CHAPTER 3

ECG Statistics, Noise, Artifacts, and Missing Data

Gari D. Clifford

3.1 Introduction

Chapter 1 presented a description of the ECG in terms of its etiology and clinical features, and Chapter 2 an overview of the possible sources of error introduced in the hardware collection and data archiving stages. With this groundwork in mind, this chapter is intended to introduce the reader to the ECG using a signal processing approach. The ECG typically exhibits both persistent features (such as the average P-QRS-T morphology and the short-term average heart rate, or average RR interval), and nonstationary features (such as the individual RR and QT intervals, and long-term heart rate trends). Since changes in the ECG are quasi-periodic (on a beat-to-beat, daily, and perhaps even monthly basis), the frequency can be quantified in both statistical terms (mean, variance) and via spectral estimation methods. In essence, all these statistics quantify the power or degree to which an oscillation is present in a particular frequency band (or at a particular scale), often expressed as a ratio to power in another band. Even for scale-free approaches (such as wavelets), the process of feature extraction tends to have a bias for a particular scale which is appropriate for the particular data set being analyzed. ECG statistics can be evaluated directly on the ECG signal, or on features extracted from the ECG. The latter category can be broken down into either morphology-based features (such as ST level) or timing-based statistics (such as heart rate variability). Before discussing these derived statistics, an overview of the ECG itself is given.

3.2 Spectral and Cross-Spectral Analysis of the ECG

The short-term spectral content for a lead II configuration and the source ECG segment are shown in Figure 3.1. Note the peaks in the power spectral density (PSD) at 1, 4, 7, and 10 Hz, corresponding approximately to the heart rate (60 bpm), T wave, P wave, and the QRS complex, respectively. The spectral content for each lead is highly similar regardless of the lead configuration, although the actual energy at each frequency may differ.

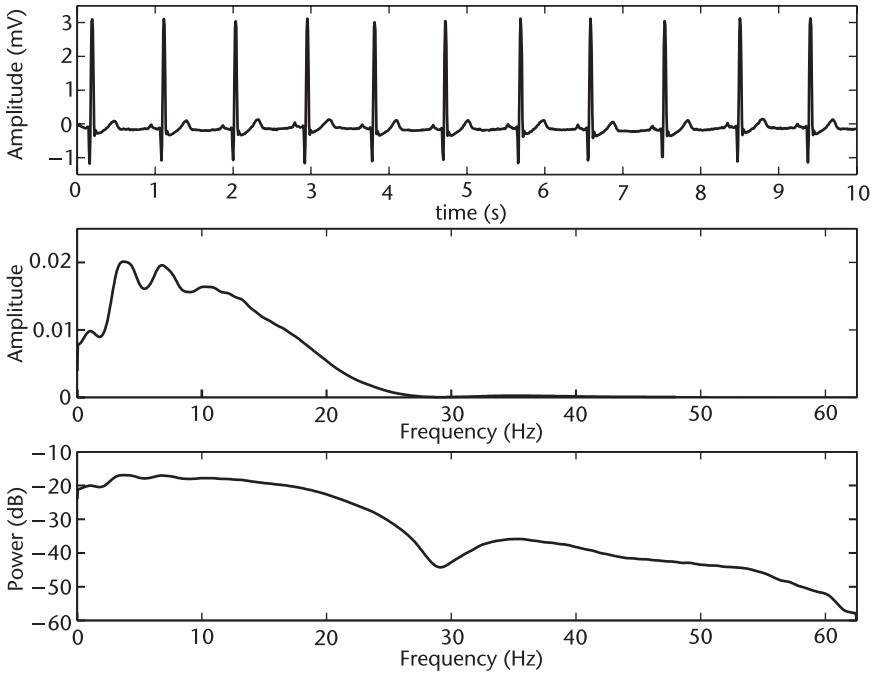

Figure 3.1 Ten seconds of 125-Hz typical ECG in sinus rhythm recorded with a lead II placement (upper plot) and associated linear and log-linear periodograms (middle and lower plots, respectively). A 256-point Welch periodogram was used with a hamming window and a 64-point overlap for the PSD calculation.

Figure 3.2 illustrates the PSDs for a typical full (12-lead) 10-second recording.[1] To estimate the spectral similarity between pairs of leads, the cross spectral coherence (CSC) can be calculated. The magnitude squared coherence estimate between two signals x and y, is

$$C_{xy} = \left| P_{xy}^2 \right| / (P_x P_y) \tag{3.1}$$

where P_x is the power spectral estimate of x, P_y is the power spectral estimate of y, and P_{xy} is the cross power spectral estimate[2] of x and y. Coherence is a function of frequency with C_{xy} ranging between 0 and 1 and indicates how well signal x corresponds to signal y at each frequency.

The CSC between any pair of leads will give values greater than 0.9 at most physiologically significant frequencies (1 to 10 Hz); see Figure 3.3. Note also that there is a significant coherent component between 12 and 50 Hz. By comparing this with the CSC between two adjacent 10-second segments of the same ECG lead, we can see that this higher frequency component is absent, indicating that it is due to some transient or incoherent phenomena, such as observation or muscle noise. Note that there is still a significant amount of coherence within the spectral band

1. [Px, Fx] = PWELCH(ECG,HAMMING(512),256,512,1000); in Matlab.
2. This operation can be achieved by using Matlab's MSCOHERE.M which uses Welch's averaged periodogram method [1], or by using COHERE.C from PhysioNet [2].

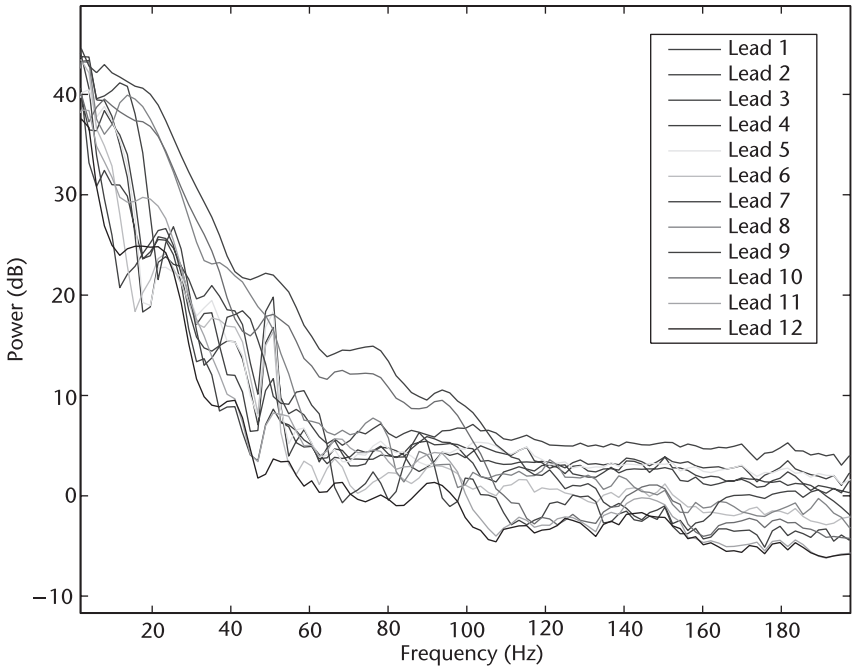

Figure 3.2 PSD (dB/Hz) of all 12 standard leads of 10 seconds of an ECG in sinus rhythm. A 512-point Welch periodogram was used with a hamming window and with a 256-point overlap. Note that the leads are numbered arbitrarily, rather than using their clinical labels.

corresponding to the heart rate (HR), T wave, P wave, and QRS complex (1 to 10 Hz). Changing heart rates (which lead to changing morphology; see Section 3.3) and varying the FFT window size and overlap will change the relative magnitude of this cross-coherence. Furthermore, different pairs of leads may show differing degrees of CSC due to dispersion effects (see Section 3.3).

3.2.1 Extreme Low- and High-Frequency ECG

Although the accepted range of the diagnostic ECG is often quoted to be from 0.05 Hz (for ST analysis) to 40 or 100 Hz, information does exist beyond these limits. Ventricular late potentials (VLPs) are microvolt fluctuations that manifest in the terminal portion of the QRS complex and can persist into the ST-T segment. They represent areas of delayed ventricular activation which are manifestations of slowed conduction velocity (resulting from ischemia or deposition of collagens after an acute myocardial infarction). VLPs, therefore, are interesting for heart disease diagnosis [3–5]. The upper frequency limit of VLPs can be as high as 500 Hz [6].

On the low frequency end of the spectrum, Jarvis and Mitra [7] have demonstrated that sleep apnea may be diagnosed by observing power changes in the ECG at 0.02 Hz.

3.2.2 The Spectral Nature of Arrhythmias

Arrhythmias, which manifest due to abnormalities in the conduction pathways of the heart, can generally be grouped into either atrial or ventricular arrhythmias. Ventricular arrhythmias manifest as gross distortions of the beat morphology since

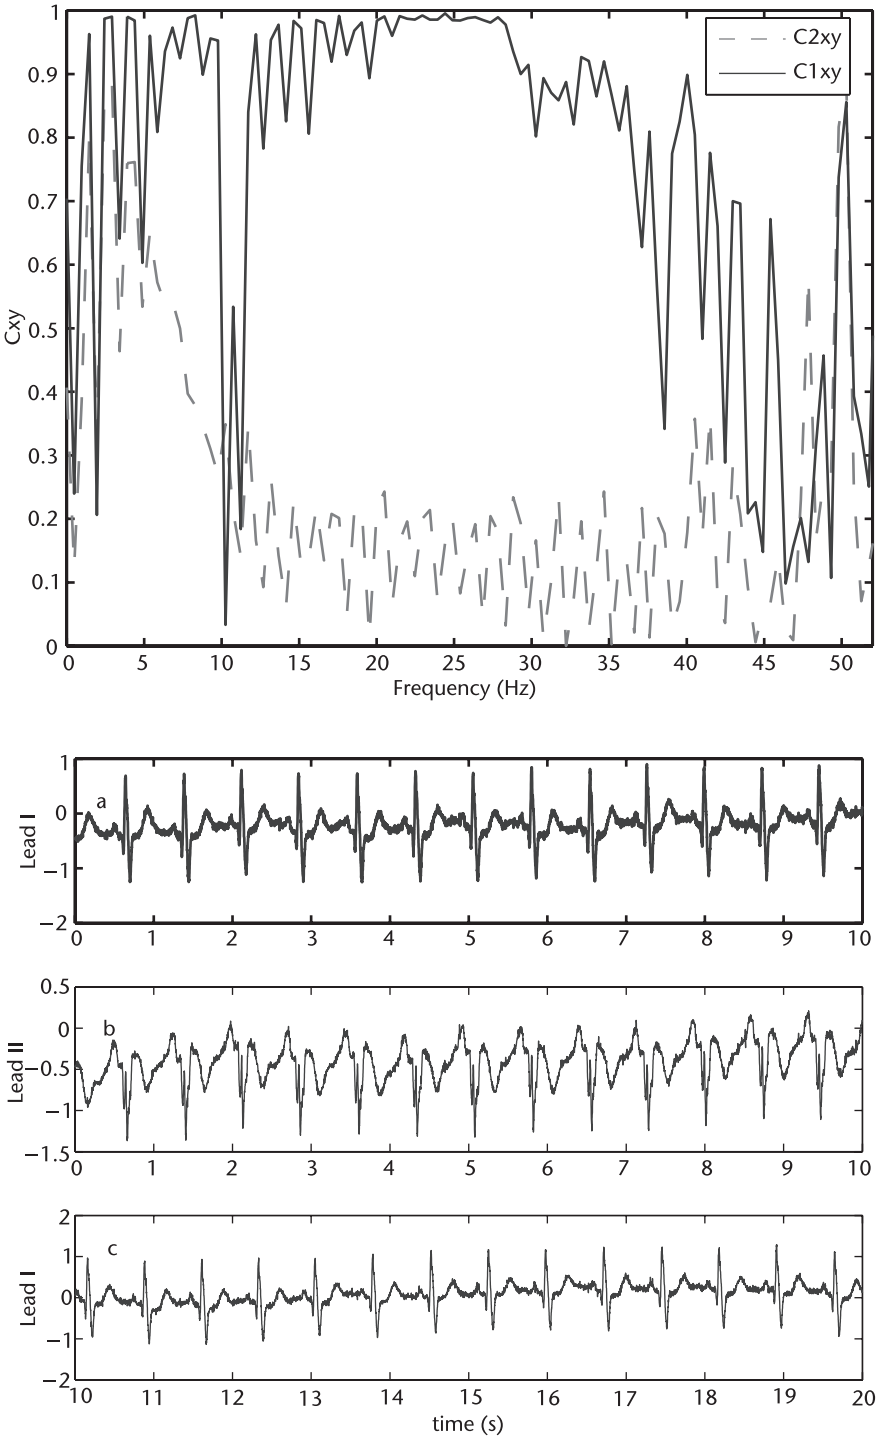

Figure 3.3 Cross-spectral coherence of two ECG sections in sinus rhythm. $C1_{xy}$ (solid line) is the CSC between two simultaneous lead I and lead II sections of ECG (plot *a* and plot *b* in the lower half of the figure). Note the significant coherence between 3 Hz and 35 Hz. $C2_{xy}$ (dashed line) is the CSC between two adjacent 10-second sections of lead I ECG (plot *a* and plot *c* in the lower half of the figure). Note that there is significantly less coherence between the adjacent signals except at 50 Hz (mains noise) and between 1 and 10 Hz.

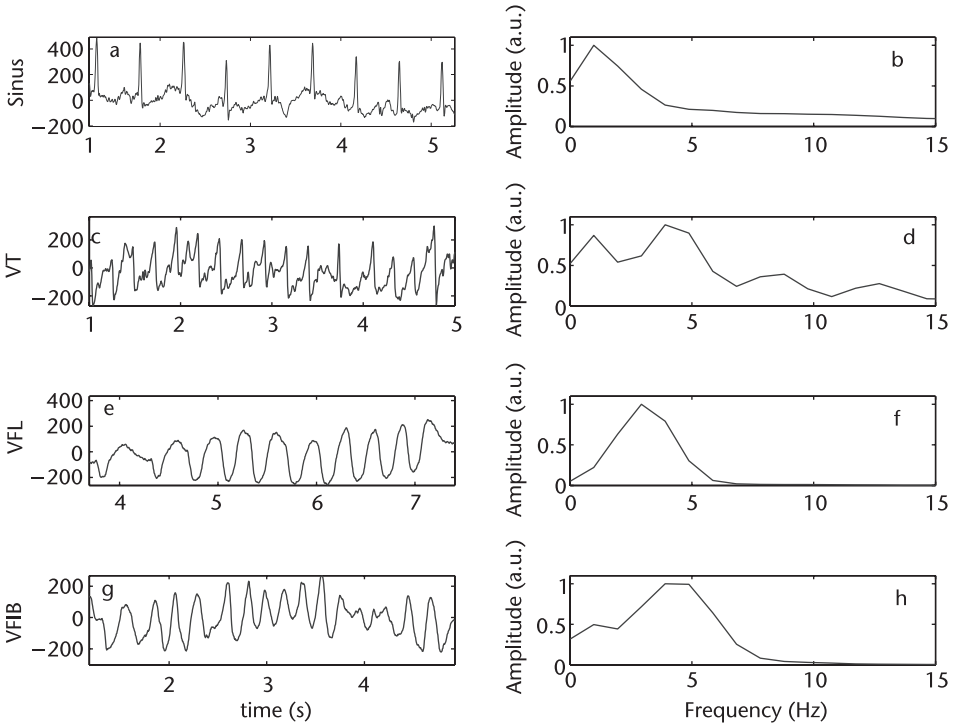

Figure 3.4 (a) Sinus rhythm and (b) corresponding PSD. (c) Ventricular tachycardia (VT) and (d) corresponding PSD. (e) Ventricular flutter (VFL) and (f) corresponding PSD. (g) Ventricular fibrillation (VFIB) and (h) corresponding PSD. Note that ventricular beats exhibit broader QRS complexes and therefore a shift in QRS energy to lower frequencies. Note also that higher frequencies (than normal) also manifest. VFL destroys many of the subtle ECG features and manifests as a sinusoidal-like oscillation around the frequency of the (rapid) heart rate. VFIB manifests as a less organized and more rapid oscillation, and therefore the spectrum is broader with more energy at higher frequencies. (All PSDs were calculated on 5-second segments with the same parameters as in Figure 3.1, but linear scales are used for clarity.)

the depolarization begins in the ventricles rather than the atria. The QRS complex becomes broader due to the depolarization occurring along an abnormal conduction path and therefore progressing more slowly, masking the latent P wave from delayed atrial depolarization. Figure 3.4(a) illustrates a 5-second segment of ventricular tachycardia (VT) with a high heart rate of around 180 bpm or 3 Hz, and the accompanying power spectral density [Figure 3.4(b)]. Although the broadening of the QRS complexes during VT causes a shift in the QRS spectral peak to slightly lower frequencies, the overall peaks are similar to the spectrum of a sinus rhythm[3] (see Figure 3.1), and therefore, spectral separation between sinus and VT rhythms is difficult. Figure 3.4(a) shows a 5-second segment of sinus rhythm ECG for the same patient before the episode of VT, with a relatively high heart rate (108 bpm). Note that although the P waves, QRS complexes, and T waves are discernible above the noise, the main spectral component is the 1- to 2-Hz baseline noise.

3. Below 60 bpm sinus rhythm is known as sinus bradycardia, and between 100 to 150 bpm it is known as sinus tachycardia. Note also that sinus rhythm is sometimes known as sinus arrhythmia if the heart rate rises and falls periodically, such as in RSA; see Section 3.7.

