Sampling and Recovery of Pulse Streams

Chinmay Hegde and Richard G. Baraniuk

Department of Electrical and Computer Engineering

Rice University

Abstract

Compressive Sensing (CS) is a new technique for the efficient acquisition of signals, images, and other data that have a sparse representation in some basis, frame, or dictionary. By sparse we mean that the N-dimensional basis representation has just $K \ll N$ significant coefficients; in this case, the CS theory maintains that just $M = \mathcal{O}(K \log N)$ random linear signal measurements will both preserve all of the signal information and enable robust signal reconstruction in polynomial time. In this paper, we extend the CS theory to *pulse stream* data, which correspond to S-sparse signals/images that are convolved with an unknown F-sparse pulse shape. Ignoring their convolutional structure, a pulse stream signal is $K = SF$ sparse. Such signals figure prominently in a number of applications, from neuroscience to astronomy. Our specific contributions are threefold. First, we propose a pulse stream signal model and show that it is equivalent to an infinite union of subspaces. Second, we derive a lower bound on the number of measurements M required to preserve the essential information present in pulse streams. The bound is linear in the total number of degrees of freedom $S + F$, which is significantly smaller than the naive bound based on the total signal sparsity $K = SF$. Third, we develop an efficient signal recovery algorithm that infers both the shape of the impulse response as well as the locations and amplitudes of the pulses. The algorithm alternatively estimates the pulse locations and the pulse shape in a manner reminiscent of classical deconvolution algorithms. Numerical experiments on synthetic and real data demonstrate the advantages of our approach over standard CS.

I. INTRODUCTION

Digital signal processing systems face two parallel challenges. On the one hand, with ubiquitous computing power, memory and communication bandwidth, the pressure is on *acquisition* devices, such as analog-to-digital converters and digital cameras, to capture signals at ever increasing sampling rates. To date, signal acquisition has been governed by the Shannon/Nyquist sampling theorem, which states that all the information contained in a signal is preserved if it is uniformly sampled at a rate twice the bandwidth of its Fourier transform. On the other hand, to counter the resulting deluge of Nyquist-rate samples, DSP systems must utilize efficient *compression* schemes that preserve the essential information contained in the signals of interest. Transform compression of a discrete-time signal $x \in \mathbb{R}^N$ involves representing the signal in a suitable basis expansion $x = \Psi\alpha$, with Ψ an $N \times N$ basis matrix, and storing only the K largest basis coefficients. The number of large coefficients in α is known as the *sparsity K* of the signal in the basis Ψ. For many classes of interesting signals, $K \ll N$, and hence efficient signal compression can be achieved.

An intriguing question can thus be asked: can a system simultaneously attain the twin goals of signal acquisition and compression? Surprisingly, the answer in many cases is *yes*. This question forms the core of the burgeoning field of Compressive Sensing (CS) [3, 4]. A prototypical CS system works as follows: a signal x of length N is sampled by measuring its inner products with $M \ll N$ vectors; the output of the sampling system is thus given by the vector $y = \Phi x = \Phi\Psi\alpha$, where $\Phi \in \mathbb{R}^{M \times N}$ is a non-invertible matrix. The CS theory states that, with high probability, x can be exactly reconstructed from y provided that (i) the elements of Φ are chosen randomly from certain probability distributions, and (ii) the number of samples M is $\mathcal{O}\left(K \log(N/K)\right)$. Further, this recovery can be carried out in polynomial time using efficient greedy approaches or optimization based methods (e.g., [5, 6]).

For some applications, there exist more selective signal models than simple sparsity that encode various types of inter-dependencies among the locations of the nonzero signal components.

Recent work has led to the development of CS theory and algorithms based on *structured sparsity* models that are equivalent to a finite union of subspaces [7, 8]. By exploiting the dependencies present among the nonzero coefficients, M can be significantly reduced; for certain structured sparsity models, with high probability the number of measurements M required for exact recovery is merely $\mathcal{O}(K)$ (without the additional logarithmic dependence on the signal length N).

Despite the utility of sparsity models, in many real-world sensing applications the assumption of sparsity itself is an oversimplification. For example, an electrophysiological recording of a neuron is often approximated as a series of narrow spikes but is more accurately modeled as a series of more elongated pulses, the pulse shape being characteristic to the particular neuron. As another example, a high-resolution image of the night sky consists of a field of points (corresponding to the locations of the stars) convolved with the point spread function of the imaging device. Such signals can be modeled as an S-sparse *spike stream* that have been convolved with an unknown F-sparse *impulse response* so that the resulting overall sparsity $K = SF$. We call such a signal a *pulse stream*. Applying standard compressive sensing and recovery to such a pulse stream would require $M = \mathcal{O}(SF\log(N/SF))$ measurements, and thus incur a corresponding multiplicative increase by a factor of F as compared to sensing merely the underlying spike stream (in which case $M = \mathcal{O}(S\log N/S)$); this can be prohibitive in some situations. Thus, it is essential to develop a CS framework that can handle not just sparse signals but also more general pulse streams.

