

needed for practical application and the binomial and Bessel kernels give reduced interference distributions with convenient discrete implementations. These have been applied to a variety of biomedical signal processing problems (Akay, 1998; Guo *et al.*, 1994b; Wood *et al.*, 1992).

Research into time-frequency distributions is continuing, with strong interest in the development of distributions with reduced interference terms and positive values (Nickel *et al.*, 1998; Sang *et al.*, 1996).

Time-frequency distributions have been applied to HRV analysis (Chan *et al.*, 1997; Pola *et al.*, 1996; van Steenis and Tulen, 1996; Novak and Novak, 1993). Since these methods may be applied to non-stationary signals, they allow the investigation of HRV during periods of physiological transition. This is an important and continuing area of research (Task Force, 1996).

The wavelet transform (WT) is an another time-frequency method. The WT decomposes a signal into dilated and translated versions of an analysing wavelet function (Strang and Nguyen, 1996; Chui, 1992). Various flavours of the WT exist: in the continuous wavelet transform, the time and time-scale parameters are continuous, and a dilated and translated version of a wavelet ψ is written:

$$\psi_{a,b}(t) = |a|^{-1/2}\psi((t-b)/a) . \qquad (2.8)$$

The wavelet series is a sampled set of continuous WT coefficients, so that time remains continuous, but the time-scale parameters are discrete, and most commonly sampled on a dyadic grid in the time-scale plane: $a = 2^j, b = k2^j$ and $j, k \in Z$.

The discrete-time wavelet transform is effectively the continuous WT of a sampled sequence $x(n)$. In the discrete wavelet transform, both time and time-scale parameters are discrete. The discrete WT has various efficient implementations, and may be used as a building block for the fast computation of other wavelet transforms.

The wavelet transform may be used as a time-frequency representation, since the dilation-in-time parameter relates inversely to frequency. While in the STFT, both time and frequency resolution are fixed, the great attraction of the WT is a variable time-frequency resolution,

providing good time resolution at high frequencies and good frequency resolution at low frequencies.

The WT may be viewed in terms of a filter bank, with filters of fixed shape and relative bandwidth (the absolute bandwidth increases with centre frequency). The WT may therefore be used to perform a frequency decomposition of a signal, in a similar manner to the decomposition provided by the AR model (Section 2.2.3.1).

A technique closely related to wavelets is that of matching pursuits, in which a broad library of wavelet functions is used, and each part of the signal is matched to the elementary wavelet best able to represent it (Akay and Mulder, 1996).

Wavelets are becoming increasingly popular for HRV analysis (Toledo *et al.*, 1998; Figliola and Serrano, 1997; Wiklund *et al.*, 1997; Yang and Liao, 1997; Sisli *et al.*, 1996; Wiklund, 1996; Marciano *et al.*, 1995; Tsuji and Mori, 1994).

2.2.3.3 Nonuniform sampling

It is possible to estimate the spectrum of the nonuniformly sampled HRV signal without first resorting to uniform resampling. The Lomb periodogram (Lomb, 1976) is perhaps the best known spectral estimator for nonuniformly sampled signals. Given samples of the signal $x(t)$ at times t_n, the Lomb periodogram may be written:

$$X_s(f) = X(f) \otimes \int_{-\infty}^{\infty} \delta(t - t_n)\, e^{-j2\pi ft} dt . \qquad (2.9)$$

The Lomb estimate is thus the convolution of the underlying spectrum with the Fourier transform of the nonuniform sampling function. This convolution results in a loss of dynamic range in the spectral estimate. The Lomb periodogram has been applied to HRV (Laguna *et al.*, 1998; Laguna *et al.*, 1995), but the reduced dynamic range suggests that this technique is of limited value.

2.2.3.4 Noise considerations

Long-term (24 hour) analysis of HRV has shown a $1/f$ structure to the HRV spectrum (Kobayashi and Musha, 1982). For short-term spectral estimation, this $1/f$ structure is a non-white source of noise which may adversely affect the estimation of narrowband components

within the spectrum (Yamamoto and Hughson, 1991). "Coarse graining spectral analysis" (Hughson *et al.*, 1994; Yamamoto and Hughson, 1991) exploits the fact that the $1/f$ component is scale-invariant, in contrast to the harmonic components present in the signal. Cross-correlation of the HRV signal with rescaled ("coarse-grained") versions of itself allows the $1/f$ component to be removed from the spectral estimate.

Alternatively, a prewhitening filter may be used to transform the $1/f$ component into white noise, followed by spectral estimation and compensation for the prewhitening. This can reduce the severity of the cross-terms in the Wigner distribution (Chan *et al.*, 1997).

The self-similar (scale invariant) nature of $1/f$ processes makes them suitable for analysis by methods aimed at nonlinear dynamics and chaotic behaviour (Hoyer *et al.*, 1997; Vibe and Vesin, 1996; Schmidt and Morfill, 1995; Signorini *et al.*, 1994; Calcagnini Jr *et al.*, 1993; Kaplan and Talajic, 1991).

2.2.3.5 Interpreting heart rate variability spectra

Each HRV spectrum contains a great deal of information, consisting of nominally 128 numbers describing the spectral amplitude in the band 0-0.5 Hz. Some form of information reduction is normally applied to the raw spectral estimates in order to facilitate interpretation by the user. Autoregressive spectral estimation inherently performs information reduction, since it models the signal using a limited number of parameters, and thus gives smooth spectral estimates.

Regardless of the spectral estimation method used, the spectrum is usually condensed into a few key features which have meaning with respect to the underlying autonomic function.

Frequencies above about 0.15 Hz are generally accepted to be purely parasympathetic in origin, but the relative contributions at lower frequencies are more complex. Nevertheless, the most common format for presenting the results of HRV spectral analysis is in terms of the powers contained in fixed bands of the spectrum. The high frequency (HF) power (usually taken as $0.15 - 0.4$ Hz) is accepted as a marker of HR modulation by variations in vagal activity. As this component of the HRV spectrum is a direct result of respiratory activity, various algorithms which employ a respiration reference signal may be used to extract the power which is coherent with respiration (Bianchi *et al.*, 1990). These techniques avoid the

necessity of choosing the fixed HF band, and are generally adaptive, allowing slow changes in the respiratory signal (Varanini *et al.*, 1996; Bianchi *et al.*, 1994).

Both the parasympathetic and sympathetic arms of the autonomic nervous system contribute to HRV frequencies below about 0.15 Hz. There is no accepted marker which correlates purely with sympathetic modulation of heart rate (Akselrod, 1995), although the ratio of LF:HF power (LF band usually taken as 0.05 – 0.15 Hz or similar) is often used for this purpose (Pagani *et al.*, 1986). This ratio is based partly on the largely complementary nature of autonomic control: sympathetic activity is accompanied by vagal withdrawal and vice versa. The LF:HF ratio does not accurately reflect sympathetic tone, nor sympathovagal balance in the general case (Akselrod, 1995).

2.2.3.6 Normalisation of spectral values

Spectral measures of heart rate variability are often expressed in normalised units. The major function of normalisation is to reduce the effect of changes in total HRV power on the band-power (LF and HF) measures.

Under sympathetic stimulation, the total variance in the HR signal is decreased; the opposite occurs under parasympathetic stimulation. The change in total power affects the absolute values of the power in the low- and high-frequency bands of the spectrum (Malliani *et al.*, 1994a), so LF activity may appear to decrease under sympathetic stimulation. A similar decrease in total variance occurs with age (Schwartz *et al.*, 1991; O'Brien *et al.*, 1986; Pagani *et al.*, 1986). Spectra are therefore commonly normalised with respect to total power (VLF components may be excluded from this total) (Pagani *et al.*, 1986).

An alternative normalisation technique is to divide the heart rate variability spectrum by mean heart rate squared (Davidson *et al.*, 1997; Akselrod *et al.*, 1985; Akselrod *et al.*, 1981). This effectively treats the HRV signal as a fractional variation of the mean value and makes the integrated powers in the various spectral bands dimensionless. A side effect of this technique is that it reduces the possible discrepancy between heart rate variability and heart interval variability analysis. This is shown in Appendix A. However, normalising by mean heart rate squared does not account for variations in total power and indeed may exacerbate this problem. Under conditions of increased sympathetic activity, mean heart rate will rise, while total variance decreases (Malliani *et al.*, 1994a). Normalising with respect to an increased mean heart rate will therefore further decrease the total normalised variance.

It should be noted that the commonly used LF/HF ratio is invariant to any normalisation, as the normalisation will affect each band equally and thus be cancelled out.

2.2.4 Multi-signal processing

Variations in heart rate are just one manifestation of the mechanisms which control the cardiovascular system. Blood pressure, respiration, baroreflex action and peripheral resistance are further examples of entities which are varied by the body in an attempt to maintain adequate blood flow under a wide range operating conditions. Two examples of multi-signal processing are given in the following sections. These are the exploitation of a respiratory reference signal to remove the respiratory sinus arrhythmia from heart rate variability, and the combination of heart rate, blood pressure and respiration in a more complete model of the cardiovascular system.

2.2.4.1 Adaptive filtering

Under conditions of stationary respiration, the respiratory component of HRV is usually easily removed using simple bandstop filtering techniques (Womack, 1971). Under normal conditions, however, respiration is rarely stationary for extended periods of time. Paced, or "metronome" breathing may help, but may also cause a stress reaction if the paced frequency is significantly different from the subject's natural respiration rate (Haaksma *et al.*, 1996; King and Reisman, 1995). Much interest has been shown in adaptive filtering techniques, as the removal of the RSA component from the HRV signal is a textbook example of the application of adaptive filtering for noise cancellation.

Separation of the respiratory influences from the remainder of the HRV offers several advantages. Since high frequency activity (>0.15 Hz) is widely considered to be caused by variations in vagal tone linked to respiration rate, the RSA component can give a direct window on this influence. Variations in respiratory tidal volume can also contribute to HRV at lower frequencies (Niccolai *et al.*, 1994; Brown *et al.*, 1993). Removing this influence allows better estimation of other contributions to the HRV signal (Varanini *et al.*, 1996).

In the classical noise cancellation scenario (Widrow *et al.*, 1975), the signal of interest (HRV) is assumed to be corrupted by additive noise (RSA). A reference signal which is correlated with the corrupting noise is assumed to be available. A filter ($H(z)$ in Figure 2.6) is designed

to minimise the squared "error" output $\hat{h}(n)$, thus removing the noise from the signal of interest.

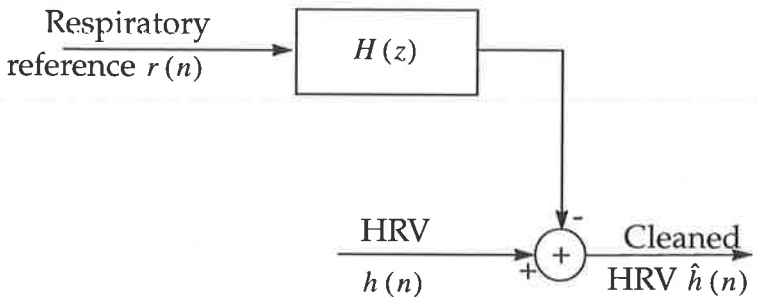

Figure 2.6: Respiratory noise cancellation scheme

The filter $H(z)$ is typically implemented using a tapped-delay or lattice filter, which may be employed recursively for on-line filtering (Bianchi *et al.*, 1994; Wiklund *et al.*, 1991). Time-frequency techniques may also be used to estimate the respiration rate over short data segments (Tazebay *et al.*, 1995), allowing an appropriate filter to be constructed for that segment.

The respiratory reference signal may be acquired from a nasal thermistor or from chest movements. If the cost or inconvenience of the extra sensor is inappropriate for a particular application, or if the external respiration signal becomes corrupted, the respiration may be derived from the ECG signal itself (Travaglini *et al.*, 1998; Zhao *et al.*, 1994). Lung inflation during respiration causes the orientation of the heart to change within the chest, causing variations in the baseline of the ECG which are correlated with respiration. No examples of the use of this technique with adaptive filtering were found in the literature.

2.2.4.2 Multivariate modelling

The natural extension to the incorporation of respiratory information is that of a multivariate model of the cardiovascular system, commonly including heart rate, blood pressure and respiratory interactions. Modelling in this manner allows investigation of the system in a manner which would not be possible through physical experiments. The mechanisms of the low-frequency (0.1 Hz) rhythm seen in heart rate and blood pressure have

been examined and modelled in this manner (Baselli *et al.*, 1994; Baselli *et al.*, 1988b; de Boer *et al.*, 1987).

There are two main approaches to such modelling: one approach is to base the model structure and functional equations on prior experimental evidence (Madwed *et al.*, 1989; Saul *et al.*, 1989; de Boer *et al.*, 1987; de Boer *et al.*, 1985a); conversely, one may employ an extremely general mathematical model and fit this to the experimental data. The former approach simplifies the computational aspects but requires more prior knowledge and assumptions in model construction. An example is shown in Figure 2.7, in which the baroreceptors, the action of the baroreflex on heart rate and peripheral resistance, the properties of the myocardium and the Windkessel effect are modelled.

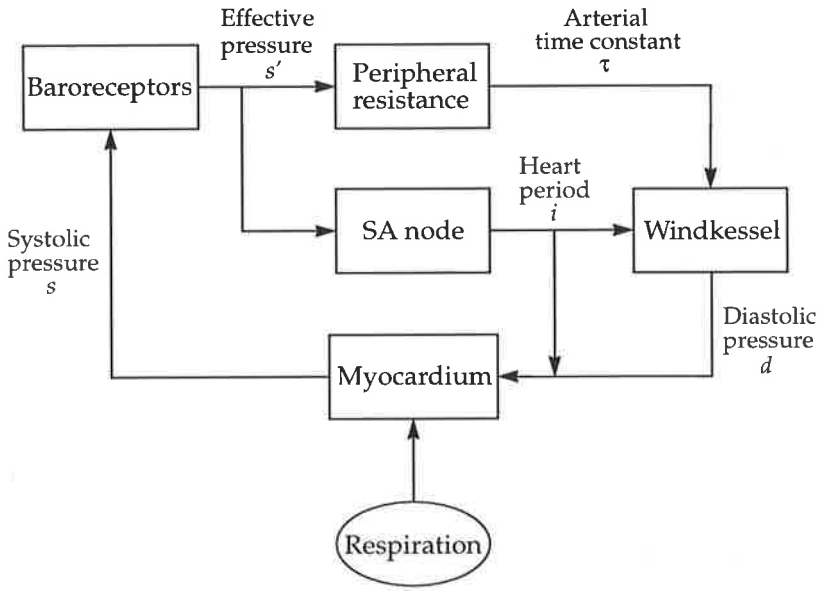

Figure 2.7: Model of the cardiovascular system. Experimental results are used to construct equations to represent the effects of each block in the model. (Adapted from de Boer *et al.*, 1987).

Each block in the model is represented by a relatively simple difference equation. For example, the effect of systolic blood pressure on heart period is modelled by (de Boer *et al.*, 1987):

$$i_n = \alpha_0 s'_n + \sum_{k>0} \alpha_k s'_{n-k}. \qquad (2.10)$$

Here, i_n is the deviation of the heart period (from its mean value) at the time of sample n, s'_n is the deviation from the mean of the effective systolic blood pressure, α_0 is the strength of the (vagally mediated) baroreflex action, and α_k $(k>0)$ are the coefficients of the sympathetically mediated control of heart period.

Model parameters are generally taken to be typical of the values seen in experimental evidence (de Boer *et al.*, 1985a), although this approach presents some difficulties:

> "A problem in modelling the cardiovascular system is that reliable values for the different parameters of the model ... are scarcely found in the literature." - de Boer *et al.*, 1987

The alternative approach is to incorporate relatively little prior knowledge into the model, employing a general mathematical form, such as the multivariate autoregressive model depicted in Figure 2.8 (Korhonen *et al.*, 1996a; Korhonen *et al.*, 1996b; Korhonen *et al.*, 1996c; Baselli *et al.*, 1994; Kalli *et al.*, 1988; Baselli *et al.*, 1988a; Baselli *et al.*, 1988b; Akselrod *et al.*, 1985).

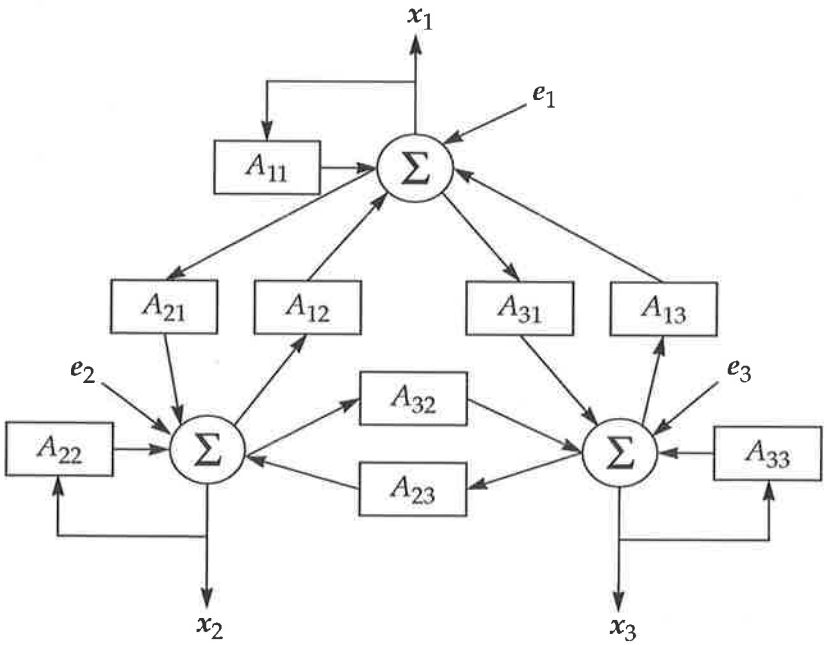

Figure 2.8: Multivariate model structure. x_1, x_2 and x_3 are the observed variables (respiration, heart rate and blood pressure). e_1, e_2 and e_3 are white noise sources. The parameters of each functional block (A_{11}, etc.) are estimated by fitting the model to experimental data. (Adapted from Korhonen *et al.*, 1996c).

The system equation for a multivariate AR model describes each variable as a linear combination of past values of all variables:

$$\boldsymbol{x}(k) = -\sum_{i=0}^{M} \boldsymbol{A}(i)\,\boldsymbol{x}(k-i) + \boldsymbol{e}_i(k) \,. \tag{2.11}$$

Here, $\boldsymbol{x}(k)$ is the three-element vector describing blood pressure, heart rate and respiratory measurements at sample k, $\boldsymbol{e}_i(k)$ is the model error and $\boldsymbol{A}(i)$ is the model coefficient matrix for lag i. The method requires little prior information, although it is sensible to disallow non-physiological mechanisms in the model (e.g. heart rate does not directly affect respiration). The model parameters may be estimated using well known methods, such as the Levinson algorithm. The method suffers similar difficulties in model order selection to the AR method for spectral estimation.

One common property of the preceding two approaches is the fact that the models are linear (de Boer's model incorporates a nonlinear sigmoid element in the baroreceptors, but is otherwise linear). Thus, these approaches suffer from an inability to describe nonlinear behaviour of the cardiovascular system. For example, for respiration at frequencies near 0.1 Hz, the low-frequency (0.1 Hz) component of HRV can become entrained with respiration (Kitney *et al.*, 1985). Entrainment causes the LF component of HRV to follow the pattern of respiration. A similar entrainment can be seen for periodic thermal stimuli (Kitney, 1975; Hyndman *et al.*, 1971).

Increasing attention is being focused on nonlinear modelling, including neural network models (Chon *et al.*, 1997; Chon *et al.*, 1996; Chon *et al.*, 1995; Vallverdú *et al.*, 1991; Saul *et al.*, 1988b). Such models are often able to account for a greater fraction of the variability seen in experimental data (Chon *et al.*, 1997), but may also be more difficult to interpret than linear models.

2.2.5 Conclusions

Techniques for processing the HRV signal have developed rapidly over the last two decades. Our understanding of the underlying physiology is becoming more refined and at the same time, computational power is becoming cheaper by the day. More sophisticated processing methods can be realised within reasonable time and cost constraints.

The early processing methods (basic spectral- and time-domain techniques) suffer from limitations which deflate the potential of HRV as a clinical tool. Modern methods of multivariate modelling and nonlinear analysis are helping to further clarify our understanding of the cardiac control mechanisms, as well as offering wider clinical applications.

2.3 Clinical applications of heart rate variability analysis

The noninvasive nature of heart rate variability analysis has generated strong interest in its use as a probe of autonomic function. The first clinical application of HRV analysis was in 1965, when foetal distress was found to correlate with a decreased variance of heart rate (Hon and Lee, 1965). Since then, HRV analysis has been employed to detect abnormalities of the autonomic system associated with a variety of conditions.

> "The enormous advantage of analyzing these fluctuations is that the information can be provided on-line, noninvasively and without interfering with the normal dynamic functioning of the control mechanisms" - Akselrod, 1995, p.147

2.3.1 Risk assessment for myocardial infarction patients

In 1978, it was documented that patients who had suffered an acute myocardial infarction (MI) and who lacked respiratory sinus arrhythmia were at a higher risk of in-hospital mortality than patients with RSA (Wolf *et al.*, 1978). An extended study (Kleiger *et al.*, 1987) concluded that a depressed variance of heart rate correlated with increased mortality in the first four years following MI. More advanced HRV analysis techniques have revealed reduced parasympathetic activity in patients with recent MI (Bigger Jr *et al.*, 1988; Craelius *et al.*, 1992).

While the causal mechanisms linking depressed HRV and cardiac mortality are still the subject of research, HRV analysis remains a valuable tool in risk assessment following MI.

More comprehensive methods have been proposed in the form of multivariate methods. Heart rate variability information in combination with other cardiovascular variables, such as the frequency of extrasystoles and late potentials in the ECG, has been reported to provide

more consistent prediction of post-MI mortality than the use of HRV alone (Task Force, 1996; Malik and Camm, 1994).

2.3.2 Hypertension

Hypertension is abnormally elevated blood pressure, a condition which carries with it increased risk of stroke and coronary heart disease (McMahon *et al.*, 1990) and myocardial infarction (Rakugi *et al.*, 1996). The World Health Organisation defines borderline hypertension as a resting diastolic blood pressure (DBP) greater than 90 mmHg and less than 95 mmHg, and systolic blood pressure (SBP) between 140 and 160 mmHg. Developed hypertension is defined by DBP greater than 95 mmHg and SBP greater than 160 mmHg. Blood pressure generally increases with age, and elderly people with hypertension are also at greater risk of cardiovascular disease (Kaplan, 1994). Hypertension is therefore likely to become an increasing problem as the world's population ages.

Measuring blood pressure is the most obvious and widely used test for hypertension, but may give misleading results due to, amongst many other factors, day-to-day variability of blood pressure (Pickering, 1993). Autonomic changes may offer additional information: a number of independent studies have shown an increased sympathetic tone in subjects with marginally elevated blood pressure (Anderson *et al.*, 1989; Guzzetti *et al.*, 1988). If the borderline hypertension progresses to established hypertension, the sympathetic drive is reduced, although blood pressure remains elevated (Julius, 1991).

Studies in humans (Malliani *et al.*, 1991b; Guzzetti *et al.*, 1988; Pagani *et al.*, 1984) found that mildly hypertensive subjects displayed increased low-frequency and decreased high-frequency components in the resting HRV spectrum when compared to controls. In addition, hypertensive subjects displayed a less marked response to tilt (from supine to standing posture). These results are consistent with findings using independent methods, such as increased sympathetic drive seen in direct neural recordings (Anderson *et al.*, 1989).

The progression from borderline to developed hypertension is not well understood. Only 30% of young subjects with borderline hypertension go on to develop permanent hypertension (Julius and Schork, 1971). Nevertheless, early detection of hypertension is of great interest. Studies with spontaneously hypertensive rats (Akselrod *et al.*, 1987) showed that the low-frequency oscillations in arterial blood pressure are reduced in hypertension, and that this reduction occurs prior to the onset of clinical signs of hypertension. The fact that this

study showed a reduction rather than an increase in LF activity is probably due to the normalisation applied to the spectral values (see Section 2.2.3.6); the interesting result is the precursory nature of the autonomic dysfunction. The detection of hypertension before the onset of increased blood pressure would be a valuable step toward efficient early treatment. Heart rate variability offers great promise for this purpose since it is a cheap, non-invasive procedure.

2.3.3 Diabetic autonomic neuropathy

Autonomic dysfunction is a common complication of diabetes, and can lead to silent myocardial ischemia, sudden cardiac death, renal disease and anaesthetic complications (Freeman *et al.*, 1991). Early non-invasive investigations of autonomic function in diabetic subjects employed relatively simple methods such as comparisons of RR interval length (Baldwa and Ewing, 1977; Bennet *et al.*, 1977). Even so, the information provided by such tests prompted the suggestion that invasive methods of assessing autonomic function in diabetes were no longer justified (Bennet *et al.*, 1978).

Spectral analysis of HRV, as a probe of autonomic function, is thus a logical tool for the investigation of this condition. In an early study, the centre frequency of the low-frequency peak (nominally 0.1 Hz) was observed to differ in diabetic subjects with autonomic neuropathy when compared to normal subjects (van den Akker *et al.*, 1983). This change in centre frequency was ascribed to the neural degradation caused by the neuropathy.

Later studies were generally consistent in finding a decline in power across all frequencies of the HRV spectrum, and less change in the HRV spectrum in response to a change in posture from supine to standing when compared to diabetics without autonomic dysfunction (Bellavere *et al.*, 1992; Freeman *et al.*, 1991; Bianchi *et al.*, 1990; Lishner *et al.*, 1987).

2.3.4 Other applications

As a probe of autonomic function, HRV analysis has found a number of other applications, including: investigating chronic heart failure (Saul *et al.*, 1988a), alcoholism (Duncan *et al.*, 1980), multiple sclerosis (Neubauer and Gundersen, 1978), glaucoma (Clark and Mapstone, 1985), sudden infant death syndrome (Gordon *et al.*, 1984; Leistner *et al.*, 1980), foetal distress (Bartlett *et al.*, 1991), and mental stress (Pagani *et al.*, 1995; Singh *et al.*, 1993; Watson *et al.*, 1993; Nandagopal *et al.*, 1985; Hyndman and Gregory, 1975; Danev and

de Winter, 1971); monitoring fatigue in vehicle drivers (Myrtek *et al.*, 1994; Egelund, 1982) and pilots (Jorna, 1993); guiding the thermal management of low birth weight infants (Davidson *et al.*, 1997); and as a tool for managing optimal drug dosage levels (Brouwer *et al.*, 1995; Alcalay *et al.*, 1992).

2.3.5 Conclusions

It is clear that, as a probe of autonomic function, HRV analysis has broad clinical application. However, HRV analysis is subject to various limitations which restrict its utility. In Section 2.4, two fundamental limitations of HRV analysis are discussed.

2.4 Fundamental limitations of heart rate variability analysis

Spectral techniques have already been shown to be unable to unambiguously recover information regarding autonomic activity from the HRV signal (Section 2.2.3). This can be considered to be a limitation of spectral processing methods rather than a fundamental limitation of HRV analysis in general. In this section, two limitations are discussed which are more basic in nature, and apply to all HRV analysis techniques.

2.4.1 As a probe of autonomic tone

The term *tone* describes the average density of traffic in a given nerve fibre or branch: autonomic tone is the average activity in the sympathetic and parasympathetic branches of the autonomic nervous system. While it is true that HRV analysis may be used as a probe of autonomic function, it cannot be said that HRV provides information relating directly to autonomic tone. Rather, it is *variations* in tone which cause variations in heart rate. A depiction of the relationship between tone and HRV is given in Figure 2.9.

For very low levels of nerve activity (Figure 2.9(a)), both the modulation depth (variation in tone) and tone are low. Increasing parasympathetic stimulus leads to Figure 2.9(b) - moderate tone, but with strong modulation, causing a large peak in the HRV spectrum. As the stimulus strengthens, the tone continues to increase, but modulation depth decreases as the nerve activity saturates (Figure 2.9(c)). Consistent with this, the LF/HF ratio has been

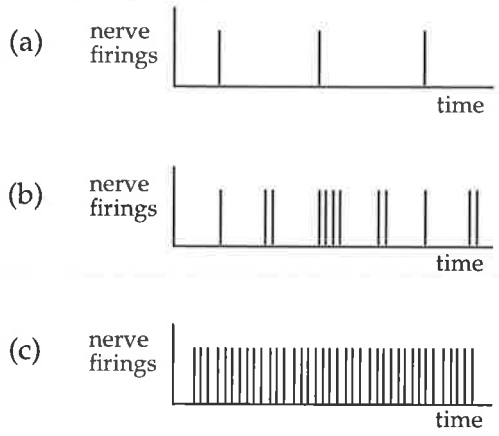

Figure 2.9: Depiction of three levels of nerve activity - (a) low tone, low variation; (b) moderate tone, high variation; (c) saturated tone, low variation. (Adapted from Malik and Camm, 1993)

observed to increase with moderate exercise, before decreasing at maximal exercise (Arai *et al.*, 1989).

This limitation of HRV well known:

> "The exact relation between the level of neural fluctuations and the mean level of neural activity is far from being known." - Akselrod, 1995, p. 159.

However, it is commonly assumed that tone is linearly related to modulation depth (at least over a limited operating range) (Akselrod, 1995). In general studies, there may be a poor correlation between tone and HRV indices (Bootsma *et al.*, 1994).

One avenue for extending the utility of HRV analysis is to combine information relating to mean heart rate with HRV (Malik and Camm, 1993; Chess *et al.*, 1975). Mean heart rate is a reflection of average autonomic activity, and is therefore an indicator of autonomic tone. On its own,

> "mean heart rate is subject to too many other control mechanisms and pathologic phenomena, and cannot be used as a reliable estimator of autonomic activity and tone." - Malik and Camm, 1993.

A combination of the two may offer improved performance over either in isolation. Experimental evidence has shown a correlation with mean HR and HRV (Coumel *et al.*, 1995;

Ungi *et al.*, 1995). Heart rate variability could alternatively be corroborated against baroreflex sensitivity, which is also related to autonomic activity (Hyndman *et al.*, 1997).

2.4.2 Arrhythmias

Disturbances of the normal heart rhythm may provide valuable information to a cardiologist but pose some difficulties for heart rate variability analysis. Since the heart is not under normal control during arrhythmias, such data is generally unsuitable for HRV analysis.

Under normal conditions, the pacemaker for the heart is the sinoatrial (SA) node. Each action potential generated by the SA node propagates through the right and left atria and on to the atrio-ventricular (AV) node. The AV node delays further propagation of the impulse for a short period to allow the blood within the atria to be expelled into the ventricles before ventricular contraction occurs. The impulse then travels along the Bundle of His, which subdivides into the left and right bundle branches, and ultimately into a complex network of conducting fibres known as the Purkinje fibres. Through these fibres, the impulse spreads across the ventricles causing contraction and the expulsion of the blood within.

If the pacemaker impulse originates in a site other than the SA node (cells within the atria, AV node or ventricles are capable of initiating these impulses) then the result is an ectopic beat. Isolated ectopic beats are quite common and may be caused by emotional stress, stimulants such as caffeine, or physical abnormalities of the myocardium (Little, 1985).

If the SA node resumes control of the heart following an ectopic beat, and the ectopic beat has not caused any alteration to the rhythm of the sinus node, then the ectopic can be removed by interpolation and the data may be used for HRV analysis as normal (Vybiral *et al.*, 1990).

An ectopic focus in the atrium causes the SA node to depolarise before it would normally reach threshold. Consequently, when the SA node resumes the pacemaking function following an atrial ectopic beat, the cardiac rhythm is slightly advanced in time. An ectopic focus in the AV node may or may not disturb the depolarisation of the SA node cells, while ventricular ectopic beats have an ectopic focus which is sufficiently removed from the SA node so as not to disturb the normal sinus rhythm. Ectopic beats which do not alter the phase of the cardiac rhythm will typically be manifest as a short interbeat interval (the heart will beat earlier than expected due to the ectopic) followed by a compensatory long interval. This

causes an impulse-like artifact in the RR interval tachogram and broadband noise in the spectral estimate if not removed.

An *ectopic rhythm* occurs when the pacemaker function is assumed by an ectopic focus for two or more beats. In this case, the heart is not under the normal control of the autonomic nervous system.

Arrhythmias may also result from alterations in the normal conduction pathway of the depolarisation impulse. Impediment or prevention of the transmission through the AV node is referred to as AV block. In its least severe form (first-degree AV block) this causes an unusually long delay in transmission of the impulse. The most severe form (third-degree AV block) prevents the impulse from travelling from the atria to the ventricles. In the latter case, the SA node acts as the pacemaker for the atria, while an ectopic pacemaker generates an independent rhythm for the ventricles (Berne and Levy, 1986). For AV block of any kind, heart rate records derived from the R-wave of the ECG are unlikely to provide an accurate representation of the SA node firing rate.

Arrhythmia-like data may also result from abnormally large P- or T-waves in the ECG, which may be mistaken for R-waves.

With the exception of isolated ectopic beats which do not alter the SA node rhythm, data which contains arrhythmias is unsuitable for HRV analysis and should be discarded. Various methods for the detection of arrhythmias are available (Liao *et al.*, 1996; Kamath and Fallen, 1995; Sapoznikov *et al.*, 1992a).

2.5 Research issues in heart rate variability processing

This thesis aims to identify and overcome existing deficiencies in the HRV information processing chain. These stem from both the fundamental limitations of HRV analysis (Section 2.4) as well as limitations of particular analysis methodologies.

One of the major shortcomings of spectral HRV analysis methods lies in the interpretation of the results. Examining fixed-band powers within the HRV spectrum is a simple and almost universally-accepted method of quantifying results of HRV analysis. However, the band-

powers themselves do not have an intuitive physiological interpretation in terms of autonomic tone. As detailed in Section 2.4.1, HRV power reflects the strength of modulation of autonomic tone, rather than tone itself. Band-power analysis also requires the choice of the band limits, for which different researchers argue different values (e.g. Takalo *et al.*, 1994; Malliani *et al.*, 1994b), making comparison of results difficult. To this end, the Task Force report (1996) suggests standardised band limits (VLF band ≤0.04 Hz, LF band 0.04-0.15 Hz, HF band 0.15-0.4 Hz), but even this is not a universal solution since some applications may be better suited to fixed bands covering other regions (e.g. Hilton *et al.*, 1998).

Fixed-band analysis is also insensitive to changes in the LF peak frequency, identified below as a potentially valuable source of information regarding autonomic tone.

The extraction of concise, relevant descriptors of the HRV signal can be considered as a form of dimension reduction or data visualisation. Many general-purpose algorithms exist for the task of data visualisation, including principal components analysis and multidimensional scaling. These algorithms can be applied to a variety of data formats, which raises the primary research question addressed in this thesis:

Are general-purpose dimension reduction and visualisation techniques of value for extracting information from the HRV signal?

This question is addressed in Chapters 3 through 5.

Another significant deficiency of current HRV analysis techniques stems from the fact that the amplitudes of the HRV components do not directly offer information regarding autonomic tone. Rectifying this deficiency would be an extremely valuable contribution, since tone is a valuable marker for evaluating the behaviour of the autonomic system.

Spectral analysis of HRV need not be restricted to the analysis of the spectral amplitudes. An intriguing possibility for improving the information available from HRV analysis stems from the low-frequency (so-called 0.1 Hz) peak in the HRV spectrum. While it is nominally found at about 0.1 Hz, the precise location has been found to vary with physiological changes. Little is known about the significance of this shift. A notable exception is the work conducted by van den Akker *et al.* (1983) into the effects of diabetic autonomic neuropathy on the autonomic nervous system. Autonomic neuropathy is a complication of diabetes in which the autonomic nervous system becomes affected. A significant decrease in the LF peak frequency

has been observed in diabetic patients with autonomic neuropathy, and this was shown to be consistent with changes in the delays in the blood pressure control loop which occur as the sympathetic nerve fibres become damaged with neuropathy (van den Akker *et al.*, 1983). The LF peak frequency has also been observed to increase during sympathetic stimulation (Raymond *et al.*, 1997; Baselli *et al.*, 1991; Kamath *et al.*, 1991; Nandagopal *et al.*, 1985) and decrease during parasympathetic stimulation (Kamath *et al.*, 1991; Weise *et al.*, 1989).

One may speculate about the relationship between the LF peak frequency and autonomic tone. The frequency of this peak is widely thought to be set by the resonant frequency of the blood pressure feedback loop, and so the changes in the LF peak frequency mentioned above may be caused by alterations in the autonomic balance. As the balance shifts towards sympathetic dominance, the baroreflex control of the SA node becomes increasingly dominated by the slow sympathetic efferent activity. Thus, the decrease in the LF peak frequency observed with sympathetic stimulation may be a consequence of this slowing of the feedback loop. If this is indeed the case, the LF peak frequency may be of clinical value as a marker of autonomic balance. This approach may even avoid the problem encountered with spectral amplitudes, which give misleading results when autonomic tone becomes saturated.

This hypothesis requires physiological validation, which is beyond the scope of this thesis. Nevertheless, throughout this work, the LF peak frequency will be considered as an important characteristic of the HRV spectrum.

Another method of obtaining a more reliable estimate of autonomic tone from HRV analysis is to combine the information from HRV analysis with information from other signals and sources. Mean heart rate has been suggested as a suitable signal with which HRV information may be combined, since mean heart rate is itself an indicator of autonomic tone (Malik and Camm, 1993). This is not the only possible combination of variables: baroreflex sensitivity has also been identified as a marker of autonomic activity (Hyndman *et al.*, 1997; Clayton *et al.*, 1995) and, like heart rate information, may be acquired non-invasively. The fusion of HRV information with other variables may also be of value in other situations; for example, the combination of HRV information with patient history may offer improved diagnostic performance compared to the use of either information source in isolation.

The fusion of HRV information with other variables may be approached from a modelling viewpoint: a model is postulated, which explains the physiological coupling between the variables of interest. This model is then fitted to experimentally-obtained data, and the

parameters of the model examined in order to ascertain the physiological significance of the data. This is a similar approach to the multivariate modelling methods discussed in Section 2.2.4.

An alternative approach to information fusion is that of classification: information from a variety of sources is combined with the objective of maximising the correct classification of data into predefined classes. This approach avoids the necessity of constructing a physiologically-plausible model of the interactions between the variables, which may be a difficult task in itself.

The classification approach to information fusion has shown merit in the assessment of the likelihood of mortality following myocardial infarction, in which HRV information can be combined with other risk indicators (Malik and Camm, 1994; Bigger Jr *et al.*, 1992b). Similarly, the combination of age with HRV information may assist the prediction of susceptibility to sudden cardiac death (Yarnold *et al.*, 1994), and combining respiratory and HRV information can aid in the prediction of infant state (Stevens *et al.*, 1988; Harper *et al.*, 1987).

However, many published studies use relatively simple methods for information fusion, such as linear discriminant analysis (Severi *et al.*, 1997; Yarnold *et al.*, 1994; Harper *et al.*, 1987). Methods which are better able to deal with non-normal or nonlinear data, such as neural networks (Wiklund, 1996) seem to have been sparsely applied to HRV analysis.

Two research questions therefore arise:

Can information fusion help improve the utility of HRV analysis?

and:

Are simple classification techniques sufficient for this task?

A preliminary investigation into the classification of heart rate variability in normal and mildly hypertensive subjects is described in Appendix E, and a more rigorous treatment is given in Chapter 6.

2.6 Summary

A summary of current HRV analysis techniques has been given in this chapter. The motivation for this work is to identify and overcome deficiencies in existing analysis methods, and it has been found that the principal deficiencies arise from the fact that HRV does not have a direct, simple correlation with autonomic tone. This causes difficulties in the interpretation of the results of HRV analysis, particularly for fixed-band spectral analysis.

The use of general-purpose data visualisation techniques for the interpretation of HRV data is investigated in the next few chapters of this thesis. Three dimension-reduction and data-visualisation techniques (multidimensional scaling, the generative topographic mapping and the self-organising map) are introduced in the next chapter and then applied to HRV data in Chapter 4.

The fact that HRV does not have a simple interpretation in terms of autonomic tone also complicates the use of HRV as a clinical diagnostic aid. The combination of HRV information with other physiological information is considered in Chapter 6 as a means of improving the reliability of HRV analysis.

Chapter 3 Data visualisation

Wheels within wheels in a spiral array
A pattern so grand and complex
Time after time we lose sight of the way
Our causes can't see their effects.

- Neil Peart, "Natural Science"

The investigation of large, high-dimensional data sets can be an intimidating task. Data visualisation techniques generally aim to extract a few salient features from a data set, which show the structure and relationships of the objects within the set. The major task of visualisation is dimension reduction, most commonly down to two dimensions to allow simple graphical depiction. This dimension reduction is carried out under certain criteria which ensure that the structure of the data is maintained during the transformation.

Visualisation techniques are becoming increasingly popular for the processing of biomedical data. An expanding and aging population, in addition to society's expectations of consistently high standard of health care, place increasing demands on biomedical engineering for reliable, cost-effective methods of assisting or automating many common medical tasks. The need to be able to process large volumes of data with a high degree of reliability makes manual methods for many procedures unrealistic. There are strong incentives to develop computer assistance for tedious tasks, particularly in areas which require specialist training, such as classification of mammograms (Sauerbrei *et al.*, 1998) or ultrasonic prostate images (Huynen *et al.*, 1994). Unfortunately, by nature, many of these tasks are not well suited to computer analysis and may also require significant expert knowledge which must be incorporated into the classifier. Reliable automation may thus be very difficult to achieve.

A compromise solution is to design systems which assist, but do not replace, the human specialist. Such a system might perform sub-tasks for which it is well suited, particularly if these tasks would be tedious for a human operator - such as object delineation in medical images (Tsapatsoulis *et al.*, 1997; Schnorrenberg *et al.*, 1994; Maes *et al.*, 1993). Alternatively,

the system might be semi-autonomous, allowing user input to direct or fine-tune the activity. An example is a measurement system for the evaluation of cardiac autonomic neuropathy (Vespasiani *et al.*, 1996), in which the clinician is guided through the administration of a standard set of tests, and then the results are collected and processed. A semi-automated process such as this is likely to be more time-efficient and also improve the consistency of measurements and results.

Visualisation techniques also fit into this "compromise" category. Rather than make an explicit classification of data, the role of the visualisation algorithm is as an aid to decision making, by providing the user with a clear depiction of hard-to-visualise data. An information processing paradigm which utilises visualisation techniques relies on the human operator as the pattern recognition and classification machine and so in this approach, the specialist domain knowledge of the user is retained in the information processing chain without the need to explicitly include or learn this expert knowledge in the algorithm.

Humans have excellent clustering and pattern recognition skills in two dimensions (Daoudi *et al.*, 1993; Siedlecki *et al.*, 1988b) and a two-dimensional graphical display (often termed a "mapping" of the data) is the most common format for the presentation of information to the user. More exotic methods of presenting data to the user include stereoscopic, three-dimensional images (Tsai *et al.*, 1998) and aural presentation (Jovanov *et al.*, 1999).

Visualisation algorithms can help the clinician in a variety of ways in the management and efficient use of diagnostic databases. Visualisation techniques may be used to explore the structure of large, unwieldy data sets (Rosario *et al.*, 1997), showing the relationships between, for example, EEG signals during different tasks (Lowe, 1997), respiratory sounds (Çagatay Güler *et al.*, 1998), or magnetic resonance imaging data (Pearlman *et al.*, 1995).

Visualisation techniques which can also generalise - that is, project data which was not used during the initial map construction - allow the clinician to relate an unknown data record to previously gathered, known, records (Ornes and Sklansky, 1997). This also allows the monitoring of the on-going condition of a patient by comparing data records in time (Lehtinen *et al.*, 1997).

The principal difference between the tasks of data visualisation and of feature extraction for classification should be clarified here. General purpose algorithms designed to perform

feature extraction for classification exist (e.g. Sherrah, 1998), but these operate with a different objective than do visualisation techniques. The purpose of the former is to extract features which explicitly maximise the performance during subsequent classification, and usually operate with more than two features. Data visualisation attempts to provide a useful depiction of the data to the user, using only two-dimensional displays.

In the next section, three visualisation algorithms are briefly reviewed.

3.1 Visualisation algorithms

An astonishing variety of algorithms exist for dimension reduction and visualisation tasks. In order to facilitate the choice of an appropriate algorithm, some desirable attributes for such algorithms are identified.

Inductive

An extremely valuable property is that of generalisation. We first construct a mapping from a set of training examples. Generalisation is the ability to project new data (which was not part of the training set) onto the mapping. The projection of new data onto a mapping allows the comparison of unknown data with known samples from a database and in the context of biomedical information processing, this has direct merit as a diagnostic aid.

Interactive

Interaction is an essential part of any practical visualisation scheme. In many situations, the data to be visualised is non-metric, from a nonlinear manifold or it displays structure at multiple scales. In these cases, a single two-dimensional representation is unlikely to describe adequately all inter-object relationships within a data set. Interaction with the user may be used to direct attention to particular regions of the mapping, in order to improve the perception of the interesting or important parts of the data set (Lee, 1998; Daoudi *et al.*, 1993).

Hierarchical clustering and visualisation methods (Coxon, 1982; Friedman and Tukey, 1974; Maltson and Dammann, 1965) offer an excellent avenue for interactivity: at the top level of the hierarchy, the entire data set is described by a small number of clusters, representing the coarse structure. At lower levels of the hierarchy, the clusters become increasingly finely partitioned, revealing local, small-scale structure of the data. Such

hierarchies may be constructed by starting either with a small number of clusters, and progressively splitting them, or with a large number of clusters (one per data point, at the extreme) and merging neighbouring clusters together. The mapping process may be made interactive by allowing the user to manually select the clusters to merge or split.

Other forms of interaction are possible depending on the formulation of the algorithm in question.

Flexible

Biomedical data can take a variety of formats, such as images, time- and frequency-domain signals and patient details such as age and prior history. Thus, the algorithm may need to marry disparate data types - for example, boolean (is the subject a smoker?) with image data.

Quantifiable

The question "how do we know when we have done a good job?" is one which cannot go unanswered. There are two issues to consider when assessing a mapping: how well the map points represent the original data, and the quality of the map as a perceptual representation of the data.

A map which poorly represents the data is unlikely to provide much useful information to the user and conclusions drawn from such a mapping may well be erroneous. If the visualisation algorithm operates by optimising some form of cost function, then the cost function value may be used to quantify the fit of the map to the data. Two mappings generated using the same cost function and same data set may be compared on the basis of their cost function values. This does not extend to the comparison of mappings generated using different algorithms, since one specifically optimises a given cost function while the other does not. Some general measures of topological preservation have been proposed (Goodhill and Sejnowski, 1997; Goodhill and Sejnowski, 1996), but it is important to recognise that particular techniques adopt different notions of structure preservation (e.g. local vs. global structure) and this affects the resulting mapping.

If class labels are available for the data, it is reasonable to suggest that a configuration which provides a good description of the objects would also perform well as a feature vector for subsequent classification. Given the same data, mappings may therefore be compared on the basis of classification performance (Lehtinen *et al.*, 1997; Joutsiniemi *et al.*, 1995). If the

algorithm is amenable to generalisation, then partitioning of the data into training and test sets could be done. Otherwise, areas of the map may be labelled by majority voting of the samples; each projected sample is then classified according to the label of the area into which it is mapped. It is emphasised that the objective of the mapping is not to perform feature extraction for subsequent classification. Two dimensional mappings generally do not provide a sufficient description of objects to enable reliable classification. Mappings in four or more dimensions may be potentially useful for classification (Lerner *et al.*, 1996) but are of no use for visualisation.

The question of perceptual quality is a separate issue. Providing a numerically good representation of the data (assessed by cost function value or classification performance) does not guarantee that the mapping will be perceptually useful. Highly clustered data, for example, may lead to degenerate mappings in which all data from one class converge to a single point in the mapping. This may give a small value for the cost function but is not a perceptually useful mapping. Thus, it is important to undertake subjective evaluation of mappings in order to gain an insight into their value to the human observer (Siedlecki *et al.*, 1988a).

Nonlinear

"Real" data often exists on a nonlinear manifold, and linear techniques may be insufficient to adequately represent the data in low-dimensional space.

Computationally tractable

For obvious reasons, computationally efficient solutions are desirable.

With these desirable features in mind, a review of three well-known visualisation algorithms is now given.

3.1.1 Multidimensional scaling

Multidimensional scaling (MDS) is a technique commonly used in the social sciences. The basic data upon which MDS operates is pairwise dissimilarity (or similarity) data, with one measure of dissimilarity for each pair of objects within a data set. The dissimilarity data may be gathered directly - for example, by asking a number of people to subjectively rate the dissimilarity of pairs of objects. In a data visualisation context, dissimilarities are typically

generated by comparing each pair of data vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, using some dissimilarity function $d^*(\boldsymbol{x}_i,\boldsymbol{x}_j) \equiv d_{ij}^*$.

Once the dissimilarity data has been gathered, the aim of MDS is to construct a set of points in low-dimensional space (commonly referred to as latent space) so that the inter-point distances match the data dissimilarities. These points form a graphical representation of the high-dimensional data, from which the structural relationships may be deduced. The representative mapping is most commonly made in a Euclidean space, although other spaces (particularly Minkowskian spaces) may also be used.

Metric multidimensional scaling

Multidimensional scaling techniques may be broadly divided into metric and non-metric algorithms. These differ in the assumptions made regarding the dissimilarity data. Classical metric scaling techniques (Torgerson, 1958; Young and Householder, 1938) make the assumption that the dissimilarity data is metric, or equivalently, that the dissimilarity function used is a metric function. The requirements for such a function are listed in Table 3.1.

1.	$d_{ij}^* \geq 0$
2.	$d_{ij}^* = 0$ if and only if $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ are identical
3.	$d_{ij}^* = d_{ji}^*$
4.	d_{ij}^* should satisfy the triangle inequality: $d_{ik}^* + d_{kj}^* \geq d_{ij}^*$

Table 3.1: Properties of a metric function d_{ij}^* applied to vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$.

The Euclidean distance is perhaps the most well-known example of a metric function.

A function which satisfies constraints 1-3 in Table 3.1 but not the triangle inequality (constraint 4) may be transformed into metric form by adding a constant to each dissimilarity (except the self-dissimilarities d_{kk}^*):

$$d^{\dagger}{}_{ij}^* = d_{ij}^* + c \qquad \forall i \neq j, \tag{3.1}$$

and

$$c \geq max_{i,j,k} \left| d_{ij}^* + d_{ik}^* - d_{jk}^* \right| . \tag{3.2}$$

However, adding a large constant will tend to swamp the dissimilarity data, obscuring the true data relationships.

