Proceedings of 20th International Congress on Acoustics, ICA 2010

23-27 August 2010, Sydney, Australia

A Deterministic Filterbank Compressive Sensing Model for Bat Biosonar

David A. Hague, John R. Buck, and Igal Bilik

University of Massachusetts Dartmouth, North Dartmouth, MA, United States

PACS: 43.60.-C, 43.80.LB, 43.60.BF, 43.60.PT, 43.60.RW

ABSTRACT

The Big Brown Bat (Eptesicus fuscus) uses Frequency Modulated (FM) echolocation calls to accurately estimate range and resolve closely spaced objects. Recent work by Fontaine and Peremans have shown that a sparse representation model for bat echolocation calls facilitates distinguishing objects spaced as closely as 2 micro-seconds in time-delay and was also robust to noise over a realistic range of signal to noise ratios (SNR). Fontaine and Peremans used the random FIR filter Compressive Sensing (CS) technique as their input method. Their study demonstrated that the undersampled data provided by the FIR filter output still contains sufficient information to accurately reconstruct and resolve sparse target signatures using L1 minimization techniques from CS. Their work raises the intriguing question as to whether under-sampled sensing approaches structured more like the bat's auditory system still contain the information necessary for the hyper-resolution observed in behavioral tests. This research investigates the ability to estimate sparse echo signatures using a downsampled filterbank for the sensing basis that is closer to a bat auditory system than randomized FIR filters. The returning echoes are sensed using a discrete-time constant-bandwidth filter bank followed by downsampling that loosely resembles the filtering and smoothing of the bat's cochlea. L1 minimization then reconstructs the sparse target return from this under-sampled signal. Initial simulations demonstrate that this filterbank CS model reconstructs sparse sonar targets with a high degree of accuracy while substantially undersampling the filter outputs. In addition, the overdecimated filterbank CS approach has better target resolution than the Matched Filter for SNR values ranging from 5-45 dB and has better detection performance than the Inverse Filter method. This is all accomplished while undersampling the return echo signal by as much as a factor of six. The deterministic sensing basis has the distinct advantage over the random sensing basis in the respect that the circulant structure of the filterbank sensing matrix can easily be implemented in electric circuits.

I. INTRODUCTION

The Big Brown Bat (Eptesicus fuscus) operates in complex acoustic environments. Using wideband Frequency Modulated (FM) echolocation calls in the 20-100 kHz band. Bats can detect target reflections or glints as closely spaced as 2 μs in the midst of noise [6]. The task that the bat's auditory system performs very well is the same challenge posed to human made active sonar systems. Active sonar transmits sound and attempts to detect and estimate the range of the ensonified objects. In addition, active sonar systems must be able to resolve closely spaced targets for complicated multi-target environments.

The Matched Filter (MF) or Auto-Correlation Filter optimizes the detection and delay estimation for a single known signal in additive white noise. However, it does not optimize the resolution of closely spaced targets. The resolution of the MF is determined by the mainlobe of the Auto-Correlation function of the transmit waveform. For bat FM transmit waveforms, the mainlobe width is approximately 10-12 μs.

Target resolution can be optimized by utilizing a linear time invariant filter that minimizes the width of the peak of the target return that results from the convolution of the transmit waveform with the impulse response of the filter. This can be obtained by using an inverse filter (IF). The frequency response of the IF is generated by taking the reciprocal of the frequency spectrum of the transmit waveform. Convolving the return signal with the impulse response of the IF results in scaled and delayed impulse functions that correspond to the amplitude and time delay of the target returns. This results in a significant improvement in resolution of closely spaced targets for high SNR. However, as the SNR decreases, the IF amplifies the noise out of the frequency band of the transmit waveform and the detection performance is significantly degraded. While tapering the out of frequency band response of the IF will improve the detection capabilities, its detection performance still will not rival that of the matched filter and the tapering degrades resolution performance.

Recent work by Fontaine and Peremans [3] has proposed a sparse representation of bat biosonar signals. This method builds from the assumption that the echo return signal can be sparsely represented as a convolution of the transmit waveform and a small number of impulse functions that model the target image return. The return signal is then reconstructed via ℓ_1 minimization methods. Simulations showed that this sparse representation of bat echolocation calls can resolve targets spaced as closely as 2 μs in range separation with performance similar to the inverse filter while also showing a significant improvement in detection capability for a range of realistic SNR's. In addition, a compressive sensing technique was ap-

plied to the sparse representation approach and was able to reconstruct the biosonar signals using fewer samples than the Nyquist rate requires while largely preserving the resolution and detection performance. Their sampling methodology was based off of work done by Baraniuk and Steeghs [11] and was shown to yield impressive target return reconstruction results. While no claim was made in regard to the biological plausibility of this method, [3] raises the intriguing question of whether a sensing methodology similar to the mammalian auditory system contains the necessary information to achieve the resolution and detection performance reported in observational studies.

