Research Article

Time-Frequency Detection of Slowly Varying Periodic Signals with Harmonics: Methods and Performance Evaluation

John M. O'Toole (EURASIP Member)[1] and Boualem Boashash[1, 2]

[1] *Perinatal Research Centre and UQ Centre for Clinical Research, Royal Brisbane and Women's Hospital, The University of Queensland, Herston, QLD 4029, Australia*
[2] *College of Engineering, Qatar University, Qatar*

Correspondence should be addressed to John M. O'Toole, j.otoole@uq.edu.au

We consider the problem of detecting an unknown signal from an unknown noise type. We restrict the signal type to a class of slowly varying periodic signals with harmonic components, a class which includes real signals such as the electroencephalogram or speech signals. This paper presents two methods designed to detect these signal types: the ambiguity filter and the time-frequency correlator. Both methods are based on different modifications of the time-frequency-matched filter and both methods attempt to overcome the problem of predefining the template set for the matched filter. The ambiguity filter method reduces the number of required templates by one half; the time-frequency correlator method does not require a predefined template set at all. To evaluate their detection performance, we test the methods using simulated and real data sets. Experiential results show that the two proposed methods, relative to the time-frequency-matched filter, can more accurately detect speech signals and other simulated signals in the presence of coloured Gaussian noise. Results also show that all time-frequency methods outperform the classical time-domain-matched filter for both simulated and real signals, thus demonstrating the utility of the time-frequency detection approach.

1. Introduction

For some applications there is a need to detect time-varying periodic signals with harmonic components. Examples include electrical power signals [1], recorded musical instruments [2], and speech signals [3]. The work in this paper was motivated by one particular biomedical problem where we encounter signals with such characteristics. This application involves the detection of seizures in newborn infants, work which is being conducted in the Signal Processing Research Concentration of the Perinatal Research Centre within the UQ-CCR located in one of the largest hospitals in Australia, the Brisbane Royal and Women's hospital. The work involves developing an automated method to detect seizure periods in electroencephalogram (EEG) signals for newborns. The goal is to develop an accurate system that can be used in a clinical setting to assist the clinician in making a diagnosis or prognosis. An effective detection scheme would allow for the improvement of health outcomes for newborns.

Analyses of EEG showed nonstationary characteristics for the seizure signals [4, 5], which lead to the development of a detection method in the time-frequency domain [6, 7]. This method used a time-frequency-matched filter which correlated a set of reference signals, in the time-frequency domain, to the EEG signal. Initial results were promising [6] but the efficacy of the method is limited somewhat by the variability of the seizure signal [7]. The crux of the problem is how to define the set of reference signals, known as the template set, to accurately represent the range of seizure types—too many templates lead to computational problems and increase the probability of error, too few templates, and the probability of missing a seizure increases.

This provided the motivation to develop new methods, still using the knowledge gained from the time-frequency characterisation of seizure periods [4, 5], without the problems associated with predefining a template set. Two methods, both based on modified versions of the time-frequency matched filter, are described in this paper. The first

method requires only one half of the template set used by the time-frequency-matched filter method; even better, the second method requires no predefined template set at all.

We evaluated the detection accuracy of these two methods, and compared this performance with the performance for the time-frequency-matched filter and the time-domain-matched filter methods. This performance evaluation used both simulated and real data sets. The results show two key points: first, we were not able to identify the most accurate detection method over all the tested data sets; and second, all three time-frequency methods outperformed the time-domain-matched filter for all, but one, data sets. (The one exception was the data set of simulated signals embedded in white Gaussian noise.) We found, unexpectedly, that for our particular EEG seizure data set the time-frequency-matched filter marginally outperformed the two proposed methods. One of the proposed methods, however, outperformed all other methods for the speech data set and for the simulated data sets with coloured Gaussian noise.

2. Background

To introduce the topic, we start with a review of time-frequency distributions and the matched filter.

2.1. Time-Frequency Distributions. Time-frequency distributions (TFDs) are often used to analyse nonstationary signals because they highlight the time-varying characteristics of these signals [8]. There are many different types of TFDs, which are grouped into classes. Probably the most commonly used class is the quadratic class of TFDs. At the core of this class is the Wigner-Ville distribution (WVD). The WVD,

$$W_s(t, f) = \int_{-\infty}^{\infty} z\left(t + \frac{\tau}{2}\right) z^*\left(t - \frac{\tau}{2}\right) e^{-j2\tau f} d\tau \qquad (1)$$

is a quadratic function of the time-domain signal $z(t)$, where $z^*(t)$ is the complex conjugate of $z(t)$. The signal-under-analysis, $s(t)$, is transformed from a real-valued to complex-valued signal $z(t)$ using the Hilbert transform [9]. Because the transformation of $z(t)$ from the time to the time-frequency domain is bilinear, the WVD contains cross-terms between the signal's components [9]. Convolving the WVD with the time-frequency kernel $\gamma(t, f)$ can suppress or attenuate these cross-terms. This smoothed WVD represents the quadratic TFD class $\rho_s(t, f; \gamma)$,

$$\rho_s(t, f; \gamma) = W_s(t, f) \underset{t}{*} \underset{f}{*} \gamma(t, f), \qquad (2)$$

where different kernels define different distributions in the class [8].

2.2. Matched Filter. We can describe the problem of detecting a seizure event from the nonseizure background in the EEG signal as the classic detection problem of detecting a known signal in noise. For the signal $x(t)$ there are two possibilities:

$$\begin{aligned} &\text{H}_0\text{:} \ \ x(t) = n(t), &&\text{signal absent,} \\ &\text{H}_1\text{:} \ \ x(t) = s(t) + n(t), &&\text{signal present,} \end{aligned} \qquad (3)$$

FIGURE 1: The epoch TFD split into 4 segments of length T_s. The bold line represents the IF law of the signal.


FIGURE 2: The ith TFD segment (a) $\rho_i(t, f)$ and (b) template TFD $\rho_i^{\mathrm{T}}(t, f; \theta_k)$. Note that the template TFD is a time-reversed version of the $(i-1)$ TFD segment. The template TFD is rotated about a set of slope values to attempt a match with the ith TFD segment. This example is taken from Figure 1 for $i = 2$.

where $s(t)$, of duration T, represents the signal to detect and $n(t)$ represents the noise. The detection problem is how to select the correct hypothesis, H_0 or H_1. If we assume H_1, and $s(t)$ is present, then this is a true detection; if we assume H_1, and $s(t)$ is not present, then this is a false detection.

The matched filter is a time-domain detection method that linearly filters $x(t)$ [10]. This method maximises the signal to noise ratio (SNR) of $s(t)$ and $n(t)$. The basic method requires that the signal $s(t)$ is known and $n(t)$ is a zero-mean, white Gaussian noise process. To make the decision as to whether $s(t)$ is present or not, we form a test statistic as follows:

$$\eta = \int_{(T)} x(t) s^*(t) dt \qquad (4)$$

TABLE 1: Parameter values to generate the slowly-varying periodic signal with K harmonic components. All parameters are generated from a uniform probability distribution $U(a, b)$, where a, b represent the lower and upper limits of the distribution. The subscript k represents the kth harmonic component and T_e is the length in seconds of the signal. The parameter g defines the relative energy in the harmonic components.

Parameter	Value
Number of harmonics, K	$U(2, 5)$
Number of turning points for IF law, M	$U(3, 9)$
Time-value turning point for amplitude law $t1_i$	$t1_i = c + d_i T_e/5$,
	with $c \sim U(-T_e/9, T_e/9)$ and $d = \{1, 2.5, 4\}$
Amplitude-value of turning point for amplitude law a_i	$U(0.9, 1.1)$
Amplitude scaling parameter e_k, for $2 \leq k \leq K+1$; ($e_1 = 1$)	$e_k = 1/(k+g) + 0.3 + b$,
	with $b \sim U(-0.15, 0.15)$
Time-value of turning point for IF law $t2_i$	$U(M^{-1} - 0.1, M^{-1} + 0.1)$
Frequency-value of turning point for IF law f_i	$U(-0.003, 0.003)$
Phase value θ_k	$U(-\pi, \pi)$

and then compare this test statistic with a predefined threshold value ζ to determine the hypothesis. Thus,

$$H_0: \eta < \zeta,$$
$$H_1: \eta > \zeta. \tag{5}$$

The matched filter is known as an optimum detector because it maximises the SNR and therefore maximises the probability of a true detection [11].

We can extend the basic matched filter method, which uses the time-domain signal in (4), with a time-frequency formulation. The inner-product of the WVDs for the signals $x(t)$ and $s(t)$,

$$\eta_{\text{TF}} = \iint_{(T)} W_x(t, f) W_s(t, f) dt \, df \tag{6}$$

is directly related to the matched filter in (4) as [12–14]

$$\eta_{\text{TF}} = |\eta|^2. \tag{7}$$

(Integral limits, unless otherwise specified, span from minus to plus infinity.) The test statistic η_{TF} is known as a locally optimal detector [13, 14]. Because of the direct relation between η_{TF} and η, the time-frequency approach provides no immediate advantage over the conventional-matched filter, apart from indirect advantages such as embedded time-frequency filtering or use of the cross-WVD [12].

