Quantifying Errors in Spectral Estimates of HRV Due to Beat Replacement and Resampling

Gari D. Clifford*, *Member, IEEE,* and Lionel Tarassenko

Abstract—Spectral estimates of heart rate variability (HRV) often involve the use of techniques such as the fast Fourier transform (FFT), which require an evenly sampled time series. HRV is calculated from the variations in the beat-to-beat (RR) interval timing of the cardiac cycle which are inherently irregularly spaced in time. In order to produce an evenly sampled time series prior to FFT-based spectral estimation, linear or cubic spline resampling is usually employed. In this paper, by using a realistic artificial RR interval generator, interpolation and resampling is shown to result in consistent over-estimations of the power spectral density (PSD) compared with the theoretical solution. The Lomb-Scargle (LS) periodogram, a more appropriate spectral estimation technique for unevenly sampled time series that uses only the original data, is shown to provide a superior PSD estimate. Ectopy removal or replacement is shown to be essential regardless of the spectral estimation technique. Resampling and *phantom* beat replacement is shown to decrease the accuracy of PSD estimation, even at low levels of ectopy or artefact. A linear relationship between the frequency of ectopy/artefact and the error (mean and variance) of the PSD estimate is demonstrated. Comparisons of PSD estimation techniques performed on real RR interval data during minimally active segments (sleep) demonstrate that the LS periodogram provides a less noisy spectral estimate of HRV.

Index Terms—Beat replacement, fast Fourier transform, heart rate variability, interpolation, irregular sampling, Lomb periodogram, resampling, sleep, uneven sampling.

I. INTRODUCTION

THE VARIATION in the timing between beats of the cardiac cycle, known as heart rate variability (HRV), has been shown to be an indicator of health [1]–[3]. Since spectral analysis was first introduced into the evaluation of HRV in 1969 [4], [5], a large body of literature has arisen concerning this topic. However, in 1993, the US Food and Drug Administration (FDA) withdrew its support of HRV as a clinical parameter due to a lack of consensus on the efficacy and applicability of HRV in the literature [6]. Although the Task Force of the European Society of Cardiology and the North American Society of Pacing Electrophysiology [3] 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 [7]). 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 is analyzed *off-line* by experts [8]. Such costly and time consuming expert analysis is required because of the limitations of algorithms in removing all nonsinus (ectopic) beats and artefacts from the data, together with the sensitivity of frequency-based HRV metrics to such outliers in the time series [3]. There is, therefore, a need to develop a method of automatic HRV analysis, which is insensitive to such artefacts and abnormal beats.

Standard methods for calculating frequency-based HRV metrics from the ECG require accurate beat detection (to locate a consistent reference point in the sinus beat's morphology). Before the time-series of beat-to-beat timing differences can be calculated, time stamps corresponding to ectopic beats and artefacts must be removed. Since ectopic beats often occur instead of a sinus beat, they are sometimes replaced by a *phantom* beat within the expected region for a sinus beat. [Confusingly, this is sometimes referred to as interpolation. In order to differentiate this from resampling, the former is referred to as phantom beat replacement, and the latter as resampling in this paper. Interpolation will be referred to by the specific method (linear or cubic) through which beat replacement or resampling is performed.]

A time series can then be formed from the successive difference between the times of occurrence of the remaining beats. If ectopic beat replacement is not used, the time interval formed by removing the beat must be discarded, since it does not represent a time between adjacent beats. The resultant time series is a list of varying intervals occurring at nonequidistant sampling times (the time difference between each R-peak on the electrocardiogram (ECG) [2]), known as an RR tachogram. Therefore, in order to perform spectral analysis using the fast Fourier transform (FFT), the time series should be resampled using interpolation at a frequency at least twice the maximum frequency of the signal (i.e., at least 7 Hz or 210 bpm [2]).

Although more complicated resampling schemes have been proposed [9], [10], the error this introduces in the evaluation of the power spectral density (PSD) has not been widely documented [11], [12]. The action of replacing unevenly sampled data (the RR intervals) with evenly sampled data requires the assumption of some underlying model which describes the relationship between each point. Stable schemes (such as linear interpolation) lead to coarse approximations to the underlying nonlinear behavior, and more complex schemes are prone to instabilities and fail to capture the full underlying dynamics of the signal. (See [23] for a more detailed discussion). Furthermore,

*G. D. Clifford is with the Department of Engineering Science, University of Oxford, Oxford OX1 3PJ U.K. and also with the Harvard-MIT Division of Health Sciences & Technology, Massachusetts Institute of Technology, Cambridge MA 02142 USA (e-mail: gari@mit.edu).

L. Tarassenko is with the Department of Engineering Science, University of Oxford, Oxford OX1 3PJ U.K.

resampling the RR tachogram at a frequency (f_s) below the original ECG (f_{ecg}) from which it is derived effectively shifts the fiducial point by up to $1/2(1/f_s - 1/f_{ecg})$ s. A more appropriate method for calculating the PSD of an unevenly sampled signal is the Lomb–Scargle (LS) periodogram [13]–[15] since no explicit data replacement is made (or model assumed) and the PSD is calculated from only the known values. This paper, therefore, presents a series of experiments to quantify the errors in spectral estimates of HRV due to resampling and beat replacement. PSD estimation performance of the LS periodogram on RR intervals with ectopic beats is compared with that of Welch's FFT [16] on the data after linear or cubic spline interpolation.

However, analysis of real data is problematic since the real (underlying) variability that we are attempting to estimate is completely unknown. Therefore, in order to ascertain the errors introduced through spectral estimation, a realistic artificial RR tachogram is initially used. The spectral methods are then implemented on real data to demonstrate consistency with results from the artificial scenarios.

II. BACKGROUND AND PREVIOUS WORK

Use of the FFT with linear or cubic spline interpolation for beat replacement and resampling are considered standard methods for spectral HRV analysis [2]. Clayton *et al.* [17] have demonstrated that FFT and autoregressive (AR) methods can provide a comparable measure of the low-frequency (LF) and high-frequency (HF) metrics (the power between 0.04 and 0.15 Hz and between 0.15 and 0.4 Hz, respectively [2]) on linearly resampled 5-min 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 modeling techniques will, therefore, not be considered within the scope of this paper.

In addition to the issues raised by the unevenly sampled nature of the RR tachogram, the necessity of removing nonsinus (ectopic) beats [2] introduces further resampling problems, increasing the variance of the interbeat intervals. Ectopic beats are routinely removed from the RR tachogram prior to HRV analysis for two main reasons. First, it is thought that an ectopic beat is not generated by the mechanism that is responsible for the variability in the RR intervals [2]. Second, ectopic beats often occur substantially earlier (or sometimes later) than when one would expect a normal beat to occur and are followed (or preceded) by a prolonged pause [2]. These unusually short RR intervals create higher than normal frequency components and lead to a significant rise in the estimation of the HF component which may distort the *true* measure of an HRV metric.