This paper builds from Fontaine and Peremans' work and exploits the sparsity of Biosonar signals using a sensing methodology that more closely resembles the physiology of the mammalian auditory system. The sensing method comprises of a filterbank of finite impulse response (FIR) filters followed by decimation. The filterbank spans the frequency band of the bat's echolocation call. The filterbank decimation factor can be adjusted to sample at the Nyquist rate (maximally decimated) or to sample below the Nyquist rate (over-decimated). Like Fontaine and Peremans, we make no claims to the biological plausibility of the compressed sensing algorithm. The aim of this research is to demonstrate that the undersampled outputs of a deterministic filterbank still contain sufficient information to reconstruct sparse target impulse responses in a manner analogous to that observed in behavioural experiments with bats. The simulation results show that this sensing method is able to resolve closely spaced targets that the MF cannot and can detect these closely spaced targets with a higher success rate than the IF while also undersampling the target return signal.

II. COMPRESSIVE SENSING

CS theory asserts that a signal that is sparsely represented in some basis Ψ can be acquired from fewer measurements than generally required, or undersampling. A common example is that signals containing only a few frequency tones (sparse in frequency spectrum) can often be accurately acquired while sampling below the Nyquist rate. Consider an N-length bandlimited discrete-time signal x[n] with no additive noise sampled at a frequency f_s such that the Nyquist criterion is met (i.e. that f_s is at least twice the highest frequency of x[n]). Further, assume that x[n] is sparse in some basis. In vector form, x can be represented in the sparse basis in vector form as

$$\underline{x} = \Psi \cdot \underline{a} \tag{1}$$

where Ψ is of dimension N x N. The N x 1 vector $\underline{a}$ contains S non-zero coefficients where S is much less than the length of the signal N. In CS terminology, we say that the sparse vector $\underline{a}$ is S-sparse or x is S-sparse in Ψ. This particular case assumes a noiseless signal. In reality almost every signal contains some additive noise and the signal is no longer strictly S-sparse. However, the S columns of Ψ represent the vast majority of the signal energy. In the CS framework this signal is then said to be compressible.

$$\underline{x} = \Psi \cdot \underline{a} + \underline{n} \tag{2}$$

The sparse or compressible signal is then sampled by a M x N sensing matrix Φ where M < N to obtain an undersampled signal y[n]. In vector form, y can be written as

$$\underline{y} = \Phi \cdot \underline{x}. \tag{3}$$

Substituting (1) in to (3) yields

$$\underline{y} = \Phi \cdot \Psi \cdot \underline{a} = \Theta \cdot \underline{a} \tag{4}$$

Where $\Theta = \Phi \cdot \Psi$. Finding the signal's sparse coefficient vector $\underline{a}$ becomes a convex optimization problem. Candes, Romberg, and Tao [4] show that ℓ_1 norm minimization exactly recovers S-sparse signals and closely approximates compressible signals with high probability. The ℓ_1 norm minimization problem can be solved by the linear program known as Basis Pursuit.

A sufficient condition that allows for sparse signal reconstruction is for the sensing matrix to obey the Restricted Isometry Property (RIP). The RIP states [5]

For each integer S=1,2,… (sparsity), define an isometry constant δ_s of the matrix Θ such that:

$$\left(1-\delta_s\right)\left\|\underline{a}\right\|_2^2 \le \left\|\Theta \cdot \underline{a}\right\|_2^2 \le \left(1+\delta_s\right)\left\|\underline{a}\right\|_2^2 \tag{5}$$

holds for all S-sparse vectors.

This theorem states that if Θ satisfies the RIP, then any subset of S columns of Θ are approximately orthogonal [5]. Consequently, no S-sparse vector has a significant amount of energy in the null space of Θ, and it should be possible to at least approximately reconstruct $\underline{a}$ from $\underline{y}$. A related measure known as coherence requires that "the rows of Φ cannot sparsely represent the columns of Ψ" [9]. The coherence of the sensing matrix is computed as

$$\mu\left(\Phi, \Psi\right) = \sqrt{N} \max_{1 \le k, j \le N} \left|\left\langle \phi_k, \psi_j \right\rangle\right| \tag{6}$$

where ψ_j are the columns of Ψ and φ_k are the rows of Φ.

The coherence measures the largest inner product between the rows of Φ and the columns of Ψ. The coherence μ can take on values between 1 and $\sqrt{N}$. If Φ and Ψ contain common values, the coherence tends towards $\sqrt{N}$. The columns of the 2 matrices are then said to be mutually coherent. When Φ and Ψ do not share common values, the coherence tends towards 1 and they are said to be mutually incoherent. A sensing matrix with low coherence will obey the RIP with high probability. Candes and Wakin [5] show that sensing matrices with i.i.d Bernoulli or Gaussian random entries with variance 1/N maintain a low coherence and thus obeys the RIP with high probability. Most of the CS literature revolves around utilizing random sensing matrices due to their ability to guarantee reconstruction of sparse or compressible signals. However, work by Devore [9] has shown that sensing matrices with deterministic structures can also obey the RIP.

