

Marquette University

e-Publications@Marquette

Dr. Dolittle Project: A Framework for
Classification and Understanding of Animal
Vocalizations

Research Projects and Grants

7-2008

Perceptually Motivated Wavelet Packet Transform for Bioacoustic Signal Enhancement

Yao Ren
Marquette University

Michael T. Johnson
Marquette University, michael.johnson@marquette.edu

Jidong Tao
Marquette University

Follow this and additional works at: https://epublications.marquette.edu/data_drdolittle



Part of the [Signal Processing Commons](#)

Recommended Citation

Ren, Yao; Johnson, Michael T.; and Tao, Jidong, "Perceptually Motivated Wavelet Packet Transform for Bioacoustic Signal Enhancement" (2008). *Dr. Dolittle Project: A Framework for Classification and Understanding of Animal Vocalizations*. 17.

https://epublications.marquette.edu/data_drdolittle/17

Perceptually motivated wavelet packet transform for bioacoustic signal enhancement

Yao Ren,^{a)} Michael T. Johnson, and Jidong Tao

Speech and Signal Processing Laboratory, Marquette University, P.O. Box 1881, Milwaukee, Wisconsin 53233-1881

(Received 14 December 2007; revised 18 April 2008; accepted 25 April 2008)

A significant and often unavoidable problem in bioacoustic signal processing is the presence of background noise due to an adverse recording environment. This paper proposes a new bioacoustic signal enhancement technique which can be used on a wide range of species. The technique is based on a perceptually scaled wavelet packet decomposition using a species-specific Greenwood scale function. Spectral estimation techniques, similar to those used for human speech enhancement, are used for estimation of clean signal wavelet coefficients under an additive noise model. The new approach is compared to several other techniques, including basic bandpass filtering as well as classical speech enhancement methods such as spectral subtraction, Wiener filtering, and Ephraim–Malah filtering. Vocalizations recorded from several species are used for evaluation, including the ortolan bunting (*Emberiza hortulana*), rhesus monkey (*Macaca mulatta*), and humpback whale (*Megaptera novaeanglia*), with both additive white Gaussian noise and environment recording noise added across a range of signal-to-noise ratios (SNRs). Results, measured by both SNR and segmental SNR of the enhanced wave forms, indicate that the proposed method outperforms other approaches for a wide range of noise conditions.

© 2008 Acoustical Society of America. [DOI: 10.1121/1.2932070]

PACS number(s): 43.60.Hj, 43.50.Rq, 43.80.Nd [WWA]

Pages: 316–327

I. INTRODUCTION

The presence of background noise and interfering signals is a fundamental problem in the collection and analysis of bioacoustic data, regardless of the specific species under study or the type of environment. This noise takes a variety of forms, including ambient background noise due to weather conditions, continuous interference from nearby vehicular or boat traffic, or the presence of numerous nontarget vocalizations from other species and individuals. Since the distance from the acoustic recording device to the individuals under study can be quite large leading to significant signal attenuation, interfering noise can create a substantial obstacle to analysis and understanding of the desired vocalization patterns.

Common techniques to reduce noise artifacts in bioacoustic signals include basic bandpass filters and related frequency-based methods for spectrogram filtering and equalization, often incorporated directly into acquisition and analysis tools (Mellinger, 2002). Other approaches in recent years have included spectral subtraction (Liu *et al.*, 2003), minimum mean-squared error (MMSE) estimation (Álvarez and García, 2004), adaptive line enhancement (Yan *et al.*, 2005; Yan *et al.*, 2006), and denoising using wavelets (Gur and Niezrecki, 2007).

In comparison, there are a wide variety of advanced techniques used for human speech enhancement, some of which form the basis for the more recent bioacoustic enhancement methods cited above. Historically the most common approaches for speech enhancement have focused on

spectral subtraction (Boll, 1979), Wiener filtering (Lim and Oppenheim, 1978), and MMSE and log-MMSE estimations using Ephraim–Malah (EM) filtering (Ephraim and Malah, 1984; 1985). Added to this in recent years are newer methods based on subspace estimation and filtering (Ephraim and Trees, 1995) and wavelet decomposition (Johnson *et al.*, 2007).

In this paper, we introduce a new bioacoustic signal enhancement technique which is based on a perceptually scaled wavelet packet decomposition, using spectral estimation methods similar to those used for human speech enhancement. The underlying goal is to obtain higher quality and more intelligible enhanced signals through the use of more perceptually meaningful frequency representations. This method is robust across a wide range of species, needing only f_{\min} and f_{\max} frequency boundary parameters to generalize for application to a new species of interest.

The new method is compared to a variety of other enhancement and denoising techniques, including simple bandpass filtering, spectral subtraction, Wiener filtering, and the EM log-MMSE estimation. To evaluate and compare its applicability across a variety of species, the method is applied to the animals of the order Passeriformes (ortolan bunting), Primates (rhesus monkey), and Cetaceans (humpback whale). Evaluation is done by using both signal-to-noise ratio (SNR) and segmental SNR (SSNR), which is known to be a more perceptually relevant quality measure for human speech (Deller *et al.*, 2000).

^{a)}Electronic mail: yao.ren@marquette.edu

II. CURRENT ENHANCEMENT METHODS

A. Bandpass filtering

Bandpass filtering removes signal energy outside of a specified frequency range. This can be applied in either the time domain or the frequency domain (e.g., applied to a spectrogram) and is effective primarily in cases where signals are predominately narrow band and are well separated from the noise spectrum.

B. Spectral subtraction

Spectral subtraction (Boll, 1979) was one of the first algorithms applied to the problem of speech enhancement. It is based directly on the additive noise model:

$$y(n) = x(n) + d(n), \quad (1)$$

where $y(n)$, $x(n)$, and $d(n)$ denote the noise-corrupted input signal, clean signal, and additive noise signal, respectively. The noise spectrum is estimated from the Fourier transform magnitude of a silence region in the wave form, so that for each frame of the signal, an estimate for the clean signal in the frequency domain can be given directly as

$$\hat{X}(\omega) = [|Y(\omega)| - |\hat{D}(\omega)|]e^{j\phi_y(\omega)}, \quad (2)$$

where $\phi_y(\omega)$ is the phase component of the noisy signal, used under the assumption that the spectral phase is much less important than the spectral magnitude for reconstruction.

Note that application of Eq. (2) may result in negative magnitude values, which are typically set to zero. This often results in some processing artifacts that are usually described by listeners as “musical tones.” The presence of such artifacts is one disadvantage of the spectral subtraction approach.

C. Wiener filtering

Wiener filtering is conceptually similar to spectral subtraction but replaces the direct subtraction with a mathematically optimal estimate for the signal spectrum in a MMSE sense (Lim and Oppenheim, 1978).

The frequency domain formulation of the Wiener filter is given as

$$H(\omega) = \frac{S_{xx}(\omega)}{S_{xx}(\omega) + S_{dd}(\omega)}, \quad (3)$$

where $H(\omega)$ is the desired filter response and $S_{xx}(\omega)$ and $S_{dd}(\omega)$ are power spectral densities (PSDs) of the desired clean signal and noise. Since these two PSDs are unknown, this filter cannot be determined directly and instead needs to be realized in an iterative fashion. In particular, $S_{dd}(\omega)$ is estimated from a silence region and $S_{xx}(\omega)$ is initialized from the noisy wave form and then updated from the output of the filter after each iteration. This process is repeated either a fixed number of times or until a convergence criterion is reached.

