Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Noise Reduction for Heart Sounds Using a Modified Minimum-Mean Squared Error Estimator with ECG Gating

Anindya S. Paul, Eric A. Wan, and Alex T. Nelson

Abstract—**In this paper, we present a method for single channel noise reduction of heart sound recordings. Multiple noise sources, such as lung sounds, muscle contraction, and background noise can contaminate the heart sound collection making subsequent analysis difficult. Our approach is based on a spectral domain Minimum-Mean Squared Error (MMSE) Estimation, originally introduced by Ephraim and Malah in the context of speech enhancement [8]. This method uses a "decision-directed" approach to estimate the noise spectrum without the need for a separate reference signal. The noise spectrum is used to compute the SNR on-line for adapting the Wiener filter gain applied to the spectral amplitudes. A number of modifications are made to the baseline algorithm to increase the level of noise reduction while simultaneously reducing signal distortion. Enhancements include the use of a "soft" threshold to determine when to update the noise spectrum, a forward-backward filtering implementation (i.e., smoothing), and a "second-pass" iterative estimation scheme in which the residual noise is used to re-estimate the SNR and update the Wiener gains. In addition, ECG analysis is used to provide gating information on when desired heart sounds may be present in order to further guide the noise spectral estimation procedure.**

The noise reduction algorithm is tested as a front-end to an automatic heart sound analysis system. The sounds are collected through two sensors that act simultaneously as microphones and ECG electrodes. The proposed algorithm demonstrates improvements over existing noise reduction approaches in terms of SNR gain, qualitative evaluations, and automatic detection of abnormalities present in the heart sounds.

I. INTRODUCTION AND PROBLEM STATEMENT

The heart is an electromechanical system, and as such its health can be characterized by both its electrical behavior (evinced by the ECG) and its mechanical functioning. The latter can be observed through echocardiography, catheterization, or analysis of the sounds emitted by the heart. Normal individuals have at least two heart sounds – the first heart sound (S1) and the second heart sound (S2). S1 and S2 are produced due to the closure of the mitral and aortic valves at the beginning and end of the systolic process respectively [6]. Depending upon an individual's health, additional S3 and S4 sounds may occur. S3 is produced due to the filling of the ventricle in the early stage of diastole. If blood enters in a relatively "non-compliant" ventricle late in diastole, it generates S4. The four heart sounds are illustrated in Figure 1. The presence of an S3, for example, is a strong indication of congestive heart failure (CHF) [12].

Historically, heart sounds have been observed through auscultation of the heart wherein the cardiologist listens to the heart through a stethoscope. Forming a diagnosis based on sounds heard through the stethoscope is subjective, varies from person to person and also depends on a cardiologist's past experience. In recent years, there has been an effort automate the heart sound diagnosis operation in order to remove subjectivity and to make heart sound auscultation more quantitative. The Audicor system, developed by Inovise Medical, Inc., is an example of such a device that allows for the automated collection and analysis of heart sounds in the context of a standard 12-lead diagnostic ECG test. The sounds are collected through two sensors that act simultaneously as microphones and ECG electrodes; these sensors are placed in standard ECG locations on the chest. The acoustical and electrical signals are then analyzed to determine the presence or absence of abnormal heart sounds.

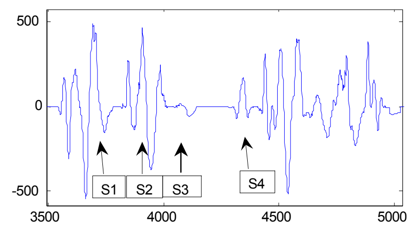

Figure 1. Heart Sound Signals with S1, S2, S3 and S4

The complex and nonstationary nature of the heart sounds makes automated heart sound analysis extremely difficult. Analysis is further complicated by noise in the audio signal. Noise sources arise due to interfering lung sounds, poor lead placement, patient motion, and other background noises. Unlike S1 and S2, S3 and S4 are low-amplitude and low-energy signals that can be buried under the noise floor. Unwanted noise may also result in the classifier producing false positives. The current system uses simple bandpass filtering with a cutoff between 1 and 100 Hz to remove high frequency noise. Our work in this project investigated more advanced signal processing techniques for noise reduction. The goal was to develop a robust noise reduction front-end to enhance the heart sounds with minimal signal distortion.

