Improved Time-Frequency Filtering of Signal-Averaged Electrocardiograms

Akbar M. Sayeed, Paul Lander and Douglas L. Jones*

Journal of Electrocardiology, vol. 28, pp. 53–58, Suppelement, 1995.

Abstract

A recently proposed time-frequency filtering technique has shown promising results for the enhancement of signal-averaged electrocardiograms. This method weights the short-time Fourier transform (STFT) of the ensemble-averaged signal, analogous to the spectral domain Wiener filtering of stationary signals. In effect, it is a self-designing time-varying Wiener filter applied to the high resolution electrocardiogram (HRECG). In this paper, we empirically show that the performance of the proposed technique is about 2-3dB lower over the critical late-potential portion of the HRECG than the optimal fixed-window time-frequency filter based on ideal *a priori* knowledge of statistics. Although this ideal knowledge and performance is unattainable in practice, these results suggest that there remains potential for modest improvement. In order to narrow this gap in performance, we propose some improvements based on alternative structures for the time-frequency filter, including time-varying STFT windows. Simulation results show that an improved fixed-window technique can potentially yield an improvement of about 1-1.5 dB. By using properly chosen time-varying windows, the performance could potentially be improved even further. Thus, the improved techniques could produce a HRECG using fewer averages than the existing method, or could tolerate a lower signal-to-noise ratio.

*A. M. Sayeed and D. L. Jones are with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801. Their work was supported by the National Science Foundation under Grant No. MIP 90-12747, the Joint Services Electronics Program under Grant No. N00014-90-J-1270, and the Schlumberger Foundation. P. Lander is with the University of Oklahoma Health Sciences Center, Oklahoma City, OK 73190.

1 Introduction

Existing techniques for recording the high resolution electrocardiogram (HRECG) require signal averaging over 200 to 600 beats in order to attain satisfactorily high signal-to-noise ratios (SNRs). This translates to 3 to 10 minutes worth of measurements, which is an inconveniently long time in many situations [1]. Thus, there is a strong need to develop new HRECG techniques which could ultimately yield a high fidelity recording with a measurement time of 30 seconds or less [1].

Recently, recognizing the nonstationary nature of the ECG signal, a time-frequency filtering technique has been proposed which has shown significant improvement in performance over conventional ensemble averaging techniques [2, 1]. It involves weighting the short-time Fourier transform (STFT) of the ensemble-averaged signal and then recovering the signal from the modified STFT. The weighting function is designed *a posteriori*, based on the recorded measurements, and no *a priori* information is used except that the underlying signal of interest is deterministic in nature. Thus, the technique has wider applicability in situations where a deterministic nonstationary signal is to be estimated from noisy measurements.

Although the proposed technique is a significant improvement, it still falls short of the ultimate SNR requirements by a few dBs. As we will see in the next section, the time-frequency filter is in effect a self-designing time-varying Wiener filter. This observation yields useful insight into the structure of the filter and suggests ways in which the performance of the technique could be further enhanced. A natural question is whether such time-frequency techniques can achieve the required performance even in theory? We contend that the answer is a qualified "yes." We demonstrate, via simulations, that the optimal fixed-window time-frequency filter based on the ideal *a priori* statistical information can yield a 2-3dB improvement in performance. Although such ideal performance is obviously not attainable in practice, it nonetheless shows that there is room for modest improvement. In Section 3 we propose some improvements based on alternative time-frequency filter structures that have the potential of improving the performance of the time-frequency filter by up to 1-1.5dB, as evidenced by simulations.

2 Review of the Time-Frequency Optimal Filter

The time-frequency filter is applied to the ensemble-averaged HRECG. Let $x_i(t) = s(t) + n_i(t)$, $i = 1 \cdots N$ be an ensemble of N aligned beats, where $s(t)$ is the deterministic ECG signal and $n_i(t)$ is measurement noise which is modeled as Gaussian. The ensemble average is given by

$$\overline{x}(t) = \frac{1}{N}\sum_{i=1}^{N} x_i(t) = s(t) + \frac{1}{N}\sum_{i=1}^{N} n_i(t) = s(t) + \overline{n}(t) \ . \tag{1}$$

The STFT of the ensemble average given by

$$S_{\overline{x}}(t, f; \gamma) = \int \overline{x}(\tau)\gamma^*(\tau - t)e^{-j2\pi f\tau} d\tau \ , \tag{2}$$

where γ is the analysis window, is multiplied with a time-frequency weight $W(t, f)$ to yield a modified time-frequency representation $M_{\overline{x}}$ [2]

$$M_{\overline{x}}(t, f; \gamma) = S_{\overline{x}}(t, f; \gamma)W(t, f) \ . \tag{3}$$

The optimal weight $W(t, f)$ is given by a generalization of the spectral domain Wiener filter for stationary signals [2]. The optimal estimate of $s(t)$ is then obtained from $M_{\overline{x}}$ by using a STFT inversion technique.