D. Ephraim–Malah filtering

The Wiener filter is an optimal linear estimator of the clean signal spectrum in a MMSE sense. Ephraim and Malah extended this idea by deriving an optimal nonlinear estimator of the clean spectral amplitude. This estimator assumes that the real and imaginary parts of the spectral magnitude have a zero-mean Gaussian probability density distribution and are statistically independent. Under this statistical model, a short time spectral amplitude estimator was derived by using the MMSE optimization criteria (Ephraim and Malah, 1984). This work was then modified to use log spectral amplitude (LSA) rather than spectra as an optimization criterion (Ephraim and Malah, 1985) since the log spectral distance is a more perceptually relevant distortion criteria, resulting in improved overall enhancement results. This estimator, known as the EM filter, can be summarized by using the following estimation formula for the clean signal Fourier transform coefficient \hat{A}_k in each frequency bin:

$$\hat{A}_k = \frac{\xi_k}{1 + \xi_k} e^{(1/2) \int_{v_k}^{\infty} (e^{-t}) dt} R_k, \quad (4)$$

In this equation, $\xi_k = \lambda_x(k) / \lambda_d(k)$, $v_k = [\xi_k / (1 + \xi_k)] \gamma_k$, and $\gamma_k = R_k^2 / \lambda_d(k)$, where R_k is the noisy speech Fourier transform magnitude in the k th frequency bin, and $\lambda_d(k)$ and $\lambda_x(k)$ are the average noise and signal powers in each bin. Similar to the spectral subtraction method, the noise power is estimated from silence regions in the wave form, while $\lambda_x(k)$ is a moving average of spectrally subtracted noisy spectra [$R_k^2 - \lambda_d(k)$]. The *a priori* SNR ξ_k is estimated via the EM well-known “decision-directed method,” which is updated from the previous amplitude estimate using a forgetting factor α as follows:

$$\hat{\xi}_k(n) = \alpha \frac{\hat{A}_k^2(n-1)}{\lambda_d(k, n-1)} + (1 - \alpha) P[\gamma_k(n) - 1], \quad (5)$$

where the indicator function P is given by

$$P[\gamma_k(n) - 1] \triangleq \begin{cases} \gamma_k(n) - 1, & \gamma_k(n) - 1 \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

The key characteristics of this estimator are that it tends to do less enhancement (i.e., less change to the noisy signal spectrum) when the SNR is high, and that musical noise artifacts are significantly reduced.

E. Wavelet denoising

Spectral subtraction, Wiener filtering, and EM filtering are all based on the same mathematical tool, the short time Fourier transform (STFT), with the waveform divided into short frames during which the signal is assumed to be stationary. The STFT is a compromise between time resolution and frequency resolution: a shorter frame length results in a better time resolution but poorer frequency resolution. The wavelet transform (WT) by comparison has the advantage of implicitly using a variable window size for different frequency components. This often results in better handling of broadband nonstationary signals, including speech and bioacoustic data.




FIG. 1. (a) Discrete WT. (b) Wavelet packet decomposition tree.

Whereas the STFT is a function of frequency for each individual signal frame, the WT is a function of two variables, time and scale. Scale is used rather than frequency because depending on the wavelet basis being used, each scale may actually represent information across a range of frequencies. Like the Fourier transform, the WT has both continuous WT and discrete WT (DWT) implementations. A DWT can be efficiently implemented by using a quadrature mirror filter decomposition, resulting in scales that are powers of 2, called a dyadic transform. A further generalization of the DWT is the wavelet packet transform (WPT). In the WPT, the filtering process is iterated on both the low frequency and high frequency components, whereas the DWT iterates only on the low frequency components. Filter decomposition structures for the DWT and WPT are shown in Fig. 1. In the decomposition tree, each node is labeled (l, n) , where l is the decomposition level and n represents a sub-band node index. The root of the tree, $(l, n) = (0, 0)$, refers to the entire signal space. The left and right branches denote low-pass and high-pass filterings followed by 2:1 downsampling, respectively.

The application of wavelets for signal enhancement, sometimes referred to as denoising, is a three step procedure involving wavelet decomposition, wavelet coefficient thresholding, and wavelet reconstruction. Given an appropriate choice of the wavelet basis function, the signal energy will be concentrated in a small number of relatively large coefficients while ambient noise will be spread out, allowing coefficients to be thresholded.

Threshold selection and implementation are two factors which significantly impact wavelet denoising methods. Common methods include hard, soft, and nonlinear thresholding approaches. Hard thresholding sets all coefficient values beneath the threshold to zero, leaving the others unchanged (Jansen, 2001); soft thresholding additionally reduces all coefficient values to maintain continuity; while nonlinear thresholding typically enforces a smoothness constraint on the coefficient mapping function as well. Typical threshold selection methods include universal thresholding and the Stein unbiased risk estimator (Donoho, 1995), both implemented by using soft thresholding.

Recently, the EM suppression rule (Ephraim and Malah, 1984) for speech enhancement has been applied to the wave-

let domain as a more advanced time-varying thresholding approach (Cohen, 2001). This method helps reduce the “musical noise” artifacts caused by uniformly applied thresholds.

III. PROPOSED METHOD

The method introduced here is based on a modified wavelet packet decomposition using a MMSE coefficient estimation for thresholding. The key element of the technique is the use of the Greenwood warping function to determine the WPT decomposition structure based on a perceptually motivated frequency axis.

Greenwood (1961) has shown that many land and aquatic mammals perceived frequency on a logarithmic scale along the cochlea, which corresponds to a nonuniform frequency resolution. This relationship can be modeled by the equation

$$A(10^{\alpha x} - k), \tag{7}$$

where α , A , and k are species-specific constants and x is the cochlea position. Transformation between true frequency f and perceived frequency f_p can be obtained through the following equation pair:

$$F_p(f) = (1/\alpha)\log_{10}(f/A + k), \tag{8}$$

$$F_p^{-1}(f_p) = A(10^{\alpha f_p} - k). \tag{9}$$

The constants α , A , and k can be found if frequency-cochlear position data are available. However, since cochlear information has never been measured for many species, an approximate solution is needed. Lepage (2003) has shown that k can be estimated as 0.88 based on both theoretical justification and experimental data acquired on a number of mammalian species. By assuming this value for k , α and A can be solved for a given approximate hearing range, $f_{\min} - f_{\max}$, of the species (Clemins, 2005; Clemins and Johnson, 2006; Clemins et al., 2006):