Previous work on noise reduction for heart sounds includes a variety of techniques. Adaptive noise cancellation approaches have been used successfully to enhance heart

Anindya S. Paul is a Ph.D. candidate in the Department of Computer Science and Electrical Engineering, Oregon Health & Sciences University, 20000 NW Walker Rd, Beaverton, OR 97006 (phone: 503-748-7521, e-mail: anindya@csee.ogi.edu).

Eric A. Wan, Ph.D., is an Associate Professor in the Department of Computer Science and Electrical Engineering, Oregon Health & Sciences University, 20000 NW Walker Rd, Beaverton, OR 97006 (phone: 503-748-1164, e-mail: ericwan@csee.ogi.edu).

Alex T. Nelson, Ph.D. is an Algorithm Development Engineer at Inovise Medical, Inc, 10565 SW Nimbus Avenue, Suite 100, Portland, OR 97223-4311 (e-mail: alex.nelson@inovise.com).

1-4244-0033-3/06/$20.00 ©2006 IEEE.

sounds, but require an additional reference to the noise signal [9,11]. The uses of short-time power spectrum analysis and wavelet filter based approaches have had moderate success with noise removal and on separating lung sounds from heart sounds [1,7]. The reduced order Kalman Filter has also been applied to improve SNR from noisy heart sounds [4], but requires an accurate model of the all heart sounds that may be of interest. Many other techniques for noise reduction have seen their origin in the field of speech enhancement. These enhancement approaches are commonly based on spectral subtraction (SS) [2,3,10], which use some estimate of the noise spectrum to subtract from the noisy signal. Classical SS often distorts the signal and introduces certain artifacts called "musical noise". Ephraim and Malah [8] proposed a minimum mean-square error (MMSE) based short-time spectral amplitude (STSA) estimator that is closely related to SS. This technique uses a "decision-directed" approach to estimate the SNR on-line in the spectral domain. The SNR is then used to determine a Wiener filter gain applied to the spectral amplitudes for estimating the signal. A number of variants on the Ephraim-Malah approach remain popular for speech enhancement [5,9,10,13,16].

We propose a number of modifications to the Ephraim-Malah approach in order to improve the noise reduction and fine tune the approach to the heart sound application. The system initializes the noise spectrum from a region in the noisy signal where the corresponding ECG values are very low. Additional ECG timing information is used to prevent the decision-directed noise updating from occurring during periods where S3 and S4 sounds may be expected. To reduce artifacts associated with rapid changes in the noise characteristics, the hard-threshold is replaced with a soft-threshold implementation in order to obtain smooth transitions. Forward-backward smoothing is also implemented to improve noise tracking and reduce transition effects before and after abrupt signal changes. Finally, a "second-pass" iterative estimation scheme uses the residual noise to re-estimate the SNR and update the Wiener gains, providing additional noise reduction without signal distortion.

The modified algorithm is used to preprocess noisy heart sound data before feeding as an input to the analysis and classification system. Performance comparisons are made between the base-line algorithms and our enhancements.

II. METHODOLOGY

A. MMSE Short-Time Spectral Amplitude Estimation

The MMSE short-time spectral amplitude (STSA) estimation technique is widely used in speech enhancement applications [5,8,13,14]. Our approach uses this as starting point.

Assuming additive noise, the observed signal is defined as:

$$\boldsymbol{y}(t) = \boldsymbol{x}(t) + \boldsymbol{n}(t), \tag{1}$$

where $\boldsymbol{x}(t)$ and $\boldsymbol{n}(t)$ denote the clean signal and noise respectively. The objective is to estimate the clean signal $\hat{\boldsymbol{x}}(t)$ from the noisy sensor observation $\boldsymbol{y}(t)$.

The noisy signal $\boldsymbol{y}(t)$ is first divided into overlapping frames $\mathbf{y}_r(t)$ (for the heart signals sampled at 1000 Hz, we use a T=40 ms Hamming window corresponding to N=40 samples with a 50% overlap). An FFT is performed on each windowed frame to compute the magnitude spectrum, $Y_r(\omega_k)$, where r denotes the frame index and ω_k the spectral bin. The phase $\theta_r(\omega_k)$ obtained from the FFT is retained for resynthesis. A spectral filter is then used to estimate the clean magnitude spectrum

$$\hat{X}_r(\omega_k) = G_r(\omega_k) Y_r(\omega_k), \tag{2}$$

where the Wiener gain $G_r(\omega_k)$ is specified as

$$G_r(\omega_k) = \frac{\mathrm{SNR}_{priori,\mathrm{r}}(\omega_k)}{1 + \mathrm{SNR}_{priori,\mathrm{r}}(\omega_k)}. \tag{3}$$

