Compressive Sensing of Streams of Pulses

Chinmay Hegde and Richard G. Baraniuk
Department of Electrical and Computer Engineering
Rice University

***Abstract*—Compressive Sensing (CS) has developed as an enticing alternative to the traditional process of signal acquisition. For a length-N signal with sparsity K, merely $M = \mathcal{O}(K \log N) \ll N$ random linear projections (measurements) can be used for robust reconstruction in polynomial time. Sparsity is a powerful and simple signal model; yet, richer models that impose additional *structure* on the sparse nonzeros of a signal have been studied theoretically and empirically from the CS perspective.**

In this work, we introduce and study a sparse signal model for *streams of pulses*, i.e., S-sparse signals convolved with an unknown F-sparse impulse response. Our contributions are threefold: (i) we geometrically model this set of signals as an infinite union of subspaces; (ii) we derive a sufficient number of random measurements M required to preserve the metric information of this set. In particular this number is linear merely in the number of degrees of freedom of the signal $S + F$, and *sublinear* in the sparsity $K = SF$; (iii) we develop an algorithm that performs recovery of the signal from M measurements and analyze its performance under noise and model mismatch. Numerical experiments on synthetic and real data demonstrate the utility of our proposed theory and algorithm. Our method is amenable to diverse applications such as the high-resolution sampling of neuronal recordings and ultra-wideband (UWB) signals.

I. Introduction

Digital signal processing systems face two parallel challenges. With the availibility of ubiquitous computing power, memory and communication bandwidth, the pressure is on *acquisition* devices, such as analog-to-digital converters, to develop the ability to capture signals arising from a plethora of sources at ever increasing rates. On the other hand, to counter the "digital data deluge," DSP systems must develop efficient *compression* schemes that preserve the essential information contained in the signals of interest.

Discrete time signal acquisition is fundamentally governed by the Shannon/Nyquist sampling paradigm, which guarantees that all the information contained in the signal is preserved if it is uniformly sampled at a rate twice as fast as the bandwidth of its Fourier transform. Conversely, *transform compression* involves representing a signal sampled at the Nyquist rate $x \in \mathbb{R}^N$ in a suitable expansion $x = \Psi\alpha$, with Ψ being an $N \times N$ basis matrix. The number of large coefficients in the basis expansion α is known as the *sparsity* of the signal in the basis Ψ; for many interesting classes of signals, $K \ll N$. The JPEG compression system for images exploits the fact that naturally occuring images are sparse (or nearly sparse) in the Fourier basis. An intriguing question can thus be asked: is it possible to address the above two challenges in one shot, i.e., can a single system attain the twin goals of signal acquisition and compression?

Surprisingly, the answer in many cases is *yes*. Addressing this issue is the central tenet in Compressive Sensing (CS) [1, 2]. A prototypical CS system works as follows: a K–sparse signal x of length N is sampled by measuring its inner product with $M \ll N$ vectors; therefore, the output of the sampling system is given by the vector $y = \Phi x = \Phi\Psi\alpha$, where $\Phi \in \mathbb{R}^{M \times N}$ is a non-invertible matrix. CS theory states that x can be *exactly* reconstructed from y, provided the elements of Φ are chosen randomly from certain probability distributions, and the number of samples $M = \mathcal{O}(K \log(N/K))$ so that it is linear in the sparsity K and only logarithmic in the signal dimension N. Further, this recovery can be carried out in polynomial time, using efficient greedy approaches or optimization based methods [3, 4].

Thus, CS can be viewed as a new information acquisition paradigm which exploits a simple signal model (sparsity) to motivate a low-complexity representation (non-adaptive random measurements) for discrete-time signals. Nevertheless, depending on the application at hand, there may exist richer signal models that encode various types of interdependencies among signal components. For instance, the sparse nonzero coefficients of a signal may be grouped according to certain pre-defined blocks of fixed sizes. Recent work has led to the development of CS theory and algorithms that are based on *structured sparsity* models that are equivalent to a *finite* union of subspaces [5, 6]. For many models, the number of random measurements M can be significantly reduced, while preserving computational efficiency and robustness to noise.