In this paper, we take some initial steps towards such a CS pulse stream framework. First, we propose a deterministic signal model for pulse streams. We show that our proposed model is equivalent to a specific instance of an *infinite union of subspaces*. Second, as our main theoretical contribution, we derive a bound on the number of random linear measurements M required to preserve the essential information contained in such signals. The proof relies on the particular high-dimensional geometry exhibited by the pulse stream model. Our derivation shows that $M = \mathcal{O}((S + F)\log N)$; i.e., M must be proportional to the number of degrees of freedom of the signal $S + F$ but *only sublinear* in the total sparsity $K = SF$. Third, we develop algorithms to recover signals from our model. Under certain additional restrictions on the signals of interest, one of the algorithms provably recovers both the spike stream and the impulse response. We analyze its convergence, computational complexity, and robustness to variations in the pulse
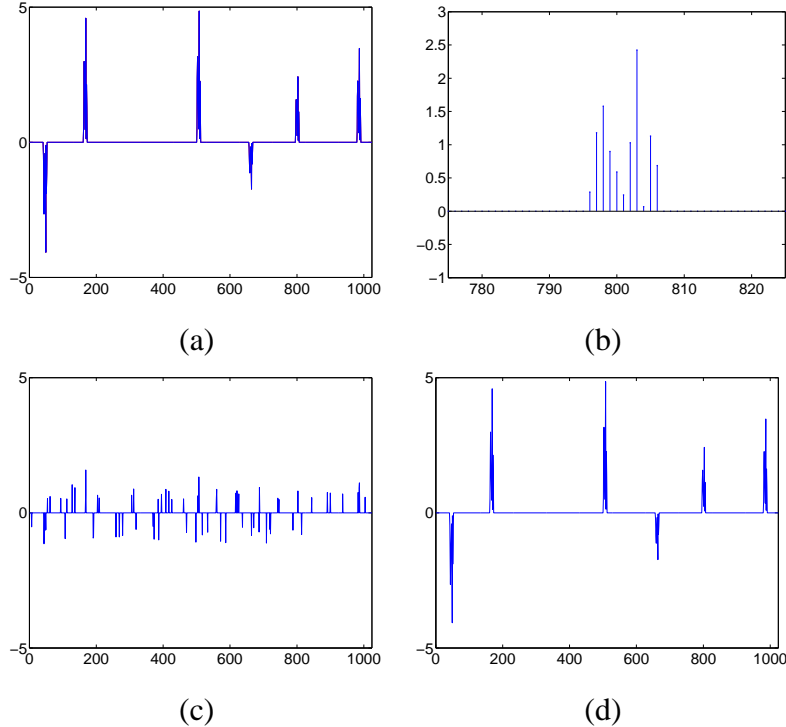
Fig. 1. *(a) Test signal of length $N = 1024$ obtained by convolving a spike stream with $S = 6$ spikes with an impulse response of length $F = 11$, so that the total signal sparsity $K = SF = 66$. (b) Profile of one pulse . Signal recovery from $M = 100$ random Gaussian measurements performed using (c) a state-of-the-art CS recovery algorithm (CoSaMP [9], MSE = 13.42), and (d) our proposed Algorithm 2 (MSE = 0.0028).*

shape. Numerical experiments on real and synthetic data sets demonstrate the benefits of the approach. As demonstrated in Figure 1, we obtain significant gains over conventional CS recovery methods, particularly in terms of reducing the number of measurements required for recovery.

This paper is organized as follows. In Section II, we review the rudiments of standard and structured sparsity-based CS. In Section III, we propose a deterministic signal model for pulse streams and discuss its geometric properties. In Section IV, we derive bounds on the number of random measurements required to sample signals belonging to our proposed model. In Section V, we develop an algorithm for stable signal recovery and analyze its convergence and robustness to model mismatch. Numerical results are presented in Section VI, followed by a concluding discussion in Section VII.

II. Background on Compressive Sensing

A. Sparse signal models

A signal $x \in \mathbb{R}^N$ is K-*sparse* in the orthonormal basis $\Psi \in \mathbb{R}^{N \times N}$ if the corresponding basis representation $\alpha = \Psi^T x$ contains no more than K nonzero elements. Without loss of generality, we assume the sparsity basis Ψ 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 $\sigma(x)$ 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 $\binom{N}{K}$ subspaces of $\mathbb{R}^N$, with each subspace the linear span of exactly K canonical unit vectors of $\mathbb{R}^N$. For a general $x \in \mathbb{R}^N$, we define its best K-sparse approximation x_K as

$$x_K = \arg \min_{u \in \Sigma_K} \|x - u\|_2. \tag{1}$$

Many signals of interest exhibit more complex dependencies in terms of their nonzero values and locations. Signals that permit only a small number of admissible support configurations can be modeled by a restricted union of subspaces, consisting only of L_K canonical subspaces (with $L_K \ll \binom{N}{K}$). If $\Sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_{L_K}\}$ denotes the restricted set of admissible supports, then a *structured sparsity model* [7] is the set

$$\mathcal{M}_K := \{x : \sigma(x) \in \Sigma\}. \tag{2}$$

B. Signal acquisition via nonadaptive linear measurements

Suppose that instead of collecting all the coefficients of a vector $x \in \mathbb{R}^N$, we merely record M inner products (measurements) of x with $M < N$ pre-selected vectors; this can be represented in terms of the linear transformation $y = \Phi x, \Phi \in \mathbb{R}^{M \times N}$. Φ is called the *sampling matrix*; it is at most rank-M and hence has a nontrivial nullspace. The central result in the Compressive Sensing (CS) theory is that, despite the non-invertible nature of Φ, if x is *sparse*, then it can be exactly recovered from y if Φ satisfies a condition known as the restricted isometry property (RIP) [10]:

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

$$(1 - \delta_K)\|x\|_2^2 \le \|\Phi x\|_2^2 \le (1 + \delta_K)\|x\|_2^2. \tag{3}$$

A matrix Φ with the K-RIP essentially ensures a *stable embedding* of the set of *all* K-sparse signals Σ_K into a subspace of dimension M. The RIP requires Φ to leave the norm of every sparse signal approximately invariant; also, Φ must necessarily not contain any sparse vectors in its nullspace. At first glance, it is unclear if a matrix Φ that satisfies the RIP should even exist when $M < N$; indeed, deterministic design of a sampling matrix having the RIP is an NP-complete problem. Nevertheless, it has been shown that provided $M \geq \mathcal{O}\left(K \log(N/K)\right)$, a matrix Φ whose elements are i.i.d. samples from a random subgaussian distribution possesses the RIP with high probability [10]. Thus, M can be linear in the sparsity of the signal set K and *only logarithmic* in the signal length N.