The computation of the spectral gain depends on the *a priori* SNR, defined as

$$\mathrm{SNR}_{priori,\mathrm{r}}(\omega_k) = \frac{E\left[|X_r(\omega_k)|^2\right]}{E\left[|N_r(\omega_k)|^2\right]} \tag{4}$$

where $X_r(\omega_k)$ is the clean signal spectrum and $N_r(\omega_k)$ is the noise spectrum at frame *r*. Computation of the *a priori* SNR assumes that the clean signal spectrum and noise spectrum are known. Since this is usually not the case, we need to estimate (and track) the SNR and noise spectrum online from the noisy signal.

The estimation of the *a priori* SNR is performed using a decision-directed method [8]. This approach updates the *a priori* SNR using the estimated signal amplitude from the previous frame, the current noise estimate, and the current *posterior* SNR estimate. Specifically,

$$\hat{\mathrm{SNR}}_{priori,r}(\omega_k) = \frac{\gamma |\hat{X}_{r-1}(\omega_k)|^2}{|\hat{N}_r(\omega_k)|^2} + (1-\gamma)\left[\mathrm{P}\left((\hat{\mathrm{SNR}})_{posteriori,r} - 1\right)\right], \tag{7}$$

where $\hat{X}_{r-1}(\omega_k)$ is the estimated clean signal spectrum at the previous frame, and $\hat{N}_r(\omega_k)$ is the current estimate of the noise magnitude spectrum. The *posterior* SNR is computed using the estimated noise spectrum as follows

$$\hat{\mathrm{SNR}}_{posteriori,\mathrm{r}}(\omega_k) = \frac{|Y_r(\omega_k)|^2}{|\hat{N}_r(\omega_k)|^2}. \tag{8}$$

In equation (7), P denotes half wave rectification,

$$\mathrm{P}(x) = \begin{Bmatrix} x, & \text{if } x \ge 0 \\ 0, & \text{otherwise} \end{Bmatrix}. \tag{9}$$

A detailed explanation and derivation of the decision directed approach is given in [8].

The noise magnitude spectral estimate, $\hat{N}_r(\omega_k)$, is needed in equations (7) and (8), and is updated according to a "spectral distance" based hard threshold rule,

$$\text{if } \sum_k \left[10\log_{10}\left\{ |Y_r(\omega_k)| - |\hat{N}_{r-1}(\omega_k)| \right\} \right] \le \text{th}$$

$$\hat{N}_r(\omega_k) = \alpha \hat{N}_{r-1}(\omega_k) + (1-\alpha) Y_r(\omega_k) \tag{10}$$

This computes the average spectral distance between the noisy signal spectrum and the previously estimated noise spectrum. If this distance is less than a predetermined threshold, the noise spectrum is updated. Initialization is typically done during known periods of signal silence (e.g., at the start of the signal for speech enhancement).

The Wiener gain is applied in Equations (2) and (3) using the estimate $\hat{\text{SNR}}_{priori,r}(\omega_k)$ to derive the clean magnitude spectrum $\hat{X}_r(\omega_k)$. Finally, $\hat{X}_r(\omega_k)$ is combined with the original noisy phase $\theta_r(\omega_k)$ and an inverse FFT is taken to reconstruct the signal $\hat{\boldsymbol{x}}(t)$.

B. Algorithm Modifications

When the spectral based filter, as described above, is applied to the noisy heart sounds, a fair amount of noise reduction is achieved (see Experiments section for examples). However, it is also observed that the estimate contains a number of sporadic high frequency noise artifacts that detract from the quality of the reconstruction. In addition, some signal distortion is apparent, which is more severe when the amplitude of the heart sound is low. While the level of distortion may be acceptable if the automatic classifier is retrained on the processed data, even minor distortion is unacceptable for presentation to a cardiologist.

The cause of the noise artifacts and distortion are attributed to errors in estimating the instantaneous noise spectrum, $\hat{N}_r(\omega_k)$. This stems from either failing to accurately track the nonstationary noise, or allowing some of the low energy heart signals of interest to leak into the noise estimation. In order to combat this, a number of enhancements are proposed to the baseline algorithm. These are described next.