$$A = \frac{f_{\min}}{1 - k}, \tag{10}$$

$$\alpha = \log_{10}\left(\frac{f_{\max}}{A} + k\right). \tag{11}$$

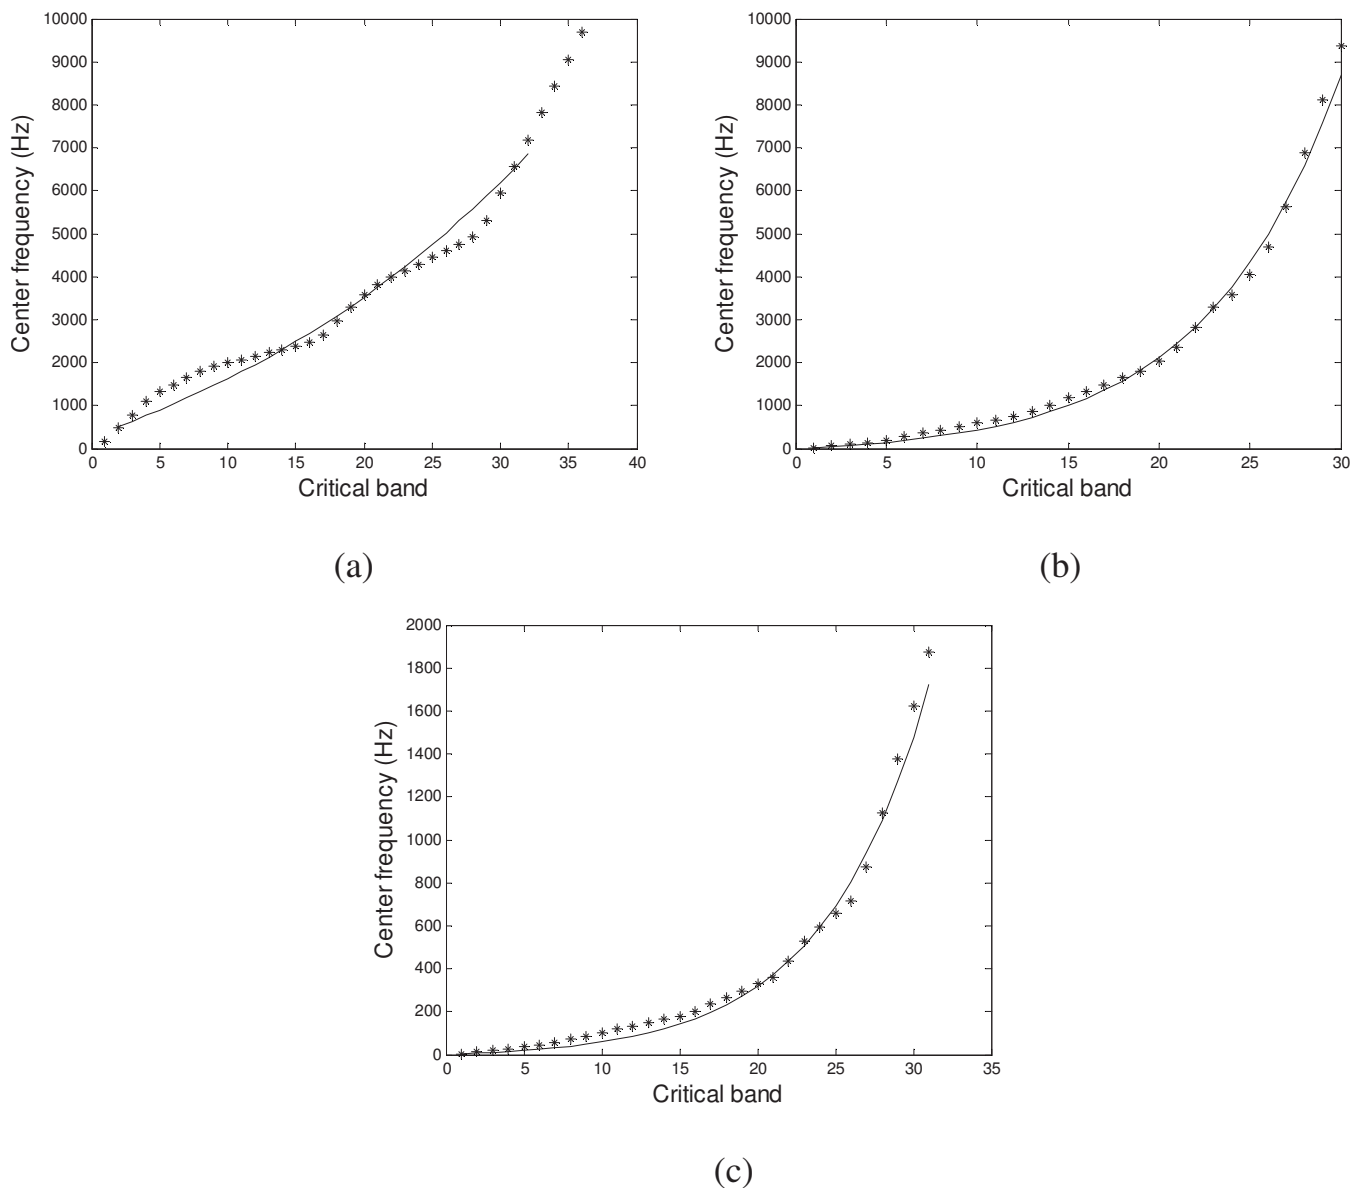


FIG. 2. Center frequencies of the Greenwood scale (solid line) and WPD critical bands. (a) Ortolan bunting. (b) Rhesus monkey. (c) Humpback whale.

Thus, a frequency warping function can be constructed by using the species-specific values of f_{\min} and f_{\max} .

A perceptually motivated WT can be designed to mimic the auditory frequency scale by using decomposition critical bands. This implementation was originally proposed by Black for coding (Black and Zeytinoglu, 1995) and has been widely used for perceptual speech enhancement (Cohen, 2001; Fu and Wan, 2003; Shao and Chang, 2006). To generalize this technique to bioacoustic signal enhancement, we propose to decompose a wavelet packet tree into the critical bands with respect to the species-specific Greenwood frequency warping curve.

Figure 2 shows an approximation of the Greenwood scale by critical-band WPD for three distinct species: ortolan bunting (*Emberiza hortulana*) downsampled to 20 kHz, rhesus monkey (*Macaca mulatta*) downsampled to 20 kHz, and the humpback whale (*Megaptera novaeanglia*) sampled at 4 kHz. The corresponding decomposition trees are illustrated in Fig. 3. The perceptual WPD splits the frequency range corresponding to different species data into critical

bands: ortolan bunting, 0 Hz–10 kHz, 36 critical bands; rhesus monkey, 0 Hz–10 kHz, 30 critical bands; humpback whale, 0 Hz–2 kHz, 31 critical bands. The bands are established automatically by optimally matching the subband center frequencies to the perceptual scale curve in the mean error sense. For the Greenwood scale calculation, the f_{\min} and f_{\max} used in Eqs. (10) and (11) are 400 and 7200 Hz for the ortolan bunting (Edward, 1943), 20 and 42 000 Hz for the rhesus monkey (Heffner, 2004), and 2 and 6000 Hz for the humpback whale (Helweg, 2000).