An analogous isometry condition holds for structured sparsity models containing L_K canonical subspaces [7, 8, 11]. This is known as the *model-based RIP* and is defined thus: Φ satisfies the $\mathcal{M}_K$-*RIP* if (3) holds for all $x \in \mathcal{M}_K$. It can be shown [11] that the number of measurements M necessary for a subgaussian sampling matrix to have the $\mathcal{M}_K$-RIP with constant δ and with probability $1 - e^{-t}$ is bounded as

$$M \geq \frac{c}{\delta^2}\left(\ln(2L_K) + K \ln \frac{12}{\delta} + t\right). \tag{4}$$

We can make two inferences from (4). First, the number of measurements M is logarithmic in the *number* of subspaces in the model; thus, signals belonging to a more concise model can be sampled using fewer random linear measurements. Second, M is *at least linear* in the sparsity K of the measured signal.

C. Recovery methods

Given measurements $y = \Phi x$, CS recovery methods aim to find the "true" sparse signal x that generated y. One possible method is to seek the sparsest x that generates the measurements y, i.e.,

$$\widehat{x} = \arg \min_{x'} \|x'\|_0 \quad \text{subject to} \quad y = \Phi x'. \tag{5}$$

where the ℓ_0 "norm" of a vector x' denotes the number of nonzero entries in x'. This method can be used to obtain the true solution x provided $M \geq 2K$. However, minimizing the ℓ_0 norm can be shown to be NP-complete and is not stable in the presence of noise in the measurements [10].

If the sampling matrix Φ possesses the RIP, then tractable algorithms for CS recovery can be developed. These broadly follow two different approaches. The first approach entails solving a convex relaxation of (5):

$$\widehat{x} = \arg\min_{x'} \|x'\|_1 \quad \text{subject to} \quad y = \Phi x', \tag{6}$$

which corresponds to a linear program and hence can be solved in polynomial time. A common variant of this formulation includes accounting for noise of bounded magnitude in the measurements [6]. The second approach entails an iterative, greedy selection of the support $\sigma(x)$ of the true solution x. This approach is employed by several algorithms such as orthogonal matching pursuit (OMP) [5], compressive sampling matching pursuit (CoSaMP) [9], and iterative hard thresholding [12].

Both optimization and greedy approaches provide powerful stability guarantees in the presence of noise while remaining computationally efficient. Given noisy measurements of any signal $x \in \mathbb{R}^N$ so that $y = \Phi x + n$, if Φ possesses the RIP, then the signal estimate $\widehat{x}$ obtained by these algorithms has bounded error:

$$\|x - \widehat{x}\|_2 \leq C_1 \|x - x_K\|_2 + \frac{C_2}{\sqrt{K}} \|x - x_K\|_1 + C_3 \|n\|_2, \tag{7}$$

where x_K is the best K-sparse approximation to x as defined in (1) and C_1, C_2 are constants. Furthermore, with a simple modification, algorithms like CoSaMP and iterative hard thresholding can be used to reconstruct signals belonging to any structured sparsity model [7].

To summarize, at the core of CS lie three key concepts: a signal model exhibiting a particular type of low-dimensional geometry in high-dimensional space, a low-rank linear mapping that provides a stable embedding of the signal model into a lower dimensional space, and algorithms that perform stable, efficient inversion of this mapping.

III. SIGNAL MODELS FOR PULSE STREAMS

Our objective is to extend the CS theory and algorithms to pulse stream signals. The conventional sparse signal model Σ_K does not take into account the dependencies between the values and locations of the nonzeros in such signals. Indeed, these dependencies cannot be precisely captured by any structured sparsity model $\mathcal{M}_K$ that merely comprises a reduced subset of the

subspaces in Σ_K. This necessitates richer models that capture the *convolutional* structure present in the nonzero coefficients of pulse streams.

A. General model

Consider the following deterministic model for signals that can be modeled by the convolution of an S-sparse spike stream $x \in \mathbb{R}^N$ with an F-sparse impulse response $h \in \mathbb{R}^N$.

Definition 2: Let $\mathcal{M}_S \subset \mathbb{R}^N$ be a union of S-dimensional canonical subspaces, as defined in (2). Similarly, let $\mathcal{M}_F \subset \mathbb{R}^N$ be a union of F-dimensional canonical subspaces. Consider the set

$$\mathcal{M}_{S,F}^z := \{z \in \mathbb{R}^N : z = x * h, \text{ such that } x \in \mathcal{M}_S \text{ and } h \in \mathcal{M}_F\}, \tag{8}$$

where $*$ denotes the circular convolution operator. Then, $\mathcal{M}_{S,F}^z$ is called a *pulse stream model*.

We make two immediate observations:

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

$$z = x * h = h * x = Hx = Xh, \tag{9}$$

where H (respectively, X) is a square circulant matrix with its columns comprising circularly shifted versions of the vector h (respectively, x). Therefore, Definition 2 remains unchanged if the roles of x and h are reversed. We exploit this property during signal recovery from CS measurements in Section V.