Previous comparative spectral HRV studies include Albrecht and Cohen's [18] experiments to compare the effects of different methods of phantom beat replacement on the FFT of (linearly) resampled tachograms. They found that beat replacement using linear interpolation produced a more accurate PSD estimate than using a predictive autocorrelation method. Birkett *et al.* [19] compared two methods of computing HRV spectra for congestive heart failure patients; in the first method, linear and cubic spline interpolation [20] was used to replace ectopic beats. In the second method, segments with ectopy were discarded. They

found that HF calculations were unaffected, but LF power was significantly higher using either interpolation method. In both methods the same FFT method was employed. In 1994 Lippman *et al.* [10] compared ectopy section removal with linear, cubic and nonlinear predictive interpolation as well as the null case (no ectopic removal). They concluded that ectopy correction is necessary for HRV analysis but that section removal performs as well as more complicated interpolation techniques. Their results indicated that linear and cubic interpolation of removed sections of the RR tachogram led to significant errors in some frequency and time domain HRV metrics.

A more appropriate PSD estimation technique for unevenly sampled data is the LS periodogram [13]–[15] and analysis by Moody *et al.* [11], [21], has shown that the LS periodogram can produce a more accurate estimation of the PSD than FFT methods for typical RR tachograms. In the even sampling limit, the LS and FFT methods are equivalent and, therefore, their performance is directly comparable when the underlying signal is completely known (such as for an artificial signal).

This paper, therefore, details experiments on the errors in estimating the LF, HF, LF/HF ratio metrics as evaluated using the FFT [3] and LS methods with and without phantom beat replacement and how resampling of the RR tachogram leads to errors in spectral HRV metrics as a function of the level of ectopy (or removed data). Results on real data of patients during sleep (the most quiescent of physiological activities where the HRV signal is the least random [22]) are then presented to illustrate how resampling leads to a significantly higher levels of noise in the HRV estimate.

III. EXPERIMENTAL METHODS

A. The Lomb-Scargle Periodogram

1) Definition: In Scargle's 1982 paper [14] on unevenly sampled time correlated data he demonstrated that the generalized N-point discrete Fourier transform (DFT) of a time series $X(t_j)$ for arbitrary times t_j (or uneven sampling $\Delta t = t_{j+1} - t_j \neq \text{constant}$) with variance σ^2 and a mean value $\overline{x}$, leads to the expression for the normalized (LS) periodogram

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

where $\tau \equiv \tan^{-1}((\sum_j \sin(2\omega t_j))/(2\omega \sum_j \cos(2\omega t_j)))$. 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 (1) 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 LS periodogram 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 ($\Delta t = t_{j+1} - t_j = \text{constant}$) (1) then reduces to the classical periodogram definition [13]. See [13]–[15] and [23] for mathematical derivations and further details. The implementation used in this paper is the *lomb.c* program available from Physionet [21].

2) Information Limits: In order to choose a sensible window size, the requirement of stationarity must be balanced against the time required resolve the information present. The European and North American Task force on standards in HRV [3] suggests that the shortest time period over which HRV metrics should be assessed is 5 min. As a result, the lowest frequency that can be resolved is $1/300 \approx 0.003$ 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 metrics. The upper frequency limit of the highest band for HRV analysis is 0.4 *Hz* [2]. Since the average time interval for N points over a time T is $\Delta t_{av} = T/N$ then the average Nyquist frequency [14] is $f'_c = 1/2\Delta t_{av} = N/2T$. Thus, a 5-min window ($T = 300$) with the Nyquist constraint of $N/2T \geq 0.4$ 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-min segment are used). Utilization of the LS periodogram, 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 s (corresponding to an instantaneous heart rate of $\text{HR}_i = 60/\text{RR}_i = 48$ 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. Heart rates with no RR intervals smaller than 1.25 s ($\text{HR}_i < 48$ bpm) can still be analyzed, but there is no power contribution at 0.4 Hz.

3) 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 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, 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 [24].

B. Resampling and Welch's FFT

Before presenting an unevenly sampled time series to an FFT algorithm the data must first be resampled to an even time se-

ries. This is achieved using the *interp1* function in the Matlab programming envoirnment. The algorithm resamples a series of unevenly sampled RR intervals and their time stamps, using a user-specified interpolation method onto a regular time axis at a given frequency. Two commonly employed methods for resampling will be considered; linear and cubic spline interpolation [2], [20]. Welch's FFT [16] is implemented using Matlab's *psd* function. Harris [25] has shown 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). This windowing function is, therefore, applied to all resampled data prior to PSD estimation via the FFT.

C. Generating Artificial Data

All data analyzed for this paper is artificially generated using a modification of a method detailed in previous papers by the authors [26], [27]. Segments of 5 min are used, with known frequency components, LF/HF ratio, average RR interval, and standard deviation of the RR interval. An artificial RR tachogram is generated by mixing two sine waves with frequencies at the LF and HF peaks. Since the LF frequency boundaries are defined in the literature to be 0.04 and 0.15 Hz [2], and the HF band lies between 0.15 and 0.4 Hz, the centre frequencies of these bands, 0.095 and 0.275 Hz, are chosen as the LF and HF frequency components $\omega_l/2\pi$ and $\omega_h/2\pi$ for the synthetic signal. The respiratory component (at 0.275 *Hz*) is given an amplitude of 2.5 bpm and the LF component (at 0.095 Hz) a 2 bpm amplitude. An expression for the HR can then be written as

$$\text{HR}(t) = \text{HR}_o + A_l \sin(\omega_l t) + A_h \sin(\omega_h t + \phi) \qquad (2)$$

where, initially, $\text{HR}_o = 60$ bpm, $\omega_l = 0.095/2\pi$, $\omega_h = 0.275/2\pi$, $A_l = 2$ and $A_h = 2.5$, with no phase difference ($\phi = 0$) between ω_l and ω_h. The RR tachogram can then be formed by sampling $\text{RR}_t = 60/HR_t$ at the required frequency f_s, such that $t_n = n/f_s$ where $(n = 1, 2, \ldots, N)$.