In this paper, we study the compressive sensing of *streams of pulses*, i.e., the set of S-sparse signals that are convolved with an *unknown* F-sparse impulse response. Such signals widely occur in practice. For instance, neuronal spike trains can be viewed as a sparse set of spikes of varying heights convolved with the signature impulse response of the particular neuron. Another example would be the low-pass filtered output of a digital ultra-wideband (UWB) receiver. The overall sparsity of such a signal $K = SF$; thus, a conventional CS system would acquire $M = \mathcal{O}(SF \log(N/SF))$ compressive measurements for signal recovery.

Our particular contributions are as follows. First, we develop a deterministic model for pulse streams that can be geometrically viewed as an *infinite* union of subspaces. Second, we derive a bound on the number of random measurements M that preserve the geometric structure of the signals, thereby enabling their stable recovery. Importantly, our derivation shows that $M = \mathcal{O}(F + S \log N)$, i.e., the number of measurements required for information preservation is *sublinear* in the sparsity $K = SF$ and proportional to the number of degrees of freedom $S + F$. Third, we develop a polynomial time algorithm that recovers any signal belonging to this set from M measurements. Numerical experiments on real and synthetic data sets demonstrate the benefits of our approach.

Our work represents the first attempt to adopt a deterministic union-of-subspaces model for streams of pulses; this enables us to derive rigorous guarantees for the number of measurements required for stable recovery. The algorithm for CS recovery developed in this paper can be linked to various concepts in the literature such as best basis compressive sensing [7], simultaneous sparse approximation and dictionary learning [8], and the classical signal processing problem of blind deconvolution [9]. We obtain significant gains over conventional CS recovery methods, particularly in terms of reducing the number of measurements required for stable recovery, as evident from the example in Figure 1.

The rest of the paper is organized as follows. In Section II, we review the rudiments of standard and structured sparsity-based CS. In Section III, we introduce our signal model and examine its geometric properties. In section IV, we derive a lower bound on the number of random measurements required to sample this signal set. In Section V, we develop an algorithm for the stable signal recovery and discuss its properties. Numerical results are presented in Section VI, followed by conclusions in Section VII.

II. Background

A. The geometry of sparse signal models

A signal $x \in \mathbb{R}^N$ is K-*sparse* in the orthonormal basis Ψ if the corresponding basis expansion $\alpha = \Psi^T x$ contains no more than K nonzero elements. In the sequel, unless otherwise noted, the sparsity basis Ψ is assumed to be the identity matrix for $\mathbb{R}^N$. The locations of the nonzeros of x can additionally be encoded by a binary vector of length N with a 1 indicating a nonzero; this vector $s(x)$ is called the *support* of x. Denote the set of all K-sparse signals in $\mathbb{R}^N$ as Σ_K. Geometrically, Σ_K can be identified as the union of $\binom{N}{K}$, K-dimensional subspaces of $\mathbb{R}^N$, with each subspace being the linear span of exactly K canonical unit vectors of $\mathbb{R}^N$.

Often, we are interested in sparse signal ensembles which exhibit more complex dependencies in terms of their nonzero values and locations. For instance, the signals of interest might admit only a small number of support configurations. Such classes of signals may also be modeled

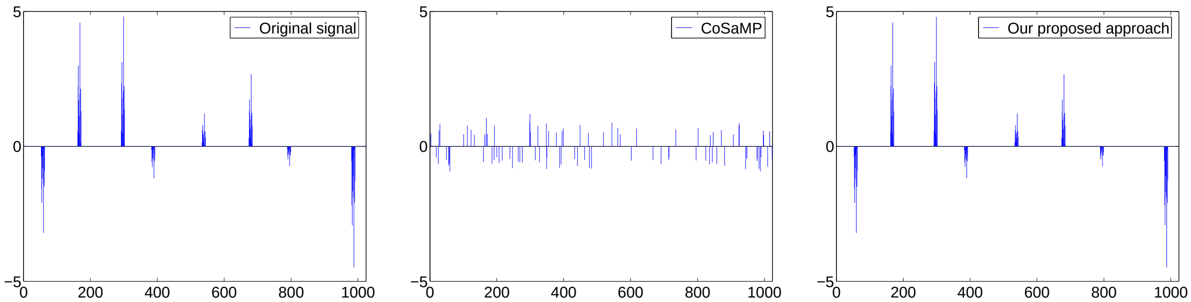

(a) (b) (c)