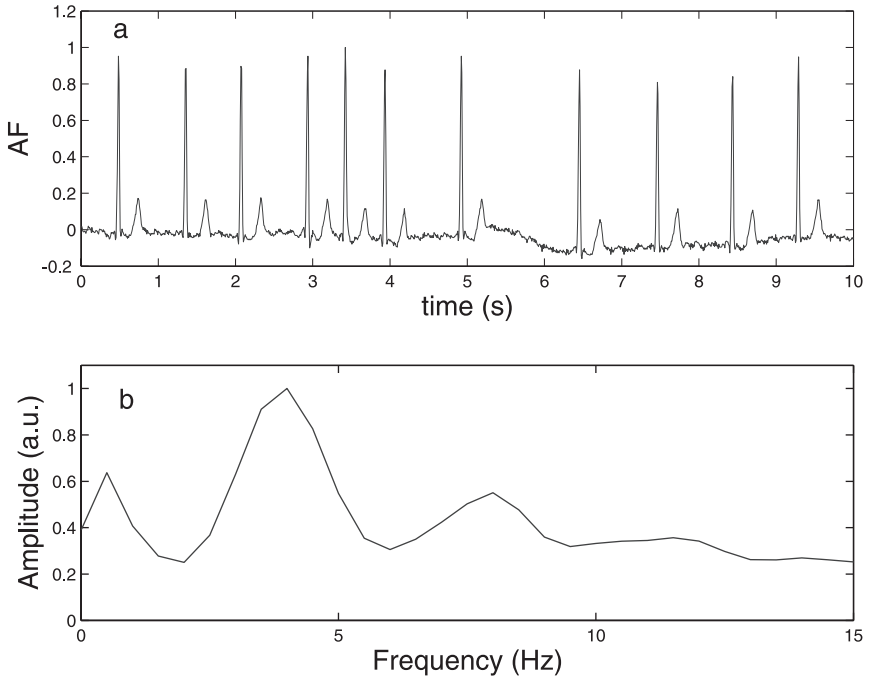

Figure 3.5 (a) Atrial fibrillation (AF) and (b) corresponding PSD. Note the similarity to sinus rhythm in Figure 3.4(a, b). (All PSDs were calculated with the same parameters as in Figure 3.4.)

When the ventricular activation time slows sufficiently, QRS complexes become severely broadened and ventricular flutter (VFL) is possible. This arrhythmia manifests as sinusoidal-like disturbances in the ECG, and is therefore relatively easy to detect through spectral methods. Figure 3.4(e) illustrates a 4-second segment of transient VFL and the corresponding power spectrum [Figure 3.4(f)]. If the ventricular arrhythmia is more erratic and manifests with a higher frequency of oscillation, then it is known as the extreme condition ventricular fibrillation (VFIB). Colloquially, the heart is said to be *squirming* "like a bag of worms," with little or no coherent activity. At this point, the heart is virtually useless as a pump and immediate physical or electrical intervention is required to encourage the cardiac cells to depolarize/repolarize in a coherent manner.

Atrial arrhythmias, in contrast to ventricular arrhythmias, manifest as small disturbances in the timing and relative position of the (relatively low amplitude) P wave and are therefore difficult to detect through spectral methods. Figure 3.5 illustrates the ECG and its corresponding power spectrum for an atrial arrhythmia. Atrial arrhythmias do, however, manifest significantly different changes in the beat-to-beat timing and can therefore be detected by collecting and analyzing statistics on such intervals [8] (see Section 3.5.3).

3.3 Standard Clinical ECG Features

Clinical assessment of the ECG mostly relies on relatively simple measurements of the intrabeat timings and amplitudes. Averaging over several beats is common to

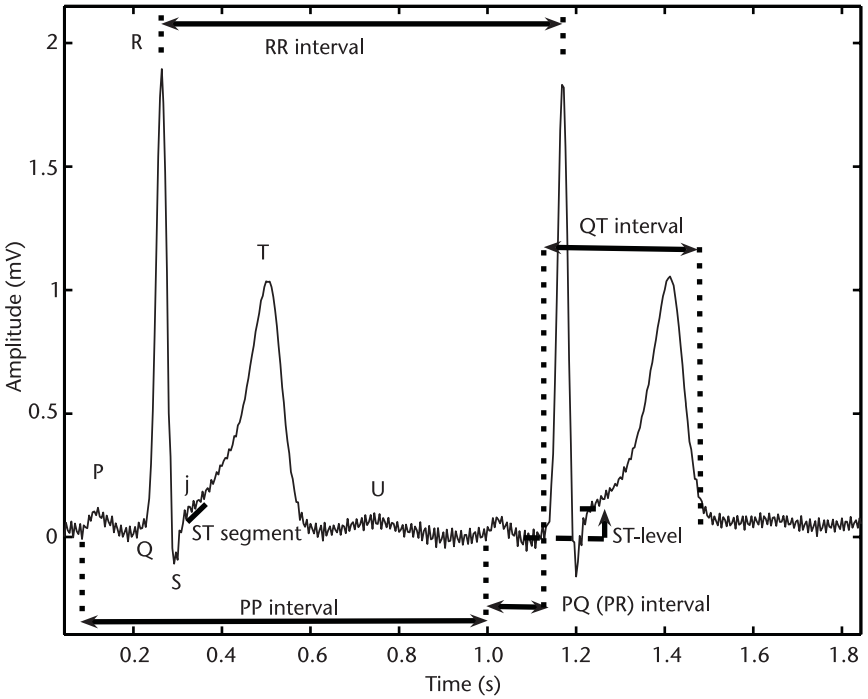

Figure 3.6 Standard fiducial points in the ECG (P, Q, R, S, T, and U) together with clinical features (listed in Table 3.1).

either reduce noise or average out short-term beat-to-beat interval-related changes. The complex heart rate-related changes in the ECG morphology (such as QT hysteresis[4]) can themselves be indicative of problems. However, a clinician can extract enough diagnostic information to make a useful assessment of cardiac abnormality from just a few simple measurements.

Figure 3.6 illustrates the most common clinical features, and Table 3.1 illustrates typical normal values for these standard clinical ECG features in healthy adult males in sinus rhythm, together with their upper and lower limits of normality. Note that these figures are given for a particular heart rate. It should also be noted that the heart rate is calculated as the number of P-QRS-T complexes per minute, but is often calculated over shorter segments of 15 and sometimes 30 seconds. In terms of modeling we can think of this heart rate as our *operating point* around which the local interbeat interval rises and falls. Of course, we can calculate a heart rate over any scale, up to a single beat. In the latter case, the heart rate is termed the instantaneous (or beat-to-beat) heart rate, $\mathrm{HR_i} = 60/RR_n$, of the nth beat. Each consecutive beat-to-beat, or RR, interval[5] will be of a different length (unless the patient is paced), and a correlated change in ECG morphology is seen on a beat-to-beat basis.

4. See Section 3.4 and Chapter 11.
5. The beat-to-beat interval is usually measured between consecutive R-peaks and hence termed the RR interval. See Section 3.7.

Table 3.1 Typical Lead II ECG Features and Their Normal Values in Sinus Rhythm at a Heart Rate of 60 bpm for a Healthy Male Adult (see text and Figure 3.6 for definitions of intervals)

Feature	*Normal Value*	*Normal Limit*
P width	110 ms	±20 ms
PQ/PR interval	160 ms	±40 ms
QRS width	100 ms	±20 ms
QTc interval	400 ms	±40 ms
P amplitude	0.15 mV	±0.05 mV
QRS height	1.5 mV	±0.5 mV
ST level	0 mV	±0.1 mV
T amplitude	0.3 mV	±0.2 mV

Note: There is some variation between lead configurations. Heart rate, respiration patterns, drugs, gender, diseases, and ANS activity also change the values. $QTc = \alpha QT$ where $\alpha = (RR)^{-\frac{1}{2}}$. About 95% of (normal healthy adult) people have a QTc between 360 ms and 440 ms. Female durations tend to be approximately 1% to 5% shorter except for the QT/QTc, which tends to be approximately 3% to 6% longer than for males. Intervals tend to elongate with age, at a rate of approximately 10% per decade for healthy adults.

Often, the RR interval will oscillate periodically, shortening with inspiration (and lengthening with expiration). This phenomenon, known as respiratory sinus arrhythmia (RSA) is partly due to the Bainbridge reflex, the expansion and contraction of the lungs and the cardiac filling volume caused by variations of intrathoracic pressure [9]. During inspiration, the pressure within the thorax decreases and venous return increases, which stretches the right atrium resulting in a reflex that increases the local heart rate (i.e., shortens the RR intervals). During expiration, the reverse of this process results in a slowing of the local heart rate. In general, the normal beat-to-beat changes in morphology are ignored, except for derivations of respiration, although the phase between the respiratory RR interval oscillations and respiratory-related changes in ECG morphology is not static; see Section 3.8.2.2 and Chapter 8. The reason for this is that the mechanisms which alter amplitude and timing on the ECG are not exactly the same (although they are coupled either mechanically or neurally with a phase delay which may change from beat to beat; see Chapter 8). Changes in the features in Table 3.1 and Figure 3.6, therefore, occur on a beat-to-beat basis as well as because of shifts in the operating point (average heart rate), although this is a second order effect.

The PR interval extends from the start of the P wave to the end of the PQ-junction at the very start of the QRS complex (that is, to the start of the R or Q wave). Therefore, this interval is sometimes known as the PQ interval. This interval represents the time required for the electrical impulse to travel from the SA node to the ventricle and normal values range between 120 and 200 ms. The PR interval has been shown to lengthen and shorten with respiration in a similar manner to the RR interval, but is less pronounced and is not fully correlated with the RR interval oscillations [10].

The global point of reference for the ECG's amplitude is the isoelectric level, measured over the short period on the ECG between the atrial depolarization (P wave) and the ventricular depolarization (QRS complex). In general, this point is thought to be the most stable marker of 0V for the surface ECG since there is a short pause before the current is conducted between the atria and the ventricles.

Interbeat segments are not usually used as a reference point because activity before the P wave can often be dominated by preceding T-wave activity.

The QRS width is representative of the time for the ventricles to depolarize, typically lasting 80 to 120 ms. The lower the heart rate, the wider the QRS complex, due to decreases in conduction speed through the ventricle. The QRS width also changes from beat-to-beat based upon the QRS axis (see Chapter 1), which is correlated with the phase of respiration (see Chapter 8) and with changes in RR interval and therefore the local heart rate. The RS segment of the QRS complex is known as the ventricular activation time (VAT) and is usually shorter (lasting around 40 ms) than the QR segment. This asymmetry in the QRS complex is not a constant and varies based upon changes in the autonomic nervous system (ANS) axis, lead position, respiration and heart rate (see Chapter 8).

The QRS complex usually rises (for positive leads) or falls to about 1 to 2 mV from the isoelectric line for normal beats. Artifacts (such as electrode movements) and abnormal beats (such as ventricular ectopic beats) can be several times larger in amplitude. In particular, baseline wander can often be the largest amplitude signal on the ECG, with the QRS complexes appearing as almost indistinguishable periodic anomalies. For this reason, it is important to allow sufficient dynamic range in the amplification (or digital storage) of ECG data; see Chapter 2.

The point of inflection after the S wave is known as the *j*-point, and is often used to define the beginning of the ST segment. In normals, it is expected to be isoelectric since it is the pause between ventricular depolarization and repolarization. The ST level is generally measured around 60 to 80 ms after the *j*-point, with adjustments for local heart rates (see Chapters 9 and 10). Abnormal changes in the ECG, defined by the Sheffield criteria [11], are ST level shifts ≥ 0.1 mV (or about 5% to 10% of the QRS amplitude for a sinus beat on a V5 lead). Since only small deviations form the isoelectric level are significant markers of cardiac abnormality (such as ischemia), the correct measurement of the isoelectric line is crucial. The interbeat segments between the end of the P wave and start of the Q wave are so short (less than 10 samples at 125 Hz), that the isoelectric baseline measurement is prone to noise. Multiple-beat averaging is therefore often employed. ST segment and *j*-point elevation, common in athletes, has been reported to normalize with exercise [12] and therefore *j*-point elevations may be difficult to distinguish from other changes seen in ECG.

The QT interval is measured between the onset of the QRS complex and the end of the T wave. It is considered to represent the time between the start of ventricular depolarization and the end of ventricular repolarization and is therefore useful as a measure of the duration of repolarization (see Chapter 11). The QT interval varies depending on heart rate, age, and gender. As with some other parameters in the ECG, it is possible to approximate the (average) heart rate dependency of the QT interval by multiplying it by a factor $\alpha = (\hat{RR})^{-\frac{1}{2}}$ where $\hat{RR}$ is the local average RR interval. The resultant QT interval is called the *corrected* QT interval, QTc [13]. However, this factor works over a limited range and is subject dependent to some degree, over and above the usual confounding variables of age, gender, and drug regime; see Section 3.4.1.

Furthermore, ANS activity shifts can also change α. In general, the last RR interval duration affects the action potential (see Chapter 1) and hence the QT interval. It is also known that the QT-RR dependence is both a function of the

average heart rate and the instantaneous interval, RR_i [14]. Note that there is some variation in these parameters between lead configurations. Although inter-lead differences are sometimes used as cardiovascular markers themselves (such as in QT dispersion [15]), it is unclear whether there is a specific physiological origin to such differences, or whether such metrics are just measuring an artifact which correlates with a clinical marker [16, 17].

One of the problems in measuring the QT interval correctly (apart from the noise in the ECG and the resultant onset and offset ambiguities) is due to the changes in the *j*-point and T wave morphology with heart rate. It has been observed that as the heart rate increases, the T wave increases in height and becomes more symmetrical [18]. Furthermore, in some subject groups (such as athletes), the T wave is often observed to be inverted [12].

To summarize, the following changes are typically observed with increasing heart rate [12, 18, 19]:

- The average RR interval decreases.
- The PR segment shortens and slopes downward (in the inferior leads).
- The P wave height increases.
- The Q wave becomes slightly more negative (at very high heart rates).
- The QRS width decreases.
- The R wave amplitude decreases in the lateral leads (e.g., V5) at and just after high heart rates.
- The S wave becomes more negative in the lateral and vertical leads (e.g., V5 and aVF). As the R wave decreases in amplitude, the S wave increases in depth.
- The *j*-point often becomes depressed in lateral leads. However, subjects with a normal or resting *j*-point elevation may develop an isoelectric *j*-point with higher heart rates.
- The ST level changes (depressed in inferior leads).
- The T wave amplitude increases and becomes more symmetrical (although it can initially drop at the onset of a heart rate increase).
- The QT interval shortens (depending on the autonomic tone).
- The U wave does appear to change significantly. However, U waves may be difficult to identify due to the short interval between the T and following beat's P waves at high heart rates.

It should be noted however, that this simple description is insufficient to describe the complex changes that take place in the ECG as the heart rate increases and decreases. These dynamics are further explored in the following section.

3.4 Nonstationarities in the ECG

Nonstationarities in the ECG manifest both in an interbeat basis (as RR interval timing changes) and on an intrabeat basis (as morphological changes). Although the former changes are often thought of as rhythm disturbances and the latter as beat abnormalities, the etiology of the changes are often intricately connected. To be clear, although we could categorize the beat-to-beat changes in the RR interval

timing and ECG morphology as nonstationary, they can actually be well represented by nonlinear models (see Section 3.7 and Chapter 4). This chapter therefore refers to these changes as stationary (but nonlinear). The transitions between rhythms is a nonstationary process (although some nonlinear models exist for limited changes). In this chapter, abnormal changes in beat morphology or rhythm that suggest a rapid change in the underlying physiology are referred to as nonstationary.

3.4.1 Heart Rate Hysteresis

So far we have not considered the dynamic effects of heart rate on the ECG morphology. Sympathetic or parasympathetic changes in the ANS which lead to changes in the heart rate and ECG morphology are asymmetric. That is, the dynamic changes that occur as the heart rate increases, are not matched (in a time symmetric manner) when the heart rate reduces and there is a (several beat) lag in the response between the RR interval change and the subsequent morphology change. One well-known form of heart rate-related hysteresis is that of *QT hysteresis*. In the context of QT interval changes, this means that the standard QT interval correction factors[6] are a gross simplification of the relationship, and that a more dynamic model is required. Furthermore, it has been shown that the relationship between the QT and RR interval is highly individual-specific [20], perhaps because of the dynamic nature of the system. In the QT-RR phase plane, the trajectory is therefore not confined to a single line and *hysteresis* is observed. That is, changes in RR interval do not cause immediate changes in the QT interval and ellipsoid-like trajectories manifest in the QT-RR plane. Figure 3.7 illustrates this point, with each of the central contours indicating a response of either tachycardia (RT) and bradycardia (RB) or normal resting. From the top right of each contour, moving counterclockwise (or anticlockwise); as the heart rate increases (the RR interval drops) the QT interval remains constant for a few beats, and then begins to shorten, approximately in an inverse square manner. When the heart rate drops (RR interval lengthens) a similar time delay is observed before the QT interval begins to lengthen and the subject returns to approximately the original point in the QT-RR phase plane. The difference between the two trajectories (caused by RR acceleration and deceleration) is the QT hysteresis, and depends not only on the individual's physiological condition, but also on the specific activity in the ANS. Although the central contour defines the limits of normality for a resting subject, active subjects exhibit an extended QT-RR contour. The 95% limits of normal activity are defined by the large, asymmetric dotted contour, and activity outside of this region can be considered abnormal.

The standard QT-RR relationship for low heart rates (defined by the Fridericia correction factor $\mathrm{QTc} = \mathrm{QT}/RR^{1/3}$) is shown by the line cutting the phase plane from lower left to upper right. It can be seen that this factor, when applied to the resting QT-RR interval relationship, overcorrects the dynamic responses in the normal range (illustrated by the striped area above the correction line and below the normal dynamic range) or underestimates QT prolongation at low heart rates

6. Many QT correction factors have been considered that improve upon Bazett's formula ($\mathrm{QTc} = \mathrm{QT}/\sqrt{RR}$), including linear regression fitting ($\mathrm{QTc} = \mathrm{QT} + 0.154(1 - RR)$), which works well at high heart rates, and the Fridericia correction ($\mathrm{QTc} = \mathrm{QT}/RR^{1/3}$), which works well at low heart rates.