To introduce the small departures from stationarity that are observed in real RR tachograms [3], the two frequencies ω_l and ω_h are incrementally changed from sample to sample to generate a Gaussian spread of frequencies each with a standard deviation of $\sigma = 3$. The number of samples for each frequency ω_j in the LF and HF bands are (respectively) given by

$$S_l(j) = \sum_{j=1}^{N} (\text{int}) \left| \frac{N}{\sqrt{(2\pi\sigma_l^2)}} e^{-1/2((\omega(j)-\omega_l)/\sigma_l)^2} \right|$$

$$S_h(j) = \sum_{j=1}^{N} (\text{int}) \left| \frac{N}{\sqrt{(2\pi\sigma_h^2)}} e^{-1/2((\omega(j)-\omega_h)/\sigma_h)^2} \right| \qquad (3)$$

where $(\text{int})||$ indicates a standard rounding procedure to produce an integer number of samples, $\sigma_h = (\text{HF}_h - \text{HF}_l)/\sigma$, $\sigma_l = (\text{LF}_h - \text{LF}_l)/\sigma$, $\text{HF}_h = 0.4$ Hz, $\text{HF}_l = 0.15$ Hz, $\text{LF}_h = 0.15$ Hz, $\text{LF}_l = 0.04$ Hz, and $\sigma = 3$. Note that these frequency ranges are the upper and lower bounds of the LF and HF bands [3] and that the centre of these bands are the centre frequencies quoted above, 0.095 Hz and 0.275 Hz. Note also that as j is incremented from 1 to N, each corresponding frequency ω_j

to be generated is assigned a value $S_l(j)$ and $S_h(j)$. These are the number of samples at which ω_j remains at that frequency before being incrementally changed. In this way, the signal quickly drifts through a range of values free of noncausal jumps. The LF/HF ratio, is the ratio of the power in each frequency band is proportional to the square of the amplitude and is therefore, theoretically $(A_l/A_h)^2 = (2/2.5)^2 = 0.64$ in this example. Fig. 1(a) is a histogram of the number of samples S_j at each frequency. Fig. 1(b) is a section of the resultant RR tachogram and Fig. 1(c) is the corresponding PSD.

Ideally the sampling frequency, f_s, for generating the RR tachogram should be as high as possible. It has been demonstrated that the inaccuracies of R-peak location due to low ECG sampling frequencies can significantly affect the RR tachogram and any metrics derived from it [2]. 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.* [28] and Abboud *et al.* [29] have shown that for such patients a sampling rate of at least 1000 Hz is required. For normal patients however, a sampling rate of 128 Hz has been found to be accurate enough to locate the R-peaks and hence compute HRV [3]. To minimize sampling errors, the artificial RR tachogram [see, (2)] is generated at 1000 Hz.

To select physiologically plausible RR intervals, the RR interval at each sample point RR_t must match the difference between the time stamp of this sample and the time of the previous sample (Δt). One must, therefore, select a point and trace along the evenly sampled time series until one *discovers* an RR interval equal to the difference between its time stamp and the time stamp of the first RR interval. If this process is repeated along the entire length of the signal, a plausible RR tachogram results.

The following algorithm is used to generate such a time series (unevenly) sampled from the original time series:

Record the first data point pair, (t_1, x_1), as the first time stamp and RR interval pair (t_1', RR_1), Then proceed through each sample, t_i, until $x_n \geq t_i - t_1'$. This defines the second RR interval ($RR_2 = x_i, t_2' = t_i$). Generalize for all the $n = N$ RR intervals s.t. $\mathrm{RR}_n \geq t_i - t_{n-1}'$. (See **Fig. 1(b)** for a graphical illustration of this procedure.)

D. Resampling for Spectral Estimation

Although the artificial data is generated at 1000 Hz, to maintain accurate locations of the R-peaks, interpolative resampling of these points to produce an evenly sampled time series prior to FFT analysis is usually performed at lower rates to reduce the computational overhead. Most of the papers in the field of HRV report on the use of resampling rates between 2 and 4 Hz (although 1 and 10 Hz have also been used) [2], [3]. Choosing a resampling rate of 7 Hz for the RR tachogram has the advantage
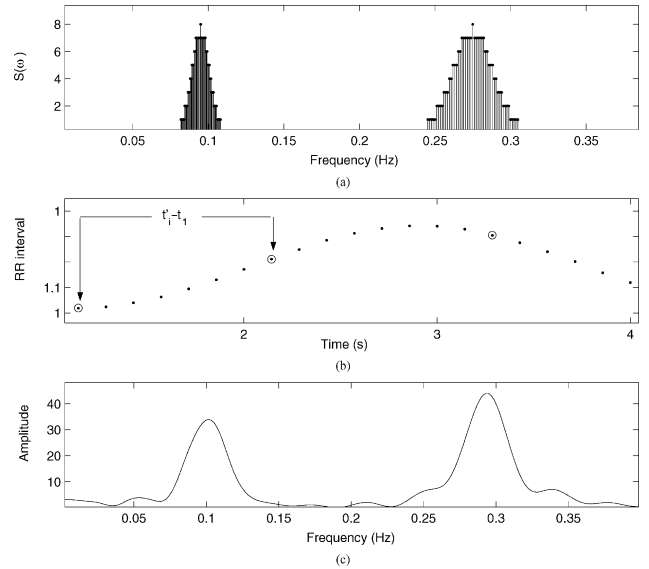

Fig. 1. (a) Histogram of the number of samples S_j generated by (3) at each frequency. (b) Section of the resultant RR tachogram using (2). (c) Corresponding PSD. For graphical clarity, the evenly sampled time series $(t_1 \ldots t_n)$ has been generated at 7 Hz rather than the 1000 Hz used in the experiments. The unevenly downsampled points $(t_i', i \geq n)$ are indicated by circles. See Section III-C.

that $\mathrm{HR}_i \leq 210$ bpm can be resolved (see Section III-A2 on the average Nyquist criterion). To compute an FFT, the number of points should be a power of 2. Therefore, 2^{11} points of each 7 Hz resampled segment (corresponding to approximately 293 s of data) are used in the following analysis.

E. Performance Metrics

In general, in order to assess spectral leakage and estimation inaccuracies the following parameters may be calculated:

- the amplitude of a particular sidelobe;
- the sum of the amplitudes of the sidelobes from the first up to some specific frequency, such as the Nyquist frequency;
- the amplitude of the peak at the Nyquist frequency (to measure the aliasing); and
- the width of the main (largest amplitude) peak to measure the resolution.

However, since the natural sampling of an RR tachogram is inherently uneven and only an *average* Nyquist frequency can be defined, a meaningful comparison between the DFT and the LS method for computing spectral components near the Nyquist frequency cannot be made. A ratio of the power between specified regions in the spectrum is nevertheless still appropriate and, thus, the LF/HF ratio makes a useful metric to assess significant differences (in terms of HRV values) between PSD estimates using different methods.