Fig. 1. *Example CS recovery of a stream of pulses. (a) Test signal of length $N = 1024$ comprising $S = 8$ pulses of length $F = 11$, so that the signal sparsity $K = 88$. Signals are recovered from $M = 90$ random Gaussian measurements using (b) an iterative sparse approximation algorithm (CoSaMP [10]) (c) Our proposed Algorithm 1. Near-perfect recovery is achieved using our approach.*

by a union of subspaces, consisting only of L_K canonical subspaces (so that $L_K \leq \binom{N}{K}$). Thus, if $\Omega = \{\Omega_1, \ldots, \Omega_{L_K}\}$ denote the set of permitted supports, a *structured sparsity model* [5] can be defined as the set:

$$\mathcal{M}_K := \{x : \text{supp}(x) \in \Omega\}. \qquad (1)$$

Any structured sparsity model $\mathcal{M}_K$ is itself contained in the set Σ_K. An intuitive interpretation is as follows: the smaller the value of L_K, the "smaller" the signal set $\mathcal{M}_K$, and the more restrictive the model.

B. Stable embedding via linear measurements

Suppose instead of collecting all the coefficients of a vector $x \in \mathbb{R}^N$, we merely record M inner products (measurements) of x with $M < N$ preselected vectors; this can be represented in terms of a linear transformation $y = \Phi x, \Phi \in \mathbb{R}^{M \times N}$. The central tenet of CS is that x can be *exactly* recovered from y, even though Φ is necessarily low-rank and has a nontrivial nullspace. In particular, a condition on Φ known as the *restricted isometry property* (RIP) can be defined as follows.

Definition 1: [11] An $M \times N$ matrix Φ has the K-RIP with constant δ_K if, for all $x \in \Sigma_K$,

$$(1 - \delta_K)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 + \delta_K)\|x\|_2^2. \qquad (2)$$

The RIP requires Φ to leave the norm of every sparse signal approximately invariant; in particular, Φ must necessarily not contain any sparse vectors in its nullspace. In a similar fashion, the $bK-$RIP can be defined in cases where Φ preserves the norms of $b-$wise differences of K-sparse signals, with b being a small integer (2 or 3). An important hallmark of CS is as follows: provided $M \geq \mathcal{O}(K \log(N/K))$, matrices whose elements are chosen as i.i.d. samples from a random subgaussian distribution work with high probability. Thus, M is linear in the sparsity of the signal set K and only logarithmic in the native dimension N. The RIP proves to be an essential component in the design of stable reconstruction algorithms as discussed below.

An analogous isometry condition can be proposed for structured sparsity models [5, 12]; Φ is said to satisfy the *$\mathcal{M}_K$-RIP* if Equation 2 holds for all $x \in \mathcal{M}_K$. For general structured sparsity models defined by L_K canonical subspaces, the RIP can be attained with high probability provided $M \geq \mathcal{O}(K + \log(L_K))$. Two observations can be made. First, the number of measurements M is logarithmic in the *number* of subspaces in the model, i.e., a "smaller" signal set can be embedded using fewer random linear measurements. Second, M has to be at least as large as the number of nonzeros K of the measured signal.