Soft-Thresholding:

The hard-thresholding approach used to determine noise updating is highly sensitive to the threshold value. In addition, an incorrect decision can results in small discontinuities in the estimate when the noise amplitude changes rapidly and is missed in the update. Noticeable distortions can also occur when the signal leaks into the noise estimate. To smooth this effect, we introduce a soft-thresholding approach, specified by

$$V_r = \sum_k \left[\left\{ |Y_r(\omega_k)| - |\hat{N}_{r-1}(\omega_k)| \right\} \right] \tag{12}$$

$$\tilde{P} = a \tanh\left\{ b\left(th - V_r\right) \right\} \tag{13}$$

$$\hat{N}_r(\omega_k) = \left(\alpha - \tilde{P}\right)\hat{N}_{r-1}(\omega_k) + \left\{ (1-\alpha) + \tilde{P} \right\} Y_r(\omega_k) \tag{14}$$

where the hyperbolic tanh function is used to approximate a hard threshold, with the constant a controlling the height and b controlling the approximate slope or severity of the threshold. With the added constraint $\alpha + a = 1$, when $\tilde{P}$ is close to *0* (noise dominant region), the noise updating process follows a weighted averaging between the previously estimated noise spectrum and current noise spectrum. When $\tilde{P}$ saturates near *-a* (signal dominant region), $\hat{N}_r(\omega_k) \approx \hat{N}_{r-1}(\omega_k)$, and no noise updating occurs. In between these two extreme, some noise updating occurs using the difference between the current noisy signal spectrum and previous estimate of noise spectrum. Nominal parameter settings chosen were a = 0.3, α = 0.7, b = 3 and th = 0.5.

ECG Gating:

The 12 channel ECG is used to first identify probable time windows for S1-S4. Using an adaptive threshold for each frequency band, Audicor determines the location of each S1 and S2 within the computed detection window. The S3 detection time windows are located using information within the ECG and the computed position of the S2 offset (S3 typically occurs within 100 ms after S2). The S4 detection time windows are located based on PQ intervals and Q-wave onset positions (S4 may be present within a window of 100 ms before S1).

Since it is critical not to confuse the signals associated with S3 and S4 with the noise statistics, the noise updating (Equations 12-14) is prevented during these time windows.

Forward-Backward Noise Estimation:

With nonstationary noise, it may be easier to track changes *backward* in time versus *forward* in time. Reversing time can also be effective for dealing with convergence before or after a signal spike (as typical with heart sounds), or to mitigate edge effects with ECG gating. Non-causal filtering is feasible since the heart sound recordings are only 10 seconds in duration. Thus we propose a smoothing scheme in which the noise spectrum is estimated along the forward and backward direction and then averaged to obtain the final noise spectrum.

Implementation of smoothing is straightforward. Equations 12-14 are first run in the forward direction to calculate $\hat{N}_r^{fwd}(\omega_k)$, which are saved for each frame. The same equations are then run with time reversed to calculate $\hat{N}_r^{back}(\omega_k)$. A weighted averaging is then used to obtain the smoothed estimate

$$\hat{N}_r(\omega_k) = \alpha \hat{N}_r^{fwd}(\omega_k) + (1-\alpha)\hat{N}_r^{back}(\omega_k). \tag{21}$$

This smoothed noise estimate is used in the (forward) decision directed update (Equation 7) to estimate the *a priori* SNR. Note that we also explored smoothing the SNR estimate directly by running Equation 7 forward and backward in time. However, this did not prove to be as effective as smoothing just the noise estimates.

Initialization:

Simply initializing the noise spectrum with the first observed frame can results in unwanted attenuation if heart sounds are present near the beginning of the recording. Instead, we first initialize to a region where the ECG levels are low. The noise update filter is then run in the forward direction to the end of the recording. All estimates are then discarded, except the noise estimate in the last frame which is used to initialize the backward filter. The backward filter's last estimate (i.e., the first frame) is then used to initialize the forward filter, which is run again for use in the final smoothed estimate.