In the case of metric dissimilarities, an exact reconstruction in low-dimensional space may be possible. Specifically, if the inner product matrix $\boldsymbol{X}^T\boldsymbol{X}$ is positive semi-definite of rank p, then a configuration in p dimensional Euclidean space can be found such that the inter-point distances match exactly the object dissimilarities (Mardia et al., 1979). In general, an exact fit is not possible and so a configuration which best fits the data is needed.

The least squares scaling (LSS) (Anderson, 1971; Sammon, 1969; Spaeth and Guthery, 1969) finds this best-fitting configuration in a least-squares sense. The error function may be written:

$$E = \sum_i \sum_{j<i} h_{ij} \left(f(d_{ij}^*) - d_{ij} \right)^2 , \tag{3.3}$$

where d_{ij} is the distance between point $\boldsymbol{y}_i$ and point $\boldsymbol{y}_j$ in the map, d_{ij}^* is the dissimilarity between data vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, and h_{ij} is a weighting function. The purpose of the function f is to transform the dissimilarities into metric form if necessary (for example, by adding a constant to each dissimilarity). Non-metric data may be used with the LSS, but strongly non-metric data may be more usefully processed using a non-metric scaling method.

Non-metric scaling

Non-metric MDS relaxes the assumption regarding the metric nature of the dissimilarities and instead assumes that only the rank order of the dissimilarities is meaningful. The dissimilarities may be considered to be related to the Minkowski metric by the function g:

$$d_{ij}^* = g\left(\left(\sum_k \left| \boldsymbol{x}_i(k) - \boldsymbol{x}_j(k) \right|^p \right)^{1/p} \right), \tag{3.4}$$

where g is an (unknown) monotone function (Davison, 1983):

$$d_{ij} < d_{kl} \Rightarrow g(d_{ij}) < g(d_{kl}) \; \forall i,j,k,l . \tag{3.5}$$

In a similar manner to the least-squares scaling, one may define a loss function (Kruskal, 1964a):

$$S = \sqrt{\frac{\sum_{i<j}\left(d_{ij} - \hat{d}_{ij}\right)^2}{\sum_{i<j}(d_{ij})^2}}, \tag{3.6}$$

where $\hat{d}_{ij}$ is most commonly referred to as the disparity between $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$. The disparities are calculated to match the distances d_{ij} as closely as is possible, subject to the constraint that the rank order of the disparities matches that of the original dissimilarities:

$$d_{ij}^* < d_{kl}^* \Rightarrow \hat{d}_{ij} \leq \hat{d}_{kl} \quad \forall i,j,k,l\,. \tag{3.7}$$

Thus, the rank order of the original dissimilarities is preserved in the mapping.

The minimisation of this objective function usually involves two steps: the minimisation with respect to the distances d, and the minimisation with respect to the disparities $\hat{d}$. This process is done iteratively and a typical algorithm is given in block form in Figure 3.1. Kruskal

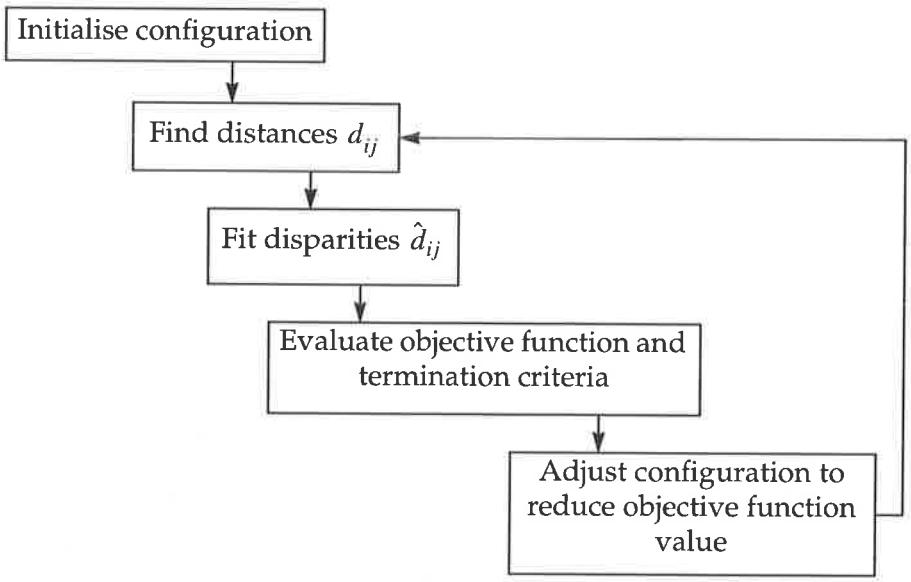

Figure 3.1: The typical steps in a non-metric multidimensional scaling algorithm.

outlined an early algorithm for minimising this objective function (Kruskal, 1964b); more recent approaches include the alternating least squares and majorisation techniques. The

optimisation is usually quite difficult and prone to entrapment in local minima in the error surface.

The following sections address various issues pertaining to the least-squares scaling. This form of MDS has been chosen for further investigation, since the cost function of this scaling is quite flexible, allowing choice of functions d^* and d, the function f and the weights h. Application to a wide variety of data types is therefore possible. Further, for data visualisation purposes, it is assumed that the dissimilarity data is of interval nature - that is, the magnitudes of the dissimilarities are meaningful, not just the rank order.

3.1.1.1 Dissimilarity functions

The function d^* strongly controls the form of the mapping and should be chosen with some care. The nature of the data will dictate appropriate dissimilarity functions. A common choice is the Euclidean distance and this may be an appropriate choice if the data represents, for example, a vectorised list of features. More complex data will require more involved measures of dissimilarity - for example, an image dissimilarity function is likely to need to be translation, rotation and dilation invariant.

Data with missing entries can be dealt with during the design of the dissimilarity function - for example, by evaluating dissimilarities based only on the elements present in both data vectors. This approach has been used with the self-organising map (Section 3.1.3).

3.1.1.2 Weighting function

With unity weighting ($h_{ij} = 1$), the LSS is equivalent to principal components analysis. More imaginative choices introduce nonlinearity into the algorithm - for example, the weighting may be used to alter the scale at which structure is preserved in the mapping. Placing emphasis on small inter-object distances will encourage locally correct preservation, in a manner similar to the small neighbourhood sizes used in the later training stages of Kohonen's self-organising map. This is important since small errors in large distances are unlikely to greatly affect the user's perception of the data, but errors of the same magnitude in small distances could be far more detrimental. This is particularly true in regions where two clusters are close or overlapping.

The Sammon mapping (Sammon, 1969) uses a weighting of $h_{ij} = 1/d_{ij}^*$ which reduces the dominance of large dissimilarities and allows the mapping to place increased emphasis on

local structure. However, small dissimilarities can cause difficulties as they have a large effect on the cost function.

Curvilinear component analysis (Demartines and Hérault, 1997) is formulated with a general weighting function which is suggested to be a bounded, monotonically decreasing function of the output distances d_{ij}. This also favours local topographical structure. As a function of the output distances, the solution is recursive and is claimed to be "the only one to unfold strongly folded data" (Demartines and Hérault, 1997).

Emphasising local structure can also avoid degenerate mappings which may occur in highly clustered data, wherein the dissimilarities between the points within each cluster are smaller than the dissimilarities between the clusters themselves. With a uniform weighting function, the relationships between the points within each cluster will be poorly represented, as the between-cluster dissimilarities will dominate the objective function (Coxon, 1982).

Other weighting strategies may be adopted - for example, consequential region scaling uses a weighting which is a function of both input dissimilarities and output distances, and which emphasises moderate distances in the mapping (Lee, 1999).

3.1.1.3 Prior information

The LSS as described is an unsupervised mapping: no information other than the object dissimilarities are used to generate the low-dimensional representation. Suppose, however, that additional information regarding the objects is available *a priori*, such as the class to which each belongs. Exploiting this knowledge during the construction of the mapping may improve the interpretability of the resulting mapping.

This may be done in a number of ways. One method is to create a second set of dissimilarities based on the additional knowledge. This "subjective dissimilarity" might be defined as zero between points of the same class, and unity between objects belonging to different classes. This subjective dissimilarity is used to supplement the objective function by replacing d_{ij}^* in Equation 3.3 with the variable δ_{ij} (Tipping, 1996):

$$\delta_{ij} = (1-\alpha)\, d_{ij}^* + \alpha s_{ij}, \tag{3.8}$$

where s_{ij} is the subjective dissimilarity between $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, and $0 \le \alpha \le 1$ controls the relative contribution of this subjective knowledge (Tipping, 1996). With $\alpha = 0$, the mapping is

unsupervised, as before. A fully supervised mapping (one in which the additional knowledge completely determines the form of the mapping) may be obtained by setting $\alpha = 1$. Different modifications to the objective function may be made, achieving for example enhanced cluster separation (Cox and Ferry, 1993) or tightened clusters of data belonging to the same class (Koontz and Fukunaga, 1972).

Another method of incorporating class information is possible if the dissimilarity function is under the control of the user. In the case of a Euclidean dissimilarity function, the contribution of each data vector element may be weighted:

$$d_{ij}^{*} = \left(\sum_{k} q(k) \{ x_i(k) - x_j(k) \}^2 \right)^{1/2}. \tag{3.9}$$

The weight vector $\boldsymbol{q}$ may be chosen to maximise the discrimination between classes (Curcie and Craelius, 1997):

$$q(k) = \frac{\sigma(k)}{\overline{\sigma}_c(k)}, \tag{3.10}$$

where $\sigma(k)$ is the inter-class variance of element k and $\overline{\sigma}_c(k)$ is the average within-class variance of element k. Thus, emphasis is placed on elements with large inter-class variance and small intra-class variance. However, this method is unlikely to work well with non-linearly separable data, since it considers only variances.

3.1.1.4 Implementation

Various methods exist by which one may minimise the objective function described in (3.3). Differentiating with respect to the low-dimensional point $\boldsymbol{y}_i$ yields:

$$\frac{\partial E}{\partial \boldsymbol{y}_i} = -2 \sum_{j<i} \frac{h_{ij}(f(d_{ij}^{*}) - d_{ij})\ (\boldsymbol{y}_i - \boldsymbol{y}_j)}{d_{ij}}. \tag{3.11}$$

Standard optimisation algorithms may then be applied, such as the Newton-Raphson gradient-descent algorithm (Sammon, 1969) or a conjugate-gradient method (Tipping, 1996). Such implementations have two significant drawbacks, however. The computational complexity scales as approximately $O(N^2)$ with the number of input patterns, as each pass

requires the computation of $N(N-1)/2$ derivatives. In addition, no protection is offered against local minima in the error surface.

The method of "pinning" (Demartines and Hérault, 1997) scales in computational complexity as approximately $O(n)$ and also admits the possibility of temporary increases in the value of the objective function during minimisation. The latter property offers some chance of escaping local minima. The algorithm is quite simple:

1. At each iteration, randomly select one $\mathbf{y}_i$ and "pin" it (hold it fixed).

2. Move all of the other $\mathbf{y}_j$ around, *without regard to the interactions amongst these* $\mathbf{y}_j$.

The rule suggested by Demartines and Hérault for the second step is:

$$\Delta \mathbf{y}_j = -\gamma(t)\,\nabla_j E, \tag{3.12}$$

where E is the objective function and $\gamma(t)$ is the step size which decreases with time. Applying this rule to Equation 3.3 yields:

$$\Delta \mathbf{y}_j = \frac{\gamma(t)\,h_{ij}\,(d_{ij}^{*} - d_{ij})\,(\mathbf{y}_j - \mathbf{y}_i)}{d_{ij}}. \tag{3.13}$$

Demartines and Hérault show that, on average, E is decreased through the action of this rule, but can momentarily increase, allowing an escape from local minima in the error surface. The duty of this mechanism is similar to that of momentum in the backpropagation training rule for multilayer perceptrons.

Stochastic techniques are another method of avoiding entrapment in local minima and have also been applied to the mapping of pairwise proximity data (Graepel and Obermayer, 1999; Klock and Buhmann, 1997). Rather than use a least-squares criterion, these mappings are formulated in a maximum likelihood sense which offers advantages for dissimilarity data subset selection.

The scaling properties of algorithms which operate on pairwise dissimilarities can become prohibitive for large data sets. One method of easing the computational demands is to choose a limited number of objects from the entire set, from which a mapping can be generated in the

standard manner (Pykett, 1978; Sammon, 1969). The subset of objects used to compute the mapping in such methods should be chosen so that the resulting mapping displays similar structure to the mapping generated using the entire object set. The subset choice is critical to the performance of the subset algorithms and so demands a significant amount of computational attention. The maximum-likelihood formulation mentioned above allows a principled approach to selecting the most relevant data dissimilarity values for subset mapping (Hofmann and Buhmann, 1995).

3.1.1.5 Neural network implementation

The least-squares scaling may be implemented in the framework of a feedforward neural network. This allows new data to be projected onto a mapping with low computational overhead. It can also offer advantages in the construction of the map itself. This is discussed in detail in Chapter 5.

Neural network implementations of the Sammon mapping have been applied to the problem of distinguishing EEG signals by task (Lowe, 1997), chromosome classification (Lerner *et al.*, 1998) and visualisation of intensive-care unit monitoring data (van Gils *et al.*, 1997).

3.1.1.6 Interaction with the user

The least-squares scaling offers several opportunities for interaction with the user. Following an initial mapping of all data, manual selection of "interesting" points may be carried out and the mapping adjusted to give the best possible representation of these points (Lee, 1998). This may be done by restricting the error function to operate only on those points of interest. Since the initial map configuration shows the best representation of all data points, this subsequent map will also tend to retain the global map structure.

The scale of the weighting function h_{ij} may also be placed under the control of the user, giving "more revealing results than the ones obtained by automatic methods" (Demartines and Hérault, 1997).

3.1.1.7 Assessing mappings

The final value of the least-squares scaling objective function (3.3) is clearly a badness-of-fit measure of the map distances with respect to the data dissimilarities. A more insightful tool

for evaluating mappings is the point-wise error - the contribution of the kth point to the final value of the objective function, evaluated using:

$$g_k = \sum_{j<k} w_{kj} (f(d_{kj}^*) - d_{kj})^2 . \tag{3.14}$$

A constant value of g across all data points indicates equal error in the representation of each, while an unusually large value of g_k for a particular data point $\boldsymbol{x}_k$ would cast doubt on the validity of that point's position $\boldsymbol{y}_k$ in the mapping.

A similar instrument is the inter-point distances reconstruction plot, in which the data dissimilarities are plotted point-wise against the reconstructed distances, as shown in Figure 3.2. A perfect reconstruction would yield a line passing through the origin with a slope of 1.

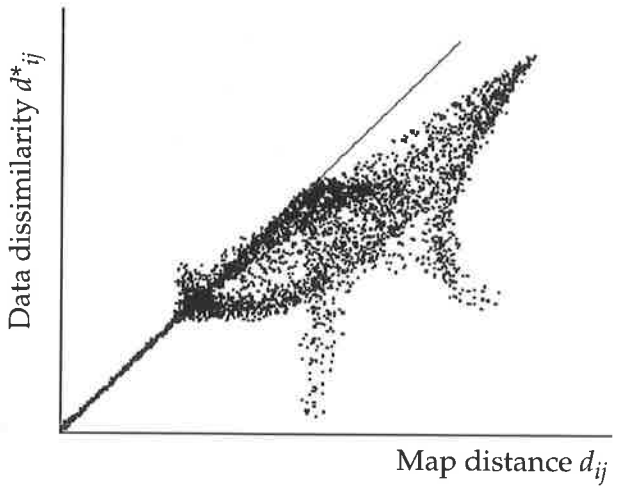

Figure 3.2: An example of an inter-point distances reconstruction plot. Small-scale structure is preserved very well in this case. (Reproduced from Demartines and Hérault, 1997).

3.1.2 The generative topographic mapping

The generative topographic mapping (GTM) (Bishop *et al.*, 1998b) has been touted as a "principled alternative" to Kohonen's self-organising map. In the GTM, a discrete grid of K points $\{\boldsymbol{u}_i\}, i = 1 \ldots K$ is defined in a low- (usually two-) dimensional latent space in a similar manner to the self-organising map. This grid is usually chosen to be rectangular and uniform, in order to aid visualisation. A set Φ of fixed nonlinear basis functions $\{\varphi_m\}, m = 1 \ldots M$ and a $D \times M$ weight matrix $\boldsymbol{W}$ is then used to define a nonlinear

transformation between latent space and D-dimensional data space. The projection of each grid point forms a reference vector $\boldsymbol{m}_i$ in data space:

$$\boldsymbol{m}_i = \boldsymbol{W}\Phi(\boldsymbol{u}_i) . \tag{3.15}$$

Each $\boldsymbol{m}_i$ is used as the centre of a Gaussian distribution in data space with variance $1/\beta$:

$$p(\boldsymbol{x}|i) = \{\beta/(2\pi)\}^{D/2} e^{-\{\beta\|\boldsymbol{m}_i - \boldsymbol{x}\|^2/2\}} . \tag{3.16}$$

Thus, the probability density function of the GTM is described by a mixture of Gaussians:

$$p(\boldsymbol{x}|\boldsymbol{W}, \beta) = \sum_{i=1}^{K} P(i)\, p(\boldsymbol{x}|i) . \tag{3.17}$$

The data is assumed to have been drawn from this mixture of Gaussians and $P(i)$ is the prior probability that the ith Gaussian component was responsible for generating a data point $\boldsymbol{x}$. This is usually chosen to be equal and constant:

$$P(i) = 1/K . \tag{3.18}$$

The model is trained using the expectation-maximisation (EM) algorithm (Bishop *et al.*, 1998b; Dempster *et al.*, 1977), giving a maximum-likelihood fit. At each iteration, one first finds the expected log-likelihood using the existing parameter values $\boldsymbol{W}_{old}$ and β_{old}:

$$\langle L(\boldsymbol{W}, \beta)\rangle = \sum_{n=1}^{N}\sum_{i=1}^{K} p(i|\boldsymbol{x}, \boldsymbol{W}_{old}, \beta_{old}) \ln p(\boldsymbol{x}|i, \boldsymbol{W}, \beta) . \tag{3.19}$$

New values of the parameters are then chosen to maximise this expectation and the process is repeated. Convergence to a local maximum is guaranteed (Bishop *et al.*, 1998b). The parameters must be given initial values which may be random or chosen to approximate a PCA mapping. An incremental algorithm has also been derived (Bishop *et al.*, 1998a).

The topographic nature of the GTM is ensured by choosing the bases Φ to provide a smooth, continuous transformation from latent to data space. Points which are close in latent space will be close in data space, thus ensuring local topographic ordering. This does not, however, guarantee that the inverse transformation from data to latent space is smooth,

although it is likely if the posterior distribution is unimodal and sharply peaked for the majority of training points (Carreira-Perpiñán and Renals, 1998). The initialisation to an approximate principal components configuration gives an initial global ordering to the map.

For visualisation, Bayes' theorem may be used to invert the transformation from latent space to data space. A given data point $\boldsymbol{x}$ has a posterior distribution in latent space given by a sum of delta functions centred at the grid points $\boldsymbol{u}_i$, each with coefficient:

$$p(\boldsymbol{u}_i|\boldsymbol{x}, \boldsymbol{W}, \beta) = \frac{p(\boldsymbol{x}|\boldsymbol{u}_i, \boldsymbol{W}, \beta)}{\sum_{j=1}^{K} p(\boldsymbol{x}|\boldsymbol{u}_j, \boldsymbol{W}, \beta)}. \tag{3.20}$$

The mean of this distribution can be computed in order to visualise a set of data more conveniently; in some cases, the distribution may be multi-modal, in which case the mode of the distribution may be more appropriate.

The strength of mixture-model based algorithms such as the GTM is the explicit modelling of the data densities. This offers principled methods of dealing with tricky problems, such as missing data and novelty detection.

One distinct disadvantage of the GTM is the computational investment: complexity scales exponentially with the number of latent dimensions and training can take almost twice the time of the SOM (Kaski, 1997). A semi-linear model which attempts to redress this disadvantage has been given (Bishop *et al.*, 1998a).

A procedure for constructing hierarchical mixtures of latent-variable models such as the GTM has been developed (Bishop and Tipping, 1998), bestowing interactivity with the user.

The GTM has been applied to visualisation of human basal ganglia responses (Branston *et al.*, 1998), and of electropalatographic data (Carreira-Perpiñán and Renals, 1998); and for exploration of large medical databases (Rosario *et al.*, 1997).

3.1.3 Kohonen's self-organising map

The self organising map (SOM) (Kohonen, 1997; Kohonen, 1990; Kohonen, 1982) is an unsupervised neural network which uses a competitive learning strategy to construct an ordered mapping of the input data. Originally developed in order to explain observed neuro-

biological behaviour, the algorithm has nonetheless found wide application in visualisation and structure exploration.

The map consists of a one- or two-dimensional array of neurons. During training, neurons within the map become sensitive to specific input stimulus patterns. If neighbouring neurons in the map are trained to represent similar input patterns, the map will form an ordered representation of the input data.

The SOM has been the subject of a great deal of research since its inception in the early 1980s. An array of variations on the basic SOM theme have been devised: different matching metrics used in choosing winning units, different neighbourhood definitions, the hierarchical SOM (Kohonen, 1997; Miikkulainen, 1990) and the supervised SOM (in which prior knowledge such as class labels may be incorporated) (Kohonen, 1988). The basic concept of the SOM - competitive learning - is described below.

Assume the neurons are located at grid points $\{\boldsymbol{u}_i\}, i = 1 \ldots K$ in two dimensional space. Associated with the ith neuron is a D-dimensional reference vector $\boldsymbol{m}_i$, where D is the dimension of the input data $\boldsymbol{x}$.

At each step of the training process, an input pattern $\boldsymbol{x}_n$ is chosen randomly from the available data. The neuron c which is closest to the input vector is chosen as the winner of the competition:

$$c = \arg\min_i \{\|\boldsymbol{x} - \boldsymbol{m}_i\|^2\} . \tag{3.21}$$

The neurons are usually arranged in a lattice formation and a neighbourhood of neurons surrounding the winning neuron are allowed to learn at each step. The neighbourhood is defined by the kernel:

$$h_{ci}(t) = h(\|\boldsymbol{u}_c - \boldsymbol{u}_i\|, t) , \tag{3.22}$$

which is a time-dependent, decreasing function of the grid distance from the ith neuron from the winning cth neuron. By choosing a wide neighbourhood kernel at the beginning of the learning process, global ordering of the map may be achieved. The choice of neighbourhood kernel can be crucial for convergence, particularly with large maps (Kohonen, 1997, page 88).

The weight update equation at each time step is given by:

$$\boldsymbol{m}_i(t+1) = \boldsymbol{m}_i(t) + h_{ci}(t)\{\boldsymbol{x} - \boldsymbol{m}_i(t)\}. \quad (3.23)$$

Training continues until the map has converged. Once global ordering has been achieved, training should continue with a small neighbourhood kernel for a "fairly long period of time" for good statistical accuracy (Haykin, 1994, p. 412).

Various optimisations may be made to reduce the computational demands of the algorithm, particularly the task of finding the winning unit at each stage. Tree-structured or table look-up algorithms can help in this regard.

Once the SOM has converged, it may be exploited for various tasks. Visualisation of data may be done by mapping each data pattern to the neuron with the most similar reference vector, providing an ordered display of the data.

The density of reference vectors in an ordered map corresponds roughly to the density of the input samples (Kohonen, 1997; Ritter, 1991). The clustering density in different regions of the data space may also be visualised using the SOM, by displaying the distances between reference vectors of neighbouring neurons (Kraaijveld *et al.*, 1995). No assumptions are required on the shapes of the clusters (Kaski, 1997), but the map tends to over-represent regions of low input density and under-represent regions of high input density (Haykin, 1994). The discrete nature of the grid locations also tends to distort the topography of the mapping (this is also true of the GTM).

Missing data (components within a data vector may be missing) may be dealt with during the step of choosing the winning neuron (Equation 3.21). The distances in this case may be evaluated based only on those entries present in the data vector, and for small amounts of missing data, this technique gives better performance than simply discarding data with missing entries (Samad and Harp, 1992).

The most commonly cited disadvantages of the SOM are the lack of cost function and convergence proof. In the general case, the SOM learning does not have a cost function (Erwin *et al.*, 1992). For the special case of discrete data and fixed neighbourhood kernel h_{ci}, the cost function has been shown to be (Ritter and Schulten, 1988):

$$E = \sum_k \sum_i h_{ci} \|\boldsymbol{x}_k - \boldsymbol{m}_i\|^2 . \tag{3.24}$$

This cost function is strikingly similar to the cost function of the *k*-means clustering algorithm and indeed, the two are closely related (Kaski, 1997).

The SOM has been applied to a diverse range of problems, many with a medical theme, such as classification of lung sounds (Kallio *et al.*, 1991) and recognition of EEG spectra (Joutsiniemi *et al.*, 1995).

3.1.4 Conclusions

The least-squares scaling is the most appealing of the three algorithms for general use as a visualisation tool. It has a flexible cost function and can accommodate various data types through a suitable definition of the dissimilarity function. Parameters of the cost function can also be adjusted interactively by the user to focus on the important structure of the mapping and prior information (if available) can also be included.

The neighbourhood size of the SOM is equivalent to the weighting function of the LSS and may also be adjusted in an interactive manner. The scale at which the GTM preserves structure is controlled indirectly by the choice of weights and bases and the algorithm is restricted mainly to focusing on local structure.

Both the GTM and the LSS have an explicit cost function. The lack of cost function and convergence proof for the SOM suggest that objective quantification of the results may be difficult.

The main deficiency of the LSS is the lack of generalisation. This can be rectified by training a neural network to perform the LSS and methods for this are discussed and developed in Chapter 5. The GTM and SOM are both capable of generalisation without modification of the algorithms.

For very large data sets, the computational demands of the LSS may be unrealistic and the GTM may be preferred in this case.

3.2 Interpreting mappings

One remaining point which should be addressed before moving on is the interpretation of mappings. In most applications, the purpose of visualisation is not simply to construct a graphical depiction of inter-data relationships, but to gain a deeper understanding of the data. The interpretation of mappings relies heavily on domain knowledge of the data in order to appreciate the structure of the mapping.

This task may be assisted by analysing the characteristics of the data which are responsible for certain aspects of the mapping distribution. For example, given clusters of normal and pathological data, by identifying the components of the data which differ between the two clusters.

Producing a "backward mapping" can assist in this regard - rather than projecting data from data to latent space, the reverse is desired. This can give the input data vector which would be required to attain a particular output map location (Figure 3.3). By examining the changes in the input vector as one moves from one cluster to another, an insight into the important data attributes can be gained.

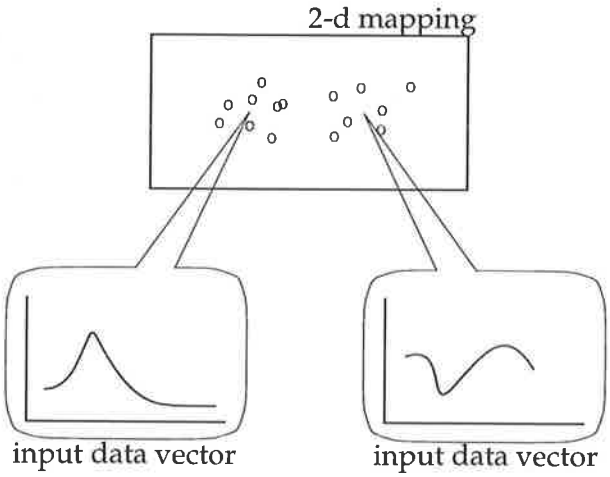

Figure 3.3: Backward mapping. The input vector corresponding to interesting map location (cluster centroids, in this case) are examined.

Implementing the backward mapping in the case of the GTM is trivial, as the algorithm explicitly builds a parametric model for this transformation. For the LSS, no such functional form exists. One option (Demartines and Hérault, 1997) is to produce a local fit to the existing

input/output point pairs: that is, for a given mapping described by $\{x\} \rightarrow \{y\}$ and a map location y_0 (which does not necessarily coincide with one of the mapped points in the set $\{y\}$), the corresponding backward-mapped vector x_0 is found by minimising the cost function (3.3) for the combined set of points, but moving *only* the unknown point x_0. This gives a local fit with low computational cost.

A related idea is to examine the sensitivity of the mapping to variations in the input data elements, by adding small deviations to the input vectors and examining the changes in the map location (Kaski, 1997; Lehtinen *et al.*, 1997). Note that this requires the ability to generalise, or project new data onto the mapping.

The interpretation of mappings may also be assisted by examining the distribution of the map with respect to other information about the data set. In a biomedical context, viewing the distribution of the mapping with respect to, say, subject age or weight may reveal significant patterns in the map.

3.3 Summary

A brief review of data visualisation has been given. The least-squares scaling and the generative topographic mapping have been identified as appropriate algorithms for the visualisation of biomedical data. The least-squares scaling is very flexible and offers a number of avenues for user interaction. The generative topographic mapping has a convenient mechanism for projecting new data in an efficient manner.

The application of the least-squares scaling and the generative topographic mapping to heart rate variability data is addressed in the following chapter.

Chapter 4 Visualisation of heart rate variability data

In this chapter the application of the least-squares scaling and the generative topographic mapping to heart rate variability data is demonstrated. Before these visualisation algorithms can be applied, a suitable format for the HRV data must be established.

4.1 Data preprocessing

Suitable preprocessing is necessary to generate the pairwise data dissimilarity measures from the raw HRV data.

4.1.1 The least-squares scaling

As discussed in Section 3.1.1.1, the choice of the dissimilarity function d_{ij}^* must be tailored to suit the data to which the LSS is to be applied. For heart rate variability data, we wish to investigate data in the spectral domain, estimated from the time-domain HRV signal by an appropriate method (see Section 2.2.3). Thus, we consider the data to be in the general format of power spectral density amplitudes.

Given data in this form, an appropriate method of computing the dissimilarity between data is needed. Each spectral profile typically shows oscillations at respiratory frequency, at approximately 0.1 Hz, and at lower frequencies again. The amplitudes of these oscillations are meaningful, as they indicate the strength of heart rate modulation by the associated underlying mechanisms. However, the significance of the frequency at which each of these oscillations occur bears some consideration.

The frequency of respiratory sinus arrhythmia is clearly linked to the frequency of respiration - in adults usually between 0.2 and 0.4 Hz (Akselrod, 1995). Changes in the frequency of the RSA component are simply a reflection of an altered breathing rate and in most cases are of no clinical interest. Two spectra in which the RSA components are centred

on different frequencies, but which are otherwise identical should thus be considered to be quite similar (Figure 4.1).

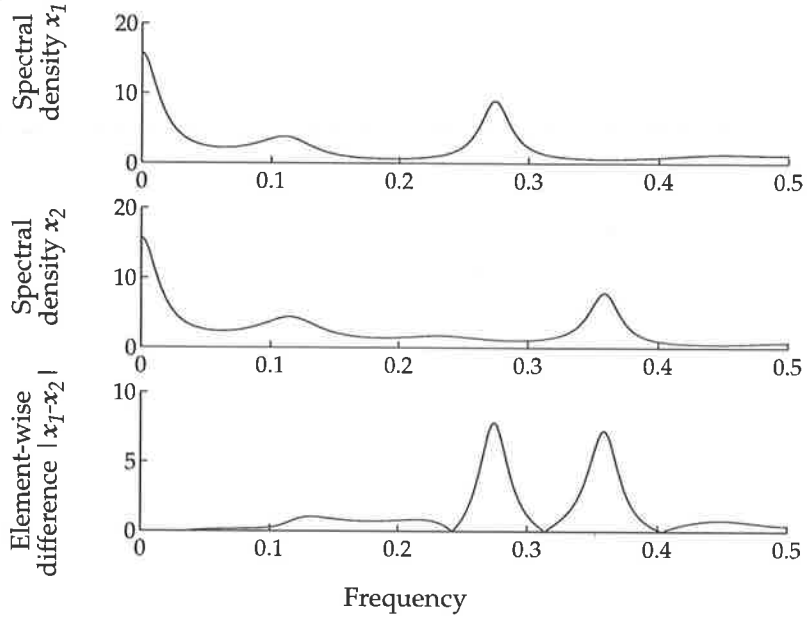

Figure 4.1: Element-wise comparison of two spectra with misaligned respiratory components. The overall dissimilarity would be undesirably large.

The LF component of the HRV spectrum, nominally located at 0.1 Hz - but which may be anywhere from 0.06 Hz to 0.14 Hz (Moser *et al.*, 1994) - is generally considered to be due in part to the baroreceptor feedback control of blood pressure. The centre frequency of this component varies and has been shown to decrease with exercise (Baselli *et al.*, 1991; Kamath *et al.*, 1991), with a change in posture from supine to standing (Kamath *et al.*, 1991; Weise *et al.*, 1989), in subjects with diabetic autonomic neuropathy (van den Akker *et al.*, 1983), and in physically fit subjects (Baselli *et al.*, 1988c). The changes in the frequency of this peak are usually quite small - of the order of 0.01 Hz - but are most certainly of clinical interest. Misaligned baroreceptor peaks should be considered to be an indication of dissimilar spectra.

A variety of mechanisms contribute to the very low frequency (<0.05 Hz) portion of the HRV spectrum. Little is known about the significance of these components: here, it was assumed that the centre frequency of oscillations in this region are of interest.

With these points in mind, the following preprocessing of each HRV spectral estimate was undertaken prior to the calculation of the pairwise dissimilarities across the data set.

1. The RSA component was first separated from the remainder, either by simple band-summation of the spectral estimate (the HF component in conventional fixed-band HRV spectral analysis) or by adaptive filtering if the appropriate time-domain signals were available. The latter offers some protection against the case where respiratory frequency wanders outside the limits of a fixed band. Additionally, explicit identification of the respiratory component encourages monitoring of the respiratory frequency. The user may easily verify that the respiratory frequency does not exceed half the average inter-beat frequency (which causes aliasing in the spectral estimate) nor become too low, as this may cause entrainment of the 0.1 Hz peak.

2. The variance (or the sum under the PSD curve) of the respiratory component was calculated.

3. The spectrum of the remainder of each signal was truncated, discarding frequencies above 0.15 Hz. Only low-frequency activity will be present in the remainder after filtering.

4. The band-limited spectral estimate was then smoothed to reduce the effects of noise. Excessive smoothing will obviously mask important information in the spectra, such as small changes in the low-frequency peak locations. Smoothing may be carried out by using a windowing function in the autocorrelation domain (Furui, 1989) or the spectral domain (Blackman and Tukey, 1958). It is worth noting that an almost identical problem to the comparison of HRV spectra has been tackled in the field of speech recognition: the comparison of an unknown speech sample with a known reference pattern. In western languages, there is normally little correlation between pitch variations and the phonetic content of a word, especially at high frequencies (Holmes, 1988). Spectral smoothing is also a common technique in speech processing, as a means of avoiding problems with pitch variations.

5. The respiratory component and remainder were then recombined into a composite vector $\boldsymbol{x}$, containing the smoothed spectral estimate from 0-0.15 Hz concatenated with the variance of the RSA component (Figure 4.2). The pairwise dissimilarities of the data may be computed directly, using an element-wise comparison of pairs of vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, such as the weighted Euclidean distance:

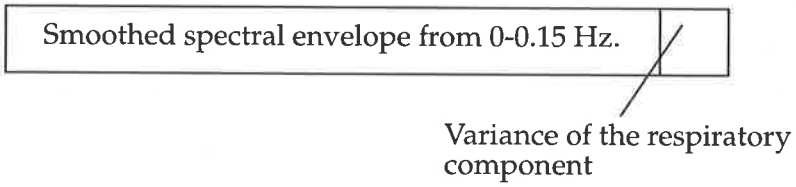

Figure 4.2: Structure of input vectors after preprocessing. The pairwise dissimilarity between vectors may now be computed using, say, a weighted Euclidean distance.

$$d_{ij}^{*} = \left(\sum_{k} q(k) \{ x_i(k) - x_j(k) \}^2 \right)^{1/2}. \tag{4.1}$$

The weight vector $\boldsymbol{q}$ is chosen to reflect the relative importance of the elements in $\boldsymbol{x}$ - which is generally unknown, and so unity weighting may be used.

Normalisation of spectra has been discussed already in Section 2.2.3.6. Since normalisation generally improves the interpretability and reduces the effect of total-power changes on the narrowband components, it is important to apply normalisation to the spectra for visualisation also. The logarithm of the data vectors may also be taken. This reduces the effects of particularly large peaks in the spectrum, which otherwise may cause far-outlying points in the map and detract from the overall perception of the data set.

It should be noted that very similar measures are used the field of speech processing, for example, the Itakura-Saito distance (Itakura and Saito, 1970) is given by:

$$d_{ij}^{IS} = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left(\log \{ \frac{x_i(\omega)}{x_j(\omega)} \} + \frac{x_j(\omega)}{x_i(\omega)} - 1 \right) d\omega. \tag{4.2}$$

This distance measure is not suitable for use with the LSS since it is not symmetric ($d_{ij}^{IS} \neq d_{ji}^{IS}$). The cosh measure (Gray Jr and Markel, 1976) is a symmetric version of the Itakura-Saito distance and is computed by simple averaging:

$$d_{ij}^{COSH} = (d_{ij}^{IS} + d_{ji}^{IS})/2. \tag{4.3}$$

However, Equations 4.2 and 4.3 make no concessions for the case where the respiratory components of the two HRV spectra are misaligned.

The above numbered procedure for calculating dissimilarities may be applied in conjunction with any spectral estimation method and was used for the work in Section 4.2. Particular spectral estimation techniques may offer other methods of computing dissimilarities based on the signal model employed. Some possibilities are discussed in Section 4.1.2.

4.1.2 Parametric spectral estimation

Parametric spectral estimation methods produce a set of model parameters from which the spectral estimate is computed. In this case, it is not necessary to explicitly compute the spectral estimates in order assess the spectral dissimilarities between two signals $\boldsymbol{s}_i$ and $\boldsymbol{s}_j$.

Autoregressive (AR) spectral estimation in particular has been widely used in the field of HRV processing. The AR model parameters $\boldsymbol{a}_i$ and $\boldsymbol{a}_j$ along with the estimates of the driving noise variances σ_i and σ_j contain all of the information necessary to compute the spectral estimates. Thus, it is possible to compute dissimilarities based on this information. For example, the Itakura-Saito distance (Equation 4.2) may be evaluated using the following form:

$$d_{ij}^{IS} = \frac{\sigma_i^2 \boldsymbol{a}_i^T \boldsymbol{V}_j \boldsymbol{a}_i}{\sigma_j^2 \boldsymbol{a}_j^T \boldsymbol{V}_j \boldsymbol{a}_j} + \log\left(\frac{\sigma_j^2}{\sigma_i^2}\right) - 1, \tag{4.4}$$

where $\boldsymbol{V}_j$ is the autocorrelation matrix of the signal $\boldsymbol{s}_j$.

Alternatively, the AR model may be explicitly decomposed into its constituent poles (Cerutti, 1995; Zetterberg, 1969), which would allow a dissimilarity measure based on the positions of the poles in the z-plane.

Comparison of HRV spectra based on parametric model coefficients has been done for the purposes of classification (Curcie and Craelius, 1997; Lee *et al.*, 1989). However, comparison in this manner does not deal appropriately with misaligned respiratory peaks. The pole-based method offers some protection in this regard, since the arguments of the appropriate poles could be ignored. Both methods also require constant model order across all data within the set, which may not be appropriate if the data range across multiple subjects and significant physiological changes (Pinna *et al.*, 1996).

Other spectral estimation techniques may offer convenient means of comparing spectra. Cepstral estimation is another parametric method, popular in the field of speech recognition. While it has been applied with some success to HRV studies (Curcie and Craelius, 1997; Watson, 1994), it has not been established as a common method in this field. The wavelet transform, on the other hand, seems destined to become a stock tool for HRV processing. In the case of orthonormal wavelets, the coefficients of the wavelet transform may be regarded as the outputs of the corresponding filter bank. Comparison of HRV spectra could therefore be made on the basis of the wavelet coefficients. Nonlinear dynamics may also offer a means of comparing HRV signals, for example, on the basis of fractal number (Bakardjian and Yamamoto, 1994).

Thus, particular spectral estimation or nonlinear analysis methods may lend themselves to various measures of dissimilarity. Here, attention has been focused on the use of spectral amplitude data, since this is a very general format and does not presuppose the use of any particular spectral estimation algorithm.

4.1.3 The generative topographic mapping

The preprocessing described for the LSS is also appropriate for the GTM, since the GTM uses the Euclidean distance in formulating probability densities.

4.2 Case study 1

In this study, heart rate variability data was collected from a simple physiological experiment, providing a dataset representing well-understood changes in the autonomic control of the heart. Six healthy volunteers were studied in the supine and tilted positions, before and after the administration of 80mg of the β-sympathetic blocker drug propranolol.

Orthostatic tilt (from supine to upright) is accompanied by a shift from vagal to sympathetic control of the heart. Mean heart rate is increased on standing and studies with HRV have shown an increased VLF (<0.05 Hz), increased LF and decreased HF spectral components, as well as increased total power (Pagani *et al.*, 1986; Weise *et al.*, 1989; Lipsitz *et al.*, 1990). The administration of β-sympathetic blockade has been shown to reduce the LF component of the HRV spectrum in both the tilted (Pomeranz *et al.*, 1985) and supine (Pagani *et al.*, 1986) postures.

The intent of the study was to demonstrate the application of the least-squares scaling and the generative topographic mapping to HRV data.

4.2.1 Methods

The subjects were allowed to breathe freely throughout the procedure and the electrocardiogram and respiration (chest diameter) were recorded. Five minute segments of RR interval data were selected for analysis. Isolated ectopic beats were removed using linear interpolation and segments containing multiple ectopic beats were discarded.

The heart rate record was interpolated using cubic spline interpolation and resampled at 10Hz. To avoid aliasing, the resampled signal was low-pass filtered using an 8th order Chebychev filter with cutoff frequency of 0.4 Hz and decimated by 10, to give a signal with uniform sampling at 1 Hz. The time series was then demeaned and detrended by removing the best straight-line fit from the data.

Adaptive filtering was used to separate the RSA component from each RR variability signal. An 8th order lattice filter was used, with forgetting factor 0.98 (Bianchi *et al.*, 1994). The validity of each filtering operation was confirmed by checking the whiteness of the cross-correlation of the filtered RRV signal and the respiratory signal (Billings and Alturki, 1992).

A time-frequency distribution employing a binomial kernel function (see Section 2.2.3) was used to estimate the time-varying power spectrum of the filtered RRV data. Spectra were subsequently averaged over non-overlapping 32-second time intervals in order to reduce cross-terms and negative values (Cohen, 1989). Negative values of the distribution which remained were set to zero. Since averaging in time was used, no further smoothing was applied to the spectral estimates.

The variance of the previously separated RSA component was calculated and scaled by $1/\sqrt{64} = 1/8$. Since the respiratory signal nominally resides in the band 0.15-0.4 Hz (which is 64 frequency bins wide), this scaling equalises the contribution of the RSA and lower-frequency components in the dissimilarity function. The scaled variance was concatenated with the spectrum of the remainder signal and the log taken, as described in Section 4.1.1.

At the completion of the preprocessing, the final data were 39-element vectors, containing the smoothed spectral density from 0-0.15 Hz (38 elements) concatenated with the scaled variance of the RSA component (39th element).

The interval variability (RRV) and rate variability (HRV) spectral data from all subjects were first subjected to conventional analysis. The total variability power, as well as the powers in three spectral sub-bands [VLF (0 – 0.05 Hz), LF (0.06 – 0.14 Hz) and HF (power coherent with respiration)] were evaluated. These quantities were evaluated without normalisation, as well as with two normalisation procedures commonly used in HRV studies.

The visualisation was carried out on a subset of the data to ease computational demands and to facilitate the interpretation and comparison of the various mappings. This data subset was taken from one subject only, giving a total of 31 input vectors (7 patterns for tilt/ propranolol, and 8 patterns each for the remaining combinations of posture and drug).

4.2.2 Results and discussion

Propranolol was effective in inhibiting the sympathetic nerve activity, with an increase in heart period from 0.94(0.09) seconds [mean(standard deviation)] to 1.12(0.14)s (p<0.01) in the supine position. Tilt produced a decrease in heart period from 0.94(0.09)s to 0.74(0.07)s (p<0.01).

Unnormalised results

The unnormalised total power of RRV decreased with tilt and increased with propranolol, as expected (see Table B.1, Appendix B). In the tilted position, HRV showed a decline in total power with propranolol. The changes in total power distort the measures obtained from band powers - unnormalised RRV LF power decreased with tilt and increased with propranolol.

Normalising by total power

Normalising the spectra by the total power (excluding very low frequency components, below 0.05 Hz) produced familiar results for posture and β-blockade: tilt caused an increase in LF and the LF/HF ratio while HF decreased, reflecting the transition to increased sympathetic control (Table B.2, Appendix B). In the tilt posture, propranolol caused a decrease in LF; a marginal decrease was also seen in LF when propranolol was administered in the supine position. In this posture the heart is largely under parasympathetic control, and this

changes toward sympathetic dominance when upright. Thus, the effects of propranolol were less evident in the supine posture.

Normalising by mean heart rate squared

The total power of the non-normalised HRV spectra varied considerably when compared to that of the RRV spectra, as expected from the discussion in Section 2.2.3.6. The difference in the two total powers, as a fraction of the total RRV power was 1.03 (1.09) (Table B.3, Appendix B). Under spectral normalisation by mean heart rate squared (or mean heart interval squared for RRV) the discrepancy in total power shown by HRV and RRV was reduced to 0.12 (0.14) . Thus, this normalisation scheme was effective in reducing the discrepancy between HRV and RRV total power, consistent with the results of Section 2.2.3.6. However, this normalisation scheme did not adequately account for variations in total power, allowing changes in total power to obscure the changes in the band-power measures. Perhaps as a consequence, fewer significant changes were seen under this normalisation - for example, none of the VLF, LF or HF measures for either RRV or HRV showed a significant change with tilt.

Ramifications for visualisation

From these results of fixed-band analysis, normalisation by total power (excluding components below 0.05 Hz) appeared to be the most sensitive to changes in HRV or RRV power spectra with tilt and propranolol. This normalisation with the RRV signal was used for the investigation of the visualisation algorithms in the next section.

4.2.2.1 Visualisation results

Principal components analysis

The projection of the data onto the first two principal components is shown in Figure 4.3. This mapping separates the data roughly by posture according to x-axis position, but with no clear separation of the placebo/propranolol clusters for each posture.

The first eigenvectors corresponding to the two largest eigenvalues are plotted in Figure 4.4. The projection of the data onto these two vectors explains about 73% of the total variance in the data. In the first principal component (the projection onto which determines the x-coordinate of points in the mapping), the RSA component (rightmost vector entry) is positive, while the other entries are negative. Since all data elements are negative (after normalising

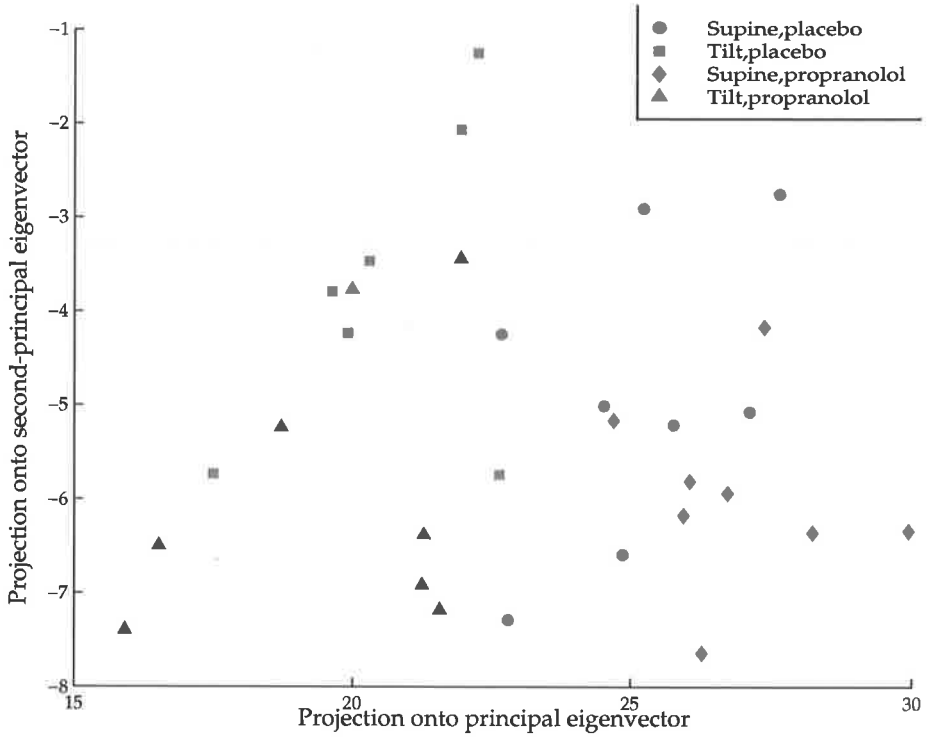

Figure 4.3: Projection of the posture/propranolol data onto the first two principal components. The effects of posture are roughly separable by x-axis position.

and taking the logarithm), data with large RSA and small lower-frequency components will tend to have a large projection onto this eigenvector. The supine data is clustered at higher x-coordinates than the tilt data, and this matches the fixed-band observation that tilt causes a reduction in RSA and an increase in the lower frequency components.

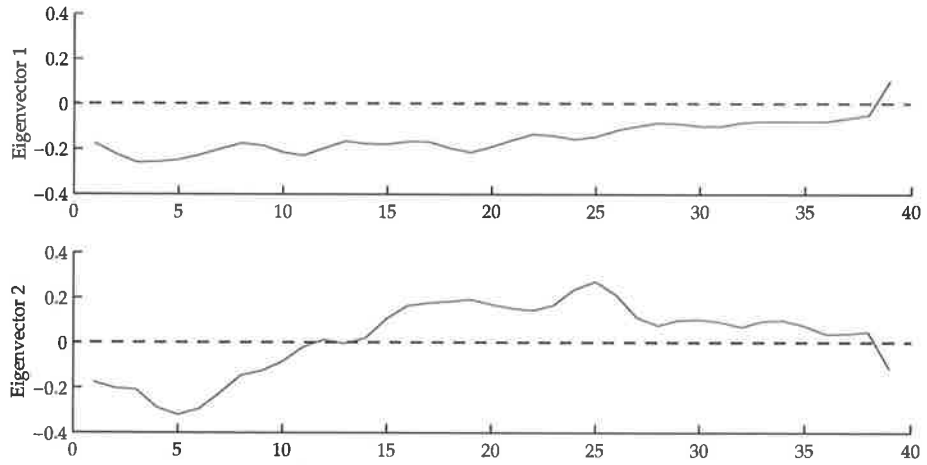

Figure 4.4: Eigenvectors corresponding to the two largest eigenvalues.