2) Geometry: It is useful to adopt the following geometric point of view: for a fixed $h \in \mathcal{M}_F$, the set $\{h * x : x \in \mathcal{M}_S\}$ forms a finite union of S-dimensional subspaces, owing to the fact that it is generated by the action of h on L_S canonical subspaces. Denote this set by $h(\mathcal{M}_S)$. Then, the pulse stream model in (8) can be written as

$$\mathcal{M}_{S,F}^z = \bigcup_{h \in \mathcal{M}_F} h(\mathcal{M}_S).$$

Thus, our signal model can be interpreted as a specific instance of an *infinite union of subspaces*.[1] Note that (4) cannot be applied in this case since it only considers finite unions of subspaces.

---

[1]A general theory for sampling signals from infinite unions of subspaces has been introduced in [13].

However, let $K = SF$ denote the maximum sparsity of the signals in Definition 2. Then, it is clear that the set $\mathcal{M}_{S,F}^z$ is a very small subset of Σ_K, the set of all SF-sparse signals. We exploit this property while proving our main sampling results in Section IV.

Note that the exact definition of the convolution operator changes depending on the domain of the signals of interest. For one-dimensional (1D) time domain signals of length N, the square matrix H is formed by all N circular shifts of the vector h; for 2D images of size N pixels, H is formed by all 2D circular shifts of h, and so forth.

B. Special case: Disjoint pulses

The model proposed in Definition 2 is general and applicable even to signals in which successive pulses overlap with each other. In Section IV we develop a lower bound on the number of samples required to preserve the essential information contained in an arbitrary pulse stream. However, feasible recovery of such general pulse streams from CS measurements is rather difficult; we examine this in detail in Section V. Therefore, we will also consider a more restrictive model where the pulses are assumed to not overlap.

For concreteness, consider 1D time domain signals as specified by (9). Note that H and x need not be unique for a given z; any ordered pair $(\alpha H, x/\alpha)$ satisfies (9), and 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 $-\tau$. In particular, to eliminate the ambiguity due to circular shifts, we make the following two assumptions:

1) the impulse response h is *concentrated*, i.e., the F nonzero coefficients of h are contiguously located in its first F indices. Thus, the structured sparsity model $\mathcal{M}_F$ for the vector h consists of the lone subspace spanned by the first F canonical unit vectors.

2) the spikes are sufficiently separated in time. In particular, any two consecutive spikes in the vector x are separated at least by Δ locations, where $\Delta \geq F$. A structured sparsity model for time-domain signals with sufficiently separated nonzeros has been introduced in [14].

The notion of disjoint pulses can be immediately generalized to signals defined over domains of arbitrary dimension. Consider S-sparse spike streams x defined over a domain of dimension n. Suppose that at most one spike in x can occur in a hypercube in $\mathbb{R}^n$ with side Δ. This defines

9

a special structured sparsity model for the spike streams of interest; denote this model as $\mathcal{M}_S^\Delta$. Further, let the F nonzero coefficients in h be concentrated within a hypercube centered at the domain origin whose side length is no greater than Δ. Then, a deterministic model for sums of non-overlapping pulses of arbitrary dimension can be proposed as follows.

Definition 3: Let $\mathcal{M}_S^\Delta$ be the structured sparsity model for spike streams as defined above. Let $\mathcal{M}_F$ be the subspace of concentrated impulse responses of sparsity F. Define the set

$$\mathcal{M}(S, F, \Delta) = \{z \in \mathbb{R}^N : z = x * h, \text{such that } x \in \mathcal{M}_S^\Delta \text{ and } h \in \mathcal{M}_F\}. \tag{10}$$

Then, $\mathcal{M}(S, F, \Delta)$ is called the *disjoint pulse stream model*.

This model eliminates possible ambiguities that arise due to the shift-invariant nature of convolution; i.e., the locations of the nonzero spikes that generate a disjoint pulse stream are *uniquely* defined. This property proves to be essential in developing and analyzing a feasible method for signal recovery (Section V). See Figure 1(a) for an example stream of disjoint pulses in 1D.

IV. Sampling Theorems for Pulse Streams

Pulse streams can be modeled as an infinite union of low-dimensional subspaces. The next ingredient in the development of a CS framework for such signals is a bound on the number of linear samples required to preserve the essential information of this signal set.

A. General pulse streams

We derive a sampling theorem for signals belonging to the model $\mathcal{M}_{S,F}^z$ proposed in Definition 2. Again, note that here the pulses are allowed to overlap with one another. Suppose that $K = SF$. As mentioned above, $\mathcal{M}_{S,F}^z$ is a subset of the set of all K-sparse signals Σ_K. On the other hand, only a small fraction of all K-sparse signals can be written as the convolution of an S-sparse spike stream with an F-sparse impulse response. Thus, intuition suggests that we should be able to compressively sample signals from this set using fewer random linear measurements than that required for the set of all K-sparse signals. The following theorem makes this precise.