C. Recovery from compressive measurements

The CS recovery problem is to perform stable, feasible inversion of the operator Φ, given M measurements $y = \Phi x$. A number of CS recovery algorithms have been proposed [3, 4] with the primary tradeoffs being the number of measurements, recovery time and robustness to noise. More recently, iterative support selection algorithms (such as CoSaMP [10]) have emerged that offer uniform, stable guarantees while remaining computationally efficient. An added advantage of these iterative algorithms is that with a simple modification, they can be used to reconstruct signals belonging to *any* structured sparsity model [5]. If $y = \Phi x + n$, where $x \in \mathcal{M}_K$ and Φ satisfies the $\mathcal{M}_K$-RIP, then the estimate $\widehat{x}$ obtained by the CS recovery algorithm satisfies the following bound:

$$\|x - \widehat{x}\|_2 \leq C\|n\|_2.$$

When there is no noise term n, the estimate $\widehat{x}$ exactly coincides with the signal x. In other words, given a sufficient number of nonadaptive linear measurements y, any signal x belonging to a particular sparse signal model can be exactly reconstructed in polynomial time.

To summarize:, at the core of CS lie three key concepts: a signal model exhibiting a particular type of geometry in high-dimensional space; the construction of a low-rank linear transformation with particular properties defined on the model; and the development of methods to perform stable, efficient inversion of this mapping onto its domain.

III. Pulse streams

In several applications, the assumption of exact sparsity under a basis transform is only approximate. An EEG recording of a single neuron may be approximated by a stream of spikes, but can be better modeled by a stream of pulses, the shape of each pulse being a characteristic of the neuron. A high-resolution image of the night sky will consist of a field of points (corresponding to the locations of the stars) convolved with the point spread function of the imaging device. A similar pulse-broadening effect can be observed in the output of high speed UWB (ultra-wideband) receivers. The Discrete Fourier Transform of a sum of nonharmonic sinusoids is not a stream of spikes, but rather the convolution of this stream with a sinc function; this increases the apparent number of nonzero coefficients of the underlying signal in the Fourier domain (DFT leakage).

Pulse streams have been briefly studied from a CS perspective [13]; our objective in this paper is to develop a comprehensive CS framework for such signals. We begin with the choice of an appropriate signal model. A general model for streams of pulses can be defined thus:

Definition 2: Let $\mathcal{M}_S^x$, $\mathcal{M}_F^h$ be structured sparsity models defined in $\mathbb{R}^N$. Define the set:

$$\mathcal{M}_{S,F}^z := \{z \in \mathbb{R}^N : z = x * h, \text{ such that } x \in \mathcal{M}_S^x \text{ and } h \in \mathcal{M}_F^h\},$$

where $*$ denotes the circular convolution operator in $\mathbb{R}^N$. Then, $\mathcal{M}_{S,F}^z$ is called a (S, F) pulse-stream model.

Owing to the commutative property of the convolution operator, an element z in $\mathcal{M}_{S,F}^z$ can be represented in multiple ways:

$$z = x * h = Hx = Xh,$$

where H (respectively, X) is a square circulant matrix with its columns comprising circularly shifted versions of the vector h (respectively, x). For a given z, H and x need not be unique. Any $(\alpha H, x/\alpha)$ satisfies the above equality; so does (H', x'), where H' is generated by a circularly shifted version of h by a time delay $+\tau$ and x' is a circularly shifted version of x by $-\tau$. A more concise pulse-stream model can be introduced by making the following two assumptions:

1) the filter coefficients are *minimum phase*, i.e. all its nonzero coefficients are concentrated at the start of the impulse response. Thus, the model for the filter vector h consists of the lone subspace spanned by the first F canonical unit vectors.

2) the sparse spikes are sufficiently separated in time. A deterministic structured sparsity model for such signals has been recently introduced in [14] and ensures that consecutive spikes are separated by at least Δ locations from one another.

This model is useful since it eliminates possible ambiguities that arise due to the shift invariant nature of convolution, i.e., a vector z belonging to this model consists of a stream of disjoint pulses and hence the locations of the nonzero spikes are uniquely defined. We denote this special pulse-stream model by $\mathcal{M}^{\Delta}_{S,F}$.