To measure the PSD estimation error as a function of the quantity of missing data, both the mean LF/HF ratio and the variance of its estimate using a Monte Carlo method (averaging the results of 1000 randomly seeded realizations of the artificial RR tachogram) for different levels of beat removal is calculated.
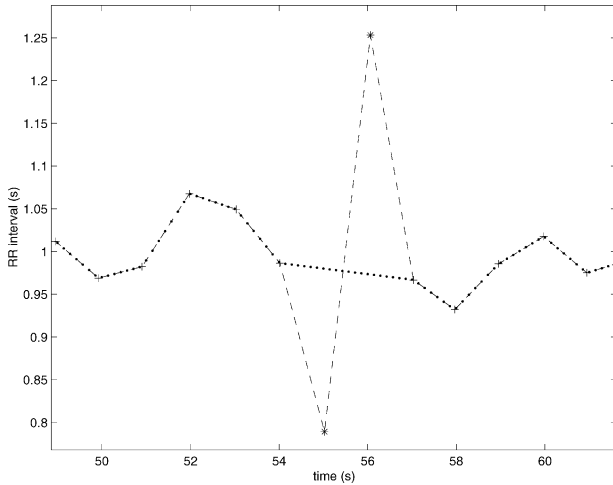
Fig. 2. *Realistic* RR tachogram (+) evenly resampled with 7 Hz linear resampling ($\cdots$). An ectopic beat is chosen to occur at about 55 s. Application of (5) moves the corresponding RR interval downwards and to the left. The following RR interval occurs at the same time as before, but has a much larger magnitude since the previous RR intervals occurred prematurely and, thus, its location in time has moved. Note that the line of dots which form a straight line through this segment (54–57 s) corresponds to the 7 Hz resampled (linearly interpolated) RR tachogram after these two points are removed (see Section V-B). The broken line represents the resampled waveform before the points are removed.

IV. EXPERIMENTS

The three methods for estimating the LF/HF ratio from the RR tachogram under consideration are; the FFT on 7 Hz linear resampled RR tachogram ($\mathrm{FFT_{lin}}$), the FFT on 7 Hz cubic spline resampled RR tachogram ($\mathrm{FFT_{cub}}$) and the LS periodogram without resampling (LS).

A. Adding an Artificial Ectopic Beat

This section deals with how the presence of ectopics and their removal affect estimates of spectral HRV metrics. Ectopic beats can be added to the artificial RR tachogram using a simple procedure. Kamath *et al.* [30] 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 2 data points as follows. The nth and $(n+1)$th beats (where n is chosen randomly) are replaced (respectively) by

$$\mathrm{RR}'_n = \gamma \mathrm{RR}_{n-1}, \qquad (4)$$
$$\mathrm{RR}'_{n+1} = \mathrm{RR}_{n+1} + \mathrm{RR}_n - \mathrm{RR}'_n \qquad (5)$$

where the ectopic beat's timing is the fraction, γ, of the previous RR interval (initially 0.8). Fig. 2 shows an artificial RR tachogram (generated as in the last section) with the effects of the ectopic beat added; a sharp dip and then an increase in the following RR interval because of the prematurity of the ectopic beat. The two points affected are marked by asterisks ($*$).

Note that the ectopic beat is introduced at random within the central 50% of the 5-min window. (Artificial ectopy is not introduced in the first or last 75 s of the 5-min window so that windowing does not significantly diminish the effect that the ectopy will have on the FFT-base estimates).
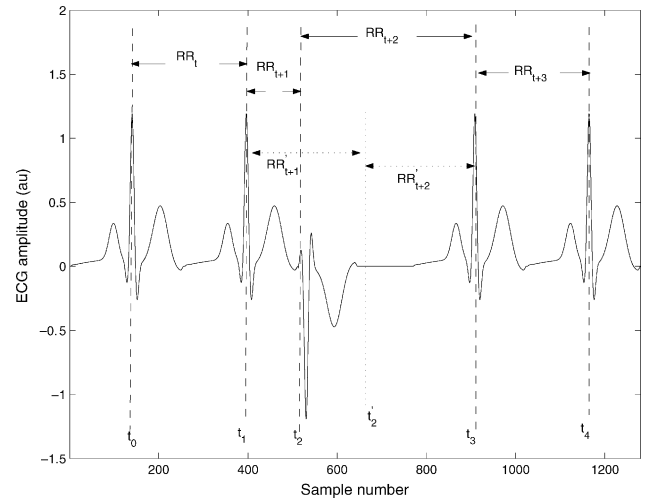

Fig. 3. Sinus beats occurring at times t_0, t_1, t_3 and t_4, and an ectopic beat occurring at t_2. The ectopic beat occurs earlier than would be expected for a sinus rhythm beat and shortens the first associated RR interval (RR_{t+1}) and lengthens the second (RR_{t+2}). To avoid introducing HF components at this point, a phantom beat can be placed at t'_2 half way between t_1 and t_3 so that RR_{t+1} and RR_{t+2} become RR'_{t+1} and RR'_{t+2}, both equal to $(t_3 - t_1)/2$.

B. Ectopic Beat Removal or Replacement

This section deals with the performance of each PSD estimation algorithm for varying levels of beat removal (the beats are assumed or defined to be ectopic). The procedure involves removing or replacing a randomly selected RR interval and the affected following datum (the asterisks in Fig. 2).

In order to evaluate the effect of beat replacement or removal on the spectral estimation methods being tested in this paper, the following four LF/HF ratio estimation techniques are compared for varying levels of ectopy:

(i)   FFT with ectopic beat removal, 7 Hz linear resampling
(ii) FFT with ectopic beat removal and insertion of phantom beat, 7 Hz linear resampling
(iii) FFT with ectopic beat removal, 7 Hz cubic spline resampling
(iv) LS periodogram with ectopic beat removal, no resampling

Fig. 4 illustrates the data after phantom beat insertion (if applicable) and resampling, prior to spectral estimation using the FFT or LS periodogram for each of these four methods for the data in Fig. 3. A Hamming window is applied to the resampled data of each 5-min segment.