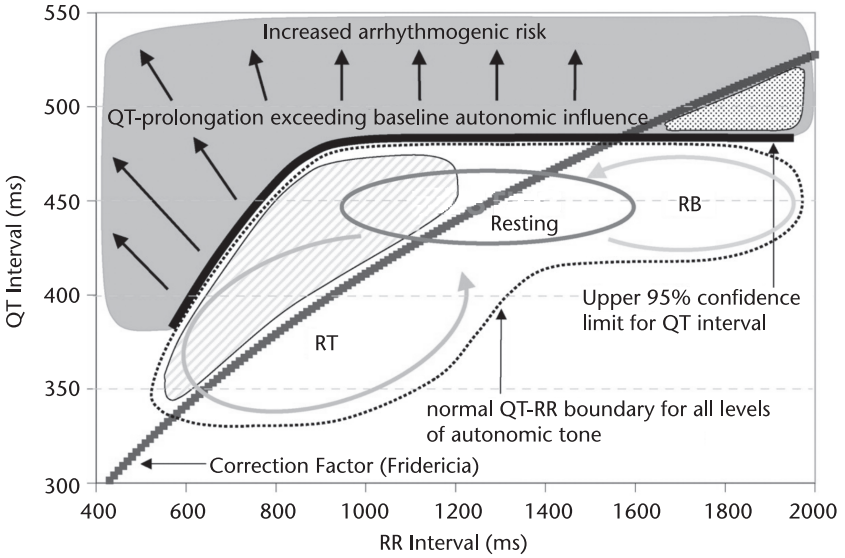

Figure 3.7 Normal dynamic QT-RR interval relationship (dotted-line forming asymmetric contour) encompasses autonomic reflex responses such as tachycardia (RT) and bradycardia (RB) with hysteresis. The statistical outer boundary of the normal contour is defined as the upper 95% confidence bounds. The Fridericia correction factor applied to the resting QT-RR interval relationship overcorrects dynamic responses in the normal range (striped area above correction line and below 95% confidence bounds) or underestimates QT prolongation at slow heart rates (shaded area above 95% confidence bounds but below Fridericia correction). QT prolongation of undefined arrhythmogenic risk (dark shaded area) occurs when exceeding the 95% confidence bounds of QT intervals during unstressed autonomic influence. (*From:* [21]. © 2005 ASPET: American Society for Pharmacology and Experimental Therapeutics. Reprinted with permission.)

(shaded area above normal range but below Fridericia correction) [21]. Abnormal QT prolongation is illustrated by the upper dark shaded area, and is defined to be when the QT-RR vector exceeds the 95% normal boundary (dotted line) during unstressed autonomic influence [21].

Another, more recently documented heart rate-related hysteresis is that of ST/HR [22], which is a measure of the ischemic reaction of the heart to exercise. If ST depression is plotted vertically so that negative values represent ST elevation, and heart rate is plotted along the horizontal axis typical ST/HR diagrams for a clinically normal subject display a negative hysteresis in ST depression against HR, (a clockwise hysteresis loop in the ST-HR phase plane during postexercise recovery). Coronary artery disease patients, on the other hand, display a positive hysteresis in ST depression against HR (a counterclockwise movement in the hysteresis loop during recovery) [23].

It is also known that the PR interval changes with heart rate, exhibiting a (mostly) respiration-modulated dynamic, similar to (but not as strong as) the modulation observed in the associated RR interval sequence [24]. This activity is described in more detail in Section 3.7.

3.4.2 Arrhythmias

The normal nonstationary changes are induced, in part, by changes in the sympathetic and parasympathetic branches of the autonomic nervous system. However,

sudden (abnormal) changes in the ECG can occur as a result of malfunctions in the normal conduction pathways of the heart. These disturbances manifest on the ECG as, sometimes subtle, and sometimes gross distortions of the normal beat (depending on the observation lead or the physiological origin of the abnormality). Such beats are traditionally labeled by their etiology, into ventricular beats, supraventricular and atrial.[7]

Since ventricular beats are due to the excitation of the ventricles before the atria, the P wave is absent or obscured. The QRS complex also broadens significantly since conduction through the myocardium is consequently slowed (see Chapter 1). The overall amplitude and duration (energy) of such a beat is thus generally higher. QRS detectors can easily pick up such high energy beats and the distinct differences in morphology make classifying such beats a fairly straightforward task. Furthermore, ventricular beats usually occur much earlier or later than one would expect for a normal sinus beat and are therefore known as VEBs, ventricular *ectopic* beats (from the Greek, meaning out of place).

Abnormal atrial beats exhibit more subtle changes in morphology than ventricular beats, often resulting in a reduced or absent P wave. The significant changes for an atrial beat come from the differences in interbeat timings (see Section 3.2.2). Unfortunately, from a classification point of view, abnormal beats are sometimes more frequent when artifact increases (such as during stress tests). Furthermore, artifacts can often resemble abnormal beats, and therefore extra information from multiple leads and beat context are often required to make an accurate classification.

3.5 Arrhythmia Detection

If conduction abnormalities are transient, then an abnormal beat manifests. If conduction problems persist, then the abnormal morphology repeats and an arrhythmia is manifest, or the ECG degenerates into an almost unrecognizable pattern. There are three general approaches to arrhythmia analysis. One method is to perform QRS detection and beat classification, labeling an arrhythmia as a quorum of a series of beats of a particular type. The common alternative approach is to analyze a section of the ECG that spans several beat intervals, calculate a statistic (such as variance or a ratio of power at different frequencies) on which the arrhythmia classification is performed. A third option is to construct a model of the expected dynamics for different rhythms and compare the observed signal (or derived features) to this model. Such model-based approaches can be divided down into ECG-based methods or RR interval statistics-based methods. Linear ECG-modeling techniques [26] are essentially equivalent to spectral analysis. Nonlinear state-space model reconstructions have also been used [27], but with varying results. This may be partly due to the sensitivity of nonlinear metrics to noise. See Chapter 6 for a more detailed description of this technique together with a discussion of the problems associated with applying nonlinear techniques to noisy data.

7. The table in [25], which lists all the beat classifications labeled in the PhysioNet databases [2] together with their alphanumeric labels, provides an excellent detailed list of beat types and rhythms.

3.5.1 Arrhythmia Classification from Beat Typing

A run of abnormal beats can be classified as an arrhythmia. Therefore, as long as consistent fiducial points can be located on a series of beats, simple postprocessing of a beat classifier's output together with a threshold on the heart rate can be sufficient for correctly identifying many arrhythmias. For example, supraventricular tachycardia is the sustained presence of supraventricular ectopic beats, at a rate over 100 bpm. Many more complex classification schemes have been proposed, including the use of principal component analysis [28, 29] (see Chapters 9 and 10) hidden Markov models [30], interlead comparisons [31], cluster analysis [32], and a variety of supervised and unsupervised neural learning techniques [33–35]. Further details of the latter category can be found in Chapters 12 and 13.

3.5.2 Arrhythmia Classification from Power-Frequency Analysis

Sometimes there is no consistently identifiable fiducial point in the ECG, and analysis of the normal clinical features is not possible. In such cases, it is usual to exploit the changes in frequency characteristics that are present during arrhythmias [36, 37]. More recently, joint time-frequency analysis techniques have been applied [38–40], to take advantage of the nonstationary nature of the cardiac cycle.

Other interesting methods that make use of interchannel correlation techniques have been proposed [31], but results from using a decision tree and linear classifier on just three AR coefficients (effectively performing a multiple frequency band thresholding) give some of the most promising results. Dingfei et al. [26] report classification performance statistics (sensitivity, specificity) on the MIT-BIH database [2] of 93.2%, 94.4% for sinus rhythm, 100%, 96.2% for superventricular tachycardia, 97.7%, 98.6% for VT, and 98.6%, 97.7% for VFIB. They also report classification statistics (sensitivity, specificity) of 96.4%, 96.7% for atrial premature contractions (APCs), and 94.8%, 96.8% for premature ventricular contractions (PVCs).[8] Sensitivity and specificity figures in the mid to upper 90s can be considered state of the art. However, these results pertain to only one database and the (sensitive) window size is prechosen based upon the prior expectation of the rhythm. Despite this, this approach is extremely promising, and may be improved by developing a method for adapting the window size and/or using a nonlinear classifier such as a neural network.

3.5.3 Arrhythmia Classification from Beat-to-Beat Statistics

Zeng and Glass [8] described a model for AV node conduction which was able to accurately model many observations of the statistical distribution of the beat-to-beat intervals during atrial arrhythmias (see Chapter 4 for a more details on this model). This model-based approach was further extended in [41] to produce a method of classifying beats based upon their statistical distribution. Later, Schulte-Frohlinde et al. [42] produced a variant of this technique that includes a dimension of time and allows the researcher to observe the temporal statistical changes. Software for this technique (known as *heartprints*) is freely available from [43].

More recent algorithms have attempted to combine both the spectral characteristics and time domain features of the ECG (including RR intervals) [44].

8. Sometimes called VPCs (ventricular premature contractions).

The integration of such techniques can help improve arrhythmia classification, but only if the learning set is expanded in size and complexity in a manner that is sufficient to provide enough training examples to account for the increased dimensionality of the input feature space. See Chapters 12 and 13 for further discussions of training, test, and validation data sets.

3.6 Noise and Artifact in the ECG

3.6.1 Noise and Artifact Sources

Unfortunately, the ECG is often contaminated by noise and artifacts[9] that can be within the frequency band of interest and can manifest with similar morphologies as the ECG itself. Broadly speaking, ECG contaminants can be classified as [45]:

1. *Power line interference:* 50 $\pm$0.2 Hz mains noise (or 60 Hz in many data sets[10]) with an amplitude of up to 50% of full scale deflection (FSD), the peak-to-peak ECG amplitude;
2. *Electrode pop or contact noise:* Loss of contact between the electrode and the skin manifesting as sharp changes with saturation at FSD levels for periods of around 1 second on the ECG (usually due to an electrode being nearly or completely pulled off);
3. *Patient–electrode motion artifacts:* Movement of the electrode away from the contact area on the skin, leading to variations in the impedance between the electrode and skin causing potential variations in the ECG and usually manifesting themselves as rapid (but continuous) baseline jumps or complete saturation for up to 0.5 second;
4. *Electromyographic (EMG) noise:* Electrical activity due to muscle contractions lasting around 50 ms between dc and 10,000 Hz with an average amplitude of 10% FSD level;
5. *Baseline drift*: Usually from respiration with an amplitude of around 15% FSD at frequencies drifting between 0.15 and 0.3 Hz;
6. *Data collecting device noise:* Artifacts generated by the signal processing hardware, such as signal saturation;
7. *Electrosurgical noise:* Noise generated by other medical equipment present in the patient care environment at frequencies between 100 kHz and 1 MHz, lasting for approximately 1 and 10 seconds;
8. *Quantization noise and aliasing*;
9. *Signal processing artifacts* (e.g., Gibbs oscillations).

Although each of these contaminants can be reduced by judicious use of hardware and experimental setup, it is impossible to remove all contaminants. Therefore, it is important to quantify the nature of the noise in a particular data set and

9. It should be noted that the terms *noise* and *artifact* are often used interchangeably. In this book *artifact* is used to indicate the presence of a transient interruption (such as electrode motion) and *noise* is used to describe a persistent contaminant (such as mains interference).
10. Including recordings made in North and Central America, western Japan, South Korea, Taiwan, Liberia, Saudi Arabia, and parts of the Caribbean, South America, and some South Pacific islands.

choose an appropriate algorithm suited to the contaminants as well as the intended application.

3.6.2 Measuring Noise in the ECG

The ECG contains very distinctive features, and automatic identification of these features is, to some extent, a tractable problem. However, quantifying the nonsignal (noise) element in the ECG is not as straightforward. This is partially due to the fact that there are so many different types of noises and artifacts (see above) that can occur simultaneously, and partially because these noises and artifacts are often transient, and largely unpredictable in terms of their onset and duration. Standard measures of noise-power assume stationarity in the dynamics and coloration of the noise. These include:

- Route mean square (RMS) power in the isoelectric region;
- Ratio of the R-peak amplitude to the noise amplitude in the isoelectric region;
- Crest factor / peak-to-RMS ratio (the ratio of the peak value of a signal to its RMS value);
- Ratio between in-band (5 to 40 Hz) and out-of-band spectral power;
- Power in the residual after a filtering process.

Except for ($16.\dot{6}$, 50, or 60 Hz) mains interference and sudden abrupt baseline changes, the assumption that most noise is Gaussian in nature is approximately correct (due to the central limit theorem). However, the coloration of the noise can significantly affect any interpretation of the value of the noise power, since the more colored a signal is, the larger the amplitude for a given power. This means that a signal-to-noise ratio (SNR) for a brown noise contaminated ECG (such as movement artifact) equates to a much cleaner ECG than the same SNR for an ECG contaminated by pink noise (typical for observation noise). Figure 3.8 illustrates this point by comparing a zero-mean unit-variance clean ECG (upper plot) with the same signal with additive noise of decreasing coloration (lower autocorrelation). In each case, the noise is set to be zero-mean with unit variance, and therefore has the same power as the ECG (SNR = 1). Note that the whiter the noise, the more significant the distortion for a given SNR. It is obvious that ECG analysis algorithms will perform differently on each of these signals, and therefore it is important to record the coloration of the noise in the signal as well as the SNR.

Determining the color of the noise in the ECG is a two-stage process which first involves locating and removing the P-QRS-T features. Moody et al. [28, 29] have shown that the QRS complex can be encoded in the first five principal components (PCs). Therefore, a good approximate method for removing the signal component from an ECG is to use all but the first five PCs to reconstruct the ECG. Principal component analysis (PCA) involves the projection of N-dimensional data onto a set of N orthogonal axes that represent the maximum directions of variance in the data. If the data can be well represented by such a projection, the p axes along which the variance is largest are good descriptors of the data. The $N - p$ remaining components are therefore projections of the noise. A more in-depth analysis of PCA can be found in Chapters 5 and 9.

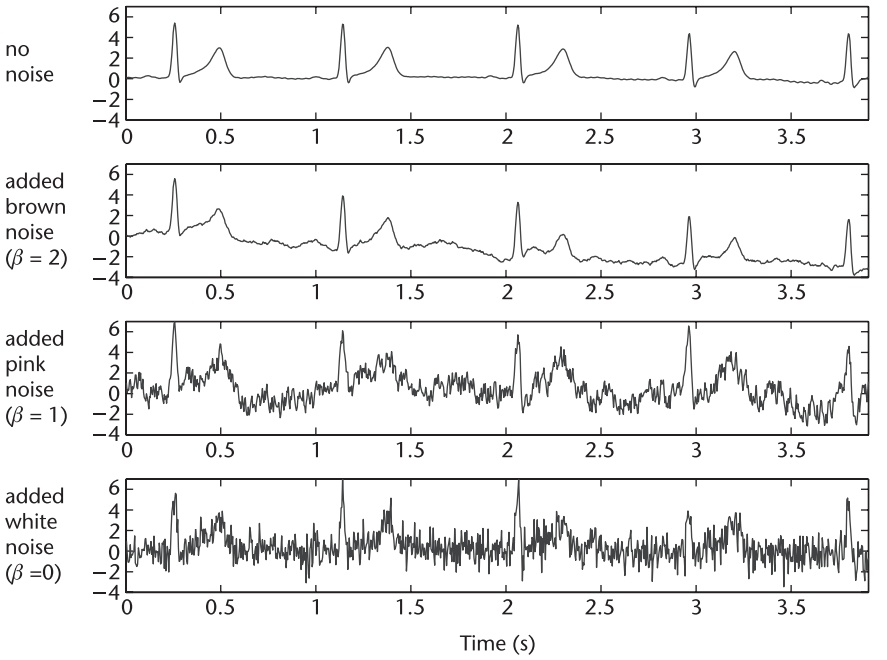

Figure 3.8 Zero-mean unit-variance clean ECG with additive brown, pink, and white noise (also zero-mean and unit-variance, and hence SNR = 1 in all cases).

Practically, this involves segmenting each beat in the given analysis window[11] such that the start of each P wave and the end of each T wave (or U wave if present) are captured in each segmentation with m-samples. The N beats are then aligned so that they form an $N \times m$ matrix denoted, X. If singular value decomposition (SVD) is then performed to determine the PCs, the five most significant components are discarded (by setting the corresponding eigenvalues to zero), and the SVD inverted, X becomes a matrix of only noise. The data can then be transformed back into a 1-D signal using the original segmentation indices.

The second stage involves calculating the log power-spectrum of this noise signal and determine its slope. The resultant spectrum has a $1/f^{\beta}$ form. That is, the slope β determines the color of the signal with the higher the value of β, the higher the auto-correlation. If $\beta = 0$, the signal is white (since the spectrum is flat) and is completely uncorrelated. If $\beta = 1$, the spectrum has a $1/f$ spectrum and is known as *pink* noise, typical of the observation noise on the ECG. Electrode movement noise has a Brownian motion-like form (with $\beta = 2$), and is therefore known as *brown* noise.