It is also useful to adopt the following geometric point of view. For a fixed $h \in \mathcal{M}^h_F$, the set $\{h * x : x \in \mathcal{M}^x_S\}$ forms a finite union of K-dimensional subspaces, owing to the fact that it is generated by the action of h on L^x_S canonical subspaces. Denote this set by $h(\mathcal{M}^x_S)$. Then, the (S, F) pulse-stream model can be written as:

$$\mathcal{M}^z_{S,F} = \bigcup_{h \in \mathcal{M}^h_F} h(\mathcal{M}^x_S).$$

Thus, our signal model is an *infinite union of subspaces*. The above interpretation remains unchanged if the roles of x and h are reversed. For simplicity of notation, we define $K = SF$, the maximum sparsity of the signals in our proposed model. Wherever possible, we will also drop the superscripts x, h, z and denote our trio of signal models in $\mathbb{R}^N$ as $\mathcal{M}_S, \mathcal{M}_F$ and $\mathcal{M}_K$, respectively. See Figure 1(a) for an example stream of pulses.

IV. Stable embedding of pulse streams

It is easy to see that $\mathcal{M}_K$, as defined above, is a subset of the set of all K-sparse signals Σ_K. On the other hand, only a minute fraction of all K-sparse signals can be written as the convolution of an S-sparse signal with an F-sparse filter. Intuition suggests that we should be able to achieve a stable embedding of this set using *fewer* random linear measurements than those required for the stable embedding of all K-sparse signals. Indeed, the following theorem makes this precise.

Theorem 1: Let $\mathcal{M}^x_S$ be a union of L^x_S canonical subspaces, and $\mathcal{M}^h_F$ be a union of L^h_F canonical subspaces. Suppose $\mathcal{M}^z_{S,F}$ is the associated pulse-stream model. Then, there exists a constant c such that for any $t > 0$ and

$$M \geq \frac{c}{\delta^2}\left((S+F)\ln\left(\frac{1}{\delta}\right) + \log(L^x_S L^h_F) + t\right),$$

an $M \times N$ i.i.d. subgaussian matrix Φ will satisfy the following property with probability at least $1 - e^{-t}$: for every pair $z_1, z_2 \in \mathcal{M}^z_{S,F}$,

$$(1-\delta)\|z_1 - z_2\|_2^2 \leq \|\Phi z_1 - \Phi z_2\|_2^2 \leq (1+\delta)\|z_1 - z_2\|_2^2.$$

The proof of this theorem is provided in the expanded version of this manuscript [15]. Theorem 1 indicates that the number of measurements required for the stable embedding of signals in $\mathcal{M}^z_K$ is proportional to $(S + F)$; thus, it is *sublinear* in the maximum sparsity of the signals SF. Existing models for structured sparsity require at least $2K = 2SF$ linear measurements to ensure approximate preservation of pairwise distances. Our proposed model improves upon such *signal support* models by introducing implicit correspondences between the *amplitudes* of the nonzero coefficients via the convolution operator. The bound in Theorem 1 is linear in the number of degrees of freedom $S + F$, and therefore is essentially optimal for the signal class $\mathcal{M}^z_K$.

Theorem 1 is valid for sparse signals and filters belonging to arbitrary models. For the special pulse-stream model $\mathcal{M}^{\Delta}_{S,F}$, the following corollary is a consequence of Theorem 1 and Theorem 1 of [14].

Corollary 1: An $M \times N$ i.i.d. subgaussian random matrix satisfies the RIP for signals belonging to $\mathcal{M}^{\Delta}_{S,F}$ with high probability if

$$M \geq \mathcal{O}\left(S + F + S\log(N/S - \Delta)\right).$$

V. Stable recovery of pulse streams

The CS recovery problem for the pulse-stream model can be stated as follows: given noisy measurements of a stream of pulses:

$$y = \Phi z + n = \Phi H x + n = \Phi X h + n,$$

the goal is to reconstruct z from y. Standard or structured sparsity methods for CS recovery are unsuitable for this problem, since both x (respectively, X) and h (respectively, H) are *unknown* and have to be simultaneously inferred. This task is similar to performing *blind deconvolution* [9], which attempts simultaneous inference of the spike locations and filter coefficients, the key difference in our case being that we are only given access to the random measurements y and not Nyquist-rate samples x.