The mapping does not clearly separate the propranolol data from the placebo, although there is a tendency for the latter to be located at higher y-coordinates. The structure of the second eigenvector in Figure 4.4 suggests that data with small RSA and very low frequency (<0.05 Hz) and large low frequency (0.05-0.15 Hz) components will tend to have high y-coordinates in the mapping. This is consistent with the known depression of LF activity by propranolol.

Least-squares scaling

An unweighted Euclidean distance was used to generate the pairwise dissimilarities from the preprocessed data. A least-squares scaling of the data with weighting $w_{ij} = 1/d^*_{ij}$ (the Sammon mapping) is shown in Figure 4.5. This is very similar to the PCA map, as might be

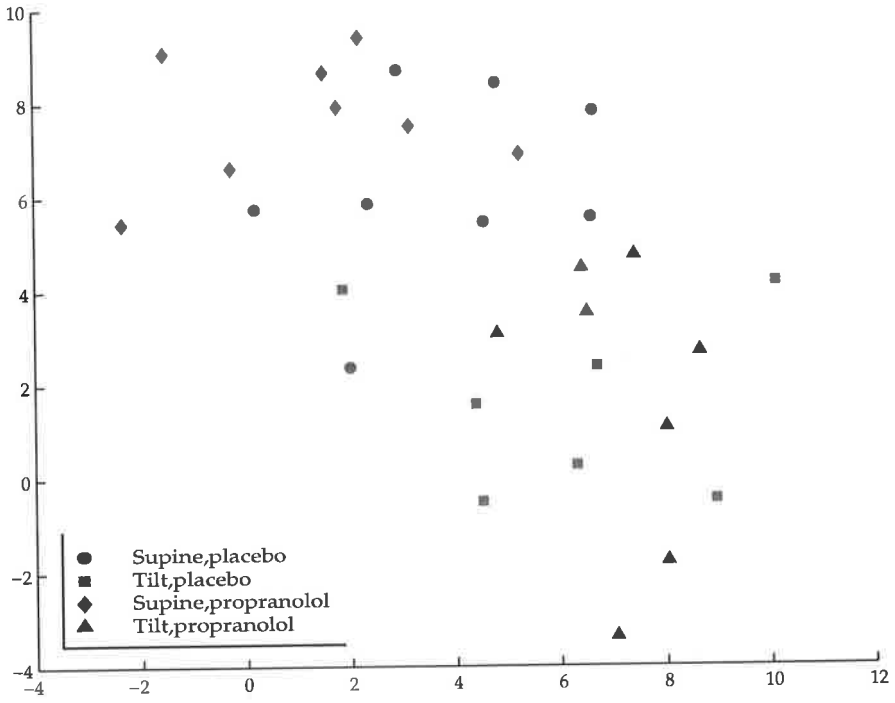

Figure 4.5: Sammon mapping of the posture/propranolol data. The error for this mapping was 0.045.

expected: LSS with unity weighting is equivalent to PCA, and the weighting used in the Sammon mapping variant simply tends to favour correct reconstruction of local structure at the expense of global structure. The contribution from each map point to the final error is shown in Figure 4.6, and indicates that, with the possible exception of points 5 and 6, the map represents the data dissimilarities reasonably well. The two questionable points belong to the

supine/placebo cluster and can be found at approximate coordinates [2.0, 2.4] and [6.6, 5.6] respectively.

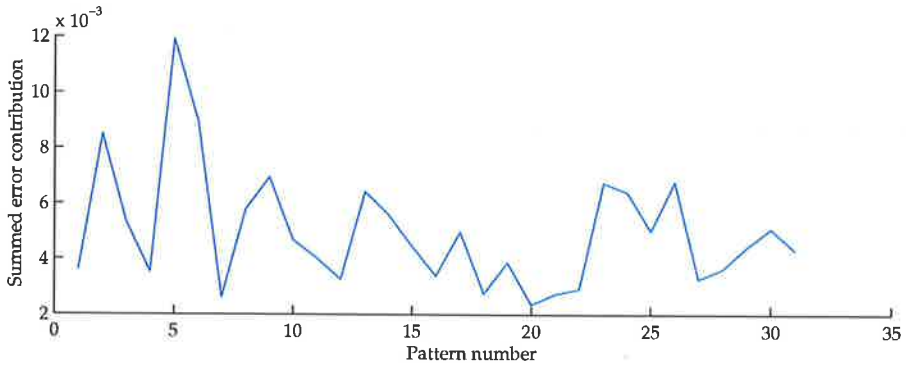

Figure 4.6: The contribution of each pattern to the overall final error of the Sammon mapping in Figure 4.5.

Like PCA, the Sammon mapping fails to clearly separate the propranolol and placebo data in either posture.

The backward mapping technique discussed in Section 3.2 was applied to the approximate centroids of the tilt and supine clusters, giving estimates of the input vectors $\boldsymbol{x}_t$ and $\boldsymbol{x}_s$ required to give points at the tilt and supine cluster centroid locations. The difference between $\boldsymbol{x}_t$ and $\boldsymbol{x}_s$ represents the changes in the input which would be required in order to move from the tilt to the supine cluster and is shown in Figure 4.7a. This describes very similar changes to those indicated by fixed-band analysis: supine data has larger RSA and smaller low-frequency components than tilt data. Similarly, the change in input required to move from the upper to lower regions of the mapping are shown in Figures 4.7b and 4.7c. Again, the change with propranolol is consistent with a decrease in LF activity.

An exponential weighting function $h_{ij} = e^{-\lambda d_{ij}^*}$ was applied in order to strongly emphasise the local structure of the mapping (Figure 4.8). The locally-weighted mapping shows considerable difference around the region of point 5, identified as being relatively poorly represented in the Sammon mapping of Figure 4.5. The error contribution plot (Figure 4.9) shows that the local structure is well represented across all data points.

The mappings in Figures 4.5 and 4.8 are purely unsupervised. Prior information may be included in the LSS in a variety of ways. A weighted Euclidean distance was used to recompute the data dissimilarities. The weights were chosen to favour those data vector

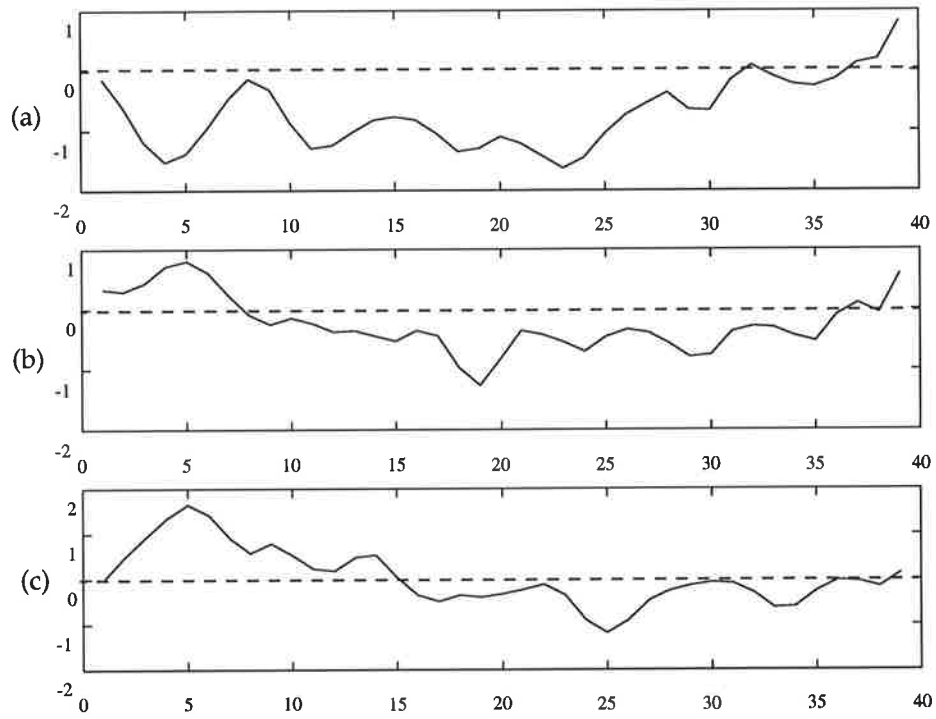

Figure 4.7: The changes in the input vector to the Sammon mapping which are required in order to move (a) from the centroid of the supine to the centroid of the tilt cluster; (b) from the upper to the lower part of the tilt cluster; (c) from the upper to the lower part of the supine cluster.

elements with small intra-class variance and large inter-class variance, as described in Section 3.1.1.3. In this case, this weighting made very little difference to the map configuration, suggesting that the variance is insufficient to describe the inter-class relationships of the data.

A second method was then employed: a class-based "subjective dissimilarity" s_{ij} between each data point was calculated: $s_{ij} = 0$ for points i and j with the same drug/posture combination, and $s_{ij} = 5$ otherwise. (The scalar 5 was chosen as the approximate mean of the raw data dissimilarities). These subjective dissimilarities were used to modify the data dissimilarities according to Equation 3.8, with coefficient α=0.4. The resulting Sammon mapping is shown in Figure 4.10. The clusters are now more clearly separated, allowing examination of the inter-cluster structure. The effects of posture are still roughly separable by x-axis position and in general, each cluster is a neighbour of the two clusters with which it shares either a posture or drug attribute. Further, the cluster corresponding to supine posture and propranolol is closest to the supine/placebo cluster. Since in the supine posture, the heart is largely under vagal control, the effects of propranolol should be minimal and so these two clusters would be expected to be the most similar. The tilt/placebo cluster is most distant

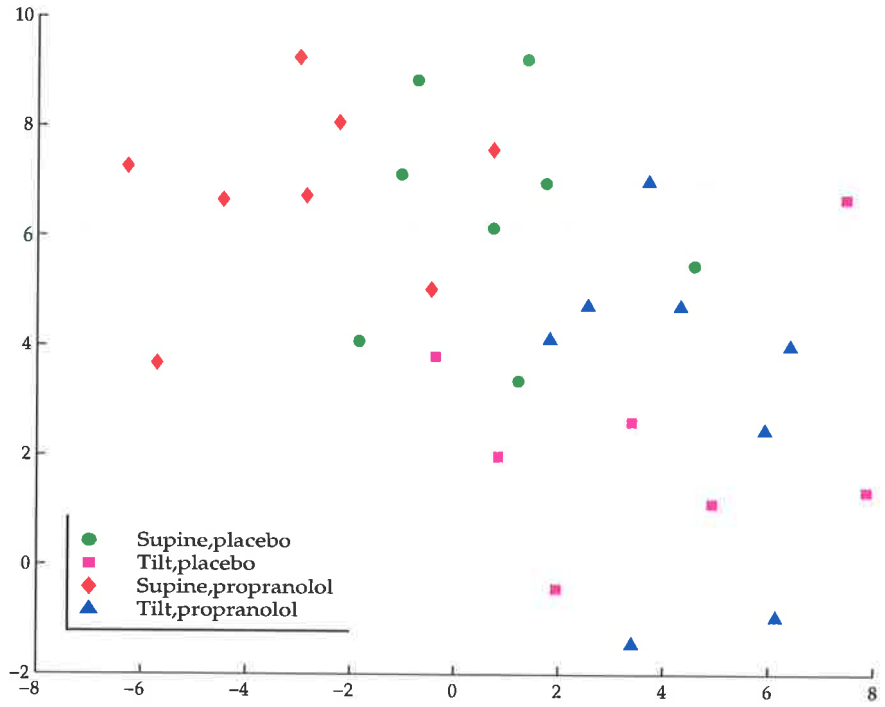

Figure 4.8: Least-squares scaling of the posture/propranolol data. An exponential weighting on the data dissimilarities was used, so the mapping represents local structure at the expense of global structure.

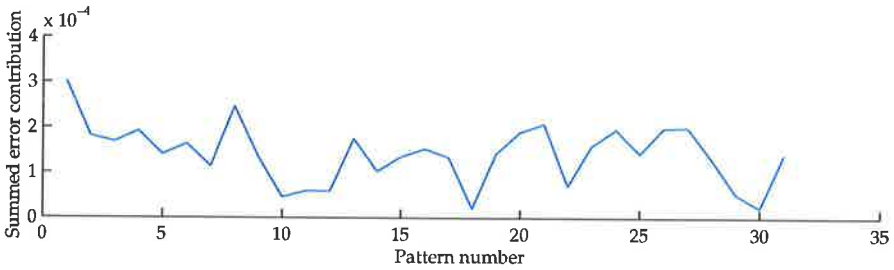

Figure 4.9: The contribution of each pattern to the overall final error of the least-squares scaling in Figure 4.8. The local structure is well preserved across all points.

from supine/propranolol, again as expected, since these two conditions are at the two extremes of sympathovagal control.

The final demonstration on the LSS is that of "fusion" of HRV and mean heart rate information. The combination of these two variables has been suggested as a better measure of autonomic tone than HRV alone (Malik and Camm, 1993). The RR interval data was segmented into windows which matched the time duration of each spectral data vector and

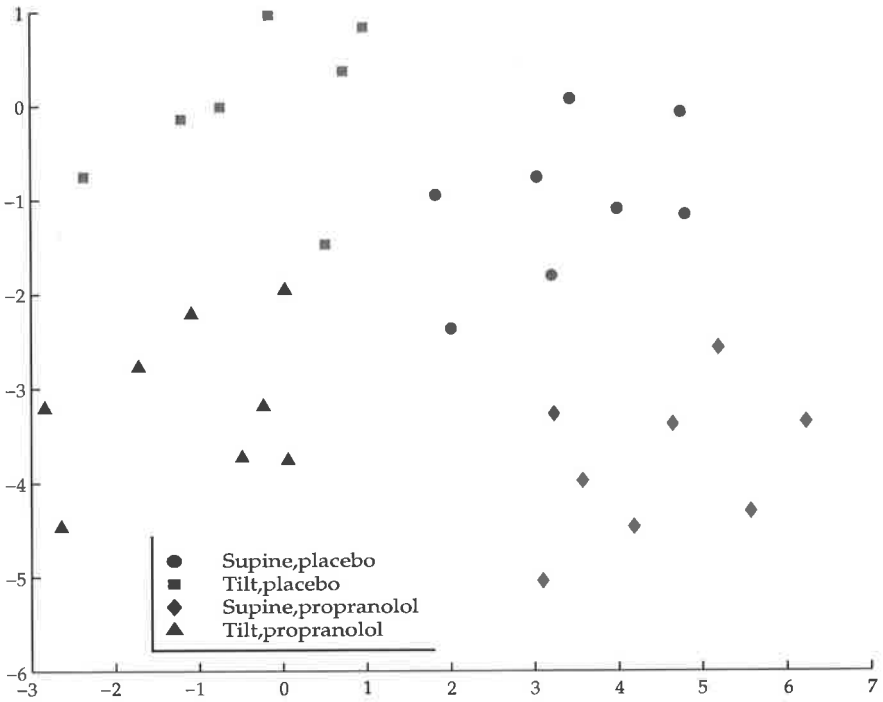

Figure 4.10: Sammon mapping of the posture/propranolol data, with class information used to partially supervise the mapping ($\alpha = 0.4$).

then averaged over these time windows. The mean heart interval data, averaged by class, is shown in Table 4.1.

Supine (placebo)	Tilt (placebo)	Supine (propranolol)	Tilt (propranolol)
0.91 (0.03)	0.75 (0.02)	1.10 (0.03)	0.98 (0.07)

Table 4.1: Mean (std) of heart period (RR interval) for the data subset used in the visualisations.

The mean RR interval was included in the data vectors and the dissimilarities recomputed. As with the inclusion of the class-label information, the cluster separation improves (Figure 4.11).

Generative topographic mapping

The posterior-mean projection of the data onto a 20x20 grid in latent space is shown in Figure 4.12. Nine uniformly-spaced Gaussian basis functions were used with moderate spread $s = 1$. The mapping was initialised using the PCA configuration and in fact the final mapping does not differ greatly. Backward mapping with the GTM is trivial since the

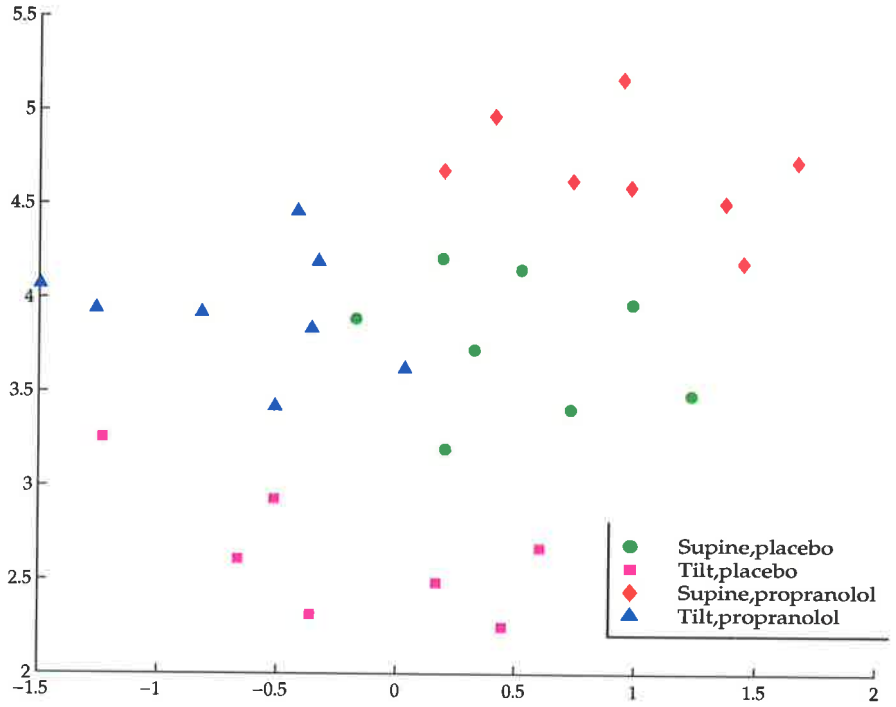

Figure 4.11: Sammon mapping of the posture/propranolol data, with mean heart rate used as additional information.

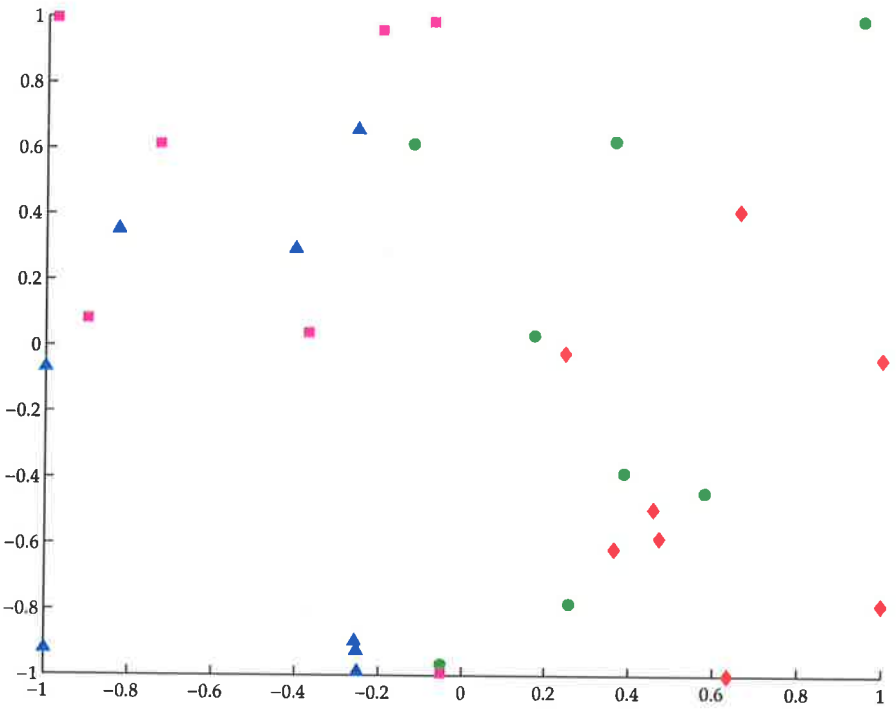

Figure 4.12: Posterior-mean projection of posture/propranolol data onto the GTM latent space.

transformation from latent to data space is explicitly modelled by the algorithm. Examining changes in the input vector with respect to changes in the map location gives similar results here to those shown with PCA and the LSS.

4.2.3 Conclusions

Section 4.2 gives a flavour of several visualisation algorithms applied to HRV data. In the purely unsupervised case, the least-squares did little better than PCA, as did the GTM. Each provided a mapping which logically represented the relationships between the data from a simple physiological experiment.

Use of prior information in the LSS, in the form of class labels, improved the cluster separation and revealed the inter-class relationships more clearly. "Fusion" of mean heart rate and HRV information had a similar effect.

4.3 Summary

The use of graphical methods for HRV analysis is not new. The well-established 'geometric techniques (Malik, 1995) use graphical techniques to assess histograms of RR interval length. The Poincaré plot (Kamen and Tonkin, 1995) uses successive interval lengths as the coordinates of points in two dimensional space. Both techniques rely on the human operator to assess the resulting figure - for example, in the Poincaré plot, one may look for fan- or comet-shaped distributions. The robustness of these methods to data outliers (Malik, 1995) stems from the inclusion of the human factor in the information processing chain (Daoudi *et al.*, 1993).

The LSS and GTM can be considered as general-purpose graphical analysis methods. The advantages of human influence are retained and the methods may be applied to data in a variety of formats. In this section, the application of these two techniques to HRV spectral data has been demonstrated, with excellent results. With suitable preprocessing, there is no need to rely on fixed-band analysis of the spectral data, thus avoiding the necessity of choosing these band limits as well as retaining sensitivity to movement of low-frequency peaks in the spectrum.

These methods do not extract information which relates directly to the sympathetic and vagal tone and the axes of the figures do not necessarily have any physiological significance. However, this is not a significant disadvantage over conventional processing since fixed-band analysis suffers the same limitations.

These visualisation techniques also offer the opportunity to discover which characteristics of the input vector cause the observed structure in the mapping - for example, by "backward mapping" points of interest in the map. In this way, the expert knowledge of the user regarding the relevance of the input vector elements can be used to extract the physiological relevance of the map structure.

The input data format is quite flexible and not restricted to simple spectral profiles. The model parameters of parametric spectral estimation methods may offer an alternative means of computing dissimilarities, as discussed in Section 4.1.2.

Fusion of data from several sources is also possible, using a weighted Euclidean distance to control the influence of each information source on the map structure.

The ability to compare data is a valuable tool in the arsenal of the HRV analyst. While the definition of "normal" HRV is yet to be firmly established (Task Force, 1996), the comparison of unknown data to a database of known examples opens up the possibility of constructing a diagnostic aid. Visualisation techniques offer direct, visual comparison of data. Data which becomes available after a map has been constructed may be projected by the GTM onto the same mapping, since this algorithm has the ability to generalise to new data. The LSS has no such function, and in the next chapter, the use of neural networks to add this capability is discussed.

Chapter 5 Generalisation and prior information

The collection of biomedical data is often a dynamic affair. As new subjects are recruited or examined, new data becomes available and the capacity to compare this new data to an existing database is a valuable one. In the context of biomedical data, this enables data from a patient of unknown status to be compared with a database of known examples. The property of generalisation - the ability to efficiently project new data onto an existing mapping - is therefore vital.

This is easily done with visualisation techniques which, as part of the mapping process, provide an explicit transformation from data to feature space. The GTM and SOM are examples of such algorithms: the LSS does not, in its original form, provide an easy way to project novel data. Sections 5.1 and 5.2 examine the implementation of the LSS in the framework of a feedforward neural network in order to provide the capability to generalise.

The use of prior information to improve mappings has already been discussed (Sections 3.1.1.3 and 4.2.2.1). Prior information may also be incorporated into the LSS in a variety of ways. The same is not true of the GTM, and a simple modification to the algorithm is given in Section 5.3 to allow the algorithm to do so. The effect of the use of prior information on generalisation ability is also examined.

The chapter concludes with a demonstration of generalisation and the use of prior information for the visualisation of HRV data from a study into sleep apnoea.

5.1 Feedforward networks for least-squares scalings

One of the major limitations of least-squares scalings is the inability to project new data. Demartines and Hérault (1997) suggest a local optimisation of the objective function for projecting new data. Here, the new data point x_0 is pooled with the existing data and the cost

function re-minimised, but moving only the corresponding map point y_0. However, this strategy is not suitable if prior information has been used during map construction, since this information will not be known for the new data sample.

A popular approach to this problem is to employ a universal approximator to learn the functional form of the mapping from data to feature space. This has been done using neural networks, in the form of multilayer perceptrons (van Wezel *et al.*, 1997; Mao and Jain, 1995; Tattersall and Limb, 1994) and radial basis function networks (Tipping, 1996; Lowe and Tipping, 1995; Webb, 1995; Lowe, 1993). Radial basis function networks are linear with respect to the output layer weights, which offers some computational advantages in training. These networks are the primary focus of attention in the next few sections.

5.1.1 Training the network

Two basic strategies are available for training a network to perform a LSS.

The most obvious is to break the process into two steps: generate a mapping directly (using appropriate optimisation techniques) and then train the network to reproduce this mapping using conventional supervised learning. This technique has been used to reproduce nonmetric multidimensional scalings with linear and quadratic models (Cox and Ferry, 1993).

The second strategy is to combine the construction of the map and the estimation of the network parameters into a single step. The general form of the cost function for the LSS is reproduced here:

$$E = \sum_i \sum_{j<i} h_{ij} (d_{ij}^* - d_{ij})^2 \quad , \tag{5.1}$$

where d_{ij} is the distance between point $\boldsymbol{y}_i$ and point $\boldsymbol{y}_j$ in the map, d_{ij}^* is the dissimilarity between data vectors $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, and h_{ij} is a weighting function.

Consider the radial basis function network which gives an output $\hat{\boldsymbol{y}}_i$ for input $\boldsymbol{x}_i$ according to the equation:

$$\hat{\boldsymbol{y}}_i = \sum_{k=1}^{K} \boldsymbol{w}_k \boldsymbol{g}_k (\|\boldsymbol{x}_i - \boldsymbol{c}_k\|) \, . \tag{5.2}$$

The network contains K hidden units, the kth of which is centred on $\boldsymbol{c}_k$, has transfer function g_k and has an associated output weight vector $\boldsymbol{w}_k$.

By differentiating 5.1 with respect to the network weights $\boldsymbol{w}_k$, one obtains:

$$\frac{\partial E}{\partial \boldsymbol{w}_k} = \sum_i \frac{\partial E}{\partial \hat{y}_i} \cdot \frac{\partial \hat{y}_i}{\partial \boldsymbol{w}_k} . \tag{5.3}$$

The first term on the right hand side is identical to that in Equation 3.11 in the original LSS formulation. The second term can be evaluated from Equation 5.2:

$$\frac{\partial \hat{y}_i}{\partial \boldsymbol{w}_k} = g_k(\|\boldsymbol{x}_i - \boldsymbol{c}_k\|) . \tag{5.4}$$

Thus, the network weight update equation has been formulated so that the construction of the map and network training are performed concurrently (the second strategy mentioned above). This technique was first suggested by Lowe (1993) and dubbed "relative supervision".

An extensive study of such networks has been made (Tipping, 1996; Lowe and Tipping, 1995; Lowe, 1993); however, it is not clear from this previous work whether relative supervision offers better network generalisation performance than do supervised (two step) training methods. Tipping (1996) made an explicit comparison of the generalisation performance of relative supervision with that of supervised training and concluded that relative supervision offered superior generalisation performance. The relative supervision algorithm was shown to be "self-regularising" - that is, the algorithm contains inherent mechanisms which encourage a smooth network and therefore good generalisation performance. However, Tipping did not consider the effect of including regularisation mechanisms in the supervised training methods. It is suggested here that the inclusion of such mechanisms should yield smoother networks and therefore better generalisation performance. This may put the generalisation performance of supervised methods on a par with that of relative supervision.

Section 5.2 compares the generalisation performance of the relative supervision algorithm with that of supervised methods, with particular attention to the regularisation mechanisms in both approaches. The principles behind regularisation and good generalisation are

reviewed briefly in the following section and then examined in the context of the least-squares scaling in Section 5.2.2.

5.2 Regularisation

The link between smooth transformations and good generalisation can be illustrated with a simple example. Say we wish to train a network to model the function $y = f(x)$, given a set of noisy training observations shown by the circles in Figure 5.1. The figure also shows three interpolations of the data, representing a highly smooth, moderately smooth and highly unsmooth network transfer function. The term "network smoothness" is often used to describe the smoothness of the network transfer function. Decreasing smoothness offers an increased flexibility in the network, so training set error will also decrease - the unsmooth network fits the training samples closely.

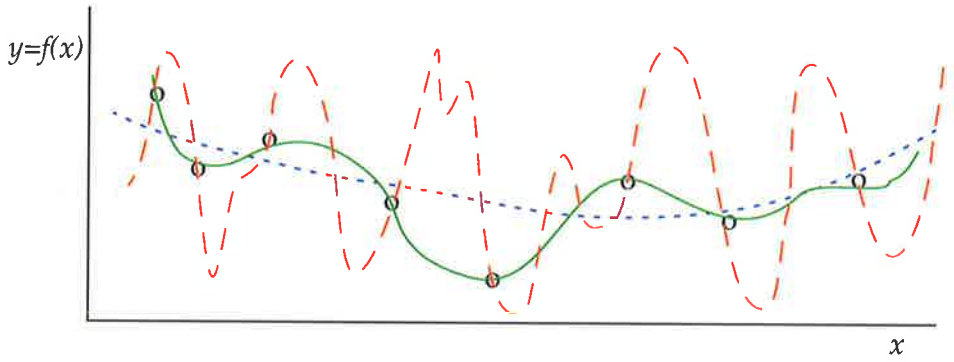

Figure 5.1: Training data (circles), with a highly smooth (dotted), smooth (solid) and highly unsmooth (dashed) fitted curve.

However, when used to project new data - essentially *interpolating* between training data points - both the highly smooth and unsmooth networks will often perform poorly. Finding the correct level of smoothness is a difficult problem which depends heavily on the nature of the data.

The complexity of the network structure plays an important role in ensuring that the network transformation is smooth. A small number of hidden units in the network will restrict the complexity of the transformation and is thus a form of smoothing. Employing wide bases in RBF networks has an equivalent effect, since this causes the responses of the units to become correlated and the effective number of units is reduced. Choosing more or

narrower bases reduces the smoothness of the network. Encouraging a smooth network transformation through choice of network structure is known as structural stabilisation.

It is also possible to include additional mechanisms in the network training which prevent overfitting. The training of feedforward neural networks is commonly carried out in a least-squares framework and a term may be added to the network cost function which penalises unsmooth transformations:

$$E_{reg} = E + \nu\Omega. \tag{5.5}$$

Here, Ω is the penalty term, the scalar ν controls the extent to which the penalty term contributes to the error function, and E is the usual squared-error term. The effect of the penalty term is to trade model fit (error) for smoothness.

The penalty term Ω may take the form of weight decay (or ridge regression):

$$\Omega_{wd} = \sum_{k=1}^{H} \|w_k\|^2, \tag{5.6}$$

operating on the premise that large weights results in unsmooth networks. One may also penalise the network curvature directly (Bishop, 1993), using:

$$\Omega_c = \sum_n \sum_m \sum_i \frac{\partial^2 y_n(m)}{\partial x_n^2(i)}. \tag{5.7}$$

Regularisation may also take the form of training with noise or early stopping (to prevent the network settling exactly on the training data) or other variants. A more detailed treatment may be found in Bishop (1995).

5.2.1 Assessing regularisation

Determining the optimal level of smoothness of the network - and therefore the optimal level of regularisation - is a difficult, data-dependent problem. The network training error offers no useful information in this regard since it is minimised by an overfitted network. A function which penalises both training error and model complexity may be of some use, but requires knowledge of the weighting to apply to the two sources of penalty.

The most common method of determining the optimal regularisation is through the use of a test set. The performance of the network is evaluated on a set of data, separate from that used for training, and the regularisation chosen to maximise performance on this test set. A large test set should be used in order to accurately assess the network performance - if data is in limited supply, then cross-validation may be used. A smaller subset of data is withheld for testing, but the training process is repeated for several different choices of test subset and the results averaged.

The criterion used to assess network performance may take various forms. The objective function of a mapping (evaluated on the test set) is a natural choice. If no such function exists (as for the SOM) but class-labelled data are available, then classification could be used as the criterion of test set performance. The justification for using classification to measure visualisation performance has been presented in Section 3.1.

5.2.2 Regularisation in least-squares scaling networks

Regularisation must be incorporated into all networks in order to provide good generalisation performance. The intent of this section is to compare the generalisation performance of the relative supervision algorithm with the generalisation performance of supervised training techniques.

The overriding source of regularisation in the relative supervision algorithm is structural stabilisation, determined by the number and width of basis function units in the network. A demonstration of this is given on the iris data set. This is well known data, used widely in the pattern classification literature as a benchmark. The data comprises four measurements made on three different types of irises, with 50 examples of each type of iris. A network was trained using relative supervision, with as many basis functions as training data. The final values of the error function (Equation 5.1) for the training and test sets (also known as training stress and test stress) are shown with varying basis width in Figure 5.2.

The gradual rise in training stress with narrowing bases (increasing σ) is likely to be due to the decrease in smoothness of the network causing the mapping to become trapped in local minima (see Section 5.2.2.2). Accompanying this graceful degradation of training stress is an alarming rise in test stress, indicating that the network generalises poorly with narrow bases. As the basis widths become narrow, the correlation among the basis function units decreases and the smoothness of the network is reduced.

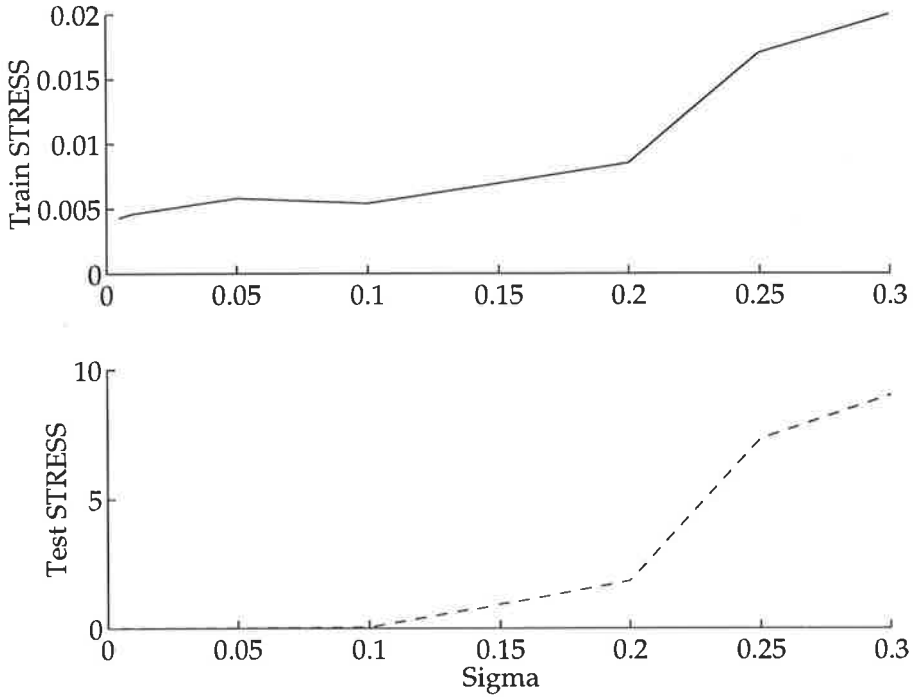

Figure 5.2: Training and test stress with narrowing basis functions in the relative supervision algorithm. As the width of the Gaussian basis functions is reduced (increased σ) the test stress (and thus generalisation error) rises sharply. This illustrates the role of structural stabilisation in the regularisation of the network.

The relative supervision algorithm also makes use of another, less obvious, regularisation mechanism: translation of solutions.

5.2.2.1 Regularisation by translation

The cost function (Equation 5.1) for a least-squares scaling is clearly unaffected by an arbitrary rotation or translation of the mapping. Tipping (1996) showed that the relative supervision algorithm tends to favour solutions with small network weights and suggested that the small network weights were achieved by exploiting the rotation- and translation-invariant nature of the least-squares cost function. Consider an RBF network which has been trained to produce a multidimensional scaling. The norm of the RBF weight matrix can easily be shown to be invariant to rotations of the output mapping. However, translating the solution does affect the norm of the weight matrix (Tipping, 1996, reproduced in Appendix C) and so it is possible to reduce the magnitude of the network weights through translation of the solution. Reducing the magnitude of the network weights while keeping the widths of the

basis functions fixed has the effect of increasing the smoothness of the network transformation.

The translation of solutions in order to reduce network weights is an unorthodox concept and is unique to networks which perform least-squares scalings. Of particular note is the fact that translation does not penalise the fit of the network to the training data. In this sense, translation differs significantly from conventional regularisation techniques. Thus far, regularisation by translation has only been found as an inherent mechanism of the relative supervision algorithm. One may then ask if it is possible to employ translation of the solutions as a form of regularisation in supervised training.

Tipping (1996) showed that it is possible to determine analytically the translation of a given mapping which will minimise the norm of the weight matrix for a given rotation of that mapping. However, Tipping claimed that the globally minimum value of the weight matrix norm is dependent on the map rotation, and not possible to determine analytically. Tipping argued that an explicit search of map rotations would be needed in order to find the network weight matrix with the smallest norm. Such a search would be a difficult, nonlinear optimisation problem and may well be practically infeasible.

In fact, the minimum value of the norm of the weight matrix does *not* depend on the rotation of the mapping. Since translation may be applied to each dimension of the output mapping in turn, it is convenient to analyse the effects of translation on each column of the weight matrix (corresponding to the output in each dimension) separately. While it is quite true that the minimum possible value for the norm of each *column* of the weight matrix depends on the rotation of the solution [this was shown by Tipping (1996)], the minimum value of the norm of the weight matrix *as a whole* is invariant to rotation, and depends only on the translation of the solution. This is shown in Appendix C.

In principle then, it is possible to incorporate translation of solutions in a supervised LSS network training scheme. It is not necessary to search rotations of the mapping in order to minimise the weights in the network and the optimal translation of the mapping may be determined analytically using Equation C.7, Appendix C. This translation should be applied at the completion of the regular training procedure. Note that this result only applies to networks which are linear with respect to the weights and so a similar trick with a multilayer perceptron would be difficult to implement.

Example 1

The use of regularisation-by-translation with a supervised training method was evaluated experimentally, using the posture/propranolol data from Section 4.2. Data from all subjects was used in order to provide sufficient training examples for the network. For simplicity, the problem was reduced to a two-class problem, pooling the drug data and concentrating on the changes with posture.

The generalisation performance of the relative supervision algorithm was compared to a network trained using a supervised method, but incorporating translation of the solution as described above. Identical networks (number and width of basis functions) were used for both training methods.

For the supervised training, the two-dimensional mapping $\boldsymbol{Y}$ was first generated independently of the network, using the pinning algorithm (Demartines and Hérault, 1997). A radial basis function network was constructed, using as many basis functions as training examples. Quite wide bases ($\sigma = 1 \times 10^{-3}$) were used to impose structural stabilisation and thus ensure network smoothness.

For a radial basis function network with K basis functions $\{\varphi_1 \ldots \varphi_K\}$, the network Jacobian is given by:

$$\boldsymbol{J} = \begin{bmatrix} \varphi_1(\boldsymbol{x}_1) & \varphi_2(\boldsymbol{x}_1) & \ldots & \varphi_K(\boldsymbol{x}_1) \\ \varphi_1(\boldsymbol{x}_2) & \varphi_2(\boldsymbol{x}_2) & \ldots & \varphi_K(\boldsymbol{x}_2) \\ \ldots & \ldots & \ldots & \ldots \\ \varphi_1(\boldsymbol{x}_N) & \varphi_2(\boldsymbol{x}_N) & \ldots & \varphi_K(\boldsymbol{x}_N) \end{bmatrix}. \tag{5.8}$$

The pseudoinverse $\boldsymbol{J}^+$ of the network Jacobian was used to compute the network weights needed to produce the desired output $\boldsymbol{Y}$:

$$\boldsymbol{W} = \boldsymbol{J}^+\boldsymbol{Y}. \tag{5.9}$$

The translation of the solution $\boldsymbol{Y}$ required to minimise the network weights was then calculated using Equation C.7, and the weight matrix adjusted accordingly.

The relative supervision algorithm was used to train a network with the same number and width of basis functions.

The smoothness of each network was assessed on the basis of the curvature of the network transfer function, using Equation 5.7. The generalisation performance of the two methods was evaluated on the basis of the test stress and also the classification score. These were computed using a test set comprising 170 samples each of supine and tilt data and averaged over 50 runs. Results are shown in Table 5.1.

Training method	Weight matrix norm	Curvature	Training stress	Test stress	7-nearest-neighbour classification	Bayes' likelihood ratio classification
Relative supervision	$1.3\text{x}10^{22}$ ($7.0\text{x}10^{21}$)	$5.2\text{x}10^{8}$ ($2.5\text{x}10^{8}$)	0.067 (0.0088)	0.098 (0.015)	0.59 (0.049)	0.60 (0.036)
Supervised training	$4.7\text{x}10^{21}$ ($2.6\text{x}10^{21}$)	$1.9\text{x}10^{8}$ ($1.3\text{x}10^{8}$)	0.11 (0.015)	0.13 (0.013)	0.61 (0.055)	0.61 (0.057)

Table 5.1: Comparison of networks trained with relative supervision and with supervised algorithms, using the posture/propranolol data set. Values are mean(std). The network trained in a supervised manner had less curvature, but higher stress.

The results show that the magnitude of the network weights and the network curvature were smaller for the supervised training method. The absolute value of the weight matrix norms was very large: this is a consequence of using very wide basis functions. The results also show that the training and test stress values were significantly better for the relative supervision algorithm, indicating a better representation of the relationships between the input data. The classification performance (assessed using both a nearest-neighbours and Bayes' likelihood rule classifier) was the same for both.

Two possible explanations may be forwarded to explain the larger stress for the supervised method compared to that obtained with relative supervision. The first explanation stems from the network smoothness, since both the network weights and curvature were lower in the supervised method than with relative supervision. The supervised network may have been insufficiently flexible and therefore unable to attain a good map configuration. The second possible explanation is that the pinning method used to construct the target map for the supervised training may simply not have reached as deep a minimum of the stress function as did relative supervision.

Examination of the stress scores for the target maps (before network fitting) revealed that the latter explanation was indeed true. The pinning algorithm simply generated target maps

with larger stress values than did relative supervision. However, this does not explain why the supervised method yielded smoother networks than relative supervision. Both methods used the same number and width of basis functions, so the difference in smoothness was not due to structural complexity.

A further investigation was made. Each of the network weight matrices given by the relative supervision training method was translated explicitly using Equation C.7. In each case, this translation resulted in a reduction of the norm of the weight matrix, with an average reduction of about 10%.

So, while relative supervision may employ translation in order to minimise the network weights, it does not necessarily exploit this mechanism fully. A reduction in the network weights was obtained by applying the translation equation at the completion of relative supervision training.

Example 2

A similar comparison of relative supervision and the fully supervised training method was carried out to ensure that the results of the first example were not simply a peculiarity of the data. In this second comparison, the iris data set was used. Twenty five examples of each of the three types of iris were used for training and the remainder for testing. The results are shown in Table 5.2.

Training method	Weight matrix norm	Curvature	Training stress	Test stress	5-nearest-neighbour classification
Relative supervision	$3.2\text{x}10^{21}$ ($6.5\text{x}10^{21}$)	$3.9\text{x}10^{16}$ ($8.4\text{x}10^{16}$)	0.0045 (0.0007)	0.0070 (0.0018)	0.96 (0.020)
Supervised training	$2.0\text{x}10^{21}$ ($3.7\text{x}10^{21}$)	$2.6\text{x}10^{16}$ ($4.8\text{x}10^{16}$)	0.054 (0.024)	0.056 (0.024)	0.95 (0.022)

Table 5.2: Comparison of networks trained with relative supervision and with supervised algorithms, using the iris data set. Values are mean(std). Stress values were quite different, although classification performance was about the same for both.

With this data, the relative supervision again achieved lower values for both the training and test stress than did the supervised method. The size of the weights and the curvature of the networks was about the same. Both achieved test stress values which were comparable to

the training values, suggesting that generalisation was good in both cases. It is interesting to note that the classification performance (using the 5-nearest-neighbours classifier) was similar for both, despite the fact that relative supervision had a much lower test stress. Since the iris data has well separated clusters, there are large distances which tend to dominate the stress measure. The higher stress measure for the supervised procedure was probably indicative of errors in the large inter-cluster distances. Since such errors have little effect on the interpretability of the resulting map, there was little difference in the resulting classification performance.

5.2.2.2 Map construction

The previous section has explored the regularisation involved in the relative supervision algorithm. The results indicate that the generalisation performance which can be obtained using supervised training methods is comparable to that from the relative supervision algorithm. However, the process of generating the mapping has not yet been addressed.

With the relative supervision algorithm, map construction is performed concurrently with network training. In the purely supervised approach to training, map construction is a separate step which must be carried out prior to the training of the network. It is here that relative supervision holds an undeniable advantage over supervised methods. The surface of the error function is fraught with local minima and presents a difficult optimisation problem. The smoothness imposed by network regularisation in the relative supervision algorithm reduces the effects of local minima in the stress cost function, aiding the construction of the map (Tipping, 1996). In a two-step procedure, network regularisation cannot help in this manner since map construction occurs independently of network training. Additionally, if conventional regularisation mechanisms such as weight decay or early stopping are used with supervised training, the fit of the network to the training data will be further penalised. For this reason, the training stress in Tables 5.1 and 5.2 are higher for the supervised network.

In contrast, regularisation by translation does not penalise the fit of the network to the training data. However, it has already been noted that structural stabilisation (through the use of wide bases) contributes to the regularisation of the relative supervision algorithm. A consequence of using wide bases is that the network Jacobian may become ill-conditioned, which can cause the training of the network to effectively cease before reaching a minimum (Tipping, 1996). Thus, relative supervision does not necessarily attain the global minimum of the stress function, but the smoothness imposed by the network aids the construction of the

map and generally achieves a lower training stress than does map construction by direct means.

The pinning algorithm (Demartines and Hérault, 1997) is probably the most appropriate tool if the map is to be constructed independently of network training. The algorithm is quite fast ($O(N)$) and allows temporary increases in the error function, offering some hope of escaping local minima. The change in error at time t is given by:

$$\Delta E(i) = -\alpha(t) \sum_{j \neq i} \left[\left\| \nabla_j E_{ij} \right\|^2 + \sum_{p \neq i,j} (\nabla_j E_{pj})^T (\nabla_j E_{ij}) \right], \tag{5.10}$$

where i is the index of the "pinned" pattern. The error will increase if the second term takes a sufficiently negative value. As might be expected, in practice the algorithm is somewhat sensitive to the choice of step size α. Too small a step size does not allow the algorithm to jump out of local minima, but too large a step size makes it difficult for the algorithm to settle into any minimum at all.

5.2.2.3 Other considerations

The computational cost is an important consideration for the practical application of any algorithm.

Each iteration of the relative supervision algorithm requires the computation of $N(N-1)/2$ derivatives, an unattractive scaling behaviour if large data sets are involved (N is the number of input patterns).

In this regard, the supervised (two-step) method is potentially more appealing, as a fast method of generating the mapping may be used (such as pinning). The cost of supervised training must then be considered, which will vary depending on the regularisation mechanism chosen.

For the relative supervision algorithm, Tipping (1996) recommends the choice of a large number of wide bases. This is a simple strategy for ensuring network smoothness, but one which may be unnecessarily wasteful of computational resources, particularly with large data sets. Smaller networks may be achieved without loss of performance, by careful choice of basis function centres and widths.

A method common of achieving this in supervised RBF network training is to perform some manner of cluster analysis, such as k-means clustering (Moody and Darken, 1989), in order to find a small set of hidden units which provide a good representation of the distribution of the input data. However, since visualisation is essentially a clustering procedure itself, this seems a somewhat circuitous approach.

Another method is that of orthogonal least squares (Chen *et al.*, 1991; Chen *et al.*, 1989). Here, basis functions are added iteratively from a pool of candidate units, choosing the unit which gives the best improvement in network error at each iteration. The procedure is terminated when a desired level of network performance (assessed by, for example, cross-validation) is achieved. This method is not applicable to the relative supervision algorithm since it requires the target training points in order to assess the improvement in network error offered by each candidate basis function.

From a practical viewpoint, the use of the relative supervision algorithm is quite straightforward (using, for example, the "shadow-targets" implementation (Tipping and Lowe, 1997) and offers good generalisation performance. However, the algorithm can become cumbersome with large data sets. Supervised methods offer greater flexibility and the potential for smaller networks, but require an increased attention to practical details. The construction of the mapping can be particularly difficult.

5.2.3 Multilayer perceptron

The discussion in Section 5.2.2 has focused exclusively on radial basis function network implementations. In a multilayer perceptron implementation, the network is trained using a variant of the backpropagation algorithm, presenting pairs of input patterns at a time. The rotation and translation required to minimise the network weights would be extremely difficult to find, since the network is not linear with respect to the weights. Regularisation is achieved through structural stabilisation (Mao and Jain, 1995). Smaller networks reduce the complexity of the network transfer function, offering improved generalisation as compared with highly complex networks. The choice of network size is "an important yet difficult problem" (Mao and Jain, 1995), since a too-small network will give an overly smooth transfer function and will not adequately capture the essence of the data.

5.2.4 Regularisation and the generative topographic mapping

In contrast to the least-squares scaling, generalisation is an integral part of the GTM algorithm. The GTM defines a parametric transformation from latent (low-dimensional) space to data space, and a set of regularly-spaced grid points in latent space. Each grid point is transformed to give the centre of a distribution in data space, and training consists of adjusting the parameters of the transformation to fit the sum of these distributions to the training data. For the visualisation of a new point in data space, Bayes' theorem is used to invert the transformation and find the corresponding distribution in latent space. This posterior distribution can be condensed into its mean or mode for convenient, simultaneous visualisation of a number of data points.