If we replace the WVD $W(t, f)$ with the more general TFD representation $\rho(t, f)$,

$$\eta_{\text{TF}} = \iint_{(T)} \rho_x(t, f) \rho_s(t, f) dt \, df \tag{8}$$

then this test statistic η_{TF} is only related to η in (7) if the Doppler-lag kernels for ρ_x and ρ_s satisfy the condition $|g(\nu, \tau)| = 1$ [13]. This condition severely constrains the type of TFDs as most useful TFDs have nonunity, real-valued kernels. The notable exception is the WVD.

The time-frequency-matched filter, using the test statistic in (8), is known as a suboptimum detector [13, 15] because

the filter does not maximise the SNR. The method, however, may prove useful for an application when the constraints on the matched filter do not hold and the optimum method is not applicable. That is, if the signal $s(t)$ is not known and can only be inferred from noisy measurements, or if $s(t)$ is randomly perturbed in some way, or if the noise $n(t)$ is not white Gaussian noise then the suboptimum method may prove useful [6, 13, 15–18].

We can shift the test statistic over time and frequency to obtain $\eta_{\text{TF}}(t, f)$, which (10) in the next section describes, and then use the generalised likelihood ratio test [13, 14] to obtain the test statistic:

$$\eta_{\text{TF}} = \max_{(t, f)} \eta_{\text{TF}}(t, f). \tag{9}$$

This approach is useful when $s(t)$ is shifted in time and/or frequency.

2.3. Time-Frequency-Matched Filter. The method in [6, 7] uses the time-frequency-matched filter method to detect seizure in newborn EEG. This method correlates a template set, a collection of seizure-like events, with the EEG signal $\text{eeg}(t)$ in time frequency as follows:

(1) form the TFD ρ_{eeg} for EEG signal $\text{eeg}(t)$ for an epoch of length T;

(2) for each reference signal $r(t)$ form the TFD template ρ_r and

   (a) produce the time-frequency matched filter test statistic $\eta_{\text{TF}}(t, f)$ shifted over time and frequency:

$$\eta_{\text{TF}}(t, f) = \iint_{(T)} \rho_{\text{eeg}}(t', f') \rho_r(t' - t, f' - f) dt' df' \tag{10}$$

   (b) using (9), take the maximum value of $\eta_{\text{TF}}(t, f)$ to obtain the final test statistic
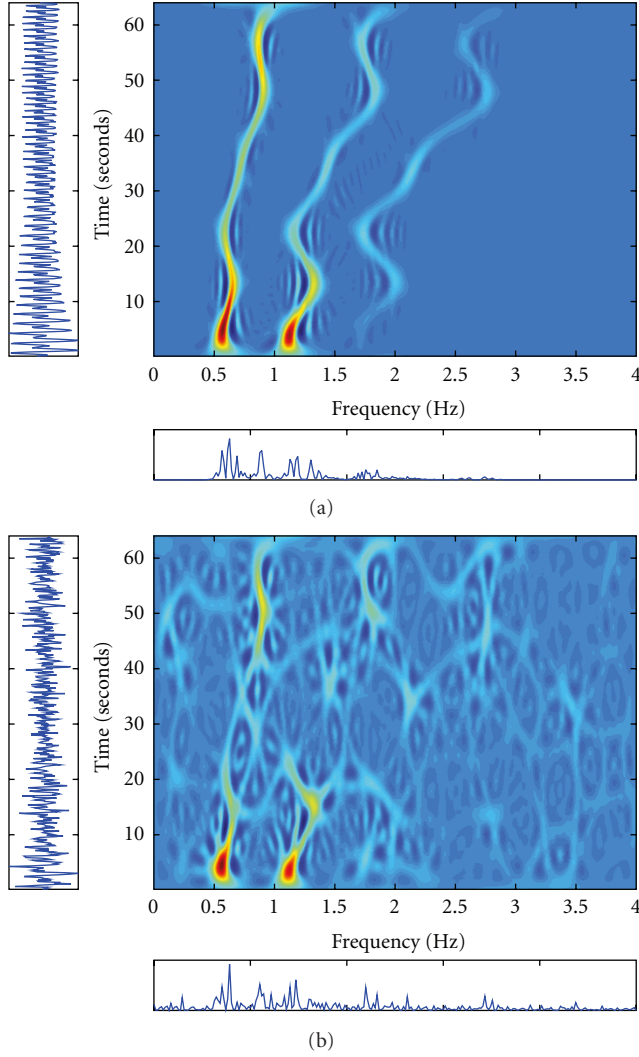
(3) iterate over all epochs for the EEG.

(a)

(b)

FIGURE 3: Simulated signal with low-energy harmonics: (a) SNR of 20 dB, and (b) SNR of 0 dB. Noise is white Gaussian noise.

(a)

(b)

FIGURE 4: Simulated signal, with high-energy harmonics: (a) SNR of 20 dB, and (b) SNR of 0 dB. Noise is coloured Gaussian noise.

The reference signal $r(t)$ is a piecewise LFM (linear-frequency modulated) signal [6]. The template set is a collection of piecewise LFM signals with different LFM slope parameters. (We use the phrase *LFM slope* to refer to the slope of the instantaneous frequency (IF) law of the LFM signal.) Note that the method in [6, 7] used a slightly different version of the process described here; we detail the difference in the appendix. Also, we show in this appendix that both approaches give similar results.

Defining the template set for the time-frequency-matched filter is problematic [19]. Although the piecewise LFM model—or piecewise LFM model with harmonic components [20]—can accurately model seizure events [20], the parameters in these models vary from newborn to newborn, or even from EEG channel to channel, in the same patient. Thus the method requires a large template set to represent patient or channel-specific seizures. The size of the template set is, however, proportional to the probability of error—as the template set size increases so does the false alarm rate [15, 19].
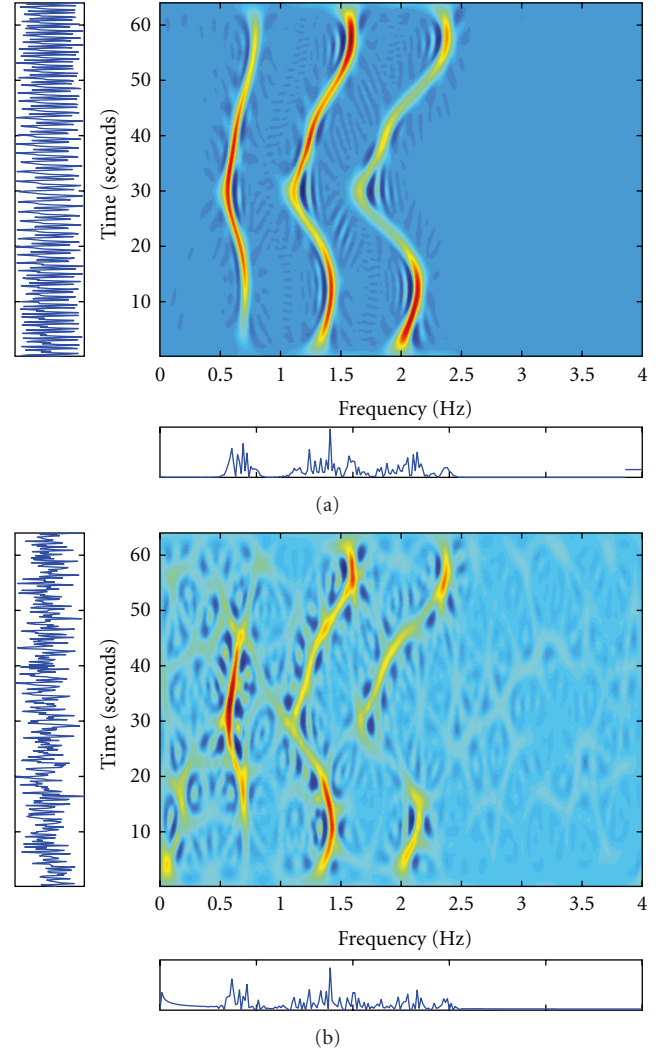
TABLE 2: Seizure detection results using newborn EEG.

Method	AUC	TDR = 1 − FDR
matched filter	0.75	72%
TF matched filter	0.99	96%
ambiguity filter	0.91	88%
TF correlator	0.95	90%

3. Methods

This section describes two extensions of the time-frequency-matched filter: the ambiguity filter method and the time-frequency correlator method. These two new methods attempt to address the problem of having to predefine the template set for unknown signals types. The following explains these two methods in more detail.

3.1. Ambiguity Filter Method. This method reduces the size of the time-frequency-matched filter's template set,

FIGURE 5: Results for detecting signals in white Gaussian noise with lower-energy harmonics. At a specific signal to noise ratio (SNR), 50 epochs of noisy signal and 50 epochs of noise are used to generate the receiver operator characteristic curve. From this curve, the area under the curve (AUC) and point of equal sensitivity and specificity (TDR = 1 − FDR) are calculated.

by defining the templates as real-valued functions in the Doppler-lag domain [19]. For the time-frequency-matched filter method, the piecewise LFM signal has $2L$ parameters, where L is the number of pieces. For example, a 2-piece LFM signal requires the parameters $[T_1, T_2, \alpha_1, \alpha_2]$, where T_i is the time duration and α_i is the slope of the ith piece, to uniquely define the signal. The modified method in [19] requires only the slope α_i parameter because, as we shall see, the templates are independent of time duration values T_i. Thus, for our 2-piece LFM example this modified method requires only the parameters $[\alpha_1, \alpha_2]$.