For method (ii), the ectopic beat is replaced by a phantom beat, half way between the two adjacent beats. For methods (i), (iii), and (iv), the two (abnormal) RR intervals associated with the ectopic beat are removed. The appropriate resampling method is then employed, noting that no resampling is performed for method (iv). The mean and variance of these 1000 runs are then calculated. This is then repeated, gradually increasing the number of ectopic beats introduced within the window, up to the average Nyquist limit (see Section III-A2), where $\Delta \mathrm{RR}_{av} \geq 0.8$, giving an upper limit at most of 60 RR intervals (from 30 ectopic beats) being removed out of a total of 300 beats.
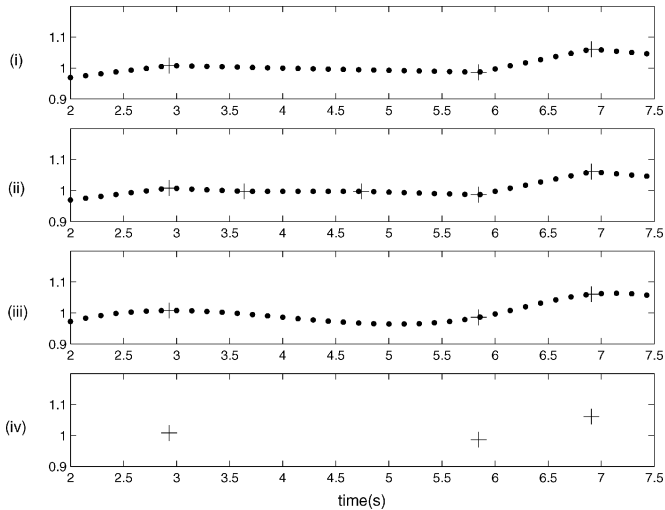
Fig. 4. Re-sampled data (evenly sampled dots) from original RR intervals (+) using the four methods detailed in the text; (i) Ectopic RR interval removal followed by 7-Hz linear resampling; (ii) Ectopic beat replacement followed by 7-Hz linear re-sampling; (iii) Ectopic RR interval removal followed by 7-Hz cubic spline re-sampling; (iv) Ectopic RR interval removal with no re-sampling.

TABLE I
LS AND FFT DERIVED FREQUENCY METRICS FOR DIFFERENT MAGNITUDES OF ECTOPY (γ)

Metric → PSD ↓	$\frac{LF}{HF}$	LF	HF	γ
LS	0.64	0.39	0.61	†
FFT_{cub}	0.69	0.41	0.59	†
FFT_{lin}	1.08	0.52	0.48	†
LS	0.60	0.37	0.62	0.8
FFT_{cub}	0.64	0.39	0.61	0.8
FFT_{lin}	0.99	0.50	0.50	0.8
LS	0.34	0.26	0.74	0.7
FFT_{cub}	0.40	0.29	0.71	0.7
FFT_{lin}	0.98	0.49	0.51	0.7
LS	0.32	0.25	0.76	0.6
FFT_{cub}	0.34	0.25	0.75	0.6
FFT_{lin}	0.53	0.35	0.65	0.6
LS	0.47	0.32	0.68	0.8 ‡
FFT_{cub}	0.50	0.33	0.67	0.8 ‡
FFT_{lin}	0.77	0.43	0.57	0.8 ‡

Subscripts indicate whether linear (lin), cubic spline (cub) or no interpolation was used. † indicates no ectopy is present. ‡ indicates two ectopic beats are present.

C. Real Data

While no *gold standard* exists for evaluating spectral HRV estimates on real data, the LF/HF ratio during sleep has been shown to be related to sleep state [2], [22]. In general, the deeper sleep a normal subject experiences, the more dominant the parasympathetic branch of the nervous system is, which is associated with an increase in the HF component and a parallel decrease in the LF/HF ratio [23], [22]. Since sleep stages can last up to 40 min, during which time there is often little change in the sympathovagal balance (except for brief arousals), and an accepted objective scale for measuring sleep depth exists, the RR tachogram of a sleeping normal subject provides an excellent signal upon which to compare the stability of the spectral estimation techniques described in this paper.

ECG and polysomnographic recordings were taken from 5 healthy males between the ages of 21 and 42 (mean age 32), weighing between 65 kg and 88 kg with no known cardiac or sleeping disorders. Subjects 1–5 slept for 7.2, 7.0, 3.3, 7.7, and 7.4 h, respectively. Subject 3 awoke early because of a nonhealth related incident. The polysomnogram was scored at 30-s intervals by two independent experts into 6 standard stages and only 5-min segments where a majority of the scores were in agreement (over 80% of the record) were used.

The fiducial points of each beat are extracted using a standard peak detector [23]. The nonsinus beats were then removed using a timing threshold technique described in [31]. Each RR tachogram was segmented into sections corresponding to rapid eye movement (REM) and slow wave sleep (SWS; stages 3 and 4). The LF/HF ratio was then estimated using the three methods (FFT_{lin}, FFT_{cub} and LS) for a sliding 5-min window (with a 80% overlap) applied to all contiguous segments for each sleep stage. The mean and variance of these estimate was then calculated on a per sleep stage basis.

V. RESULTS

A. The Effect of Ectopy on the LF/HF-Ratio

Table I details the results of the three frequency-domain HRV metrics calculation methods being investigated in this paper for different strengths of ectopy (γ) of a single beat in a 5-min section of an unevenly sampled RR tachogram (300 beats). The interpolation (in the case of FFT derived results) is performed at 7 Hz, as before. Note that for the no ectopy case (†) the LS method produces the theoretical value for the LF/HF ratio of 0.64 (see Section III-C). For decreasing values of γ, corresponding to increasingly earlier arrival times of the ectopic beat, both the values of the LF/HF ratio and LF decrease while the HF increases. In particular the LF/HF ratio drops from 1.08, 0.69, and 0.64 for the FFT_{lin}, FFT_{cub}, and LS methods, respectively, to approximately 50% of these values in the case of $\gamma = 0.6$ (a relatively common occurrence in normal humans [2]).

For more than one ectopic beat in a 5-min section of data (not occurring sequentially), the same trend is observed. The last 3 lines in Table I are the results for HRV metric estimation with two ectopics and $\gamma = 0.8$. Again, the power in the LF region is under-estimated and over-estimated in the HF region resulting in a value for the LF/HF ratio of approximately 50% of the no-ectopy case. The removal of ectopy for calculating spectral HRV metrics is, therefore, essential.

B. Resampling Methods With Beat Replacement or Removal: Metric Performance

Fig. 5 presents results on LF/HF ratio estimates from methods (i)–(iv) for Monte Carlo runs over 1000 5-min artificial RR tachograms each with an LF/HF ratio of 0.64. Linear interpolative beat replacement of ectopic beats [method (ii)] results in a similar performance to method (i), where no beat replacement is performed, due to the small difference in the value of the relevant RR intervals and the data inserted from linear resampling. However, as the incidence of ectopic beats increases, it
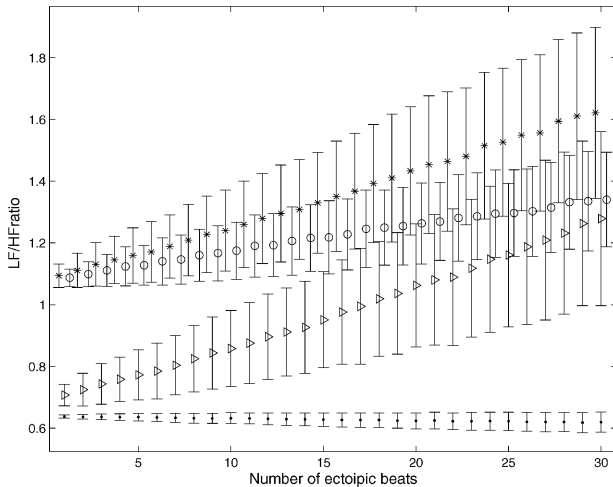
Fig. 5.   LF/HF ratio estimates for methods i ($*$), ii (o), iii ($\triangleright$) and iv ($\cdot$) when dealing with different numbers of ectopic beats in a 5-min RR tachogram with an actual LF/HF ratio of 0.64. Each point is an average of 1000 randomly seeded runs with error bars indicating ± 1 standard deviation. Note that each method has been plotted with a small offset on the x-axis for reasons of clarity.