The parametric transformation mentioned above is defined by a set of basis functions Φ and a weight matrix $\boldsymbol{W}$, and the topographic nature of the GTM is ensured by choosing a small number of relatively wide bases, in order to provide a smooth, continuous transformation from latent to data space.

The regularisation of the GTM is provided principally through the choice of basis functions, although a weight-decay term may be included during training if desired (Bishop *et al.*, 1998b). The choice of basis functions and regularisation coefficient may be made through trial and error, although a recent extension of the GTM algorithm offers a method of controlling the smoothness in a more elegant manner through Bayesian formalism (Bishop *et al.*, 1998a).

5.2.5 Conclusions

Radial basis function networks are a popular means of adding the capability of generalisation to the least-squares scaling algorithm. Training the networks may be done concurrently with the map construction, or in a separate step. The former approach is known as relative supervision. This is generally easy to use and offers good generalisation performance, but the computational demands can become prohibitive for large data sets. These demands arise from the map construction itself, which requires the evaluation of $N(N-1)/2$ derivatives at each iteration of the algorithm (N is the number of input patterns). Separating the steps of map construction and network training offers the possibility of using a more efficient technique for the map construction.

In Section 5.2, the generalisation performance of the two training methods was compared. Regularisation in the relative supervision algorithm is provided by structural stabilisation and, to some extent, translation of solutions. Translation is a novel regularisation mechanism which reduces the magnitude of the network weights but which, unlike conventional regularisation, does not penalise the fit of the network to the training data.

The possibility of incorporating regularisation-by-translation into supervised training methods has previously been raised (Tipping, 1996). The translation which is required to achieve a minimum value of the weight matrix norm (for a particular map rotation) may be found analytically. However, it was previously thought (Tipping, 1996) that the minimum value of the weight matrix norm was dependent on the rotation of the solution. Thus, if one wished to incorporate translation as a mechanism of regularisation into a supervised method, a search of map rotations would be needed in order to find the minimum weight matrix norm. This would be a computationally expensive search procedure, and so regularisation by translation was considered an impractical approach.

In Section 5.2.2.1 it was shown that the globally minimum value of the weight matrix norm is *not* dependent on the rotation of the solution. It is therefore feasible to incorporate this form of regularisation into networks which are trained in a supervised manner. It was also shown that relative supervision does not necessarily make full use of the translation mechanism, and that a reduction in network weight amplitudes may be obtained by explicitly applying translation after relative supervision training.

It was concluded in Section 5.2.2.1 that, provided appropriate regularisation is incorporated into each approach, supervised training methods can be expected to offer comparable generalisation performance to relative supervision.

However, the construction of least-squares scalings through conventional optimisation is a difficult task. Combining the steps of map construction and network training (as with relative supervision) eases this task significantly, since the network smoothness reduces the chance of becoming trapped in suboptimal minima in the map cost function. The trade-off is the scaling behaviour of the algorithm with the number of input data, which can become prohibitive. Separating the two steps offers the possibility of using a faster method of constructing the map (such as pinning).

5.3 Prior information

The concept of incorporating prior knowledge during map construction was introduced in Section 3.1.1.3. This prior knowledge is assumed to be known only for the training data and commonly takes the form of class labels. The incorporation of this information generally serves to clarify or enhance clusters which are not well resolved in a purely unsupervised mapping.

Poor cluster separation in the map may be due to several factors: the most obvious is that the data is simply non-separable, consisting of overlapping clusters in data space. Alternatively, the data may be inherently of dimension larger than two, and unable to be represented adequately in a two-dimensional mapping, or it may be two-dimensional, but exist on a nonlinear manifold in higher dimensional space which cannot be adequately "unfolded" by the basic unsupervised scaling algorithm.

The incorporation of class information during map construction can help to unravel strongly folded data, producing a more useful mapping. An example is given on a mapping of two rings, interlinked in three-dimensional space as depicted in Figure 5.3a. A typical unsupervised Sammon mapping (Figure 5.3b) does not untangle the two rings, and the rings are distorted as a result. The inter-point distance reconstruction plot is shown in Figure 5.3c. The same data is mapped using an element of supervision, accomplished by replacing the data dissimilarities d_{ij}^* in Equation 5.1 with δ_{ij}:

$$\delta_{ij} = (1-\alpha)\, d_{ij}^* + \alpha s_{ij}, \tag{5.11}$$

where α is a scalar controlling the degree of supervision, and

$$s_{ij} = \begin{cases} 0 \text{ for data points } i \text{ and } j \text{ belonging to the same ring (class)} \\ 1 \text{ otherwise} \end{cases}. \tag{5.12}$$

The mapping with $\alpha = 0.5$ (Figure 5.3d) shows clear separation of the rings, although some distortion remains. The interpoint distances reconstruction plot (Figure 5.3e) indicates that local structure is better represented in the partially supervised mapping.

In this example, the data occupy distinct volumes in three dimensional space with no cluster overlap. It is the intricate nature of the cluster arrangement which prevents the

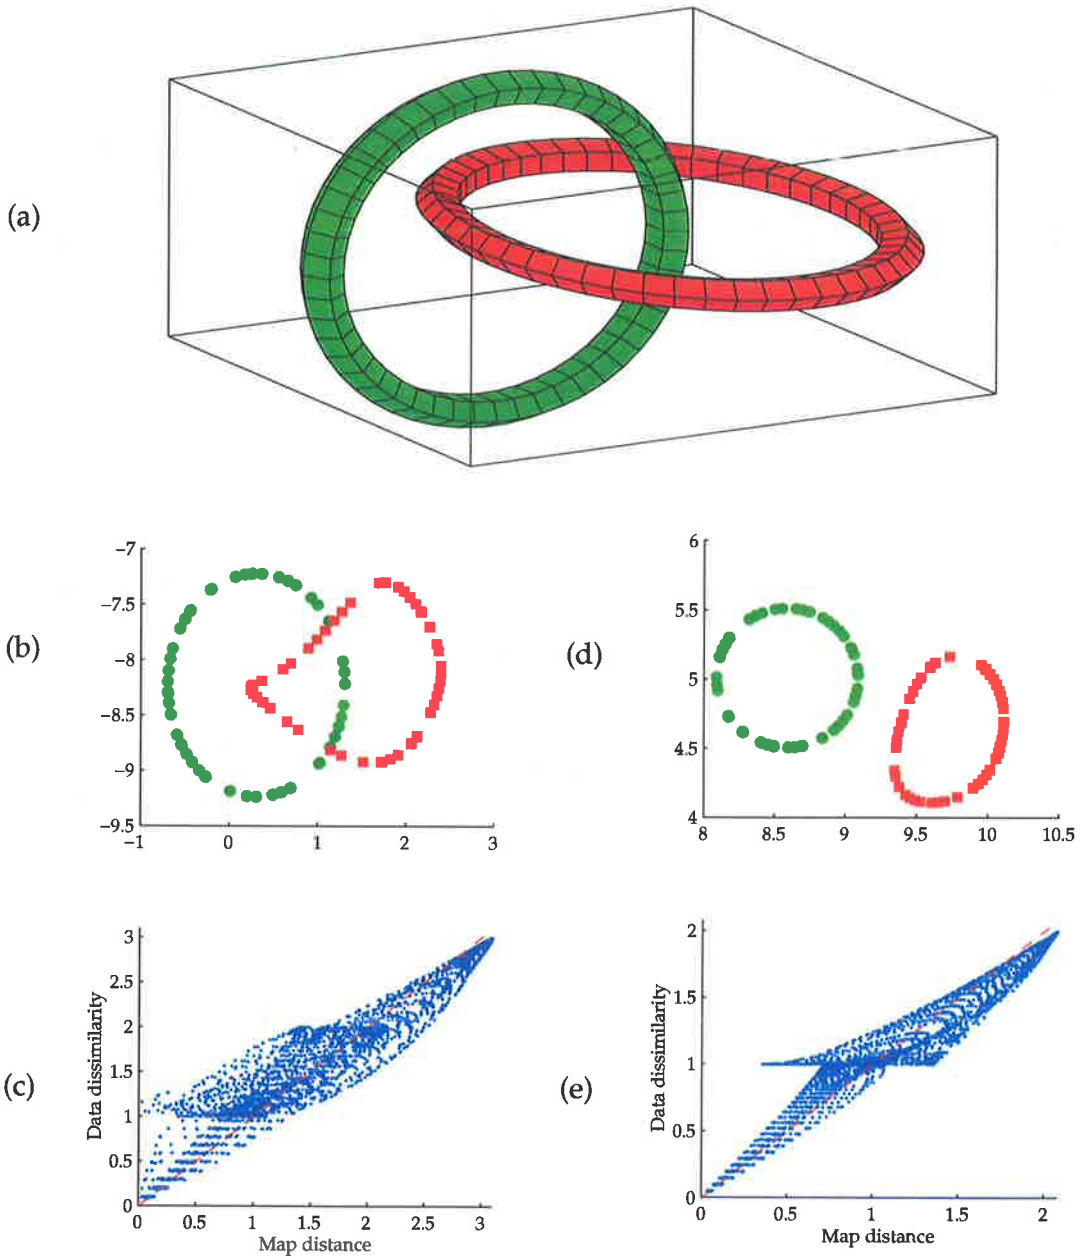

Figure 5.3: (a) Two interlinked rings; (b) Unsupervised Sammon mapping and (c) the associated interpoint distances reconstruction plot; (d) Partially supervised Sammon mapping ($\alpha = 0.5$) and (e) the interpoint distances reconstruction plot.

unsupervised algorithm from separating the two clusters in the mapping. The interlinked rings are also an example of data which cannot be unambiguously represented in two dimensions, and thus can be considered have an "inherent dimensionality" greater than two. The use of prior information is one method of forcing the mapping into a preferred

orientation (one which is more meaningful to the observer, even though it may be a numerically inferior solution by the unsupervised rule).

Prior information regarding the data may also take the form of knowledge of the scale at which interesting structure exists in the data. This information may be incorporated through the use of an alternative weighting function h_{ij} in the cost function (Equation 5.1). Most commonly, weightings which emphasise local structure in the mapping are used. Demartines and Hérault (1997) advocate the use of a weighting function which is a non-increasing function of the distance in map space d, such as the simple step function:

$$h_{ij} = \begin{cases} 1 & \text{if } d_{ij} \leq \lambda \\ 0 & \text{otherwise} \end{cases}. \tag{5.13}$$

A mapping of the interlinked rings using this weighting is shown in Figure 5.4. Demartines and Hérault claim that this weighting function is the "only one to unfold strongly folded data". Since the weighting is a function of the output distances, the solution is recursive. This can make the map construction considerably more difficult.

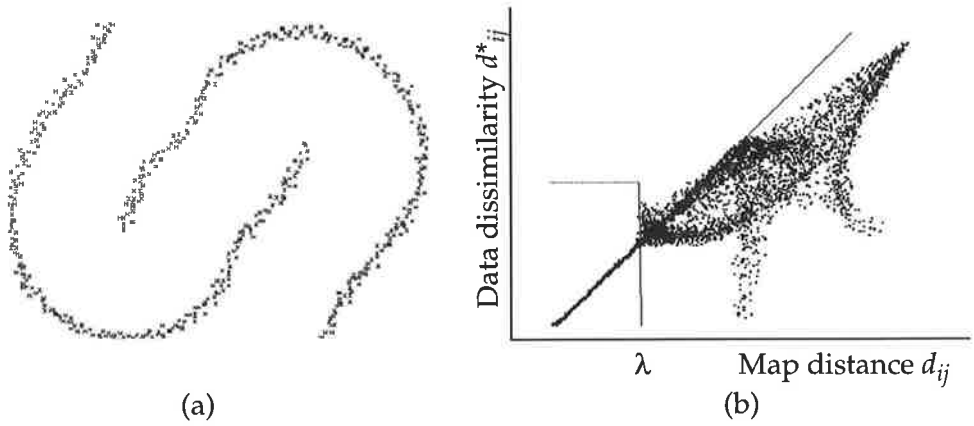

Figure 5.4: (a) An unsupervised least-squares scaling of the two interlinked rings, using a modified weighting function; (b) The interpoint distances reconstruction plot. Local structure (distance less than λ) is preserved well in the mapping. Reproduced from Demartines and Hérault, 1997.

An alternative weighting strategy, dubbed "consequential region scaling" (Lee, 1999), emphasises moderate distances since these are most likely to connect points on the boundary of the same cluster.

5.3.1 Prior information and the generative topographic mapping

The generative topographic mapping (GTM) algorithm (Bishop *et al.*, 1998a; Bishop *et al.*, 1998b) relies on prior information in the selection of the number of latent grid units, the number and width of basis functions, and the regularisation coefficient. However, the basic formulation of the algorithm does not offer any mechanism for incorporating class-label knowledge. A modification to the original algorithm is given below, in which such information may be included by assigning a class "preference" to each grid unit *a priori*.

5.3.1.1 The modified algorithm

Assume we have two classes of data, with a class label $c_n \in \{0, 1\}$ for each vector $\boldsymbol{x}_n$ in the training data. The prior information takes the form of a scalar $0 \leq \gamma_i \leq 1$ for each grid unit i. If a particular unit j is chosen to prefer data with class label $c = 0$, this is incorporated by assigning $\gamma_j = 0$. Units k with no particular class preference are assigned $\gamma_k = 0.5$; units l chosen to prefer class 1 are assigned $\gamma_l = 1$. Intermediate values may be used.

The distribution of data x, given component i, weights W and variance $1/\beta$ is the same as in the original formulation of the GTM:

$$p(\boldsymbol{x}|i,\boldsymbol{W},\beta) = \{\beta/(2\pi)\}^{D/2} e^{-\{\beta\|\boldsymbol{m}_i - \boldsymbol{x}\|^2/2\}} . \tag{5.14}$$

The class-conditional distribution in data space is given by summing over all components:

$$p(\boldsymbol{x}|c,\boldsymbol{W},\beta) = \sum_{i=1}^{I} P(i|c)\, p(\boldsymbol{x}|i,\boldsymbol{W},\beta) , \tag{5.15}$$

where the priors $P(i|c)$ reflect the prior "preferences":

$$P(i|c) = (1/I)\,|1 - c - \gamma_i| . \tag{5.16}$$

In the case of equal priors ($\gamma_i = 0.5, \forall i$), this reverts to the standard GTM formulation. The priors need not be constant throughout the algorithm and could initially be chosen to be equal to ensure global ordering, before evolving with time to reflect the class-label information.

The joint likelihood is given by:

$$p(\boldsymbol{x},c|\boldsymbol{W},\beta) = P(c|\boldsymbol{W},\beta)\,p(\boldsymbol{x}|c,\boldsymbol{W},\beta)\,, \tag{5.17}$$

and

$$P(c|\boldsymbol{W},\beta) = N_c/N\,. \tag{5.18}$$

The parameters $\boldsymbol{W}$ and β are estimated by maximising the expected log-likelihood:

$$\langle L(\boldsymbol{W},\beta)\rangle = \log \prod_{n=1}^{N} p(\boldsymbol{x}_n,c_n|\boldsymbol{W},\beta)\,, \tag{5.19}$$

which, under the usual assumption of independent data, becomes:

$$\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N} \log p(\boldsymbol{x}_n,c_n|\boldsymbol{W},\beta)\,. \tag{5.20}$$

In keeping with the original formulation of the GTM algorithm, this likelihood is maximised using the expectation-maximisation (EM) algorithm. At each iteration of the algorithm, the expected log-likelihood is evaluated using available parameter estimates and then new parameters $\boldsymbol{W}_{new}$ and β_{new} are chosen to maximise this expectation.

The expected log-likelihood (Equation 5.20) may be differentiated with respect to $\boldsymbol{W}$ and set to zero, giving:

$$\sum_{n=1}^{N}\sum_{i=1}^{I} \{\boldsymbol{W}_{new}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\}\,\Phi^T(\boldsymbol{u}_i)\,R_{ni} = 0\,, \tag{5.21}$$

with:

$$R_{ni} = \frac{|1 - c_n - \gamma_i| \exp\left(-\frac{\beta}{2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\right)}{\sum_{j=1}^{I} |1 - c_n - \gamma_j| \exp\left(-\frac{\beta}{2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_j) - \boldsymbol{x}_n\|^2\right)}, \tag{5.22}$$

evaluated at each iteration using the available estimates of $\boldsymbol{W}$ and β.

Similarly, differentiating $\langle L(\boldsymbol{W},\beta)\rangle$ with respect to $1/\beta$ leads to the update equation:

$$\frac{1}{\beta_{new}} = \frac{1}{ND}\sum_{n=1}^{N}\sum_{i=1}^{I} R_{ni}\left\|\boldsymbol{W}_{new}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\right\|^2. \tag{5.23}$$

The complete derivation of these equations is given in Appendix D. These update equations are in fact identical to those of the standard GTM algorithm, with the exception of the definition of R_{ni}.

Once the parameters $\boldsymbol{W}$ and β have been estimated, new data may be visualised by computing:

$$p(i|\boldsymbol{x},\boldsymbol{W},\beta) = \frac{p(\boldsymbol{x}|i,\boldsymbol{W},\beta)\,P(i)}{p(\boldsymbol{x})}, \tag{5.24}$$

with:

$$P(i) = \sum_{c=1}^{C} P(c)\,P(i|c). \tag{5.25}$$

5.3.1.2 Demonstration

The standard generative topographic mapping applied to the interlinked rings gives the mappings in Figures 5.5a and 5.5b, for strong and weak regularisation respectively. The smoothness imposed by the strong regularisation does not allow the rings to become separated. For this particular example, the unsupervised GTM algorithm actually produces quite sensible and useful mappings.

Figures 5.5c and 5.5d show the GTM with prior information incorporated into the mapping. The units i on the left half of the grid were allocated to the ring shown with green circles ($\gamma_i = 1$), the units j on the right half of the grid to the ring shown with red squares ($\gamma_j = 0$). With weak regularisation (Figure 5.5c), the points become well-separated by class, but with severe distortion (particularly for the green ring). Moderate regularisation (Figure 5.5d) reduces both the distortion and the cluster separation.

This simple example not only serves to demonstrate the modification to the algorithm, but also shows its principal weakness: the inclusion of the prior information is done in a such a

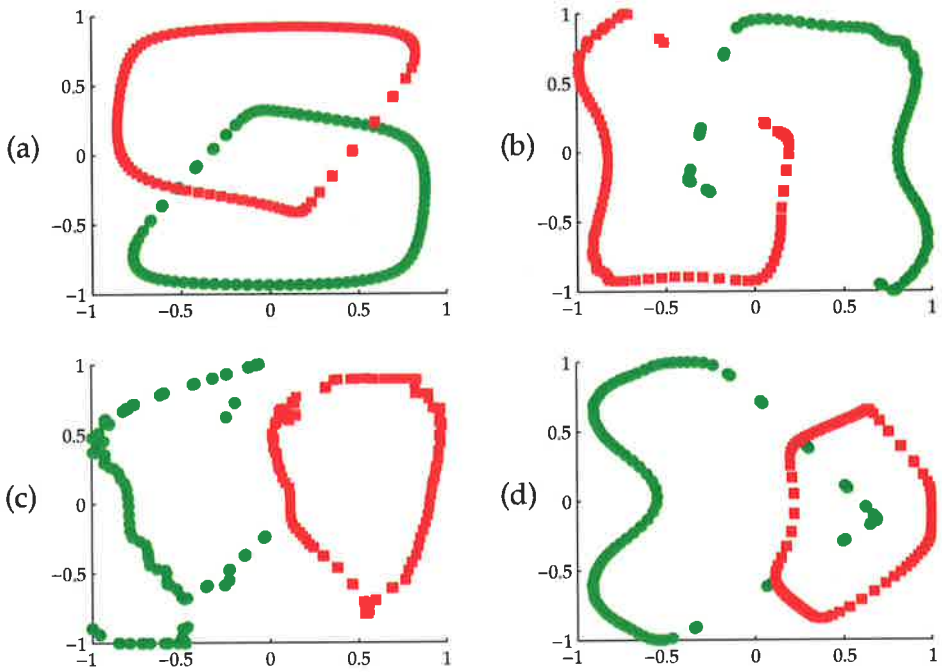

Figure 5.5: (a) Equal priors (standard GTM), strong weight-decay regularisation ($\nu = 0.1$) **; (b) equal priors, weak regularisation** ($\nu = 1 \times 10^{-6}$) **; (c) priors** $\gamma_i = 1$ **for units** i **on the left half of the grid,** $\gamma_j = 0$ **for units** j **on the right half of the grid, and weak regularisation** ($\nu = 1 \times 10^{-6}$) **; (d) priors as for (c) and stronger regularisation** ($\nu = 5 \times 10^{-4}$) **.**

way as to encourage a *specific* output configuration (in this case, one cluster on each side of the grid). This suggests that the desired cluster arrangement is known in advance. For complex data sets (particularly with multiple classes), it may be difficult to pose a suitable prior distribution. One also runs the risk of creating an artefactual cluster arrangement which does not reflect the true topography of the data.

An alternative method of incorporating prior information into the GTM, and an avenue for further study, is to include a penalty term in the likelihood equation which encourages cluster separation. This penalty term could be formulated as a function of map distances, encouraging large inter-class distances and small intra-class distances, in a similar manner to the method of incorporating prior information into the least-squares scaling.

5.3.2 Prior information and generalisation

Poor cluster separation in the mapping may be indicative of data which simply overlaps in data space. Incorporating class information in this case will result in an artefactual separation

of the clusters and any network trained to produce such a mapping will generalise poorly since, by necessity, the training must overfit in the region of the cluster overlap.

Interesting data from real-world situations is likely to be drawn from a nonlinear manifold and also likely to show true cluster overlap. Incorporation of prior information may help to untangle the manifold and so an improvement in generalisation performance may occur with moderate influence from prior information. With insufficient regularisation, one would expect to see a decline in generalisation performance with strong influence from prior information, since the network would begin to overfit the training data in the region of the cluster overlap. If the smoothing effect of regularisation were too strong there may be no increase in generalisation at all, since the incorporation of prior information is likely to require a decrease in network smoothness, particularly near cluster boundaries (Figure 5.6).

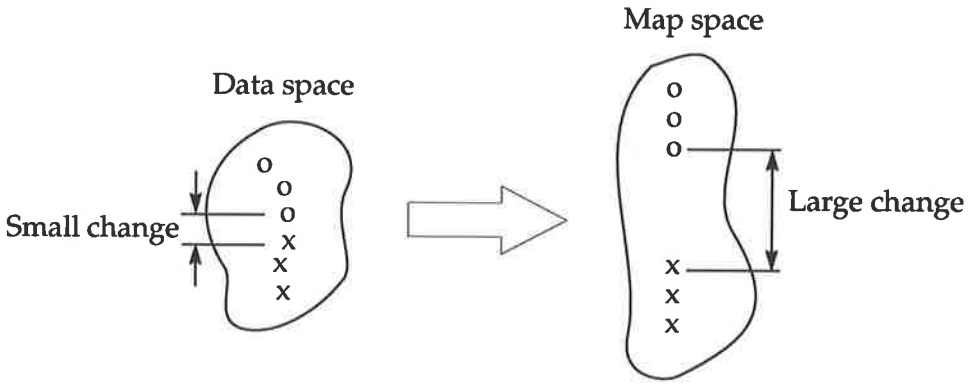

Figure 5.6: In order to enhance cluster separation in feature space, the mapping must be less smooth near the cluster boundaries: a small change in data space location gives a relatively large change in the feature location.

Finding the optimal level of prior influence is a difficult problem which depends heavily on the nature of both the data and the prior information. The purely unsupervised mapping should naturally lead to a fairly smooth network since points which are similar in data space will be close to each other in the mapping. A purely supervised mapping is likely to require a less smooth transformation. The hybrid mapping, incorporating elements of both, allows a gradual trade-off of smoothness for map interpretability. Thus, the inclusion of prior information is linked with the network regularisation and both are probably most conveniently determined on the basis of validation set performance.

5.3.3 Conclusions

The incorporation of prior knowledge during map construction is one method of improving the separation of clusters in the least-squares scaling. A similar effect may be obtained with the GTM, using the algorithm modification suggested above. These techniques are of use in untangling data on a nonlinear manifold. Care must be taken not to create artefactual cluster separations if good generalisation performance is desired.

5.4 Case study 2: Sleep apnoea

This section uses data from a sleep apnoea study (Hilton *et al.*, 1998) to demonstrate generalisation of mappings in a biomedical context, using both the least-squares scaling and the generative topographic mapping. The use of class-label information during map construction is also demonstrated on the least-squares scaling.

5.4.1 The sleep apnoea/hypopnoea syndrome

The sleep apnoea/hypopnoea syndrome (SAHS) is respiratory disorder which occurs during sleep. Apnoea describes the complete cessation of airflow during sleep, while hypopnoea is a reduction in airflow by 50% or more. While infrequent apnoea episodes are quite common in normal subjects, recurrent episodes are abnormal. The traditional definition of sleep apnoea is five or more episodes of blocked airflow in an hour, each lasting ten seconds or longer (Parkes, 1985). Sleep apnoea may be central or obstructive in origin: central apnoea stems from a problem with the central respiratory drive and is characterised by a lack of respiratory movement. Obstructive apnoea is the most common form of apnoea and is caused by closure of the upper airway. The chest and diaphragm move with changes in intrathoracic pressure induced by respiratory effort, but there is no airflow at the nose or mouth. Obstructive apnoea is particularly prevalent amongst males, the obese and older individuals, and may be linked with other disorders: systemic hypertension has been reported in up to 50% of apnoea sufferers (Fletcher, 1995).

Each episode of apnoea is terminated by an arousal, often accompanied (in the case of obstructive apnoea) with a loud snort as the airway is cleared. Sleep is fragmented by these frequent arousals, so patients often display daytime drowsiness, difficulty in concentrating and headaches. In severe cases, patients may develop cardiovascular complications such as

right sided congestive heart failure. Sleep apnoea has also been identified as a contributor to cot death (Hudgel, 1986; McCaffree, 1986; Gordon *et al.*, 1984).

Sleep apnoea may be reliably detected by polysomnography in a specialised sleep laboratory. Typically, polysomnography involves the recording of EEG (electroencephalogram - electrical activity of the brain), EMG (electromyogram - electrical activity of the chin muscle) and EOG (electrooculogram - movement of the eye) to aid in identification and staging of sleep, breathing pattern (flow or temperature sensors at the nose and mouth), snoring intensity (microphone), ventilatory effort (by oesophageal balloon or by strain gauge around chest), ECG and arterial oxygen saturation (by an oximeter placed on the ear or finger). Such an array of sensors requires significant setup time, incurs high cost of disposables, and requires maintenance throughout the night to ensure that all the sensors are properly positioned. The cost of polysomnography has been estimated at £126 per study (Bradley *et al.*, 1995). There is strong interest in more convenient screening procedures for sleep apnoea.

> "The laboratory time and expense involved in performing polysomnography has prompted investigators to search for less demanding methods of detecting apnea that will allow screening of large populations for epidemiologic data." (Fletcher, 1986, p. 6)

Oximetry is currently a widely accepted alternative to polysomnography for apnoea screening. The interruption in breathing caused by apnoea also causes a decrease in the arterial oxygen saturation and in turn the oxyhaemoglobin saturation (SaO_2) (Hudgel, 1986). Oximetry is an inexpensive, noninvasive method and typically shows excellent specificity (fraction of normal subjects diagnosed correctly) but often poor sensitivity (fraction of apnoea subjects diagnosed correctly) (Ryan *et al.*, 1995; Douglas *et al.*, 1992). Hypopnoeic episodes which cause arousal from sleep but not significant oxygen desaturation will not be detected through oximetry (Ryan *et al.*, 1995). Many subjects who have undergone oximetry are therefore further investigated with polysomnography (Hilton *et al.*, 1998).

During all types of apnoea, the heart rate has been observed to slow (bradycardia), with the extent of slowing strongly related to the decrease in oxyhaemoglobin saturation (Zwillich *et al.*, 1982). Arousal is accompanied by sympathetic stimulation and an increase in heart rate (tachycardia) (Leuenberger *et al.*, 1995). This cyclic change in the heart rate associated with apnoea has long been recognised and has been used as the core of various screening methods

(Keyl *et al.*, 1997; Tilkian *et al.*, 1978). This pattern is not specific to sleep apnoea (Hudgel, 1986) and so in isolation, a given brady-tachycardia cycle is not specific indicator of SAHS. However, repeated occurrences throughout the night is a strong indicator of sleep apnoea (Hilton *et al.*, 1999).

The respiratory component of HRV may also offer useful information for apnoea screening. During central apnoea, there is a loss of respiratory drive as well as respiratory effort, and so no RSA component would be expected in the HRV signal. During obstructive apnoea, the neural signals controlling respiration remain intact, as does respiratory effort, and the RSA component of HRV will persist.

5.4.2 Methods

Polysomnography and Holter monitoring were performed on 40 male subjects, with a mean (standard deviation) age of 40(10). Each polysomnography record was scored blinded by a technician for apnoeas and hypopnoeas. Heart rate data was taken from a 15-minute segment during light sleep for each subject.

Spectra were estimated using a time-frequency representation with binomial kernel (Williams, 1996) and then averaged over non-overlapping 32 second segments. Segments known to contain an episode of apnoea were identified and collated and an equal number of control data were added, chosen at random from the control subjects. The resulting data set contained 108 averaged spectral densities, with spectra from normal sleep and from episodes of apnoea each numbering 54 samples.

A typical apnoea spectral density is shown in Figure 5.7, with the very low frequency activity clearly visible. An example of the spectral density of the HRV signal during normal sleep is shown in Figure 5.8.

The preprocessing of the data was similar to that used in Section 4.2. The respiratory sinus arrhythmia was isolated by taking the mean of each spectrum from 0.15-0.4 Hz. The remainder (0-0.15 Hz band) of each spectrum was smoothed using a windowed averaging function. Each final 39-dimensional data vector comprised 38 elements containing the smoothed spectral envelope from 0-0.15 Hz, and one element containing the RSA component. This last element was scaled by the square root of the number of frequency bins in the band 0.15-0.4 Hz, in order to maintain the contribution of the RSA to dissimilarity measures.

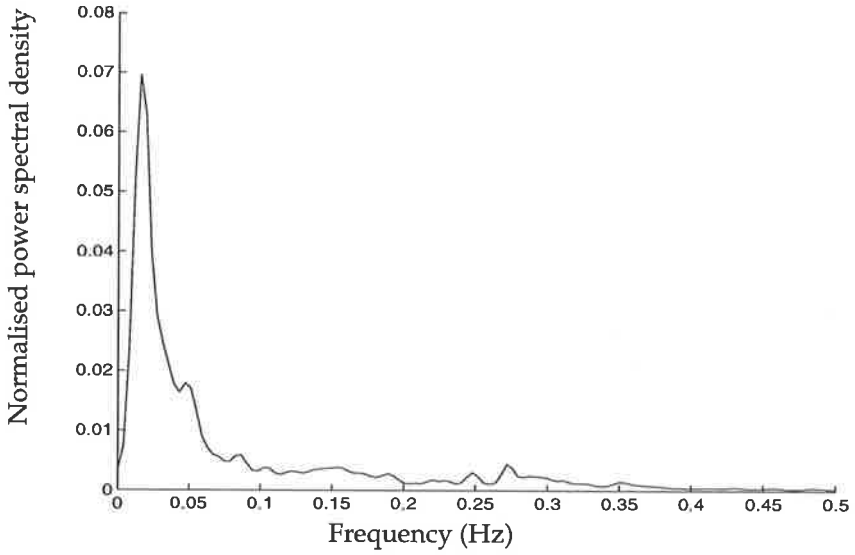

Figure 5.7: Typical spectrum associated with an episode of sleep apnoea. The very low frequency peak is caused by the brady-tachycardia cycle of the heart during apnoea.

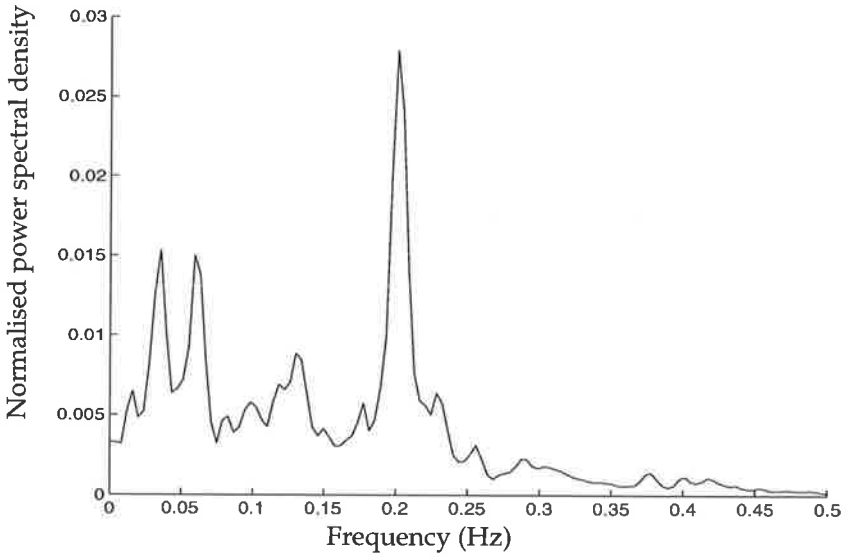

Figure 5.8: Example spectrum of the HRV signal during normal sleep. The respiratory component is clearly visible at about 0.2 Hz, along with lower-frequency components.

5.4.3 Results and discussion

The LSS and GTM were used to visualise the structural relationships in this data set.

5.4.3.1 Visualisation

The generative topographic mapping

The posterior-mean distribution of the data onto a 20x20 grid is shown in Figure 5.9. The

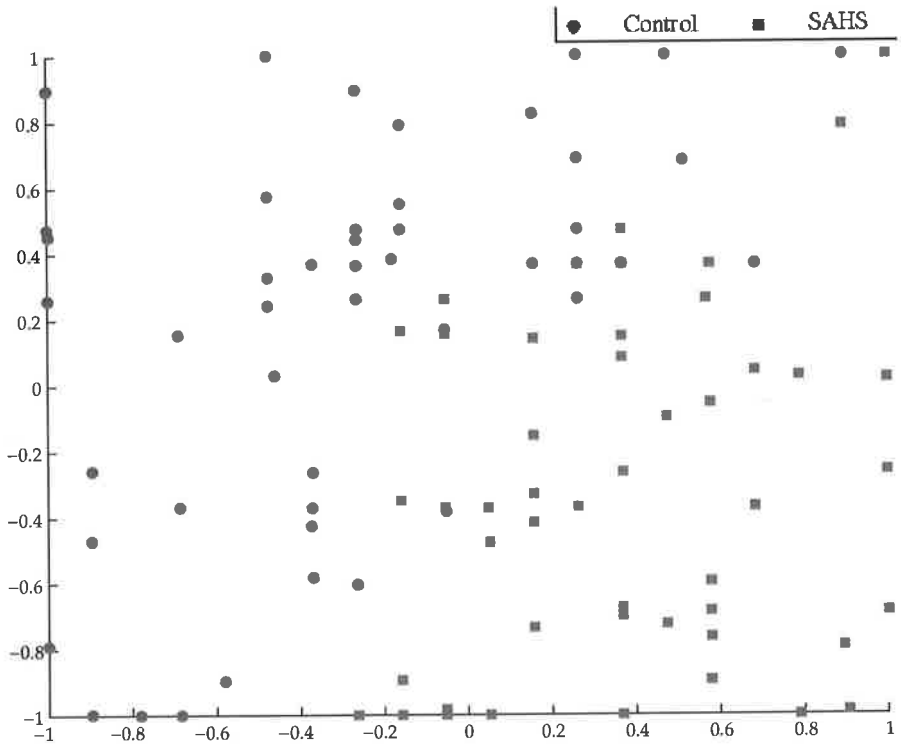

Figure 5.9: Generative topographic mapping of the non-normalised data on a 20x20 grid.

data were smoothed lightly in frequency (averaged across a sliding window of width 0.03 Hz, which was 7 frequency bins) but not normalised.

The map shows two clusters, with the control and apnoea data roughly separable by the rising diagonal (bottom left to top right). Figure 5.10 uses colour to indicate the distribution of the map points with respect to various characteristics of the input data: VLF power (taken as the mean of the PSD in the band 0-0.08 Hz), RSA power, the centre frequency of the spectrum in the band 0-0.15 Hz, and the 3dB bandwidth of the spectrum about this centre frequency. Figure 5.11 shows the distribution of the map points with respect to subject age, body mass index (ratio of weight to the square of height), neck circumference and oximetry score corresponding to each map point. Body mass index and neck circumference offer information regarding the obesity of the subject, and the oximetry (SaO_2) index offers a comparison with the results of oximetry, the most popular low-cost alternative to full polysomnography.

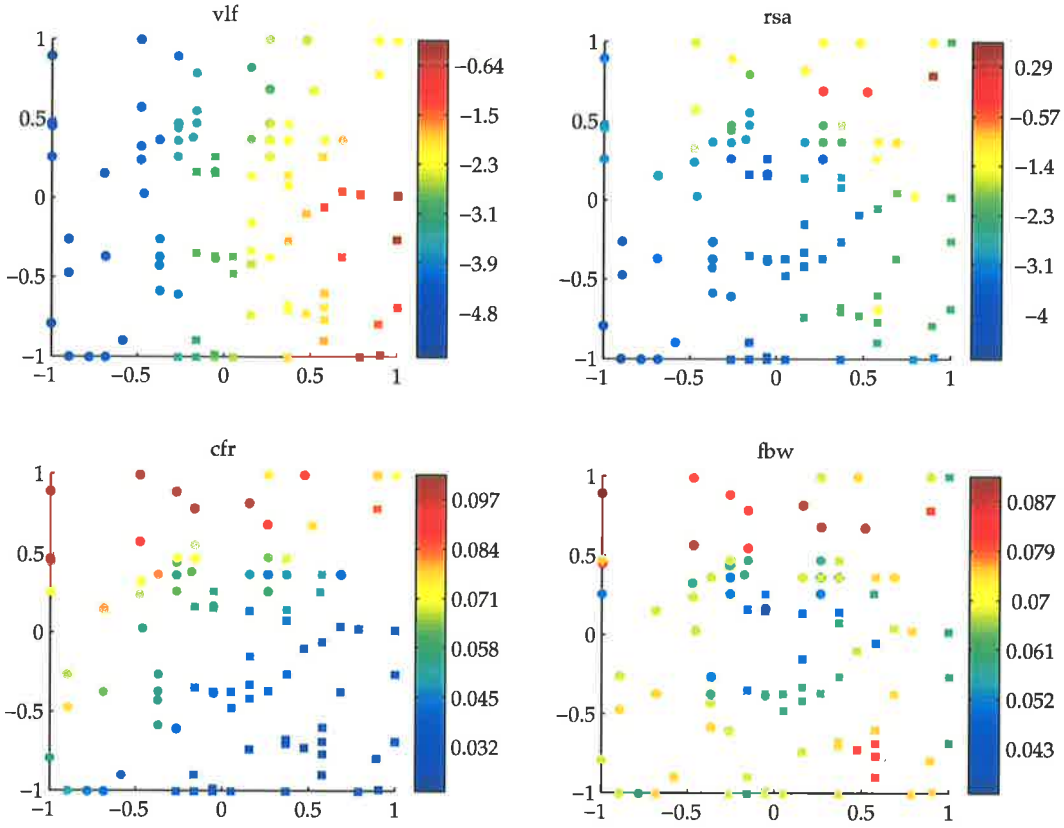

Figure 5.10: The distribution of the GTM points (for non-normalised data) with respect to log VLF power (vlf), log RSA power (rsa), centroid of PSD in band 0-0.15 Hz (cfr) and bandwidth about centroid (fbw).

As expected, the distribution with respect to VLF power shows a clear ordering of the data, with VLF power increasing with x-coordinate (correlation coefficient 0.97). In fact, the VLF power is almost sufficient to discriminate apnoea and control subjects: with the exception of a pocket of control subjects in the top right hand corner of the mapping, the subjects which display a (log) VLF power of greater than about -3 are SAHS subjects. Those control subjects which display a similarly large VLF component are all aged 35 or less (see the distribution with respect to age, Figure 5.10). It is known that total HRV power declines with age (Lipsitz *et al.*, 1990; O'Brien *et al.*, 1986) and so these control points show large VLF components by virtue of age rather than as an indicator of apnoea. The centroid of the PSD in the band 0-0.15 Hz is also higher for these subjects (about 0.08 Hz) than for apnoea subjects (about 0.05 Hz).

The distribution with respect to RSA power shows a weak radial ordering (lowest values in the lower left hand quadrant of the map). The centroid of the PSD in the band 0-0.15 Hz also showed a strong ordering with respect to map position.

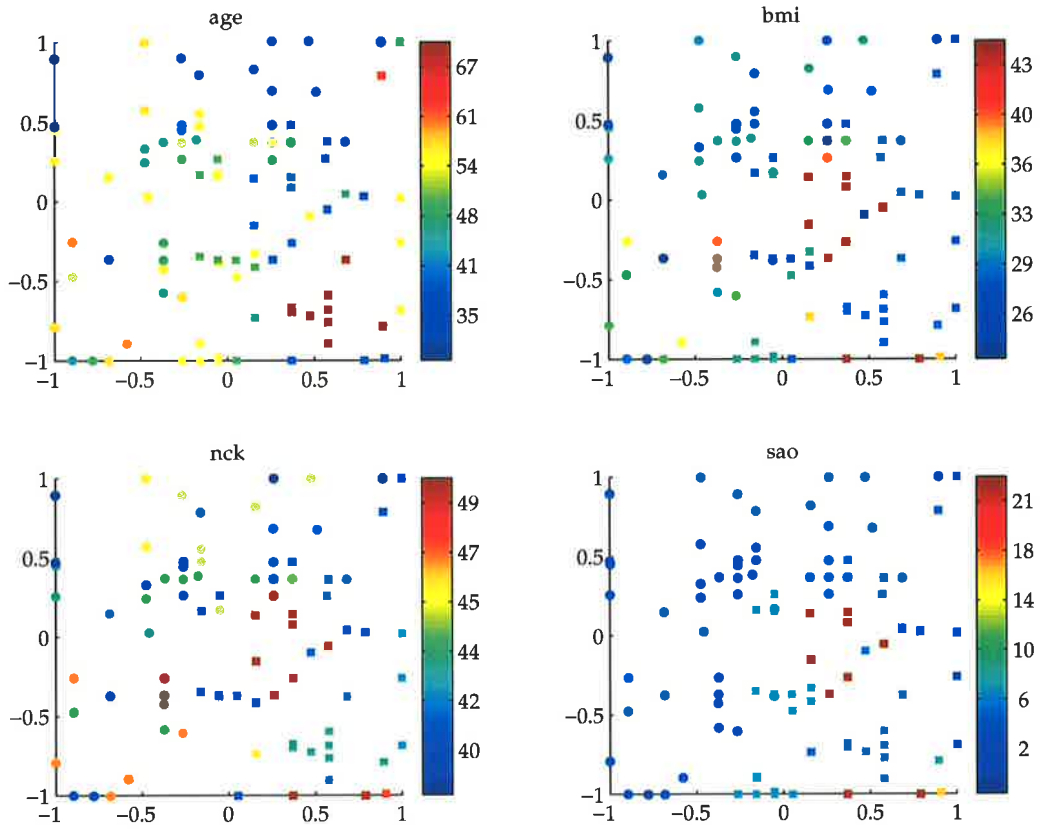

Figure 5.11: The distribution of the GTM points (for non-normalised data) with respect to age, body mass index (bmi), neck circumference (nck) and oximetry score (sao).

The distributions in Figure 5.11 offer an appreciation of the map structure with respect to variables not included in the input data. The distribution with respect to age does not show strong ordering but there is a tendency for similar ages to cluster together. Since no normalisation has been applied, this probably reflects the decline of total HRV power with age.

The BMI for control data shows a rough increase from the top right of the cluster to the bottom left. Higher values of BMI would suggest a decrease in physical fitness, which is accompanied by a decline in parasympathetic control with respect to fitter subjects. The ordering with respect to BMI is roughly opposite to that with respect to RSA, which is a marker of parasympathetic activity. No clear pattern with BMI is evident for the SAHS data.

Figure 5.11 also shows the distribution with respect to the SaO_2 index. By definition, control subjects have low SaO_2; the SAHS points show groupings of similar SaO_2 values, but the arrangement of these groupings does not show any striking pattern.

Normalised data

The data was then normalised with respect to the total power in each spectrum (excluding components below 0.05 Hz). The GTM projection is shown in Figure 5.12, again with clustering of control and SAHS points, but with apparently more overlap between clusters than was the case without normalisation. Several SAHS points are now surrounded by control points.

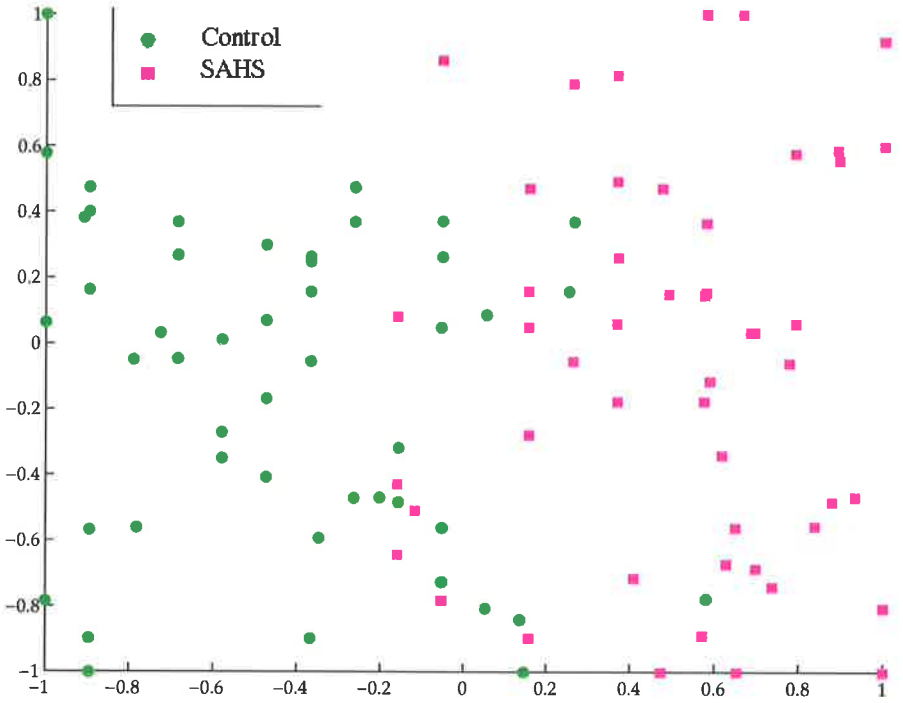

Figure 5.12: Generative topographic mapping of the apnoea data on a 20x20 grid. The data was normalised with respect to the total power (excluding VLF components).

The distribution with respect to VLF power (Figure 5.13) again correlates very strongly with x-coordinate. The structure with respect to RSA was much clearer than with no normalisation. The projection of each map point onto the falling (top left to bottom right) diagonal correlated with the RSA power (see the distribution with respect to RSA power in Figure 5.13). The normalisation reduced the effects of changes in the total power, allowing the RSA component an increased contribution to the structure of the mapping.

The centroid frequency and bandwidth showed a strong ordering with respect to the map points: SAHS subjects with wide, low-frequency components were mapped to the top of the

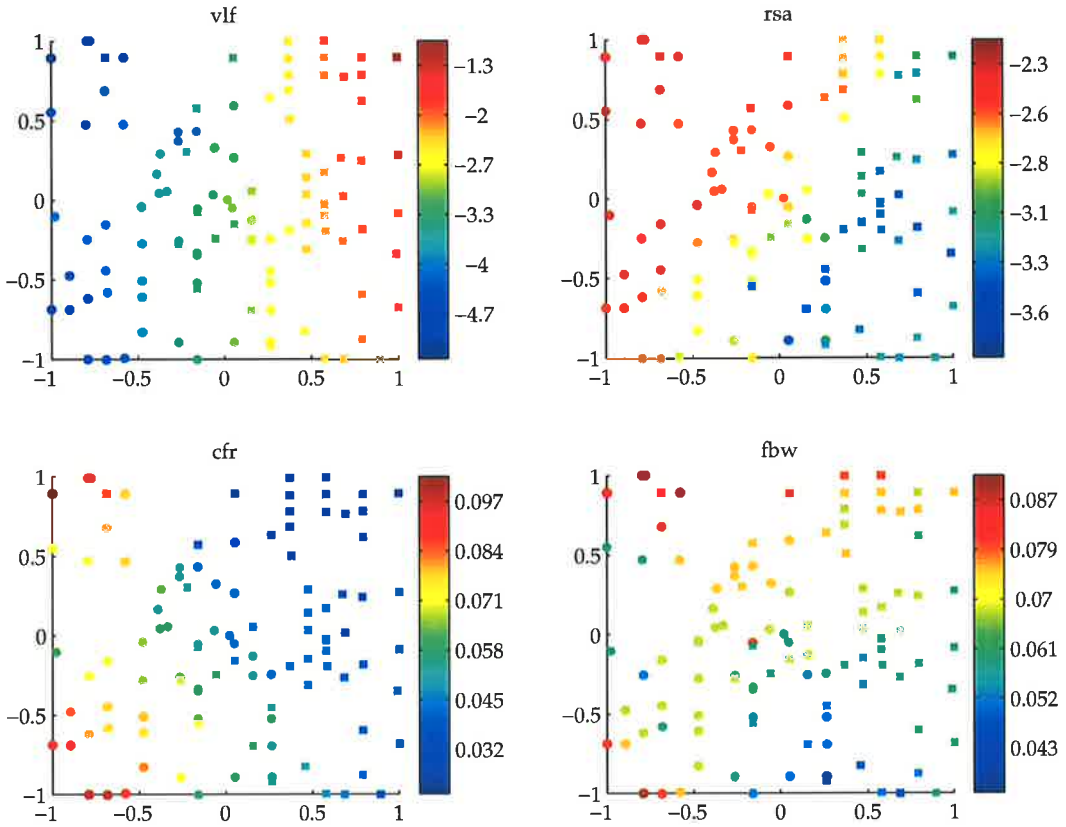

Figure 5.13: The distribution of the GTM points (for normalised data) with respect to log VLF power (vlf), log RSA power (rsa), centroid of PSD in band 0-0.15 Hz (cfr) and bandwidth about centroid (fbw).

cluster, and points at the bottom were characterised by narrow components at higher slightly frequencies.