To explain how this method works, let us start with the test statistic $\eta_{\mathrm{TF}}(t, f)$, which equates to [19]

$$\eta_{\mathrm{TF}}(t, f) = \mathcal{F}^{-1}_{\nu \to t} \left\{ \mathcal{F}_{\tau \to f} \left\{ A_{\mathrm{eeg}}(\nu, \tau) \hat{A}_r(\nu, \tau) \right\} \right\}, \qquad (11)$$

where $\mathcal{F}$ represents the Fourier transform; $\mathcal{F}^{-1}$ represents the inverse Fourier transform; and $A_{\mathrm{eeg}}(\nu, \tau)$ is the ambiguity function (AF). The AF is a two-dimensional Fourier transform of the TFD:

$$A_{\mathrm{eeg}}(\nu, \tau) = \mathcal{F}_{t \to \nu} \left\{ \mathcal{F}^{-1}_{f \to \tau} \left\{ \rho_{\mathrm{eeg}}(t, f) \right\} \right\}. \qquad (12)$$

(Also, ν represents the Doppler direction, and τ represents the lag direction.) The symbol $\hat{A}_r$ represents an AF of sorts: the two-dimensional Fourier transform of the *time- and frequency-reversed* TFD:

$$\hat{A}_r(\nu, \tau) = \mathcal{F}_{t \to \nu} \left\{ \mathcal{F}^{-1}_{f \to \tau} \left\{ \rho_r(-t, -f) \right\} \right\}. \qquad (13)$$

We need this reversed-TFD transform to satisfy the relation in (11) [19].

The time-frequency-matched filter method uses TFD templates $\rho_r(t, f)$ of the piecewise LFM signal. These template TFDs contain both auto- and cross-terms. The reference AF templates $\hat{A}_r$, however, use only the autoterms. We define $\hat{A}_r$, the reference template, as a sum of window functions $h(t)$ located along the $(\nu - \alpha_i \tau)$ axis, as this is where the autoterms reside [21]. (Recall that α is the slope of the LFM signal.) Each function $h(\nu - \alpha_i \tau)$ is meant to model the auto-term for the ith piece of the piecewise LFM signal. We construct the AF of L auto-term components as

$$\hat{A}_r(\nu, \tau) = \sum_{i=0}^{L-1} h(\nu - \alpha_i \tau). \qquad (14)$$

Thus, the AF of the reference signal is independent of T_i, the length of the pieces in the piecewise LFM model. For this method, we used a Gaussian window for $h(t)$ in (14).

The final stage of the detection process is to extract a continuous IF from the test statistic $\eta_{\mathrm{TF}}(t, f)$ in (11) and then use the length of the IF as the final test statistic. We do this because the reference templates $\hat{A}_r(\nu, \tau)$ are similar to a smoothing Doppler-lag kernel, and therefore $\eta_{\mathrm{TF}}(t, f)$ is similar to a TFD. Thus, if the continuous IF is long enough, greater than a predefined threshold, then this would indicate the presence of a slowly varying periodic signal. We shall refer to this method as the *ambiguity filter method* because the templates set could be described as a set of filters in the ambiguity domain.
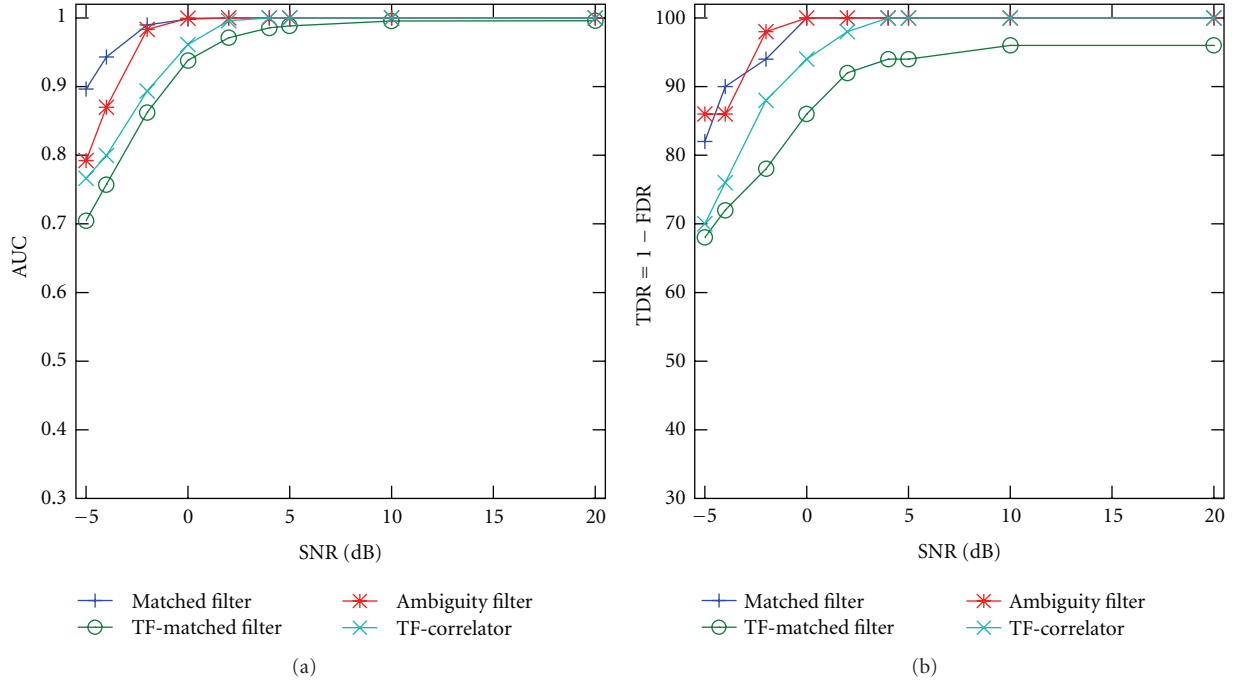
FIGURE 6: Results for detecting signals in white Gaussian noise with high-energy harmonics.

There are three advantages, comparative to the time-frequency-matched filter method, for the ambiguity filter method: (1) the template set is reduced by one half which reduces the probability of error in the method [15]; (2) by defining the AF as a sum of smoothing functions located on the autoterms of the piecewise LFM model, the ambiguity filter method is more robust to differences between the template and EEG seizure epoch because the method needs only match the autoterms and not the cross-terms; (3) the computational load is reduced by one half because of the smaller template set size. These first two advantages are reflected by the results in [19] which show how the ambiguity filter method outperforms the time-frequency-matched filter method. Although these results reflect an improvement, the challenging problem still remains—how to predefine the template set when the signal to detect is not exactly known? The next method addresses this problem.

3.2. Time-Frequency Correlator Method. The following method does not require a predefined template set [17]. This method uses the principle that because our slowly-varying periodic signal is repetitive, a short-time segment of the seizure should correlate well with an adjacent short-time segment. That is the principle that Navakatikyan et al. [22] use for their EEG seizure detection method. The proposed method here uses additional prior information, however. We know, from time-frequency analysis, that a seizure signal can be represented by a piecewise LFM with additional harmonics [20, 23]. The proposed method correlates time-adjacent, short-time segments in the time-frequency domain to match the slowly varying IF laws. This method correlates two TFDs, not WVDs, in the time-frequency domain and therefore differs from the time-domain-matched filter, as the test statistic

in (8) does not, in general, reduce to (6). We shall refer to this method as the time-frequency correlator method.

The outline of this method is as follows. First, generate a TFD for the received signal of length T_e seconds. Then, split this TFD up into 4 segments each of length $T_s = T_e/4$. The method assumes that the IF laws in each segment is linear, although these IF segments can have different slope values. Next, correlate segment one with segment two, rotating segment one to allow for a difference in slopes between the two segments. Continue this procedure for all the segments; that is, segment 2 correlated with segment 3, and segment 3 correlated with segment 4. Finally, find the minimum test statistic from these correlations and compare this with a threshold to produce the hypothesis. The method, in more detail, now follows.