Iterative Residual Noise Estimation (IRNE):

When applying the noise reduction procedure as described so far, it is observed that the filtered signal has a small residual noise floor. To reduce the residual noise further, we propose an iterative residual noise estimation step. The idea is that after filtering, the noise spectrum should be consistent with the estimate and given directly by

$$\hat{N}_r^{res}(\omega_k) = Y_r(\omega_k) - \hat{X}_r(\omega_k) \quad (22)$$

This new noise estimate is then used to update the SNR and re-evaluate the Wiener gain to find the clean spectral estimate

$$\hat{\mathrm{SNR}}_{priori,r}^{res}(\omega_k) = \frac{|\hat{X}_r(\omega_k)|^2}{|\hat{N}_r^{res}(\omega_k)|^2} \quad (23)$$

$$G_r^{res}(\omega_k) = \frac{\beta \hat{\mathrm{SNR}}_{priori,r}^{res}(\omega_k)}{1+\beta \hat{\mathrm{SNR}}_{priori,r}^{res}(\omega_k)} \quad (24)$$

$$\hat{X}_r(\omega_k) = G_r^{res}(\omega_k) Y(\omega_k) \quad (25)$$

Because the initial filtering procedure reduces most, but not all of the noise, the estimate in (22) can results in a slight over estimation of the SNR. Thus we add a multiplicative gain term $\beta = .6$, which can be used to adjust the desired amount of noise reduction.

In summary, modifications over the baseline algorithm include ECG windowing, soft-thresholding, forward-backward smoothing, smart initialization, and iterative residual noise estimation. Figure 2 shows an overall block diagram of the modified spectral based Wiener filter approach.

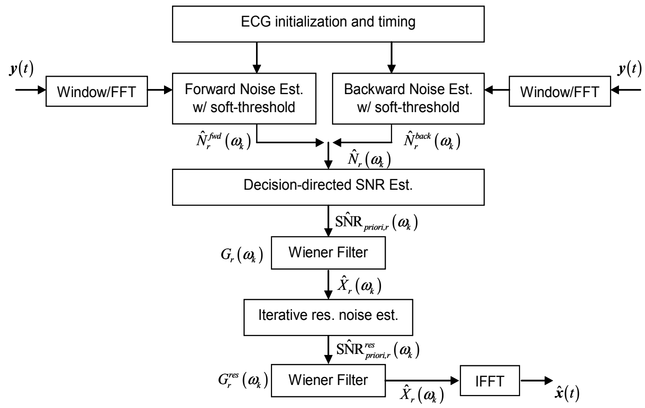

Figure 2: Modified MMSE STSA Algorithm

III. EXPERIMENTS

A. Comparisons

In this section, the performance of three different variants of our modified MMSE filter are compared with Ephraim-Malah's MMSE based STSA estimator, as well as a basic implementation of spectral subtraction [2]. Results are illustrated in Figure 3. The basic spectral subtraction algorithm not only severely distorts the heart sound but also eliminates some low amplitude components of the heart signal (see markers 1 and 2 in Figure 3b). The MMSE based STSA algorithm also distorts the heart signal by removing certain harmonics. As the algorithm estimates the noise using a hard thresholding approach, it is particularly severe

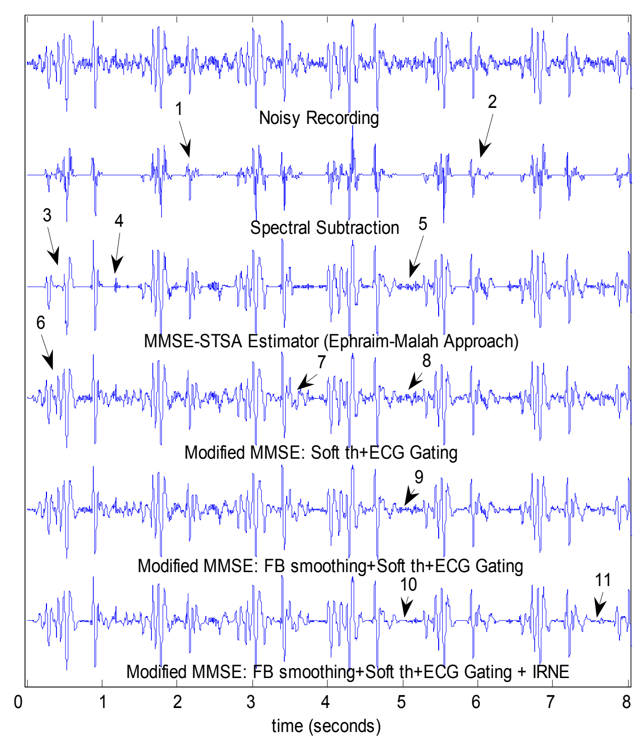

Figure 3: Comparison of noise reduction filters on a short segment of noisy heart sounds recorded using Audiocor.