III. SPARSE REPRESENTATION AND SENSING METHOD

The sparse representation exploited in this paper assumes that the echolocation return signal can be expressed as a linear superposition of the echoes. The echoes are modelled as impulse functions with amplitude corresponding to the strength of the target return. It is also assumed that the phase of the transmitted signal is not altered when reflected. The return target image can be modelled as a sum of scaled and time-delayed impulse functions and the return signal x[n] in turn can be written as

$$x\left(t\right) = \sum_{i=1}^{S} a_i s\left[t - 2r_i / c\right] + w\left(t\right) \tag{7}$$

$$x\left[n\right] = x\left(n \cdot T_s\right) \tag{8}$$

where a_i represents the amplitude of the echo reflected by target i with delay $2r_i/c$ where r_i is the range of target i,

$s(t-2r_i/c)$ is the transmit waveform with time delay $2r_i/c$, S is the number of reflections, w(t) is additive white Gaussian noise, and c is the velocity of sound in air (~350 m/s). The amplitude here is constrained to being positive given that the reflection coefficient is positive. This is a realistic assumption for bat biosonar. Bats are attempting to locate insects and other solid objects. The acoustic energy is transferred from a medium of low acoustic impedance (air) to a medium of high acoustic impedance (solid object). This results in a positive reflection coefficient.

Given that the target returns are modelled as impulse functions, the discrete-time return signal x[n] can also be expressed as the convolution of the transmit waveform with the target impulse response. Because the target return image in (7) is a superposition of impulse functions, it can be considered to be sparse. The return signal x[n] can then be written as (1) where the N x N sparse basis Ψ is constructed as

$$\Psi = \begin{bmatrix} s[0T_s] & 0 & \cdots & 0 \\ s[1T_s] & s[0T_s] & \cdots & 0 \\ s[2T_s] & s[1T_s] & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ s[(P-1)T_s] & s[(P-2)T_s] & \cdots & 0 \\ 0 & s[(P-1)T_s] & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & s[0T_s] \end{bmatrix} \tag{9}$$

Where T_s is the discrete-time sampling period, s[n] is the discrete-time representation of the bat echolocation signal, and P is the length of s[n]. The sparse basis is essentially a time shifted dictionary of the transmit waveform.

The work by Fontaine and Peremans [3] makes use the decimating random Finite Impulse Response (FIR) filter sensing matrix described by Baraniuk and Steeghs [11]. This work differs from Fontaine and Peremans approach on two fronts. First, the sensing matrix is composed of a bank of constant bandwidth FIR filters. Second, the filter taps are no longer populated by pseudo random values, but with deterministic coefficients used in traditional FIR filter design. The filterbank comprises of L constant bandwidth FIR filters of impulse response length N_h that cover the frequency band of the bat echolocation call (20-100 kHz) and can be decimated by an integer factor D. The sensing matrix carries out the convolution of the return waveform with each of the filters as seen in Figure 1.

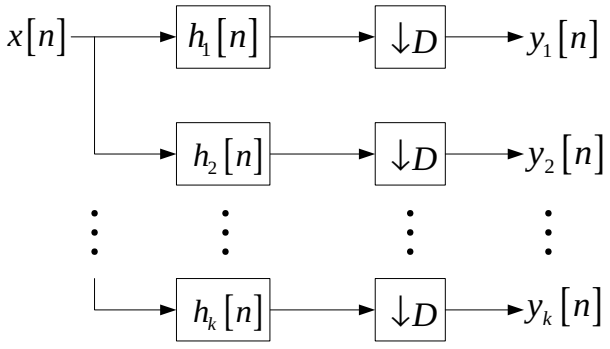

Figure 1. Filterbank system block diagram.

Figure 1 implements (3) if we choose

$$y = \left[Y[0]\ Y[1] \ldots Y[M/L] \right]^T \tag{10}$$

$$Y[n] = \left[y_1[n]\ y_2[n] \cdots y_L[n] \right]^T \tag{11}$$

$$\underline{x} = \left[x[0]\ x[1] \cdots x[N-1] \right]^T \tag{12}$$

$$\Phi = \begin{bmatrix} H_0 & O & \cdots & \cdots & \cdots & \cdots & O \\ O & & H & & O & O & \vdots \\ \vdots & \ddots & & H & & \ddots & \vdots \\ \vdots & & \ddots & & \ddots & O & O \\ O & \cdots & \cdots & O & & H & \end{bmatrix} \tag{13}$$

where O is an L x D matrix of zeros and H_0 is the L x 1 vector given by

$$H_0 = \left[h_1[0] \cdots h_L[0] \right]^T \tag{14}$$

and

$$[H]_{ij} = h_i[N_h - j] \tag{15}$$

for i=1,…,L and j=1,…,N_h. The sensing matrix is of size M x N with M calculated by

$$M = \frac{(N_h + N - 1)L}{D} \tag{16}$$

where N_h is the length of the filters' impulse response, L is the number of filters, and the integer D is the decimation factor. From here the under-sampling factor α is defined as the ratio of the number of filters to the decimation factor:

$$\alpha = L/D. \tag{17}$$

For the case of D = L, the under-sampling factor is 1 and thus the size of Φ and the sampled signal y is (N_h + P – 1) which is the length of the return signal after it is convolved the filterbank impulse response. When D > L, the under-sampling factor is less than 1 and the filterbank has over-decimated the filtered version of the return signal x[n]. Over-decimation results in aliasing of the discrete-time signal which destroys information contained in x[n].