3.2.1. Detection Method:

(1) split the received signal up into epochs $\mathrm{eeg}_j(t)$ of length T_e;

(2) form the TFD for $\mathrm{eeg}_j(t)$ as $\rho_{\mathrm{eeg}}(t, f)$;

(3) divide this TFD, in time, into four segments of length $T_s = T_e/4$, known as $\rho_i(t, f)$ for $i = 0, 1, 2, 3$. This segmentation process is illustrated in Figure 1;

(4) iterate the following over $i = 1, 2, 3$:

    (a) define the template TFD $\hat{\rho}$ for the ith segment as a time inverted TFD $(i - 1)$ segment; that is, let

$$\hat{\rho}_{(i-1)}(t, f) = \rho_{(i-1)}(iT_s - t, f). \quad (15)$$

FIGURE 7: Results for detecting signals in coloured Gaussian noise with low-energy harmonics.
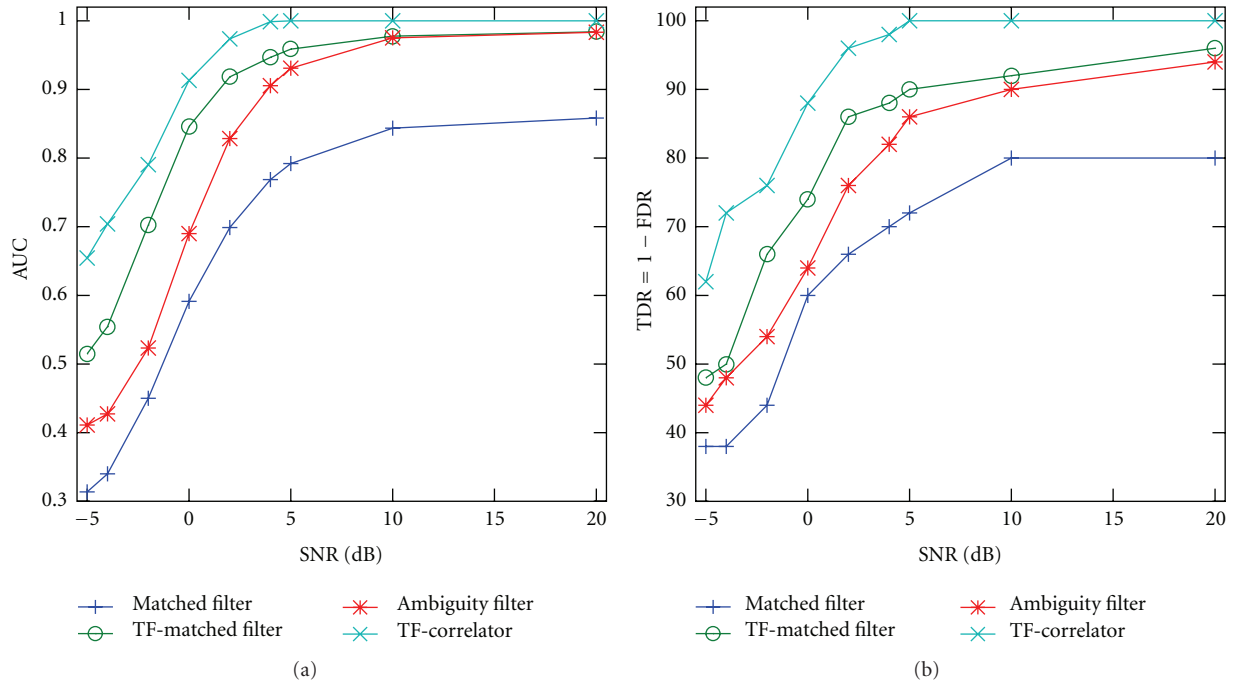

FIGURE 8: Results for detecting signals in coloured Gaussian noise with high-energy harmonics.

Thus, the time-inverting process for $\rho_{(i-1)}(t, f)$ is equivalent to turning the TFD segment upside down in time,

(b) rotate the template TFD $\hat{\rho}_{(i-1)}$ over a set of discrete rotations $\Theta = \{\theta_1, \theta_2, \ldots, \theta_K\}$; let

$$\rho_i^{\mathrm{T}}(t, f; \theta_k) = \hat{\rho}_{(i-1)}(t, f) \underset{f}{*} W_{m_k}(t, f), \qquad (16)$$

where $m_k(t) = \mathrm{e}^{\mathrm{j}2\pi(\theta_k/2)t^2}$ and θ_k is from the set Θ. This results in a set of rotated TFD templates.

(c) correlate the template set with the ith segment TFD,

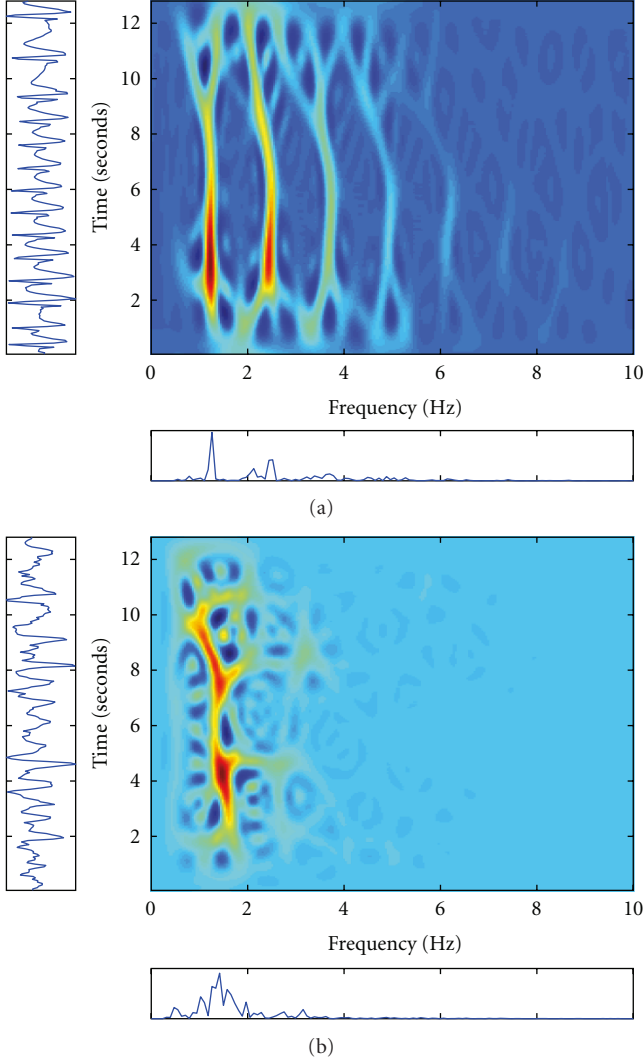$$\hat{\eta}(\theta_k) = \iint_{(T_s)} \rho_i(t, f)\rho_i^{\mathrm{T}}(t, f; \theta_k)\,\mathrm{d}t\,\mathrm{d}f \qquad (17)$$

FIGURE 9: Epoch of newborn EEG with seizure (a) and nonseizure (b) activity.

FIGURE 10: Speech segment (recorded spoken word) in 20 dB (a) and 0 dB (b) of babble noise.

and then take the maximum test statistic

$$\hat{\eta}_i = \max_{\theta_k \in \Theta} \hat{\eta}_i(\theta_k). \tag{18}$$

Figure 2 illustrates this process.

Why is this rotation process necessary? Recall that we have assumed a piecewise LFM-type signal is present in $\rho_{\text{eeg}}(t, f)$ and that the turning points of the IF in the TFD segments are located at $t = t_s + iT_s$. Because the IF law for the continuous component passes through the time-frequency point $(t_s + iT_s, f_i)$, then $\overline{\rho}_{(i-1)}(t, f)$ and $\rho_i(t, f)$ will be equal around $t = 0$, as Figure 1 illustrates. If the slope of the LFM in the $(i-1)$ segment, $\alpha_{(i-1)}$, does not equal the slope, α_i, in LFM of the ith segment, then the correlation between the $(i-1)$ and i segments will be small. Thus, if we rotate the template TFD $\overline{\rho}_{(i-1)}(t, f)$ about the point
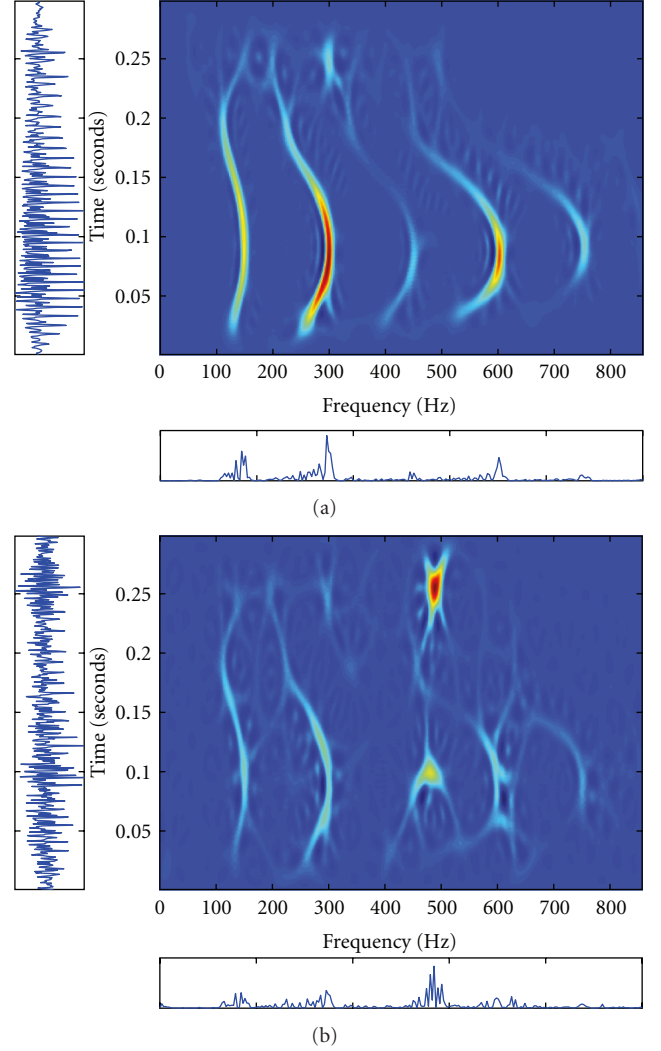
$(0, f_i)$ to the angle $\alpha_k = \alpha_{(i-1)} + \alpha_i$, then the two TFD segments would match and produce a large correlation (see Figure 2);

(5) the test statistic for the epoch is

$$\hat{\eta} = \min_{i \in \{1,2,3\}} \hat{\eta}_i. \tag{19}$$

The rationale for this is that if the LFM component is continuous and present throughout the four segments, then each $\hat{\eta}_i$ will remain relatively large; likewise, if the LFM component is not present over all segments then $\hat{\eta}_i$ will be reduced. Thus, the size of the epoch T_e should reflect some lower limit on the duration of the EEG seizure;

(6) although the seizure LFM-type components can have slope values as large as ± 0.1 Hz per second, the rate of slope change is rather small over an epoch of less than 20 seconds [6, 20].