when the heart sound amplitude is low. The filter eliminates some weak S2 and also affects S1 at the beginning of the signal (see markers 3 and 4 in Figure 3c). It also retains some high frequency noise artifacts (see markers 5 in Figure 3c). The modified MMSE filter with ECG gating and soft thresholding reduces the noise to an acceptable level and also preserves the low amplitude heart sound signals (see markers 6 and 7 in Figure 3d). The only concerns are the continued presence of a few high frequency residual noise artifacts (see marker 8). Adding the forward-backward smoothing further reduces the noise artifacts and overall noise floor (see marker 9 in Figure 3e). As a last step, adding the iterative residual noise estimation (IRNE) eliminates much of the remaining noise (see marker 10 and 11 in Figure 3f).

As evident from these plots, the modified MMSE filter with all the enhancements provides the best results. Significant noise removal is achieved with minimal signal distortion. These findings are consistent on multiple examples across a wide range of SNR levels.

B. Analysis Results

Figure 4 demonstrates the full application of the noise reduction approach using the Audicor system. The figure shows outputs of the system, corresponding to printed reports displaying automated ECG and heart sound findings. Clinicians would use these reports to make diagnostic and treatment decisions. Two different examples are shown. For each example we display the same report before and after processing with the noise reduction filter. In the top row, an S3 sound is correctly identified (along with S1 and S2) after noise reduction. In the bottom row, a significant improvement in noise reduction is achieved while the weak S4 sound is still preserved.