The map also showed more interesting structure with respect to age, particularly for the SAHS points. Apnoea data were located on the right hand side of the map, at increasing y-coordinate with age. Total HRV is reduced with age, but the decline of (absolute) RSA with age is less than of non-respiratory components (Simpson and Wicks, 1988). Normalised RSA thus increases with age - following the pattern of the map.

The distribution of points with respect to BMI shows no obvious pattern for the SAHS data. Control subjects with high BMI tended to be located near the adjoining region of the two clusters.

The "mismapped" SAHS points (those apparently located in the control cluster) appear to fit the general patterns of age, VLF and RSA. However, the distribution with respect to the SaO_2 index (Figure 5.14) offers some clue to the reason for their location: all points on the left

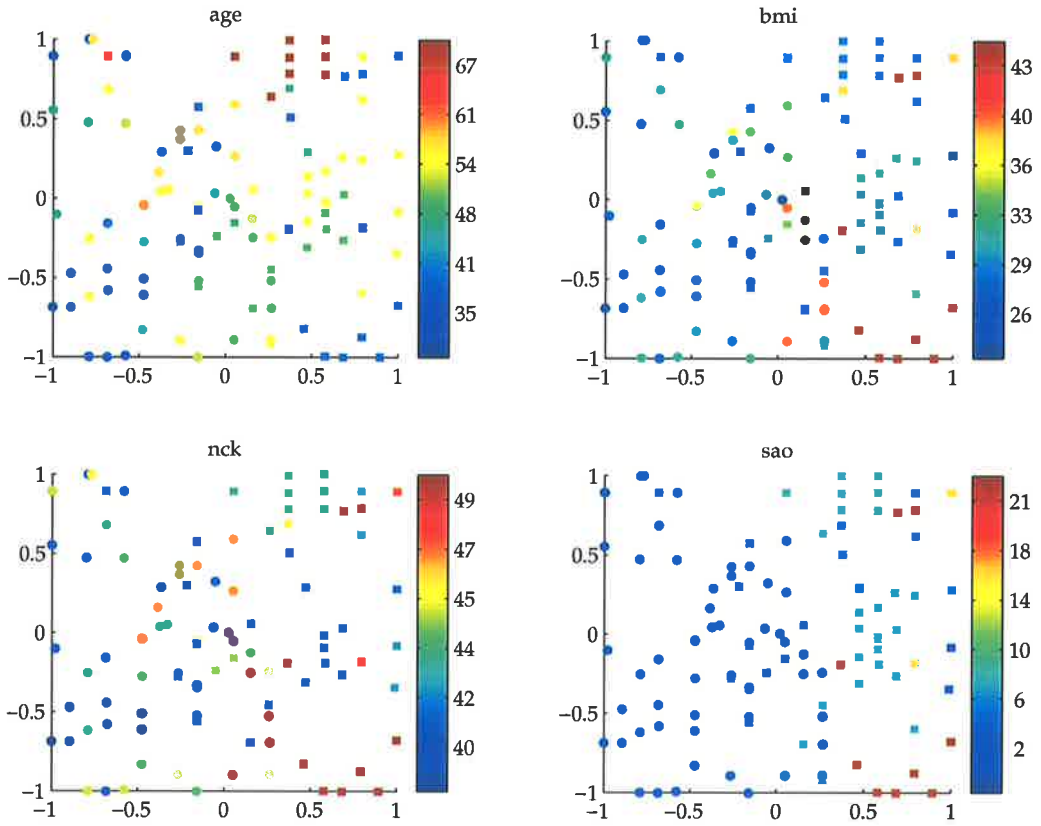

Figure 5.14: The distribution of the GTM points (for normalised data) with respect to age, body mass index (bmi), neck circumference (nck) and oximetry score (sao).

hand side of the map (including true controls, as well as these mismapped points) have low SaO_2 values. Oximetry is known to miss hypopnoeas and apnoeas which do not cause significant desaturation of the arterial oxyhaemoglobin since this is the sole information used by oximetry.

The extent of bradycardia during apnoea and hypopnoea is strongly related to the degree of oxyhaemoglobin desaturation (Zwillich *et al.*, 1982). Since these subjects did not show significant desaturation, the VLF component is correspondingly smaller. In addition, two of the mismapped points were from a subject who displayed a high number of hypopnoeas and relatively low number of apnoeas (clearly visible in Figure 5.14). During hypopnoea, the airflow is reduced but not completely blocked, giving a larger RSA than during apnoea. The

remaining mismapped SAHS points are from subjects with low body mass index and fit subjects will also show a more pronounced RSA component. Thus, the increased RSA (due to the normalisation) and the smaller VLF component (due to the decreased degree of desaturation) for each of these subjects causes them to be mapped amongst control points.

A simple classification scheme was used to objectively assess the cluster separation of the mappings with and without normalisation. Each mapping was trained on the complete data set and then each point in the mapping classified by voting k-nearest neighbours and Bayes likelihood rule classifiers. Note that this procedure does not assess generalisation properties since the classifiers were constructed and then applied to the same data. Low classification scores in this test simply indicate a high degree of overlap between the control and apnoea clusters in a mapping. The cluster-separation scores are tabulated in Table 5.3, and confirm the subjective observation that non-normalised data shows better cluster separation than normalised data.

Classifier	7-nearest-neighbour		Bayes' likelihood ratio	
	Specificity	Sensitivity	Specificity	Sensitivity
Non-normalised data	0.81	0.91	0.87	0.87
Normalised data	0.89	0.81	0.83	0.83

Table 5.3: Cluster separation for the GTM mappings of the normalised and non-normalised data, assessed by training set classification. The overlap between the control and SAHS clusters increases with normalisation.

The Sammon mapping

The form of the Sammon mapping was very similar to that of the GTM, for both normalised and non-normalised data. The Sammon mapping of the non-normalised data is shown in Figure 5.15, with additional information plots in Figures 5.16 and 5.17. The cluster separation is presented in Table 5.4; and again, is very similar to that seen in the GTM.

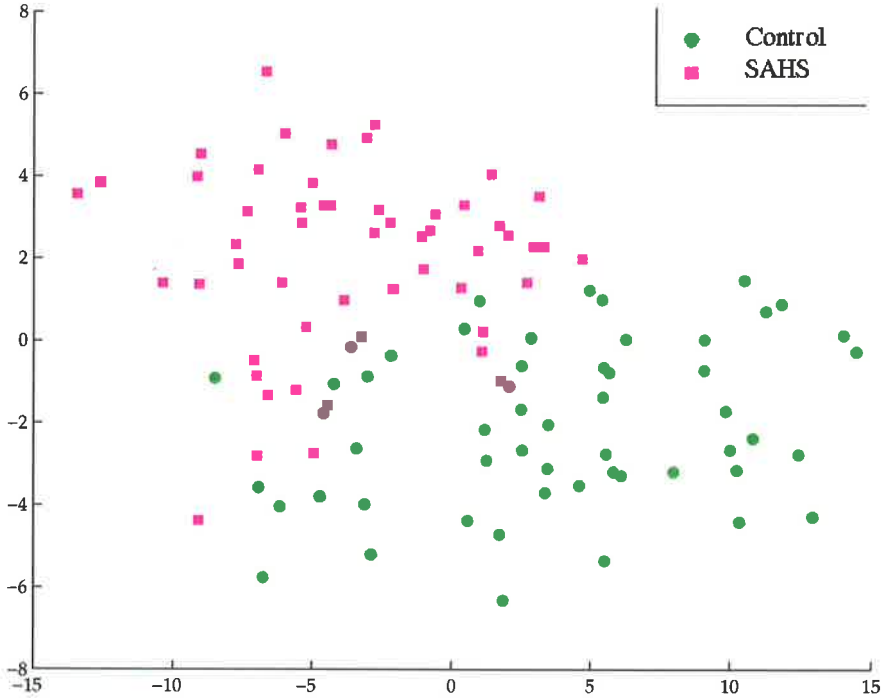

Figure 5.15: The Sammon mapping of the non-normalised SAHS data.

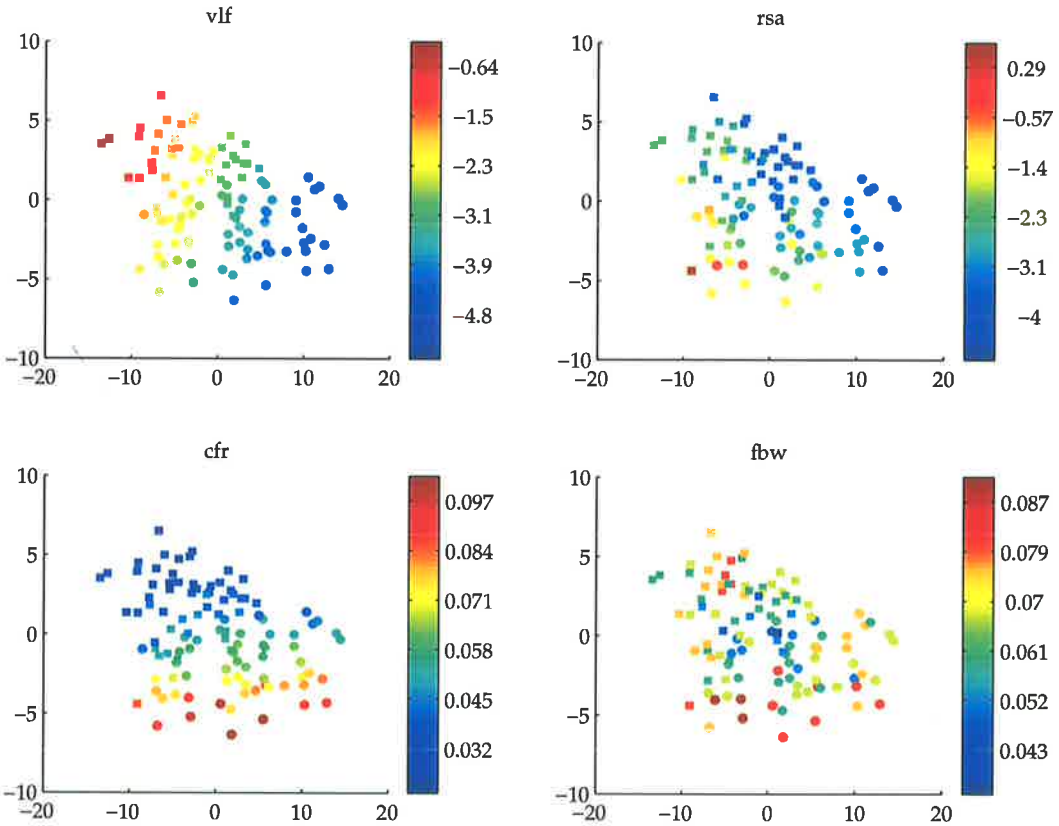

Figure 5.16: The distribution of the Sammon mapping (of non-normalised data) with respect to log VLF power (vlf), log RSA power (rsa), centroid of PSD in band 0-0.15 Hz (cfr) and bandwidth about centroid (fbw).

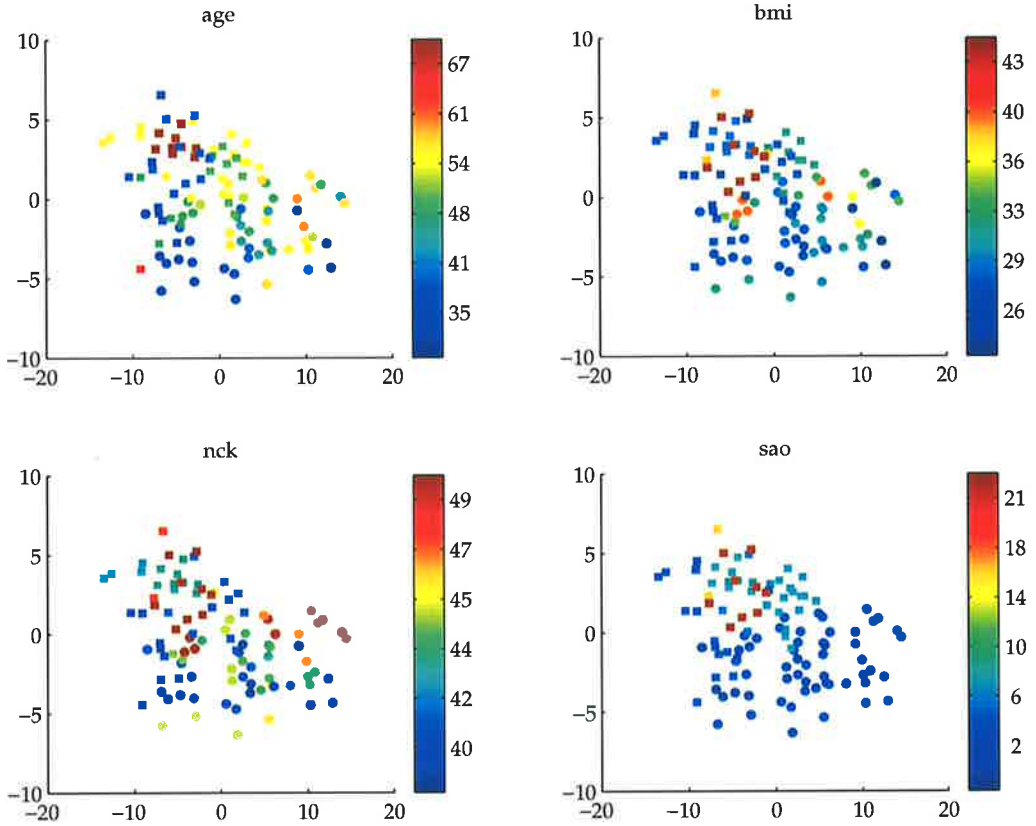

Figure 5.17: The distribution of the Sammon mapping (of non-normalised data) with respect to age, body mass index (bmi), neck circumference (nck) and oximetry score (sao).

Classifier	7-nearest-neighbour		Bayes' likelihood ratio	
	Specificity	Sensitivity	Specificity	Sensitivity
Non-normalised data	0.82(0.04)	0.89(0.03)	0.86(0.03)	0.87(0.05)
Normalised data	0.91(0.001)	0.80(0.007)	0.81(0.006)	0.82(0.003)

Table 5.4: Mean(std) cluster separation for the Sammon mappings of the normalised and non-normalised data, assessed by training set classification and averaged over 50 runs. As was the case with the GTM, the overlap between the control and SAHS clusters increases with normalisation.

5.4.3.2 Interactive visualisation

Interactivity has been identified in Section 3.1 as a very desirable attribute for data visualisation algorithms. A common device for interaction is the manual selection of points of interest by the user. The mapping is then adjusted to give the best representation of these particular points.

Figure 5.18 shows such an interaction with the Sammon mapping. The points within the box in Figure 5.18a were selected, and the error function (Equation 5.1) re-minimised with respect only to the selected points. The resulting mapping (Figure 5.18b) shows the local structure in this region of interest more accurately.

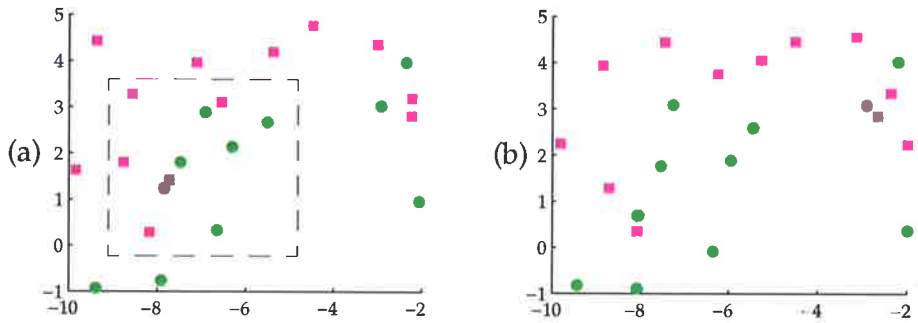

Figure 5.18: Interactive visualisation using the Sammon mapping: (a) zoomed view of Figure 5.15. The error function was re-minimised, considering only those points enclosed in the box; (b) the resulting mapping.

5.4.3.3 Generalisation

The general forms of the mappings of the data have been established in Section 5.4.3.1. The generalisation properties of the various maps are now explored.

The GTM was run again, with one sample withheld from the training data. The GTM was trained as usual and then the retained point was projected onto the mapping. This "test" position in the map was recorded and the procedure repeated, withholding each data point in turn. The GTM starts from an approximate principal components projection of the data each time and this starting configuration was very similar across iterations. (It was not identical, since a different data point was withheld from a relatively small data set each time. If the number of available data patterns was very large, then the removal of one point would have had negligible effect on the statistics of the set). The GTM is deterministic and so the test

position of each point should be close to the training position (the position in the mapping for that point, when the training data comprised the entire set).

The mapping of the collated "test" points is shown in Figure 5.19, with lines used to indicate the displacement of each point from its training position. The majority of the points are displaced little from their training positions and those points which are displaced significantly are generally on the edges of the grid. The overall picture of the data set conveyed by the mapping is similar to that of the map shown in Figure 5.9, and so it is fair to surmise that the GTM generalises well in this case. Similarly good performance was obtained with the other combinations of mapping algorithm and normalisation regime.

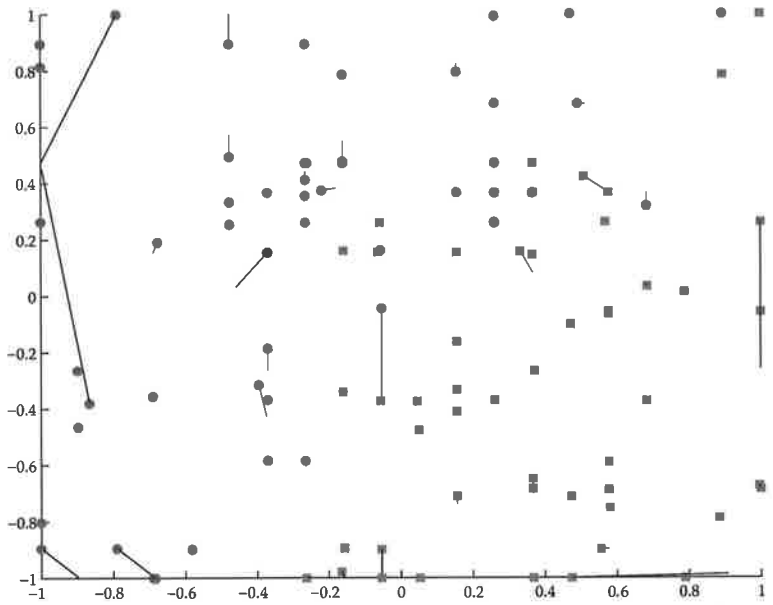

Figure 5.19: "Generalised" GTM projection of the non-normalised data. Each data point was withheld from the training set in turn, and projected as a test point. The lines indicate the displacement of this "test" position from the original position of the point.

To gain a more objective assessment of the generalisation properties of the mappings, classification of the test samples was also used. A test set comprising 10 randomly chosen samples each of control and SAHS data were excluded from the training data. Two simple classifiers were used: the voting *k*-nearest neighbours and Bayes' likelihood ratio. Here, we are not looking for high classification rates, but rather test set classification which is comparable to training set classification. Comparable test and training set classification indicates that the test points are being projected in approximately the same manner as the training points - and thus is a marker of good generalisation.

The classification scores on the test set are shown in Table 5.5 for the GTM and Table 5.6 for the Sammon mapping. Only one test set classification result was significantly different from the training set value (nearest-neighbour specificity in the Sammon mapping with normalised data - the classification was higher for the test set than the training set). Generalisation is thus excellent for both the GTM and the Sammon mapping, with both normalisation schemes.

	7-nearest-neighbour		Bayes' likelihood ratio	
	Specificity	Sensitivity	Specificity	Sensitivity
Non-normalised data	0.80(0.14)	0.89(0.09)	0.89(0.11)	0.87(0.10)
Normalised data	0.89(0.09)	0.81(0.12)	0.83(0.11)	0.82(0.12)

Table 5.5: Mean(std) GTM test set classification scores. The mean classification in each case was not significantly different from the training set classification (assessed using t-test with $p < 0.02$), indicating good generalisation.

	7-nearest-neighbour		Bayes' likelihood ratio	
	Specificity	Sensitivity	Specificity	Sensitivity
Non-normalised data	0.83(0.12)	0.89(0.12)	0.87(0.11)	0.85(0.13)
Normalised data	0.95(0.07)*	0.80(0.13)	0.84(0.10)	0.82(0.13)

Table 5.6: Mean (std) Sammon mapping test set classification scores. Only one result (marked *) was significantly different from the training set classification (assessed using a t-test and $p < 0.02$), indicating good generalisation.

As shown in Section 5.2.2, translation of least-squares solutions may be used to reduce the size of the network weights. Table 5.7 lists the weight matrix norm and the network curvature with and without translation. Under normalisation, both the network weights and the curvature showed a decrease with translation (Table 5.7). However, classification did not improve significantly with translation (results not shown). Translation does not alter the fit of the network to the training data and so does not reduce cluster overlap any more so than does the basic relative supervision. The curvature and weight-norm measures reflect the smoothness of the network as a whole, and not necessarily the smoothness of the transformation in the region of the cluster overlap.

With no normalisation of the data, there was no significant change in the network weights, curvature or classification performance with translation.

Training method	Weight matrix norm		Curvature	
	Normalised	Non-normalised	Normalised	Non-normalised
Relative supervision	$9.5\text{x}10^{22}$ ($5.6\text{x}10^{23}$)	$9.4\text{x}10^{23}$ ($2.8\text{x}10^{24}$)	$1.4\text{x}10^{14}$ ($8.2\text{x}10^{14}$)	$3.5\text{x}10^{15}$ ($1.0\text{x}10^{16}$)
Relative supervision with explicit translation	$5.8\text{x}10^{19}$ ($9.8\text{x}10^{19}$)	$9.5\text{x}10^{23}$ ($3.6\text{x}10^{24}$)	$4.7\text{x}10^{10}$ ($7.4\text{x}10^{10}$)	$3.3\text{x}10^{15}$ ($1.2\text{x}10^{16}$)
Supervised training	$1.0\text{x}10^{5}$ ($4.3\text{x}10^{5}$)	$5.3\text{x}10^{5}$ ($1.5\text{x}10^{6}$)	$1.4\text{x}10^{6}$ ($1.4\text{x}10^{6}$)	$1.2\text{x}10^{6}$ ($9.8\text{x}10^{5}$)

Table 5.7: Mean(std) network weights and curvature with different training methods. Results were averaged over 50 runs.

The training of the network may also be done in a purely supervised manner. For simplicity, relative supervision was used to construct a map and this configuration was then used as the desired output configuration for the supervised training of a second network. Forward selection and simple weight decay were used to provide regularisation. A range of basis functions with moderate to narrow widths were used in the pool of available units during forward selection. Training was terminated when a minimum in the generalised cross-validation error was reached. Re-estimation of the weight-decay parameter was also used (Orr, 1997; Orr, 1995).

Both the size of the network weights and the network curvature dropped dramatically with the supervised training (Table 5.7), a consequence of the use of moderate-width bases (rather than the very wide bases used with relative supervision). The classification results, however, were no better than those obtained from the relative-supervision-trained networks (Table 5.8).

5.4.3.4 Prior information

The effect of prior information was investigated for the Sammon mapping. Equation 5.11 was used to modify the data dissimilarities, using class labels as the prior information.

	7-nearest-neighbour		Bayes' likelihood ratio	
	Specificity	Sensitivity	Specificity	Sensitivity
Non-normalised data	0.85(0.13)	0.90(0.10)	0.88(0.11)	0.88(0.08)
Normalised data	0.92(0.09)	0.82(0.11)	0.81(0.13)	0.85(0.10)

Table 5.8: Mean(std) classification scores with supervised network training. The lower curvature of the network (see Table 5.7) compared to training with relative supervision did not yield better classification performance.

Increasing values of the parameter α reflect an increasing level of influence of the class labels on the mapping.

Under data normalisation, both the cluster separation and the curvature of the network rise with increasing supervisorial coefficient α (Figure 5.20). The unsupervised (α=0) configuration is not particularly well separated by class, and the class-label information serves to force the clusters to become more clearly separated. The network must become less smooth in order to achieve the necessary transformation, and so the curvature rises. The classification does not improve with the influence of class information, and in fact decreases with strong class influence (Figure 5.21).

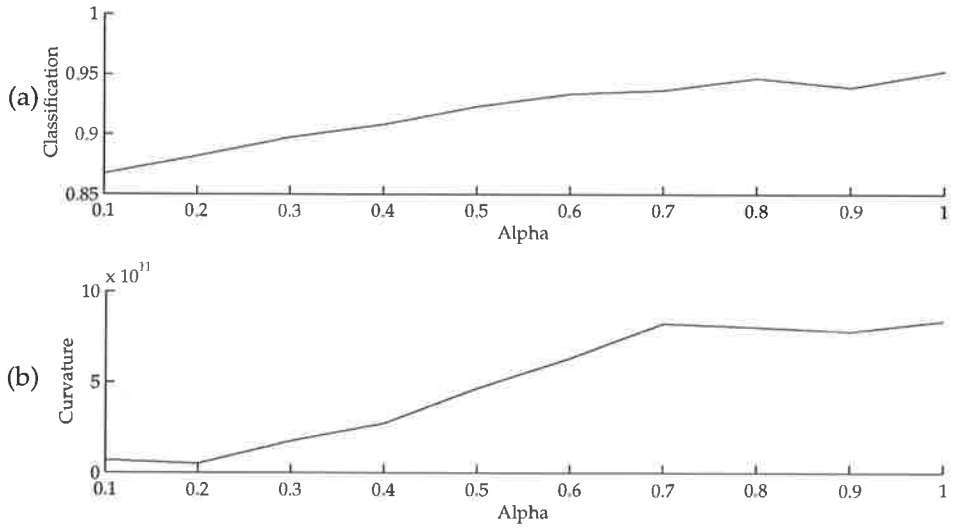

Figure 5.20: (a) Cluster separation (assessed by training set nearest-neighbour classification) and (b) network curvature for increasing class-label influence α and normalised data.

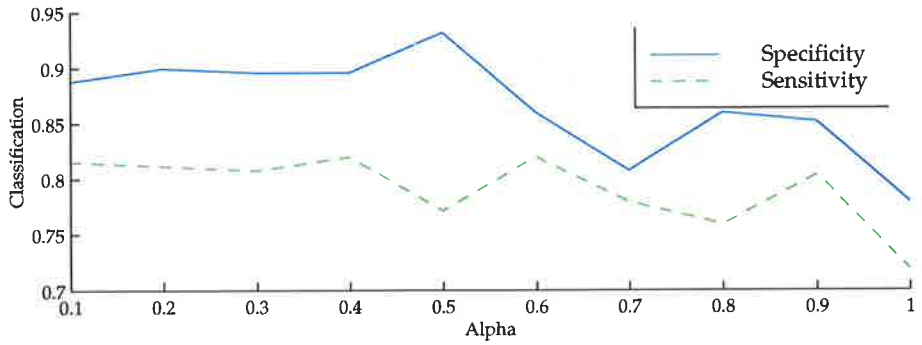

Figure 5.21: Nearest neighbour (test set) classification for increasing class-label influence α and normalised data.

The behaviour of the unnormalised mapping with class-label information is more interesting. With no normalisation, the cluster separation again improves with α as expected, but the curvature shows an initial, marked decrease for small α (Figure 5.22). The classification (both specificity and sensitivity) also rises, peaking at α=0.3 (Figure 5.23). For further increases in α, the curvature remains relatively constant, but the classification (particularly the sensitivity) degrades before settling to a relatively constant value.

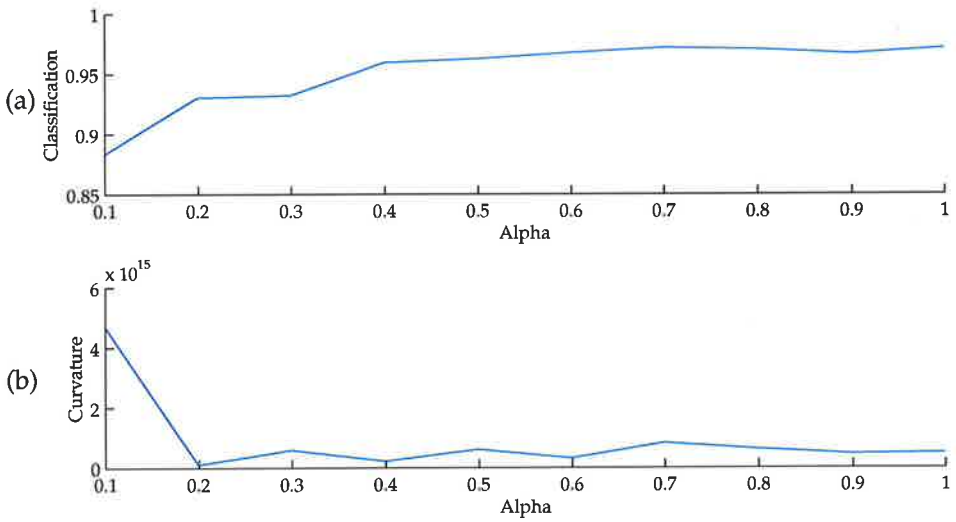

Figure 5.22: (a) Cluster separation (assessed by training set nearest-neighbour classification) and (b) network curvature for increasing class-label influence α and non-normalised data.

Consider for a moment the effect of non-zero α on the data dissimilarities. The incorporation of class-label information serves to increase small dissimilarities between data of different classes. Large dissimilarities between different-class data and all dissimilarities

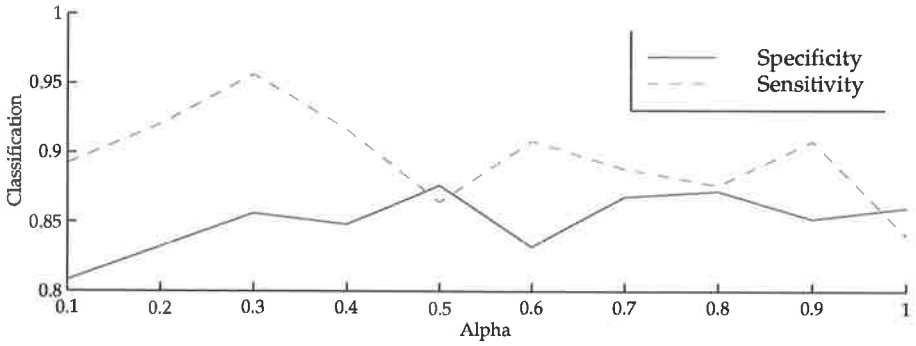

Figure 5.23: Nearest neighbour (test set) classification for increasing class-label influence α and non-normalised data.

between same-class data are reduced with non-zero α. Hence, the cluster separation increases with α and at the same time the contributions of the larger dissimilarities to the error function are reduced. This decrease in the importance of large dissimilarities (of which there will be a relatively small number) allows the more numerous small and moderate dissimilarities to be better represented - with a corresponding decrease in curvature. The increase in cluster separation gives the increased classification scores.

For large values of α, the cluster separation can become artificially large, leading to an increase in the curvature of the network (in this region of the map) and poor generalisation. In this case, however, the map construction is constrained by the network regularisation, thus reducing the severity of this effect. Nevertheless, there is a decrease in the classification with $\alpha > 0.3$ for this reason.

The Sammon mapping configuration with $\alpha = 0.3$ is shown in Figure 5.24, with the usual plots of additional information in Figures 5.25 and 5.26.

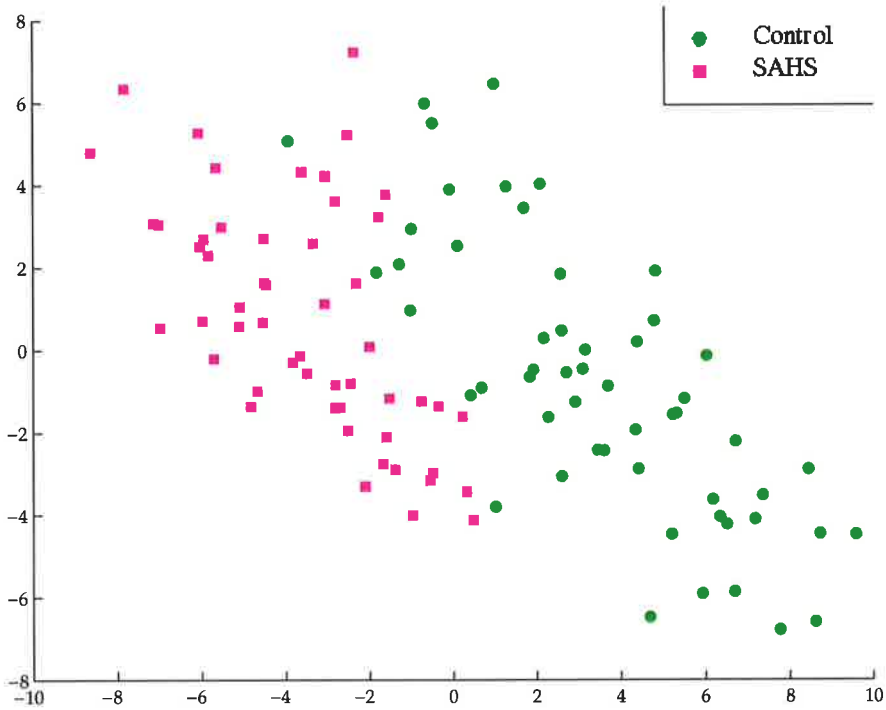

Figure 5.24: Sammon mapping of non-normalised data, using prior class-label knowledge and $\alpha = 0.3$. The use of prior information reduces the cluster overlap, as compared with Figure 5.15.

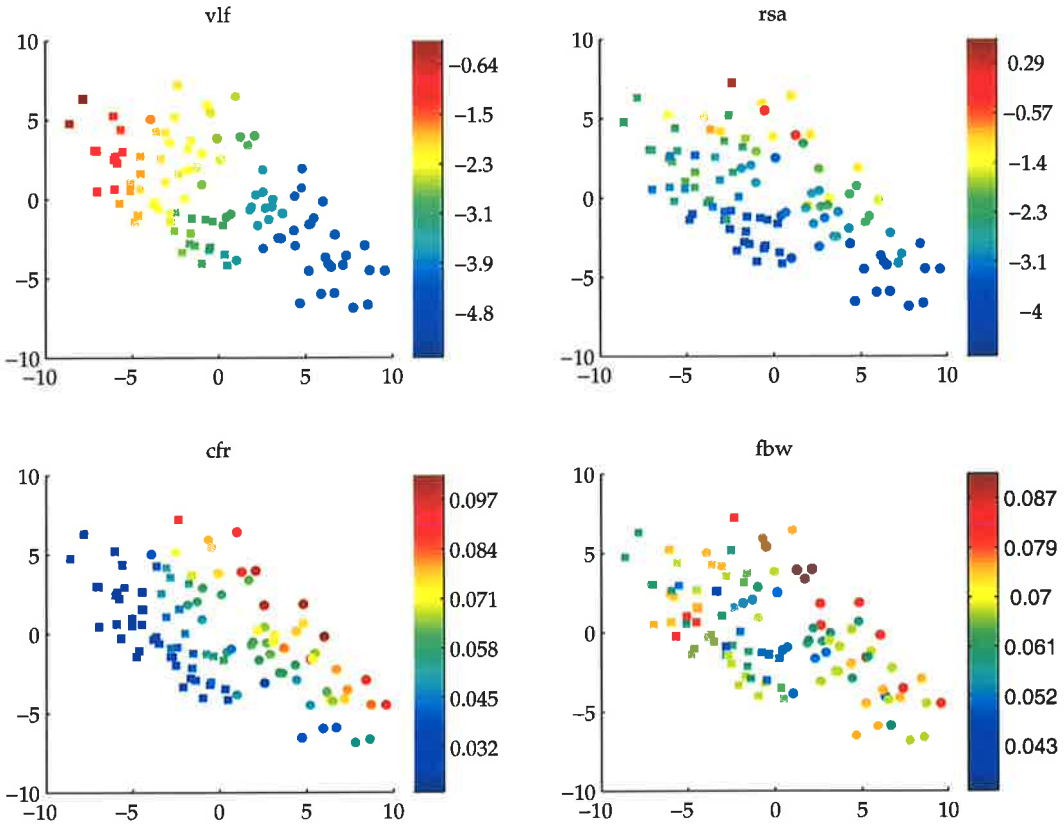

Figure 5.25: The distribution of the Sammon mapping with respect to log VLF power (vlf), log RSA power (rsa), centroid of PSD in band 0-0.15 Hz (cfr) and bandwidth about centroid (fbw). The mapping used no normalisation and class-label information with $\alpha = 0.3$.

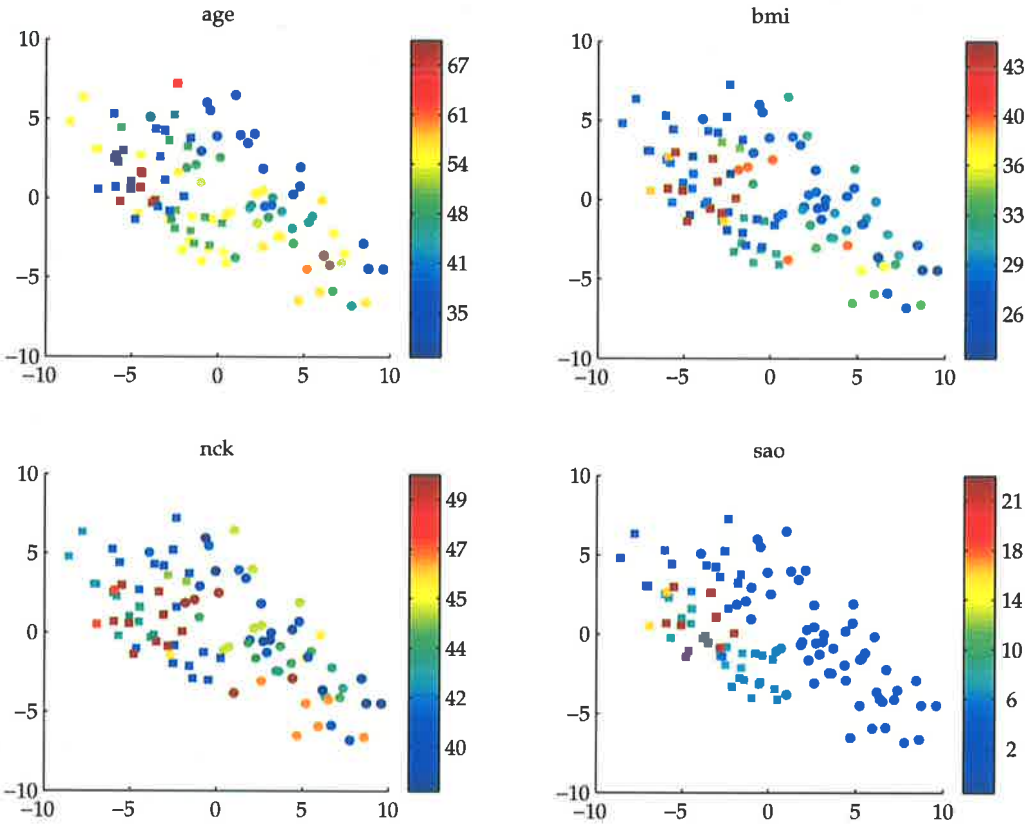

Figure 5.26: The distribution of the Sammon mapping with respect to age, body mass index (bmi), neck circumference (nck) and oximetry score (sao). The mapping used no normalisation and class-label information with $\alpha = 0.3$.

5.4.4 Conclusions

The least-squares scaling and the generative topographic mapping have been applied to data from a study of sleep apnoea. These techniques have allowed the interactive exploration of the data set and the identification of physiologically significant characteristics of the data.

The GTM and LSS were very similar in structure. Regardless of whether or not the data were normalised, the ordering of the points in the maps was strongly related to the VLF power in the HRV spectra. The brady-tachycardia cycle associated with sleep apnoea is known to cause this VLF activity. The results of Section 5.4 show that similar VLF activity may also be seen in young control subjects. The peak in the power spectrum associated with this VLF component tends to be located at higher frequencies in young control subjects when compared to subjects during apnoea.

If the data are normalised, the RSA component of the HRV spectrum has an increased influence on the form of the mapping. In this case, SAHS subjects can be mistaken as normal, particularly with fit SAHS subjects or during hypopnoea. Although not investigated in this study, the normalised form of the data may be useful in discriminating central and obstructive apnoea.

The LSS has the advantage over the GTM of flexibility in the cost function and a convenient means of incorporating prior (class label) information. The generalisation performance of a network trained to perform a least-squares scaling of non-normalised data was shown to improve when a moderate level of class-label influence was incorporated during map construction.

5.5 Summary

This chapter has investigated the use of neural networks to add generalisation capability to the least-squares scaling. Training may be carried out in conjunction with map construction, or as separate steps. The generalisation performance of networks is much the same for both training approaches, provided appropriate regularisation is incorporated during network training. For relative supervision, translation of solutions is an inherent regularisation mechanism, although it is not necessarily exploited fully by the algorithm. It is not difficult to employ this form of regularisation during supervised training, but it is applicable only to networks which utilise the least-squares scaling cost function.

Relative supervision makes the task of constructing the map considerably easier than is the case with direct optimisation methods, since the network smoothness helps avoid local minima in the map error surface. The parameterised nature of this approach also reduces the number of variables which need to be adjusted, when compared to methods which operate on the data directly.

The least-squares scaling also offers a convenient mechanism for including prior information during map construction. This can help to untangle data on a nonlinear manifold and improve visualisation performance.

The application of these techniques to data from a sleep apnoea study also demonstrated the value of incorporating information from a variety of sources. The subject age and body

mass index, in particular, have direct relevance to the diagnosis and understanding of sleep apnoea. In principle, the use of visualisation techniques for information fusion is simple: all relevant variables can be collated into one composite input vector and applied to the visualisation algorithm as before. Unfortunately, this is difficult to do efficiently, since the contribution from each variable to the dissimilarity function must be weighted appropriately. The cost functions of the least-squares scaling and generative topographic mapping are concerned simply with fitting the map or model to the data, and do not offer any means of finding an appropriate weighting.

The next chapter considers the use of classification techniques as a more appropriate tool for information fusion. The construction of classifiers is governed by the goal of maximising classification performance on the training data, usually subject to some regularisation constraints in order to ensure a smooth model. The weighting of the input variables is an efficient, inherent part of the classifier training.

Chapter 6 Information fusion and classification

In this chapter, the fusion of heart rate variability with information from other signals and sources is investigated. We assume that for each subject, we have information from several different sources and we wish to find a useful way to combine that information. This is a typical scenario in condition monitoring, where data from a variety of sensors is combined and an overall decision on the state of the subject is made (Navabi *et al.*, 1991). This could be done by concatenating all information into a single feature vector for each subject, with an appropriate weighting on each information source, and then applying a visualisation algorithm (van Gils *et al.*, 1997). Unfortunately, visualisation algorithms are concerned simply with constructing a two-dimensional representation of the data and do not offer a convenient means of finding appropriate weightings.

Many classification techniques operate by first defining a model of the data distributions and then fitting the model parameters so as to maximise classification accuracy on a set of training data. The contribution of each information source to the final decision function is determined by the classifier. Classification techniques thus offer a convenient means of performing information fusion. The notion of a "useful combination" of information in this case may be considered to be a combination which gives reliable classification performance.

A few studies have employed classification as a means of combining heart rate variability with other information (Vainamo *et al.*, 1996; Voss *et al.*, 1996; Yarnold *et al.*, 1994; Harper *et al.*, 1987). Several have used the linear discriminant analysis classifier (Yarnold *et al.*, 1994; Harper *et al.*, 1987), as have other classification studies using heart rate or heart rate variability data (Felgueiras *et al.*, 1998; Curcie and Craelius, 1997). Linear discriminant analysis makes the assumption that the data from each class may be described using a normal distribution. In order to ensure acceptable generalisation and therefore test set classification performance, all classification algorithms incorporate some measure of regularisation. The simple model structure of the linear discriminant classifier imposes structural stabilisation and is thus the source of regularisation in this classifier. However, the model structure in this

case may be a simplistic approximation to the true distributions and can lead to an over-smoothed classifier and poor performance.

There are many classifiers which use more general models for the data distributions, such as mixture model classifiers. In this family of classifiers, mixtures of simple distributions are used to build models of more complex distributions. This offers more appropriate modelling of non-normal distributions while retaining computational tractability.

In this chapter, the use of mixture model and linear discriminant analysis classification techniques are investigated for the task of combining heart rate variability with other information.

6.1 The shared mixture classifier

We assume that we have a training set of N feature vectors $\{\boldsymbol{x}_n\}, n = 1...N$, and a corresponding list of class labels $\{c_n\}, c \in \{1, 2, ..., C\}$. For classification, we wish to be able to predict the class of an unknown feature vector, based on the information provided by the training set data. In order to approach the problem statistically, we construct a model for the class-conditional probability density function of the data $p(\boldsymbol{x}|c,\Theta)$. A Gaussian approximation to this density function is given by:

$$p(\boldsymbol{x}|c,\Theta) = (2\pi)^{-N/2}|\boldsymbol{B}_c|^{-1/2}\exp\{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{m}_c)^T\boldsymbol{B}_c(\boldsymbol{x}-\boldsymbol{m}_c)\}, \quad (6.1)$$

where Θ denotes the model parameters: in this case $\boldsymbol{B}_c$ (the inverse of the covariance matrix of the distribution), and $\boldsymbol{m}_c$ (the mean of the distribution).

In practice, this may give a poor approximation to the true density, as the data may be skewed or multimodal. In this case, a mixture of K such densities may be used to more accurately model the density:

$$p(\boldsymbol{x}|c,\Theta) = \sum_{k=1}^{K} P(k|c,\Theta)\, p(\boldsymbol{x}|k,\Theta). \quad (6.2)$$

Here, $P(k|c,\Theta)$ is the prior probability that component k was responsible for data from class c. This approximation by mixtures of simple densities is known as the mixture model

approach (McLachlan and Basford, 1988). The representation given above is correctly known as the shared mixture model, since the components $k = 1 \ldots K$ are shared amongst all classes. Non-shared mixture models are also possible, but the shared representation is computationally more efficient (Bishop, 1995, p. 180) and thought to perform better on sparse data (Jarrad, 1998).

Classification may be done by computing the class posterior distribution:

$$p(c|\boldsymbol{x},\Theta) = \frac{\sum_{k=1}^{K} P(c,k|\Theta)\, p(\boldsymbol{x}|k,\Theta)}{\sum_{\tilde{c}=1}^{C}\sum_{k=1}^{K} P(\tilde{c},k|\Theta)\, p(\boldsymbol{x}|k,\Theta)} \tag{6.3}$$

and then classifying according to the estimated class $\hat{c}$:

$$\hat{c} = \arg\max_{c} p(c|\boldsymbol{x},\Theta) . \tag{6.4}$$

This is known as the shared mixture classifier (SMC). The well-known linear discriminant analysis (LDA) classifier is essentially a non-shared mixture classifier, with one Gaussian component per class.

In practice, the model parameters Θ are unknown and must be estimated from the data. In the maximum likelihood approach, we wish to find the most likely set of parameters, given the sets of features $\boldsymbol{X} = \{\boldsymbol{x}\}$ and class labels $\boldsymbol{C} = \{c\}$. This may be done by maximising the parameter posterior density $p(\Theta|\boldsymbol{X},\boldsymbol{C})$:

$$p(\Theta|\boldsymbol{X},\boldsymbol{C}) = \frac{p(\boldsymbol{X},\boldsymbol{C}|\Theta)\, p(\Theta)}{\int p(\boldsymbol{X},\boldsymbol{C}|\Theta)\, p(\Theta)\, d\Theta} . \tag{6.5}$$

Commonly, $p(\Theta)$ is unknown and replaced by a constant, so we may equivalently maximise the joint log-likelihood function:

$$L = \log p(\boldsymbol{X},\boldsymbol{C}|\Theta) = \sum_{n=1}^{N} \log \sum_{k=1}^{K} P(c_n,k|\Theta)\, p(\boldsymbol{x}_n|c_n,k,\Theta) . \tag{6.6}$$

The optimal value of Θ is typically estimated using the expectation-maximisation (EM) (Dempster *et al.*, 1977) or gradient-ascent (Bow, 1984) algorithms.

The mixture-model representation is almost identical to the mixture-of-Gaussians used in the generative topographic mapping. There are two principal differences: firstly, the GTM does not consider the class information since it strives to produce a representation of the data rather than a classification. Secondly, the means of the Gaussian components are adjusted indirectly in the GTM, through changes to the weight matrix.

The number of Gaussian components in the shared mixture classifier governs the smoothness of the classification boundary and is the dominant form of regularisation. One may also explicitly encourage diagonal covariance matrices and thus more spherical Gaussians, giving a smoother estimate of the density surface.

6.2 Heart rate variability and mean heart rate

In this section, the fusion of heart rate variability information with mean heart rate is explored. The combination of HRV and mean heart rate information has been suggested as a method of improving the practical utility of heart rate analysis, compared to the use of HRV information alone (Malik and Camm, 1993; Chess *et al.*, 1975).

The individual reproducibility of mean heart rate is good, but inter-subject variability can be large (Huikuri *et al.*, 1990; Appel *et al.*, 1989). The group reproducibility of HRV tends to be better than for mean heart rate, particularly in cardiac disease (Kautzner *et al.*, 1995; Bigger Jr *et al.*, 1992a; van Hoogenhuyze *et al.*, 1991).

Mean heart rate is an indicator of average autonomic activity and therefore tone, but is not a reliable indicator on its own (Malik and Camm, 1993). For low to moderate levels of autonomic activity, it is commonly assumed that HRV is proportional to tone (Akselrod, 1995). Thus, both mean heart rate and HRV are indicators of autonomic tone. The information provided by HRV analysis can be skewed by a variety of mechanisms, including measurement error in the beat times, spectral estimation errors due to transient signals or non-white noise and inappropriate normalisation regimes. Additionally, as autonomic tone saturates, the modulation depth and thus HRV values drop, even though tone is increasing. These sources of error are different to those which affect the mean heart rate signal.

The correlation between mean heart and HRV has been studied in a variety of contexts (Coumel *et al.*, 1995; Murakawa *et al.*, 1993). Examining the two signals together allows uncorrelated noise in the two to be averaged out and may therefore provide improved information regarding autonomic tone. The two signals have been studied in conjunction in the monitoring of foetal state (Stevens *et al.*, 1988; Chess *et al.*, 1975).

6.2.1 Changes with posture and propranolol

The data set collected for the study of the effects of posture and propranolol on HRV, first introduced in Section 4.2, was again used. Samples from all six subjects were used, so as to include inter-subject variations into the data.