Given this perceptual decomposition structure, a MMSE estimator for performing thresholding can be derived in the wavelet domain (Cohen, 2001; Cohen and Berdugo, 2001). Using an additive time-domain model, the resulting wavelet domain model is

$$Y_{l,n}(k) = X_{l,n}(k) + D_{l,n}(k), \quad (12)$$

where $Y_{l,n} = \langle y, \varphi_{l,n,k} \rangle$, $X_{l,n}(k) = \langle x, \varphi_{l,n,k} \rangle$, $D_{l,n}(k) = \langle d, \varphi_{l,n,k} \rangle$, k is the index of the coefficients in each subband, l is the




FIG. 4. Spectrograms of ortolan bunting signals: Clean signal, -10 dB SNR noisy signals, and signals enhanced by bandpass filtering, spectral subtraction, Wiener filtering, EM log-MMSE filtering, and perceptual WPT filtering (the left column is for white noise and the right is for environment noise).

$$\hat{q}_{l,n}(k) = 1 - \begin{cases} \frac{\log[\xi_{l,n}(k)/\xi_{\min}]}{\log(\xi_{\max}/\xi_{\min})} & \text{if } \xi_{\min} \leq \xi_{l,n}(k) \leq \xi_{\max} \\ 0 & \text{if } \xi_{l,n}(k) \leq \xi_{\min} \\ 1 & \text{otherwise,} \end{cases} \quad (16)$$

where ξ_{\min} and ξ_{\max} are empirical constants, $\xi_{\min} = -10$ dB, and $\xi_{\max} = -5$ dB.

An estimate for the clean speech, which minimizes the mean-square error, results in

$$\hat{X}_{l,n}(k) = \frac{\lambda_{l,n}(k)p_{l,n}(k)}{\lambda_{l,n}(k) + \sigma_{l,n}^2(k)} Y_{l,n}(k), \quad (17)$$

where the signal variance is given by using the decision-directed method of Ephraim and Malah:

$$\hat{\lambda}_{l,n}(k) = \alpha |\hat{X}_{l,n}(k-1)| + (1 - \alpha) \max[|Y_{l,n}(k)| - \sigma_{l,n}(k), 0]. \quad (18)$$

IV. EXPERIMENTAL SETUP AND RESULTS

The proposed method and comparative baseline approaches were applied to ortolan bunting (*Emberiza hortu-*

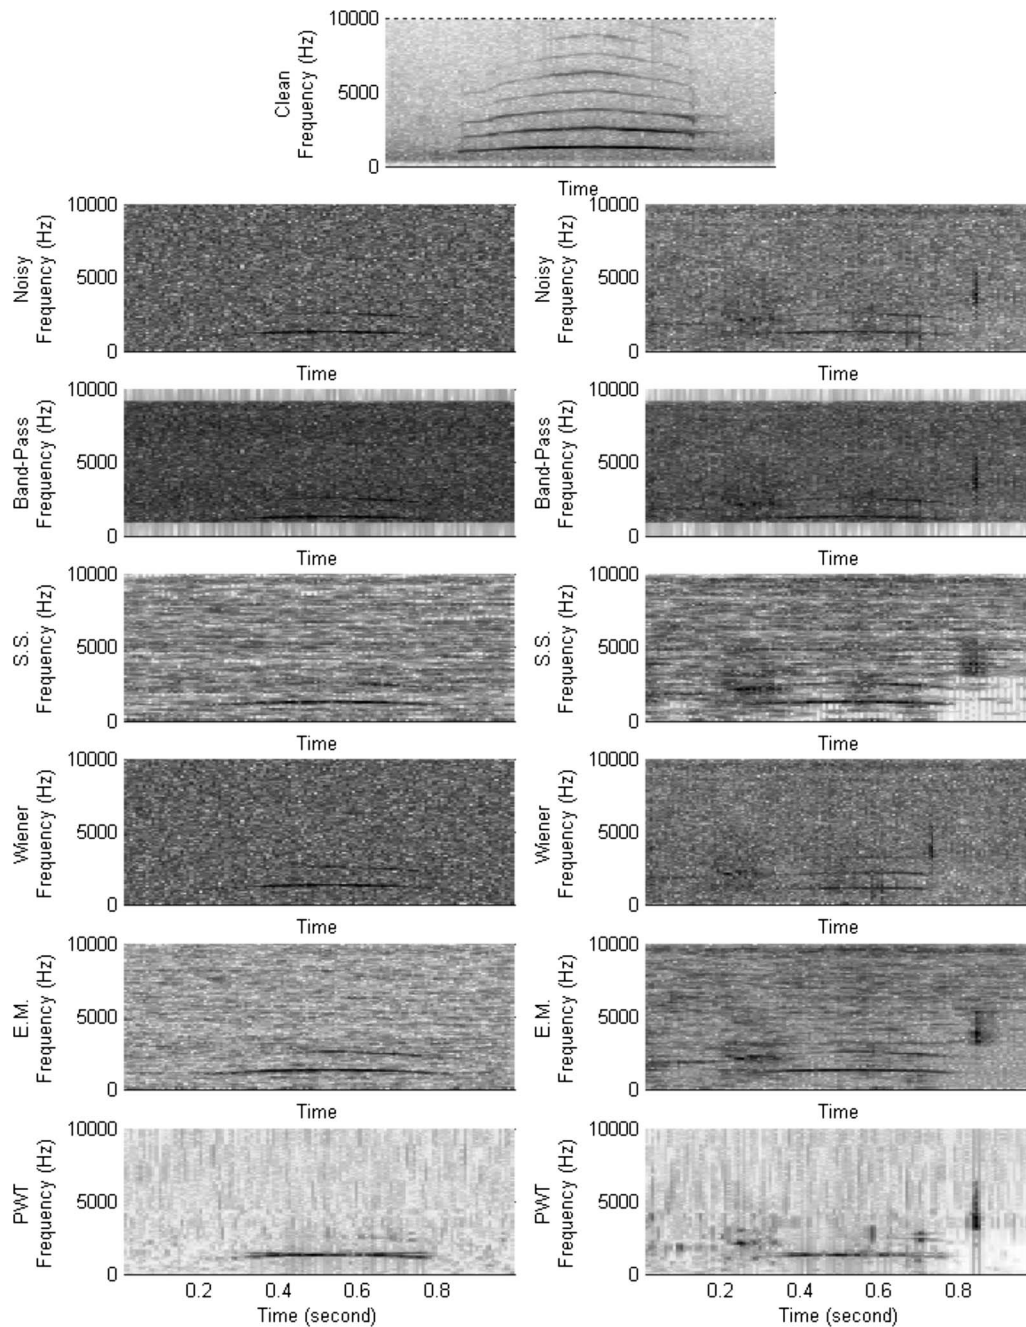


FIG. 5. Spectrograms of rhesus monkey signals: Clean signal, -10 dB SNR noisy signals, and signals enhanced by bandpass filtering, spectral subtraction, Wiener filtering, EM log-MMSE filtering, and perceptual PWT filtering (the left column is for white noise and the right is for environment noise).

lana), rhesus monkey (*Macaca mulatta*) and humpback whale (*Megaptera novaeanglia*). Norwegian ortolan bunting vocalization data were collected from County Hedmark, Norway in May of 2001 and 2002 (Osiejuk *et al.*, 2003). Rhesus data were recorded on the island of Cayo Santiago, Puerto Rico by Joseph Solitis and John D. Newman (Li *et al.*, 2007). Humpback whale data (Payne and McVay, 1971) was provided by MobySound (Mellinger and Clark, 2006), a database for research in automatic recognition of marine animal calls. These data were collected in March 1994 off the north coast of the island of Kauai, HI. Ten clean vocalizations from each species were segmented from the original recording data.