becomes noticeable that the increase in over-estimation of the LF/HF ratio using method (i) is twice that of method (ii). This is as expected; twice the number of RR intervals are changed using linear resampling in method (i), since the ectopic beat is not replaced, but removed [see Fig. 4(i) and (ii)]. A similar trend can be observed in method (iii), but since cubic spline resampling results in a more accurate representation of the underlying signal than linear resampling [see Fig. 4(iii)], cubic spline resampling produces only a small over-estimate of the LF/HF ratio with a low incidence of ectopy. Interpolation of ectopic beats with resampling prior to FFT analysis [method (ii)] produces an estimate with accuracy between methods (i) and (iii); no beat replacement and resampling with linear and cubic spline interpolation, respectively.

Method (iv) produces an accurate estimate of the LF/HF ratio to within 1% of the theoretical value of 0.64, with a standard deviation over the 1000 Monte Carlo runs of less than 1%. At the limit of 30 ectopic beats, method (iv) continues to result in an accurate estimate [to within 3% of the theoretical value and with a standard deviation of less than 2.8%, an order of magnitude lower than for methods (i)–(iii)]. This is an important point; not only does resampling produce, on average, an over-estimate in the LF/HF ratio compared to the theoretical value, it also produces an increase in the uncertainty of the estimate. Both of these errors increase with the quantity of missing data. However, the LS method produces little error with respect to the theoretical value, and little increase in variance on the estimate at even high levels of missing data (20% in the case illustrated here).

It should be noted that these results are only for one particular heart rate and LF/HF ratio. For lower heart rates, a lower density of data will cause the linear interpolation-based FFT method to further over-estimate the LF frequencies and the cubic spline-based methods will behave increasing more erratically, oscillating between the RR intervals. The LS periodogram will provide only a slight increase in variance of the estimate, as for higher missing data. For higher heart rates, *ceteris paribus*

TABLE II
LS AND FFT DERIVED MEAN LF/HF-RATIOS ($\pm 1\sigma$) IN SWS

PSD $\rightarrow$ Subject $\downarrow$	LS	FFT_{cub}	FFT_{cub}
1	0.62 ± 0.77	0.83 ± 1.30	1.10 ± 1.25
2	1.92 ± 1.12	2.27 ± 1.87	3.46 ± 2.93
3	1.18 ± 0.52	1.30 ± 0.74	1.94 ± 1.14
4	0.95 ± 0.40	1.07 ± 0.68	1.48 ± 0.96
5	2.39 ± 1.31	2.53 ± 1.76	3.65 ± 2.50

Subscripts indicate interpolative technique; linear (lin), cubic spline (cub) or none. note that the mean LF/HF ratio and its variance (σ^2) are significantly lower for the ls periodogram (using the F-test and student's t-test, respectively).
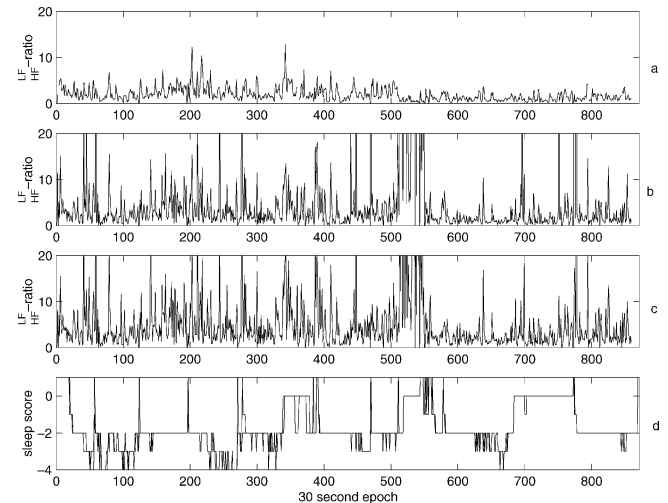

Fig. 6.   LF/HF ratio estimates (a-c) and sleep scores (d) during 7.2 h of sleep for a normal healthy human male (subject 01). LF/HF ratio is estimated using the three methods; (a) LS, (b) FFT_{cub} and (c) FFT_{lin}. Note the FFT methods produce estimates with a higher variance. The lower graph (d; the corresponding hypnogram over this segment) is plotted in the conventional manner. 1 indicates Awake, 0 REM sleep and -1 to -4 represent sleep states 1 to 4. SWS is defined to be stages 3 and 4.

one would expect an improvement in the performance of the FFT-based metrics since the density of data (in time) increases and the interbeat estimates of the interpolative techniques should be more accurate (for a given resampling frequency). However, since an increase in heart rate is associated with an increase in the incidence of ectopy [2] and artefact [31] and, therefore, more nonsinus RR intervals that require removal. FFT-based techniques may, therefore, perform less well at higher heart rates whereas the LS periodogram will result in an improved estimate from the overall larger number of data points available. Another factor that should affect the performance of the FFT-based metrics is the overall HRV. High variance will lead to increases in the overestimate of the LF frequencies by linear interpolation, and unstable oscillations by cubic spline interpolation.

C. Results on Real Data

Table II summarizes the results for all five subjects and Fig. 6 illustrates the fluctuations of the three spectral HRV metrics being tested for subject 1 over the entire night's record (7.2 h).

From visual inspection, the LF/HF ratio calculated by the LS method (a) appears to be less noisy (has a lower variance). This observation is verified using the F-test [15] ($P < 0.01$) to compare between the LS and each FFT-based method over the entire course of the record. Note that between the 500th and 600th 30 s epoch there is an increase in mental activity [REM and awake stages 0 and 1, respectively, in Fig. 6(d)] which is accompanied by an increase in heart rate and incidence of nonsinus (artefactual) data. The FFT-based estimates [Fig. 6(b) and (c)] have large (and differing) LF/HF ratio over these epochs, due to the large number of corrupted beats that must be removed and the resultant over-estimation of the LF component of the LF/HF ratio. Student's t-test [15] also demonstrates that the mean is significantly lower for the LS method. These observations are consistent for every subject (see Table II). See [32] for further details of experiments on real data.