Applying this to the target image problem results in the inverse problem of estimating the target return signal x[n] using the sparse basis Ψ and the measurements y. Given the assumption that the target return signal is sparse in Ψ, this estimation problem can be solved via convex optimization. The ℓ_1 minimization problem can be stated as

$$\min \left\| \underline{a}' \right\|_1 \text{ subject to } \underline{y} = \Theta \cdot \underline{a} \tag{18}$$

where $\underline{a}'$ is the estimate of the sparse vector and $\underline{r}$ is the return signal vector. For a noisy target return signal, the ℓ_1 minimization problem can be reformulated as follows:

$$\min \left\| \underline{a}' \right\|_1 \text{ subject to } \left\| \underline{y} - \Theta \cdot \underline{a} \right\|_2 \le \varepsilon \tag{19}$$

where ε is an adjustable tolerance parameter. This equation implies that the sparse solution obtained is a close fit to the noisy observations.

IV. SIMULATION RESULTS

As described earlier, the columns of Ψ contain time shifted version of the original transmit waveform. The echolocation call is modelled as a downward sweeping hyperbolic FM sinusoidal signal with a fundamental component sweeping from 50 kHz to 20 kHz and a second harmonic sweeping from 100 kHz to 40 kHz in a 1 ms duration. The transmit waveform s[n] is a loosely modelled version of the types of calls emitted by FM bats, specifically the big brown bat. The instantaneous frequency and phase of this signal is represented as

$$f_i(t) = a \cdot i \Big/ (t+b) \tag{20}$$

$$\theta_i(t) = i \cdot 2\pi a \log(t+b) \tag{21}$$

Where i=1, 2 and represents the 1st (fundamental) and 2nd harmonic of the emitted call. A plot of the transmit waveform in both the time and time-frequency domain is given in Figure 2. The amount of noise present in the return signal is measured by the SNR given by

$$SNR(dB) = 10\log_{10}\left(\frac{2E}{N_0}\right) \tag{22}$$

Where E is the energy of the return signal and N_0 is the noise energy both in units of Joules. In bat echolocation scenarios, a SNR of 45 dB is considered high, 25 dB is moderate, and 5 dB is low.

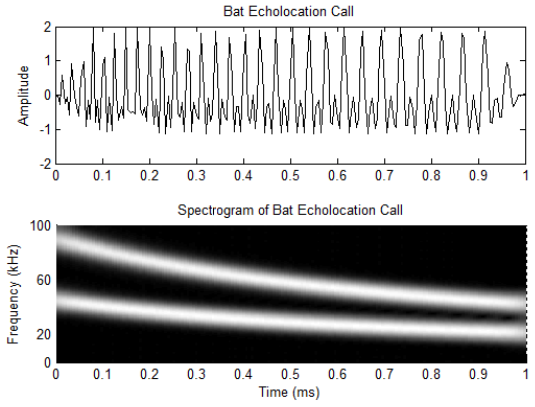

Figure 2. Plot of the bat echolocation call in the time and time-frequency domain.

In addition to the sparse representation method, the simulations also implemented both the MF and IF for comparison. As stated earlier, the MF is the optimum detection receiver but has a target resolution limited to the main lobe width of the Auto-Correlation function of the transmit waveform. The IF has significantly better target resolution capabilities at the cost of reduced detection performance.

The simulations implemented the detection system seen in Figure 3. The system is comprised of 4 main components; the receiver, a magnitude squared block, a peak selector, and a threshold comparator. The return signal x[n] is passed through either a MF, IF, or the sparse representation receiver. The square of the magnitude of the raw target image t[n] is then taken. This is analogous to the full-wave rectifier which is used in many human made active sonar receivers. The squared magnitude signal u[n] is next passed through a peak selection algorithm. This algorithm creates v[n] which now maintains only the peaks of u[n]. This is especially useful for the MF as the mainlobe of the Auto-Correlation function spans several discrete-time data points and would result in the detection system indentifying several targets at a time delay where there is only one true target. Lastly, a range of detection thresholds γ can be applied to test whether a target is present.

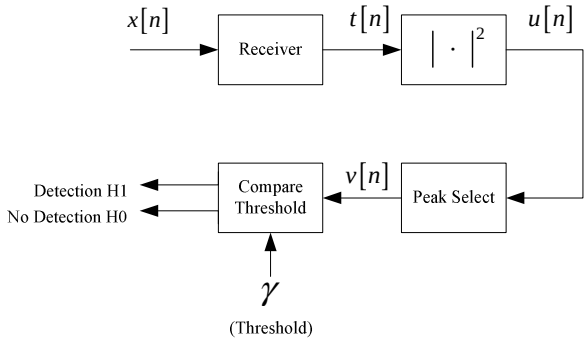

Figure 3. Block diagram of the detection system used for all detection simulations (after Figure 5 of [8]).