Suppose that the spikes are separated by a minimum separation distance Δ and that the filter is minimum phase. An algorithm for reconstruction can be proposed thus: we fix a candidate support configuration Ω for the spike domain. Then, we form the circulant matrix H from the current estimate of our filter $\widehat{h}$ (so that $\widehat{H} = \mathbb{C}(\widehat{h})$), calculate the spike *dictionary* $\Phi\widehat{H}$ and select only those columns that correspond to the assumed spike locations Ω. This transforms our problem into an overdetermined system, which can be solved using least-squares. Once the spike coefficients have been inferred, we may use the commutativity property of the convolution operator, form the filter dictionary $\Phi\widehat{X}$ and solve a similar least-squares problem for the filter coefficients. This process is repeated until convergence. The overall reconstruction problem can be solved by repeating this process for every support configuration Ω belonging to the spike model. If the matrix Φ contains a sufficient number of rows M (such as specified by Theorem 1), it can be shown that if this algorithm converges, it recovers the correct solution (the proof is detailed in the expanded version of this manuscript [15]).

However, this algorithm involves iteratively solving a combinatorial number of estimation problems and is infeasible for large N. A simpler (suboptimal) method is to leverage a recent algorithm for CS recovery that uses the separated spike train model. As opposed to cycling through every possible support configuration for the spikes, we instead update the support configuration at each step based on the current estimates of the sparse signal and filter coefficients. This can be shown to be equivalent to solving a suitable linear program (for details, refer [14]). Denote this support update by $\mathbb{D}(\cdot)$. Once a support has been chosen, we repeat the least squares steps as above to solve for the sparse signal coefficients and filter coefficients respectively. This process is iterated until convergence. The modified algorithm can be viewed as an iterative sparse approximation process that continually updates its estimate of the sparsifying dictionary. The procedure is detailed in pseudocode form in Algorithm 1.

VI. Numerical experiments

Figure 1 demonstrates the considerable advantages that our proposed algorithm offers in terms of the number of compressive measurements required for reliable reconstruction. The test signal was generated by choosing $S = 8$ spikes with random amplitudes and locations and convolving this spike stream with a minimum phase filter ($F = 11$) with a randomly chosen impulse response. The overall sparsity of the signal $K = SF = 88$; thus, even the best sparsity-based CS algorithm would require $2K = 176$ measurements. Our approach (Algorithm 1) returns an accurate estimate of both the spike signal as well as the filter impulse response using merely $M = 90$ measurements.

Figure 2 displays the averaged results of a Monte Carlo simulation of our algorithm over 200 trials. Each trial was conducted by generating a sample signal belonging to $\mathcal{M}_K$, computing M linear random Gaussian measurements, reconstructing with different algorithms and recording the magnitude of the recovery error for different values of the overmeasuring factor M/K. It is clear that our proposed approach outperforms both conventional CS recovery (CoSaMP [10]) with target sparsity $K = SF$, as well as block-based reconstruction [5] with knowledge of the size and number of blocks (resp. F and S). In fact, our algorithm performs nearly as well as the oracle decoder that possesses perfect prior knowledge of the filter coefficients and solves only for the sparse signal coefficients.

Further, we show that our algorithm is stable to noise in the signal and measurement domains. We generate a length-1024 signal comprising $S = 8$ pulses of width $F = 11$, add a small amount of Gaussian noise to all its components, compute 150 noisy linear measurements and reconstruct using Algorithm 1. As is evident from Figure 3, our proposed algorithm provides a good approximation of the original (noisy) stream of pulses.

Figure 4(a) shows the electrochemical spiking potential of a single neuron measured using an EEG. The shape of the pulses is characteristic of the neuron; however, there exist minor fluctuations in the amplitudes, locations and profiles of the pulses. Despite the apparent model mismatch, our algorithm recovers a good approximation (Figure 4(b)) to the original signal. The inferred impulse response can be viewed as an a scale-invariant average across different pulses and can be potentially used to distinguish between the firings of various neurons.