VI. Discussion

It is clear that the interpolation techniques used to replace ectopic beats add LF components and reduce the high-frequency content. HF is, therefore, under-estimated whilst LF and the LF/HF ratio are over-estimated with resampling PSD estimation methods. Furthermore, not all ectopic beats occur *instead* of a sinus-beat. Uncertainty as to whether beat replacement or removal is appropriate, coupled with the negligible advantage offered by beat replacement in cases where it is appropriate, implies that the LS method, which is not plagued by such uncertainties, is a superior unsupervised method. The results presented in this paper are in agreement with the results on real data of Birkett *et al.* [19] and Lippman *et al.* [10] (see Section IV-A) noting that the LF over-estimation dominates the calculation of the LF/HF ratio.

By analyzing the data from a theoretical stand-point an information bound ($\Delta t_{av} = 1.25$) is found, above which the average Nyquist frequency cannot be resolved and an accurate estimate of the HRV at 0.4 Hz cannot be made. [15, Chapter 13] presents an interesting discussion of the estimation of the periodogram accuracy at each frequency based upon the sample statistics. A low number of RR intervals corresponding to the frequencies near 0.4 Hz will lead to erroneous estimates at these frequencies. The power above $1/2RR_{min}$ Hz, where RR_{min} is the shortest RR interval, should be assumed to be zero, since we do not have any information above this frequency.

Resampling of the inherently unevenly sampled RR tachogram adds erroneous data which cause significant errors in PSD estimates of the RR tachogram. Furthermore, replacement or removal of nonsinus beats (ectopic or artefacts) prior to FFT analysis using standard methods, compound these errors and produce significant increases in both the mean and variance of the estimate, which increase linearly with the incidence of nonsinus beats (such as artefacts, data dropouts and ectopy) which must be removed. The LS periodogram, a method which requires no resampling of unevenly sampled signals, is shown to provide a superior estimate of HRV metrics which is stable and accurate even for a high percentage (20%) of RR interval removal. Beat removal is shown to be essential, but replacement of ectopic beats is shown only marginally to reduce the error introduced by resampling (and incorrect

replacement increases the error). It should be noted that an optimal adaptive threshold timing for removing abnormal beats, prior to the use of the LS periodogram, is now appropriate. Since the number of missing beats does not significantly affect the LS periodogram-based estimate of the LF/HF ratio, a timing threshold, to exclude abnormal beats from the RR tachogram can be reduced (see [23] and [31]).

The LS method, therefore, makes an excellent on-line algorithm for analyzing data where high levels of artefact or ectopy are found, or where beat classification is uncertain (since removing a suspect beat does not affect the performance of the LS periodogram significantly, up to the average Nyquist limit). Recent work by Teich *et al.* [33] has shown that all current scale independent spectral HRV metrics (such as Wavelets and Detrended Fluctuation Analysis) evaluated on resampled data are theoretically equivalent to FFT methods. The analysis in this paper is, therefore, relevant to all existing frequency domain HRV techniques. A more detailed discussion of this can be found in [23].

Many of these techniques employed to calculate HRV metrics can only be evaluated on an evenly sampled time series. In order to compare the performance of such techniques, the LS method could be inverted to provide an optimal resampling scheme. Scargle [34] has demonstrated that calculation of a Lomb DFT is possible and hence its inverse could be used to aid resampling of a time series.

An interesting treatment of statistical procedures related to linear prediction and optimal filtering for unevenly sampled data sets by Rybicki *et al.* [35] provides alternatives to the above approach and could be compared. However, all of these can at best, deal with weakly nonstationary processes. A more recent method developed by Qi *et al.* [36] uses Kalman filtering to estimate the spectral components of unevenly sampled nonstationary data. This approach could be compared to other approaches that deal with nonstationary data after the application of various interpolation schemes to develop an improved HRV metric.

VII. Conclusion

Using a realistic artificial RR tachogram generator traditional FFT methods for calculating the PSD of resampled tachogram have been evaluated on a signal with completely known components and compared to the LS periodogram, which does not require resampling. This paper demonstrates that nonsinus information must be removed from the RR tachogram and that the estimation of HRV through conventional spectral methods (that require a resampling step) introduce significant systematic errors. The LS periodogram is shown to mitigate for these errors even in the presence of a large amount of missing data (either from corruption, ectopy or dropout).

By analyzing real data in sleep, it is demonstrated that the LS periodogram provides a lower variance estimate of the LF/HF ratio for a given sleep stage, particularly SWS. Since one expects the LF/HF ratio to be relatively stable within SWS, the lower variance of the LS estimate over time in SWS indicates that the LS periodogram provides a more robust and accurate method for calculating spectral HRV.