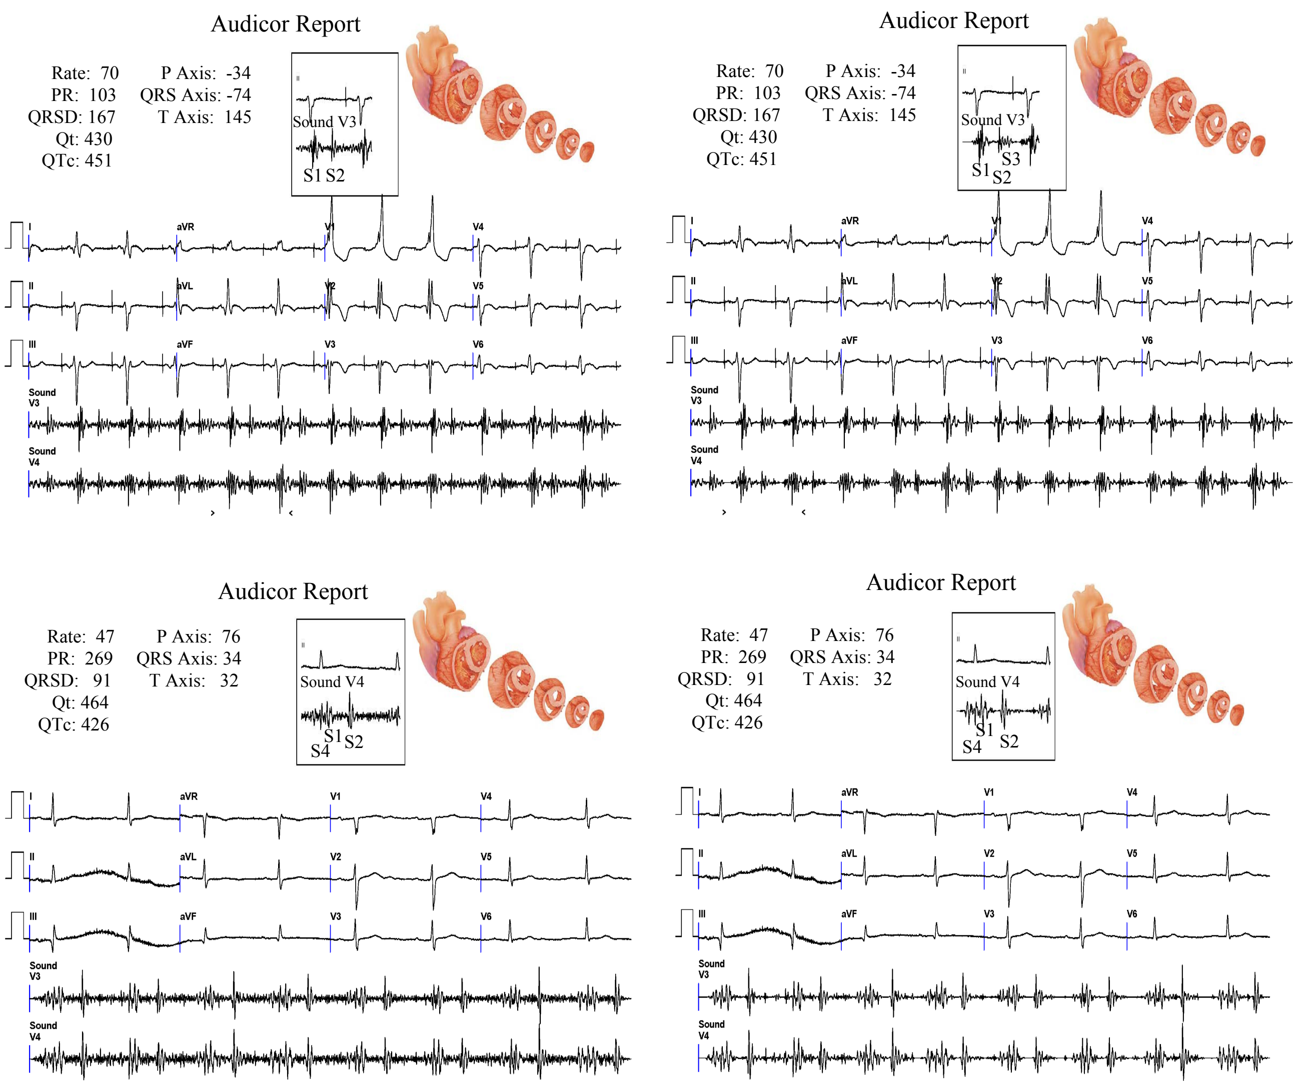

Figure 4: Two AUDICOR evaluations. a) Top left figure – noisy heart sound data, b) Top right figure - filtered heart sound data. (Note that Audicor is able to detect an S3 on the data following noise reduction), c) Bottom left figure - noisy heart sound data, d) Bottom right figure – filtered heart sound data (note that AUDICOR detects S4 on both the noisy and clean data). Minor formatting of standard printouts was performed to accommodate space.

These experiments indicate both improved diagnosis and signal preservation of important heart sounds. Note that the noise reduction also allows for analysis of some recordings that would have otherwise been discarded due to low SNR levels. In addition to the automatic diagnosis, the resulting plots appear more "pleasing" from a subjective perspective for displaying to a clinician.

C. *Clinical/Objective Evaluations*

The initial feedback from cardiologists indicates a preference with the added noise reduction filtering. For an objective evaluation of diagnostic improvement, we are in the process of testing a dataset for which we have a gold-standard diagnosis with clinical correlates for CHF. Initial findings show that the noise reduction results in a number of heart sound recordings became analyzable that would have been discarded because of low SNR. The noise reduction process also allows for the correct identification of new S3 and S4 sounds, which were previously hidden under the noise floor. On the downside, an increase in the number of false positives has also been observed. However, these are early results obtained without retraining the automatic classifier using the noise reduction front-end, as we expect will be necessary to fully optimize the system. Further tests should allow for statistical performance differences with and without noise reduction to be determined. In addition, changes in diagnosis that may result from the noise processing will be re-evaluated by a cardiologist to confirm (or reject) the diagnosis.

IV. Conclusions

In this paper, a noise reduction filter was developed for enhancing heart sounds. Starting with the Ephraim-Malah MMSE based STSA estimator, a number of enhancements were introduced. These include ECG gating, a soft thresholding rule, forward-backward noise estimation, and iterative residual noise estimation. Combined, these techniques are able to effectively remove various interfering noise sources while retaining critical low amplitude heart sounds. The approach was optimized for use as a front-end to the Audicor system. Initial tests indicate the improved ability to automatically detect S3 and S4 sounds. More extensive objective evaluations are now planned.

Furthermore, as two channels of acoustic signals are available, our next step involves fusing the two sources of information directly (as opposed to processing the channels independently). We are currently investigating both beamforming and Independent Component Analysis (ICA) techniques. These approaches will need to be integrated with the current modified MMSE approach and further optimized for the heart sound application.