Because the method assumes that the IF laws will vary slowly over time, which is correct for a slowly varying periodic signal, the method has a penalisation measure to prevent false detections of components that do not conform to this signal type. This is achieved by first specifying the rotations selected from $\hat{\eta}_i(\theta_k)$,

$$\tilde{\theta}_i = \arg\max_{\theta_k \in \Theta} \hat{\eta}_i(\theta_k), \tag{20}$$

and then defining the penalisation function as

$$c(\sigma) = \left(1 - \frac{\sigma}{w}\right), \tag{21}$$

where σ is the variance of $\tilde{\theta}_i$ over $i = 1, 2, 3$. The value w in (21) is a predefined weighting parameter in the range $\sigma < w < \infty$. Limiting w ensures that $0 < c(\sigma) < 1$. Within the range $\sigma < w < \infty$, as $w \to \infty$ then $c(\sigma) \to 0$ and as $w \to \sigma$ then $c(\sigma) \to 1$.
We then use the function $c(\sigma)$ to weight the epoch's test statistic $\hat{\eta}$; that is, let

$$\eta = c(\sigma)\hat{\eta}. \tag{22}$$

If the variance of the slope values σ is large, then $c(\sigma)$ will be small and thus reduce the value of η. Simply put, a large value for σ will penalise the test statistic. This is desirable as a large σ value indicates a signal type that is not slowly varying, and for the EEG signal, not a seizure signal [6, 20]. Conversely, when σ is small then $c(\sigma)$ will be small and the test statistic η will not be heavily penalised. A small σ value is indicative of a signal with a slowly varying IF, such as a seizure signal;

(7) is the seizure signal present?

$$\begin{aligned} \eta < \zeta, & \quad \text{no seizure,} \\ \eta > \zeta, & \quad \text{seizure present,} \end{aligned} \tag{23}$$

where ζ is the predefined detection threshold;

(8) iterate this whole process over the different epochs with an overlap window.

3.2.2. Limitations and Assumptions for Method. The method assumes that the EEG signal is a piecewise LFM signal where the end time points of the IF for the pieces, known as turning points, are located at the end of the TFD segments $t = iT_s$. The turning points for the EEG will not be an abrupt or sudden change in the IF law because, as others have observed [6, 20], EEG seizure typically has a continuous IF law that varies slowly and smoothly over time [6] and because the TFD provides some time-frequency smoothing of the components. Hence the method should be able to cope with the situation when the turning points are not located at values of iT_s. The results in the next section support this statement.

EEG background may have discontinuous LFM-like components, which could result in a false detection for the

TABLE 3: Detection results using newborn EEG.

Method	AUC	TDR = 1 − FDR
time-frequency matched filter	0.99	96%
time-frequency filter	1.00	96%

method. Two scenarios could cause this: if the discontinuous components are centred in time around iT_s, or if the discontinuous components are equidistant in time from iT_s, for $i = 1, 2, 3$. To ensure that these scenarios do not produce a large η value, the method uses a sliding window on the data with a significant overlap, larger than 75% of T_e. Thus, by shifting the EEG by a fraction of T_s, the LFM-like components will no longer be centred around the turning points or equidistant from the turning points and therefore the method should not produce a large ζ value. Again, the results in the next section support this statement.

The description of the method shows, in Figures 1 and 2, a piecewise LFM signal model without harmonic components. Whether the piecewise LFM model has harmonic components or not, the method will produce a large test statistic for both these signal types. This is because the harmonic components have IF laws that are parallel to the main component's IF law. Therefore, when the main component's IF laws are matched in the correlation process the harmonic components will also match and the method will produce a large η value. Also note that although this method was developed specifically for EEG it could be applied to other signals. Other applications, however, may require some slight modifications, such as adjusting the number of segments in Figure 1.

In the next section, the results section, we compare the two methods from this section—the ambiguity filter and the time-frequency correlator—with the two methods from the previous section—the time-frequency matched filter and the time-domain matched filter. For the time-frequency-matched filter, we use the method in Section 2.3. We do not use the version of the time-frequency-matched filter method proposed in [6], which we call the *time-frequency filter* method, because its performance is very similar to that of the time-frequency-matched filter. (The appendix shows the relative performance of these two methods.) There is no compelling reason to use the more complicated, and therefore more computational demanding, time-frequency filter instead of the classical time-frequency-matched filter method.

4. Results

The two methods from the previous section were initially developed to detect seizure events in newborn EEG data. We are interested in evaluating how they perform with other signals loosely described as slowly varying periodic signals with harmonic components. To do so, we test the methods on both simulated and real data and compare with the classical time-frequency and time-domain-matched filter methods.

Our methods use the following parameters. The time-frequency-matched filter method uses a template set with 20
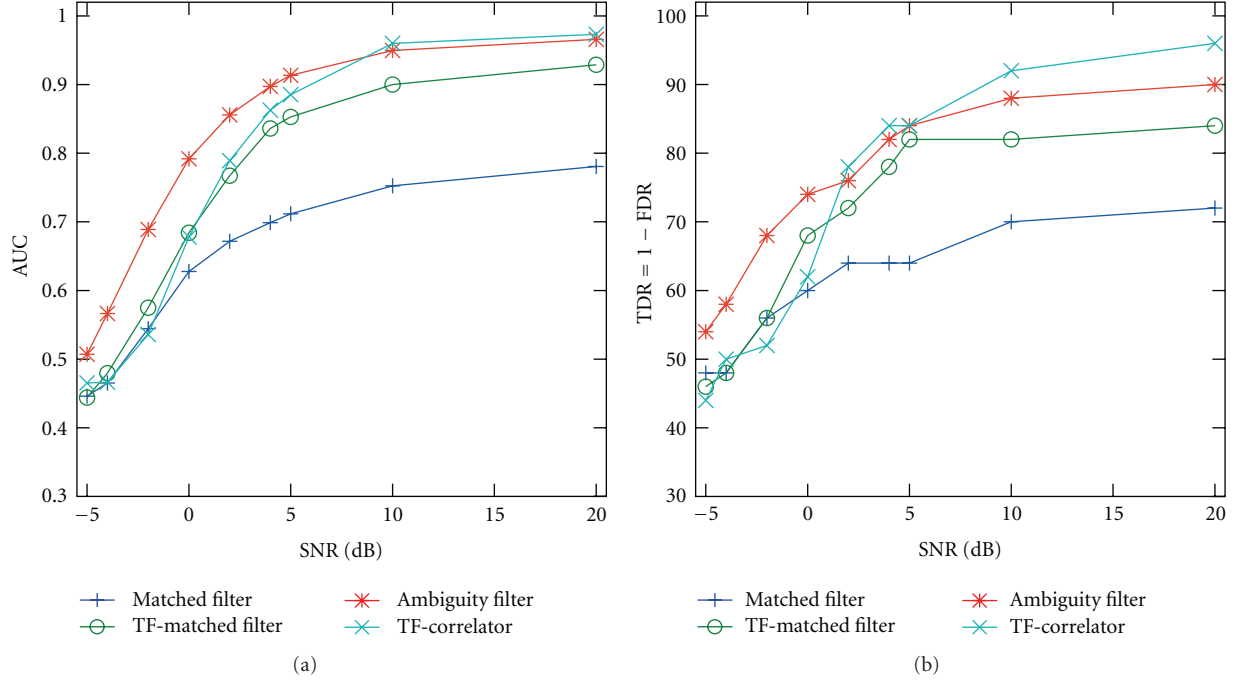
FIGURE 11: Detection results using speech signals. The test data is divided up into 50 epochs of noisy speech, for different levels of dB, and 50 epochs of noise.

template TFDs. Each template TFD uses a piecewise-LFM signal with two pieces: the parameter set is $[P_1, P_2, \alpha_1, \alpha_2]$ as $L = 2$. The length of the two pieces are both equal to one half of the epoch length; $P_1 = P_2 = T_e/2$; and the α (slope) values are from the set $[-0.06, -0.04, -0.02, 0, 0.02, 0.04, 0.06]$.

The ambiguity filter method uses 10 templates each with two LFM components: the parameter set is $[\alpha_1, \alpha_2]$, with α (slope) values in the same range as the previous method. The time-frequency correlator method uses 10 Θ (rotation) values in (16) using the same set as α from the previous methods.

The time-frequency-matched filter method and the time-frequency correlator methods both use TFDs. For these tests, we used a separable-kernel TFD [24]; this choice of kernel was justified from the analysis in [25] for finding a suitable kernel for an EEG seizure detection method. The separable time-frequency kernel used was $\gamma(t, f) = g_1(t)G_2(f)$, where $g_1(t)$ is a Hamming window of length $T_e/12$ and $g_2(\tau)$, the inverse Fourier transform of $G_2(f)$, is a Hanning window of length $T_e/3$. (T_e is the length of the epoch.)