There are many ways in which a signal $x(t)$ can be recovered from its STFT; some common ones are:

$$x(t) = \frac{1}{\gamma^*(0)} \int S_x(t,f;\gamma) e^{j2\pi ft} df \ , \quad \gamma(0) \neq 0 \tag{4}$$

$$x(t) = \frac{1}{\int \gamma^*(u)du} \int\int S_x(t',f;\gamma) e^{j2\pi ft} df dt' \ . \tag{5}$$

$$x(t) = \frac{1}{\int |\gamma(u)|^2 du} \int\int S_x(t',f;\gamma)\gamma(t-t') e^{j2\pi ft} df dt' \ . \tag{6}$$

Although these inversion formulas give identical answers for a valid STFT, they yield different answers, in general, when applied to the modified STFT, $M_{\overline{x}}$; in particular, the inversion formula (6) is optimal in a least-squares sense.

Interpretation as a time-varying filter. Simplifying the inversion formulas, as applied to the modified STFT $M_{\overline{x}}$, we find that the time-frequency filter can be equivalently expressed as a time-varying filter; that is

$$\hat{s}(t) = \int h(t,\tau)\overline{x}(\tau)d\tau \ , \tag{7}$$

where, for example, the filter $h(t,\tau)$ corresponding to the inversion formula (4) is given by

$$h_1(t,\tau) = \frac{1}{\gamma^*(0)}\gamma^*(\tau - t)w(t,\tau - t) \tag{8}$$

where $w(t,\tau)$ is related to $W(t,f)$ through a Fourier transform

$$w(t,\tau) = \int W(t,f)e^{-j2\pi f\tau}df \ . \tag{9}$$

3 Ideal Filter Based on Known Statistics

As we mentioned in the previous section, the role of the time-frequency weight $W(t,f)$ is analogous to the spectral domain Wiener filter. More precisely, the notion of a power spectral density is replaced by that of a time-varying spectrum; namely, the ensemble-averaged spectrogram (magnitude squared of the STFT). Since we are trying to estimate $s(t)$ from $\overline{x}(t) = s(t) + \overline{n}(t)$, the optimal weight $W_{opt}(t,f)$, analogous to the classical Wiener filter, is given by

$$W_{opt}(t,f) = \frac{|S_s(t,f;\gamma)|^2}{|S_s(t,f;\gamma)|^2 + E\{|S_{\overline{n}}(t,f;\gamma)|^2\}} = \frac{|S_s(t,f;\gamma)|^2}{|S_s(t,f;\gamma)|^2 + \frac{1}{N}E\{|S_n(t,f;\gamma)|^2\}} \ , \tag{10}$$

where E is the expectation operator and

$$E\{|S_n(t,f;\gamma)|^2\} = E\{|S_{n_i}(t,f;\gamma)|^2\} = \int\int R_n(\tau_1,\tau_2)\gamma^*(\tau_1 - t)\gamma(\tau_2 - t)e^{-j2\pi f(\tau_1-\tau_2)}d\tau_1 d\tau_2 \ , \tag{11}$$

$R_n(\tau_1,\tau_2) = E\{n(\tau_1)n^*(\tau_2)\}$ being the correlation function of measurement noise.

In [2], an estimate of $W_{opt}(t,f)$ obtained from the recorded ensembles is used in the time-frequency filter. Since an estimate of W_{opt} is used, the performance is not ideal. To obtain an upper bound on performance, we generated synthetic data using additive colored Gaussian noise and computed W_{opt}, using (10), from a known signal $s(t)$ and the noise statistics estimated directly from the noise realizations.[1] We compared the performance of this ideal filter with the one in [2], based on estimating W_{opt} from the noisy

[1] We *estimated* the noise statistics from the actual noise realizations instead of using an analytic *a priori* model for the noise statistics because it approximates more closely the ideal filtering achievable in practice.

signal ensemble, by using the performance criteria used in [1]. The results showed that the ideal filter provided a 2.65dB improvement in SNR over the self-designing filter. This indicates that modest improvements in performance are conceivable by using improved versions of the time-frequency filter proposed in [2].