3.7 Heart Rate Variability

The baseline variability of the heart rate time series is determined by many factors including age, gender, activity, medications, and health [46]. However, not only

11. The window must contain at least five beats, and preferably at least 30 to capture respiration and ANS-induced changes in the ECG morphology; see Section 3.3.

does the mean beat-to-beat interval (the heart rate) change on many scales, but the variance of this sequence of each heartbeat interval does so too. On the shortest scale, the time between each heartbeat is irregular (unless the heart is paced by an artificial electrical source such as a pacemaker, or a patient is in a coma). These short-term oscillations reflect changes in the relative balance between the sympathetic and parasympathetic branches of the ANS, the *sympathovagal balance*. This heart rate irregularity is a well-studied effect known as *heart rate variability* (HRV) [47]. HRV metric values are often considered to reflect the competing actions of these different branches of the ANS on the sinoatrial (SA) node.[12] Therefore, RR intervals associated with abnormal beats (that do not originate from the SA node) should not be included in a HRV metric calculation and the series of consecutive normal-to-normal (NN) beat intervals should be analyzed.[13]

It is important to note that, the fiducial marker of each beat should be the onset of the P wave, since this is a more accurate marker than the R peak of the SA node stimulation (atrial depolarization onset) for each beat. Unfortunately, the P wave is usually a low-amplitude wave and is therefore often difficult to detect. Conversely, the R wave is easy to detect and label with a fiducial point. The exact location of this marker is usually defined to be either the highest (or lowest) point, the QRS onset, or the center of mass of the QRS complex. Furthermore, the competing effects of the ANS branches lead to subtle changes in the features within the heartbeat. For instance, a sympathetic innervation of the SA node (from exercise, for example) will lead to an increased local heart rate, and an associated shortening of the PR interval [10], QT interval [21], QRS width [48], and T wave [18]. Since the magnitude of the beat-to-beat modulation of the PR interval is correlated with, and much less significant than that of the RR interval [10, 49], and the R peak is well defined and easy to locate, many researchers choose to analyze only the *RR tachogram* (of normal intervals). It is unclear to what extent the differences in fiducial point location affects measures of HRV, but the sensitivity of the spectral HRV metrics to sampling frequencies below 1 kHz indicates that even small differences may have a significant effect for such metrics under certain circumstances [50].

If we record a typical RR tachogram over at least 5 minutes, and calculate the power spectral density,[14] then two dominant peaks are sometimes observable; one in the low frequency (LF) range ($0.015 < f < 0.15$ Hz) and one in the high frequency (HF) region ($0.15 \leq f \leq 0.4$ Hz). In general, the activity in the HF band is thought to be due mainly to parasympathetic activity at the sinoatrial node. Since respiration is a parasympathetically mediated activity (through the vagal nerve), a peak corresponding to the rate of respiration can often be observed in this frequency band (i.e., RSA). However, not all the parasympathetic activity is due to respiration. Furthermore, the respiratory rate may drop below the (generally accepted) lower bound of the HF region and therefore confound measures in the LF region. The LF region is generally thought to reflect sympathetically mediated activity[15] such as

12. See Chapter 1 for more details.
13. The temporal sequence of events is therefore known as the NN tachogram, or more frequently the RR tachogram (to indicate that each point is between each normal R peak).
14. Care must be taken at this point, as the time series is unevenly sampled; see section 3.7.2.
15. Although there is some evidence to show that this distinction does not always hold [46].

blood pressure-related phenomena. Activity in bands lower than the LF region are less well understood but seem to be related to myogenic activity, physical activity, and circadian variations. Note also that these frequency bands are on some level quite ad hoc and should not be taken as the exact limits on different mechanisms within the ANS; there are many studies that have used variants of these limits with practical results.

Many metrics for evaluating HRV have been described in the literature, together with their varying successes for discerning particular clinical problems. In general, HRV metrics can be broken down into either statistical time-based metrics (e.g., variance), or frequency-based metrics that evaluate power, or ratios of power, in certain spectral bands. Furthermore, most metrics are calculated either on a short time scale (often about 5 minutes) or over extremely long periods of time (usually 24 hours). The following two subsections give a brief overview of many of the common metrics. A more detailed analysis of these techniques can be found in the references cited therein. A comprehensive survey of the field of HRV was conducted by Malik et al. [46, 51] in 1995, and although much of the material remains relevant, some recent notable recent developments are included below, which help clarify some of the problems noted in the book. In particular, the sensitivity (and lack of specificity) of HRV metrics in many experiments has been shown to be partly due to activity-related changes [52] and the widespread use of resampling [53]. These issues, together with some more recent metrics, will now be explored.

3.7.1 Time Domain and Distribution Statistics

Time domain statistics are generally calculated on RR intervals without resampling, and are therefore robust to aggressive data removal (of artifacts and ectopic beats; see Section 3.7.6). An excellent review of conventional time domain statistics can be found in [46, 51]. One recently revisited time domain metric is the $pNN50$; the percentage of adjacent NN intervals differing by more than 50 ms over an entire 24-hour ECG recording. Mietus et al. [54] studied the generalization of this technique; the $pNNx$—the percentage of NN intervals in a 24-hour time series differing by more than $x$ms ($4 \leq x \leq 100$). They found that enhanced discrimination between a variety of normal and pathological conditions is possible by using a value of x as low as 20 ms or less, rather than the standard 50 ms threshold. This tool, and many of the standard HRV tools, are freely available from PhysioNet [2]. This work can be considered similar to recent work by Grogan et al. [55], who analyzed the predictive power of different bins in a smoothed RR interval histogram and termed the metric *cardiac volatility*. Histogram bins were isolated that were more predictive of deterioration in the ICU than conventional metrics, despite the fact that the data was averaged over many seconds. These results indicate that only certain frequencies of cardiac variability may be indicative of certain conditions, and that conventional techniques may be including confounding factors, or simply noise, into the metric and diminishing the metric's predictive power.

In Malik and Camm's collection of essays on HRV [51], metrics that involve a quantification of the probability distribution function of the NN intervals over a long period of time (such as the TINN, the "triangular index"), were referred to as *geometrical indices*. In essence, these metrics are simply an attempt at calculating

robust approximations of the higher order statistics. However, the higher the moment, the more sensitive it is to outliers and artifacts, and therefore, such "geometrical" techniques have faded from the literature.

The fourth moment, kurtosis, measures how peaked or flat a distribution is, relative to a Gaussian (see Chapter 5), in a similar manner to the TINN. Approximations to kurtosis often involve entropy, a much more robust measure of non-Gaussianity. (A key result of information theory is that, for a set of independent sources, with the same variance, a Gaussian distribution has the highest entropy, of all the signals.) It is not surprising then, that entropy-based HRV measures are more frequently employed that kurtosis.

The third moment of a distribution, *skewness*, quantifies the asymmetry of a distribution and has therefore been applied to patients in which sudden accelerations in heart rate, followed by longer decelerations, are indicative of a clinical problem. In general, the RR interval sequence accelerates much more quickly than it decelerates.[16] Griffin and Moorman [56] have shown that a small difference in skewness (0.59 ± 0.10 for sepsis and 0.51 ± 0.012 for sepsis-like illness, compared with -0.10 ± 0.13 for controls) can be an early indicator (up to 6 hours) of an upcoming abrupt deterioration in newborn infants.

3.7.2 Frequency Domain HRV Analysis

Heart rate changes occur on a wide range of time scales. Millisecond sympathetic changes stimulated by exercise cause an immediate increase in HR resulting in a lower long-term baseline HR and increased HRV over a period of weeks and months. Similarly, a sudden increase in blood pressure (due to an embolism, for example) will lead to a sudden semipermanent increase in HR. However, over many months the baroreceptors will reset their operating range to cause a drop in baseline HR and blood pressure (BP). In order to better understand the contributing factors to HRV and the time scales over which they affect the heart, it is useful to consider the RR tachogram in the frequency domain.

3.7.3 Long-Term Components

In general, the spectral power in the RR tachogram is broken down into four bands [46]:

1. Ultra low frequency (ULF): $0.0001 \text{ Hz} \geq \text{ULF} < 0.003 \text{ Hz}$;
2. Very low frequency (VLF): $0.003 \text{ Hz} \geq \text{VLF} < 0.04 \text{ Hz}$;
3. Low frequency (LF): $0.04 \text{ Hz} \geq \text{LF} < 0.15 \text{ Hz}$;
4. High frequency (HF): $0.15 \text{ Hz} \geq \text{HF} < 0.4 \text{ Hz}$.

Other upper- and lower-frequency bands are sometimes used. Frequency domain HRV metrics are then formed by summing the power in these bands, taking ratios,

16. Parasympathetic withdrawal is rapid, but is damped out by either parasympathetic activation or a much slower sympathetic withdrawal.

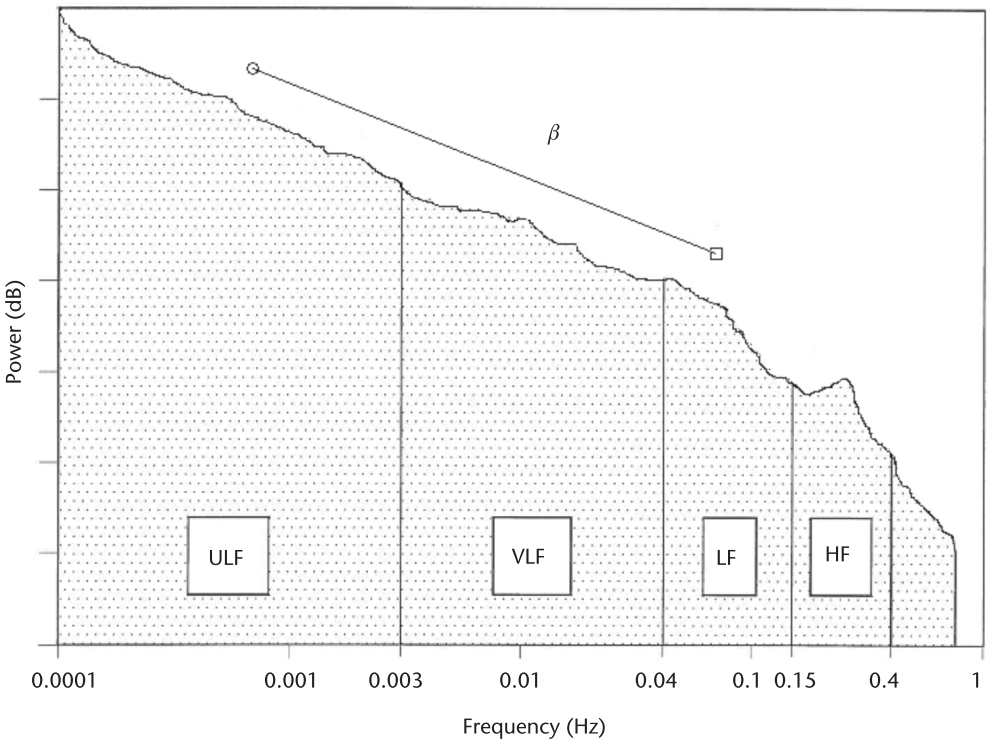

Figure 3.9 Typical periodogram of a 24-hour RR tachogram where power is plotted vertically and the frequency plotted horizontally on a log scale. Note that the gradient β of the $log - log$ plot is only meaningful for the longer scales. (*After:* [46].)

or calculating the slope,[17] β, of the $log - log$ power spectrum; see Figure 3.9. The motivation for splitting the spectrum into these frequency bands lies in the belief that the distinct biological regulatory mechanisms that contribute to HRV act at frequencies that are confined (approximately) within these bands. Fluctuations below 0.04 Hz in the VLF and ULF bands are thought to be due to long-term regulatory mechanisms such as the thermoregulatory system, the reninangiotensin system (related to blood pressure and other chemical regulatory factors), and other humoral factors [57]. In 1998 Taylor et al. [58] showed that the VLF fluctuations appear to depend primarily on the parasympathetic outflow. In 1999 Serrador et al. [59] demonstrated that the ULF band appears to be dominated by contributions from physical activity and that HRV in this band tends to increase during exercise. They therefore assert that any study that assesses HRV using data (even partially) from this frequency band should always include an indication of physical activity patterns. However, the effect of physical (and moreover, mental) activity on HRV is so significant that it has been suggested that controlling for activity for all metrics is extremely important [52].

Since spectral analysis was first introduced into HRV analysis in the late 1960s and early 1970s [60, 61], a large body of literature has arisen concerning this topic.

17. In the HRV literature, this slope is sometimes denoted by α.

In 1993, the U.S. Food and Drug Administration (FDA) withdrew its support of HRV as a useful clinical parameter due to a lack of consensus on the efficacy and applicability of HRV in the literature [62]. Although the Task Force of the European Society of Cardiology and the North American Society of Pacing Electrophysiology [46] provided an extensive overview of HRV estimation methods and the associated experimental protocols in 1996, the FDA has been reluctant to approve medical devices that calculate HRV unless the results are not explicitly used to make a specific medical diagnosis (e.g., see [63]). Furthermore, the clinical utility of HRV analysis (together with FDA approval) has only been demonstrated in very limited circumstances, where the patient undergoes specific tests (such as paced breathing or the Valsalva Maneuver) and the data are analyzed *off-line* by experts [64].

Almost all spectral analysis of the RR tachogram has been performed using some variant of autoregressive (AR) spectral estimation[18] or the FFT [46], which implicitly requires stationarity and regularly spaced samples. It should also be noted that most spectral estimation techniques such as the FFT require a windowing technique (e.g., the hamming window[19]), which leads to an implicit nonlinear distortion of the RR tachogram, since the value of the RR tachogram is explicitly joined to the time stamp.[20]

To mitigate for nonstationarities, linear and polynomial detrending is often employed, despite the lack of any real justification for this procedure. Furthermore, since the time stamps of each RR interval are related to the previous RR interval, the RR tachogram is inherently unevenly (or irregularly) sampled. Therefore, when using the FFT, the RR tachogram must either be represented in terms of power per cycle per beat (which varies based upon the local heart rate, and it is therefore extremely difficult, if not impossible, to compare one calculation with another) or a resampling method is required to make the time series evenly sampled.

Common resampling schemes involve either linear or cubic spline interpolative resampling. Resampling frequencies between 2 and 10 Hz have been used, but as long as the Nyquist criterion is satisfied, the resampling rate does not appear to have a serious effect on the FFT-based metrics [53]. However, experiments on both artificial and real data reveal that such processes overestimate the total power in the LF and HF bands [53] (although the increase is marginal for the cubic

18. Clayton et al. [65] have demonstrated that FFT and AR methods can provide a comparable measure of the low-frequency LF and high-frequency HF metrics on linearly resampled 5-minute RR tachograms across a patient population with a wide variety of ages and medical conditions (ranging from heart transplant patients who have the lowest known HRV to normals who often exhibit the highest overall HRV). AR models are particularly good at identifying line spectra and are therefore perhaps not an appropriate technique for analyzing HRV activity. Furthermore, since the optimal AR model order is likely to change based on the activity of the patient, AR spectral estimation techniques introduce an extra complication in frequency-based HRV metric estimation. AR modeling techniques will therefore not be considered in this chapter. As a final aside on AR analysis, it is interesting to note that measuring the width of a Poincaré plot is the same as treating the RR tachogram as an AR1 process and then estimating the process coefficient.
19. In the seminal 1978 paper on spectral windowing [66], Harris demonstrated that a hamming window (given by $W(t_j) = 0.54 - 0.46\cos(\omega t_j), [j = 0, 1, 2, \ldots, N-1]$) provides an excellent performance for FFT analysis in terms of spectral leakage, side lobe amplitude, and width of the central peak (as well as a rapid computational time).
20. However, the window choice does not appear to affect the HRV spectral estimates significantly for RR interval variability.

spline resampling if the RR tachogram is smoothly varying and there are no missing or removed data points due to ectopy or artifact; see Section 3.7.6). The FFT overestimates the $\frac{\text{LF}}{\text{HF}}$-ratio by about 50% with linear resampling and by approximately 10% with cubic spline resampling [53]. This error can be greater than the difference in the $\frac{\text{LF}}{\text{HF}}$-ratio between patient categories and is therefore extremely significant (see Section 3.7.7). One method for reducing (and almost entirely removing) this distortion is to use the Lomb-Scargle periodogram (LSP) [67–71], a method of spectral estimation which requires no explicit data replacement (nor assumes any underlying model) and calculates the PSD from only the known (observed) values in a time series.

3.7.4 The Lomb-Scargle Periodogram

Consider a physical variable X measured at a set of times t_j where the sampling is at equal times ($\Delta t = t_{j+1} - t_j = \text{constant}$) from a stochastic process. The resulting time series data, $\{X(t_j)\}$ $(i = 1, 2, ..., N)$, are assumed to be the sum of a signal X_s and random observational errors,[21] R;

$$X_j = X(t_j) = X_s(t_j) + R(t_j) \tag{3.2}$$

Furthermore, it is assumed that the signal is periodic, that the errors at different times are independent ($R(t_j) \neq f(R(t_k))$ for $j \neq k$) and that $R(t_j)$ is normally distributed with zero mean and constant variance, σ^2.