Both white noise and true environment noise were added to the clean data at SNR levels of -15 , -10 , -5 , 0 , $+5$, and $+10$ dB. The environment noise came from ambient noise regions of appropriate domain recordings for each species, spectrally flattened with a low order filter to preserve the basic noise characteristics while ensuring that the energy is spread through the entire frequency band. For the rhesus monkey vocalizations, background noise was taken from a Vervet monkey data set (Seyfarth and Cheney, 2004). For the ortolan bunting vocalizations, background noise came directly from the data set. For the humpback whale, marine noise was taken from a Beluga whale vocalization data set (Scheifele *et al.*, 2005), downsampled to 4000 Hz.




FIG. 6. Spectrograms of humpback whale: Clean signal, -10 dB SNR noisy signals, and signals enhanced by bandpass filtering, spectral subtraction, Wiener filtering, EM log-MMSE filtering, and perceptual WPT Filtering (the left column is for white noise and the right is for environment noise).

Based on visual examination of the clean data from Figs. 4–6, tight passbands are chosen around the vocalizations. Selected ranges are 2600–5600, 1000–10 000, and 200–2000 Hz for the ortolan bunting, rhesus monkey, and humpback whale data, respectively. For the spectral subtraction, Wiener filter, and EM filter approaches, the signal is divided into 32 ms windows with 75% overlap between frames. This frame length was chosen empirically, as it is sufficiently long for good spectral estimation in each frame but not so long as to affect temporal change in the signals, and adjustments to this value cause only minor changes to the overall enhancement results. Frequency analysis is done using a Hanning window and noise estimation is accomplished using the first three frames of the signal. For wavelet

analysis, the discrete Meyer wavelet is used as the mother wavelet, which was chosen to provide good separation of subbands due to their regularity property (Cohen, 2001). The decomposition was done as illustrated in Fig. 3. The forgetting factor α used in Eqs. (5) and (18) is set to 0.98 for the EM filter and 0.92 for the wavelet denoising.

SNR and SSNR are used as objective measurement criteria for all sets of experiments. SSNR is computed by calculating the SNR on a frame-by-frame basis over the signal and averaging these values. This permits the measure to assign equal weights to the loud and soft portions of the signal,