Results of the simulation are illustrated in Figure 1. Figure 1(a) shows the underlying ECG signal which is to be estimated from an ensemble of 64 noisy beats. The simulated late potentials (small bumps) are primarily in the 140-200ms portion of the signal. Figure 1(b) shows a typical noisy beat (the *SNR* is -16dB) and Figure 1(c) shows the estimate based on simply averaging the 64 signal records, which is the estimation method currently used in most cases. The estimate based on the time-frequency (TF) filter in [2] is shown in Figure 1(e). Note that the TF filter processes the ensemble-averaged signal in Figure 1(c). And finally, Figure 1(d) shows the estimate based on the ideal filter. It is clear even visually that the ideal estimate is more refined in the crucial late-potential region of the signal. In this case, the ensemble average has an *SNR* of 1.3dB which is improved to 7.6dB by the self-designing filter of [2], and the ideal filter has the best performance with a 11dB *SNR*. The superiority of the ideal estimate is more clearly evident in Figure 2 in which the errors in the various signal estimates are displayed.

4 Improved Time-Frequency Filters

How can we refine the time-frequency filter proposed in [2] to narrow the gap in performance compared to the ideal filter? We propose three approaches based on the time-varying filter interpretation, alternative techniques for inverting the STFT, and time-varying STFT windows in order to better track the nonstationarities in the crucial late potential portion of the signal.

Estimating the time-varying filter directly. Since the overall operation is simply equivalent to that of a time-varying filter (see (7)), why not design the filter $h(t,\tau)$ directly? We assume that the signal characteristics are slowly time-varying and, thus, to estimate the signal at a particular time t_o we design a locally optimal *stationary* Wiener filter based on the statistics computed from a windowed portion of the recorded signals around t_o. More specifically, local estimates of the autocorrelation $R^{t_o}_{\overline{x},\overline{x}}(\tau)$ and the cross-corrrelation $R^{t_o}_{s,\overline{x}}(\tau) = R^{t_o}_{s,s}(\tau)$ are obtained from the recorded ensembles, using standard techniques, and then $h(t_o,\tau)$ is computed locally in the frequency domain by using

$$H(t_o,f) = \frac{S^{t_o}_{s,s}(f)}{S^{t_o}_{\overline{x},\overline{x}}(f)} \;, \tag{12}$$

where $S^{t_o}_{s,s}(f)$ and $S^{t_o}_{\overline{x},\overline{x}}(f)$ are the Fourier transforms of $R^{t_o}_{s,s}(\tau)$ and $R^{t_o}_{\overline{x},\overline{x}}(\tau)$, respectively.

Alternative techniques for inverting the STFT. As we mentioned in Section 2, there are many ways in which a signal can be recovered from its STFT which give different results when inverting a modified STFT. For example, the equivalent time-varying filter corresponding to the inversion formula (5) is given by

$$h_2(t,\tau) = \frac{1}{\int \gamma^*(t)dt} \int \gamma^*(\tau - t')w(t',\tau - t)dt' \tag{13}$$

which is different from the filter h_1 given in (8) corresponding to the inversion in (4). We note, by comparing (8) with (13), that, in effect, h_2 performs a slightly different smoothing in addition to that performed by h_1, and thus could potentially reduce the effect of noise even further. We note that in the original code developed for the filter in [2], an inversion of the form (5) is used and thus the filter is of the form (13).