The N-point discrete Fourier transform (DFT) of this sequence is

$$FT_X(\omega) = \sum_{j=0}^{N-1} X(t_j) e^{-i\omega t_j} \tag{3.3}$$

($\omega_n = 2\pi f_n$, $n = 1, 2, ..., N$) and the power spectral density estimate is therefore given by the standard method for calculating a periodogram:

$$P_X(\omega) = \frac{1}{N} \sum_{j=0}^{N-1} \left| X(t_j) e^{-i\omega t_j} \right|^2 \tag{3.4}$$

Now consider arbitrary t_j's or uneven sampling ($\Delta t = t_{j+1} - t_j \neq \text{constant}$) and a generalization of the N-point DFT [68]:

$$F : T_X(\omega) = \left(\frac{N}{2} \right)^{\frac{1}{2}} \sum_{j=0}^{N-1} X(t_j)[A\cos(\omega t_j) - iB\sin(\omega t_j)] \tag{3.5}$$

where $i = \sqrt{-1}$, j is the summation index, and A and B are as yet unspecified functions of the angular frequency ω. This angular frequency may depend on the

21. Due to the additive nature of the signal and the errors in measuring it, the errors are often referred to as noise.

vector of sample times, $\{t_j\}$, but not on the data, $\{X(t_j)\}$, nor on the summation index j. The corresponding (normalized) periodogram is then

$$P_X(\omega) = \frac{1}{N}|FT_X(\omega)|^2 = \frac{A^2}{2}\left[\sum_j X(t_j)\cos(\omega t_j)\right]^2 + \frac{B^2}{2}\left[\sum_j X(t_j)\sin(\omega t_j)\right]^2 \tag{3.6}$$

If $A = B = \left(\frac{2}{N}\right)^{\frac{1}{2}}$, (3.5) and (3.6) reduce to the classical definitions [(3.3) and (3.4)] For even sampling (Δt = constant) FT_X reduces to the DFT and in the limit $\Delta t \to 0, N \to \infty$, it is proportional to the Fourier transform. Scargle [68] shows how (3.6) is not unique and further conditions must be imposed in order to derive the corrected expression for the LSP:

$$P_N(\omega) \equiv \frac{1}{2\sigma^2}\left\{\frac{\left[\sum_j (x_j - \overline{x})\cos(\omega(t_j - \tau))\right]^2}{\sum_j \cos^2(\omega(t_j - \tau))} + \frac{\left[\sum_j (x_j - \overline{x})\sin(\omega(t_j - \tau))\right]^2}{\sum_j \sin^2(\omega(t_j - \tau))}\right\} \tag{3.7}$$

where $\tau \equiv \tan^{-1}\left(\frac{\sum_j \sin(2\omega t_j)}{2\omega\sum_j \cos(2\omega t_j)}\right)$. τ is an offset that makes $P_N(\omega)$ completely independent of shifting all the t_j's by any constant. This choice of offset makes (3.7) exactly the solution that one would obtain if the harmonic content of a data set, at a given frequency ω, was estimated by linear least-squares fitting to the model $x(t) = A\cos(\omega t) + B\sin(\omega t)$. Thus, the LSP weights the data on a *per-point* basis instead of weighting the data on a *per-time interval* basis. Note that in the evenly sampled limit ($\Delta t = t_{j+1} - t_j$ = constant), (3.7) reduces to the classical periodogram definition [67]. See [67–72] for mathematical derivations and further details. C and Matlab code (*lomb.c* and *lomb.m*) for this routine are available from PhysioNet [2, 70] and the accompanying book Web site [73]. The well-known numerical computation library *Numerical Recipes in C* [74] also includes a rapid FFT-based method for computing the LSP, which claims not to use interpolation (rather *extirpolation*), but an implicit interpolation is still performed in the Fourier domain. Other methods for performing spectral estimation from irregularly sampled data do exist and include the min-max interpolation method [75] and the well-known geostatistical technique of *krigging*[22] [76]. The closely related fields of missing data *imputation* [77] and *latent variable discovery* [78] are also appropriate routes for dealing with missing data. However, the LSP appears to be sufficient for HRV analysis, even with a low SNR [53].

22. Instead of weighting nearby data points by some power of their inverted distance, krigging uses the spatial correlation structure of the data to determine the weighting values.

3.7.5 Information Limits and Background Noise

In order to choose a sensible window size, the requirement of stationarity must be balanced against the time required to resolve the information present. The European and North American Task Force on standards in HRV [46] suggests that the shortest time period over which HRV metrics should be assessed is 5 minutes. As a result, the lowest frequency that can be resolved is $\frac{1}{300} \approx 0.003$ Hz (just above the lower limit of the VLF region). Such short segments can therefore only be used to evaluate metrics involving the LF and HF bands. The upper frequency limit of the highest band for HRV analysis is 0.4 Hz [51]. Since the average time interval for N points over a time T is $\Delta t_{\text{av}} = \frac{T}{N}$, then the average Nyquist frequency [68] is then $f_c' = \frac{1}{2\Delta t_{\text{av}}} = \frac{N}{2T}$. Thus, a 5-minute window ($T = 300$) with the Nyquist constraint of $\frac{N}{2T} \geq 0.4$ for resolving the upper frequency band of the HF region, leads to a lower limit on N of 240 beats (an average heart rate of 48 bpm if all beats in a 5-minute segment are used). Utilization of the LSP, therefore, reveals a theoretical lower information threshold for accepting segments of an RR tachogram for spectral analysis in the upper HF region. If RR intervals of at least 1.25 seconds (corresponding to an instantaneous heart rate of $\text{HR}_\text{i} = \frac{60}{\text{RR}_\text{i}} = 48$ bpm) exist within an RR tachogram, then frequencies up to 0.4 Hz do exist. However, the accuracy of the estimates of the higher frequencies is a function of the number of RR intervals that exist with a value corresponding to this spectral region. Tachograms with no RR intervals smaller than 1.25s ($\text{HR}_\text{i} < 48$ bpm) can still be analyzed, but there is no power contribution at 0.4 Hz.

This line of thought leads to an interesting viewpoint on traditional short-term HRV spectral analysis; interpolation adds extra (erroneous) information into the time series and pads the FFT (in the time domain), tricking the user into assuming that there is a signal there, when really, there are simply not enough samples within a given range to allow the detection of a signal (in a statistically significant sense). Scargle [68] shows that at any particular frequency, f, and in the case of the null hypothesis, $P_X(\omega)$, has an exponential probability distribution with unit mean. Therefore, the probability that $P_X(\omega)$ will be between some positive value z and dz is $e^{-z}dz$, and hence, for a set of M independent frequencies, the probability that none give values larger than z is $(1 - e^{-z})^M$. The false alarm probability of the null hypothesis is therefore

$$P(> z) \equiv 1 - (1 - e^{-z})^M \tag{3.8}$$

Equation (3.8) gives the significance level for any peak in the LSP, $P_X(\omega)$ (a small value, say, $P < 0.05$ indicates a highly significant periodic signal at a given frequency). M can be determined by the number of frequencies sampled and the number of data points, N (see Press et al. [69]). It is therefore important to perform this test on each periodogram before calculating a frequency-based HRV metric, in order to check that there really are measurable frequencies that are not masked by noise or nonstationarity. There is one further caveat: Fourier analysis assumes that the signals at each frequency are independent. As we shall see in the next chapter on modeling, this assumption may be approximately true at best, and in some cases the coupling between different parts of the cardiovascular system may render Fourier-based spectral estimation inapplicable.

3.7.5.1 A Note on Spectral Leakage and Window Carpentry

The periodogram for unevenly spaced data allows two different forms of spectral adjustment: the application of time-domain (data) windows through weighting the signal at each point, and adjustment of the locations of the sampling times. The time points control the power in the window function, which leaks to the Nyquist frequency and beyond (the aliasing), while the weights control the side lobes. Since the axes of the RR tachogram are intricately linked (one is the first difference of the other), applying a windowing function to the amplitude of the data implicitly applies a nonlinear stretching function to the sample points in time. For an evenly sampled stationary signal, this distortion would affect all frequencies equally. Therefore, the reductions in LF and HF power cancel when calculating the $\frac{LF}{HF}$-ratio. For an irregularly sampled time series, the distortion will depend on the distribution of the sampling irregularity. A windowing function is therefore generally not applied to the irregularly sampled data. Distortion in the spectral estimate due to edge effects will not result as long as the start and end point means and first derivatives do not differ greatly [79].

3.7.6 The Effect of Ectopy and Artifact and How to Deal with It

To evaluate the effect of ectopy on HRV metrics, we can add artificial ectopic beats to an RR tachogram using a simple procedure. Kamath et al. [80] define ectopic beats (in terms of timing) as those which have intervals less than or equal to 80% of the previous sinus cycle length. Each datum in the RR tachogram represents an interval between two beats and the insertion of an ectopic beat therefore corresponds to the replacement of two data points as follows. The nth and $(n+1)$th beats (where n is chosen randomly) are replaced (respectively) by

$$RR'_n = \gamma\, RR_{n-1} \tag{3.9}$$

$$RR'_{n+1} = RR_{n+1} + RR_n - RR'_n \tag{3.10}$$

where the ectopic beat's timing is the fraction, γ, of the previous RR interval (initially 0.8). Note that the ectopic beat must be introduced at random within the central 50% of the 5-minute window to avoid windowing effects. Table 3.2 illustrates the effect of calculating the LF, HF, and $\frac{\text{LF}}{\text{HF}}$-ratio HRV metrics on an artificial RR tachogram with a known $\frac{\text{LF}}{\text{HF}}$-ratio (0.64) for varying levels of ectopy (adapted from [53]). Note that increasing levels of ectopy lead to an increase in HF power and a reduction in LF power, significantly distorting the $\frac{\text{LF}}{\text{HF}}$-ratio (even for just one beat).

It is therefore obvious that ectopic beats must be removed from the RR tachogram. In general, FFT-based techniques require the replacement of the removed beat with a *phantom* beat at a location where one would have expected the beat to have occurred if it was a sinus beat. Methods for performing phantom beat replacement range from linear and cubic spline interpolation,[23] AR model prediction, segment removal, and segment replacement.

23. Confusingly, phantom beat replacement is generally referred to as interpolation. In this chapter, it is referred to as phantom beat insertion, to distinguish it from the mathematical methods used to either place the phantom beat, or resample the unevenly sampled tachogram.

Table 3.2 LSP Derived Frequency Metrics for Different Magnitudes of Ectopy (γ)

Metric →	$\frac{LF}{HF}$	LF	HF	γ
Actual Value ↓				
0.64	0.64	0.39	0.61	†
0.64	0.60	0.37	0.62	0.8
0.64	0.34	0.26	0.74	0.7
0.64	0.32	0.25	0.76	0.6
0.64	0.47	0.32	0.68	0.8 ‡

† indicates no ectopy is present.
‡ indicates two ectopic beats are present.
Source: [52].

Although more robust and promising model-based techniques have been used [81], Lippman et al. [82] found that simply removing the signal around the ectopic beat performed as well as these more complicated methods. Furthermore, resampling the RR tachogram at a frequency (f_s) below the original ECG ($f_{ecg} > f_s$) from which it is derived effectively shifts the fiducial point by up to $\frac{1}{2}(\frac{1}{f_s} - \frac{1}{f_{ecg}})$s. The introduction of errors in HRV estimates due to low sampling rates is a well-known problem, but the additive effect from resampling is underappreciated. If a patient is suffering from low HRV (e.g., because they have recently undergone a heart transplant or are in a state of coma) then the sampling frequency of the ECG must be higher than normal. Merri et al. [83], and Abboud et al. [84] have shown that for such patients a sampling rate of at least 1,000 Hz is required. Work by Clifford et al. [85] and Ward et al. [50] demonstrate that a sampling frequency of 500 Hz or greater is generally recommended (see Figure 4.9 and Section 4.3.2).

The obvious choice for spectral estimation for HRV is therefore the LSP, which allows the removal of up to 20% of the data points in an RR tachogram without introducing a significant error in an HRV metric [53]. Therefore, if no morphological ECG is available, and only the RR intervals are available, it is appropriate to employ an aggressive beat removal scheme (removing any interval that changes by more than 12.5% on the previous interval [86]) to ensure that ectopic beats are not included in the calculation. Of course, since the ectopic beat causes a change in conduction, and momentarily disturbs the sinus rhythm, it is inappropriate to include the intervals associated with the beats that directly follow an ectopic beat (see Section 3.8.3.1) and therefore, all the affected beats should be removed at this nonstationarity. As long as there is no significant change in the phase of the sinus rhythm after the run of affected beats, then the LSP can be used without seriously affecting the estimate. Otherwise, the time series should be segmented at the nonstationarity.

3.7.7 Choosing an Experimental Protocol: Activity-Related Changes

It is well known that clinical investigations should be controlled for drugs, age, gender, and preexisting conditions. One further factor to consider is the activity of the patient population group, for this may turn out to be the single largest confounder of metrics, particularly in HRV studies. In fact, some HRV studies may be doing little more than identifying the difference in activity between two

patient groups, something that can be more easily achieved by methods such as actigraphy, direct electrode noise analysis [87], or simply noting of the patient's activity using an empirical scale. Bernardi et al. [88] demonstrated that HRV in conscious patients (as measured by the $\frac{LF}{HF}$-ratio) changes markedly depending on a subject's activity. Their analysis involved measuring the ECG, respiration, and blood pressure of 12 healthy subjects, all aged around 29 years, for 5 minutes during a series of simple physical (verbal) and mental activities. Despite the similarity in subject physiology and physical activity (all remained in the supine position for at least 20 minutes prior to, and during the recording), the day-time $\frac{LF}{HF}$-ratio had a strong dependence on mental activity, ranging from 0.7 for controlled breathing to 3.6 for free talking. It may be argued that the changes in these values are simply an effect of changing breathing patterns (that modify the HF component). However, significant changes in both the LF component and blood pressure readings were also observed, indicating that the feedback loop to the central nervous system (CNS) was affected. The resultant change in HRV is therefore likely to be more than just a respiratory phenomenon.

Differences in mental as well as physical activity should therefore be minimized when comparing HRV metrics on an interpatient or intrapatient basis. Since it is probably impossible to be sure whether or not even a willing subject is controlling their thought processes for a few minutes (the shortest time window for traditional HRV metrics [46]), this would imply that HRV is best monitored while the subject is asleep, during which the level of mental activity can be more easily assessed.

Furthermore, artifact in the ECG is significantly reduced during sleep (because there is less physical movement by the subject) and the variation in $\frac{LF}{HF}$-ratio with respect to the mean value is reduced within a sleep state [52, 53, 72]. Sleep stages usually last more than 5 minutes [89], which is larger than the minimum required for spectral analysis of HRV [51]. Segmenting the RR time series according to sleep state basis should therefore provide data segments of sufficient length with minimal data corruption and departures from stationarity (which otherwise invalidate the use of Fourier techniques).

The standard objective scale for CNS activity during sleep was defined by Rechtschaffen and Kales [90], a set of heuristics known as the *R&K rules*. These rules are based partially on the frequency content of the EEG, assessed by expert observers over 30-second epochs. One of the five defined stages of sleep is termed dream, or rapid eye movement (REM), sleep. Stages 1–4 (light to deep) are non-REM (NREM) sleep, in which dreaming does not occur. NREM sleep can be further broken down into drowsy sleep (stage 1), light sleep, (stages 1 and 2), and deep sleep (stages 3 and 4), or slow wave sleep (SWS). Healthy humans cycle through these five sleep stages with a period of around 100 minutes, and each sleep stage can last up to 20 minutes during which time the cardiovascular system undergoes few changes, with the exception of brief arousals [89].

When loss of consciousness occurs, the parasympathetic nervous system begins to dominate with an associated rise in HF and decrease in $\frac{LF}{HF}$-ratio. This trend is more marked for deeper levels of sleep [91, 92]. PSDs calculated from 5 minutes of RR interval data during wakefulness and REM sleep reveal similar spectral components and $\frac{LF}{HF}$-ratios [92]. However, stage 2 sleep and SWS sleep exhibit a shift towards an increase in percentage contributions from the HF components

Table 3.3 $\frac{LF}{HF}$-Ratios During Wakefulness, NREM and REM Sleep

Activity →	*Awake*	*REM*	*NREM*
Condition ↓		*Sleep*	*Sleep*
Normal [92]	N/A	2→2.5	0.5→1
Normal [46]	3.9	2.7	1.7
Normal [91]	4.0 ± 1.4	3.1 ± 0.7	1.2 ± 0.4
CNS Problem [93]	N/A	3.5→5.5	2→3.5
Post-MI [91]	2.4 ± 0.7	8.9 ± 1.6	5.1 ± 1.4

Note: N/A = not available; Post-MI = a few days after myocardial infarction; CNS = noncardiac related problem. Results quoted from [46, 91–93].

(above 0.15 Hz) with $\frac{LF}{HF}$-ratio values around 0.5 to 1 in NREM sleep and 2 to 2.5 in REM sleep [92]. In patients suffering from a simple CNS but noncardiac related problem, Lavie et al. [93] found slightly elevated NREM $\frac{LF}{HF}$-ratio values of between 2 and 3.5 and between 3.5 and 5.5 for REM sleep. Vanoli et al. [91] report that myocardial infarction (MI) generally results in a raised overall $\frac{LF}{HF}$-ratio during REM and NREM sleep with elevated LF and $\frac{LF}{HF}$-ratio (as high as 8.9) and lower HF. Values for all subjects during wakefulness in these studies (2.4 to 4.0) lie well within the range of values found during sleep (0.5 to 8.9) for the same patient population (see Table 3.3). This demonstrates that comparisons of HRV between subjects should be performed on a sleep-stage specific basis.

Recent studies [52, 53] have shown that the segmentation of the ECG into sleep states and the comparison of HRV metrics between patients on a per-sleep stage basis increases the sensitivity sufficiently to allow the separation of subtly different patient groups (normals and sleep apneics[24]), as long as a suitable spectral estimation technique (the LSP) is also employed. In particular, it was found that deep sleep or SWS gave the lowest variance in the $\frac{LF}{HF}$-ratio both in an intrapatient and interpatient basis, with the fewest artifacts, confirming that SWS is the most stable of all the sleep stages. However, since certain populations do not experience much SWS, it was found that REM sleep is an alternative (although slightly more noisy) state in which to compare HRV metrics. Further large-scale studies are required to prove that sleep-based segmentation will actually provide patient-specific assessments from HRV, although recent studies are promising.

3.8 Dealing with Nonstationarities

It should be noted at this point that all of the traditional HRV indices employ techniques that assume (weak) stationarity in the data. If part of the data in the window of analysis exhibits significant changes in the mean or variance over the length of the window, the HRV estimation technique can no longer be trusted. A cursory analysis of any real RR tachogram reveals that shifts in the mean or variance are a frequent occurrence [94]. For this reason it is common practice to *detrend* the signal by removing the linear or parabolic baseline trend from the window prior to calculating a metric.

24. Even when all data associated with the apneic episodes were excluded.

However, this detrending does not remove any changes in variance over a stationarity change, nor any changes in the spectral distribution of component frequencies. It is not only illogical to attempt to calculate a metric that assumes stationarity over the window of interest in such circumstances, it is unclear what the meaning of a metric taken over segments of differing autonomic tone could be. Moreover, changes in stationarity of RR tachograms are often joined by transient sections of heart rate overshoot and an accompanying increased probability of artifact on the ECG (and hence missing data) [86, 95].

In this section we will explore a selection of methods for dealing with nonstationarities, including multiscale techniques, detrending, segmentation (both statistically and from a clinical biological perspective), and the analysis of change points themselves.