Theorem 1: Suppose $\mathcal{M}_{S,F}^z$ is the pulse stream model from Definition 2. Let $t > 0$. Choose an $M \times N$ i.i.d. Gaussian matrix Φ with

$$M \geq \mathcal{O}\left(\frac{1}{\delta}\left((S+F)\ln\left(\frac{1}{\delta}\right) + \log(L_S L_F) + t\right)\right). \tag{11}$$

10

Then, Φ satisfies the following property with probability at least $1 - e^{-t}$: for every pair $z_1, z_2 \in \mathcal{M}_{S,F}^z$,

$$(1 - \delta)\|z_1 - z_2\|_2^2 \le \|\Phi z_1 - \Phi z_2\|_2^2 \le (1 + \delta)\|z_1 - z_2\|_2^2. \tag{12}$$

The proof of this theorem is presented in Appendix A. An important consequence of the theorem is that, by definition, $\mathcal{M}_S$ is a subset of the set of all S-dimensional canonical subspaces. In particular,

$$L_S \le \binom{N}{S} \approx \left(\frac{eN}{S}\right)^S. \tag{13}$$

Similarly, $L_F \le \left(\frac{eN}{F}\right)^F$. Therefore, the logarithmic term in the expression for M in (11) scales as

$$\log(L_S L_F) \le S + S \log(N/S) + F + F \log(N/F) \le 2(S + F) \log N \tag{14}$$

Thus, (11) indicates that the number of measurements M required for sampling signals in $\mathcal{M}_{S,F}^z$ is proportional to $(S + F)$. Therefore, M is *sublinear* in the total sparsity of the signals $K = SF$. In contrast, conventional structured sparsity models would require at least $2K = 2SF$ linear measurements to ensure a stable embedding of the signal set [11]. In addition, the number of degrees of freedom of each signal can be considered to be $\mathcal{O}(S + F)$, corresponding to the positions and locations of the coefficients of the sparse signal and impulse response. Therefore, the bound in Theorem 1 is essentially optimal for the signal class $\mathcal{M}_{S,F}^z$.

B. Special case: Disjoint pulse streams

Theorem 1 is valid for signals belonging to the general model $\mathcal{M}_{S,F}^z$. In the case of disjoint pulse streams, we can derive a more stringent lower bound. By definition, the F nonzero coefficients of h are concentrated in a hypercube around the domain origin. Therefore, h lies in a lone subspace spanned by F basis vectors of $\mathbb{R}^N$, and hence $L_F = 1$. Further, a simple modification of Theorem 1 of [14] states that the number of subspaces in the structured sparsity model $\mathcal{M}_S^\Delta$ is given by

$$L_S = \binom{N - S\Delta + S - 1}{S - 1}. \tag{15}$$

Thus, for the disjoint pulse stream model $\mathcal{M}(S, F, \Delta)$, we obtain the following easy corollary to Theorem 1.

Corollary 1: If $t > 0$ and

$$M \geq \mathcal{O}\left(\frac{1}{\delta}\left((S + F)\ln\left(\frac{1}{\delta}\right) + S\log(N/S - \Delta) + t\right)\right), \tag{16}$$

then an $M \times N$ i.i.d. Gaussian matrix Φ will satisfy (12) with probability at least $1 - e^{-t}$ for any pair of signals z_1, z_2 belonging to the $\mathcal{M}(S, F, \Delta)$ model.

Note that the parameter Δ can be at most N/S, since S spikes must be packed into N coefficient locations with at least Δ locations separating any pair of spikes. A higher value of Δ implies that the model $\mathcal{M}_S^\Delta$ admits a smaller number of configurations; thus, (16) implies that fewer measurements are needed to sample pulse streams in which the pulses are widely separated.

V. RECOVERY OF PULSE STREAMS

The final ingredient in our extended CS framework for pulse streams consists of new algorithms for the stable recovery of the signals of interest from compressive measurements. This problem can be stated as follows. Suppose $z \in \mathcal{M}_{S,F}^z$. If w are given the noisy measurements

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

then we aim to reconstruct z from y. The main challenge stems from the fact that *both* x (respectively, X) and h (respectively, H) are unknown and have to be simultaneously inferred.

This problem is similar to performing sparse approximation with *incomplete* knowledge of the dictionary in which the target vector (either x or h) is sparse. This problem has received some interest in the literature [15, 16]; the common approach has been to first assume that a training set of vectors $\{x_i\}$ exists for a fixed impulse response h, and then to infer the coefficients of h using a sparse learning algorithm (such as basis pursuit [6]), and finally to solve for the coefficients $\{x_i\}$. In the absence of training data, we must infer both the spike locations and the impulse response coefficients. Therefore, our task is also similar to *blind deconvolution* [17]; the main differences are that we are only given access to the random linear measurements y as opposed to the Nyquist rate samples z and that our primary aim is to reconstruct z as faithfully as possible as opposed to merely reconstructing x.

Our general approach will be to fix an estimate of h, obtain the "best possible" estimate of x, update our estimate of h, and iterate. This is commonly known as *alternating minimization* (AM) and has been shown to be suitable for blind deconvolution settings [18]. As demonstrated below in the proof of Theorem 2, we require that the best possible estimate of the spike stream x and the impulse response h at each iteration are unique. For this reason, we will assume that our target signal z belongs to the disjoint pulse stream model $\mathcal{M}(S, F, \Delta)$.

A. Alternating minimization with exhaustive search

Consider $z \in \mathcal{M}(S, F, \Delta)$, so that $z = x * h$. This implies that the spikes in x are separated by a minimum separation distance Δ and that the impulse response h is concentrated. Suppose first that we are given noiseless CS measurements $y = \Phi z$. We fix a candidate support configuration σ for the spike stream (so that σ contains S nonzeros.) Then, we form the circulant matrix $\widehat{H}$ from all possible shifts of the current estimate of the impulse response $\widehat{h}$ (denote this operation as $\widehat{H} = \mathbb{C}(\widehat{h})$). Further, we calculate the dictionary $\Phi\widehat{H}$ for the spike stream x and select the submatrix formed by the columns indexed by the assumed spike locations σ (denote this submatrix as $(\Phi\widehat{H})_\sigma$). This transforms our problem into an overdetermined system, which can be solved using least-squares. In summary, we use a simple matrix pseudoinverse to obtain the estimate