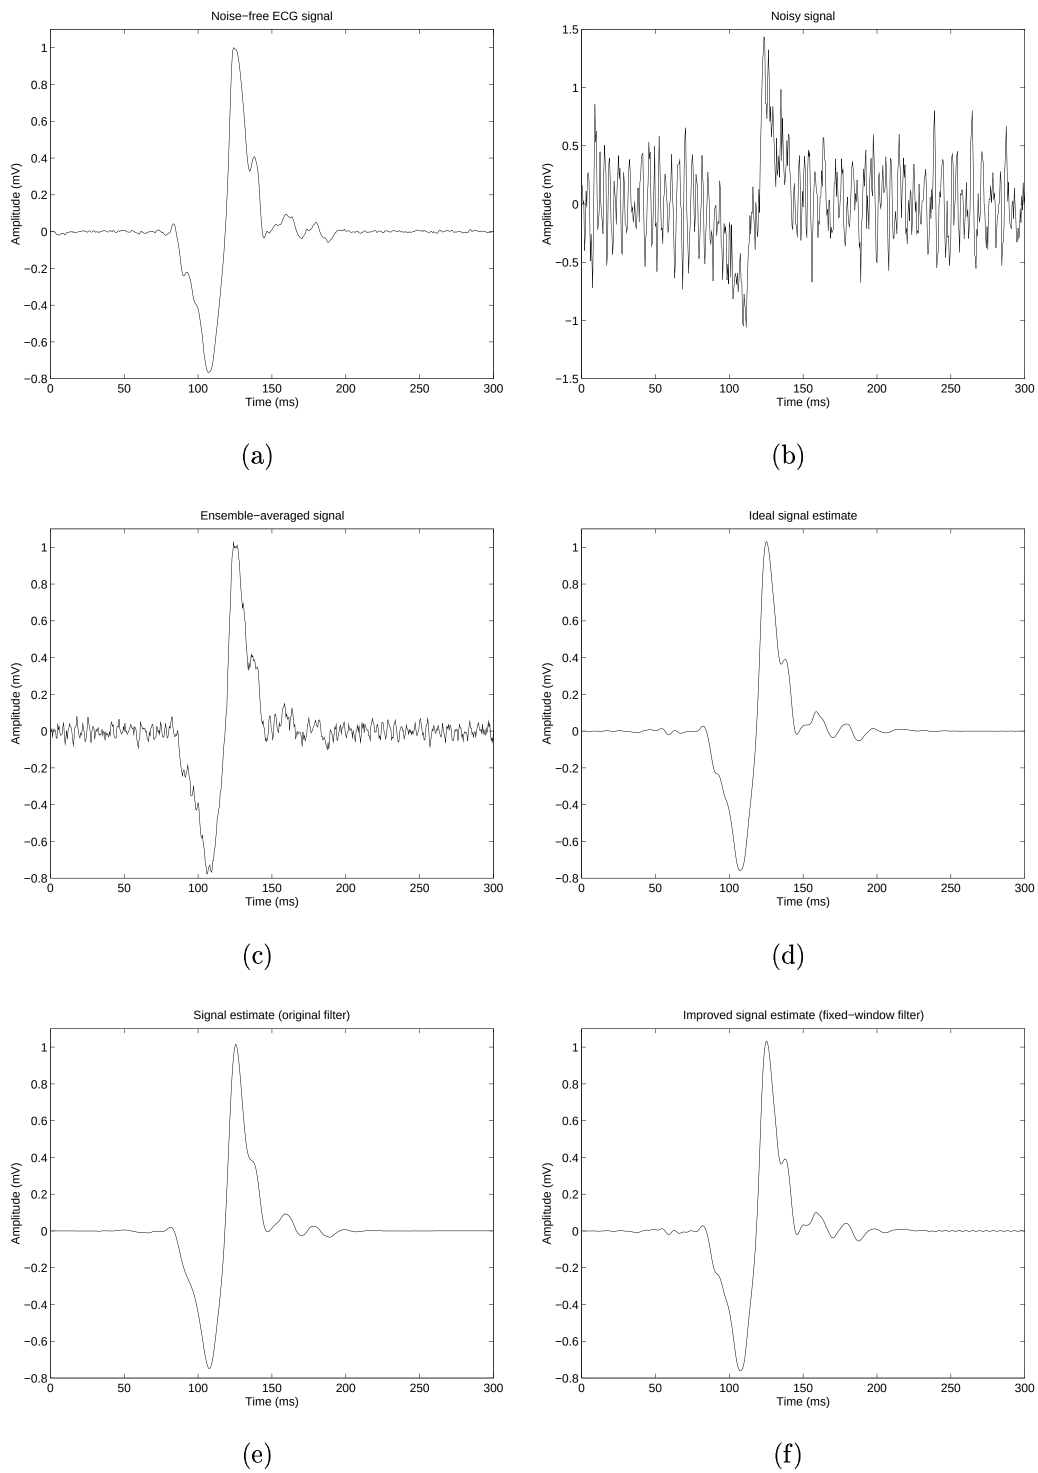

Figure 1: Simulated HRECG estimation using time-frequency filtering; $SNR = -16dB$, number of beats = 64. (a) The underlying noise-free ECG signal to be estimated. (b) A single noisy record. (c) Ensemble average of 64 noisy records. (d) ECG signal estimate based on the ideal TF filter . (e) Signal estimate based on the original TF filter proposed in [2]. (f) Signal estimate based on the improved fixed-window TF filter.