3.8.1 Nonstationary HRV Metrics and Fractal Scaling

Empirical analyses employing detrending techniques can lead to metrics that appear to distinguish between certain patient populations. Such techniques include multiscale power analysis such as detrended fluctuation analysis (DFA) [96, 97]. Such techniques aid in the quantification of long-range correlations in a time series, and in particular, the *fractal scaling* of the RR tachogram. If a time series is self-similar over many scales, then the $log - log$ power-frequency spectrum will exhibit a $1/f^{\beta}$ scaling, where β is the slope of the spectrum. For a white noise process the spectrum is flat and $\beta = 0$. For pink noise processes, $\beta = 1$, and for Brownian processes, $\beta = 2$. Black noise has $\beta > 2$.

DFA is an alternative variance-based method for measuring the fractal scaling of a time series. Consider an N-sample time series x_k, which is integrated to give a time series y_k that is divided into boxes of equal length, m. In each box a least squares line fit is performed on the data (to estimate the trend in that box). The y coordinate of the straight line segments is denoted by $y_k^{(m)}$. Next, the integrated time series, y_k, is detrended by subtracting the local trend, $y_k^{(m)}$, in each box. The root-mean-square fluctuation of this integrated and detrended time series is calculated by

$$F(m) = \sqrt{\frac{1}{N}\left(\sum_{k=1}^{N}\left[y_k - y_k^{(m)}\right]^2\right)} \tag{3.11}$$

This computation is repeated over all time scales (box sizes) to characterize the relationship between $F(m)$, the average fluctuation, as a function of box size. Typically, $F(m)$ will increase with box size m. A linear relationship on a $log - log$ plot indicates the presence of power law (fractal) scaling. Under such conditions, the fluctuations can be characterized by a scaling exponent α, the slope of the line relating $\log F(m)$ to $\log m$, that is, $F(m) \sim m^{\alpha}$.

A direct link between DFA and conventional spectral analysis techniques and other fractal dimension estimation techniques exists [98–101]. These techniques include semivariograms (to estimate the Hausdorf dimension, H_a, [98]), the rescaled range (to estimate the Hurst exponent, H_u [98, 102]), wavelet transforms

(to estimate the variance of the wavelets H_w [100, 103]), the Fano factor α_F [102, 104], and the Allan factor, α_A [102]. Their equivalences can be summarized as [105]

$$\begin{aligned} \beta &= 2\alpha - 1 \\ \beta &= 2H_a + 1 \\ \beta &= 2H_u - 1 \\ \beta &= H_w \\ \beta &= \alpha_F \\ \beta &= \alpha_A \end{aligned} \tag{3.12}$$

However, it is interesting to note that each of these fractal measures has limited ranges of applicability and suffer from differing problems [106]. In particular, the Fano factor is unsuitable for estimating $\beta > 1$, and the Allan factor (a ratio of the variance to the mean) is confined to $0 < \beta < 3$ [106]. Recently McSharry et al. [100] performed an analysis to determine the sensitivity of each of these metrics for determining fractal scaling in RR interval time series. They demonstrated that for a range of colored Gaussian and non-Gaussian processes ($-2 < \beta < 4$), H_w provided the best fractal scaling range ($-2 < \beta < 4$ for Gaussian and $-0.8 < \beta < 4$ for non-Gaussian processes).

3.8.1.1 Multiscale Entropy

Multiscale entropy (MSE) is a nonlinear variant of these multiscale metrics that uses an entropy-based metric known as the sample entropy.[25] For a time series of N points, $\{u(j) : 1 \leq j \leq N\}$ forms the $N - m + 1$ vectors $x_m(i)$ for $\{i|1 \leq i \leq N - m + 1\}$, where $x_m(i) = u(i + k) : 0 \leq k \leq m - 1$ is the vector of m data points from $u(i)$ to $u(i + m - 1)$. If A_i is the number of vectors $x_{m+1}(j)$ within a given tolerance r of $x_{m+1}(i)$, B_i is the number of vectors $x_m(j)$ within r of $x_m(i)$ and $B(0) = N$, is the length of the input series, the sample entropy is given by

$$\text{SampEn}(k, r, N) = -\ln \frac{A(k)}{B(k-1)} \qquad (k = 0, 1, \ldots, m-1) \tag{3.13}$$

Sample entropy is the negative natural logarithm of an estimate of the conditional probability that subseries (epochs) of length m that match point-wise within a tolerance r also match at the next point.

The algorithm for calculating sample entropy over many scales builds up runs of points matching within the tolerance r until there is not a match, and keeps track of template matches in counters $A(k)$ and $B(k)$ for all lengths k up to m. Once all the matches are counted, the sample entropy values are calculated by $\text{SampEn}(k, r, N) = -ln(\frac{A(k)}{B(k-1)})$ for $k = 0, 1, \ldots, m-1$ with $B(0) = N$, the length of the input series.

25. Sample entropy has been shown to be a more accurate predictor of entropy in the RR tachogram than other traditional entropy estimation methods.

MSE does not change linearly with scale and therefore cannot be quantified by one exponent. In general, MSE increases (nonlinearly) with increasing N (or decreasing scale factor), reflecting the reduction in long-term coherence at longer and longer scales (shorter scale factors). This metric has been shown to be an independent descriptor of HRV to the fractal scaling exponent β [95]. An open-source implementation of this algorithm can be found on the PhysioNet Web site [107].

3.8.2 Activity-Related Changes

3.8.2.1 Segmentation of the Cardiac Time Series

Another possibility when dealing with nonstationarities is to simply segment the time series at an identifiable point of change and analyze the segments in isolation.[26] An early approach by Moody [110] involved a metric of nonstationarity that included mean heart rate and HRV. Later, Fukada et al. [111] used a modified t-test to identify shifts in the mean RR interval. If we assume that the RR tachogram is a series of approximately stationary states and measure the distribution of the frequency of length and size of the switching between states, we find that the distributions approximately fit specific power laws which vary depending on a subject's condition. Fukada et al. [111] achieved the segmentation of the time series by performing a t-test[27] to determine the most significant change in the mean RR interval. This process is repeated in a recursive manner on each bisection until the statistics of small numbers prevents any further divisions. One interesting result of the application of this method to the RR tachogram is the discovery that the scaling laws differ significantly depending on whether a subject is asleep or not. It is unclear if this is a reflection of the fact that differing parts of the human brain control these two major states, but the connections between the cardiovascular system and the mechanisms that control the interplay between sleep and arousals is rapidly becoming a research field of great interest [108, 109, 112, 113]. Models that reproduce this activity in a realistic manner are detailed in Chapter 4.

Unfortunately, empirical methods for segmenting the cardiac time series based purely on the RR tachogram have shown limited success and more detailed information is often needed. It has recently been shown [52] that by quantifying HRV only during sleep states, and comparing HRV between patients *only for* a particular sleep state, the sensitivity of a particular HRV metric is significantly increased. Furthermore, the deeper the sleep state, the more stationary the signal, the lower the noise, and the more sensitive is the HRV metric. Another method for segmenting the cardiac time series into active and inactive regions is based upon the work of Mietus et al. for quantifying sleep patterns from the ECG [87, 114].

26. Or some property derived from the frequency distribution of the means or lengths of the segments [108, 109].
27. The t-test is modified to account for the fact that each sample is not independent. This may not actually be necessary if each state is independent, although the success of modeling 24-hour fluctuations with hidden Markov models may indicate that there is some correlation between states, at least in the short term. However, the success of simple t-tests demonstrate that independence may be a reasonable approximation under certain circumstances [86].

3.8.2.2 Sleep Staging from the ECG

Respiratory rate may be derived from the body surface ECG by measuring the fluctuation of the mean cardiac electrical axis [115] or peak QRS amplitudes which accompany respiration. This phenomenon is known as ECG-derived respiration (EDR); see Chapter 8 for an in-depth analysis of this technique. The changes in the sequence of RR intervals during RSA are also heavily correlated with respiration through neurological modulation of the SA node. However, since the QRS morphology shifts due to respiration are mostly mechanically mediated, the phase difference between the two signals is not always constant. Recently Mietus et al. [114] demonstrated that by tracking changes in this coupling through cross-spectral analysis of the EDR and RSA time series, they were able to quantify the type and depth of sleep that humans experience into cyclic alternating pattern (CAP) and non-CAP sleep (rather than the traditional Rechtschaffen and Kales [90] scoring).

Following [114], frequency coupling can be measured using the cross-spectral density between RSA and EDR. There are two slightly different measures: coupling frequency with respect to magnitude of the sinusoidal oscillations $A(f)$ and the consistency in phase of the oscillations $\Theta(f)$. These are calculated separately such that

$$A(f) = \mathcal{E}\left[|P^i_{xy}(f)|^2\right] \tag{3.14}$$

and

$$\Theta(f) = \left|\mathcal{E}[P^i_{xy}(f)]\right|^2 \tag{3.15}$$

where $\mathcal{E}[.]$ denotes averaging across all the $i = 1, \ldots, N$ segments and $P^i_{xy}(f)$ is the cross-periodogram of the ith segment.

In general, $P_{xy}(f)$ is complex even if $X(t)$ and $Y(t)$ are real. Since $A(f)$ is calculated by taking the magnitude squared of $P_{xy}(f)$ in each block followed by averaging, it corresponds to the frequency coupling of the two signals due to the oscillations in amplitude only. Similarly, since $\Theta(f)$ is computed by first averaging the real and imaginary parts of $P_{xy}(f)$ across all blocks followed by magnitude squaring, it measures the consistency in phase of the oscillations across all blocks. $A(f)$ and $\Theta(f)$ are normalized and multiplied together to obtain the *cardiorespiratory coupling* (CRC), a measure of the strength of coupling between RSA and EDR as follows:

$$CRC(f) = \frac{A(f)}{\max[A(f)]} * \frac{\Theta(f)}{\max[\Theta(f)]} \tag{3.16}$$

CRC ranges between 0 and 1 with a low CRC indicating poor coupling and therefore increased activity. A high CRC (>0.4) indicates decreased activity that can be interpreted as sleep or sometimes sedation [87]. A value closer to 1 means strong coupling of RSA and EDR at a given frequency. It should be noted that this method,

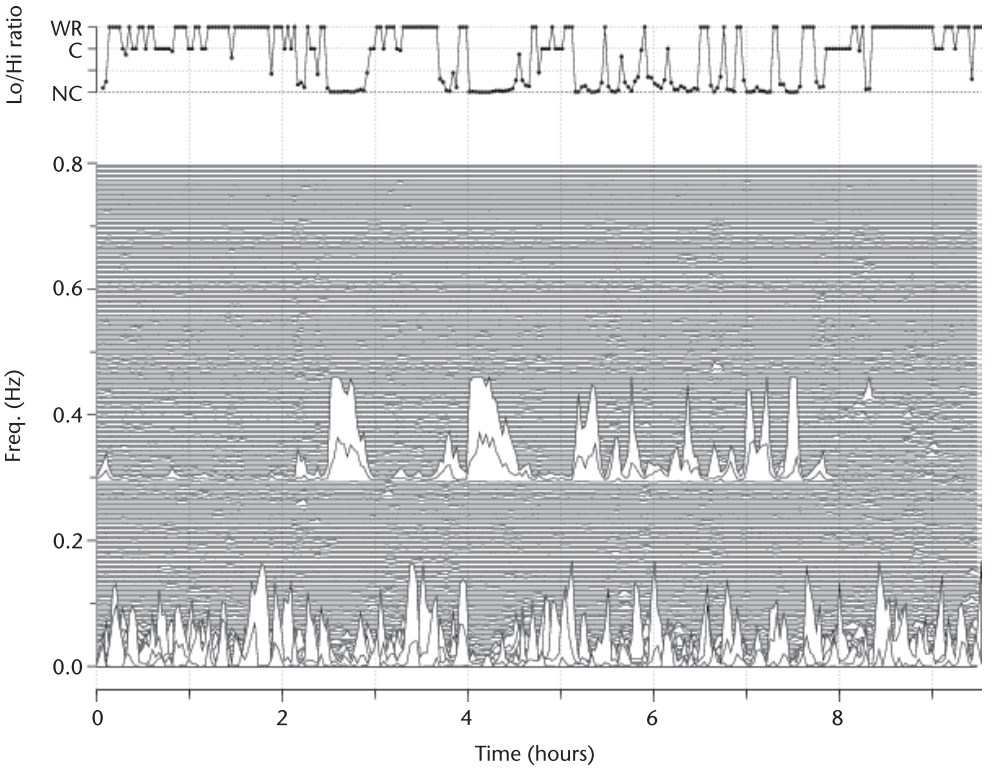

Figure 3.10 Spectrogram (lower) of EDR-RSA coherence, with associated sleep stability (upper) derived from thresholding a ratio of low frequency and high frequency (Lo/Hi) power for each segment. Note that this patient is on a ventilator so the respiratory frequency is sharp and fixed at 0.3 Hz. Stable non-CAP (NC) sleep is still observable despite the ventilation. (© 2005 J. Mietus. Reprinted with permission.)

is a slight modification of the one used in [114] (called cardiopulmonary coupling, or CPC), where the squaring of the phase is taken before the averaging.[28]

Figure 3.10 illustrates the application of this technique to a sedated and ventilated patient in an intensive care unit. The lower plot is a spectrogram; a time series of the cross spectral density between the EDR and RSA. The upper plot represents a stability of sleep from WR (wakefulness or REM sleep) to CAP sleep (C) to stable non-CAP (NC) sleep. This time series is derived by thresholding the ratio of the low (Lo) to high (Hi) frequency regions of the cross-spectral coherence. Note that despite the fact that this patient is ventilated (and hence the respiratory frequency is sharp and fixed at 0.3 Hz), stable (NC) sleep is still observable.

Coupling between RSA and EDR is more evident or easily obtainable when the subject is at rest (or in stable sleep, or perhaps, deep sleep) where there are fewer factors that may significantly influence changes in the respiratory rate or heart

28. These differences do not lead to significant differences in the metric as a predictor of stable (coupled high frequency) activity however. Furthermore, in CPC, the cross-power is thresholded at different frequencies to produce an output of wakefulness/REM sleep (WR), unstable/cyclic alternating pattern (CAP) sleep, or stable/non-CAP (NC) sleep. NC sleep is correlated with low sedation/agitation (Riker) levels [87, 116], and WR is correlated with medium to high agitation (Riker) scores.

rate. Therefore, this technique has also been employed to detect changes in activity or stationarity in patients [87]. Furthermore, the strongest coupling frequency is directly correlated with respiration, which is also a good index of activity, as well as an estimate of the prevailing respiratory rate. A sensitivity analysis of this technique also shows that the CPC metric is extremely robust to noise [87], since the presence of noise on the ECG is correlated with changes in activity [86].

It should be noted that the analysis of synchronization between the cardiac cycle and the respiratory frequency has been an area of interest for few years now [117], with promising results for determining the health of certain patient groups.

3.8.3 Perturbation Analysis

An alternative to detrending, or segmenting, the cardiac time series at nonstationary changes and analyzing the segments in isolation, is to perform an analysis of the ephemeral changes in the signal at the point of change. This type of *perturbation analysis* is a standard technique in clinical medicine and includes stress testing[29] and the Valsalva Maneuver.[30]

However, interventionist tests may not always be possible or appropriate, and a more passive analysis of change points is sometimes required. These include analyzing the periodicity of these changes, and the transient changes due to passive intrinsic shifts in cardiac activity, such as transitions between sleep states, changes due to arousals during sleep, or changes due to ectopy.

3.8.3.1 Heart Rate Turbulence

The changes in quasi-stationarity of the sinus rhythm due to the biphasic physiological change in SA node activity from PVCs is known as heart rate turbulence[31] (HRT) [118, 119]. In HRV analysis, this disturbance is removed from the RR tachogram, and an assumption is made that the phase of the RR tachogram is unchanged after the signal returns to the "undisturbed rhythm." In HRT, the changes in the "disturbed" section are analyzed to extract metrics to quantify the changes.

In general, HRT manifests as a short initial acceleration of the heart rate for a few beats, followed by a deceleration back to the basal value from before the PVC. HRT is usually quantified by two numerical parameters: turbulence onset (TO) and turbulence slope (TS). TO is defined as the percentage difference between the average value of the first two normal RR intervals following the PVC ($RR^n, n = 2, 3$) and of the last two normal intervals preceding the PVC ($RR^{-n}, n = 2, 1$) and is given by [119]:

$$TO = \frac{(RR^{+2} + RR^{+3}) - (RR^{-2} + RR^{-1})}{RR^{-2} + RR^{-1}} \times 100 \tag{3.17}$$

29. A series of exercise tests that attempt to induce heart-related problems which manifest on the ECG at high heart rates or due to strong sympathetic innervation.
30. A pressure-inducing respiratory procedure which is thought to provide a rough guide to the integrity of the autonomic neural pathways involved in the response [51].
31. It should be noted that the term *turbulence* is a misnomer, since there is no strict evidence of actual turbulence in the neural modulation, electrophysical activity, or the resultant hemodynamic flow. A more appropriate term may be heart rate perturbation (HRP).