It is useful to analyse the signal v[n] that results from the peak detection algorithm. This gives a visual perspective of how well each receiver is performing. Figure 4 illustrates the performance of all 3 receivers with the mean (over 10 trials) of the target return signal v[n] with SNR of 25 dB. The original target response or sparse vector is comprised of 10 echoes with time delays of 200, 208, 400, 412, 600, 616, 800, 824, 1000, and 1032 µs. This results in pairs of targets with spacing 8, 12, 16, 24, and 32 µs. It is important to note here that the sparse representation implemented a sensing matrix that did not undersample the return signal (i.e. the number of filters L equals the decimation factor D).

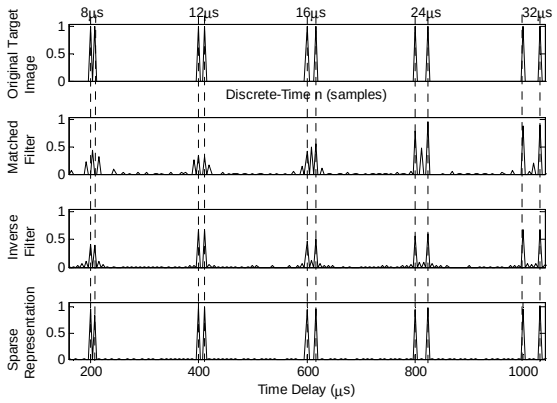

Figure 4. Original target impulse response from a return signal with 25 dB SNR and the returns v[n] of the MF, IF, and the sparse representation.

As can be seen by comparing the mean target images, both the IF and sparse representation methods successfully reconstruct all targets at their proper location. The MF is unable to resolve the targets spaced 8 µs apart in time delay. While the MF still resolves the targets spaced 12 µs and 16 µs apart, the amplitude of the targets with 12 µs spacing are well below half the original target amplitudes and the targets with 16 µs have sidelobes from the Auto-Correlation function that add to create a false detection between the two original targets with an amplitude higher than the original 2 targets. It is important to point out that while the IF successfully resolves all targets, there are more spurious spikes in the return than in the sparse representa-

tion. In addition, the sparse representation reconstructed target image has more accurate amplitude estimates than the IF. From this it can be concluded that the sparse representation appears to have better detection capabilities than the IF while maintaining the same target resolution. To assess the performance of these detectors more systematically, a range of detection thresholds can be applied to v[n] to determine when the true targets will be detected and when the spurious returns will be incorrectly determined as a true target.

Receiver Operating Characteristic (ROC) curves plot the probability of detection versus the probability of false alarms [12] quantifying the performance of a sonar receiver. Figure 5 below illustrates the ROC curves for the 3 receivers for SNR values of 45, 25, 15, and 5 dB respectively for 500 trials of the return target impulse response shown in the upper panel of Figure 4. As with the previous simulation, the sparse representation did not undersample the return signal before performing reconstruction. As expected from the second panel in Figure 4, the MF rarely detects the two most closely space peaks at around 200 µs, giving a maximum P_D of about 0.8. As the thresholds increases, some true targets are rejected by the MF system. Thus the MF ROC curves exhibit a stair case curve at SNR's 45, 25, and 15 dB that occur for P_D values of approximately 0.8, 0.7, 0.6, and 0.5.

The IF and sparse representation have ideal sharp ROC curves over a wide range of thresholds for 25 and 45 dB SNR. The probability of false alarm is always zero and the probability of detection is always unity. This results in an ROC curve which is on top of the left and top axis of the plot. This implies that the strongest false detection amplitude is still lower than the weakest true target return and thus there are a range of thresholds γ such that all true targets can be detected while acquiring 0 false detections. At 15 dB, the IF no longer has an ideal curve but it still retains solid detection performance with a P_D of 0.95 for a P_{FA} of 10^{-4}. The sparse representation curve at 15 dB is significantly sharper than the IF's ROC curve for a P_D of 0.95 for a P_{FA} of 10^{-6}. For 5 dB, the inverse filter's detection performance is significantly diminished and can only muster at best a $P_D \approx 0.85$ for a $P_{FA} \geq 0.3$. The ROC curve for the sparse representation at 5 dB SNR maintains a P_D of 1.0 for a P_{FA} of approximately 1.2×10^{-3}. It is clear from these ROC curves that not only does the sparse representation resolve the closely spaced targets where the MF cannot, but also has superior detection capabilities compared to the MF and IF.