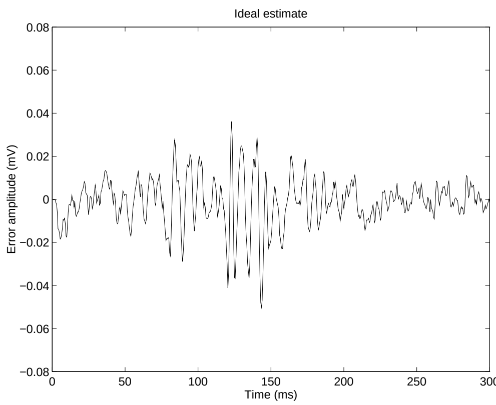

(a)

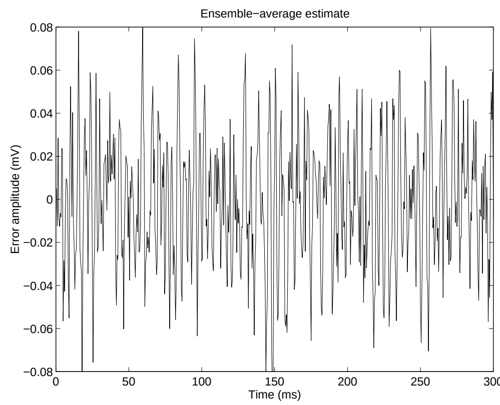

(b)

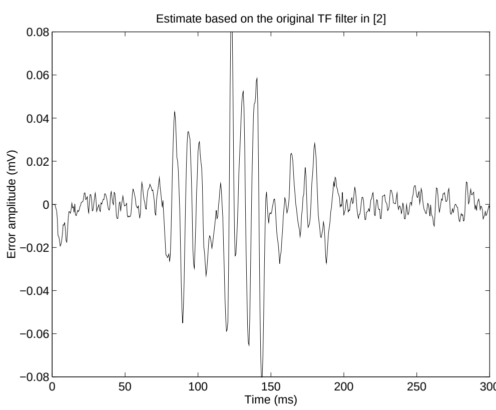

(c)

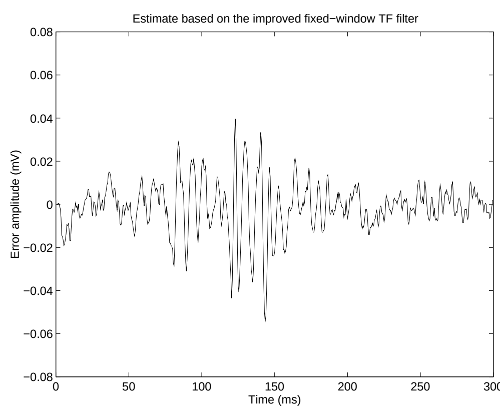

(d)

Figure 2: Error in the various time-frequency-filtered estimates. (a) Ideal estimate. (b) Ensemble-averaged signal (c) Estimate based on the original TF filter proposed in [2]. (c) Estimate based on the improved fixed-window TF filter.

The performance of such a fixed-window time-frequency filter based on the STFT inversion (5) is illustrated in Figures 1(f) and 2(d). It is evident visually from the two figures, in comparison with Figures 1(e) and 2(c), that the proposed filter performs better than the original filter proposed in [2], particularly in the late-potential region. For this particular simulation, the *SNR* of the improved estimate is 10.4dB as compared to 7.6dB for the original filter proposed in [2]. A more quantitative comparison is presented in the discussion of results in Section 5.

Time-varying STFT windows. A fixed window STFT has constant time and frequency resolution at all points in the time-frequency plane. However, from the very nature of the ECG signal, it is very clear that the time variation is relatively more rapid in the crucial late potential portion than the rest of the signal. Thus, in order to preserve such rapid time variation, it seems reasonable to use shorter windows in that portion of the ensemble-averaged signal. We ran a series of simulations in which the window lengths which optimized the performance criteria given in [1] were computed. The results, averaged over a sufficiently large number of experiments, indeed showed a time-varying optimal window pattern. Since such window length sequences cannot be estimated in practice, we inferred a monotonic window length sequence which fit the patterns reasonably well over the late potential region. The performance of this monotonically time-varying window length filter was then compared to the original filter proposed in [2], and the improved fixed-window filters proposed above.