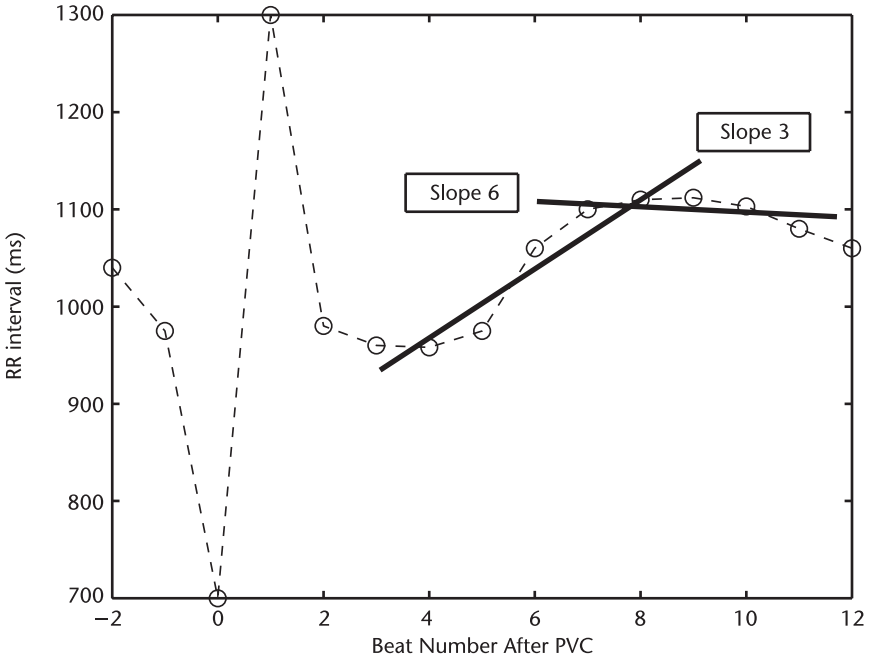

Figure 3.11 Example of HRT disturbance after PVC. Beat index denotes the RR interval associated with the beats surrounding the PVC (index zero). Note that the following beat after the PVC (RR^{+1}) has an associated elongated RR interval; the compensatory pause after the early (premature) preceding beat. Slopes are calculated for each following set of five consecutive beats (only slopes 3 and 6 are illustrated for clarity). (*After:* [120].)

where RR^{-2} and RR^{-1} are the first two normal intervals preceding the PVC and RR^{+2} and RR^{+3} the first two normal intervals following the PVC. (Note that there are two intervals associated with the PVC; the normal-PVC interval, RR^{0}, and the following PVC-normal interval RR^{+1}.) Therefore, positive TO values indicate deceleration and negative TO values indicate acceleration of the postectopic sinus rhythm. Although the TO can be determined for each individual PVC (and should be if performed online), TO has shown to be more (clinically) predictive if the average value of all individual measurements is calculated. For the calculation of the mean TO in the current version (1.11) of the freely available HRT-algorithm [121], at least 15 normal intervals after each single PVC are required.

TS is calculated by constructing an average ectopy-centered time series[32] and determining the steepest (linearly regressed) slope each possible sequence of five consecutive normal intervals in the post-PVC "disturbed" tachogram (usually taken to be up to RR^{+16}; the first 15 normal RR intervals). That is, all possible slopes from the set $\{RR^{+2,+3,\ldots,+6}, \ldots, RR^{+17,+18,\ldots,+16}\}$ are calculated, and the maximum of the 13 possible slopes is taken to be the TS. These values are usually averaged over all acceptable candidates, and can run into the tens of thousands for 24-hour Holter recordings. Figure 3.11 illustrates the calculation of TS after the disturbance

32. This averaging reduces the short-term noise from RSA and it is wise to use 20 or more ectopic examples to increase the statistical power of the technique. However, HRT estimation can be performed on as few as a single example.

in the RR tachogram due to a PVC. Note that the following beat after the PVC (RR^{+1}) has an associated elongated RR interval, the compensatory pause after the early (premature) preceding beat. Although slopes are calculated for each following set of five consecutive beats, only the third and sixth slopes are illustrated for clarity.

The criteria for excluding PVCs must be stringent, since HRT quantification can only deliver usable results if the triggering event was a true PVC (and not an artifact, T wave or similar non-PVC event). In addition, the sinus rhythm immediately preceding and following the PVC must be checked to ensure that it is free from arrhythmia, artifacts, and false beat classifications due to artifact. A useful set of exclusion criteria are:

- Remove all RR intervals < 300 ms or $> 2,000$ ms;
- Remove all RR^n where $\|RR^{n-1} - RR^n\| > 200$ ms;
- Remove all RR intervals that change by more than $> 20\%$ with respect to the mean of the five last sinus intervals (the reference interval);
- Only use PVCs with a minimum prematurity of 20%;
- Exclude PVCs with a postextrasystole interval which is at least 20% longer than the normal interval.

TO values below 0 and TS values above 2.5 are considered normal, and abnormal otherwise (i.e., a healthy response to PVCs is a strong sinus acceleration followed by a rapid deceleration).

Although the exact mechanism that leads to HRT is unknown [120], it is thought to be an autonomous baroreflex whereby the PVC causes a brief disturbance of the arterial blood pressure [122]. Therefore, HRT can be measured around the point of atrial ectopy as well as ventricular ectopy. When the autonomic control system is healthy, this rapid change causes an instantaneous response in the following RR intervals. If the autonomic control system is impaired, the magnitude of this response is either diminished or possibly completely absent. From a clinical perspective, HRT metrics have been shown to be predictive of mortality and possibly cardiac arrest. Furthermore, HRT metrics have been shown to be particularly useful when the patients are taking beta-blockers. In one study [123], combined TO and TS was found to be the only independent predictor of mortality compared to otherwise predictive markers such as mean HR, previous myocardial infarction, and a low ejection fraction.

3.8.3.2 QT Turbulence and Other Measures

Many other studies indicate that this technique is useful [121], and in particular the application of HRT analysis to study of the QT interval variation after PVCs has shown that QT interval turbulence occurs in association with HR turbulence [124]. QT TO has been defined as the relative difference between the QT interval of the first sinus cycle after an induced PVC and the mean of the QT intervals of the two sinus cycles preceding the premature beat. Savelieva et al. [124] found that patients with ischemic VT and left ventricular dysfunction exhibited significantly lower QT TO values than those with nonischemic VT and normal ventricular function. Interestingly, neither induced APCs nor PVCs produced late QT dynamics equivalent to TS of HR turbulence, and the underlying mechanism of QT turbulence

remains unclear. However, few studies in QT turbulence have been performed and difficulties in measuring and interpreting the QT interval may confound this type of analysis (see Section 3.4.1 and Chapter 11). It is therefore uncertain whether QT TS will prove to be a useful metric. Furthermore, many HRT-related metrics have been proposed with varying degrees of success, including:

- *Turbulence dynamics* [125]: The largest regressed slope between TS and heart rate in a particular individual over a sliding window of 10 bpm;
- *Turbulence timing* [126]: The beat number of the 5-RR interval sequence used in the TS calculation at which the largest slope occurs;
- *Correlation coefficient of TS* [127]: The correlation coefficient of the regression line fitted to the 5-RR intervals giving the maximum slope (i.e., where TS is defined);
- *Turbulence jump* [128]: The maximum difference between adjacent RR intervals;
- *Turbulence frequency decrease* [129]: A frequency-domain metric obtained by fitting a sine function to the postcompensatory pause.

Although many of these metrics have shown promise as independent disease predictors, further studies are required to determine whether any of these metrics provide significantly superior risk stratification to TO and TS. Furthermore, TO and TS are simpler to measure, and have been validated in large prospective studies. Further information is available from Schneider et al. [119, 121, 130].

3.9 Summary

Given the finite number of pages allowed in this book, all the possible metrics that researchers have chosen to quantify the EGG (and derived beat-to-beat timing sequences) cannot be detailed, nor all the possible applications of each metric be covered. Moreover, such a summary would be a transient document of only partial relevance in a few years from publication.

Rather, this chapter is intended to introduce the reader to the many different linear stationary and nonstationary qualities of the ECG, together with a selection of relevant metrics for evaluating these properties. This chapter is also intended to give an insight into possible approaches which are relevant to the different recording situations one may encounter (that vary based on activity, demographics and medical condition). It is important, although difficult, to differentiate between the concept of the nonstationary and the nonlinear nature of the ECG, since the application of a particular methodology or model will depend on prior beliefs concerning the relevance of these paradigms. In general, it is sufficient to apply linear techniques to quasi-stationary segments of ECG. However, to improve on such measures, it is important to understand the nonlinear nature of the ECG. Techniques for nonlinear analysis are therefore presented in Chapter 6. However, before reading the chapter on nonlinear analysis, it is advisable to read the following chapter, in which a selection of practically useful nonlinear and nonstationary models for both the ECG and the RR tachogram are presented.

References

[1] Welch, P. D., "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging over Short, Modified Periodograms," *IEEE Trans. Audio Electroacoust.*, Vol. AU-15, June 1967, pp. 70–73.

[2] Goldberger, A. L., R. G. Mark, and G. B. Moody, "PhysioNet: The Research Resource for Complex Physiologic Signals," http://www.physionet.org.

[3] Pettersson, J., O. Pahlm, and E. Carro, "Changes in High-Frequency QRS Components Are More Sensitive Than ST-Segment Deviation for Detecting Acute Coronary Artery Occlusion," *J. Am. Coll. Cardiol.*, Vol. 36, 2000, pp. 1827–1834.

[4] Schlegel, T. T., et al., "Real-Time 12-Lead High-Frequency QRS Electrocardiography for Enhanced Detection of Myocardial Ischemia and Coronary Artery Disease," *Mayo Clin Proc.*, Vol. 79, 2004, pp. 339–350.

[5] Spackman, T. N., M. D. Abel, and T. T. Schlegel, "Twelve-Lead High-Frequency QRS Electrocardiography During Anesthesia in Healthy Subjects," *Anesth. Analg.*, Vol. 100, No. 4, 2005, pp. 1043–1047.

[6] Hunt, A. C., "T Wave Alternans in High Arrhythmic Risk Patients: Analysis in Time and Frequency Domains: A Pilot Study," *BMC Cardiovasc. Disord.*, Vol. 2, No. 6, March 2002.

[7] Jarvis, M. R., and P. P. Mitra, "Apnea Patients Characterized by 0.02 Hz Peak in the Multitaper Spectrogram of Electrocardiogram Signals," *Computers in Cardiology*, Vol. 27, 2000, pp. 769–772.

[8] Zeng, W., and L. Glass, "Statistical Properties of Heartbeat Intervals During Atrial Fibrillation," *Phys. Rev. E*, 1996, pp. 1779–1784.

[9] Guyton, A. C., and J. E. Hall, *Textbook of Medical Physiology*, Philadelphia, PA: W. B. Saunders Company, 2001.

[10] Shouldice, R., et al., "Modulating Effect of Respiration on Atrioventricular Conduction Time Assessed Using PR Interval Variation," *Med. Biol. Eng. Comput.*, Vol. 40, No. 6, November 2002, pp. 609–617.

[11] Wagner, G. S., *Marriott's Practical Electrocardiography*, 9th ed., Baltimore, MD: Williams & Wilkins, 1994.

[12] Estes, N. A. M., et al., "ECG Findings in Active Patients; Differentiating the Benign from the Serious," *The Physician and Sportsmedicine*, Vol. 29, No. 3, 2001.

[13] Davey, P., "A New Physiological Method for Heart Rate Correction of the QT Interval," *Heart*, Vol. 82, 1999, pp. 183–186.

[14] Seed, W. A., et al., "Relation of Human Cardiac Action Potential Duration to the Interval Between Beats: Implications for the Validity of Rate Corrected QT Interval (QTc)," *British Heart Journal*, Vol. 57, 1987, pp. 32–37.

[15] Zabel, M., et al., "Assessment of QT Dispersion for Prediction of Mortality or Arrhythmic Events After Myocardial Infarction: Results of a Prospective, Long-Term Follow-Up Study," *Circulation*, No. 97, 1998, pp. 2543–2550.

[16] Coumel, P., P. Maison-Blanche, and F. Badilini, "Dispersion of Ventricular Repolarization: Reality? Illusion? Significance?" *Circulation*, No. 97, 1998, pp. 2491–2493.

[17] Rautaharju, P. M., "QT and Dispersion of Ventricular Repolarization: The Greatest Fallacy in Electrocardiography in the 1990s," *Circulation*, No. 99, 1999, pp. 2476–2479.

[18] di Bernardo, D., P. Langley, and A. Murray, "Effect of Changes in Heart Rate and in Action Potential Duration on the Electrocardiogram T Wave Shape," *Physiological Measurement*, Vol. 23, No. 2, 2002, pp. 355–364.

[19] Fletcher, G. F., et al., "Exercise Standards: A Statement for Healthcare Professionals from the American Heart Association," *Circulation*, Vol. 91, No. 2, 2001, pp. 580–615.

[20] Malik, M., et al., "Relation Between QT and RR Intervals Is Highly Individual Among Healthy Subjects: Implications for Heart Rate Correction of the QT Interval," *Heart*, Vol. 87, No. 3, 2002, pp. 220–228.

[21] Fossa, A. A., et al., "Dynamic Beat-to-Beat Modeling of the QT-RR Interval Relationship: Analysis of QT Prolongation During Alterations of Autonomic State Versus Human Ether A-Go-Go-Related Gene Inhibition," *J. Pharmacol. Exp. Ther.*, Vol. 312, No. 1, 2005, pp. 1–11.

[22] Lehtinen, R., "ST/HR Hysteresis: Exercise and Recovery Phase ST Depression/Heart Rate Analysis of the Exercise ECG," *J. Electrocardiol.*, No. 32, 1999, pp. 198–204.

[23] Lehtinen, R., "Diagnostic and Prognostic Value of ST/HR Hysteresis," *International Journal of Bioelectromagnetism*, Vol. 2, No. 1, 2000.

[24] Shouldice, R., C. Heneghan, and P. Nolan, "Methods of Quantifying Respiratory Modulation in Human PR Electrocardiographic Intervals," *IEEE Engineering in Medicine and Biology Conference*, IEEE Engineering in Medicine and Biology Conference, Houston, TX, October 2002.

[25] Moody, G. B., "PhysioNet: The Research Resource for Complex Physiologic Signals: Physiobank Annotations," http://www.physionet.org/physiobank/annotations.shtml.

[26] Dingfei, G., N. Srinivasan, and S. M. Krishnan, "Cardiac Arrhythmia Classification Using Autoregressive Modeling," *Biomed. Eng. Online*, Vol. 1, No. 5, November 2002.

[27] Owis, M. I., et al., "Study of Features Based on Nonlinear Dynamical Modeling in ECG Arrhythmia Detection and Classification," *IEEE Trans. Biomed. Eng.*, Vol. 49, No. 7, July 2002, pp. 733–736.

[28] Moody, G. B., and R. G. Mark, "QRS Morphology Representation and Noise Estimation Using the Karhunen-Loève Transform," *Computers in Cardiology*, Vol. 16, 1989, pp. 269–272.

[29] Mark, R. G., and G. B. Moody, "ECG Arrhythmia Analysis: Design and Evaluation Stratergies," Chapter 18, in I. Gath and G. F. Inbar, (eds.), *Advances in Processing and Pattern Analysis of Biological Signals*, New York: Plenum Press, 1996, pp. 251–272.

[30] Coast, D. A., et al., "An Approach to Cardiac Arrhythmia Analysis Using Hidden Markov Models," *IEEE Trans. Biomed. Eng.*, No. 37, 1990, pp. 826–836.

[31] Caswell, S. A., K. S. Kluge, and C. M. J. Chiang, "Pattern Recognition of Cardiac Arrhythmias Using Two Intracardiac Channels," *Computers in Cardiology*, Vol. 20, September 1993, pp. 181–184.

[32] Zhou, S. H., P. M. Rautaharju, and H. P. Calhoun, "Selection of a Reduced Set of Parameters for Classification of Ventricular Conduction Defects by Cluster Analysis," *Computers in Cardiology*, Vol. 20, 1993, pp. 879–882.

[33] Guvenir, H. A., et al., "Selection of a Reduced Set of Parameters for Classification of Ventricular Conduction Defects by Cluster Analysis," *Computers in Cardiology*, Vol. 24, 1997, pp. 433–436.

[34] Maglaveras, N., et al., "ECG Pattern Recognition and Classification Using Non-Linear Transformations and Neural Networks: A Review," *International Journal of Medical Informatics*, Vol. 52, No. 1, October 1998, pp. 191–208.

[35] Chazal, P., "Automatic Classification of the Frank Lead Electrocardiogram," Ph.D. dissertation, University of New South Wales, 1998.

[36] Barro, S., et al., "Algorithmic Sequential Decision-Making in the Frequency Domain for Life Threatening Ventricular Arrhythmias and Imitative Artefacts: A Diagnostic System," *J. Biomed. Eng.*, No. 11, 1989, pp. 320–328.

[37] Jekova, I., A. Cansell, and I. Dotsinsky, "Noise Sensitivity of Three Surface ECG Fibrillation Detection Algorithms," *Physiological Measurement*, Vol. 22, No. 2, 2001, pp. 287–297.

[38] Afonoso, V. X., and W. J. Tompkins, "Detecting Ventricular Fibrillation: Selecting the Appropriate Time-Frequency Analysis Tool for the Application," *IEEE Eng. Med. Biol. Mag.*, No. 14, 1995, pp. 152–159.

[39] Dokur, Z., T. Olmez, and E. Yazgan, "Comparison of Discrete Wavelet and Fourier Transforms for ECG Beat Classification," *Electronics Letters*, Vol. 35, No. 18, September 1999, pp. 1502–1504.

[40] deChazal, P., and R. B. Reilly, "A Comparison of the Use of Different Wavelet Coefficients for the Classification of the Electrocardiogram," *Proc. of the 14th International Conference on Pattern Recognition*, Barcelona, Spain, September 2000.

[41] Schulte-Frohlinde, V., et al., "Noise Effects on the Complex Patterns of Abnormal Heartbeats," *Phys. Rev. Lett.*, Vol. 87, No. 6, August 6, 2001, 068104.

[42] Schulte-Frohlinde, V., et al., "Complex Patterns of Abnormal Heartbeats," *Phys. Rev. E*, Vol. 66, No. 3, Pt. 1, September 2002, 031901.

[43] Schulte-Frohlinde, V., et al., "Heartprints: A Dynamical Portrait of Cardiac Arrhythmia," http://www.physionet.org/physiotools/heartprints/.

[44] deChazal, P., and R. B. Reilly, "Automatic Classification of ECG Beats Using Waveform Shape and Heart Beat Interval Features," *International Conference on Acoustics, Speech and Signal Processing (ICASSP'03)*, Hong Kong, Vol. 2, April 2003, pp. 269–272.

[45] Friesen, G. M., et al., "A Comparison of the Noise Sensitivity of Nine QRS Detection Algorithms," *IEEE Trans. Biomed. Eng.*, Vol. 37, No. 1, 1990, pp. 85–98.

[46] Malik, M., "Heart Rate Variability: Standards of Measurement, Physiological Interpretation, and Clinical Use," *Circulation*, Vol. 93, 1996, pp. 1043–1065.