The generative topographic mapping was used to extract two latent variables from the spectral data, as was done in Section 4.1.3. The classification performance using linear discriminant analysis and shared mixture classifiers was evaluated, using the latent variables alone and also in combination with mean heart rate information. All variables were demeaned and scaled to unity variance prior to classification.

6.2.1.1 Results with posture

The change from supine to upright posture (under placebo) is accompanied by an increase in sympathetic tone and withdrawal of vagal tone. This was shown with a mean beat interval across all subjects of 0.93(0.10)s while supine and 0.74(0.07)s during tilt.

Linear discriminant analysis was used to classify samples as either supine or upright, using 10 randomly selected samples of each for testing and the remainder for classifier construction. The fraction of samples correctly classified, using only mean heart interval information and averaged across 20 runs, was approximately 0.75 (Table 6.1). Using the latent variable information extracted from the HRV spectra by the GTM, the LDA test set classification averaged 0.69. Combining the two sources of information yielded a correct classification fraction of 0.83.

The SMC, using 10 shared clusters and trained using the EM algorithm, yielded considerably more accurate classification than LDA, with mean correct classification of 0.86 using mean heart interval alone, 0.83 using HRV information alone, and 0.93 for the two together (Table 6.2).

	Mean heart interval information only		HRV information only		Both mean heart interval and HRV information	
	Train	Test	Train	Test	Train	Test
Supine	0.76(0.03)	0.80(0.10)	0.58(0.05)	0.55(0.16)	0.76(0.03)	0.75(0.12)
Tilt	0.75(0.01)	0.69(0.13)	0.82(0.03)	0.82(0.14)	0.91(0.02)	0.91(0.10)

Table 6.1: Mean(std) fraction of samples correctly classified, using the linear discriminant classifier on the supine/tilt data, with placebo only.

	Mean heart interval information only		HRV information only		Both mean heart interval and HRV information	
	Train	Test	Train	Test	Train	Test
Supine	0.98(0.04)	0.96(0.08)	0.85(0.03)	0.81(0.10)	0.97(0.01)	0.93(0.08)
Tilt	0.72(0.02)	0.75(0.12)	0.80(0.04)	0.84(0.10)	0.95(0.03)	0.92(0.06)

Table 6.2: Mean(std) fraction of samples correctly classified, using the shared mixture classifier on the supine/tilt data, with placebo only.

The reason for the better performance of the SMC over LDA can be illustrated by considering the estimated probability density functions of the two classifiers. Figure 6.1a

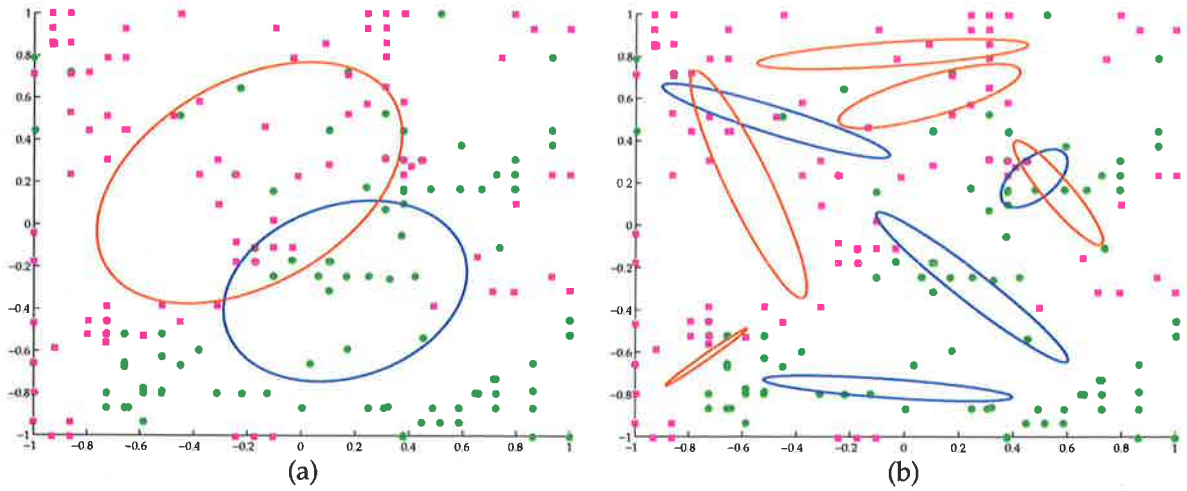

Figure 6.1: Classification of the GTM output. (a) Contours of equal probability for the linear discriminant analysis classifier; (b) contours for a non-shared mixture classifier. The blue contours in each case correspond to supine data (green circles), the red contours to tilt data (magenta squares).

shows contours of equal probability for the LDA classifier, superimposed on the GTM latent variable data. Clearly, the single-Gaussian model for each class distribution is a coarse approximation. The contours of equal probability for a mixture of 9 Gaussians is shown in Figure 6.1b, using a non-shared representation for clarity. The flexibility of the mixture model approach offers better modelling of non-normal distributions.

For both the LDA and SMC classifiers, classification using both HRV and mean heart interval information exceeded that achieved using either in isolation.

6.2.1.2 Results with propranolol

Propranolol blocks the effects of sympathetic activity on the SA node, effectively causing the sympathetic tone to be reduced or abolished. Mean(std) heart interval was 0.93(0.07) with placebo and 1.10(0.15) with propranolol, in the supine position. In the supine position, the HRV spectrum principally offers information regarding the parasympathetic modulation of heart rate. Thus, HRV is not particularly sensitive to propranolol in the supine position, as seen in Section 4.2 on data from a single subject. The discrimination between heart rate after administration of placebo or propranolol, using the LDA and SMC classifiers was investigated.

The LDA classification score was much the same using mean heart interval, HRV and both together, each giving about 0.7 correct classification (Table 6.3). For the SMC, mean heart interval information yielded a correct classification of 0.84, which dropped to 0.62 when HRV information was used and rose again to 0.86 with both information sources included (Table 6.4).

	Mean heart interval information only		HRV information only		Both mean heart interval and HRV information	
	Train	Test	Train	Test	Train	Test
Placebo	0.68(0.01)	0.70(0.14)	0.56(0.03)	0.62(0.11)	0.65(0.03)	0.71(0.13)
Propranolol	0.67(0.01)	0.67(0.13)	0.81(0.03)	0.73(0.10)	0.78(0.02)	0.70(0.18)

Table 6.3: Mean(std) fraction of samples correctly classified, using the linear discriminant classifier on the placebo/propranolol data, in the supine position.

	Mean heart interval information only		HRV information only		Both mean heart interval and HRV information	
	Train	Test	Train	Test	Train	Test
Placebo	1.0(0.00)	1.0(0.00)	0.69(0.05)	0.68(0.16)	0.89(0.02)	0.88(0.11)
Propranolol	0.62(0.02)	0.68(0.17)	0.73(0.04)	0.56(0.17)	0.92(0.02)	0.83(0.12)

Table 6.4: Mean(std) fraction of samples correctly classified, using the shared mixture classifier on the placebo/propranolol data, in the supine position.

The use of mean heart interval information in conjunction with HRV did improve the discrimination between placebo and propranolol when compared to HRV alone. However, the classification performance was no better than that seen using only mean heart interval information.

6.2.2 Conclusions

This section has demonstrated the concurrent use of mean heart interval and heart rate variability information in classification. In the tasks of discriminating between heart rate records from subjects in the supine and tilted positions, and before and after the administration of propranolol, the use of mean heart interval information and HRV outperformed the use of HRV information in isolation.

The data investigated here did not display the saturation behaviour which occurs with extreme sympathetic or vagal stimulation. In that scenario, HRV values drop as the nerve tone begins to saturate, but the mean heart rate continues to drop (in the case of vagal stimulation) or rise (with sympathetic stimulation). A specific model of this behaviour may need to be adopted for the concurrent processing of HRV and mean heart rate information in this case.

6.3 Other diagnostic information

Mean heart rate information is not the only candidate for improving the performance of HRV analysis. In most practical situations, additional clinical information is readily available -

from other sensors or imaging devices, from the patient history, or even simple personal details such as patient age.

A well-established application of this type which involves HRV is that of risk stratification following myocardial infarction. A variety of predictors of mortality following infarction are known, including depressed HRV (Wolf *et al.*, 1978). On its own, HRV has a positive predictive value (defined as the number of true positives, divided by the combined number of true and false positives) of around 0.4. When combined with other indicators, such as information regarding left ventricular ejection fraction and ventricular premature complexes, the predictive value rises to 0.5 (Bigger Jr *et al.*, 1992b). Similarly, in a study of infant sleep/wake state employing linear discriminant analysis, HRV measures alone offered a correction classification of 0.82, which rose to 0.85 when respiratory measures were also included (Harper *et al.*, 1987).

6.3.1 Application to sleep apnoea

It is known that age, obesity and gender have a direct correlation with risk of apnoea, with older, overweight males showing the most susceptibility. There may therefore be benefit in including such information during classification of HRV during apnoea.

The LDA and SMC classifiers were applied to the apnoea data set from Section 5.4. A radial basis function implementation of the least-squares scaling was used to extract features from the HRV spectral data.

An unsupervised scaling was used ($\alpha = 0$ in Equation 5.11). Classification was first performed using HRV information only, and then repeated, using HRV information in conjunction with subject age and body mass index (all subjects were male). Previously, it was noted that the use of visualisation techniques as feature extractors for classification may require more than two output variables in order to achieve reliable classification. Accordingly, classification results using up to 10 latent variables were evaluated. The classification results of the LDA and SMC classifiers are plotted in Figure 6.2.

The addition of BMI and age information (Figure 6.2b) did not produce an increase in the classification performance compared to HRV alone (Figure 6.2a). The reason for this stems from the demographic nature of the patient base used in the study: age and BMI were in fact

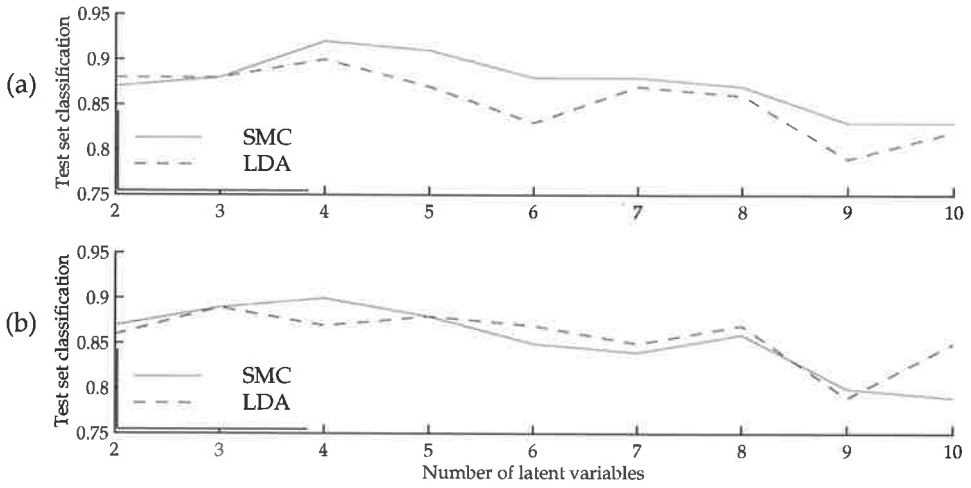

Figure 6.2: Test set performance of the SMC and LDA classifiers operating on (a) least-squares scaling output only, for 2-10 output dimensions; (b) least-squares scaling output, subject age and body mass index. Performance peaks at about 4 latent dimensions in each case, and the addition of age and BMI information does not improve the classification.

very similar for controls (49(9.5) and 30(4.2)) and apnoea subjects (46(9.6) and 32(7.0)). Therefore, this information did not add extra discriminatory power to the classifiers.

The performance of both classifiers increased with the number of latent variables, peaking at around 0.92 with 4 variables, before decaying again. It can also be seen that the simple LDA classifier offered comparable performance to the more computationally demanding SMC (using 4 shared clusters in this case), in contrast to the results of Section 6.2.1. Apparently, a simple normal distribution was sufficient to describe the class probability densities in this case. This is probably due in part to the use of the LSS rather than the GTM: the output of the GTM is confined to a discrete, square grid, giving distributions which are generally not well approximated by single-Gaussian models. No such restriction is present in the LSS, and here, the data were reasonably well modelled by LDA.

The use of class-label information with this data was previously found to yield better cluster separation in the map (Section 5.4). The classification performances of LDA and the SMC, based on a partially-supervised least-squares scaling ($\alpha = 0.3$), are shown in Figure 6.3.

The inclusion of a moderate level of class-label influence again improved the cluster separation. This lifted classification performance (using 2 latent variables) by approximately 5% to about 0.92, for both the SMC and LDA. For both classifiers, the performance with 2

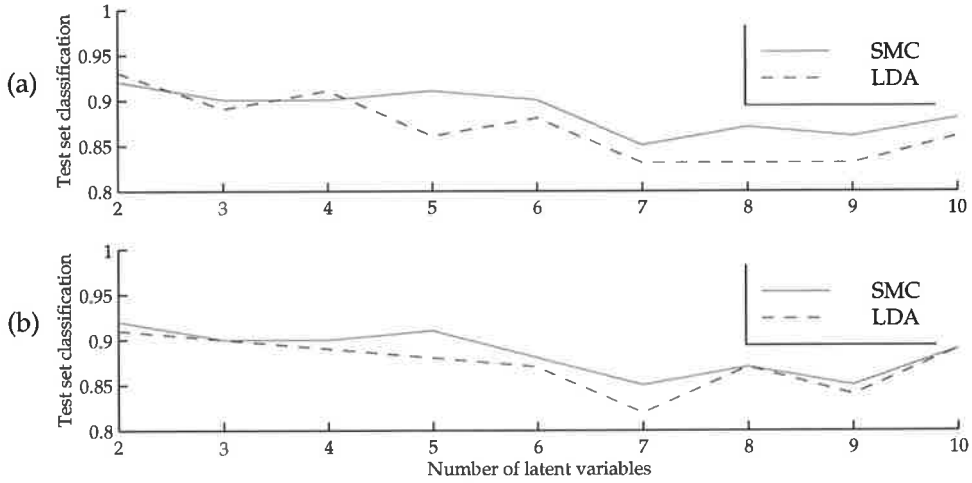

Figure 6.3: Test set performance of the SMC and LDA classifiers operating on (a) least-squares scaling output only, for 2-10 output dimensions, and with prior class-label knowledge incorporated into the mapping ($\alpha = 0.3$); (b) as for (a), but with the addition of subject age and body mass index information. Performance does not improve with increasing number of latent variables, nor with the addition of age and BMI information.

latent variables and class-label influence was about equivalent to the performance with 4 latent variables and no class-label influence (compare Figure 6.3a for two latent variables with Figure 6.2a for four latent variables). Thus, the inclusion of class-label information in this case serves to make the information extracted from the HRV spectra more concise. With the inclusion of class-label information the number of latent variables was reduced from 4 to 2 without sacrificing classification performance.

For α=0.3, classification performance did not improve with more numerous latent variables. Two features in this case were sufficient. Again, the inclusion of age and BMI information did not improve the classification results.

6.3.2 Heart rate variability and oximetry data

A widely accepted alternative to full polysomnography for apnoea screening is oximetry. Oximetry shows excellent specificity, but potentially low sensitivity, as short apnoeas or hypopnoeas which do not cause significant arterial oxygen desaturation may not be detected. The threshold of significant desaturation is commonly set at 4%. The sensitivity may be improved by lowering the threshold, to say 2%, but with a corresponding increase in the false alarm rate and thus a drop in specificity.

It is thought that the use of HRV in conjunction with oximetry may help in this regard (Hilton *et al.*, 1998; Barschdorff *et al.*, 1994). The high false alarm rate of lowered-threshold oximetry could be reduced while retaining high sensitivity, by validating each event against HRV (Figure 6.4).

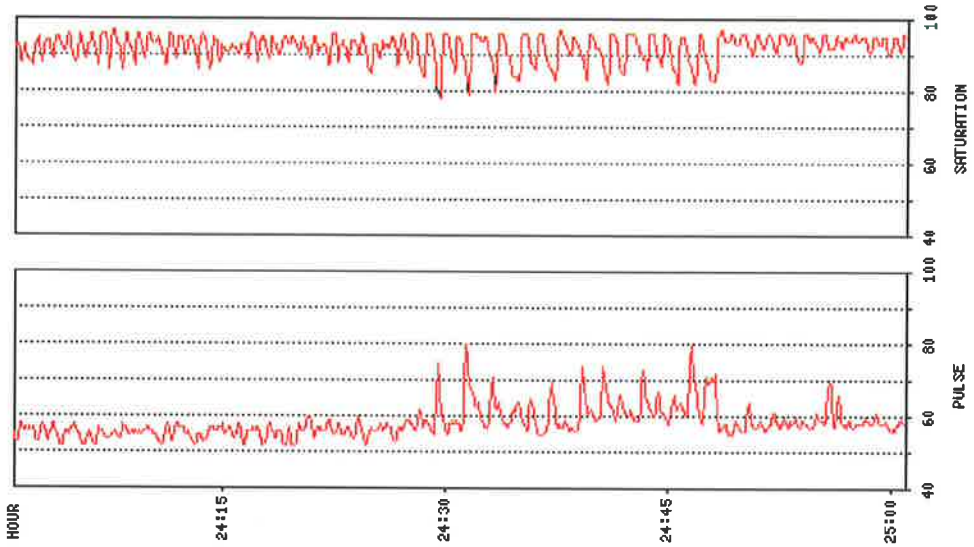

Figure 6.4: Typical traces of oximetry and heart rate.

Ideally, the oximetry data should be in the form of time-series data, which allows both the exact time of each desaturation event as well as the level of desaturation to be used. However, the only oximetry data available from this study was in summary form. This consisted of the number of recorded desaturation events during the 15-minute light sleep segment, with the desaturation threshold set at 4%.

6.3.2.1 Classification results

The shared mixture classifier with 4 shared clusters was used. Based on the oximetry data alone, the specificity and sensitivity were 0.98(0.04) and 0.85(0.09).

An unsupervised ($\alpha = 0$) least-squares scaling using 2 latent variables was applied to the HRV spectral data. When combined with the oximetry information, the specificity dropped to 0.87(0.10) and the sensitivity rose marginally to 0.87(0.13). Increasing the number of latent variables to four improved the classification, giving a specificity of 0.99(0.03) and sensitivity of 0.86(0.11). The inclusion of class-label information during map construction ($\alpha = 0.3$) did not provide any further improvement in performance.

Thus, with 4 latent variables, classification on the combination of HRV and oximetry data showed a high specificity. The low sensitivity figure (0.86) would be improved through the use of lower-threshold oximetry data. Additionally, the oximetry information used here was in the form of time-averaged summaries by subject. A simultaneous event in both signals (a brady-tachycardia cycle at the same time as a desaturation event) would be a strong indicator of apnoea. If the times of the events in the two signals are available, this would allow comparison of simultaneous events in the two signals.

6.3.3 Conclusions

For the classification of HRV spectra from apnoea subjects, the addition of age and body mass index information did not improve classification performance. While these are known to be significant risk factors for apnoea, these variables were not significantly different in the control and apnoea subjects in the study. It is likely that this would also be the case for practical screening of apnoea, since subjects who progress to an apnoea screening test are likely to be predominantly older, overweight males. The increased likelihood of apnoea in these subjects is taken into consideration by the practitioner when making the recommendation for the screening. Including this information again during the classification stage may not be beneficial.

In contrast, the combination of HRV and oximetry information promises reliable detection of apnoea events. The oximetry data available from this study was not available in an appropriate format for proper fusion with HRV information. Nevertheless, the initial results are encouraging.

6.4 Summary

In this chapter, classification techniques have been used to combine heart rate variability with other signal information. The combination of HRV and mean heart rate was shown to provide better discrimination of heart rate data by posture and drug treatment than was possible through the use of HRV information alone. The combination of HRV information with oximetry data for the detection of sleep apnoea was also investigated, but a complete treatment was not possible due to limitations on the availability of data.

This chapter has shown that it is possible to use visualisation algorithms as general-purpose feature extractors. These features can be subsequently passed to classification algorithms. In Chapter 3, it was argued that limiting the number of features to two is most convenient for visualisation purposes. However, if the features are to be passed to a classification algorithm, more than two features may be required for reliable classification. This was shown in Section 6.3.1.

It was also shown that the inclusion of class-label information into the mapping may improve classification performance (Section 6.3.1), since such information can improve the cluster separation in the mapping.

Chapter 7 Concluding remarks

There was a point to this story, but it has temporarily escaped the chronicler's mind.

- Douglas Adams

This thesis has addressed some pressing issues in the field of heart rate variability analysis. A summary of the main contributions of this work and suggestions for the direction of future research are given in this chapter.

7.1 Recapitulation

The analysis of heart rate variability data in the spectral domain is conventionally performed by considering the power contained within fixed sub-bands of the frequency spectrum. The physiological interpretation of these fixed-band powers can be difficult. While the high frequency band, synchronous with respiration, is widely accepted to be mediated solely by the parasympathetic branch of the autonomic nervous system, no such consensus has been reached for the lower frequency components. Further, the band powers reflect the depth of modulation of autonomic tone, rather than tone itself. The choice of the frequency limits of the bands can also be difficult.

7.1.1 Data visualisation

The deficiencies of the fixed-band approach to heart rate variability analysis have provided the motivation for this work, in which visualisation techniques have been used for the processing of heart rate variability spectral data. The algorithms of choice throughout this thesis have been the least-squares scaling, which is a form of multidimensional scaling, and the generative topographic mapping. These algorithms provide a graphical description of the relationships between data, allowing visual interpretation and exploration of the data. A brief review of these algorithms may be found in Chapter 3.

The least-squares scaling and the generative topographic mapping can be used to analyse spectral heart rate variability data without the need to define fixed bands of interest, and they

are also sensitive to changes in the location of the low-frequency peak in the spectrum (Section 4.1). The expert knowledge of the user can be exploited to determine the physiological significance of the elements within the input data vectors (Section 3.2).

7.1.2 Generalisation

The property of generalisation was identified as being of great value for visualisation. This property allows the efficient projection of new data, offering a means of comparing unknown data samples against a known database. The generative topographic mapping is naturally amenable to generalisation but the least-squares scaling has no mechanism for this. The use of neural networks is a popular means of adding this capability to the least-squares scaling. For radial basis function networks, which are linear in the weights, it has previously been shown that transformation of the solution may be used to minimise the norm of the weight matrix and is thus a form of regularisation. This regularisation mechanism is unique to the rotation- and translation-invariant cost function of the least-squares scaling. The translation required to achieve a minimum of the weight matrix norm can be calculated analytically, and it was shown in Section 5.2.2.1 that the minimum value possible is independent of any rotation of the solution. This mechanism may therefore be incorporated efficiently into a supervised training algorithm of radial basis function networks for least-squares scaling.

7.1.3 Prior information

Visualisation may be carried out solely on the basis of the data dissimilarities, or, if knowledge of the data (such as class labels) is available prior to map construction, then this information may be incorporated into the least-squares scaling. This can improve the user's perception of the data, by encouraging separation of class clusters in the mapping, or by forcing the map into a more informative orientation. A simple modification to the generative topographic mapping algorithm was given in Section 5.3.1, allowing similar use of prior information.

7.1.4 Classification

Visualisation techniques may also be employed as the front end of an information fusion and classification system. The fusion of heart rate variability information with information from other sources is one method of improving the reliability of heart rate variability analysis.

Classification algorithms offer a convenient framework in which this fusion may be accomplished.

For visualisation purposes, we wish to extract two latent variables in order to be able to easily construct a visual display. It was shown in Section 6.3.1 that if the visualisation algorithm is to be used as the feature extraction part of a classification system, then more than two latent variables may be required. Class-label influence during map construction can also improve the classification performance of the overall system.

The fusion of mean heart rate and heart rate variability information was shown (Section 6.2) to be of benefit in discriminating data by posture and propranolol, in both cases outperforming the classification on the basis of heart rate variability information alone. The combination of HRV information with other variables for the detection of sleep apnoea was also investigated in Section 6.3. The combination of HRV with age and body mass index information proved to do no better than the use of HRV alone. However, the combination of HRV and oximetry data was more promising, and warrants further investigation. In a similar vein, the combination of time- and frequency-domain features from the HRV signal was also shown to be beneficial in the classification of hypertensive subjects (Appendix E).

7.1.4.1 A note on classifiers

Linear discriminant analysis is a simple classification technique which is popular for the classification of biomedical data. In Section 6.2 it was shown that linear discriminant analysis can perform poorly on data which is clearly non-normal. This is hardly a surprising result, but is one which is worth noting given the proliferation of linear discriminant analysis in the literature. A classifier which is better able to model non-normal distributions (such as the shared mixture classifier, described in Section 6.1) may provide superior classification performance over linear discriminant analysis with non-normal data.

7.2 Contributions of this thesis

This thesis makes the following contributions to the fields of heart rate variability analysis and neural networks:

- a concise review of HRV processing techniques (Chapter 2);

- a concise review of multidimensional scaling, the generative topographic mapping and the self-organising map (Chapter 3);

- the use of visualisation techniques (in particular the least-squares scaling and the generative topographic mapping) for the analysis of heart rate variability data (Chapter 4). This is a new approach to the investigation of heart rate variability data, and the major contribution of this thesis. An appropriate preprocessing procedure for spectral data was given in Section 4.1. Demonstrations of visualisation were given on data from studies into the effects of posture and propranolol (Section 4.2) and sleep apnoea (Section 5.4);

- for the training of radial basis function networks, the generalisation performance of the relative supervision algorithm was compared with that of two-step (supervised) training methods. This comparison extends previous work (Tipping, 1996) by including regularisation mechanisms into the supervised training methods (Section 5.2);

- it was shown that the translation of solutions is a viable method of regularisation for the supervised training of radial basis function networks for least-squares scalings (Section 5.2). Previously, it was thought that the global minimum of the network weight magnitudes was a function of both the translation and the rotation of the map, and that minimising the network weights would involve a difficult nonlinear optimisation. Here it has been shown that it is not necessary to search rotations of the solution in order to find the global minimum of the network weight magnitudes (Appendix C). The use of translation as a mechanism of regularisation is in fact very simple to implement;

- a modification was made to the original generative topographic mapping algorithm to allow prior information to influence the mapping (Section 5.3.1). The use of prior information can improve the interpretability of the resulting mapping;

- linear discriminant analysis and the shared mixture classifier were used to combine heart rate variability with mean heart rate information (Section 6.2). The fusion of these two information sources was shown to provide better discrimination of heart rate data than the use of heart rate variability alone;

- heart rate variability was combined with oximetry data for the detection of sleep apnoea (Section 6.3.2);

- it was shown that normalisation of rate and interval spectra by the square of mean heart rate and mean heart interval reduces the discrepancy in results from the two signals (Appendix A); and

- a demonstration of the classification of heart rate variability was given in subjects with mild hypertension (Appendix E).

7.3 Implications for heart rate variability analysis

Visualisation techniques are an excellent tool for the investigation of HRV data. The relationships between data are made graphically apparent, and the characteristics of the data responsible for these patterns may be identified. While expert knowledge may be incorporated into the algorithm design (for example, in the choice of dissimilarity function), detailed modelling of domain knowledge is unnecessary, saving a great deal of complexity in the algorithm design and computational requirements.

One of the poorly-addressed problems from the clinical perspective of heart rate variability processing is the definition of "normal" heart rate variability (Task Force, 1996). The concept of normal heart rate variability is evolving as our understanding of the field grows. No rigid definition of normal heart rate variability has been proposed in the time- or frequency-domain, although continuing developments in nonlinear processing may offer a suitable characterisation of the signal in this regard (Felgueiras *et al.*, 1998; Mammoliti *et al.*, 1998; Pompe *et al.*, 1998; Signorini *et al.*, 1998).

Visualisation offers a different approach to the definition and characterisation of normal data, through the comparison of data samples. A data sample of unknown origin may be projected onto a map, allowing comparison of that sample with a range of other samples. These samples may be from a range of normal and diseased subjects. The similarity of the unknown sample to this range of examples may be easily assessed. Thus, in the visualisation-based approach, there is no need to explicitly characterise "normal" data, nor data which is typical of a particular disease.

7.4 Implications for future work

7.4.1 Visualisation of heart rate variability data

The visualisation algorithms of Chapter 3 are not restricted to the exploration of spectral data. With the definition of an appropriate dissimilarity function, these algorithms may be applied to data of any type.

The visualisation of data from nonlinear and chaotic HRV analysis techniques is one area of application which warrants investigation. Another potential avenue for application is the visualisation of the results of system modelling (Section 2.2.4.2). The fitting of a model to a particular set of physiological data is typically achieved by adjusting a large number of parameters. The parameter values corresponding to different sets of data may be compared, but the large number of parameters makes this a difficult task. Visualisation techniques could be used to investigate the relationships between a number of sets of parameters, and this also warrants investigation.

7.4.2 Other aspects of heart rate variability processing

The most attractive aspect of heart rate variability is the inexpensive, non-invasive nature of the data acquisition process. This advantage is offset to some extent by the limitations of existing processing methods. Improved methods of heart rate variability analysis, which will allow application to a wider range of haemodynamic conditions, are therefore of prime research interest.

The analysis of transient signals is an area of strong research interest. Linear models are insufficient to capture the dynamic behaviour of the cardiovascular control systems, and so nonlinear models are increasing in popularity. One such area of application, which has not been addressed in this thesis, is nonlinear filtering methods for the extraction of the respiratory component from heart rate variability. The extraction of the respiratory sinus arrhythmia component from the HRV signal by noise-cancellation filtering is well established. Adaptive linear filters are intended for use on stationary or slowly-varying data, and for application to transient signals, nonlinear filters may offer superior performance. An early study into the use of nonlinear finite impulse response filters for respiratory noise cancellation concluded that the large amount of data required for estimation of the filter

parameters rendered this approach impractical (Varanini *et al.*, 1994). The number of filter parameters and thus the required volume of training data may be reduced by utilising infinite impulse response filters, although parameter estimation is more complex. Efficient implementations of nonlinear finite- and infinite-impulse response filters are possible using radial basis function neural networks (Billings and Fung, 1995), and these may offer some improvements in the processing of transient cardiovascular signals.

7.4.2.1 As a probe of autonomic tone

The significance of changes in the centre frequency of the low-frequency peak are not well understood but, as discussed in Section 2.5, may offer valuable information regarding autonomic tone. Modelling and physiological validation of this possibility is a promising and worthwhile area of future research.

Another approach to improving heart rate variability analysis as a probe of autonomic tone involves the combination with mean heart rate information. This was addressed in Section 6.2 using classification; one may also approach the problem from a modelling perspective.

Assume for the moment that we are considering vagal control of the heart rate. For low to moderate levels of vagal activity, it seems reasonable to assume that the amplitude of the respiratory sinus arrhythmia in the heart rate variability signal is proportional to vagal tone (Akselrod, 1995), and indeed this relationship is the reason for the popularity of heart rate variability analysis as a probe of autonomic function. However, when vagal activity is high, the tone becomes saturated and the respiratory sinus arrhythmia amplitude decreases, introducing an ambiguity. The combination of mean heart rate and heart rate variability information in this situation may benefit from an explicit model of this saturation behaviour (Figure 7.1).

Figure 7.1 shows that as vagal tone increases, the high-frequency peak in the HRV spectrum increases in amplitude. This occurs until tone begins to saturate, at which point the HF peak amplitude drops. This introduces an ambiguity: two different levels of tone produce the same HRV component amplitude. It may be possible to resolve this ambiguity with the use of mean heart rate information, since mean heart rate (in the absence of other influences) is a decreasing function of vagal tone. The difficulty in such an approach lies in the fitting of the model, since the characteristics of the curve in Figure 7.1 will vary considerably from

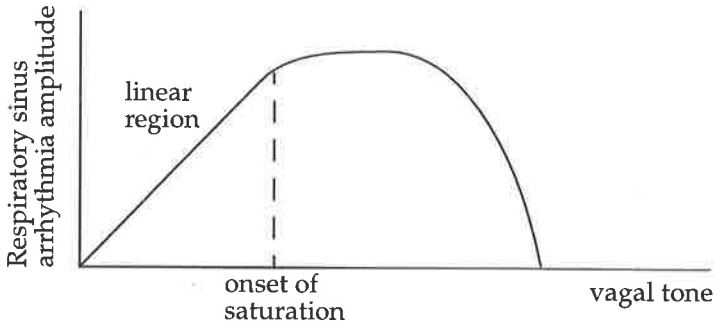

Figure 7.1: Schematic depiction of the relationship between tone and heart rate variability.

subject to subject. Calibration and fitting of the model may require careful design of the experimental protocol - for example, using progressive tilt from supine to upright, followed by exercise, in order to progressively stimulate the sympathetic control of the heart.

7.4.2.2 Clinical applications of heart rate variability analysis

Sleep apnoea

Polysomnography is a comprehensive screening tool for sleep apnoea; however, the cost makes it impractical for widespread screening. Heart rate variability and oximetry have both been proposed as low cost alternatives to polysomnography, but neither offers an acceptable level of reliability.

The combination of HRV with oximetry promises improved apnoea detection when compared to the use of either information source alone. This is a very exciting area of application, and the full potential of HRV analysis in the field of sleep apnoea is yet to be reached. This application of heart rate variability analysis has been touched on during this thesis, but has not been rigorously investigated.

Hypertension

The application of HRV analysis to the detection of hypertension also shows promise. It is possible that, through HRV, one can detect hypertension during its early stages by monitoring changes in the behaviour of the autonomic nervous system. However, the progression from mild to developed hypertension is poorly understood, so much work is needed before the prediction of this progression could be made reliably.

Appendix A Normalisation of heart rate and heart period

With no normalisation, the total power of the heart rate variability and heart period variability signals (derived from the same raw beat-time data) can differ considerably. Normalisation by total power obviously reduces this discrepancy. Here, it is shown that normalisation by mean heart rate and mean interval length has a similar effect.

Assume the beat to beat interval signal $i(t)$ shows a simple sinusoidal variation:

$$i(t) = \mu_i + A\cos\omega t. \tag{A.1}$$

The instantaneous heart rate may then be written:

$$r(t) = 1/(i(t)). \tag{A.2}$$

The average total power in the interval variability signal $i'(t) = i(t) - \overline{i(t)}$ is simply

$$P_{i'} = \frac{1}{2}A^2, \tag{A.3}$$

whereas for the heart rate variability $r'(t) = r(t) - \overline{r(t)}$ and the average total power is given by (Castiglioni, 1995):

$$P_{r'} = \frac{1 - \sqrt{1 - (A/\mu_i)^2}}{\mu_i^2\sqrt{(1 - (A/\mu_i)^2)^3}}. \tag{A.4}$$

$P_{r'}$ is obviously quite different from $P_{i'}$ and so with no normalisation, the results from HRV analysis may vary considerably depending on whether rate or interval variability is examined.

The rate and interval variability signals may be normalised by heart rate and heart period squared (Davidson *et al.*, 1997; Akselrod *et al.*, 1985; Akselrod *et al.*, 1981):

$$i''(t) = \frac{i(t) - \overline{i(t)}}{\overline{i(t)}}, \text{ and} \tag{A.5}$$

$$r''(t) = \frac{r(t) - \overline{r(t)}}{\overline{r(t)}}. \tag{A.6}$$

Mean heart rate $\overline{r(t)}$ is given by (Castiglioni, 1995):

$$\mu_r = (\mu_i^2 - A^2)^{-1/2}. \tag{A.7}$$

The total average power in the normalised interval variability signal is now:

$$P_{i''} = \frac{1}{2}\left(\frac{A}{\mu_i}\right)^2. \tag{A.8}$$

The total average power in the normalised heart rate variability signal may be calculated in a similar manner to that used to determine the (unnormalised) total average power:

$$P_{r''} = \frac{\omega}{2\pi}\int_0^{2\pi/\omega} |r''(t)|^2 dt, \tag{A.9}$$

$$= \frac{\omega}{2\pi\mu_r^2}\int_0^{2\pi/\omega} (r(t) - \mu_r)^2 dt \tag{A.10}$$

$$= \frac{1}{\mu_r^2}\left\{\frac{1 - \sqrt{1 - (A/\mu_i)^2}}{\mu_i^2\sqrt{(1 - (A/\mu_i)^2)^3}}\right\}, \tag{A.11}$$

which simplifies to

$$P_{r''} = \sqrt{1 - (A/\mu_i)^2} - (1 - (A/\mu_i)^2). \tag{A.12}$$

Now, the total average power of the heart rate variability signal matches more closely the total average power of the interval variability signal. Figure A.1a shows the discrepancy in total power between rate and interval variability signals: up to 55dB with strong modulation and fast beat rate. With normalisation by mean beat rate and interval length (Figure A.1b), there is little discrepancy between normalised rate variability power and normalised interval variability power. For strongly modulated signals ($A = 0.45$, mean heart interval $\overline{i(t)} = 0.5$) the two average powers differ by about 5 dB.

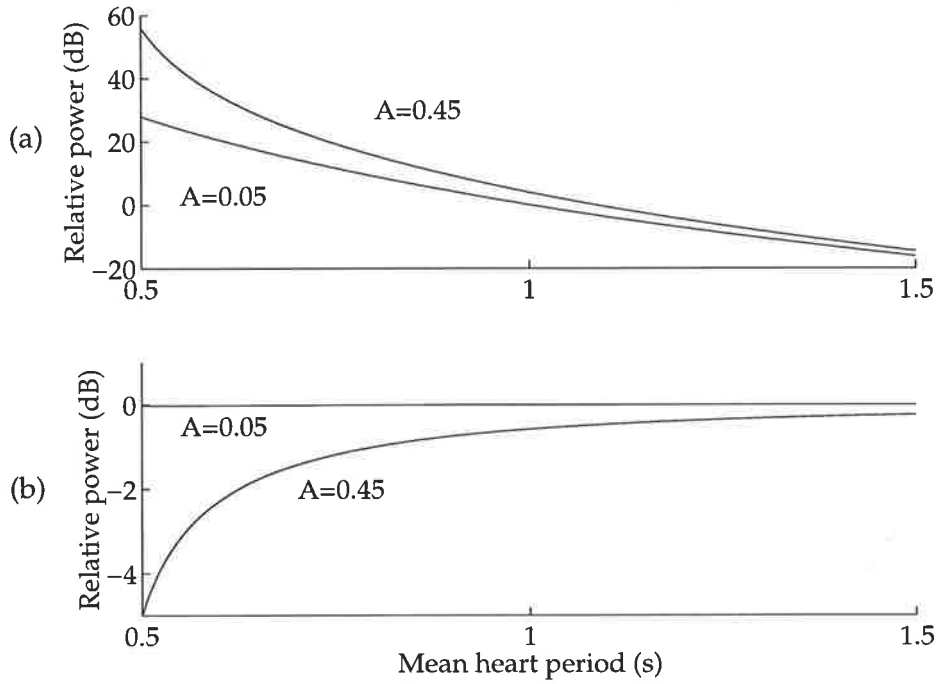

Figure A.1: Total average power of the heart rate signal, as a fraction of the total average power in the heart period signal. (a) No normalisation used, and weak ($A = 0.05$) and strong ($A = 0.45$) modulation; (b) Heart rate and heart period signals normalised by mean heart rate and mean heart period, respectively.

A comparison of results obtained using no normalisation, normalisation by total power, and normalisation by mean rate and interval is given in Section 4.2.2.

Appendix A Normalisation of heart rate and heart period

Appendix B Case study 1 results

This appendix contains the tables of results of the conventional band-power analysis of the spectral data discussed in Section 4.2. Differences were evaluated using a one-way analysis of variance. Significant differences are marked: *: significantly different from (supine, placebo) value ($p < 0.01$); ♣: significantly different from (tilt, placebo) value ($p < 0.01$); ♥: significantly different from (supine, propranolol) value ($p < 0.01$).

Measure	Supine Placebo (n=111)	Tilt Placebo (n=101)	Supine Propranolol (n=109)	Tilt Propranolol (n=121)
RRV total power	3.32(2.28)	1.75(1.86)*	5.97(3.90)*	1.97(1.50)*♥
RRV VLF power	0.53(0.46)	0.28(0.23)*	0.62(0.59)	0.44(0.44)♣♥
RRV LF power	0.62(0.56)	0.33(0.24)*	1.24(1.26)*	0.40(0.38)*♥
RRV HF power	1.39(0.86)	0.90(1.46)*	1.93(1.37)*	0.83(0.96)*♥
RRV LF/HF ratio	0.49(0.38)	1.36(1.30)*	1.19(2.03)*	0.96(1.07)*
HRV total power	4.11(2.29)	4.94(4.16)	4.16(2.82)	3.65(3.23)♣
HRV VLF power	0.75(0.64)	0.94(0.97)	0.50(0.54)*	0.96(1.07)♥
HRV LF power	0.78(0.64)	1.09(0.94)*	0.81(0.82)	0.76(0.88)♣
HRV HF power	1.72(0.87)	2.25(3.18)	1.58(1.15)	1.38(1.83)
HRV LF/HF ratio	0.49(0.38)	1.41(1.27)*	0.97(1.48)*	1.19(1.84)*

Table B.1: Mean(std) fixed-band measures of heart interval variability (RRV) and heart rate variability (HRV) power spectra, with no normalisation.

Measure	Supine Placebo (n=111)	Tilt Placebo (n=101)	Supine Propranolol (n=109)	Tilt Propranolol (n=121)
RRV total power	1.31(0.26)	1.45(0.33)*	1.19(0.14)*	1.53(0.37)*♥
RRV VLF power	0.23(0.16)	0.36(0.26)*	0.14(0.11)*	0.41(0.30)*♥
RRV LF power	0.23(0.10)	0.40(0.19)*	0.21(0.09)	0.33(0.16)*♣♥
RRV HF power	0.59(0.21)	0.48(0.24)*	0.51(0.29)	0.54(0.23)
RRV LF/HF ratio	0.49(0.38)	1.36(1.30)*	1.19(2.03)*	0.96(1.07)*
HRV total power	1.28(0.20)	1.40(0.28)*	1.17(0.12)*	1.49(0.35)*♥
HRV VLF power	0.23(0.16)	0.37(0.28)*	0.15(0.11)*	0.45(0.35)*♥
HRV LF power	0.23(0.10)	0.38(0.18)*	0.20(0.09)	0.32(0.15)*♣♥
HRV HF power	0.58(0.18)	0.44(0.23)*	0.52(0.27)	0.49(0.23)*
HRV LF/HF ratio	0.49(0.38)	1.41(1.27)*	0.97(1.48)*	1.19(1.84)*

Table B.2: Mean(std) fixed-band measures of heart interval variability (RRV) and heart rate variability (HRV) power spectra, normalised by total power (excluding components below 0.05Hz)

Measure	Supine Placebo (n=111)	Tilt Placebo (n=101)	Supine Propranolol (n=109)	Tilt Propranolol (n=121)
RRV total power	3.78(2.55)	2.98(2.82)	5.05(3.51)*	2.66(2.17)*♥
RRV VLF power	0.62(0.52)	0.51(0.40)	0.53(0.49)	0.59(0.56)
RRV LF power	0.71(0.65)	0.61(0.45)	1.05(1.09)*	0.54(0.55)♥
RRV HF power	1.57(0.87)	1.45(2.18)	1.71(1.25)	1.12(1.42)*♥
RRV LF/HF ratio	0.49(0.38)	1.36(1.30)*	1.19(2.03)*	0.96(1.07)*
HRV total power	3.68(2.22)	2.90(2.65)	4.95(3.28)*	2.65(2.09)*♥
HRV VLF power	0.66(0.57)	0.53(0.52)	0.58(0.67)	0.70(0.76)
HRV LF power	0.69(0.58)	0.60(0.49)	0.97(0.97)*	0.53(0.51)♥
HRV HF power	1.55(0.88)	1.40(2.14)	1.76(1.17)	1.02(1.24)*♥
HRV LF/HF ratio	0.49(0.38)	1.41(1.27)*	0.97(1.48)*	1.19(1.84)*

Table B.3: Mean(std) fixed-band measures of heart interval variability (RRV) and heart rate variability (HRV) power spectra, normalised by mean heart interval squared and mean heart rate squared, respectively.

Appendix C Rotation and translation of least-squares scalings

The least-squares scaling cost function is invariant to rotation and translation of the output mapping. However, the norm of the weight matrix of a radial basis function network (trained to produce a least-squares scaling) will vary with translation of the mapping. Here, it is shown that the translation required to minimise the weight matrix norm may be found analytically.

For a radial basis function network with K basis functions $\{\varphi_1 \ldots \varphi_K\}$, the network Jacobian is given by:

$$\boldsymbol{J} = \begin{bmatrix} \varphi_1(\boldsymbol{x}_1) & \varphi_2(\boldsymbol{x}_1) & \ldots & \varphi_K(\boldsymbol{x}_1) \\ \varphi_1(\boldsymbol{x}_2) & \varphi_2(\boldsymbol{x}_2) & \ldots & \varphi_K(\boldsymbol{x}_2) \\ \ldots & \ldots & \ldots & \ldots \\ \varphi_1(\boldsymbol{x}_N) & \varphi_2(\boldsymbol{x}_N) & \ldots & \varphi_K(\boldsymbol{x}_N) \end{bmatrix}. \tag{7.1}$$

For network weights $\boldsymbol{W}$, the network output is given by:

$$\boldsymbol{Y} = \boldsymbol{JW}, \tag{C.1}$$

and so for a network trained using simple supervised learning,

$$\boldsymbol{W} = \boldsymbol{J}^{+}\boldsymbol{Y}, \tag{C.2}$$

where $\boldsymbol{J}^{+}$ is the pseudo-inverse of the Jacobian.

A rotated solution $\tilde{\boldsymbol{Y}}$ is given by $\tilde{\boldsymbol{Y}} = \boldsymbol{YR}$, where $\boldsymbol{R}$ is an orthogonal rotation matrix, and the corresponding weight matrix $\tilde{\boldsymbol{W}} = \boldsymbol{J}^{+}\boldsymbol{YR}$. The norm of this weight matrix is:

$$\begin{aligned} tr[\tilde{\boldsymbol{W}}\tilde{\boldsymbol{W}}^T] &= tr[\boldsymbol{J}^{+}\boldsymbol{YRR}^T\boldsymbol{Y}^T(\boldsymbol{J}^{+})^T] \\ &= tr[\boldsymbol{J}^{+}\boldsymbol{YY}^T(\boldsymbol{J}^{+})^T] \\ &= tr[\boldsymbol{WW}^T] \end{aligned} \tag{C.3}$$

Thus, rotation of the solution does not affect the norm of the weight matrix.

Consider one dimension of the solution $\boldsymbol{y}$, and the corresponding column of the weight matrix $\boldsymbol{w}$. Translation of the solution in this dimension yields $\boldsymbol{y}' = \boldsymbol{y} + k\mathbf{1}$, where k is a scalar and $\mathbf{1}$ is a $N \times 1$ vector of ones. The weights are adjusted accordingly:

$$\begin{aligned} \boldsymbol{w}' &= \boldsymbol{J}^{+}(\boldsymbol{y} + k\mathbf{1}) \\ &= \boldsymbol{J}^{+}\boldsymbol{y} + k\boldsymbol{J}^{+}\mathbf{1} \\ &= \boldsymbol{w} + k\boldsymbol{j} \end{aligned} \quad , \tag{C.4}$$

where $\boldsymbol{j} = \boldsymbol{J}^{+}\mathbf{1}$.

Then the norm of this weight vector is:

$$\begin{aligned} \boldsymbol{w}'^{T}\boldsymbol{w}' &= (\boldsymbol{w} + k\boldsymbol{j})^{T}(\boldsymbol{w} + k\boldsymbol{j}) \\ &= \boldsymbol{w}^{T}\boldsymbol{w} + k^{2}\boldsymbol{j}^{T}\boldsymbol{j} + 2k\boldsymbol{w}^{T}\boldsymbol{j} \end{aligned} \quad , \tag{C.5}$$

which clearly varies with translation.

To find the translation which minimises $\boldsymbol{w}'^{T}\boldsymbol{w}'$, we differentiate:

$$\frac{d}{dk}(\boldsymbol{w}'^{T}\boldsymbol{w}') = 2k\boldsymbol{j}^{T}\boldsymbol{j} + 2\boldsymbol{w}^{T}\boldsymbol{j}, \tag{C.6}$$

and set this to zero, giving:

$$k = -\frac{\boldsymbol{w}^{T}\boldsymbol{j}}{\boldsymbol{j}^{T}\boldsymbol{j}}. \tag{C.7}$$

The preceding expressions were originally derived by Tipping (1996). The required translation is dependent on $\boldsymbol{w}$, and thus will vary depending on the orientation (rotation) of the initial solution.

The minimum achievable value norm of the weight matrix as a whole, however, is shown below to be invariant to rotation.