$$\widehat{x} = (\Phi\widehat{H})_\sigma^\dagger y.$$

This provides an estimate of the spike coefficients $\widehat{x}$ for the assumed support configuration σ. We now exploit the commutativity of the convolution operator $*$. We form the circulant matrix $\widehat{X}$, form the dictionary $\Phi\widehat{X}$ for the impulse response and select the submatrix $(\Phi\widehat{X})_f$ formed by its first F columns. Then, we solve a least-squares problem to obtain an estimate $\widehat{h}$ for the impulse response coefficients:

$$\widehat{h} = (\Phi\widehat{X})_f^\dagger y.$$

Finally, we form our signal estimate $\widehat{z} = \widehat{x} * \widehat{h}$. The above two-step process is iterated until a suitable halting criterion (e.g., convergence in norm for the estimated signal $\widehat{z}$). This process is akin to the Richardson-Lucy algorithm for blind deconvolution [19].

13

Algorithm 1 Alternating minimization with exhaustive search

Inputs: Sampling matrix Φ, measurements $y = \Phi x$, model parameters Δ, S, F, threshold ϵ

Output: $\widehat{z} \in \mathcal{M}(S, F, \Delta)$ such that $y - \Phi\widehat{z}$ is small

$\widehat{x} = 0$, $\widehat{h} = (\mathbf{1}_F^T, 0, \ldots, 0)/\sqrt{F}$; $i = 0$ {initialize}

for $\sigma \in \mathcal{M}_S^\Delta$ **do**

while halting criterion false **do**

    1. $\widehat{H} = \mathbb{C}(\widehat{h})$, $\Phi_h = (\Phi\widehat{H})_\sigma$ {form dictionary for spike stream}

    2. $\widehat{x} \leftarrow \Phi_h^\dagger y$ {update spike stream estimate}

    3. $\widehat{X} = \mathbb{C}(\widehat{x})$, $\Phi_x = (\Phi\widehat{X})_f$ {form dictionary for impulse response}

    4. $\widehat{h} \leftarrow \Phi_x^\dagger y$ {update impulse response estimate}

    5. $\widehat{z} \leftarrow \widehat{x} * \widehat{h}$ {form signal estimate}

end while

if $\|y - \Phi\widehat{z}\|_2 < \epsilon$ {check for energy in residual}

   return $\widehat{z}$

end if

end for

The overall reconstruction problem can be solved by repeating this process for every support configuration σ belonging to the structured sparsity model $\mathcal{M}_S^\Delta$ and picking the solution with the smallest norm of the residual $r = y - \Phi\widehat{z}$. The procedure is detailed in pseudocode form in Algorithm 1. Thus, Algorithm 1 performs alternating minimization for a given estimate for the support of the underlying spike stream x and exhaustively searches for the best possible support. Under certain conditions on the sampling matrix Φ, we can study the convergence of Algorithm 1 to the correct answer z, as encapsulated in the following theorem.

Theorem 2: Let $z \in \mathcal{M}(S, F, \Delta)$ and suppose that Φ satisfies (12) with constant δ for signals belonging to $\mathcal{M}(S, F, \Delta)$. Suppose we observe $y = \Phi z$ and apply Algorithm 1 to reconstruct the signal. Let $\widehat{z}_i$ be an intermediate estimate of z at iteration i of Algorithm 1. Then:

1) The norm of the residual $\|y - \Phi\widehat{z}_i\|_2$ monotonically decreases with iteration count i.

14

2) If at any iteration i

$$\|y - \Phi\widehat{z}_i\|_2 \leq \epsilon,$$

then we are guaranteed that

$$\|z - \widehat{z}_i\|_2 \leq c\epsilon,$$

where c depends only on δ.

The proof of this theorem is provided in Appendix B. The first part of the theorem implies that for any given support configuration σ, Steps 1 through 4 in Algorithm 1 are guaranteed to converge to a generalized fixed point [20]. The second part of the theorem provides a condition on the detection of the true support configuration σ in the following weak sense: if the energy of the residual of the signal estimate is small, then the signal has been accurately reconstructed.

Algorithm 1 involves iteratively solving a combinatorial number of estimation problems. This becomes infeasible for even moderate values of N. A simpler algorithm is described in Section V-C.

B. Model mismatch

In practical situations, we would expect to have minor variations in the shapes of the S pulses in the signal z. In this case, z can no longer be expressed as Hx where H is a circulant matrix. Let $\{h_1, h_2, \ldots, h_S\}$ be length-F vectors corresponding to each of the S pulses in the signal, and let the length-S spike stream $\widetilde{x} = (\alpha_1, \alpha_2, \ldots, \alpha_S)$. Further, let $\mathbb{S}_i$ be the circular shift operator that maps the i^{th} pulse shape h_i into its corresponding vector in $\mathbb{R}^N$. Then, we have

$$z = \sum_{i=1}^{S} \alpha_i \mathbb{S}_i(h_i). \tag{17}$$

However, Algorithm 1 attempts to estimate a single pulse shape $\widehat{h}$. Thus, following the notation in Section V-A, the matrix $\widehat{H}_\sigma$ is an $N \times S$ matrix. Assuming that the spikes in x are separated by at least Δ locations, the matrix $\widehat{H}_\sigma$ is *quasi-Toeplitz* [21], i.e., the columns of $\widehat{H}_\sigma$ are circular shifts of one another with no more than one nonzero entry in every row. Further, assume that the measurement matrix Φ equals the *identity*, i.e., we are given Nyquist-rate samples of z. Then the matrices Φ_h and Φ_x in Step 2 of Algorithm 1 are also quasi-Toeplitz. Thus, given an estimate

of the pulse shape $\widehat{h}_0$, we can derive a closed-form expression for the next impulse response estimate. Additionally, we can obtain an intermediate estimate for the spike stream $\widetilde{x}$. Suppose that the innermost loop of Algorithm 1 converges to a fixed point estimate $\widehat{h}$. We dub $\widetilde{h}$ the *anchor pulse* for the set of pulse shapes $\{h_1, h_2, \ldots, h_S\}$. The following theorem provides an expression relating the anchor pulse to the component pulse shapes.