The signals considered here are slowly-varying periodic signals (with harmonics). We use the term slowly-varying to mean that the variation of the IF over the length of the analysis window (epoch) is slow. That is, the IF does not vary or oscillate rapidly within the epoch. In the following tests, we ensure that the epoch is of a sufficient length to incorporate this slowly-varying behaviour.

4.1. Simulated Signals. The following procedure describes the method we used to generate a simulated slowly-varying periodic signal, with harmonic components. The test signal,

with K harmonic components, is given as

$$s(t) = \sum_{k=1}^{K+1} \sqrt{e_k} a_k(t) \cos\left(2\pi \int_0^t k f(\tau) d\tau + \theta_k\right), \quad (24)$$

where $a_k(t)$ is the time-varying amplitude modulation, e_k is an amplitude-scaling constant, $f(t)$ is the IF law, and θ_k is the initial phase of the component. The fundamental component of the signal is when $k = 1$, and the harmonic components are for $k = 2, 3, \ldots, K + 1$; always $e_1 = 1$, as we scale only the harmonic components and not the fundamental component.

To generate the time-varying amplitude modulation, we pick a series of time-amplitude location points $(t1_i, a_i)$, for $i = \{1, 2, 3\}$, and then create a smooth function $p(t)$ using cubic spline interpolation from this set of 3 points [20]. For the kth harmonic, we let $a_k(t) = p(t)$. To generate the IF law, we follow a similar procedure: generate a series of M time-frequency location points $(t2_i, f_i)$, for $i = \{1, 2, \ldots, M\}$ and interpolate using a cubic spline to produce the IF law $f(t)$.

We generate 50 realisations of the signal $s(t)$ using the random parameters in Table 1. For one set of 50 epochs, we set the parameter $g = 1$ in Table 1 and for another set of 50 epochs, we set the parameter $g = -0.1$. For example, if $K = 4$, then $e_k = \{0.5, 0.3333, 0.25, 0.2\}$ when $g = 1$ and $e_k = \{1.11111, 0.52632, 0.34483, 0.25641\}$ when $g = -0.1$. We shall refer to the first data set, using $g = 1$, as the *low-energy harmonic data set*, and the other data set, using $g = -0.1$, as the *high-energy harmonic data set*. We present two specific data sets for the signal as we want to assess the ability of the methods to detect signals with either low- or high-energy harmonic components. Example signals are plotted in time–frequency domain in Figures 3 and 4.
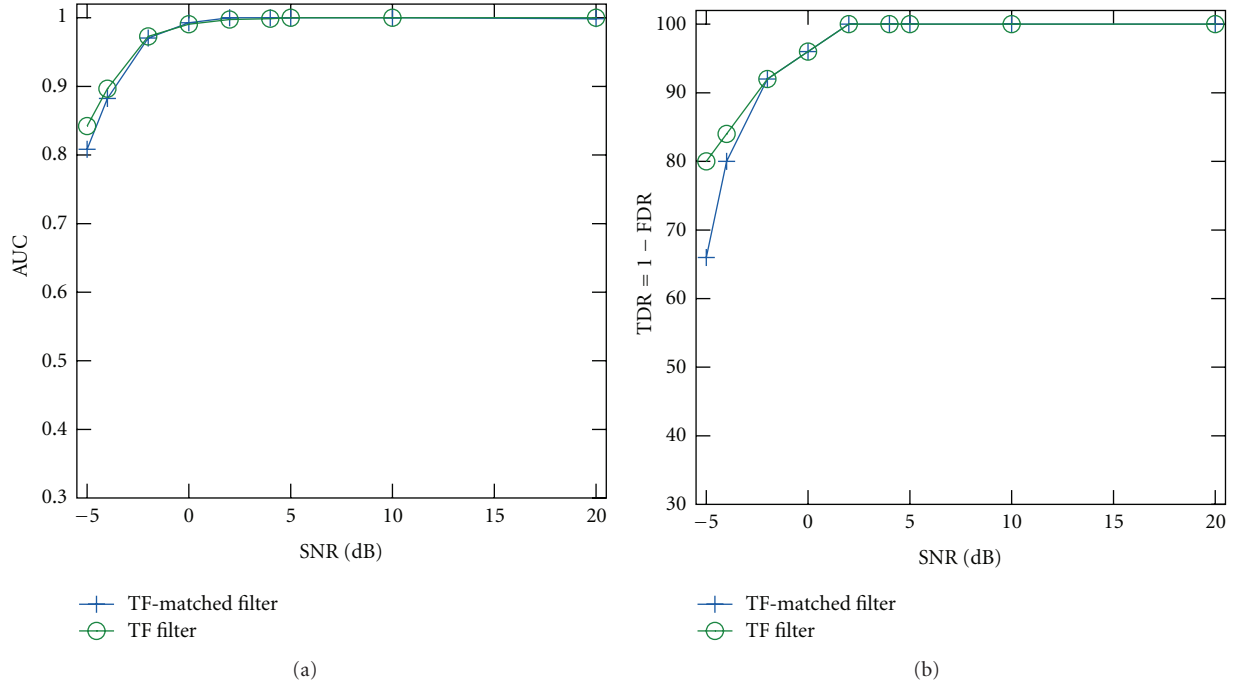
FIGURE 12: Detection results for time-frequency filter method and time-domain-matched filter in white Gaussian noise with low-energy harmonics.

We generated both white Gaussian noise and coloured Gaussian noise. For the coloured noise, we used the method described in [26] to generate the random process with a $1/f^{0.5}$ spectral power law. We combined these two noise processes with two signal sets, generating four data sets in total.

The first data set we test uses the *low-energy harmonic* signal model with white Gaussian noise. We tested 50 epochs of noisy signal and 50 epochs of noise to produce a receiver operating characteristic curve. This receiver operating characteristic is produced by comparing the test statistic with a range of threshold values, as described in (5). Each method defines the test statistic differently: for the time-frequency matched filter we used the maximum value of $\eta_{TF}(t, f)$ in Section 2.3 as the test statistic; for the ambiguity filter method we used the length of the IF extracted from $\eta_{TF}(t, f)$ in (11); for the time-frequency correlator method, we used the value η in (22).

The receiver operator characteristic is a plot of true detection rate, the rate of correctly detected events, often called sensitivity versus the false detection rate, the rate of incorrectly detected events, often called specificity. To summarise the receiver operator characteristic function, we use two measures: the area under the curve (AUC) [27] and the point where the sensitivity equals the specificity, or where the true detection rate equals one minus the false detection rate, assuming that the false detection rate is in the range $[0, 1]$. The AUC is a common summary measure for the receiver operator characteristic plot used to compare two or more detection methods [28]. The AUC has a range from 0 to 1: a perfect detection method has an AUC value

of 1, with 0% false detection rate (FDR) and 100% true detection rate (TDR) for all threshold values; and a random-guessing method has an AUC value of 0.5, a lower limit for a realistic detector [28]. The point of equal sensitivity and specificity, or TDR = 1 − FDR, represents one point on the receiver operator characteristic curve—again a summary measurement to compare two or more detection methods.

We varied the SNR ratio of the noisy signal from 20 dB down to −5 dB, and calculated the AUC and point of equal sensitivity specificity at each SNR value. The detection results for this first data set are in Figure 5. For the second data set we again used white Gaussian noise, but this time with the high-energy harmonic signal model. These results are in Figure 6.

The next two sets used coloured Gaussian noise. The results for the set with the low-energy harmonic signals are plotted in Figure 8 and for the high energy harmonic signals are plotted in Figure 6.

For the signals with low-energy harmonic components embedded in white Gaussian noise, the time-domain-matched filter and the ambiguity filter has the most accurate detection performance. If the reference signals from the template sets match the signals exactly, then the time-domain-matched filter is the optimum detection method for this data set [11]. Because of this method's accurate detection performance, with AUC > 0.95 and equal sensitivity specificity >95% over all SNR values, we assume that the template set represents a good match to the signals in the data set.