5 Results

The proposed improved filters were applied to simulated ensembles of 64 realizations each and their performance was compared to the original time-frequency filter proposed in [2]. The performance criteria of bias, variance and SNR introduced in [1] were used for the comparison. The expected values of the performance measures were estimated by taking averages over 100 independent experiments. We discovered that designing the time-varying filter directly, as described in Section 4, did not yield improvements comparable to those based on using alternative STFT inversion techniques. Moreover, such a direct approach is computationally more expensive as well. Thus, we do not report the results for the approach based on directly designing the time-varying filter.

The results are summarized in Table 1. As expected, the ideal filter has the best performance with respect to all the performance measures. The two most relevant measures are *SNR* and variance. *SNR* measures the refinement in the signal estimate in the critical late-potential region, and variance measures the noise reduction achieved in the portion where the signal is not present. The measures are important because the critical parameters are the onset and end of late potentials as estimated from the filtered ensemble-averaged signal. As a rule of thumb, the bias should be less than 10%, which is true for all the TF filtering methods in Table 1.

As evident from Table 1, the ideal filter yields an *SNR* which is 2.65dB higher than that achieved by the original, fixed-window TF filter proposed in [2]. Our improved fixed-window filter, based on using the STFT inversion (5), narrows this gap in performance by about 1.3dB. We noted that the filter based on the inversion formula (5) yielded best results, although we have not exactly identified the differences between the implementation in [2] and our technique. Note that the improved fixed-window filter outperforms the original filter in all the measures. As evident from the table, the filter based on time-varying windows used in this study does not perform substantially better than the fixed-window filter. However, we believe that,

	Ideal	Signal average	Original	Improved (fixed-window)	Improved (time-varying)
SNR (dB)	11.30	2.15	8.65	9.97	10.01
Variance	0.0024	0.1095	0.0165	0.0065	0.0094
Bias	0.0593	0.1711	0.0802	0.0694	0.0688
MSE	0.0133	0.1083	0.0243	0.0182	0.0184

Table 1: Performance comparison of the various time-frequency filters. The numbers represent the expected values of the performance measures [1] estimated by averaging over 200 experiments.

with a proper choice of the time-varying window sequence, some further improvement is possible.

6 Conclusions

In this paper we partially addressed two questions: 1) What is the theoretical upper bound on the performance of time-frequency-based filters? and 2) Can such an upper bound be achieved with practical TF-based filters? To characterize the upper bound, we used the *ideal* (in the given framework) filter based on *a priori* knowledge of signal and noise statistics. Using simulation results, based on a particular, representative set of data, we demonstrated that the performance of the ideal filter is 2-3dB better than the original TF filter proposed in [2], thus indicating potential room for improvement. Using alternative STFT inversion techniques, we developed a slightly different TF filter which achieved a gain of about 1-1.5dB in performance. Furthermore, recognizing the time-varying structure of the ECG signal in the late-potential region, we proposed a TF filter based on time-varying window lengths. Although the filter based on time-varying windows, used in this study, did not yield significant improvements, we believe that such filters have potential and deserve further study, particularly regarding the choice of filter parameters. However, considering that the ideal performance is only 2-3dB higher than the original filter, gains beyond this may be difficult to achieve in practice. Nevertheless, such improved techniques could produce a HRECG using even fewer averages than the proposed method in [2], or could tolerate even lower signal-to-noise ratios.

A few remarks about the scope of our conclusions are warranted. First, we must emphasize that the reported results were based on simulations using a particular, representative model of the ECG signal corrupted by colored Gaussian noise which closely matches the real-life scenarios. Although the data set used were representative, a more detailed study is needed to verify the improvement in performance for real data. Second, there are certain parameters in the TF filters which have to be chosen appropriately based on the characteristics of the data being analyzed. For example, in the case of time-varying-window filter, we noticed that the particular choice of the monotonically increasing window sequence depended on the *SNR* of the data set for optimum performance. A more detailed study of the sensitivity of such filter parameters to real data characteristics would be useful.

References

[1] P. Lander and E. J. Berbari, "Performance assessment of optimal filtering of the high resolution electrocardiogram," *preprint.*

[2] P. Lander and E. J. Berbari, "Enhanced ensemble averaging using the time-frequency plane," in *IEEE Int'l Symp. on Time-Frequency and Time-Scale Analysis*, pp. 241–243, Philadelphia, 1994.