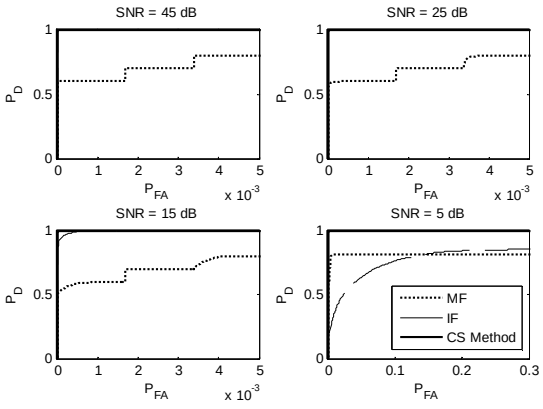

Figure 5. ROC curves of all 3 receivers for SNR values of 45, 25, 15, and 5 dB. Note the change in the P_{FA} axis in the lower right plot.

Thus far all experimental results created by the sparse representation were implemented without undersampling the return signal. To illustrate the performance of this method when undersampling the return signal, Figures 6-9 show ROC curves for 500 randomly generated sparse vectors $\underline{a}$ with a SNR of 25 dB with sparsity values of 10, 15, 20, and 25 for the MF, IF, and sparse representation with undersampling ratios α of 1/1 (not undersampled), ½, ¼, and 1/6. The targets in each sparse vector can be spaced as closely as 8 µs or as far away as 100 µs. The amplitude for each nonzero element of the sparse vector varies randomly between 0.5 and 1.0. Therefore, knowing that the MF can't resolve these closely spaced targets, it is expected that the IF and sparse representation with α=1 will outperform the matched filter.

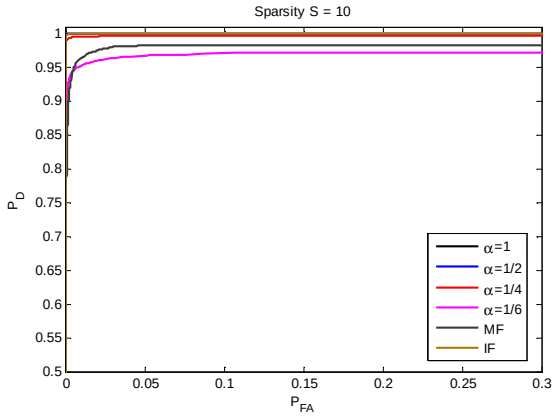

Figure 6. ROC curves for the MF, IF, and the sparse representation for undersampling ratios α of 1, ½, ¼, and 1/6 for 500 trials with randomly generated sparse vectors of sparsity S=10.

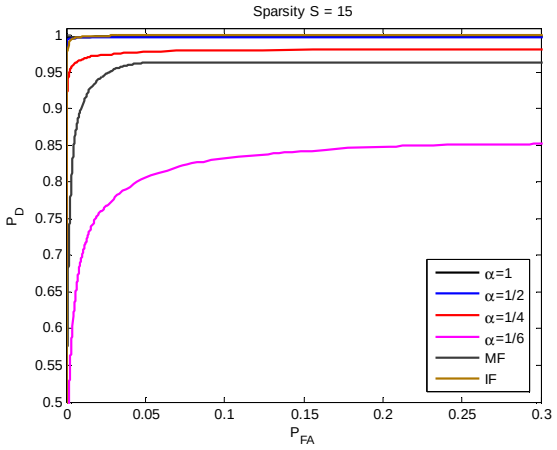

Figure 7. ROC curves for the MF, IF, and the sparse representation for undersampling ratios α of 1, ½, ¼, and 1/6 for 500 trials with randomly generated sparse vectors of sparsity S=15. Note the change in both the P_{FA} axis.

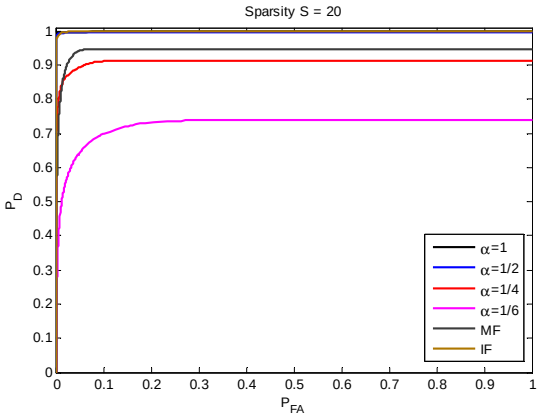

Figure 8. ROC curves for the MF, IF, and the sparse representation for undersampling ratios α of 1, ½, ¼, and 1/6 for 500 trials with randomly generated sparse vectors of sparsity S=20. Note the change in both the P_{FA} and P_D axis.