Algorithm 1

Inputs: Projection matrix Φ, measurements y, model parameters Δ, S, F.
Output: $\mathcal{M}_{S,F}^{\Delta}$-sparse approximation $\widehat{z}$ to true signal z
$\widehat{x} = 0$, $\widehat{h} = (\mathbf{1}_F^\top, 0, \ldots, 0)$; $i = 0$ {initialize}
while halting criterion false **do**
1. $i \leftarrow i + 1$
2. $\widehat{z} \leftarrow \widehat{x} * \widehat{h}$ {current signal estimate }
3. $\widehat{H} = \mathbb{C}(\widehat{h}), \Phi_h = \Phi \widehat{H}$ {form dictionary for spike domain}
4. $e \leftarrow \Phi_h^T (y - \Phi_h \widehat{x})$ {residual}
5. $\Omega \leftarrow \mathrm{supp}(\mathbb{D}_2(e))$ {prune residual according to $(2S, 2\Delta)$ model}
6. $T \leftarrow \Omega \cup \mathrm{supp}(\widehat{x}_{i-1})$ {merge supports}
7. $b|_T \leftarrow (\Phi_h)_T^\dagger y$, $b|_{T^C} = 0$ {update spike estimate}
8. $\widehat{x} \leftarrow \mathbb{D}(b)$ {prune spike estimate according to (S, Δ) model}
9. $\widehat{X} = \mathbb{C}(\widehat{x}), \Phi_x = \Phi \widehat{X}$ {form dictionary for filter domain }
10. $\widehat{h} \leftarrow \Phi_x^\dagger y$ {update filter estimate }

end while
return $\widehat{z} \leftarrow \widehat{x} * \widehat{h}$

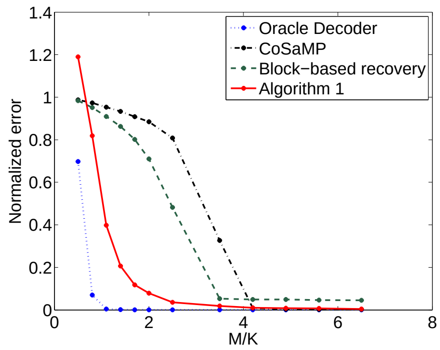

Fig. 2. *Reconstruction error vs. M/K for different reconstruction algorithms averaged over 200 sample trials. $N = 1024$, $S = 8$, $F = 11$. Our method outperforms standard and structured sparsity-based methods, particularly in low-measurement regimes.*

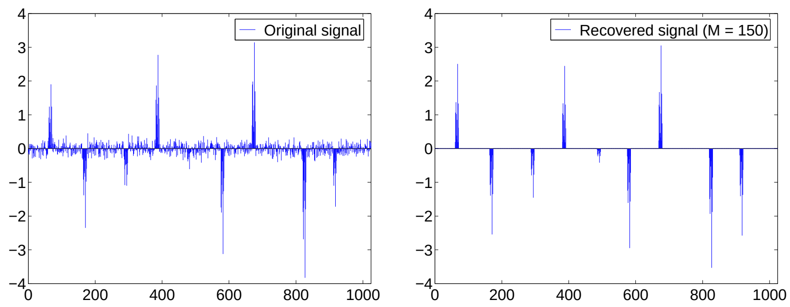

Fig. 3. *(a) Synthetic noisy signal. (b) Recovery using 150 measurements. The quality of the reconstructed signal demonstrates that our algorithm is robust to signal and measurement domain noise.*