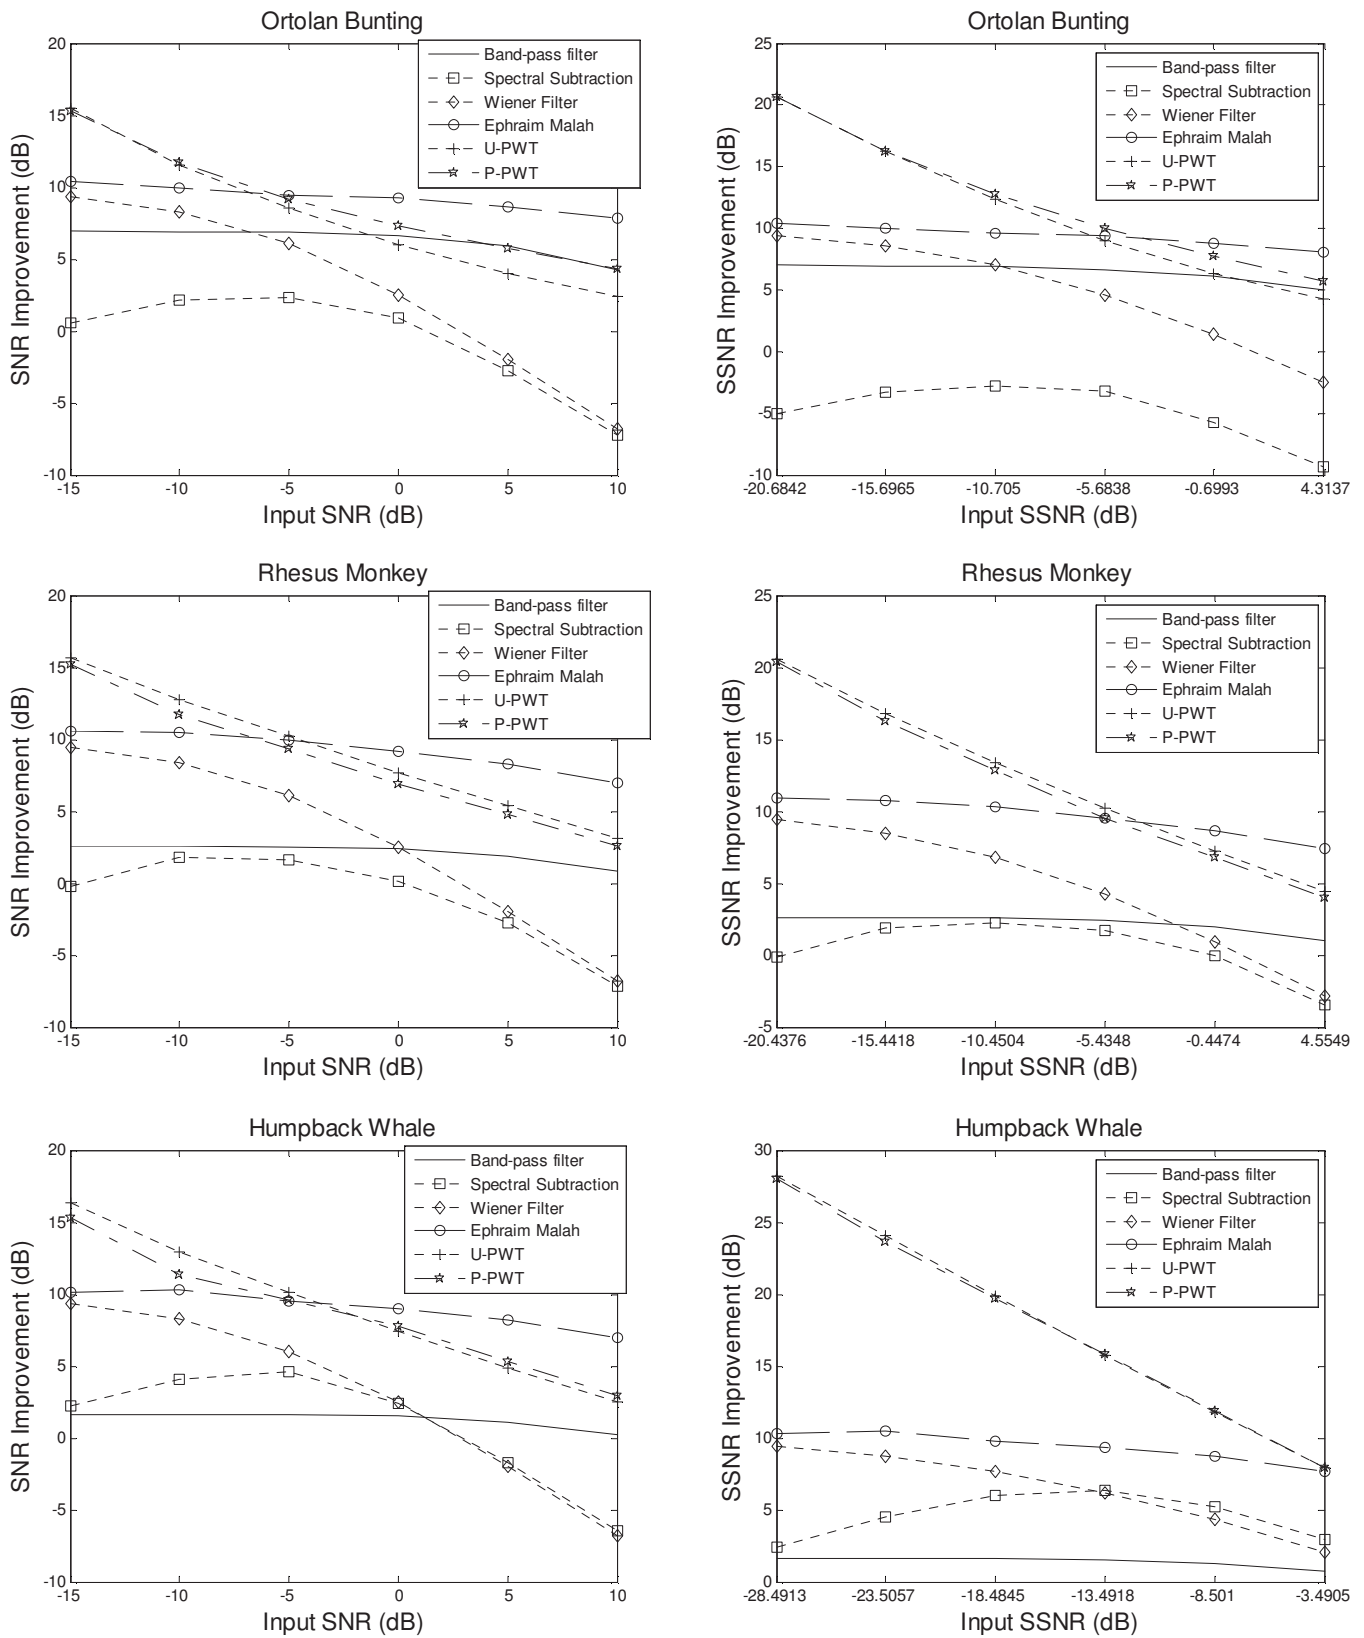


FIG. 7. SNR and SSNR results for white noise at -15 , -10 , -5 , 0 , $+5$, and $+10$ dB SNR levels.

which has been shown to have a higher correlation with perceived quality in human speech evaluation (Deller *et al.*, 2000). The formulas for SNR and SSNR are

$$\text{SNR} = 10 \log_{10} \frac{\sum_n x^2(n)}{\sum_n [x(n) - \hat{x}(n)]^2}, \quad (19)$$

$$\text{SSNR} = \frac{1}{M} \sum_{j=0}^{M-1} 10 \log_{10} \left[\frac{\sum_{n=Nj+1}^{N(j+1)} x^2(n)}{[\sum_{n=Nj+1}^{N(j+1)} (x(n) - \hat{x}(n))]^2} \right], \quad (20)$$

where M is the number of frames, each of length N , and $x(n)$ and $\hat{x}(n)$ are the original and enhanced signals, respectively.