REFERENCES

[1] R. I. Kitney I and O. Rompelman, *The Study of Heart Rate Variability*. Oxford, U.K.: Oxford Univ., 1980.
[2] M. Malik and A. J. Camm, *Heart Rate Variability*. New York: Futura, 1995.
[3] M. Malik, "Heart rate variability: Standards of measurement, physiological interpretation, and clinical use," *Circulation*, vol. 93, pp. 1043–1065, 1996.
[4] D. G. Galloway and B. F. Womack, "An Application of Spectral Analysis and Digital Filtering to the Study of Respiratory Sinus Arrhythmia," Univ. Texas, Bioeng. Res. Lab., Electron. Res. Centre, Austin, TX, Tech. Rep. 71, 1969.
[5] B. M. Sayers, "Analysis of heart rate variability," *Ergonomics*, vol. 16, no. 1, pp. 17–32, 1973.
[6] M. H. Crawford, S. Bernstein, and P. Deedwania, "ACC/AHA guidelines for ambulatory electrocardiography," *Circulation*, vol. 100, pp. 886–893, 1999.
[7] (2003) Medical Predictive Science Corp. 510(k) Application no. K021 230. Food and Drug Administration, FDA. [Online]. Available: http://www.fda.gov/cdrh/pdf2/k021 230.pdf
[8] (2001) Boston Medical Technologies Inc. 510(k) Application no. K010 955. Food and Drug Administration, FDA. [Online]. Available: http://www.fda.gov/cdrh/pdf/k010 955.pdf,
[9] R. D. Berger, S. Akselrod, D. Gordon, and R. J. Cohen, "An efficient algorithm for spectral analysis of heart rate variability," *IEEE Trans Biomed Eng.*, vol. 33, no. 9, pp. 900–904, Sep. 1986.
[10] N. Lippman, K. M. Stein, and B. B. Lerman, "Comparison of methods for removal of ectopy in measurement of heart rate variability," *Am. J. Physiol.*, vol. 267, pp. H411–H418, 1994.
[11] G. B. Moody, "Spectral analysis of heart rate without resampling," *Comput. Cardiol.*, pp. 715–718, 1993.
[12] S. S. Abeysekera, U. R. Abeyratne, and S. B. Goh, "Spectral information changes in obtaining heart rate variability from tachometer R-R interval signals," *Crit. Rev. Biomed. Eng.*, vol. 28, no. 1-2, pp. 149–55, 2000.
[13] N. R. Lomb, "Least-squares frequency analysis of unequally spaced data," *Astrophysical Space Sci.*, vol. 39, pp. 447–462, 1976.
[14] J. D. Scargle, "Studies in astronomical time series analysis. II. statistical aspects of spectral analysis of unevenly spaced data," *Astrophysical J.*, vol. 263, pp. 835–853, 1982.
[15] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992, pp. 615–619.
[16] P. D. Welch, "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, pp. 70–73, Jun. 1967.
[17] R. H. Clayton, S. W. Lord, J. M. McComb, and A. Murray, "Comparison of autoregressive and Fourier transform based techniques for estimating RR interval spectra," *Comput. Cardiol.*, pp. 379–382, 1997.
[18] P. Albrecht and R. J. Cohen, "Estimation of heart rate power spectrum bands from real-world data: Dealing with ectopic beats and noisy data," *Comput. Cardiol.*, pp. 311–314, 1989.
[19] C. L. Birkett, M. G. Kienzle, and G. A. Myers, "Interpolation over ectopic beats increases low frequency power in heart rate variability," *Comput. Cardiol.*, pp. 257–259, 1991.
[20] P. DeBoor, *A Practical Guide to Splines*. Cambridge, MA: Springer, 1978.
[21] P. Laguna, 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. 45, no. 6, pp. 698–715, Jun. 1998.
[22] P. Lavie, *The Enchanted World of Sleep*. New Haven, CT: Yale Univ Press, 1996.
[23] G. D. Clifford, "Signal processing methods for heart rate variability," D.Phil. thesis, Oxford University, Oxford, U.K., 2002.
[24] A. Schmitz and T. Schreiber, "Testing for nonlinearity in unevenly sampled time series," *Phys. Rev. E*, vol. 59, p. 4044, 1999.
[25] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proc. IEEE*, vol. 66, pp. 51–83, 1978.
[26] P. E. McSharry, G. D. Clifford, L. Tarassenko, and L. A. Smith, "Method for generating an artificial RR tachogram of a typical healthy human over 24-hours," *Comput. Cardiol.*, pp. 225–228, 2002.
[27] ——, "A dynamical model for generating synthetic electrocardiogram signals," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 3, pp. 289–294, Mar. 2003.
[28] M. Merri, D. C. Farden, J. G. Mottley, and E. L. Titlebaum, "Sampling frequency of the electrocardiogram for spectral analysis of the heart rate variability," *IEEE Trans. Biomed. Eng.*, vol. 37, no. 1, pp. 99–106, Jan. 1990.
[29] S. Abboud and O. Barnea, "Errors due to sampling frequency of electrocardiogram in spectral analysis of heart rate signals with low variability," *Comput. Cardiol.*, pp. 461–463, 1995.
[30] M. V. Kamath and F. L. Fallen, "Correction of the heart rate variability signal for ectopics and missing beats," in *Heart Rate Variability*, M. Malik and A. J. Camm, Eds. Armonk: Futura, 1995, pp. 75–85.
[31] G. D. Clifford, P. E. McSharry, and L. Tarassenko, "Characterizing artefact in the normal human 24-hour RR time series to aid identification and artificial replication of circadian variations in human beat to beat heart rate using a simple threshold," *Comput. Cardiol.*, pp. 129–132, 2002.
[32] G. D. Clifford and L. Tarassenko, "Segmenting cardiac-related data using sleep stages increases separation between normal subjects and apnoeic patients," *IOP Physiol. Meas.*, vol. 25, pp. N27–N35, 2004.
[33] M. C. Teich, S. B. Lowen, B. M. Jost, K. Vibe-Rheymer, and C. Heneghan, "Heart rate variability: Measures and models," in *Nonlinear Biomedical Signal Processing Vol. II: Dynamic Analysis and Modeling*, M. Akay, Ed. Piscataway, NJ: IEEE Press, 2000.
[34] J. D. Scargle, "Studies in astronomical time series analysis. III – Fourier transforms, auto correlation functions and cross-correlation functions of unevenly spaced data," *Astrophysical J.*, vol. 343, pp. 874–887, 1989.
[35] G. B. Rybicki and W. H. Press, "Interpolation, realization and reconstruction of noisy, irregularly sampled data," *Astrophysical J.*, vol. 398, pp. 169–176, 1992.
[36] Y. Qi, T. P. Minka, and R. W. Picard, "Bayesian spectrum estimation of unevenly sampled nonstationary data," presented at the Int. Conf. Acoustics, Speech and Signal Processing (ICASSP), Orlando, Fl, May 13–17, 2002.
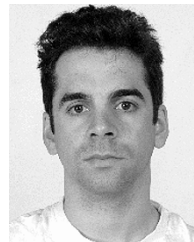
Gari D. Clifford (S'01–M'02) was born in Alton, Hampshire, U.K., in 1971. He received the B.Sc. degree in physics and electronics from Exeter University, Devon, U.K., the M.Sc. degree in mathematics and theoretical physics from Southampton University, Southampton, U.K., and the Ph.D. degree in biomedical signal processing and neural networks from Oxford University, Oxford, U.K., in 1992, 1995, and 2003, respectively.

He has also worked for DERA (now Qinetiq), Oxford Biosignals, Oriel College, and St. Cross College. He is currently a Research Scientist in the Harvard-MIT Division of Health Sciences & Physionet, Cambridge, MA. His research interests include biomedical signal processing, linear and nonlinear time series analysis, and mathematical modeling of the cardiovascular system.

Lionel Tarassenko was born in Paris, France, in 1957. He received the B.A. degree in engineering science in 1978, and the Ph.D. degree in medical engineering in 1985, both from Oxford University, Oxford, U.K..

After graduating, he worked for Racal Research Ltd. on the development of digital signal processing techniques, principally for speech coding. He then held a number of positions in academia and industry, before taking up a University Lecturership at Oxford in 1988. Since then, he has devoted most of his research effort to the development of neural network techniques and their application to signal processing, diagnostic systems, and parallel architectures. He has held the Chair in Electrical Engineering at Oxford University since October 1997.

Dr. Tarassenko was elected to a Fellowship of the Institution of Electrical Engineers (IEE) in 1996, when he was also awarded the IEE Mather Premium for his work on neural networks, and to a Fellowship of the Royal Academy of Engineering (RAE) in 2000.