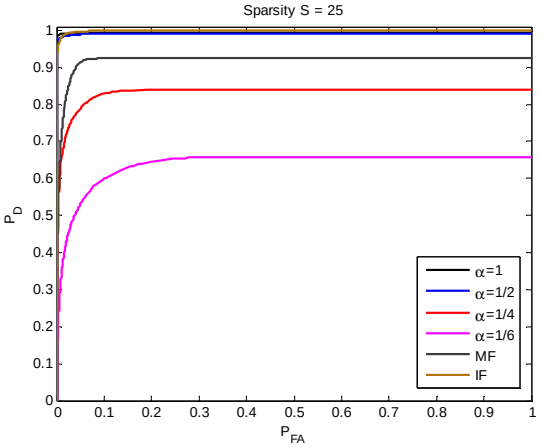

Figure 9. ROC curves for the MF, IF, and the sparse representation for undersampling ratios α of 1, ½, ¼, and 1/6 for 500 trials with randomly generated sparse vectors of sparsity S=25.

Based on Figures 6 and 7, it is clear that for sparse vectors with sparsity up to 15, the return signal x[n] can be undersampled by as much as ¼ and maintain better detection performance than the MF. The IF and the sparse representation for undersampling ratios of 1/1 and ½ have a nearly ideal ROC curve for all sparsity values. Using an undersampling value of 1/6 yields a detection performance that is inferior to the MF even for S=10, but it is still maintains acceptable detection performance. Undersampling at ¼ and lower does not yield acceptable detection probabilities for sparsity 20 and higher.

V. DISCUSSION

The results of the previous section clearly show that the sparse representation method has target resolution capabilities far beyond that of the MF as well as maintaining superior detection performance to that of the IF or MF for closely spaced target returns. There are 2 main reasons for this. First, the ℓ_1 minimization allows for accurate reconstruction of the sparse target return impulse response and thus provides excellent target return resolution. Second, the detection performance is greatly aided by the nature of the sensing methodology. Utilizing a filterbank of FIR filters spanning only the frequency band of the transmit waveform removes any noise located out of the frequency band of the transmit waveform. This results in a further improvement of the SNR and thus should provide an improvement in P_D. The IF utilizes taper weights that attenuate the frequency response out of the frequency band of the transmit waveform and provides an improvement in detection performance at a cost of reduced resolution. This results in less accurate target impulse response amplitudes and more spurious sidelobes located around a target return.

Figures 6-9 illustrate the sparse representation's ability to achieve high P_D while undersampling the return signal. This demonstrates a deterministic sensing basis that can reconstruct a sparsely represented signal at undersampling rates as low as 1/6 for a sparsity values up to S=10. It is interesting to note that this sensing matrix has high coherence values on the order of $\sqrt{N}$ yet still allows for sparse reconstruction. This surprising result has been noted in [7] where intense numerical simulation verified that the decimating random filter allowed for successful reconstruction of sparse signals in spite of high coherence. Work by DeVore [9] shows that sensing matrices of a deterministic structure can obey the RIP with high probability. A thorough mathematical proof was introduced for sensing matrices of a circulant nature. A circulant matrix is determined by its first K columns and has the property

$$\Phi_{i+1,j+K} = \Phi_{i,j} \tag{23}$$

where i and j are respectively the rows and columns of the matrix Φ. The sensing matrix comprised of the filterbank also is circulant in nature and has the property

$$\Phi_{i+L,j+D} = \Phi_{i,j} \tag{24}$$

where L and D are the number of filters and the decimation factor respectively. It appears that the filterbank sensing matrix possesses similar structure to what was analysed in [9] and may be the reason why the filterbank sensing matrix allows for undersampling and reconstructing a sparsely represented signal.

Of the 3 receivers analysed, the MF was the most straightforward to implement. The impulse response of the MF is a time-reversed version of the transmit waveform. It is often implemented in the frequency domain to exploit the computational efficiency of the FFT especially when the return signal x[n] is already in the frequency domain. When implementing an IF however, several issues need to be addressed. The discrete-time version of the transmit waveform has zeros inside and outside the unit circle. A discrete-time signal with zeros outside the unit circle cannot have a stable and causal inverse system [2, Sec. 5.6]. This implies that a stable impulse response for the IF is non-causal, infinitely long, and exponentially-decaying. The exponential decay rate is determined by the location of the zeros of the transmit waveform. The farther the zeros are from the unit circle the faster the decay.

The IF is typically implemented in the frequency domain so that the efficiency of the FFT can be exploited. Frequency domain sampling will result in periodic aliasing in the time domain of the receiver impulse response. Therefore the IF frequency response must be sampled often enough so that the periodic aliasing of the IF impulse response is negligible. This is done by selecting an FFT size N_{FFT} such that the IF impulse response decays to nearly zero by the discrete-time index N/2. Thus the location of zeros of the transmit waveform signal heavily determines N_{FFT}. For these simulations, it was determined that in order to assure the IF impulse response decayed to approximately 0.001, a FFT size needed to be on the order of about 2^{20} and therefore N_{FFT} was set to 2^{21} to ensure sufficient decay. This large FFT size resulted in a computation time for the IF that was several orders of magnitude larger than the MF.