Theorem 3: Consider z as defined in (17). Let $\widetilde{h}$ be the anchor pulse for the set of pulse shapes $\{h_1, h_2, \ldots, h_S\}$. Define $c_i = \langle h_i, \widetilde{h} \rangle$ for $i = 1, \ldots, S$. Then, we have that

$$\widetilde{h} = \frac{\sum_{i=1}^{S} c_i \alpha_i^2 h_i}{\sum_{i=1}^{S} c_i^2 \alpha_i^2}. \tag{18}$$

The proof of this theorem is provided in Appendix C. Equation (18) implies that the anchor pulse $\widetilde{h}$ is a weighted linear combination of the component pulses h_i, $i = 1, \ldots, S$, with the weights defined by the corresponding spike coefficients α_i and the inner products c_i. The anchor pulse remains unchanged if the spike coefficient vector $\widetilde{x}$ is multiplied by any constant C. Therefore, the anchor pulse can be viewed as a *scale-invariant average* of the component pulse shapes.

Theorem 3 applies to Nyquist-rate samples of the signal z. In the case of low-rate CS measurements $y = \Phi z$, the convergence analysis of Algorithm 1 for the general case of S different pulse shapes becomes more delicate. If Φ possesses the RIP only for $z \in \mathcal{M}(S, F, \Delta)$, then it could be that two different pulse streams z_1, z_2 (each with varying shapes across pulses) are mapped by Φ to the same vector in $\mathbb{R}^M$, i.e., $\Phi z_1 = \Phi z_2$; thus, the unique mapping argument employed in the proof of Theorem 2 cannot be applied in this case. One way to analyze this case is to recognize that by allowing arbitrary pulse shapes $\{h_1, h_2, \ldots, h_S\}$, our space of signals of interest is equivalent to a special structured sparsity model that consists of all K-sparse signals whose non-zeros are arranged in S blocks of size F and the starting locations of consecutive blocks are separated by at least Δ locations. As discussed in Section II, stable CS reconstruction for signals from this model requires at least $M = 2SF = 2K$ measurements; thus, Algorithm 1 converges in the general case given that M is proportional to K. Thus, in the case of arbitrary pulse shapes, the number of measurements required by Algorithm 1 is on the same order as the number of measurements required for conventional structured sparsity-based CS recovery.

16

C. Iterative support estimation

The runtime of Algorithm 1 is exponential in N. Fortunately there is a simpler means to the same end. Instead of cycling through every possible support configuration σ_i for the spike stream x, we instead retain an *estimate* of the support configuration, based on the current estimates of the spike stream $\widehat{x}$ and impulse response $\widehat{h}$, and update this estimate with each iteration. In order to ensure that the support estimate belongs to $\mathcal{M}_S^\Delta$, we leverage a special CS recovery algorithm for signals belonging to $\mathcal{M}_S^\Delta$ that is based on CoSaMP [9]. We provide an outline of the algorithm here for completeness; see [14] for details.

At each iteration, given an estimate of the spike coefficients x, we need to solve for the best $\mathcal{M}_S^\Delta$-approximation to x. Let $x = (x_1, x_2, \ldots, x_N)^T$. Given any binary vector $s = (s_1, s_2, \ldots, s_N)^T$ of length N, let

$$x_{|s} := (s_1 x_1, s_2 x_2, \ldots, s_N x_N),$$

so that $x_{|s}$ is the portion of the signal x lying within the support s. Our goal is to solve for the choice of support s so that $x_{|s}$ belongs to $\mathcal{M}_S^\Delta$ and $\|x - x_{|s}\|_2$ is minimized. The following constraints on the support vector s follow from the definition of $\mathcal{M}_S^\Delta$:

$$s_1 + s_2 + \ldots + s_N \leq S, \tag{19}$$

$$s_j + s_{j+1} \ldots + s_{j+\Delta-1} \leq 1, \text{ for } j = 1, \ldots, N, \tag{20}$$

where the subscripts are computed modulo N. The first inequality (19) specifies that the solution contains at most S nonzeros; the other N inequalities (20) specify that there is at most one spike within any block of Δ consecutive coefficients in the solution.

It can be shown that minimizing $\|x - x_{|s}\|_2$ is equivalent to maximizing $c^T s$ where $c = (x_1^2, x_2^2, \ldots, x_N^2)$, i.e., maximizing the portion of the energy of x that lies within s. Define $W \in \mathbb{R}^{(N+1)\times N}$ as a binary indicator matrix that captures the left hand side of the inequality constraints (19) and (20). Next, define $u \in \mathbb{R}^{N+1} = (S, 1, 1, \ldots, 1)$; this represents the right hand side of the constraints (19) and (20). Thus, we can represent (19) and (20) by the following binary integer program:

$$s^* = \arg\min_{s \in \{0,1\}} c^T s, \text{ subject to } Ws \leq u.$$

17

Next, we relax the integer constraints on s to obtain a computationally tractable linear program. Denote this linear program by $\mathbb{D}(\cdot)$. In [14], it is shown that the solutions to the integer program and its relaxed version are identical. Thus, we have a computationally feasible method to obtain an estimate of the support of the best $\mathcal{M}_S^{\triangle}$-approximation to x.