[47] Kitney, R. I., and O. Rompelman, (eds.), *The Study of Heart Rate Variability*, London, U.K.: Oxford University Press, 1980.

[48] Yosefy, C., et al., "The Diagnostic Value of QRS Changes for Prediction of Coronary Artery Disease During Exercise Testing in Women: False-Positive Rates," *Coronary Artery Disease*, Vol. 15, No. 3, May 2004, pp. 147–154.

[49] Shouldice, R., C. Heneghan, and P. Nolan, "Methods of Quantifying Respiratory Modulation in Human pr Electrocardiographic Intervals," *Proc. 24th Annual IEEE International Conference on Engineering in Medicine and Biology (EMBS-BMES)*, October 2002.

[50] Ward, S., et al., "Electrocardiogram Sampling Frequency Errors in PR Interval Spectral Analysis," *Proc. IEEE PGBIOMED'04*, Southampton, U.K., August 2004.

[51] Malik, M., and A. J. Camm, *Heart Rate Variability*, Armonk, NY: Futura Publishing, 1995.

[52] Clifford, G. D., and L. Tarassenko, "Segmenting Cardiac-Related Data Using Sleep Stages Increases Separation Between Normal Subjects and Apnoeic Patients," *IOP Physiol. Meas.*, No. 25, 2004, pp. N27–N35.

[53] Clifford, G. D., and L. Tarassenko, "Quantifying Errors in Spectral Estimates of HRV Due to Beat Replacement and Resampling," *IEEE Trans. Biomed. Eng.*, Vol. 52, No. 4, April 2005.

[54] Mietus, J. E., et al., "The pNNx Files: Re-Examining a Widely Used Heart Rate Variability Measure," *Heart*, Vol. 88, No. 4, 2002, pp. 378–380.

[55] Grogan, E. L., et al., "Reduced Heart Rate Volatility: An Early Predictor of Death in Trauma Patients," *Annals of Surgery*, Vol. 240, No. 3, September 2004, pp. 547–556.

[56] Griffin, M. P., and J. R. Moorman, "Toward the Early Diagnosis of Neonatal Sepsis and Sepsis-Like Illness Using Novel Heart Rate Analysis," *Pediatrics*, Vol. 107, No. 1, 2001, pp. 97–104.

[57] Cerutti, S., A. M. Bianchi, and L. T. Mainardi, "Spectral Analysis of Heart Rate Variability Signal," in M. Malik and A. J. Camm, (eds.), *Heart Rate Variability*, Armonk, NY: Futura Publishing, 1995.

[58] Taylor, J. A., et al., "Mechanisms Underlying Very-Low-Frequency RR-Interval Oscillations in Humans," *Circulation*, Vol. 98, 1998, pp. 547–555.

[59] Serrador, J. M., H. C. Finlayson, and R. L. Hughson, "Physical Activity Is a Major Contributor to the Ultra Low Frequency Components of Heart Rate Variability," *Heart*, Vol. 82, No. e9, 1999, pp. 547–555.

[60] Galloway, D. G., and B. F. Womack, *An Application of Spectral Analysis and Digital Filtering to the Study of Respiratory Sinus Arrhythmia*, Technical Report 71, Bioengineering Research Lab, University of Texas, Electronics Research Center, Austin, TX, 1969.

[61] Sayers, B. McA., "Analysis of Heart Rate Variability," *Ergonomics*, Vol. 16, No. 1, 1969, pp. 17–32.

[62] Crawford, M. H., S. Bernstein, and P. Deedwania, "ACC/AHA Guidelines for Ambulatory ECG," *Circulation*, Vol. 100, 1999, pp. 886–893.

[63] Medical Predictive Science Corporation, 510(k) patent application no. K021230, FDA, 2003.

[64] Boston Medical Technologies, Inc., 510(k) patent application no. K010955, FDA, 2001, http://www.fda.gov/cdrh/pdf/k010955.pdf.

[65] Clayton, R. H., et al., "Comparison of Autoregressive and Fourier Transform Based Techniques for Estimating RR Interval Spectra," *Computers in Cardiology*, 1997, pp. 379–382.

[66] Harris, F. J., "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proc. IEEE*, Vol. 66, No. 1, January 1978.

[67] Lomb, N. R., "Least-Squares Frequency Analysis of Unequally Spaced Data," *Astrophysical and Space Science*, Vol. 39, 1976, pp. 447–462.

[68] Scargle, J. D., "Studies in Astronomical Time Series Analysis. ii. Statistical Aspects of Spectral Analysis of Unevenly Spaced Data," *Astrophysical Journal*, Vol. 263, 1982, pp. 835–853.

[69] Press, W. H., et al., *Numerical Recipes in C*, 2nd ed., Cambridge, U.K.: Cambridge University Press, 1992.

[70] Moody, G. B., "Spectral Analysis of Heart Rate Without Resampling," *Computers in Cardiology*, No. 20, 1993, pp. 715–718.

[71] Laguna, P., G. B. Moody, and R. G. Mark, "Power Spectral Density of Unevenly Sampled Data by Least-Square Analysis: Performance and Application to Heart Rate Signals," *IEEE Trans. Biomed. Eng*, Vol. BME-45, 1998, pp. 698–715.

[72] Clifford, G. D., "Signal Processing Methods for Heart Rate Variability," Ph.D. dissertation, University of Oxford, 2002.

[73] Clifford, G. D., F. Azuaje, and P. E. McSharry, "Advanced Tools for ECG Analysis," http://www.ecgtools.org/.

[74] Press, W. H., et al., *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed., Cambridge, U.K.: Cambridge University Press, 1992.

[75] Fessler, J. A., and B. P. Sutton, "Nonuniform Fast Fourier Transforms Using Min-Max Interpolation," *IEEE Trans. Sig. Proc.*, Vol. 51, No. 2, 2003, pp. 560–574.

[76] den Hertog, D., J. P. C. Kleijnen, and A. Y. D. Siem, "The Correct Kriging Variance Estimated by Bootstrapping," Discussion Paper 46, Tilburg University, Center for Economic Research, 2004, available at http://ideas.repec.org/p/dgr/kubcen/200446.html.

[77] Brown, M. L., and J. F. Kros, "The Impact of Missing Data on Data Mining," in John Wang, (ed.), *Data Mining: Opportunities and Challenges*, Hershey, PA: IRM Press, 2003, pp. 174–198.

[78] Vermunt, J. K., "Causal Log-Linear Modeling with Latent Variables and Missing Data," in U. Engel and J. Reinecke, (eds.), *Analysis of Change: Advanced Techniques in Panel Data Analysis*, Berlin/New York: Walter de Gruyter, 1996, pp. 35–60.

[79] Schmitz, A., and T. Schreiber, "Testing for Nonlinearity in Unevenly Sampled Time Series," *Phys. Rev. E*, Vol. 59, 1999, p. 4044.

[80] Kamath, M. V., and F. L. Fallen, "Correction of the Heart Rate Variability Signal for Ectopics and Missing Beats," in M. Malik and A. J. Camm, (eds.), *Heart Rate Variability*, Armonk, NY: Futura Publishing, 1995, pp. 75–85.

[81] Mateo, J., and P. Laguna, "Improved Heart Rate Variability Signal Analysis from the Beat Occurrence Times According to the IPFM Model," *IEEE Trans. Biomed. Eng.*, Vol. 47, No. 8, 2000. pp. 985–996.

[82] Lippman, N., K. M. Stein, and B. B. Lerman, "Comparison of Methods for Removal of Ectopy in Measurement of Heart Rate Variability," *Am. J. Physiol., Heart Circ. Physiol. 36*, Vol. 267, No. 36, 1994, pp. H411–H418.

[83] Merri, M., et al., "Sampling Frequency of the Electrocardiogram for Spectral Analysis of the Heart Rate Variability," *IEEE Trans. Biomed. Eng.*, Vol. 37, No. 1, January 1990, pp. 99–106.

[84] Abboud, S., and O. Barnea, "Errors Due to Sampling Frequency of Electrocardiogram in Spectral Analysis of Heart Rate Signals with Low Variability," *Computers in Cardiology*, Vol. 22, 1995, pp. 113–116.

[85] Clifford, G. D., and P. E. McSharry, "Method to Filter ECGs and Evaluate Clinical Parameter Distortion Using Realistic ECG Model Parameter Fitting," *Computers in Cardiology*, Vol. 32, 2005.

[86] Clifford, G. D., P. E. McSharry, and L. Tarassenko, "Characterizing Abnormal Beats in the Normal Human 24-Hour RR Tachogram to Aid Identification and Artificial Replication of Circadian Variations in Human Beat to Beat Heart Rate," *Computers in Cardiology*, Vol. 29, 2002, pp. 129–132.

[87] Clifford, G. D., et al., "Segmentation of 24-Hour Cardiovascular Activity Using ECG-Based Sleep/Sedation and Noise Metrics," *Computers in Cardiology*, Vol. 32, 2005.

[88] Bernardi, L., et al., "Effects of Controlled Breathing, Mental Activity and Mental Stress with or without Verbalization on Heart Rate Variability," *Journal of the American College of Cardiology*, Vol. 35, No. 6, May 2000, pp. 1462–1469.

[89] Lavie, P., *The Enchanted World of Sleep*, New Haven, CT: Yale University Press, 1996.

[90] Rechtschaffen, A., and A. Kales, *A Manual of Standardized Terminology, Techniques and Scoring System for Sleep Stages of Human Subjects*, Public Health Service, Washington, D.C.: U.S. Government Printing Office, 1968.

[91] Vanoli, V., et al., "Heart Rate Variability During Specific Sleep Stages," *Circulation*, Vol. 91, 1995, pp. 1918–1922.

[92] Otzenberger, H., et al., "Dynamic Heart Rate Variability: A Tool for Exploring Sympathovagal Balance Continuously During Sleep in Men," *Am. J. Physiol.: Heart Circ. Physiol.*, Vol. 275, 1998, pp. H946–H950.

[93] Lavie, P., A. Shlitner, and R. Nave, "Cardiac Autonomic Function During Sleep in Psychogenic and Organic Erectile Dysfunction," *J. Sleep Res.*, Vol. 8, 1999, pp. 135–142.

[94] Clifford, G. D., and P. E. McSharry, "Generating 24-Hour ECG, BP and Respiratory Signals with Realistic Linear and Nonlinear Clinical Characteristics Using a Nonlinear Model," *Computers in Cardiology*, Vol. 31, 2004, pp. 709–712.

[95] McSharry, P. E., and G. D. Clifford, "A Statistical Model of the Sleep-Wake Dynamics of the Cardiac Rhythm," *Computers in Cardiology*, Vol. 32, 2005.

[96] Peng, C. -K., et al., "Quantification of Scaling Exponents and Crossover Phenomena in Nonstationary Heartbeat Time Series," *Chaos*, Vol. 5, 1995, pp. 82–87.

[97] Goldberger, A. L., et al., "Fractal Dynamics in Physiology: Alterations with Disease and Ageing" *Proc. Natl. Acad. Sci.*, Vol. 99, 2002, pp. 2466–2472.

[98] Schroeder, M. R., *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise*, New York: W. H. Freeman, 1991.

[99] Heneghan, C., and G. McDarby, "Establishing the Relation Between Detrended Fluctuation Analysis and Power Spectral Density Analysis for Stochastic Processes," *Phys. Rev. E*, Vol. 62, No. 5, 2000.

[100] McSharry, P. E., and B. D. Malamud, "Quatifying Self-Similarity in Cardiac Inter-Beat Interval Time Series," *Computers in Cardiology*, Vol. 32, September 2005.

[101] Willson, K., and D. P. Francis, "A Direct Analytical Demonstration of the Essential Equivalence of Detrended Fluctuation Analysis and Spectral Analysis of RR Interval Variability," *Physiological Measurement*, Vol. 24, No. 1, 2003, pp. N1–N7.

[102] Teich, M. C., et al., *Heart Rate Variability: Measures and Models*, *Nonlinear Biomedical Signal Processing, Vol. II: Dynamic Analysis and Modeling*, Piscataway, NJ: IEEE Press, 2000.

[103] Ivanov, P. C., et al., "Scaling Behaviour of Heartbeat Intervals Obtained by Wavelet-Based Time-Series Analysis," *Nature*, Vol. 383, September 1996, pp. 323–327.

[104] Turcott, R. G., and M. C. Teich, "Fractal Character of the Electrocardiogram: Distinguishing Heart-Failure and Normal Patients," *Ann. Biomed. Eng.*, Vol. 2, No. 24, March–April 1996, pp. 269–293.

[105] Malamud, B. D., and D. L. Turcotte, "Self-Affine Time Series: Measures of Weak and Strong Persistence," *Journal of Statistical Planning and Inference*, Vol. 80, No. 1–2, August 1999, pp. 173–196.

[106] Lowen, S. B., et al., "Fractal Features of Dark, Maintained, and Driven Neural Discharges in the Cat Visual System," *AMethods*, Vol. 4, No. 24, August 2001, pp. 377–394.

[107] Costa, M., A. L. Goldberger, and C. K. Peng, "Open Source Software for Multiscale Entropy Snalysis," http://www.physionet.org/physiotools/mse/.

[108] Ivanov, P. C., et al., "Sleep Wake Differences in Scaling Behavior of the Human Heartbeat: Analysis of Terrestrial and Long-Term Space Flight Data," *Europhys. Lett.*, Vol. 48, No. 5, 1999, pp. 594–600.

[109] Lo, C. C., et al., "Common Scale-Invariant Patterns of Sleep-Wake Transitions Across Mammalian Species," *PNAS*, Vol. 101, No. 50, 2004, pp. 17545–17548.

[110] Moody, G. B., "ECG-Based Indices of Physical Activity," *Computers in Cardiology*, No. 19, 1992, pp. 403–406.

[111] Fukada, K., H. E. Stanley, and L. A. N. Amaral, "Heuristic Segmentation of a Nonstationary Time Series," *Phys. Rev. E*, Vol. 69, 2004, pp. 021108.

[112] Ferini-Strambi, L., et al., "The Impact of Cyclic Alternating Pattern on Heart Rate Variability During Sleep in Healthy Young Adults," *Clinical Neurophysiology*, Vol. 111, No. 1, January 2000, pp. 99–101.

[113] Brown, E. N., et al., "A Statistical Model of the Human Core-Temperature Circadian Rhythm," *Am. J. Physiol. Endocrinol. Metab.*, Vol. 279, No. 3, 2000, pp. E669–683.

[114] Thomas, R. J., et al., "An Electrocardiogram-Based Technique to Assess Cardiopulmonary Coupling During Sleep," *Sleep*, Vol. 28, No. 9, October 2005, pp. 1151–1161.

[115] Moody, G. B., et al., "Derivation of Respiratory Signals from Multi-Lead ECGs," *Proc. Computers in Cardiology*, IEEE Computer Society Press, 1986, pp. 113–116.

[116] Riker, R. R., J. T. Picard, and G. L. Fraser, "Prospective Evaluation of the Sedation-Agitation Scale for Adult Critically Ill Patients," *Crit. Care Med.*, Vol. 27, 1999, pp. 1325–1329.

[117] Hoyer, D., et al., "Validating Phase Relations Between Cardiac and Breathing Cycles During Sleep," *IEEE Eng. Med. Biol.*, March/April 2001.

[118] Schmidt, G., et al., "Chronotropy Following Ventricular Premature Beats Predicts Mortality After Acute Myocardial Infarction," *Circulation*, Vol. 98, No. 17, 1998, pp. 1–1016.

[119] Schneider, R., P. Barthel, and G. Schmidt, "Methods for the Assessment of Heart Rate Turbulence in Holter-ECGs," *JACC*, Vol. 33, No. 2, 1999, p. 315A.

[120] Watanabe, M. A., "Heart Rate Turbulence: A Review," *Indian Pacing Electrophysiol. J.*, Vol. 3, No. 1, 2003, p. 10.

[121] Schmidt, G., et al., "Heart Rate Turbulence," http://www.h-r-t.org.

[122] Watanabe, M. A., and G. Schmidt, "Heart Rate Turbulence: A 5-Year Review," *Heart Rhythm*, Vol. 1, No. 6, December 2004, pp. 732–738.

[123] Schmidt, G., et al., "Heart Rate Turbulence in Post-MI Patients on and off β-Blockers," *PACE*, Vol. 23, No. II, 2000, p. 619.

[124] Savelieva, I., D. Wichterle, and J. A. Camm, "QT-Interval Turbulence Induced by Atrial and Ventricular Extrastimuli in Patients with Ventricular Tachycardia," *Pacing and Clinical Electrophysiology*, Vol. 28, No. s1, 2005, pp. S187–S192.

[125] Bauer, A., et al., "Turbulence Dynamics: An Independent Predictor of Late Mortality After Acute Myocardial Infarction," *International Journal of Cardiology*, No. 107, 2006, pp. 42–47.

[126] Watanabe, M. A., et al., "Effects of Ventricular Premature Stimulus Coupling Interval on Blood Pressure and Heart Rate Turbulence," *Circulation*, Vol. 106, 2002, pp. 325–330.

[127] Schneider, R., G. Schmidt, and P. Barthel, "Correlation Coefficient of the Heart Rate Turbulence Slope: New Risk Stratifier in Post-Infarction Patients," *European Heart Journal*, Vol. 22, 2001, p. 72.

[128] Berkowitsch, A., et al., "Turbulence Jump: A New Descriptor of Heart-Rate Turbulence After Paced Premature Ventricular Beats, A Study in Dilated Cardiomyopathy Patients," *European Heart Journal*, Vol. 22, 2001, p. 547.

[129] Schmidt, G., et al., "Heart-Rate Turbulence After Ventricular Premature Beats as a Predictor of Mortality After Acute Myocardial Infarction," *Lancet*, Vol. 353, No. 9162, 1999, pp. 1390–1396.

[130] Schneider, R., et al., "Heart Rate Turbulence: Rate of Frequency Decrease Predicts Mortality in Chronic Heart Disease Patients," *PACE*, Vol. 22, No. 2, 1999, p. 879.