Acknowledgements

We would like to thank Inovise Medical, Inc. for supplying the data and Audicor analysis software used in this paper.

References

[1] D. Barschdorff, U. Femmer and E. Trowitzsch, "*Automatic Phonocardiogram signal analysis in infants based on wavelet transforms and artificial neural networks*," In Computers in Cardiology, pp. 753-756, 1995.

[2] M. Berouti, R. Schwartz and J. Makhoul, "*Enhancement of Speech Corrupted By Acoustic Noise*," In Proceedings of IEEE Conference on Acoustic, Speech and Signal Processing (ICASSP), pp. 208-211, 1979.

[3] S.F. Boll, "*Suppression of Acoustic Noise in Speech using Spectral Subtraction*," In IEEE Transactions on Acoustic, Speech and Signal Processing, Vol. 27, No. 2, pp. 113-120, 1979.

[4] S. Charleston and M.R. Azimi-Sadjadi, "*Reduced Order Kalman Filtering for the enhancement of respiratory sounds*," In IEEE Transactions on Biomedical Engineering, Vol. 43(4), pp. 421-424, 1996.

[5] I. Cohen, "*Speech Enhancement Using a Noncausal A Priori SNR Estimator*," In IEEE Signal Processing Letters, Vol. 11, No. 9, pp. 725-728, September 2004.

[6] S.P. Collins, P. Arand, CJ Lindsell, WF Peacock, AB Storrow, "*Prevalence of the third and fourth heart sound in asymptomatic adults*," In Congestive Heart Failure, Vol. 11(5), pp. 242-247, 2005.

[7] L.G. Durand and P. Pibarot, "*Digital Signal Processing of the phonocardiogram: review of the most recent advances*," In Critical Reviews in Biomedical Engineering , Vol. 23(4), pp. 163-219, 1995.

[8] Y. Ephraim and D. Malah, "*Speech Enhancement Using a Minimum Mean Square Error Short-Time Spectral Amplitude Estimator*," In IEEE Transactions on Acoustics, Speech and Signal Processing, Vol. 32, No. 6, December 1984.

[9] Q. Fu and E.A. Wan, "Perceptual Wavelet Adaptive Denoising of Speech," In Proceedings of EUROSPEECH-2003, pp. 577-580, 2003.

[10] J. Gnitecki, Z. Moussavi and H. Pasterkamp, "*Recursive Least Squares Adaptive Noise Cancellation Filtering for Heart Sound Reduction in Lung Sound Recordings*," In Proceedings of the Annual International Conference on IEEE EMBS , pp. 2416-2419, 2003.

[11] R. Gemello, F. Mana and R.D. Mori, "*A Modified Ephraim-Malah Noise Suppression Rule For Automatic Speech Recognition*," In Proceedings of IEEE Conference on Acoustic, Speech and Signal Processing (ICASSP), pp. 957-960, 2004.

[12] M. Kompis and E. Russi, "*Adaptive Heart-Noise Reduction of Lung Sounds Recorded by a Single Microphone*," In Proceedings of International Conference on Engineering in Medicine and Biology Society , Vol. 2, pp. 691-692, Oct-Nov 1992.

[13] T. Lotter and P. Vary, "Noise Reduction By Maximum a Posteriori Spectral Amplitude Estimation with SuperGaussian Speech Modeling," In Proceedings of International Workshop on Acoustic Echo and Noise Control (IWAENC 2003), pp. 83-86, September 2003.

[14] G. Marcus, J. Vessey, M. Jordan, M. Huddlestom, B. McKeown, I.L. Gerber, E. Foster, K. Chatterjee, C. McCulloch, A.D. Michaels, "*Accurate Auscultation of a Clinically Useful Third Heart Sound Improves with Advancing Level of Experience*," In Archives of Internal Medicine. 2006 (In Press).

[15] R.J. Mcaulay and M.L. Malpass, "*Speech Enhancement Using a Soft-Decision Noise Suppression Filter,*" In IEEE Transactions on Acoustics, Speech and Signal Processing, Vol. 28, No. 2, April 1980.

[16] P. Scalart and J. V. Filho, "*Speech Enhancement Based On a Priori Signal to Noise Estimation*," In Proceedings of IEEE Conference on Acoustic, Speech and Signal Processing (ICASSP), pp. 629-632, 1995.