The trend is repeated for the test data set with high-energy harmonic components (again embedded in white Gaussian noise): the ambiguity filter and the time-domain-matched filter methods outperform the other two methods.
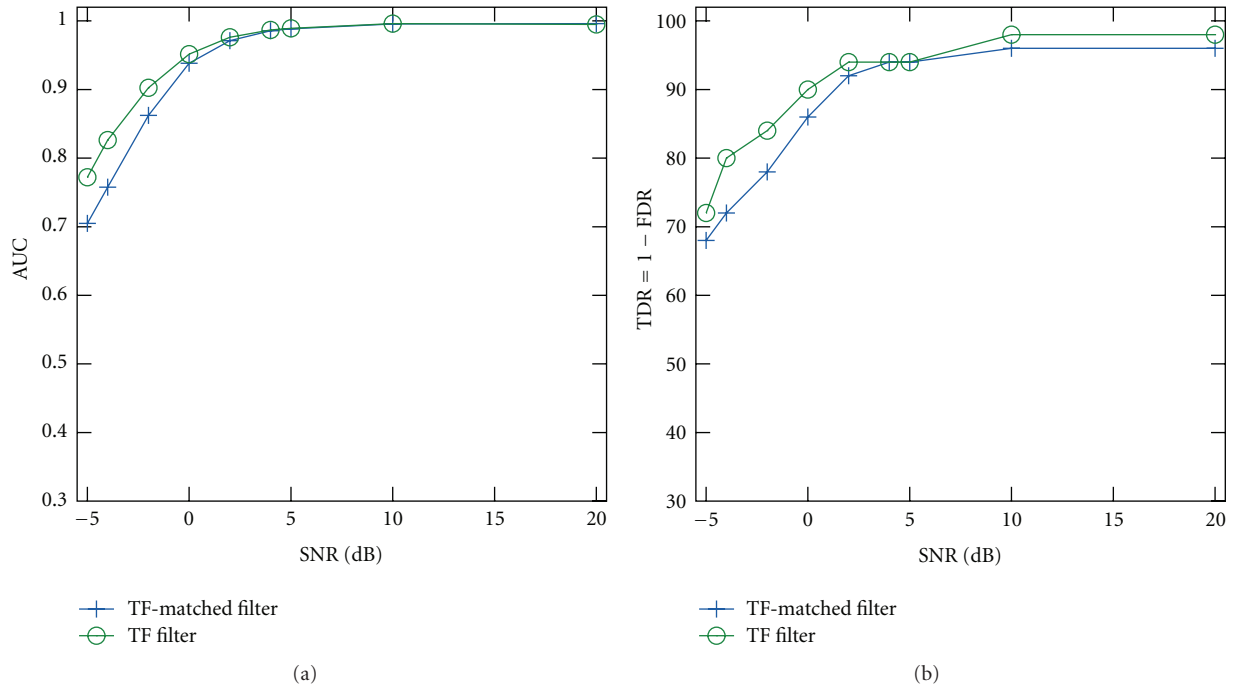
(a)

(b)

FIGURE 13: Detection results for white Gaussian noise with high-energy harmonics.


(a)

(b)

FIGURE 14: Detection results for coloured Gaussian noise with low-energy harmonics.

The result is surprising: these two methods (matched filter and ambiguity methods) use template sets without harmonic components. Thus even for signals with high-energy harmonic components, methods which do not account for these harmonic components give comparable, or even more accurate results, compared with the performance of the method which does account for these harmonic components, that is, the time-frequency correlator method.

We conclude for these two data sets, with white Gaussian noise, that the methods which perform time-frequency smoothing—the time-frequency matched filter and time-frequency correlator methods—do not perform as well as the

FIGURE 15: Detection results for coloured Gaussian noise with high-energy harmonics.
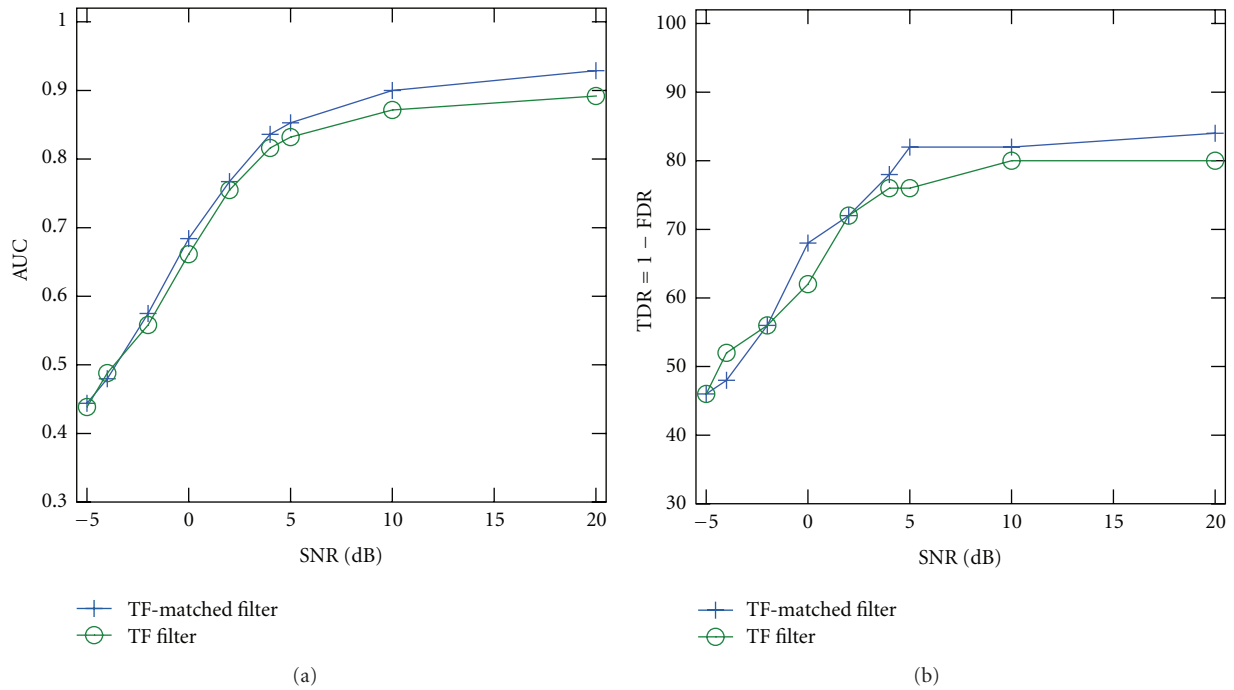

FIGURE 16: Detection results for the speech data set.

methods with no time-frequency smoothing—the ambiguity filter and the time-domain-matched filter methods.

The situation for the data sets with coloured Gaussian noise, in Figures 7 and 8, is somewhat different. For the data sets with both low- and high-energy harmonic components, the time-frequency correlator method has more

accurate detection performance compared with the detection performance for the other three methods. This difference is prominent for the high-energy harmonic component case in Figure 8: the time-frequency correlator method is significantly more accurate, approximately 20% more accurate, than the time-domain-matched filter.

4.2. Real-World Signals. As these methods were initially developed for the specific purpose of detecting seizure events in newborn EEG, we evaluate the performance of the methods on an EEG data set. This EEG data set consists of 100 epochs of EEG seizure and nonseizure events, recorded from 6 babies at the Royal Brisbane and Women's hospital, Australia. (See Figure 9 for an example epoch in the time-frequency domain.) The EEG was bandpass filtered in the range 0.5 to 10 Hz and downsampled from 256 Hz to 20 Hz. We segmented the EEG into epochs of 12.8 seconds with 256 sample points per epoch. The results for the three methods are in Table 2.

All three time-frequency methods detect seizure events accurately, with AUC values greater than 0.9 for the three methods. A surprise result, however, is the performance of the time-frequency-matched filter: a near perfect AUC score of 0.99, which is larger than the ambiguity filter's AUC (0.91) and time-frequency correlator's AUC (0.95). We note that although EEG seizure can be modelled by piecewise LFM signal with harmonic components [20], it does not always contain harmonic components [5, 7]. This, coupled with the explanation that energy in the harmonic components for EEG seizure could be relatively low, may be why the time-frequency-matched filter performs well. (Recall that from simulated results this same method was performed poorly when the energy of the harmonic components were high.) A useful outcome from this test would be an investigation to further explore this method as an automated EEG seizure detection method for use in a clinical setting.

Although the ambiguity filter and the time-frequency correlator methods were developed for detecting EEG seizures, they can be used to detect any signal described as a slowly-varying periodic signal with harmonics, as shown with the simulated data previously. We therefore also tested the methods using another real-world signal—recorded speech.

The speech test data set consists of 50 epochs of speech, where we manually selected a spoken word from speech signals taken from http://www.voxforge.org/. For the noise signal, we used a recording of a large crowd talking, often referred to as *babble* noise, produced by TNO, Soesterberg, The Netherlands, and obtained from http://spib.rice.edu/. We randomly segmented this long signal into 50 epochs and added to the speech signals at a range of SNR values.

The speech signal has a sampling frequency of 48,000 Hz and the babble noise has a sampling frequency of 19,980 Hz. We downsampled both speech and noise signals to 3,428.6 Hz after low-pass filtering with a cutoff frequency of 1,500 Hz. Each epoch contained 1,024 sampled points. See Figure 10 for an example epoch in the time-frequency domain. The detection results are plotted in Figure 11.

Note that the detection results rapidly decline from approximately 0 dB to −5 dB for all methods. For larger SNR values, however, the ambiguity filter and time-frequency correlator methods' detection performance is more accurate compared to the performance of the time-frequency-matched filter. This may be because the speech signals typically have many harmonic components and therefore the ratio of energy of the fundamental to the total energy in the harmonic components is small.

These methods could be used as a voice activity detector but would need to be compared with existing methods, such as [29], to assess their relative efficacy.

5. Conclusions

Moving away from the restrictions and assumptions necessary for theoretical analysis, we may find for applications using real signals that detection performance differs to theoretical predictions. The time-domain-matched filter method is an optimum method for detecting a known signal embedded in white Gaussian noise, but as we found in this paper, performance varies significantly for this method when the assumptions and restrictions are not satisfied. The time-frequency-matched filter has detection performance gains over the time-domain-matched filter when the signal is not precisely known or when the noise is not white Gaussian noise [13, 15, 16]. Yet, not satisfied with this performance gain, we aimed to further improve the time-frequency-matched filter by proposing two new methods.