Consider a solution given by $\boldsymbol{Y} = \boldsymbol{J}\boldsymbol{W}$. This is rotated using the rotation matrix:

$$\boldsymbol{R} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}, \tag{C.8}$$

giving the "rotated weights":

$$\tilde{\boldsymbol{W}} = \boldsymbol{WR} = \begin{bmatrix} w_{11}\cos\theta - w_{12}\sin\theta & w_{11}\sin\theta + w_{12}\cos\theta \\ w_{21}\cos\theta - w_{22}\sin\theta & w_{21}\sin\theta + w_{22}\cos\theta \end{bmatrix}. \tag{C.9}$$

This solution is then translated to minimise the sum-of-squared weights, giving the weight matrix $\tilde{\boldsymbol{W}}'$. Take each column of the weight matrix $\tilde{\boldsymbol{W}}$ in turn:

$$\begin{aligned} \tilde{w}'_1 &= \tilde{w}_1 - \frac{\tilde{w}_1^T \boldsymbol{j}}{\boldsymbol{j}^T\boldsymbol{j}}\boldsymbol{j} \\ &= \tilde{w}_1 - \frac{[\,(w_{11}\cos\theta - w_{12}\sin\theta)j_1 + \ldots + (w_{k1}\cos\theta - w_{k2}\sin\theta)j_k]}{\boldsymbol{j}^T\boldsymbol{j}}\boldsymbol{j} \\ &= [\,\{(w_{11}\cos\theta - w_{12}\sin\theta) - j_1a_1\} \ldots \{(w_{k1}\cos\theta - w_{k2}\sin\theta) - j_ka_1\}\,]^T, \end{aligned} \tag{C.10}$$

where

$$a_1 = \frac{[\,(w_{11}\cos\theta - w_{12}\sin\theta)j_1 + \ldots + (w_{k1}\cos\theta - w_{k2}\sin\theta)j_k]}{\boldsymbol{j}^T\boldsymbol{j}}. \tag{C.11}$$

Similarly,

$$\tilde{w}'_2 = [\,\{(w_{11}\sin\theta + w_{12}\cos\theta) - j_1a_2\} \ldots \{(w_{k1}\sin\theta + w_{k2}\cos\theta) - j_ka_2\}\,]^T \tag{C.12}$$

and

$$a_2 = \frac{[\,(w_{11}\sin\theta + w_{12}\cos\theta)j_1 + \ldots + (w_{k1}\sin\theta + w_{k2}\cos\theta)j_k]}{\boldsymbol{j}^T\boldsymbol{j}}. \tag{C.13}$$

The sum-squared value of all elements in the rotated and translated weight matrix is then

$$\begin{aligned} \tilde{S}' &= (\tilde{w}'_{11})^2 + \ldots + (\tilde{w}'_{k1})^2 + (\tilde{w}'_{12})^2 + \ldots + (\tilde{w}'_{k2})^2 \\ &= ((w_{11}\cos\theta - w_{12}\sin\theta) - j_1a_1)^2 + \ldots + ((w_{k1}\cos\theta - w_{k2}\sin\theta) - j_ka_1)^2 \\ &\quad + ((w_{11}\sin\theta + w_{12}\cos\theta) - j_1a_2)^2 + \ldots + ((w_{k1}\sin\theta + w_{k2}\cos\theta) - j_ka_2)^2 \end{aligned}. \tag{C.14}$$

We expand and simplify the first and third terms shown above:

$$\tilde{S}' = w_{11}^2 + w_{12}^2 + (j_1 a_1)^2 + (j_1 a_2)^2 + 2w_{12} j_1 (a_1 \sin\theta - a_2 \cos\theta) \quad . \qquad \text{(C.15)}$$
$$- 2w_{11} j_1 (a_1 \cos\theta + a_2 \sin\theta) + \text{second and fourth terms}$$

Now consider the original solution weights W, translated without rotation to give W''. Through a similar procedure to the above, the sum-squared weights may be written:

$$S'' = (w_{11} - j_1 b_1)^2 + \ldots + (w_{k1} - j_k b_1)^2 + (w_{21} - j_1 b_2)^2 + \ldots + (w_{k2} - j_k b_2)^2$$

$$= w_{11}^2 + w_{12}^2 + (j_1 b_1)^2 + (j_1 b_2)^2 - 2w_{11} j_1 b_1 - 2w_{12} j_1 b_2$$
$$+ \text{second and fourth terms} \qquad , \qquad \text{(C.16)}$$

where

$$b_1 = \frac{[w_{11} j_1 + \ldots + w_{k1} j_k]}{j^T j}, \qquad \text{(C.17)}$$

and

$$b_2 = \frac{[w_{12} j_1 + \ldots + w_{k2} j_k]}{j^T j}. \qquad \text{(C.18)}$$

It is easily shown that:

$$b_1 = a_1 \cos\theta + a_2 \sin\theta, \qquad \text{(C.19)}$$

$$b_2 = a_2 \cos\theta - a_1 \sin\theta, \text{ and} \qquad \text{(C.20)}$$

$$a_1^2 + a_2^2 = b_1^2 + b_2^2. \qquad \text{(C.21)}$$

Substituting and comparing terms in C.15 and C.16 reveals the two sum-squared weight expressions to be equal. Thus, an arbitrary rotation of a solution does not change the minimum weight norm which may be achieved by subsequent translation. The required translation may be found analytically using C.7.

Appendix D The modified generative topographic mapping

This appendix gives the derivation of the expectation-maximisation equations for the re-formulated generative topographic mapping algorithm presented in Section 5.3.1.1. This reformulation of the algorithm allows prior information to be incorporated into the mapping.

The full expression for the joint likelihood $\langle L(\boldsymbol{W},\beta)\rangle$ is given by:

$$\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N} \log\left[\frac{N}{N_c}\sum_{i=1}^{I}\frac{1}{I}\left|1-c_n-\gamma_i\right|\left(\frac{\beta}{2\pi}\right)^{D/2}\exp\left(-\frac{\beta}{2}\left\|\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\|^2\right)\right], \quad \text{(D.1)}$$

which, for notational convenience, we write as:

$$\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N}\log\Omega. \quad \text{(D.2)}$$

Differentiating with respect to the weights $\boldsymbol{W}$ gives:

$$\begin{aligned}\frac{\partial}{\partial \boldsymbol{W}}\langle L(\boldsymbol{W},\beta)\rangle &= \sum_{n=1}^{N}\frac{1}{\Omega}\cdot\frac{N_{c_n}}{NI}\cdot\sum_{i=1}^{I}\frac{\partial}{\partial \boldsymbol{W}}\left\{\left|1-c_n-\gamma_i\right|\left(\frac{\beta}{2\pi}\right)^{D/2}\exp\left(-\frac{\beta}{2}\left\|\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\|^2\right)\right\}\\ &= -\sum_{n=1}^{N}\frac{1}{\Omega}\cdot\frac{N_{c_n}}{NI}\cdot\sum_{i=1}^{I}\left|1-c_n-\gamma_i\right|\left(\frac{\beta}{2\pi}\right)^{D/2}\exp\left(-\frac{\beta}{2}\left\|\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\|^2\right)\times\\ &\qquad \beta\left(\left\{\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\}\Phi^T(\boldsymbol{u}_i)\right)\\ &= -\beta\sum_{n=1}^{N}\frac{\sum_{i=1}^{I}\left|1-c_n-\gamma_i\right|\exp\left(-\frac{\beta}{2}\left\|\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\|^2\right)\left\{\boldsymbol{W}\Phi(\boldsymbol{u}_i)-\boldsymbol{x}_n\right\}\Phi^T(\boldsymbol{u}_i)}{\sum_{j=1}^{I}\left|1-c_n-\gamma_j\right|\exp\left(-\frac{\beta}{2}\left\|\boldsymbol{W}\Phi(\boldsymbol{u}_j)-\boldsymbol{x}_n\right\|^2\right)}\end{aligned} \quad \text{(D.3)}$$

This may be written:

$$\frac{\partial}{\partial \boldsymbol{W}}\langle L(\boldsymbol{W},\beta)\rangle = -\beta \sum_{n=1}^{N}\sum_{i=1}^{I} R_{ni}\{\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\}\,\Phi^T(\boldsymbol{u}_i)\,, \tag{D.4}$$

where

$$R_{ni} = \frac{|1 - c_n - \gamma_i| \exp\left(-\frac{\beta}{2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\right)}{\sum_{j=1}^{I} |1 - c_n - \gamma_j| \exp\left(-\frac{\beta}{2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_j) - \boldsymbol{x}_n\|^2\right)}. \tag{D.5}$$

Thus, the new weights at each iteration are given by the solution of:

$$\sum_{n=1}^{N}\sum_{i=1}^{I} R_{ni}\{\boldsymbol{W}_{new}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\}\,\Phi^T(\boldsymbol{u}_i) = 0\,. \tag{D.6}$$

Similarly, we may differentiate with respect to the variance $1/\beta$:

$$\frac{\partial}{\partial(1/\beta)}\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N}\frac{1}{\Omega}\cdot\frac{N_{c_n}}{NI}\cdot\sum_{i=1}^{I}\frac{\partial}{\partial\lambda}\{|1 - c_n - \gamma_i|\left(\frac{1}{2\pi\lambda}\right)^{D/2}\exp\left(-\frac{1}{2\lambda}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\right)\} \tag{D.7}$$

where we have used $\lambda = 1/\beta$ for clarity. Evaluating this derivative, we obtain:

$$\frac{\partial}{\partial(1/\beta)}\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N}\frac{1}{\Omega}\cdot\frac{N_{c_n}}{NI}\cdot\sum_{i=1}^{I}|1 - c_n - \gamma_i|\left(\frac{1}{2\pi\lambda}\right)^{D/2}\exp\left(-\frac{1}{2\lambda}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\right)\times$$
$$\{-\frac{D}{2\lambda} + \frac{1}{2\lambda^2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\} \tag{D.8}$$

Substituting R_{ni} (Equation D.5) gives:

$$\frac{\partial}{\partial(1/\beta)}\langle L(\boldsymbol{W},\beta)\rangle = \sum_{n=1}^{N}\sum_{i=1}^{I} R_{ni}\{-\frac{\beta D}{2} + \frac{\beta^2}{2}\|\boldsymbol{W}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\}\,, \tag{D.9}$$

and setting this derivative to zero gives the update equation:

$$\frac{1}{\beta_{new}} = \frac{1}{ND}\sum_{n=1}^{N}\sum_{i=1}^{I} R_{ni}\|\boldsymbol{W}_{new}\Phi(\boldsymbol{u}_i) - \boldsymbol{x}_n\|^2\,. \tag{D.10}$$

Appendix E Classification of heart rate variability in normal and hypertensive subjects

This appendix describes a preliminary study in which simple classification algorithms were applied to HRV data from normal and mildly hypertensive subjects. The methods used in this study do not follow closely the work in the remainder of the thesis and so it has been relegated to an appendix. Nevertheless, this work demonstrates the potential of HRV analysis for the detection of early hypertension. This study was published in: Raymond, B., Taverner, D., Nandagopal, D., and Mazumdar, J. (1997). Classification of heart rate variability in patients with mild hypertension. *Australasian Physical & Engineering Sciences in Medicine*, **20**(4):207–213.

E.1 Methods

A brief description of hypertension and of the relevance of HRV may be found in Section 2.3.2.

E.1.1 Patient population and experimental protocol

Data was originally collected for an investigation into the effects of angiotensin converting enzyme (ACE) inhibition on autonomic function (Nunan *et al.*, 1994). Hypertensives were chosen on the basis of mild to moderate primary hypertension without significant left ventricular hypertrophy (based on ECG voltage criteria) or significant renal discharge dysfunction (based on creatinine clearance calculated from plasma creatinine concentration). All hypertensive subjects were receiving therapy: ACE inhibitors, beta-adrenoceptor antagonists (beta blockers) or calcium channel antagonists. Normotensive controls were recruited on the basis of no history or treatment of hypertension, and a normal blood pressure on at least two occasions while on no treatment.

Each subject underwent on one occasion a variety of tests commonly used to evaluate autonomic function, including 3 minutes of isometric handgrip at 30% of maximum voluntary contraction. An electrocardiogram was recorded throughout the procedure.

The subject base comprised 20 normotensive subjects and 29 hypertensive subjects. A two minute sample of resting heart rate data was collected after 5 minutes rest when subjects were fasting, and had abstained from caffeine and cigarettes for at least 12 hours. A record of heart rate during handgrip was also taken for each subject, using the latter 2 minutes of the 3 minute handgrip test. The handgrip test was effective at stimulating sympathetic reflexes as measured by a mean systolic blood pressure rise of 22 ± 4 (SEM) mmHg from resting after 3 minutes.

E.1.2 Data preprocessing

The heart rate record was interpolated using cubic spline interpolation and resampled at 10Hz. To avoid aliasing, the resampled signal was low-pass filtered using an 8th order Chebychev filter with cutoff frequency of 0.4 Hz and decimated by 10, to give a uniform sampling rate of 1Hz. The time series was then demeaned and detrended by removing the best straight-line fit from the data.

Ectopic beats were avoided wherever possible when selecting data segments for analysis. Data segments with more than two ectopic beats were not used; isolated ectopic beats were removed using linear interpolation where necessary.

The frequency spectrum of the HRV signal was generated using an autoregressive (AR) model, with coefficients determined using the Yule-Walker method. A model order of 10 was chosen, slightly higher than that indicated using the Akaike information criterion (AIC) for model order selection. The AIC tends to underestimate model order for a noisy or non-AR process (Kay, 1988, pp. 234-240). For simplicity, fixed frequency bands were used to extract the low and high frequency components. The LF band was chosen as 0.04 to 0.14 Hz, and the HF band as 0.14 to 0.32 Hz. Figure E.1 shows a typical resting HRV signal and its power spectrum. The LF and HF bands are shaded.

E.1.3 Feature selection

Features were extracted from the HRV signal for subsequent classification. A number of candidate features were identified and combined into feature vectors. The effectiveness of each feature vector was then evaluated by considering its performance with various classifiers.

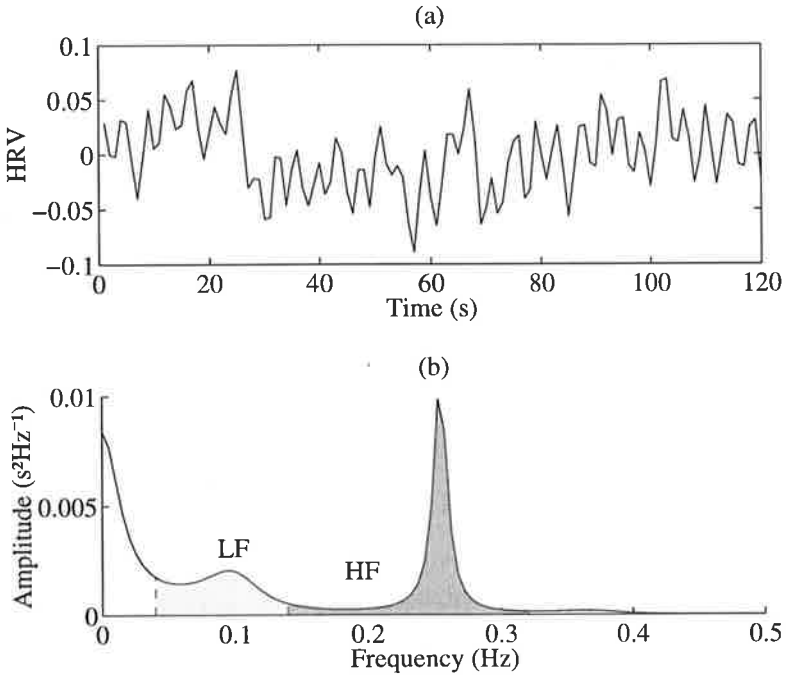

Figure E.1: (a) HRV signal and (b) estimated power spectrum of subject at rest. The LF and HF bands have been shaded.

The spectral components of the HRV signal reflect the modulation of the heart rate by the autonomic nervous system and hence are altered by a change in autonomic function such as is associated with hypertension. Indeed, an increased LF component and decreased HF component have been observed in the HRV spectra of hypertensive subjects (Malliani *et al.*, 1991b). Clearly, the spectral components are important features of the HRV signal and so made up the majority of the candidate features. Two methods of quantifying the spectral content of the predefined LF and HF frequency bands were used. The area under the power spectral density (PSD) curve in a band gives a measure of the power contained within that band. The root-mean-square (RMS) measure of the PSD within each band was also considered, where the RMS value was computed by taking the square root of the mean of the square of the PSD components within that band. This measure tends to accentuate peaks in the PSD. The area measures were normalised by dividing by the total area of the spectrum; while the RMS value of the spectrum was used to normalise the RMS measures. Figure E.2 shows the HRV and (unnormalised) power spectrum of a subject during isometric handgrip. The LF and HF areas have been shaded and the RMS values are indicated by the solid horizontal lines across each band.

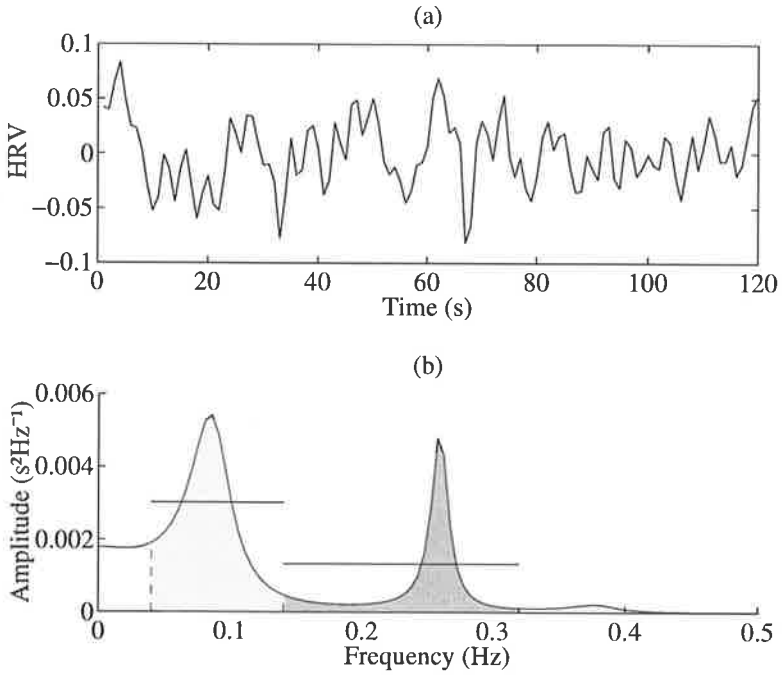

Figure E.2: (a) HRV signal and (b) estimated spectrum of subject during handgrip. The horizontal lines indicate the RMS value of the LF and HF bands (see text).

Time domain measures of the heart rate variation, such as variance, have been employed in other HRV applications (Kautzner *et al.*, 1995). Some common time domain measures have been included in the features here; see Table E.1 for a complete list of features used.

Features were then combined into feature vectors to be passed to the classification stage. Vectors of up to 6 features were used; an extensive but not exhaustive search of the possible combinations of features was made. Vectors consisting of time-domain features, frequency-domain features and a combination of both were investigated.

Two classes of data were defined: subjects diagnosed as normotensive, denoted C_n, and subjects diagnosed as hypertensive, denoted C_h, with each subject described by a feature vector, x_i.

E.1.4 Classifiers

A simple probabilistic classifier was one of the classifiers used in the study. Assuming the class probability density functions can each be adequately described by a Gaussian function, a

Feature label	Description
SDrest	Standard deviation of resting HRV signal
SDhg	Standard deviation of handgrip HRV signal
pNN50rest	Ratio of R-R intervals which differ from previous interval by >50 ms, resting heart tachogram
pNN50hg	Same, but for handgrip tachogram
rMSSDrest	RMS value of the difference between adjacent RR intervals in resting tachogram
rMSSDhg	Same, but for handgrip tachogram
LFrest	Power contained in LF band of resting HRV spectrum, as a fraction of the total power in the signal
LFhg	Fractional power contained in LF band of resting HRV spectrum
HFrest	Fractional power contained in HF band of resting HRV spectrum
HFhg	Fractional power contained in HF band of handgrip HRV spectrum
rmsLFrest	RMS value of resting HRV spectrum in LF band, normalised by the overall RMS value of the HRV spectrum
rmsLFhg	Same, but for handgrip HRV signal
rmsHFrest	RMS value of resting HRV spectrum in HF band, normalised by the overall RMS value of the HRV spectrum
rmsHFhg	Same, but for handgrip HRV signal
ARrest	AR parameters of resting HRV signal
ARhg	AR parameters of handgrip HRV signal

Table E.1: Description of diagnostic features for HRV signal.

discriminant function which gives minimum probability of error may be written (Fukunaga, 1990):

$$g(\boldsymbol{x}) = (\boldsymbol{x}-\boldsymbol{m}_h)^T K_h^{-1}(\boldsymbol{x}-\boldsymbol{m}_h) - (\boldsymbol{x}-\boldsymbol{m}_n)^T K_n^{-1}(\boldsymbol{x}-\boldsymbol{m}_n) + \ln\left(\frac{|K_h|}{|K_n|}\right) - 2\ln\left(\frac{P(C_h)}{P(C_n)}\right) \tag{E.1}$$

where the class-wise means of the feature data are denoted by $\boldsymbol{m}_h$ and $\boldsymbol{m}_n$, and the covariances by K_h and K_n.

The training data was used to estimate the means and covariances above. The classifier was then tested on the test data; an unknown subject described by a feature vector $\boldsymbol{x}_u$ was classified as normotensive if $g(\boldsymbol{x}_u) > 0$.

The second classifier used was the voting k-nearest neighbours classifier. The k vectors in the training set closest to some unknown vector $\boldsymbol{x}_u$ were found using either the Euclidean or city-block distance measure. The unknown vector was then classified as belonging to the class which made up the majority of these neighbouring vectors. Odd numbers of neighbours k were investigated, up to $k = 11$. Leave-one-out cross validation was used to evaluate classification performance.

E.2 Results and discussion

E.2.1 Comparison with previous work

The trends of the data set were first compared to published results. An unpaired t-test was used to evaluate between-group differences.

The LF component in the normotensive group (NT) at rest was greater than that of hypertensive (HT) subjects ($p < 0.001$), and the HF component for the NT group was less than that of the HT group ($p < 0.01$). These results are consistent with previously reported effects of drug therapy on the HRV spectrum. In particular, β-sympathetic blockade has been shown to reduce the LF and increase the HF component of the HRV spectrum (Tuininga *et al.*, 1994; Pagani *et al.*, 1986), while ACE-inhibitors increase parasympathetic action in patients with chronic heart failure, and calcium antagonists may reduce the LF component in post-myocardial infarction patients (Tuininga *et al.*, 1994). A study of hypertensive subjects receiving no drug therapy (Malliani *et al.*, 1991b) found an increased LF and decreased HF component in resting hypertensive subjects, in agreement with the finding of increased sympathetic and reduced vagal tone in early hypertension. In the current study, the effects of the drug treatment in resting hypertensive subjects would appear to outweigh the changes in autonomic function due to hypertension.

In response to isometric handgrip, the HT group showed a larger increase in LF ($p < 0.005$) and decrease in HF ($p < 0.001$) component compared to the NT group. Published results

(McAllister Jr, 1979) indicate that the pressor effects of handgrip in patients with hypertension may be the same or greater than those in normal subjects. The response to handgrip is less likely to have been affected by the antihypertensive drugs: ramipril (an ACE inhibitor) has been shown to have no effect on the maximum change in mean blood pressure or heart rate, or on the change in plasma NE concentration (Sugimoto *et al.*, 1989).

These results indicate that classification results based on features relating to the resting measures of the HRV signal are likely to have been affected by the antihypertensive drugs, and so must be treated with some caution. Features corresponding to handgrip, or changes in response to handgrip are less likely to have been affected by the drug therapy of the hypertensive subjects.

Figures E.1 and E.2 show the HRV signal and PSD of the same hypertensive subject at rest and during handgrip, respectively. The increased LF and decreased HF components with handgrip can be seen.

The ratio of the LF to HF components of the HRV spectrum has been cited as a marker of sympathovagal balance (Pagani *et al.*, 1986). Figure E.3 shows a plot of LFrest/HFrest against LFhg/HFhg for the two subject groups. On this plot, the normal and hypertensive points are almost linearly separable.

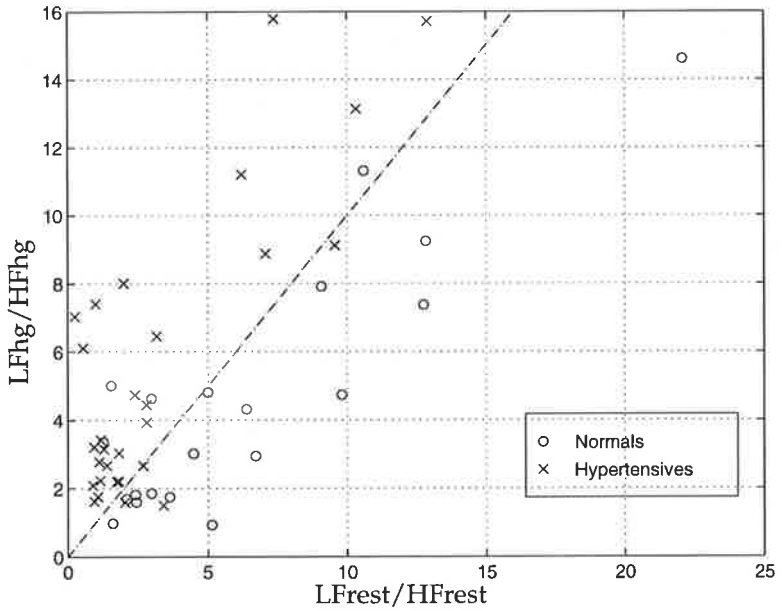

Figure E.3: Plot of LFrest/HFrest against LFhg/HFhg for normal and hypertensive subjects. The dashed line indicates an example linear decision boundary (see text).

E.2.2 Classification results

The Bayes and nearest neighbour classifiers were applied to each feature vector extracted from the data. The best performance of each classifier is listed in Table E.2, with separate results for the two different distance measures for the nearest neighbour classifier. The number of neighbours k in these cases is also listed. A false positive result is a normal subject classified as hypertensive; a false negative is a hypertensive subject classified as normotensive. The best classification result of 90% correct classifications was achieved using the nearest neighbour classifier, with the Euclidean distance metric, k=1, and a feature vector containing the change in the RMS LF value with handgrip (rmsLFrest-rmsLFhg) along with the standard deviation of the resting HRV signal, SDrest. The best result for the Bayes classifier of 84% was obtained with three different feature vectors.

The better distance metric for the nearest neighbour classifier was the Euclidean distance, providing both the single best classification result and the highest number of classifications of at least 80% correct, with a total of 24 (for k=1,3,5,7). The city block metric produced 20 such results. The mean classification rate (over all feature vectors tested) was 68% for both the Euclidean and city block distances. For the Bayes classifier, the average correct classification rate was also 68%.

The performance of the nearest-neighbour classifiers increased with increasing k up to a value of $k = 7$, with further increases in k producing deteriorating results. Figure E.4 shows a plot of mean correct classifications against k. As the number of neighbours increases beyond 7, the classifier becomes over-smoothed and performance deteriorates.

Feature vectors composed only of time-domain features proved ineffective for classification purposes, with an average result of 56% correct classification. However, many of the best feature vectors contained both frequency-domain and time-domain features, indicating that while time-domain features alone may not provide sufficient information, the information they provide complements that of the frequency-domain features.

For the Bayes classifier, 9 feature vectors produced 80% or better classification; each of these 9 vectors included either a term quantifying a change invoked by handgrip (e.g. LFrest-LFhg) or the same feature in both the rest and handgrip context. For the nearest neighbour classifiers, a total of 51 combinations of feature vector, k and distance metric produced a result of 85% or better; either the change in the LF/HF ratio or the change in the LF or HF

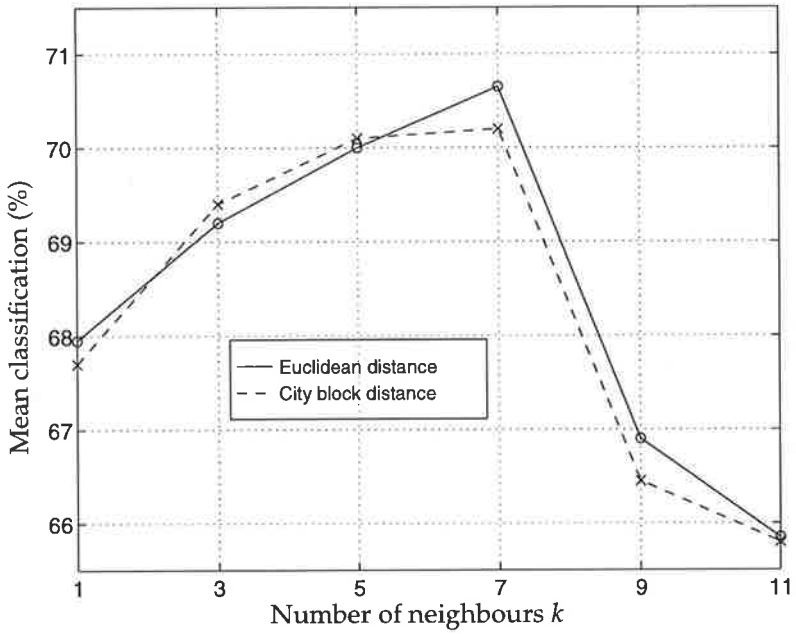

Figure E.4: Mean classifier performance (over all feature vectors) for increasing k.

component with handgrip appeared in 42 of these. These results suggest that the response of the cardiovascular system to handgrip is altered in a consistent manner by the presence of hypertension. If the results pertaining to area and RMS measures are considered separately, the RMS values constitute 4 of the 9 vectors associated with the Bayes classifier, and 40 of the 42 feature vectors associated with the nearest neighbour classifiers. The remaining vectors use area measures to quantify the LF and HF components. For the nearest-neighbour classifier, the RMS measures of the LF and HF bands of the HRV power spectrum are a better feature than the area measures. For the Bayes classifier, the two types of measurement gave much the same level of discrimination.

During the study, it was noted that the results were quite sensitive to the choice of LF and HF band limits.

E.3 Conclusions

The work presented in this section shows promise for the detection of the autonomic disturbance which accompanies and precedes the hypertensive state. Simple classifiers were investigated, and correct classification performance of up to 90% was obtained using these

Classifier	Elements of feature vector	False positives (max. 20)	False negatives (max. 29)	Correct classific -ations (%)
Bayes	LFrest-LFhg	3	5	84
	LFrest/HFrest, LFhg/ HFhg, rMSSDrest-rMSSDhg	4	4	84
	rmsLFrest/ rmsHFrest, rmsLFhg/ rmsHFhg, rMSSDrest/ rMSSDhg	3	5	84
k-nearest neighbour, Euclidean distance metric (k=1)	rmsLFrest-rmsLFhg, SDrest	2	3	90
k-nearest neighbour, city block distance metric (k=1)	rmsLFrest-rmsLFhg, SDrest	3	3	88
(k=1)	rmsLFrest-rmsLFhg, SDhg	3	3	88
(k=3)	rmsHFrest-rmsHFhg, SDrest, SDhg	4	2	88
(k=7)	rmsLFrest-rmsLFhg, rmsHFrest-rmsHFhg	4	2	88

Table E.2: Best classification results by classifier type.

techniques. The hypertensive subjects used in the study were all taking antihypertensive drugs and these may have affected the findings presented here. These or similar techniques applied to untreated subjects with only marginally elevated blood pressure may assist the prediction of which subjects are destined to develop sustained hypertension.

Reference list

Adrian, E., Bronk, D., and Phillips, G. (1932). Discharges in mammalian sympathetic nerves. *Journal of Physiology*, **74**:115–133.

Akay, M. and Mulder, E. (1996). Examining fetal heart-rate variability using matching pursuits. *IEEE Engineering in Medicine and Biology Magazine*, **15**(5):64–67.

Akay, M., editor (1998). *Time frequency and wavelets in biomedical signal processing*. IEEE, New York, USA, New York, USA.

Akselrod, S., Gordon, D., Ubel, F., Shannon, D., Barger, A., and Cohen, R. (1981). Power spectrum analysis of heart rate fluctuation: a quantitative probe of beat-to-beat cardiovascular control. *Science*, **213**:220–222.

Akselrod, S., Gordon, D., Madwed, J., Snidman, N., Shannon, D., and Cohen, R. (1985). Hemodynamic regulation: investigation by spectral analysis. *American Journal of Physiology*, **249**:H867–H875.

Akselrod, S., Eliash, S., Oz, O., and Cohen, S. (1987). Hemodynamic regulation in SHR: investigation by spectral analysis. *American Journal of Physiology*, **253**:H176–H183.

Akselrod, S. (1995). Components of heart rate variability: basic studies. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 147–164. Futura Publishing Company, Inc., Armonk, USA.

Alcalay, M., Izraeli, S., Wallach-Kapon, R., Tochner, Z., Benjamini, Y., and Akselrod, S. (1992). Paradoxical pharmacodynamic effect of atropine on parasympathetic control: a study by spectral analysis of heart rate fluctuations. *Clinical Pharmacology and Therapeutics*, **52**(5):518–527.

Anan, T., Sunagawa, K., Araki, H., and Nakamura, M. (1990). Arrhythmia analysis by successive RR plotting. *Journal of Electrocardiology*, **23**:243–248.

Anderson, A. (1971). Ordination methods in ecology. *Journal of Ecology*, **59**:713–726.

Anderson, E., Sinkey, C., Lawton, W., and Mark, A. (1989). Elevated sympathetic nerve activity in borderline hypertensive humans: evidence from direct intraneural recordings. *Hypertension*, **14**:177–183.

Appel, M., Berger, R., Saul, J., Smith, J., and Cohen, R. (1989). Beat to beat variability in cardiovascular variables: noise or music? *Journal of the American College of Cardiology*, **14**:1139–1148.

Arai, Y., Saul, J., Albrecht, P., Hartley, L., Lilly, L., Cohen, R., and Colucci, W. (1989). Modulation of cardiac autonomic activity during and immediately after exercise. *American Journal of Physiology*, **256**:H132–H141.

Bakardjian, H. and Yamamoto, K. (1994). Morphology evaluation of heart heart variability power spectrum. *Proceedings of the 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 2, pages 1332–1333, Baltimore, USA. IEEE, New York, USA.

Baldwa, V. and Ewing, D. (1977). Heart rate response to Valsalva manoeuvre: reproducibility in normals, and relation to variation in resting heart rate in diabetics. *British Heart Journal*, **39**:641–644.

Barschdorff, D., Jaeger, A., and Trowitzch, E. (1994). Automatic assessment of an "apnoea-severity-factor" combined with heart rate analysis during polysomnographic examinations in infants. *Computers in Cardiology 1994*, pages 101–104, Bethesda, USA. IEEE Computer Society Press, Los Amitos, USA.

Bartlett, M., Murray, A., and Dunlop, W. (1991). Properties of fetal heartbeat intervals during labour. *Journal of Biomedical Engineering*, **13**(2):169–172.

Bartoli, F., Baselli, G., and Cerutti, S. (1983). Application of identification and linear filtering algorithms to the R-R interval measurements. *Computers in Cardiology 1983*, pages 485–488, Seattle, USA. IEEE Computer Society Press, Los Amitos, USA.

Baselli, G., Cerutti, S., Civardi, S., Malliani, A., Orsi, G., Pagani, M., and Rizzo, G. (1988a). Parameter extraction from heart rate and arterial blood pressure variability signals in dogs for the validation of a physiological model. *Computers in Biology and Medicine*, **18**(1):1–15.

Baselli, G., Cerutti, S., Civardi, S., Malliani, A., and Pagani, M. (1988b). Cardiovascular variability signals: towards the identification of a closed-loop model of the neural control mechanisms. *IEEE Transactions on Biomedical Engineering*, **35**(12):1033–1046.

Baselli, G., Cerutti, S., Interdonato, T., Orizio, C., Perini, R., and Veicsteinas, A. (1988c). HRV during exercise in sedentary subjects and athletes. *Computers in Cardiology 1988*, pages 319–322, Washington DC, USA. IEEE Computer Society Press, Los Amitos, USA.

Baselli, G., Biancardi, L., Perini, R., Milesi, S., Veicsteinas, A., and Cerutti, S. (1991). Heart rate variability during dynamic exercise in supine position in sedentary subjects. *Computers in Cardiology 1991*, pages 437–440, Venice, Italy. IEEE Computer Society Press, Los Amitos, USA.

Baselli, G., Cerutti, S., Badilini, F., Biancardi, L., Porta, A., Pagani, M., Lombardi, F., Rimoldi, O., Furlan, R., and Malliani, A. (1994). Model for the assessment of heart period and arterial pressure variability interactions and of respiration influences. *Medical & Biological Engineering & Computing*, **32**:143–152.

Bayly, E. (1968). Spectral analysis of pulse frequency modulation in the nervous systems. *IEEE Transactions on Biomedical Engineering*, **BME-15**(4):257–265.

Bellavere, F., Balzani, I., de Masi, G., Carraro, M., Carenza, P., Cobelli, C., and Thomaseth, K. (1992). Power spectral analysis of heart rate variation improves assessment of diabetic cardiac autonomic neuropathy. *Diabetes*, **41**:633–640.

Bennet, T., Fentem, P., Fitton, D., Hampton, J., Hosking, D., and Riggott, P. (1977). Assessment of vagal control of the heart in diabetes: measures of R-R interval variation under different conditions. *British Heart Journal*, **39**:25–28.

Bennet, T., Farquhar, I., Hosking, D., and Hampton, J. (1978). Assessment of methods for estimating autonomic nervous control of the heart in patients with diabetes mellitus. *Diabetes*, **27**:1167–1174.

Bentley, P., McDonnell, J., and Grant, P. (1995). Classification of native heart valve sounds using the Choi-Williams time-frequency distribution. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 1083–1084, Montreal, Canada. IEEE, New York, USA.

Bentley, P., Grant, P., and McDonnell, J. (1998). Time-frequency and time-scale techniques for the

classification of native and bioprosthetic heart valve sounds. *IEEE Transactions on Biomedical Engineering*, **45**(1):125–128.

Berger, R., Akselrod, S., Gordon, D., and Cohen, R. (1986). An efficient algorithm for spectral analysis of heart rate variability. *IEEE Transactions on Biomedical Engineering*, **BME-33**(9):900–904.

Berne, R. and Levy, M. (1986). *Cardiovascular physiology*. The C.V. Mosby Company, St Louis, USA, 5th edition.

Bianchi, A., Bontempi, B., Cerutti, S., Gianoglio, P., Comi, G., and Natali Sora, M. (1990). Spectral analysis of heart rate variability signal and respiration in diabetic subjects. *Medical & Biological Engineering & Computing*, **28**:205–211.

Bianchi, A., Mainardi, L., Petrucci, E., Signorini, M., Mainardi, M., and Cerutti, S. (1993). Time-variant power spectrum analysis for the detection of transient episodes in HRV signal. *IEEE Transactions on Biomedical Engineering*, **40**(2):136–144.

Bianchi, A., Scholz, U., Mainardi, L., Orlandini, P., Pozza, G., and Cerutti, S. (1994). Extraction of the respiratory influence from the heart rate variability signal by means of a lattice adaptive filter. *Proceedings of the 16th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, pages 121–122, Baltimore, USA. IEEE, New York, USA.

Bigger Jr, J., Kleiger, R., Fleiss, J., Rolnitzky, L., Steinman, R., Miller, J., and The Multicenter Post-Infarction Research Group (1988). Components of heart rate variability measured during healing of acute myocardial infarction. *American Journal of Cardiology*, **61**:208–215.

Bigger Jr, J., Fleiss, J., Rolnitzky, L., and Steinman, R. (1992a). Stability over time of heart period variability in patients with previous myocardial infarction and ventricular arrhythmias. *American Journal of Cardiology*, **67**:718–723.

Bigger Jr, J., Fleiss, J., Steinman, R., Rolnitzky, L., Kleiger, R., and Rottman, J. (1992b). Frequency domain measures of heart period variability and mortality after myocardial infarction. *Circulation*, **85**:164–171.

Billings, S. and Alturki, F. (1992). Performance monitoring in non-linear adaptive noise cancellation. *Journal of Sound and Vibration*, **157**(1):161–175.

Billings, S. and Fung, C. (1995). Recurrent radial basis function networks for adaptive noise cancellation. *Neural Networks*, **8**(2):273–290.

Bishop, C. (1993). Curvature-driven smoothing: a learning algorithm for feed-forward networks. *IEEE Transactions on Neural Networks*, **4**(5):882–884.

Bishop, C. (1995). *Neural networks for pattern recognition*. Oxford University Press, New York, USA.

Bishop, C. and Tipping, M. (1998). A hierarchical latent variable model for data visualisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. To appear. Available online from `http://www.ncrg.aston.ac.uk/`

Bishop, C., Svensén, M., and Williams, C. (1998a). Developments of the generative topographic mapping. *Neurocomputing*. To appear. Available online from `http://www.ncrg.aston.ac.uk/`

Bishop, C., Svensén, M., and Williams, C. (1998b). GTM: the generative topographic mapping. *Neural Computation*, **10**(1):215–234.

Blackman, R. and Tukey, J. (1958). *The Measurement of Power Spectra*. Dover, New York, USA.

Bootsma, M., Swenne, C., Cats, V., and Bruschke, A. (1994). Pharmacological validation of heart rate variability measures that assess sympathetic or vagal tone or the sympathovagal balance (ECG analysis). *Computers in Cardiology 1994*, pages 301–304, Bethesda, USA. IEEE Computer Society Press, Los Amitos, USA.

Bow, S.-T. (1984). *Pattern recognition: applications to large data-set problems*. Marcel Dekker, Inc., New York, USA.

Bradley, P., Mortimore, I., and Douglas, N. (1995). Comparison of polysomnography with ResCare Autoset in the diagnosis of the sleep apnoea/hypopnoea syndrome. *Thorax*, **50**:1201–1203.

Branston, N., El-Deredy, W., Lisboa, P., and Thomas, D. (1998). Visualisation of human basal ganglia neuron responses using the generative topographic mapping (GTM) algorithm. *Proceedings of the IEEE World Congress on Computational Intelligence*, volume 1, pages 326–331, Anchorage, USA. IEEE, New York, USA.

Brouwer, J., Viersma, J., van Veldhuisen, D., Man in 't Veld, A., Sijbring, P., Haaksma, J., Dijk, W., and Lie, K. (1995). Usefulness of heart rate variability in predicting drug efficacy (Metaprolol vs Diltiazem) in patients with stable angina pectoris. *American Journal of Cardiology*, **76**:759–763.

Brown, T., Beightol, L., Koh, J., and Eckberg, D. (1993). Important influence of respiration on human R-R interval power spectra is largely ignored. *Journal of Applied Physiology*, **75**(5):2310–2317.

Calcagnini Jr, G., Lino, S., Strano, S., Cerutti, S., and Calcagnini, G. (1993). 1/f spectrum slope of 24 hour heart rate variability signal in normal and hypertensive subjects. *Computers in Cardiology 1993*, pages 543–545, London, UK. IEEE Computer Society Press, Los Amitos, USA.

Carreira-Perpiñán, M. and Renals, S. (1998). Dimensionality reduction of electropalatographic data using latent variable models. *Speech Communication*. To appear. Available online from `http://www.dcs.shef.ac.uk/~miguel/`

Castiglioni, P. (1995). Evaluation of heart rhythm variability by heart rate or heart period: differences, pitfalls and help from logarithms. *Medical & Biological Engineering & Computing*, **33**:323–330.

Çagatay Güler, E., Sankur, B., Kahya, Y., and Raudys, S. (1998). Visual classification of medical data using MLP mapping. *Computers in Biology and Medicine*, **28**(3):275–287.

Cerutti, S. (1995). Spectral analysis of the heart rate variability signal. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 63–74. Futura Publishing Company, Inc., Armonk, USA.

Chan, H.-L., Lin, J.-L., Du, C.-C., and Wu, C.-P. (1994). Lower frequency decomposition of 24-hour heart rate by Wigner-Ville spectral analysis. *Proceedings of the 8th International Conference on Biomedical Engineering*, pages 236–238, Singapore.

Chan, H.-L., Lin, J.-L., Huang, H.-H., and Wu, C.-P. (1997). Elimination of interference component in Wigner-Ville distribution for the signal with 1/f spectral characteristic. *IEEE Transaction on Biomedical Engineering*, **44**(9):903–907.

Chen, S., Billings, S., and Luo, W. (1989). Orthogonal least squares methods and their application to non-linear system identification. *International Journal of Control*, **50**(5):1873–1896.

Chen, S., Cowan, C., and Grant, P. (1991). Orthogonal least squares learning algorithm for radial

basis function networks. *IEEE Transactions on Neural Networks*, **2**(2):302–309.

Chess, G., La Belle, K., Milne, J., and Calaresu, F. (1975). Spectral analysis as a diagnostic aid in the management of high-risk pregnancy. *American Journal of Obstetrics and Gynecology*, **121**(4):471–474.

Choi, H. and Williams, W. (1989). Improved time-frequency representation of multicomponent signals using exponential kernels. *IEEE Transactions on Acoustics, Speech and Signal Processing*, **37**(6):862–71.

Chon, K., Jung, V., and Cohen, R. (1995). Nonlinear system identification of heart rate variability via an artificial neural network model. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 161–162, Montreal, Canada. IEEE, New York, USA.

Chon, K., Mullen, T., and Cohen, R. (1996). A dual-input nonlinear system analysis of autonomic modulation of heart rate. *IEEE Transactions on Biomedical Engineering*, **43**:530–544.

Chon, K., Mukkamala, R., Toska, K., Mullen, T., Armoundas, A., and Cohen, R. (1997). Linear and nonlinear system identification of autonomic heart-rate modulation. *IEEE Engineering in Medicine and Biology Magazine*, **16**(5):96–105.

Christini, D., Bennett, F., Lutchen, K., Ahmed, H., Hausdorff, J., and Oriol, N. (1995). Application of linear and nonlinear time series modeling to heart rate dynamics analysis. *IEEE Transactions on Biomedical Engineering*, **42**(4):411–415.

Chui, C. (1992). *An Introduction to Wavelets*. Academic Press, Inc., San Diego, USA.

Claasen, T. and Mecklenbrauker, W. (1980). The Wigner distribution - a tool for time-frequency signal analysis; part III: relations with other time-frequency signal transformations. *Philips Journal of Research*, **35**:372–389.

Clark, C. and Mapstone, R. (1985). Autonomic neuropathy in ocular hypertension. *Lancet*, **27**:185–187.

Clayton, R., Bowman, A., Ford, G., and Murray, A. (1995). Measurement of baroreflex gain from heart rate and blood pressure spectra: a comparison of spectral estimation techniques. *Physiological Measurement*, **16**(2):131–139.

Clayton, R., Lord, S., McComb, J., and Murray, A. (1997). Comparison of autoregressive and Fourier transform based techniques for estimating RR interval spectra. *Computers in Cardiology 1997*, pages 379–382, Lund, Sweden. IEEE Computer Society Press, Los Amitos, USA.

Clayton, R. and Murray, A. (1998). Comparison of techniques for time-frequency analysis of the ECG during human ventricular fibrillation. *IEE Proceedings Science, Measurement and Technology*, **145**(6):301–306.

Cohen, L. (1989). Time-frequency distributions - a review. *Proceedings of the IEEE*, **77**(7):941–981.

Coumel, P., Maison-Blanche, P., and Catuli, D. (1995). Heart rate and heart rate variability. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 207–222. Futura Publishing Company, Inc., Armonk, USA.

Courtemanche, M., Talajic, M., and Kaplan, D. (1992). Multiparameter quantification of 24 hour heart rate variability. *Computers in Cardiology 1992*, pages 299–302, Durham, USA. IEEE Computer Society Press, Los Amitos, USA.

Cox, T. and Ferry, G. (1993). Discriminant analysis using non-metric multidimensional scaling. *Pattern Recognition*, **26**(1):145–153.

Coxon, A. (1982). *The User's Guide to Multidimensional Scaling*. Heinemann Educational Books Ltd., London, UK.

Craelius, W., Akay, M., and Tangella, M. (1992). Heart rate variability as an index of autonomic imbalance in patients with recent myocardial infarction. *Medical & Biological Engineering & Computing*, **30**:385–388.

Curcie, D. and Craelius, W. (1997). Recognition of individual heart rate patterns with cepstral vectors. *Biological Cybernetics*, **77**:103–109.

Danev, S. and de Winter, C. (1971). Heart rate deceleration after erroneous responses. *Psychologische Forschung*, **35**:27–34.

Daoudi, M., Hamad, D., and Postaire, J.-G. (1993). A display oriented technique for interactive pattern recognition by multilayer neural network. *Proceedings of the 1993 IEEE International Conference on Neural Networks*, pages 1633–1637, San Francisco, USA. IEEE, New York, USA.

Davidson, S., Reina, N., Shefi, O., Hai-Tov, U., and Akselrod, S. (1997). Spectral analysis of heart rate fluctuations and optimum thermal management for low birth weight infants. *Medical & Biological Engineering & Computing*, **35**:619–625.

Davison, M. L. (1983). *Multidimensional Scaling*. John Wiley & Sons, New York, USA.

de Boer, R., Karemaker, J., and Strackee, J. (1984). Comparing spectra of a series of point events particularly for heart rate variability data. *IEEE Transactions on Biomedical Engineering*, **BME-31**(4):384–387.

de Boer, R., Karemaker, J., and Strackee, J. (1985a). Relationships between short-term blood-pressure fluctuations and heart-rate variability in resting subjects II: a simple model. *Medical & Biological Engineering & Computing*, **23**:359–364.

de Boer, R., Karemaker, J., and Strackee, J. (1985b). Spectrum of a series of point events, generated by the integral pulse frequency modulation model. *Medical & Biological Engineering & Computing*, **23**:138–142.

de Boer, R., Karemaker, J., and Strackee, J. (1987). Hemodynamic fluctuations and baroreflex sensitivity in humans: a beat-to-beat model. *American Journal of Physiology*, **253**:680–689.

de Molina, F. and Perl, E. (1965). Sympathetic activity and the systemic circulation in the spinal cat. *Journal of Physiology*, **181**:82–102.

Demartines, P. and Hérault, J. (1997). Curvilinear components analysis: a self-organizing neural network for nonlinear mapping of data sets. *IEEE Transactions on Neural Networks*, **8**(1):148–154.

Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B*, **39**(1):1–38.

Douglas, N., Thomas, S., and Jan, M. (1992). Clinical value of polysomnography. *Lancet*, **339**:347–350.

Duncan, G., Lambie, D., Johnson, R., and Whiteside, E. (1980). Evidence of vagal neuropathy in chronic alcoholics. *Lancet*, **15**:1053–1057.

Egelund, N. (1982). Spectral analysis of heart rate variability as an indicator of driver fatigue. *Ergometrics*, **25**:663.

Erwin, E., Obermayer, K., and Schulten, K. (1992). Self-organizing maps: ordering, convergence properties and energy functions. *Biological Cybernetics*, **67**:47–55.

Farrell, T., Bashir, Y., Cripps, T., Malik, M., Poloniecki, J., Bennett, E., Ward, D., and Camm, A. (1991). Risk stratification for arrhythmic events in post infarction patients based on heart rate variability, ambulatory electrocardiographic variables and the signal-averaged electrocardiogram. *Journal of the American College of Cardiology*, **18**:687–697.

Felgueiras, C., Marques de Sa, J., Bernardes, J., and Gama, S. (1998). Classification of foetal heart rate sequences based on fractal features. *Medical & Biological Engineering & Computing*, **36**(2):197–201.

Figliola, A. and Serrano, E. (1997). Analysis of physiological time series using wavelet transforms. *IEEE Engineering in Medicine and Biology Magazine*, **16**(3):74–79.

Fletcher, E. (1986). History, techniques, and definitions in sleep related disorders. In Fletcher, E., editor, *Abnormalities of Respiration during Sleep*. Grune & Stratton, Inc., Orlando, USA.

Fletcher, E. (1995). The relationship between systemic hypertension and obstructive sleep apnea: facts and theory. *American Journal of Medicine*, **98**:118–128.