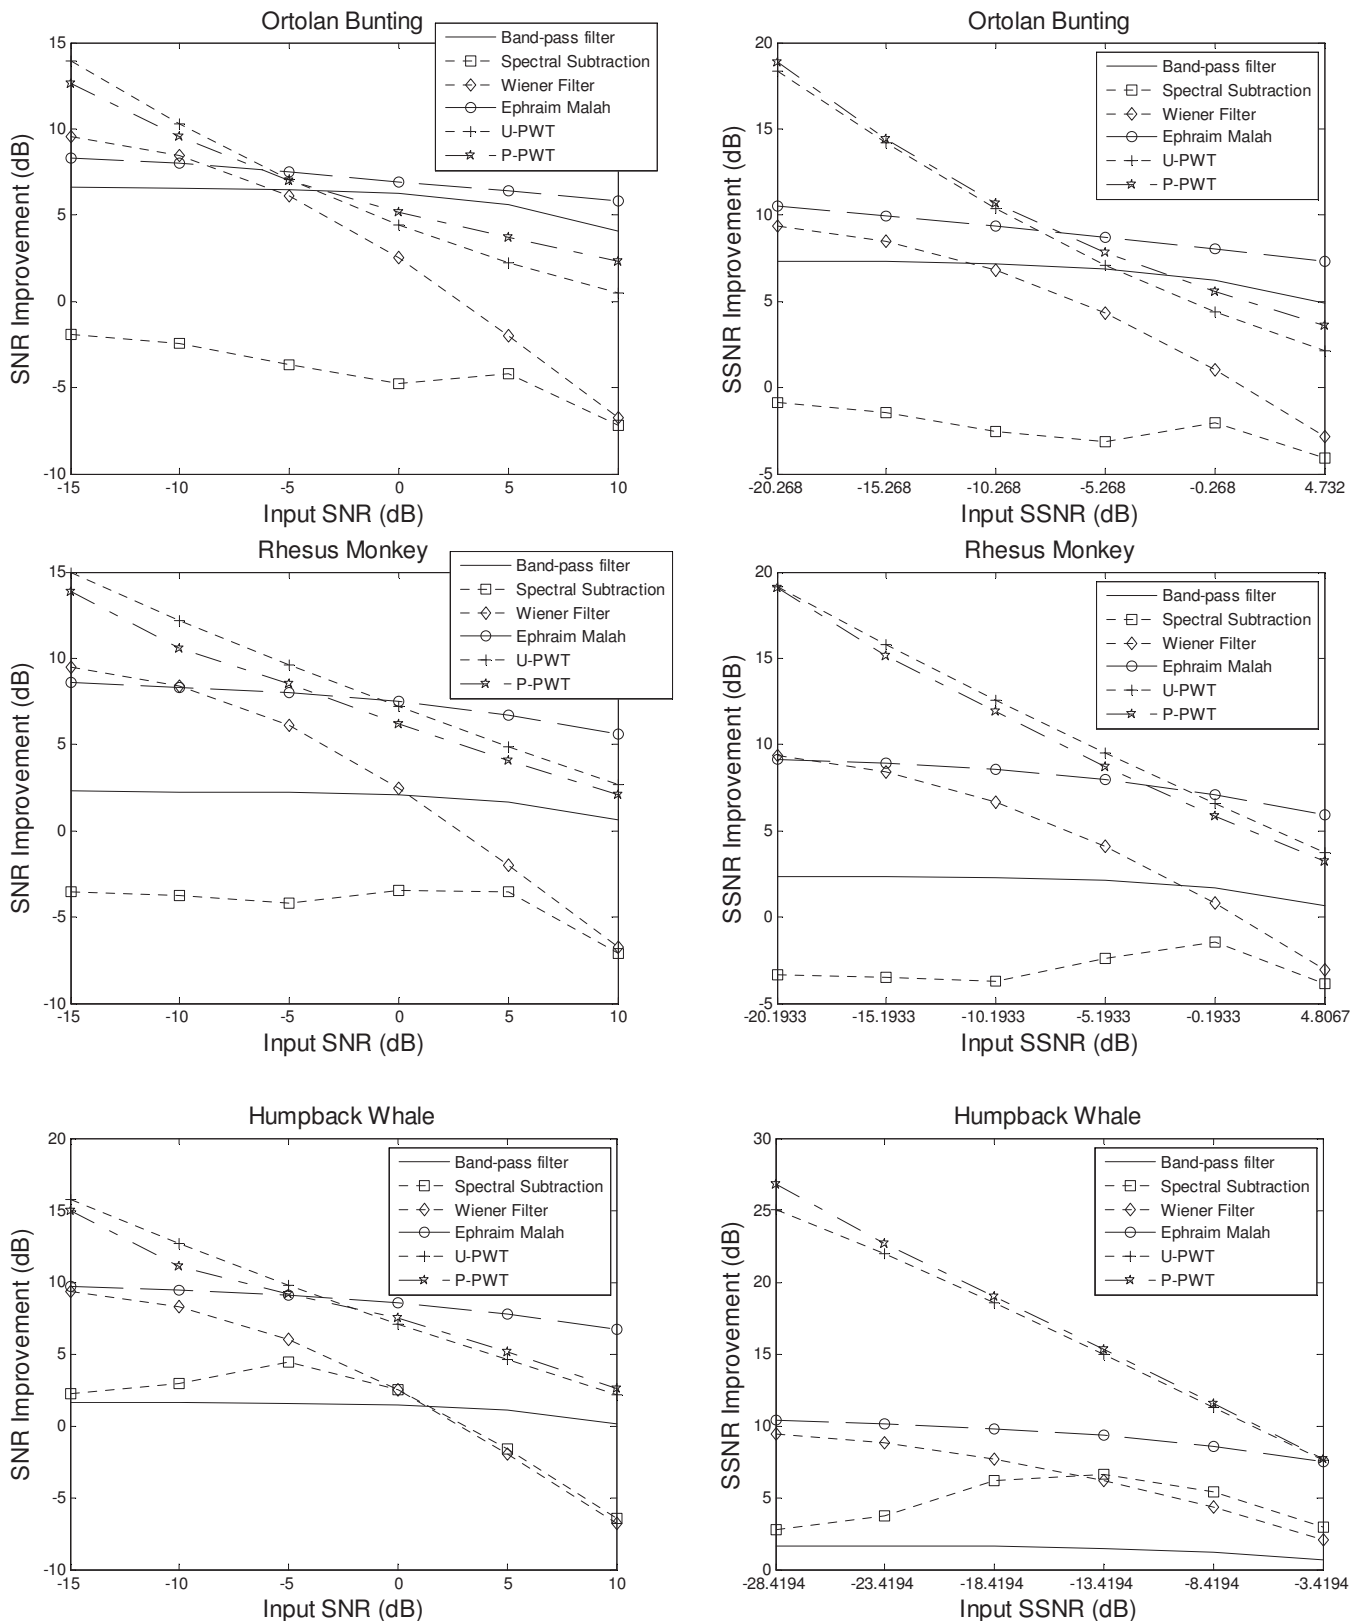


FIG. 8. SNR and SSNR results for environment noise at -15 , -10 , -5 , 0 , $+5$, and $+10$ dB SNR levels.

For visualization, spectrograms of the enhanced signals for the white noise and environment noise conditions at -10 dB SNR can be seen in Figs. 4–6.

SNR and SSNR results for the white noise and environment noise are shown in Figs. 7 and 8. The SNR and SSNR

values are given as amount of improvement over the original input noisy values. The methods shown in these figures include bandpass filtering, spectral subtraction, Wiener filtering, EM filtering, the proposed perceptual wavelet packet transform (P-WPT), as well as a uniform band wavelet

packet transform (U-WPT), which is identical to the proposed method except that it utilizes uniformly spaced frequency bands rather than the perceptual scaling.

From reviewing the spectrograms and the SNR and SSNR plots, several conclusions can be drawn. It is clear that the proposed perceptual wavelet denoising method and the EM filtering method have the best overall performance in both the white noise and the environment noise conditions. The proposed method shows better enhancement performance for the higher noise (lower original SNR) cases, in particular. By comparing the SNR improvement to the SSNR improvement in Figs. 7 and 8, it can be seen that the SSNR, which is generally considered to be a more perceptually meaningful metric, shows greater superiority for the proposed method over the other methods than does SNR. Wiener filtering and spectral subtraction have moderate enhancement performance overall, while bandpass filtering results are a little sporadic, giving generally moderate results with good results in a few specific environment cases. Specifically, as expected, bandpass filtering works relatively well in the ortolan case where the vocalization frequency range is narrow and has limited overlap with the environment noise spectrum. By comparing the P-WPT and U-WPT results, it can be seen that the use of the perceptual scale has little overall impact. In the white noise case, the SNR is slightly higher for the uniform scaling, and SSNR measures show little difference. For environmental noise, the SNR is again slightly higher for the uniform scaling, and SSNR is again similar, showing a slight benefit for the perceptual scaling in two of the three examples. Under the noisiest conditions, the two wavelet-based enhancement techniques significantly outperform all of the baseline methods.

One interesting thing to note is that each of the different enhancement methods has unique characteristics, as seen in the spectrograms of Figs. 4–6. Bandpass filtering has the expected look, keeping all noise in the target range and eliminating nearly everything out of band. Spectral subtraction shows some temporal streaking due to the fact that the noise spectrum being removed is fixed. Wiener filtering and EM filtering have similar looks, except that the EM provides better overall results. The proposed method has the best noise removal but can also be seen to possess an artifact (most noticeable in Fig. 5), seen as a faint reflection of the primary signal. This artifact, which is not audible and does not contain enough energy to significantly impact the SNR or SSNR metrics, illustrates some of the processing differences between a frequency domain approach such as the EM and a wavelet domain approach such as the proposed method. Because the mother wavelet used for analysis is somewhat broadband, each of the nodes in the decomposition trees shown in Fig. 3 contains more than a single frequency component. Thus the nodes that are given primary emphasis for reconstruction have energy at more than one frequency. However, since the nature of this wavelet representation is also more compact, coefficients not given primary emphasis can be more strongly thresholded, yielding less energy throughout the entire background frequency range, as can also be seen in the spectrograms. The selection of the mother wavelet also impacts the degree of this artifact. The overall

effect is that while the residual noise for the EM and perceptual wavelet approaches have similar total energy (with the perceptual wavelet having a little less in high noise situations), this residual noise in the EM approach is spread more evenly across the frequency range, while in the perceptual wavelet approach, it is more concentrated.