From the performance evaluation in this paper comparing detection accuracy, we can conclude the following. As expected [13, 15, 16], the three time-frequency detection methods offer significant performance gain over the classical time-domain method for detecting slowly-varying signals with harmonic components. But the relative performance of these time-frequency methods—predictably—depend on the signal and noise type. Thus our proposed methods, for specific signals in specific noise types, will have a detection performance gain over the existing time-frequency-matched filter method.

Finally, on the basis of our findings in this work, we recommend some future research directions for the time-frequency methods. First, performance could be improved: (a) incorporating time-frequency filtering, by masking the TFD of the receiver signal in the time-frequency domain [18]; (b) using the correct discrete TFD definitions [30, 31]; (c) by optimising the parameters of the methods to suit particular data types, for example, the parameters of the signal model used in the template set. Second, the methods could be made more computationally efficient by incorporating recently proposed algorithms for the time-frequency-matched filter [18]. Lastly, the methods could be tested on detecting slowly-varying periodic signals without harmonic components, as all three methods can detect either signal type.

Appendix

Variant of the Time-Frequency-Matched Filter

The method detailed in [6, 7] varies somewhat to the method we called the time-frequency-matched filter in Section 2.3, as we now show:

   (1) form the TFD ρ_{eeg} for EEG signal eeg(t) for an epoch of length T;

(2) for each reference signal $r(t)$ form the TFD template ρ_r and

    (a) produce the time-frequency-matched filter test statistic $\eta_{\text{TF}}(t, f)$ using (10);

    (b) threshold $\eta_{\text{TF}}(t, f)$ to a predefined constant c, that is, if $\eta_{\text{TF}}(t, f) < c$ at the point (t_0, f_0) then let $\eta_{\text{TF}}(t_0, f_0) = 0$;

    (c) extract the instantaneous frequency (IF) from $\eta_{\text{TF}}(t, f)$;

    (d) the length of the continuous IF function is the test statistic;

(3) iterate over all epochs for the EEG.

The difference of the preceding method compared to the time-frequency-matched filter is in how the methods define the final test statistic: the method here defines the test statistics as the length of the (continuous) IF is extracted from $\eta_{\text{TF}}(t, f)$; the time-frequency-matched filter defines the test statistic as the maximum from $\eta_{\text{TF}}(t, f)$. To distinguish the two methods, we call the method here the *time-frequency filter* as the method is more akin to a TFD than a matched filter.

We do not present this time-frequency filter method for evaluation in Section 4 because of two reasons: first, as the following shows, the two methods (time-frequency filter and time-frequency-matched filter) have similar performance; and second, because extracting the IF is a more complicated procedure compared with just taking the maximum point, the time-frequency filter method is the more computationally expensive of the two methods.

The results here use the same signals and procedure from Section 4. Both methods, time-frequency filter and time-frequency-matched filter, use the same template set containing 20 sets of piecewise LFM signals; for each template the number of pieces is two, that is $L = 2$. Figures 12, 13, 14, and 15 show the results for the simulated signals from Section 4.1; Table 3 and Figure 16 show the results for the real signals from Section 4.2. These results justify our decision to include only the classical time-frequency-matched filter in the main comparisons with the two new proposed methods.

Acknowledgments

References

[1] M. Karimi-Ghartemani and M. R. Iravani, "Measurement of harmonics/inter-harmonics of time-varying frequencies," *IEEE Transactions on Power Delivery*, vol. 20, no. 1, pp. 23–31, 2005.

[2] A. B. Horner, J. W. Beauchamp, and R. H. Y. So, "Detection of time-varying harmonic amplitude alterations due to spectral interpolations between musical instrument tones," *Journal of the Acoustical Society of America*, vol. 125, no. 1, pp. 492–502, 2009.

[3] D. J. Liu and C. T. Lin, "Fundamental frequency estimation based on the joint time-frequency analysis of harmonic spectral structure," *IEEE Transactions on Speech and Audio Processing*, vol. 9, no. 6, pp. 609–621, 2001.

[4] P. Celka, B. Boashash, and P. Colditz, "Preprocessing and time-frequency analysis of newborn EEG seizures," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 5, pp. 30–39, 2001.

[5] B. Boashash and M. Mesbah, "A time-frequency approach for newborn seizure detection," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 5, pp. 54–64, 2001.

[6] B. Boashash and M. Mesbah, "Time-frequency methodology for newborn electroencephalographic seizure detection," in *Applications in Time-Frequency Signal Processing*, A. Papandreou-Suppappola, Ed., chapter 9, pp. 339–369, CRC Press, Boca Raton, Fla, USA, 2003.

[7] B. Boashash and M. Mesbah, "Using DSP to detect seizures in newborns," *IEE Electronics Systems and Software*, vol. 1, no. 6, pp. 34–37, 2003.

[8] B. Boashash, Ed., *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, Elsevier, Oxford, UK, 2003.

[9] B. Boashash, "Note on the use of the Wigner Distribution for time-frequency signal analysis," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 36, no. 9, pp. 1518–1521, 1988.

[10] C. W. Helstrom, *Elements of Signal Detection and Estimation*, Prentice Hall, Englewood Cliffs, NJ , USA, 1995.

[11] P. Z. Peebles, *Probability, Random Variables, and Random Signal Principles*, McGraw-Hill, Sinapore, 1993.

[12] B. Boashash and P. O'Shea, "A methodology for detection and classification of some underwater acoustic signals using time-frequency analysis techniques," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 11, pp. 1829–1841, 1990.

[13] P. Flandrin, "Time-frequency formulation of optimum detection," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 36, no. 9, pp. 1377–1384, 1988.

[14] A. M. Sayeed, "Optimal time-frequency detectors," in *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, B. Boashash, Ed., chapter 12, pp. 500–509, Elsevier, Oxford, UK, 2003.

[15] C. Richard, "Time-frequency-based detection using discrete-time discrete-frequency Wigner distributions," *IEEE Transactions on Signal Processing*, vol. 50, no. 9, pp. 2170–2176, 2002.

[16] E. Chassande-Mottin and P. Flandrin, "On the time-frequency detection of chirps," *Applied and Computational Harmonic Analysis*, vol. 6, no. 2, pp. 252–281, 1999.

[17] J. M. O' Toole, M. Mesbah, B. Boashash, and P. Colditz, "A new neonatal seizure detection technique based on the time-frequency characteristics of the electroencephalogram," in *Proceedings of the International Symposium on Signal Processing and its Applications (ISSPA '07)*, vol. 3, pp. 132–135, Sharjah, United Arab Emirates, February 2007.

[18] J. M. O' Toole, M. Mesbah, and B. Boashash, "Accurate and efficient implementation of the time-frequency matched filter," *IET Signal Processing*, vol. 4, no. 4, pp. 428–437, 2010.

[19] J. O'Toole, M. Mesbah, and B. Boashash, "Neonatal EEG seizure detection using a time-frequency matched filter with a reduced template set," in *Proceedings of the 8th International Symposium on Signal Processing and its Applications (ISSPA '05)*, pp. 215–218, August 2005.

[20] L. Rankine, N. Stevenson, M. Mesbah, and B. Boashash, "A nonstationary model of newborn EEG," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 1, article no. 6, pp. 19–28, 2007.

[21] G. F. Boudreaux-Bartels, "Mixed time-frequency signal transformations," in *The Transforms and Applications Handbook*, A. D. Poularikas, Ed., chapter 12, CRC Press, Boca Raton, Fla, USA, 2nd edition, 2000.

[22] M. A. Navakatikyan, P. B. Colditz, C. J. Burke, T. E. Inder, J. Richmond, and C. E. Williams, "Seizure detection algorithm for neonates based on wave-sequence analysis," *Clinical Neurophysiology*, vol. 117, no. 6, pp. 1190–1203, 2006.

[23] P. Celka and P. Colditz, "Nonlinear nonstationary Wiener model of infant EEG seizures," *IEEE Transactions on Biomedical Engineering*, vol. 49, no. 6, pp. 556–564, 2002.

[24] B. Boashash and G. R. Putland, "Design of high-resolution quadratic TFDs with separable Kernels," in *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, B. Boashash, Ed., chapter 5, pp. 213–222, Elsevier, Oxford, UK, 2003.

[25] J. M. O' Toole, *Discrete quadratic time-frequency distributions: definition, computation, and a newborn electroencephalogram application*, Ph.D. thesis, School of Medicine, The University of Queensland, Queensland, Australia, November 2009, http://espace.library.uq.edu.au/view/UQ:185537.

[26] N. J. Kasdin, "Discrete simulation of colored noise and stochastic processes and $1/f^\alpha$ power law noise generation," *Proceedings of the IEEE*, vol. 83, no. 5, pp. 802–827, 1995.

[27] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.

[28] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.

[29] R. Tahmasbi and S. Rezaei, "Change point detection in GARCH models for voice activity detection," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 16, no. 5, pp. 1038–1046, 2008.

[30] J. M. O'Toole, M. Mesbah, and B. Boashash, "A new discrete analytic signal for reducing aliasing in the discrete Wigner-Ville distribution," *IEEE Transactions on Signal Processing*, vol. 56, no. 11, pp. 5427–5434, 2008.

[31] J. M. O'Toole, M. Mesbah, and B. Boashash, "Improved discrete definition of quadratic time-frequency distributions," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, Article ID 5229150, pp. 906–911, 2010.