The sparse representation's superior detection performance comes at a cost of a significant computational bottleneck via utilization of the Basis Pursuit algorithm. Basis Pursuit has computational complexity $O(N^3)$ [10] which is much less efficient than the FFT implementations for the MF and IF. Creating more computationally efficient ℓ_1 minimization methods is

a current topic of research. Recent work [13] has been able to reduce the complexity to *O(N)* for sensing matrices with specific properties. The main result from the sparse representation is that sparse reconstruction of the undersampled return signal allows for resolving closely spaced targets at a higher resolution than that of the return signal measurements. Reducing the computational complexity of the Basis Pursuit algorithm for sparse reconstruction coupled with undersampling the return signal would facilitate a practical real-time implementation of this receiver.

VI. CONCLUSION

The sparse representation of bat biosonar allows for reconstructing sparse target impulse responses with high resolution and superior detection performance than what can be accomplished using either an MF or IF receiver. The sensing method allows for undersampling the return signal as much as 1/6 of the Nyquist rate while largely maintaining resolution and detection performance for a modest range of sparsity values.

This research demonstrates a proof of concept approach applying a deterministic sensing methodology that loosely models the mammalian auditory system. The return signal can be reconstructed via ℓ_1 minimization methods such as Basis Pursuit. The advantage of this method is that closely spaced targets can be detected and resolved at a precision much higher than that of the measurements. The simulations demonstrate that even undersampled filterbank outputs can contain sufficient information to allow high resolution reconstruction of target impulse responses using basis pursuit to exploit the sparsity of the impulse response. This result is qualitatively similar to the temporal resolution bats achieve on the order of a few microseconds using auditory system neural processing whose precision can be on the order of hundreds of microseconds. This suggests that bats may also be exploiting sparsity assumption in processing their echolocation signals. We wish to make it quite clear that we are not suggesting the bat auditory system implements basis pursuit or another sparse reconstruction algorithm. Quite the opposite- we feel it is extraordinarily improbable that such computation is implemented by the bat's auditory system. What we do note is that coupling the degraded time resolution of an undersampled filterbank with an assumption of sparsity produces temporal hyper-acuity analogous to that of bat's natural echolocation systems.

Based on the evidence presented in this paper, the authors conclude that the application of a deterministic sensing methodology to the sparse bat biosonar problem is quite plausible. There are a number of areas of focus to fully determine if bat echolocation is indeed based on the sparse representation model. The first is adding more complexity to the sensing method so as to more accurately model the mammalian auditory system. Second, performing analysis at higher resolution levels shown and/or replicating the results of [3] for the deterministic sensing method would further solidify the legitimacy of this model. If these efforts yield successful results, behavioural experiments could then be designed to assess the accuracy of the model to what is observed in the natural world.

VII. ACKNOWLEDGEMENTS

This research was supported by the US Office of Naval Research grant N00014-07-1-0591. The authors wish to thank Jim Simmons and Herbert Peremans for their intriguing and inspiring conversations.

REFERENCES

1 A.N. Popper and R.R. Fay (Editors), *Hearing by Bats* (Springer-Verlag, New York, NY 1995)
2 A. V. Oppenheim and R. W. Schafer with J. R. Buck, *Discete-Time Signal Processing, Second Edition* (Prentice-Hall, Englewood Cliffs, NJ, 1999)
3 B. Fontaine and H. Peremans, "Determining Biosonar Images using Sparse Approximations" *J. Acoust. Soc. Am*. 125**,** 3052-3059 (2009)
4 E. Candes, J. Romberg, and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information" *IEEE Trans. Inform. Theory*, vol. 52, no. 2, 489-509 (2006)
5 E. Candes and M. B. Wakin, "An Introduction to Compressive Sampling" *IEEE Signal Processing Magazine*, 21-30 (March, 2008)
6 J. A. Simmons, M. J. Ferragamo, and C. F. Moss "Echo-Delay Resolution in Sonar Images of the Big Brown Bat, *Eptesicus Fuscus*" *Proc. Natl. Acad. Sci. USA*, Vol. 95, 12647-12652 (1998)
7 J. A. Tropp, M. B. Wakin, M. F. Duarte, D. Baron, and R. Baraniuk, "Random Filters for Compressive Sampling and Reconstruction" *Proc. IEEE ICASSP*, 872-875 (2005)
8 N. Sharma and J. R. Buck, "A Generalized Linear Filter Approach for Sonar Receivers" *IEEE Digital Signal Processing Workshop*, 507-512 (2009)
9 R. A. DeVore, "Deterministic Constructions of Compressed Sensing Matrices" *Journal of Complexity*, 918-925 (2007)
10 R. G. Baraniuk, "Compressive Sensing" *IEEE Signal Processing Magazine*, 118-124 (July, 2007)
11 R. G. Baraniuk and P. Steeghs "Compressive Radar Imaging" *Proceedings of IEEE Radar Conference*, 128-133, (2007)
12 W. S. Burdic, *Underwater Acoustic Systems Analysis* (Prentice-Hall, Englewood Cliffs, NJ, 1991)
13 W. Xu and B. Hassibi "Efficient Compressive Sensing with Deterministic Guarantees Using Expander Graphs", *ITW 2007*, 414-419 (2007)