V. CONCLUSIONS

Enhancement techniques taken from the field of speech processing have been generalized and applied to noise reduction of bioacoustic vocalizations. Four baseline methods, including spectral subtraction, Wiener filtering, and EM filtering, as well as simple bandpass filtering, were compared to a new technique based on perceptual wavelet decomposition. Results indicate improved performance of the new method, particularly for the most noisy conditions. The new approach can be easily applied to any species, requiring only upper and lower frequency limits for the species to create the appropriate Greenwood function frequency warping curve.

ACKNOWLEDGMENTS

This material is based on work supported by National Science Foundation under Grant No. IIS-0326395. The authors also want to express their thanks to Joseph Solitis and John D. Newman for providing the rhesus monkey vocalizations, T. S. Osiejuk for providing the ortolan bunting vocalizations, and Mobysound for providing the humpback whale vocalizations.

- Álvarez, B. D., and García, C. F. (2004). "System architecture for pattern recognition in eco systems," ESA Special Publication No. 553, Madrid, Spain.
- Black, M., and Zeytinoglu, M. (1995). "Computationally efficient wavelet packet coding of wide-band stereo audio signals," in Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, Detroit, MI, Vol. 5, pp. 3075–3078.
- Boll, S. F. (1979). "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Trans. Acoust., Speech, Signal Process.* **ASSP-27**, 113–120.
- Clemins, P., and Johnson, M. T. (2006). "Generalized perceptual linear prediction (gPLP) features for animal vocalization analysis," *J. Acoust. Soc. Am.* **120**, 527–534.
- Clemins, P. J. (2005). "Automatic speaker identification and classification of animal vocalizations," Ph.D. thesis, Marquette University.
- Clemins, P. J., Trawicki, M. B., Adi, K., Tao, J., and Johnson, M. T. (2006). "Generalized perceptual feature for vocalization analysis across multiple species," in Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, Paris, France, Vol. 1, pp. 253–256.
- Cohen, I. (2001). "Enhancement of speech using bark-scaled wavelet packet decomposition," in Proceedings of Eurospeech, Aalborg, Denmark, pp. 1933–1936.
- Cohen, I., and Berdugo, B. (2001). "Speech enhancement for non-stationary noise environments," *Signal Process.* **81**, 2403–2418.
- Deller, J. R., Hansen, J. H. L., and Proakis, J. G. (2000). "Speech quality assessment," in *Discrete-Time Processing of Speech Signals* (IEEE, Piscataway, NJ), Chap. 9, pp. 584–587.
- Donoho, D. L. (1995). "De-noising by soft-thresholding," *IEEE Trans. Inf. Theory* **41**, 613–627.
- Edward, E. P. (1943). "Hearing ranges of four species of birds," *Auk* **60**, 239–241.
- Ephraim, Y., and Malah, D. (1984). "Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator," *IEEE Trans. Acoust., Speech, Signal Process.* **ASSP-32**, 1109–1121.
- Ephraim, Y., and Malah, D. (1985). "Speech enhancement using a minimum mean-square error log-spectral amplitude estimator," *IEEE Trans. Acoust., Speech, Signal Process.* **ASSP-33**, 443–445.

- Ephraim, Y., and Trees, H. L. V. (1995). "A signal subspace approach for speech enhancement," *IEEE Trans. Speech Audio Process.* **3**, 251–266.
- Fu, Q., and Wan, E. A. (2003). "Perceptual wavelet adaptive denoising of speech," in *Proceedings of EuroSpeech*, Geneva, Switzerland, pp. 577–580.
- Greenwood, D. D. (1961). "Critical bandwidth and the frequency coordinates of the basilar membrane," *J. Acoust. Soc. Am.* **33**, 1344–1356.
- Gur, B. M., and Niezrecki, C. (2007). "Autocorrelation based denoising of manatee vocalizations using the undecimated discrete wavelet transform," *J. Acoust. Soc. Am.* **122**, 188–199.
- Heffner, R. S. (2004). "Primate hearing from a mammalian perspective," *Anat. Rec.* **281A**, 1111–1122.
- Helweg, D. A. (2000). "An integrated approach to the creation of a humpback whale hearing model," Technical Report No. 1835, San Diego, CA.
- Jansen, M. (2001). *Noise Reduction by Wavelet Thresholding* (Springer, New York).
- Johnson, M. T., Yuan, X., and Ren, Y. (2007). "Speech signal enhancement through adaptive wavelet thresholding," *Speech Commun.* **49**, 123–133.
- Lepage, E. L. (2003). "The mammalian cochlear map is optimally warped," *J. Acoust. Soc. Am.* **114**, 896–906.
- Li, X., Tao, J., Johnson, M. T., Solitis, J., Savage, A., Leong, K. M., and Newman, J. D. (2007). "Stress and emotion classification using jitter and shimmer features," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Honolulu, HI, Vol. **IV**, pp. 1081–1084.
- Lim, J., and Oppenheim, A. V. (1978). "All-pole modeling of degraded speech," *IEEE Trans. Acoust., Speech, Signal Process.* **26**, 197–210.
- Liu, R. C., Miller, K. D., Merzenich, M. N., and Schreiner, C. E. (2003). "Acoustic variability and distinguishability among mouse ultrasound vocalizations," *J. Acoust. Soc. Am.* **114**, 3412–3422.
- Mellinger, D. K. (2002). *Ishmael 1.0 User's Guide*, Pacific Marine Environmental Laboratory, Seattle, WA.
- Mellinger, D. K., and Clark, C. W. (2006). "MobySound: A reference archive for studying automatic recognition of marine mammal sounds," *Appl. Acoust.* **67**, 1226–1242.
- Osiejuk, T. S., Ratynska, K., Cygan, J. P., and Svein, D. (2003). "Song structure and repertoire variation in ortolan bunting (*Emberiza hortulana* L.) from isolated Norwegian population," *Ann. Zool. Fenn.* **40**, 3–19.
- Payne, R. S., and McVay, S. (1971). "Songs of humpback whales," *Science* **173**, 585–597.
- Scheifele, P. M., Andrew, S., Cooper, R. A., and Darre, M. (2005). "Indication of a Lombard vocal response in the St. Lawrence River beluga," *J. Acoust. Soc. Am.* **117**, 1486–1492.
- Seyfarth, R. M., and Cheney, D. L. (2004). *TalkBank Ethology Data: Field Recordings of Vervet Monkey Calls* (Linguistic Data Consortium, Philadelphia).
- Shao, Y., and Chang, C.-H. (2006). "A generalized perceptual time-frequency subtraction method for speech enhancement," in *Proceedings of ISCAS 2006*, pp. 2537–2540.
- Yan, Z., Niezrecki, C., and Beusse, D. O. (2005). "Background noise cancellation for improved acoustic detection of manatee vocalizations," *J. Acoust. Soc. Am.* **117**, 3566–3573.
- Yan, Z., Niezrecki, C., Cattafesta, L.N., III, and Beusse, O. D. (2006). "Background noise cancellation of manatee vocalizations using an adaptive line enhancer," *J. Acoust. Soc. Am.* **120**, 145–152.