Freeman, R., Saul, J., Roberts, M., Berger, R., Broadbridge, C., and Cohen, R. (1991). Spectral analysis of heart rate in diabetic autonomic neuropathy. *Archives of Neurology*, **48**:185–90.

Friedman, J. and Tukey, J. (1974). A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers*, **C-23**:881–889.

Fukunaga, K. (1990). *Introduction to statistical pattern recognition*. Academic Press, Boston, USA, 2nd edition.

Furui, S. (1989). *Digital Speech Processing, Synthesis and Recognition*. Marcel Dekker, Inc., New York, USA.

Gabor, D. (1946). Theory of communication. *Journal of the IEE*, **93**:429–457.

Goodhill, G. and Sejnowski, T. (1996). Quantifying neighbourhood preservation in topographic mappings. *Proceedings of the 3rd Joint Symposium on Neural Computation*, volume 6, pages 61–82, Pasadena, USA. Avaliable online from `http://www.giccs.georgetown.edu/~geoff/`

Goodhill, G. and Sejnowski, T. (1997). A unifying objective function for topographic mappings. *Neural Computation*, **9**:1291–1303.

Gordon, D., Cohen, R., Kelly, D., Akselrod, S., and Shannon, D. (1984). Sudden infant death syndrome: abnormalities in short term fluctuations in heart rate and respiratory activity. *Pediatric Research*, **18**(10):921–926.

Graepel, T. and Obermayer, K. (1999). A stochastic self-organizing map for proximity data. *Neural Computation*. To appear. Available online from `http://ni.cs.tu-berlin.de/publications/`

Gray Jr, A. and Markel, J. (1976). Distance measures for speech processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, **ASSP-24**(5):380–391.

Guo, Z., Durand, L.-G., and Lee, H. (1994a). Comparison of time-frequency distribution techniques

for analysis of simulated Doppler ultrasound signals of the femoral artery. *IEEE Transactions on Biomedical Engineering*, **41**(4):332–342.

Guo, Z., Durand, L.-G., and Lee, H. (1994b). The time-frequency distributions of nonstationary signals based on a Bessel kernel. *IEEE Transactions on Signal Processing*, **42**(7):1700–1707.

Guzzetti, S., Piccaluga, E., Casati, R., Cerutti, S., Lombardi, F., Pagani, M., and Malliani, A. (1988). Sympathetic predominance in essential hypertension: a study employing spectral analysis of heart rate variability. *Journal of Hypertension*, **6**:711–717.

Haaksma, J., Brouwer, J., van den Berg, M., Dijk, W., Dassen, W., Mulder, L., Mulder, G., and Crijns, H. (1996). Effects of metronome breathing on the assessment of autonomic control using heart rate variability. *Computers in Cardiology 1996*, pages 97–100, Indianapolis, USA. IEEE Computer Society Press, Los Amitos, USA.

Harper, R., Schechtman, V., and Kluge, K. (1987). Machine classification of infant sleep state using cardiorespiratory measures. *Electroencephalography and Clinical Neurophysiology*, **67**(4):379–387.

Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation*. Macmillan Publishing Company, Englewood Cliffs, USA.

Hilton, M., Beattie, J., Chappell, M., and Bates, R. (1997). Heart rate variability: measurement error or chaos. *Computers in Cardiology 1997*, pages 125–128, Lund, Sweden. IEEE Computer Society Press, Los Amitos, USA.

Hilton, M., Bates, R., Godfrey, K., and Cayton, R. (1998). A new application for heart rate variability: diagnosing the sleep apnoea syndrome. *Computers in Cardiology 1998*, pages 1–4, Cleveland, USA. IEEE Computer Society Press, Los Amitos, USA.

Hilton, M., Bates, R., Godfrey, K., and Cayton, R. (1999). Spectal analysis of heart rate variability as a diagnostic marker or the sleep apnoea syndrome. *Medical & Biological Engineering & Computing*. (Submitted).

Hofmann, T. and Buhmann, J. (1995). Multidimensional scaling and data clustering. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems 7*, pages 104–111. MIT Press, Cambridge, USA. Available online from `http://www.ai.mit.edu/people/hofmann/publications.html`

Holmes, J. (1988). *Speech Synthesis and Recognition*. Van Nostrand Reinhold (UK) Co. Ltd., Berkshire, England.

Hon, E. and Lee, S. (1965). Electronic evaluation of the fetal heart rate patterns preceding fetal death: further observations. *American Journal of Obstetrics and Gynecology*, **87**:814–26.

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, **24**:417–441,498–520.

Hoyer, D., Schmidt, K., Bauer, R., Zwiener, U., Köhler, M., Lüthke, B., and Eiselt, M. (1997). Nonlinear analysis of heart rate and respiratory dynamics. *IEEE Engineering in Medicine and Biology Magazine*, **16**(1):31–39.

Hudgel, D. (1986). Clinical manifestations of the sleep apnea syndrome. In Fletcher, E., editor, *Abnormalities of Respiration during Sleep*, pages 21–38. Gratton & Strune, Inc., Orlando, USA.

Hughson, R., Yamamoto, Y., McCullough, R., Sutton, J., and Reeves, J. (1994). Sympathetic and

parasympathetic indicators of heart rate control at altitude studied by spectral analysis. *Journal of Applied Physiology*, **77**(6):3527–2542.

Huikuri, H., Kessler, K., Terracall, E., Castellanos, A., Linnaluoto, M., and Myerburg, R. (1990). Reproducibility and circadian rhythm of heart rate variability in healthy subjects. *American Journal of Cardiology*, **65**:391–393.

Huynen, A., Giesen, R., de la Rosette, J., Aarnink, R., Debruyne, F., and Wijkstra, H. (1994). Analysis of ultrasonographic prostate images for the detection of prostatic carcinoma: the automated urologic diagnostic expert system. *Ultrasound in Medicine & Biology*, **20**(1):1–10.

Hyndman, B., Kitney, R., and Sayers, B. (1971). Spontaneous rhythms in physiological control systems. *Nature*, **233**:339–341.

Hyndman, B. and Mohn, R. (1975). A model of the cardiac pacemaker and its use in decoding the information content of cardiac intervals. *Automedica*, **1**:239–252.

Hyndman, B. and Gregory, J. (1975). Spectral analysis of sinus arrhythmia during mental loading. *Ergonomics*, **18**:255–270.

Hyndman, B., Swenne, C., Bootsma, M., Voogd, J., and Bruschke, A. (1997). Noninvasive measurement of baroreflex sensitivity. A better indicator of cardiac vagal tone than heart rate variability? *Proceedings of the 18th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 4, pages 1582–1583, Amsterdam, The Netherlands. IEEE, New York, USA.

Itakura, F. and Saito, S. (1970). A statistical method for estimation of speech spectral density and formant frequencies. *Electronics and Communications in Japan*, **53-A**:36–43.

Janssen, M., Swenne, C., de Bie, J., Rompelman, O., and van Bemmel, J. (1993). Methods in heart rate variability analysis: which tachogram should we choose? *Computer Methods and Programs in Biomedicine*, **41**:1–8.

Jarrad, G. (1998). Mixture model classifiers. Technical report, Centre for Sensor, Signal and Information Processing, Mawson Lakes, South Australia.

Jeong, J. and Williams, W. (1992). Kernel design for reduced interference distributions. *IEEE Transactions on Signal Processing*, **40**(2):402–412.

Jorna, P. (1993). Heart rate and workload variations in actual and simulated flight. *Ergonomics*, **36**(9):1043–1054.

Joutsiniemi, S.-H., Kaski, S., and Larsen, T. (1995). Self-organizing map in recognition of topographic patterns of EEG spectra. *IEEE Transactions on Biomedical Engineering*, **42**(11):1062–1068.

Jovanov, E., Starcevic, D., Radivojevic, V., Samardzic, A., and Simeunovic, V. (1999). Perceptualization of biomedical data. *IEEE Engineering in Medicine and Biology Magazine*, **18**(1):50–55.

Julius, S. and Schork, M. (1971). Borderline hypertension - a critical review. *Journal of Chronic Diseases*, **23**:723–754.

Julius, S. (1991). Autonomic nervous system dysregulation in human hypertension. *The American Journal of Cardiology*, **67**:3B–7B.

Kalli, S., Grönlund, J., Ihalainen, H., Siimes, A., Välimäki, I., and Antila, K. (1988). Multivariate

autoregressive modelling of autonomic cardiovascular control in neonatal lamb. *Computers and Biomedical Research*, **21**:512–530.

Kallio, K., Haltsonen, S., Paajanen, E., Rosqvist, T., Katila, T., Karp, P., Malmberg, P., Piirila, P., and Sovijarvi, A. (1991). Classification of lung sounds by using self-organizing feature maps. *Proceedings of the 1991 International Conference on Artificial Neural Networks - ICANN '91*, pages 803–808, Espoo, Finland. Elsevier Science Publishing Company, Amsterdam, The Netherlands.

Kamath, M., Fallen, E., and Ghista, D. (1988). Microcomputerized on-line evaluation of heart rate variability power spectra in humans. *Computers in Biology and Medicine*, **18**(3):165–171.

Kamath, M., Fallen, E., and McKelvie, R. (1991). Effects of steady state exercise on the power spectrum of heart rate variability. *Medicine and Science in Sports and Exercise*, **23**(4):428–434.

Kamath, M. and Fallen, E. (1995). Correction of the heart rate variability signal for ectopics and missing beats. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 75–86. Futura Publishing Company, Inc., Armonk, USA.

Kamen, P. and Tonkin, A. (1995). Application of the Poincaré plot to heart rate variability: a new measure of functional status in heart failure. *Australian and New Zealand Journal of Medicine*, **25**:18–26.

Kaplan, D. and Talajic, M. (1991). Dynamics of heart rate. *Chaos*, **1**(3):251–256.

Kaplan, N. (1994). *Clinical Hypertension*. Williams & Watkins, Baltimore, USA, 6th edition.

Kaski, S. (1997). Data exploration using self-organising maps. Thesis for the degree of Doctor of Technology, Helsinki University of Technology, Finland. Available online at `http://www.cis.hut.fi/ sami/thesis/thesis_tohtml.html`

Kautzner, J., Hnatkova, K., Staunton, A., Camm, A., and Malik, M. (1995). Day-to-day reproducibility of time-domain measures of heart rate variability in survivors of acute myocardial infarction. *American Journal of Cardiology*, **76**:309–311.

Kay, S. (1988). *Modern Spectral Estimation: Theory and Application*. Prentice Hall, Englewood Cliffs, USA.

Kennedy, H. (1995). Heart rate variability instruments from commercial manufacturers. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 127–132. Futura Publishing Company, Inc., Armonk, USA.

Keyl, C., Lemberger, P., Pfeifer, M., Hochmuth, K., and Geisler, P. (1997). Heart rate variability in patients with daytime sleepiness suspected of having sleep apnoea syndrome: a receiver-operating characteristic analysis. *Clinical Science*, **92**:335–343.

King, C. and Reisman, S. (1995). The influence of paced breathing exercises on the autonomic response to stress. *Proceedings of the 1995 IEEE Annual Northeast Bioengineering Conference*, pages 124–125, Bar Harbor, USA. IEEE, New York, USA.

Kitney, R. (1975). An analysis of the nonlinear behaviour of the human thermal vasomotor control system. *Journal of Theoretical Biology*, **52**:231–248.

Kitney, R., Fulton, T., McDonald, A., and Linkens, D. (1985). Transient interactions between blood pressure, respiration and heart rate in man. *Journal of Biomedical Engineering*, **7**:217–224.

Kitney, R. and Seydnejad, S. (1996). Heart rate variability system identification using different signal

descriptors. *Computers in Cardiology 1996*, pages 101–104, Indianapolis, USA. IEEE Computer Society Press, Los Amitos, USA.

Kleiger, R., Miller, J., Bigger Jr, J., Moss, A., and The Multicenter Post-Infarction Research Group (1987). Decreased heart rate variability and its association with increased mortality after acute myocardial infarction. *American Journal of Cardiology*, **59**:256–262.

Kleiger, R., Stein, P., Bosner, M., and Rottman, J. (1995). Time-domain measurements of heart rate variability. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 33–46. Futura Publishing Company, Inc., Armonk, USA.

Klock, H. and Buhmann, J. (1997). Multidimensional scaling by deterministic annealing. *Proceedings of the International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 246–260, Venice, Italy. Springer-Verlag, Berlin. Available online from `http://www-dbv.informatik.uni-bonn.de/~joerg/`

Kobayashi, M. and Musha, T. (1982). 1/f fluctuation of heartbeat period. *IEEE Transactions on Biomedical Engineering*, **BME-29**(6):456–457.

Koeleman, A., Ros, H., and van den Akker, T. (1985). Beat-to-beat interval measurement in the electrocardiogram. *Medical & Biological Engineering & Computing*, **23**:213–219.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, **43**:59–69.

Kohonen, T. (1988). The 'neural' phonetic typewriter. *Computer*, **21**(3):11–22.

Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, **78**:1464–1480.

Kohonen, T. (1997). *Self-Organizing Maps*. Springer-Verlag, Berlin, 2nd edition.

Koontz, W. and Fukunaga, K. (1972). A nonlinear feature extraction algorithm using distance transformation. *IEEE Transactions on Computers*, **C-21**(1):56–63.

Korhonen, I., Mainardi, L., Baselli, G., Bianchi, A., Loula, P., and Carrault, G. (1996a). Linear multivariate models for physiological signal analysis: applications. *Computer Methods and Programs in Biomedicine*, **51**:121–130.

Korhonen, I., Mainardi, L., Loula, P., Carrault, G., Baselli, G., and Bianchi, A. (1996b). Linear multivariate models for physiological signal analysis: theory. *Computer Methods and Programs in Biomedicine*, **51**:85–94.

Korhonen, I., Takalo, R., and Turjanmaa, V. (1996c). Multivariate autoregressive model with immediate transfer paths for assessment of interactions between cardiopulmonary variability signals. *Medical & Biological Engineering & Computing*, **34**:199–206.

Kraaijveld, M., Mao, J., and Jain, A. (1995). A nonlinear projection method based on Kohonen's topology preserving maps. *IEEE Transactions on Neural Networks*, **6**:548–559.

Kruskal, J. (1964a). Multidimensional scaling by optimizing goodness-of-fit to a nonmetric hypothesis. *Psychometrika*, **29**(1):1–27.

Kruskal, J. (1964b). Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, **29**(2):115–129.

Lacoss, R. (1971). Data adaptive spectral analysis methods. *Geophysics*, **36**(4):661–675.

Laguna, P., Caminal, P., Jane, R., and Rix, H. (1991). Evaluation of HRV by PP and RR interval analysis using a new time delay estimate. *Computers in Cardiology 1991*, pages 63–66, Venice, Italy. IEEE Computer Society Press, Los Amitos, USA.

Laguna, P., Moody, G., and Mark, R. (1995). Power spectral density of unevenly sampled heart rate data. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 157–158, Montreal, Canada. IEEE, New York, USA.

Laguna, P., Moody, G., and Mark, R. (1998). Power spectral density of unevenly sampled data by least-square analysis: performance and application to heart rate signals. *IEEE Transactions on Biomedical Engineering*, **45**(6):698–715.

Lee, M. (1998). Interactive visualisation of similarity structures. *Proceedings of the 1998 Australasian Computer Human Interaction Conference*, pages 292–299, Adelaide, Australia. IEEE Computer Society Press, Los Amitos, USA.

Lee, M. (1999). Consequential region scaling. *Submitted for publication.*

Lee, Y., Lee, K., and Yoon, H. (1989). Discriminant analysis of HRV's autoregressive model coefficient for cardiac arrhythmia. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 1, pages 47–48, Seattle, USA. IEEE, New York, USA.

Lehtinen, J.-C., Forsström, J., Koskinen, P., Penttilä, T.-A., Järvi, T., and Anttila, L. (1997). Visualization of clinical data with neural networks, case study: polycystic ovary syndrome. *International Journal of Medical Informatics*, **44**:145–155.

Leistner, H., Haddad, G., Epstein, R., Lai, T., Epstein, M., and Mellins, R. (1980). Heart rate and heart rate variability during sleep in aborted sudden infant death syndrome. *Journal of Pediatrics*, **97**(1):51–55.

Lerner, B., Guterman, H., Aladjem, M., Dinstein, I., and Romem, Y. (1996). Feature extraction by neural network nonlinear mapping for pattern classification. *Proceedings of 13th International Conference on Pattern Recognition*, pages 320–324, Vienna, Austria. IEEE Computer Society Press, Los Amitos, USA.

Lerner, B., Guterman, H., Aladjem, M., Dinstein, I., and Romem, Y. (1998). On pattern classification with Sammon's nonlinear mapping - an experimental study. *Pattern Recognition*, **31**(4):371–381. Available online from `http://www.ncrg.aston.ac.uk/`

Leuenberger, U., Jacob, E., Sweer, L., Waravdekar, N., Zwillich, C., and Sinoway, L. (1995). Surges of muscle sympathetic nerve activity during obstructive apnea are linked to hypoxemia. *Journal of Applied Physiology*, **79**(2):581–588.

Liao, D., Barnes, R., Chambless, L., and Heiss, G. (1996). A computer algorithm to impute interrupted heart rate data for the spectral analysis of heart rate - the ARIC study. *Computers and Biomedical Research*, **29**:140–151.

Lipsitz, L., Mietus, J., Moody, G., and Goldberger, A. (1990). Spectral characteristics of heart rate variability before and during postural tilt. *Circulation*, **81**(6):1803–1810.

Lishner, M., Akselrod, S., Moore, A., Oz, O., Divon, M., and Ravid, M. (1987). Spectral analysis of heart rate fluctuations: a noninvasive, sensitive method for the early diagnosis of autonomic

neuropathy in diabetes mellitus. *Journal of the Autonomic Nervous System*, **19**:143–153.

Little, R. (1985). *Physiology of the heart and circulation*. Year Book Medical Publishers Inc., Chicago, USA, 3rd edition.

Lomb, N. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophysical and Space Science*, **39**:447–462.

Lowe, D. (1993). Novel 'topographic' nonlinear feature extraction using radial basis functions for concentration coding in the 'artifical nose'. *3rd IEE International Conference on Artifical Neural Networks*, Brighton, UK. IEE, London.

Lowe, D. and Tipping, M. (1995). A novel neural network technique for exploratory data analysis. *Proceedings of the International Conference on Artificial Neural Networks*, volume 1, pages 339–344, Paris, France. EC2 & Cie, Paris.

Lowe, D. (1997). Extracting structure from wake EEG using neural networks. *Proceedings of the SPIE*, volume 3077, pages 17–26, Orlando, Florida. SPIE, Bellingham, USA.

Luczak, H. and Laurig, W. (1973). An analysis of heart rate variability. *Ergonomics*, **16**:85–97.

Madwed, J., Albrecht, P., Mark, R., and Cohen, R. (1989). Low-frequency oscillations in arterial pressure and heart rate: a simple computer model. *American Journal of Physiology*, **256**:H1573–H1579.

Maes, L., Bijnens, B., Suetens, P., and van de Werf, F. (1993). Automated contour detection of the left ventricle in short axis view in 2D echocardiograms. *Machine Vision and Applications*, **6**(1):1–9.

Mainardi, L., Bianchi, A., and Cerutti, S. (1994). On-line beat-to-beat monitoring of spectral parameters of heart rate variability using a pole-tracking algorithm. *Methods of Information in Medicine*, **33**:85–88.

Malik, M., Farrell, T., Cripps, T., and Camm, A. (1989). Heart rate variability in relation to prognosis after myocardial infarction-selection of optimal processing techniques. *European Heart Journal*, **10**:1060–1074.

Malik, M. and Camm, A. (1993). Components of heart rate variability - what they really mean and what we really measure. *American Journal of Cardiology*, **72**:821–822.

Malik, M., Xia, R., Odemuyiwa, O., Staunton, A., Poloniecki, J., and Camm, A. (1993). Influence of the recognition artefact in the automatic analysis of long-term electrocardiograms on time-domain measurement of heart rate variability. *Medical & Biological Engineering & Computing*, **31**:539–544.

Malik, M. and Camm, A. (1994). Heart rate variability and clinical cardiology. *British Heart Journal*, **71**:3–6.

Malik, M. (1995). Geometric methods for heart rate variability assessment. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 47–62. Futura Publishing Company, Inc., Armonk, USA.

Malliani, A., Pagani, M., Lombardi, F., and Cerutti, S. (1991a). Cardiovascular neural regulation explored in the frequency domain. *Circulation*, **84**(2):482–492.

Malliani, A., Pagani, M., Lombardi, F., Furlan, R., Guzzetti, S., and Cerutti, S. (1991b). Spectral analysis to assess increased sympathetic tone in arterial hypertension. *Hypertension, Supplement III*, **17**(4):36–42.

Malliani, A., Lombardi, F., and Pagani, M. (1994a). Power spectrum analysis of heart rate variability: a tool to explore neural regulatory mechanisms. *British Heart Journal*, **71**:1–2.

Malliani, A., Pagani, M., and Lombardi, F. (1994b). Importance of appropriate spectral methodology to assess heart rate variability in the frequency domain. *Hypertension*, **24**(1):140–141.

Maltson, R. and Dammann, J. (1965). A technique for determining and coding subclasses in pattern recognition problems. *IBM Journal*, **9**:294–302.

Mammoliti, R., Landini, L., Pola, S., Raciti, M., Santarcangelo, E., and Emdin, M. (1998). Recurrence quantification analysis describes the complex and deterministic behavior of heart rate variability in healthy subjects. *Computers in Cardiology 1998*, pages 145–148, Cleveland, USA. IEEE, New York, USA.

Mao, J. and Jain, A. (1995). Artificial neural networks for feature extraction and multivariate data projection. *IEEE Transactions on Neural Networks*, **6**(2):296–317.

Marciano, F., Bonaduce, D., Petretta, M., Valva, G., and Migaux, M. (1995). Spectral behaviour of heart rate variability in acute ischemic episodes. *Computers in Cardiology 1995*, pages 111–114, Vienna, Austria. IEEE Computer Society Press, Los Amitos, USA.

Mardia, K., Kent, J., and Bibby, J. (1979). *Multivariate Analysis*. Academic Press, London, UK.

Marple Jr, S. (1976). *Conventional Fourier, Autoregressive and Special ARMA Methods of Spectrum Analysis*. PhD thesis, Stanford University, USA. Cited in Kay, S.M. (1988).

McAllister Jr, R. (1979). Effect of adrenergic receptor blockade on the responses to isometric handgrip: studies in normal and hypertensive subjects. *Journal of Cardiovascular Pharmacology*, **1**:253–263.

McCaffree, M. (1986). Sudden infant death syndromes: respiratory mechanisms. In Fletcher, E., editor, *Abnormalities of Respiration during Sleep*, pages 229–240. Grune & Stratton, Inc., Orlando, USA.

McDonald, A. (1980). Mechanisms affecting heart-rate. In Kitney, R. and Rompelman, O., editors, *The Study of Heart-Rate Variability*, pages 3–12. Oxford University Press, New York, USA.

McLachlan, G. and Basford, K. (1988). *Mixture models: inference and applications to clustering*. Marcel Dekker, Inc., New York, USA.

McMahon, S., Peto, R., Cutler, J., Collins, R., Sorlie, P., Neaton, J., Abbott, R., Godwin, J., Dyer, A., and Stamler, J. (1990). Blood pressure, stroke, and coronary heart disease. Part I. Prolonged differences in blood pressure: prospective observational studies corrected for the regression dilution bias. *Lancet*, **335**:765–774.

Miikkulainen, R. (1990). Script recognition with hierarchical feature maps. *Connection Science*, **2**:83–101.

Moody, J. and Darken, C. (1989). Fast learning in networks of locally-tuned processing units. *Neural Computation*, **1**(2):281–294.

Moser, M., Lehofer, M., Sedminek, A., Lux, M., Zapotoczky, H.-G., Kenner, T., and Noordergraaf, A. (1994). Heart rate variability as a prognostic tool in cardiology. *Circulation*, **90**:1078–1082.

Murakawa, Y., Ajiki, K., Usui, M., Yamashita, T., Oikawa, N., and Inoue, H. (1993). Parasympathetic activity is a major modulator of the circadian variability of heart rate in healthy subjects and in

patients with coronary artery disease or diabetes mellitus. *American Heart Journal*, **126**(1):108–114.

Myrtek, M., Deutschmann, J., Strohmaier, H., Zimmerman, W., Lawerenz, S., Brugner, G., and Muller, W. (1994). Physical, mental, emotional and subjective workload components in train drivers. *Ergonomics*, **37**(7):1195–1203.

Nakao, M., Norimatsu, M., Mizutani, Y., and Yamamoto, M. (1997). Spectral distortion properties of the integral pulse frequency modulation model. *IEEE Transactions on Biomedical Engineering*, **44**(5):419–426.

Nandagopal, D., Fallen, E., Ghista, D., and Connally, S. (1985). Reproducibility of resting HRV spectrum and its changes following physiological perturbations. *Automedica*, **5**:1–13.

Navabi, M., Watt, R., Hameroff, S., and Mylrea, K. (1991). Integrated monitoring can detect critical events and improve alarm accuracy. *Journal of Clinical Engineering*, **16**(4):295–306.

Neubauer, B. and Gundersen, J. (1978). Analysis of heart rate variations in patients with multiple sclerosis. A simple measure of autonomic nervous disturbances using an ordinary ECG. *Journal of Neurology, Neurosurgery and Psychiatry*, **41**:417–419.

Niccolai, M., Emdin, M., Carpeggiani, C., Raciti, M., Maielli, M., Bertinelli, M., Varanini, M., Macerata, A., and Marchesi, C. (1994). Multichannel spectral analysis of R-R and respiratory series from 24-hour Holter recordings. *Computers in Cardiology 1994*, pages 757–760, Bethesda, USA. IEEE Computer Society Press, Los Amitos, USA.

Nickel, R., Sang, T.-H., and Williams, W. (1998). A new signal adaptive approach to positive time-frequency distributions with suppressed interference terms. *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 3, pages 1777–1780, Seattle, USA. IEEE, New York, USA.

Novak, P. and Novak, V. (1993). Time-frequency mapping of the heart rate, blood pressure and respiratory signals. *Medical & Biological Engineering & Computing*, **31**:103–110.

Nunan, T., Tonkin, A., and Taverner, D. (1994). Effects of antihypertensive drugs on cardiovascular autonomic reflexes. *Proceedings of ASCEPT*, volume 1, page 92.

O'Brien, I., O'Hare, P., and Corrall, R. (1986). Heart rate variability in healthy subjects: effect of age and the derivation of normal ranges for tests of autonomic function. *British Heart Journal*, **55**:348–354.

Ornes, C. and Sklansky, J. (1997). A visual multi-expert neural classifier. *Proceedings of the 1997 IEEE International Conference on Neural Networks*, pages 1448–1453, Houston, USA. IEEE, New York, USA.

Orr, M. (1995). Regularization in the selection of radial basis function centers. *Neural Computation*, **7**:606–623.

Orr, M. (1997). MATLAB Routines for Subset Selection and Ridge Regression in Linear Neural Networks. Technical report, Centre for Cognitive Science, Edinburgh University. Available online from `http://www.cns.ed.ac.uk/`

Pagani, M., Lombardi, F., Guzzetti, S., Sandrone, G., Rimoldi, O., Malfatto, G., Cerutti, S., and Malliani, A. (1984). Power spectral density of heart rate variability as an index of sympatho-vagal interaction in normal and hypertensive subjects. *Journal of Hypertension (Supplement)*, **2**(3):S383–S385.

Pagani, M., Lombardi, F., Guzzetti, S., Rimoldi, O., Furlan, R., Pizzinelli, P., Sandrone, G., Malfatto, G., Dell'Orto, S., Piccaluga, E., Turiel, M., Baselli, G., Cerutti, S., and Malliani, A. (1986). Power spectral analysis of heart rate and arterial pressure variabilities as a marker of sympatho-vagal interaction in man and conscious dog. *Circulation Research*, **59**:178–193.

Pagani, M., Lucini, D., Rimoldi, O., Furlan, R., Piazza, S., and Biancardi, L. (1995). Effects of physical and mental exercise on heart rate variability. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 245–266. Futura Publishing Company, Inc., Armonk, USA.

Parkes, J. (1985). *Sleep and its disorders*. W.B. Saunders Company, London, UK.

Pearlman, J., Chuang, M., and Dewagan, N. (1995). Maps for visualization and analysis of perfusion. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 487–488, Montreal, Canada. IEEE, New York, USA.

Pickering, T. (1993). Blood pressure variability and ambulatory monitoring. *Current Opinion in Nephrology and Hypertension*, **2**:380–385.

Pinna, G., Maestri, R., and Di Cesare, A. (1996). Application of time series spectral analysis theory: analysis of cardiovascular variability signals. *Medical & Biological Engineering & Computing*, **34**:142–148.

Pola, S., Macerata, A., Emdin, M., and Marchesi, C. (1996). Estimation of the power spectral density in nonstationary cardiovascular time series: assessing the role of the time-frequency representations (TFR). *IEEE Transactions on Biomedical Engineering*, **43**(1):46–59.

Pomeranz, B., Macaulay, R., Caudill, M., Kutz, I., Adam, D., Gordon, D., Kilborn, K., Barger, A., Shannon, D., Cohen, R., and Benson, H. (1985). Assessment of autonomic function in humans by heart rate spectral analysis. *American Journal of Physiology*, **248**:H151–H153.

Pompe, B., Blidh, P., Hoyer, D., and Eiselt, M. (1998). Using mutual information to measure coupling in the cardiorespiratory system. *IEEE Engineering in Medicine and Biology Magazine*, **17**(6):32–39.

Preiss, G. and Polosa, C. (1974). Patterns of sympathetic neuron activity associated with Mayer waves. *American Journal of Physiology*, **226**(3):724–730.

Pykett, C. (1978). Improving the efficiency of Sammon's nonlinear mapping by using clustering archetypes. *Electronics Letters*, **14**(25):799–800.

Rakugi, H., Yu, H., Kamitani, A., Nakamura, Y., Ohishi, M., Kamide, K., Nakata, Y., Takami, S., Higaki, J., and Ogihara, T. (1996). Links between hypertension and myocardial infarction. *American Heart Journal*, **132**:213–221.

Raymond, B., Mazumdar, J., and Nandagopal, D. (1997). Modelling the shift of the low frequency component of heart rate variability. *Computers in Cardiology 1997*, pages 407–410, Lund, Sweden. IEEE Computer Society Press, Los Amitos, USA.

Ritter, H. and Schulten, K. (1988). Kohonen's self-organizing maps: exploring their computational capabilities. *Proceedings of the IEEE International Conference on Neural Networks*, volume I, pages 109–116, San Diego, USA. IEEE, New York, USA.

Ritter, H. (1991). Asymptotic level density for a class of vector quantization processes. *IEEE Transactions on Neural Networks*, **2**:173–175.

Rompelman, O., Snijders, J., and van Spronsen, C. (1982). The measurement of heart rate variability

spectra with the help of a personal computer. *IEEE Transactions on Biomedical Engineering*, **BME-29**(7):503–510.

Rosario, B., Lovell, D., Niranjan, M., Prager, R., Dalton, K., and Derom, R. (1997). Self-organization with a large medical database: using GTM for prediction and clustering. *Proceedings of the 9th Italian Workshop on Neural Nets*, pages 257–262, Salerno, Italy. Springer-Verlag, London, UK.

Rosenblum, M. and Kurths, J. (1995). A model of neural control of the heart rate. *Physica A*, **215**:439–450.

Ryan, P., Hilton, M., Boldy, D., Evans, A., Bradbury, S., Sapiano, S., Prowse, K., and Cayton, R. (1995). Validation of British Thoracic Society guidelines for the diagnosis of the sleep apnoea/ hypopnoea syndrome: can polysomnography be avoided? *Thorax*, **50**:972–975.

Samad, T. and Harp, S. (1992). Self-organization with partial data. *Network: Computation in Neural Systems*, **3**:205–212.

Sammon, J. (1969). A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, **18**:401–409.

Sang, T.-H., Williams, W., and O'Neill, J. (1996). An algorithm for positive time-frequency distributions. *Proceedings of the IEEE Symposium on Time-Frequency and Time-Scale Analysis*, pages 165–168, Paris, France. IEEE, New York, USA.

Sapoznikov, D., Luria, M., Mahler, Y., and Gotsman, M. (1992a). Computer processing of artifact and arrhythmias in heart rate variability analysis. *Computer Methods and Programs in Biomedicine*, **39**:75–84.

Sapoznikov, D., Talajic, M., and Kaplan, D. (1992b). Multiparameter methodologies of heart rate variability analysis. *Computer Methods and Programs in Biomedicine*, **41**(2):69–75.

Sato, S., Doi, S., and Nomura, T. (1994). Bonhoeffer-van der Pol oscillator of the sino-atrial node: a possible mechanism of heart rate regulation. *Methods of Information in Medicine*, **33**:116–119.

Sauerbrei, W., Madjar, H., and Prompeler, H. (1998). Differentiation of benign and malignant breast tumors by logistic regression and a classification tree using Doppler flow signals. *Methods of Information in Medicine*, **37**(3):226–234.

Saul, J., Arai, Y., Berger, R., Lily, L., Wilson, S., Colucci, W., and Cohen, R. (1988a). Assessment of autonomic regulation in chronic congestive heart failure by the heart rate spectral analysis. *American Journal of Cardiology*, **61**:1292–1299.

Saul, J., Kaplan, D., and Kitney, R. (1988b). Nonlinear interactions between respiration and heart rate: classical physiology or entrained nonlinear oscillators. *Computers in Cardiology 1988*, pages 299–302, Washington DC, USA. IEEE Computer Society Press, Los Amitos, USA.

Saul, J., Berger, R., and Cohen, R. (1989). A simple analytical model mimics complex physiological behaviour. *Computers in Cardiology 1989*, pages 335–338, Jerusalem, Israel. IEEE Computer Society Press, Los Amitos, USA.

Sava, H., Pibarot, P., and Durand, L.-G. (1998). Application of the matching pursuit method for structural decomposition and averaging of phonocardiographic signals. *Medical & Biological Engineering & Computing*, **36**:302–308.

Sayers, B. M. (1973). Analysis of heart rate variability. *Ergonomics*, **16**(1):17–32.

Schmidt, G. and Morfill, G. (1995). Nonlinear methods for heart rate variability assessment. In Malik, M. and Camm, A., editors, *Heart Rate Variability*, pages 87–98. Futura Publishing Company, Inc., Armonk, USA.

Schnorrenberg, F., Pattichis, C., Kyriacou, K., and Schizas, C. (1994). Detection of cell nuclei in breast biopsies using receptive fields. *Proceedings of the 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 649–650, Baltimore, USA. IEEE, New York, USA.

Schwartz, J., Gibb, W., and Tran, T. (1991). Aging effects on heart rate variation. *Journal of Gerontology: Medical Sciences*, **46**(3):M99–M106.

Scott, E. (1986). *Cardiovascular physiology: an integrative approach*. Manchester University Press, Manchester, UK.

Selman, A., McDonald, A., Kitney, R., et al. (1982). The interaction between heart rate and respiration. I. Experimental studies in man. *Automedica*, **4**:131–139.

Severi, S., Cavalcanti, S., and Avanzolini, G. (1997). Heart rate variability spectral indices for haemodynamic classification of haemodialysis patients. *Physiological Measurement*, **18**(4):339–353.

Sherrah, J. (1998). *Automatic feature extraction for pattern recognition*. PhD thesis, The University of Adelaide, Australia.

Siedlecki, W., Siedlecka, K., and Sklansky, J. (1988a). Experiments on mapping techniques for exploratory pattern analysis. *Pattern Recognition*, **21**(5):431–438.

Siedlecki, W., Siedlecka, K., and Sklansky, J. (1988b). An overview of mapping techniques for exploratory pattern analysis. *Pattern Recognition*, **21**(5):411–429.

Signorini, M., Cerutti, S., Guzzetti, S., and Parola, R. (1994). Non-linear dynamics of cardiovascular variability signals. *Methods of Information in Medicine*, **33**(1):81–84.

Signorini, M., Lombardi, F., and Cerutti, S. (1998). Nonlinearity parameters for the classification of high risk myocardial infarction subjects. *Computers in Cardiology 1998*, pages 545–548, Cleveland, USA. IEEE, New York, USA.

Simpson, D. and Wicks, R. (1988). Spectral analysis of heart rate indicates reduced baroreceptor-related heart rate variability in elderly persons. *Journal of Gerontology*, **43**:M21–M24.

Singh, J., Dunt, D., and Nandagopal, D. (1993). Stress response studies in meditators using heart rate variability. *Automedica*, **15**:273–279.

Sisli, G., Akin, A., and Onaral, B. (1996). Multi-scale analysis of HRV data. *Proceedings of the IEEE Symposium on Time-Frequency and Time-Scale Analysis*, pages 101–104, Paris, France. IEEE, New York, USA.

Spaeth, H. and Guthery, S. (1969). The use and utility of the monotone criterion in multidimensinal scaling. *Multivariate Behavioural Research*, **4**:501–515.

Stevens, V., Wilson, A., Franks, C., and Southall, D. (1988). Techniques for the analysis of long-term cardiorespiratory recordings from infants. *Medical & Biological Engineering & Computing*, **26**(3):282–288.

Strang, G. and Nguyen, T. (1996). *Wavelets and Filter Banks*. Wellesey-Cambridge Press, Wellesey, USA.

Strohmer, T. (1991). Irregular sampling, frames and pseudo-inverse. Master's thesis, University of Vienna, Austria. Available online from `http://tyche.mat.univie.ac.at/`

Sugimoto, K., Kumagai, Y., Tateishi, T., Seguchi, H., and Ebihara, A. (1989). Effects on autonomic function of a new angiotensin converting enzyme inhibitor, Ramipril. *Journal of Cardiovascular Pharmacology*, **13**(Supplement 3):S40–S44.

Takalo, R., Korhonen, I., Turjanmaa, V., Majahalme, S., Tuomisto, M., and Uusitalo, A. (1994). Short-term variability of blood pressure and heart rate in borderline and mildly hypertensive subjects. *Hypertension*, **23**:18–24.

Task Force (1996). Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology. Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. *Circulation*, **93**:1043–1065.

Tattersall, G. and Limb, P. (1994). Visualisation techniques for data mining. *BT Technology Journal, British Telecom Labs, Ipswich I57RE, UK*, **12**(4):23–31.

Tazebay, M., Saliba, R., and Reisman, S. (1995). Adaptive time-frequency analysis of autonomic nervous system. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 1073–1074, Montreal, Canada. IEEE, New York, USA.

Tilkian, A., Motta, J., and Guilleminault, C. (1978). Cardiac arrhythmias in sleep apnea. In Guilleminault, C. and Dement, W., editors, *Sleep apnea syndromes*, pages 197–210. Alan R Liss.

Tipping, M. (1996). *Topographic mappings and feed-forward neural networks*. PhD thesis, The University of Aston in Birmingham, UK. Available online from `http://www.ncrg.aston.ac.uk/`

Tipping, M. and Lowe, D. (1997). Shadow targets: a novel algorithm for topographic projections by radial basis functions. *Fifth International Conference on Artificial Neural Networks*, pages 7–12, Cambridge, UK. IEE, London, UK. Available online from `http://www.ncrg.aston.ac.uk/`

Toledo, E., Gurevich, O., Hod, H., Eldar, M., and Akselrod, S. (1998). The use of a wavelet transform for the analysis of nonstationary heart rate variability signal during thrombolytic therapy as a marker of reperfusion. *Computers in Cardiology 1998*, pages 609–612, Cleveland, USA. IEEE Computer Society Press, Los Amitos, USA.

Torgerson, W. (1958). *Theory and Methods of Scaling*. John Wiley & Sons, New York, USA.

Travaglini, A., Lamberti, C., DeBie, J., and Ferri, M. (1998). Respiratory signal derived from eight-lead ECG. *Computers in Cardiology 1998*, pages 65–68, Cleveland, USA. IEEE Computer Society Press, Los Amitos, USA.

Tsai, C., Wang, J.-H., and Sun, Y.-N. (1998). Liver stereoscopic visualisation by using Hopfield neural nets. *Neural Computing & Applications*, **7**(3):229–237.

Tsapatsoulis, N., Schnorrenberg, F., Pattichis, C., Kollias, S., and Schizas, C. (1997). A neural network assisted algorithm for automated detection of breast cancer nuclei. *1997 13th International Conference on Digital Signal Processing Proceedings. DSP 97*, pages 269–272, Santorini, Greece. IEEE, New York, USA.

Tsuji, H. and Mori, H. (1994). New analysis of HRV through wavelet transform. *International Journal of Human Computer Interaction*, **6**(2):205–217.

Tuininga, Y., van Veldhuisen, D., Brouwer, J., Haaksma, J., Crijns, H., Man in 't Veld, A., and Lie, K.

(1994). Heart rate variability in left ventricular dysfunction and heart failure: effects and implications of drug treatment. *British Heart Journal*, **72**:509–513.

Ungi, I., Thury, A., and Csanady, M. (1995). Investigation of the correlation between heart rate and heart rate variability. *Computers in Cardiology 1995*, pages 189–191, Vienna, Austria. IEEE Computer Society Press, Los Amitos, USA.

Vainamo, K., Nissila, S., Makikallio, T., Tulppo, M., and Ronig, J. (1996). Artificial neural networks for aerobic fitness approximation. *The 1996 IEEE International Conference on Neural Networks*, volume 4, pages 1939–1944, Washington DC, USA. IEEE, New York, USA.

Vallverdú, M., Korenberg, M., and Caminal, P. (1991). Model identification of the neural control of the cardiovascular system using NARMAX models. *Computers in Cardiology 1991*, pages 585–588, Venice, Italy. IEEE Computer Society Press, Los Amitos, USA.

van den Akker, T., Koeleman, A., Hogenhuis, L., and Rompelman, O. (1983). Heart rate variability and blood pressure oscillations in diabetics with autonomic neuropathy. *Automedica*, **4**:201–208.

van Gils, M., Jansen, H., Nieminen, K., Summers, R., and Weller, P. (1997). Using artificial neural networks for classifying ICU patient states. *IEEE Engineering in Medicine and Biology Magazine*, **16**(6):41–47.

van Hoogenhuyze, D., Weinstein, N., Martin, G., Weiss, J., Schaad, J., Sahyouni, X., Fintel, D., Remme, W., and Singer, D. (1991). Reproducibility and relation to mean heart rate of heart rate variability in normal subjects and in patients with congestive heart failure secondary to coronary artery disease. *American Journal of Cardiology*, **68**:1668–1676.

van Steenis, H. and Tulen, J. (1996). The exponential distribution applied to nonequidistantly sampled cardiovascular time series. *Computers and Biomedical Research*, **29**:174–193.

van Wezel, M., Kok, J., and Kosters, W. (1997). Two neural network methods for multidimensional scaling. *Proceedings of the European Symposium on Artificial Neural Networks (ESANN'97)*, pages 97–102, Bruges, Belgium. D Facto, Brussels, Belgium. Available online from `http://www.cwi.nl/~michiel/index.html`

Varanini, M., Macerata, A., Emdin, M., and Marchesi, C. (1994). Nonlinear filtering for the estimation of the respiratory component in heart rate. *Computers in Cardiology 1994*, pages 565–568, Bethesda, USA. IEEE Computer Society Press, Los Amitos, USA.

Varanini, M., Pola, S., Emdin, M., Macerata, A., Cipriani, M., and Marchesi, C. (1996). Effectiveness of adaptive filtering in cancelling the respiratory component from cardiovascular series. *Computers in Cardiology 1996*, pages 517–520, Indianapolis, USA. IEEE Computer Society Press, Los Amitos, USA.

Venturi, M., Conforti, F., Macerata, A., Varanini, M., Emdin, M., and Marchesi, C. (1990). Analysis of variability: a system for comparing classical, parametric, adaptive and Wigner-Ville power spectral estimators. *Computers in Cardiology 1990*, pages 247–250, Chicago, USA. IEEE Computer Society Press, Los Amitos, USA.

Vespasiani, G., Bruni, M., Meloncelli, I., Clementi, L., Amoretti, R., Branca, S., Carinci, F., Lostia, S., Nicolucci, A., Romagnoli, F., Verga, S., and Bendetti, M. (1996). Validation of a computerised measurement system for guided routine evaluation of cardiovascular autonomic neuropathy. *Computer Methods and Programs in Biomedicine*, **51**:211–216.

Vibe, K. and Vesin, J. (1996). On chaos detection methods. *International Journal of Bifurcation and*

Chaos in Applied Sciences and Engineering, **6**(3):529–543.

Vila, J., Palacios, F., Presedo, J., Fernandez-Delgado, M., Felix, P., and Barro, S. (1997). Time-frequency analysis of heart-rate variability. *IEEE Engineering in Medicine and Biology Magazine*, **16**(5):119–126.

Vogel, S. (1992). *Vital circuits*. Oxford University Press, New York, USA.

Voss, A., Wessel, N., Kleiner, H., Osterziel, K., and Dietz, R. (1996). Characterization of cardiovascular diseases by measures of non-linear dynamics. *Medical & Biological Engineering & Computing*, **34**(Supplement 1):371–372.

Vybiral, T., Bryg, R., and Maddens, M. (1990). Impact of arrhythmias on heart rate variability-strategies to deal with imperfect clinical data. *Computers in Cardiology 1990*, pages 251–254, Chicago, USA. IEEE Computer Society Press, Los Amitos, USA.

Warner, H. and Cox, A. (1962). A mathematical model of heart rate control by sympathetic and vagus efferent information. *Journal of Applied Physiology*, **17**:349–355.

Watson, M., Nandagopal, D., and Singh, J. (1993). Modelling of heart rate variability under neuropsychic load. *Automedica*, **15**:259–271.

Watson, M. (1994). *Detection of temporal lobe seizures by mathematical modelling of heart rate variability.* PhD thesis, The University of Melbourne, Australia.

Webb, A. (1995). Multidimensional scaling by iterative majorization using radial basis functions. *Pattern Recognition*, **28**(5):753–759.

Weise, F., Baltrusch, K., and Heydenreich, F. (1989). Effect of low-dose atropine on heart rate fluctuations during orthostatic load: a spectral analysis. *Journal of the Autonomic Nervous System*, **26**:223–230.

Weise, F. and Heydenreich, F. (1989). Effects of modified respiratory rhythm on heart rate variability during active orthostatic load. *Biomedica Biochimica Acta*, **48**(8):549–556.

Wesseling, K., Settels, J., Walstra, H., van Esch, H., and Donders, J. (1982). Baromodulation as the cause of short term blood pressure variability. In Alberi, G., Bajzev, Z., and Baxa, P., editors, *Proceedings of the International Conference of Applications of Physics to Medicine and Biology*, pages 247–276, Trieste, Italy. World Scientific, Singapore. Cited in van den Akker *et al.*, 1983.

Widrow, B., Glover Jr, J., McCool, J., Kaunitz, J., Williams, C., Hearn, R., Zeidler, J., Dong Jr, E., and Goodlin, R. (1975). Adaptive noise cancelling: principles and applications. *Proceedings of the IEEE*, **63**(12):1692–1716.

Wiklund, U., Niklasson, U., and Bjerle, P. (1991). Adaptive cancellation of respiratory sinus arrhythmia. *Computers in Cardiology 1991*, pages 259–262, Venice, Italy. IEEE Computer Society Press, Los Amitos, USA.

Wiklund, U. (1996). Analysis of heart rate variability in patients with severe autonomic dysfunction. *Medical & Biological Engineering & Computing*, **34**(Supplement 1):383–384.

Wiklund, U., Akay, M., and Niklasson, U. (1997). Short-term analysis of heart-rate variability by adapted wavelet transforms. *IEEE Engineering in Medicine and Biology Magazine*, **16**(5):113–118.

Williams, W. and Jeong, J. (1992). Reduced interference time-frequency distributions. In Boashash,

B., editor, *Time frequency signal analysis*, chapter 3. Longman and Cheshire/Wiley, New York.

Williams, W. (1996). Reduced interference distributions: biological applications and interpretations. *Proceedings of the IEEE*, **84**(9):1264–1280.

Wolf, M., Varigos, G., and Hunt, D. (1978). Sinus arrhythmia in acute myocardial infarction. *Medical Journal of Australia*, **2**:52.

Womack, B. (1971). The analysis of respiratory sinus arrhythmia using spectral analysis and digital filtering. *IEEE Transactions on Biomedical Engineering*, **BME-18**(6):399–409.

Woo, M., Stevenson, W., Moser, D., Trelease, R., and Harper, R. (1992). Patterns of beat-to-beat heart rate variability in advanced heart failure. *American Heart Journal*, **123**:704–710.

Wood, J., Buda, A., and Barry, D. (1992). Time-frequency transforms: a new approach to first heart sound frequency dynamics. *IEEE Transactions on Biomedical Engineering*, **39**(7):730–740.

Yamamoto, Y. and Hughson, R. (1991). Coarse-graining spectral analysis: new method for studying heart rate variability. *Journal of Applied Physiology*, **71**(3):1143–1150.

Yan, X. and Zheng, C. (1995). Frequency-domain techniques for heart rate variability analysis. *1995 IEEE Engineering in Medicine and Biology 17th Annual Conference*, pages 961–962, Montreal, Canada. IEEE, New York, USA.

Yang, F. and Liao, W. (1997). Modeling and decomposition of HRV signals with wavelet transform. *IEEE Engineering in Medicine and Biology Magazine*, **16**(4):17–22.

Yarnold, P., Soltysik, R., and Martin, G. (1994). Heart rate variability and susceptibility for sudden cardiac death: an example of multivariable optimal discriminant analysis. *Statistics in Medicine*, **13**:1015–1021.

Young, G. and Householder, A. (1938). Discussion of a set of points in terms of their mutual distances. *Psychometrika*, **3**:19–22.

Zebrowski, J., Poplawska, W., and Baranowski, R. (1994). Entropy, pattern entropy, and related methods for the analysis of data on the time intervals between heartbeats from 24-h electrocardiograms. *Physical Review E*, **50**(5):4187–4205.

Zetterberg, L. (1969). Estimation of parameters for a linear difference equation with application to EEG analysis. *Mathematical Biosciences*, **5**:227–275.

Zhao, L., Reisman, S., and Findley, T. (1994). Derivation of respiration from electrocardiogram during heart rate variability studies. *Computers in Cardiology 1994*, pages 53–56, Bethesda, USA. IEEE Computer Society Press, Los Amitos, USA.

Zwillich, C., Devlin, T., and White, D. (1982). Bradycardia during sleep apnea: characteristics and mechanism. *Journal of Clinical Investigation*, **69**:1286–1292.