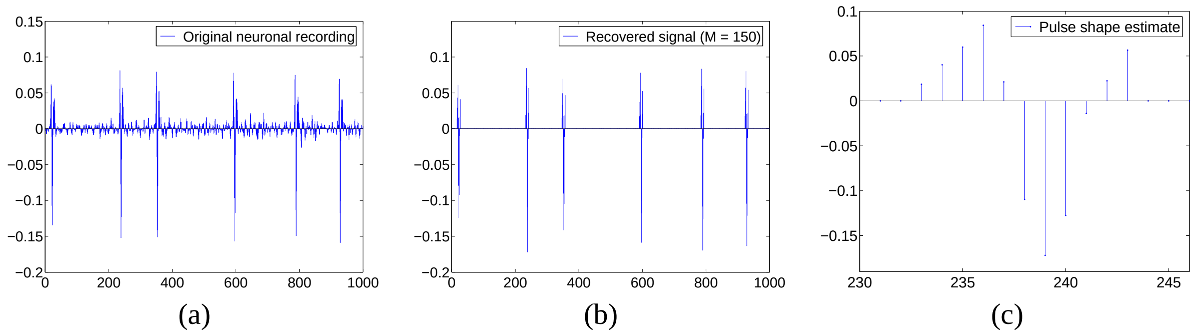

(a) (b) (c)

Fig. 4. *Experiment with real EEG data. (a) Single neuron recording (N = 1000). The signal consists of a series of pulses of approximately the same shape. (b) Recovered signal using $M = 150$ measurements. (c) Estimated pulse profile. Our algorithm is robust to small variations in pulse shapes.*

VII. Conclusions

In this paper, we have introduced and analyzed the compressive sensing of sparse pulse streams. This signal set can be modeled as an infinite union of subspaces which exhibits a particular geometric structure. This enables us to quantitatively deduce that the number of measurements needed for the stable embedding and recovery of such signals is much smaller than that required for conventional or structured sparsity-based CS. We motivate an efficient algorithm that performs signal recovery from this reduced set of measurements and numerically demonstrate its benefits over state-of-the-art methods for CS recovery. Though our theoretical and empirical results are promising, we do not yet possess a precise theoretical characterization of the convergence of our proposed algorithm under non-idealities such as noise and model mismatch; we defer this to future research.

Acknowledgments

This work was supported by NSF grants CCF-0431150, CCF-0728867, CNS-0435425, and CNS-0520280, DARPA/ONR N66001-08-1-2065, ONR N00014-07-1-0936, N00014-08-1-1067, N00014-08-1-1112, and N00014-08-1-1066, AFOSR FA9550-07-1-0301, ARO MURI W311NF-07-1-0185, W911NF-09-1-0383, and the Texas Instruments Leadership University Program.

References

[1] D. L. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, pp. 1289–1306, September 2006.

[2] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Info. Theory*, vol. 52, pp. 489–509, Feb. 2006.

[3] J. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," *IEEE Trans. Info. Theory*, vol. 53, pp. 4655–4666, Dec. 2007.

[4] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Journal on Scientific Computing*, vol. 20, p. 33, 1998.

[5] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," 2008. Preprint. Available at http://dsp.rice.edu/cs.

[6] Y. Eldar and M. Mishali, "Robust recovery of signals from a union of subspaces," 2008. Preprint.

[7] G. Peyre, "Best-basis compressed sensing," 2007. Preprint.

[8] M. Yaghoobi, T. Blumensath, and M. E. Davies, "Dictionary learning for sparse approximations using the majorization method," 2009. IEEE Trans. Sig. Proc.

[9] S. Haykin, "Blind deconvolution," 1994. Prentice Hall.

[10] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, June 2008.

[11] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematicians*, vol. 3, (Madrid, Spain), pp. 1433–1452, 2006.

[12] T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Trans. Info. Theory*, Dec. 2008.

[13] F. Naini, R. Gribonval, L. Jacques, and P. Vandergheynst, "Compressive sampling of pulse trains: spread the spectrum!," 2008. Preprint.

[14] C. Hegde, M. F. Duarte, and V. Cevher, "Compressive sensing recovery of spike trains using a structured sparsity model," in *SPARS*, (Saint Malo, France), April 2009.

[15] C. Hegde and R. G. Baraniuk, "Sampling and recovery of a sum of sparse pulses," 2009. In preparation.