Once an updated support estimate has been obtained, we repeat Steps 2, 3 and 4 in Algorithm 1 to solve for the spike stream x and impulse h. This process is iterated until a suitable halting criterion (e.g., convergence in norm for the estimated pulse stream $\widehat{z}$.) The overall algorithm can be viewed as an iterative sparse approximation procedure for the $\mathcal{M}_S^{\triangle}$ model that continually updates its estimate of the sparsifying dictionary. The procedure is detailed in pseudocode form in Algorithm 2.

D. Stability and convergence

Like many algorithms for blind deconvolution, the analysis of Algorithm 2 is not straightforward. The dictionaries $\Phi\widehat{X}$ and $\Phi\widehat{H}$ are only approximately known at any intermediate iteration, and hence the proof techniques employed for the analysis for CoSaMP do not apply. In principle, given access to a sufficient number of measurements, we may expect similar convergence behavior for Algorithm 2 as Algorithm 1. Empirically, Algorithm 2 can be shown to be stable to small amounts of noise in the signal as well as in the CS measurements and to minor variations in the pulse shape. We demonstrate this with the help of various numerical experiments in Section VI.

E. Computational complexity

The primary runtime cost of Algorithm 2 is incurred in solving the linear program $\mathbb{D}(\cdot)$. For a length-N signal, the computational complexity of solving a linear program is known to be $\mathcal{O}\left(N^{3.5}\right)$. The total computational cost also scales linearly in the number of iterations T of the outer loop. An overall upper bound runtime of the algorithm is thus $\mathcal{O}\left(N^{3.5}T\right)$.

VI. NUMERICAL EXPERIMENTS

We now present the results of a number of numerical experiments that validate the utility of our proposed theory and methods. All experiments reported in this section have been performed using Algorithm 2 for the recovery of disjoint pulse streams.

18

Algorithm 2 Iterative support estimation

Inputs: Sampling matrix Φ, measurements $y = \Phi z + n$, model parameters Δ, S, F.

Output: $\mathcal{M}(S, F, \Delta)$-sparse approximation $\widehat{z}$ to true signal z

Initialize $\widehat{x} = 0$, $\widehat{h} = (\mathbf{1}_F^T, 0, \ldots, 0)$, $i = 0$

while halting criterion false **do**

    1. $i \leftarrow i + 1$

    2. $\widehat{z} \leftarrow \widehat{x} * \widehat{h}$ {current pulse stream estimate}

    {estimate spike locations and amplitudes}

    3. $\widehat{H} = \mathbb{C}(\widehat{h}), \Phi_h = \Phi\widehat{H}$ {form dictionary for spike stream}

    4. $e \leftarrow \Phi_h^T(y - \Phi_h\widehat{x})$ {residual}

    5. $\omega \leftarrow \sigma(\mathbb{D}(e))$ {obtain model-approximation of residual}

    6. $\sigma \leftarrow \omega \cup \sigma(\widehat{x}_{i-1})$ {merge supports}

    7. $x|_\sigma \leftarrow (\Phi_h)_\sigma^\dagger y, \, x|_{\sigma^C} = 0$ {update spike stream estimate}

    8. $\widehat{x} \leftarrow \mathbb{D}(x)$ {prune spike stream estimate}

    {estimate impulse response}

    9. $\widehat{X} = \mathbb{C}(\widehat{x}), \Phi_x = (\Phi\widehat{X})_f$ {form dictionary for impulse response}

    10. $\widehat{h} \leftarrow \Phi_x^\dagger y$ {update impulse response estimate}

end while

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

A. Synthetic 1D pulse streams

Figure 1 demonstrates the considerable advantages that Algorithm 2 can offer 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 randomly chosen impulse response of length $F = 11$. The overall sparsity of the signal $K = SF = 88$; thus, standard sparsity-based CS algorithms would require at least $2K = 176$ measurements. Our approach (Algorithm 2) returns an accurate estimate of both the spike stream as well as the impulse response using merely $M = 90$ measurements.

Figure 2 displays the averaged results of a Monte Carlo simulation of Algorithm 2 over
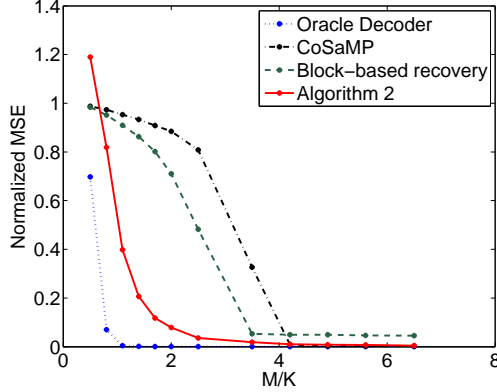
Fig. 2. *Normalized reconstruction MSE vs. M/K for different reconstruction algorithms averaged over 200 sample trials. Signal parameters: $N = 1024, S = 8, F = 11$. Algorithm 2 outperforms standard and structured sparsity-based methods, particularly when M/K is small.*

200 trials. Each trial was conducted by generating a sample signal belonging to $\mathcal{M}(S, F, \Delta)$, computing M linear random Gaussian measurements, reconstructing with different algorithms, and recording the magnitude of the recovery error for different values of M/K. Each test signal was generated by convolving a signal x belonging to $\mathcal{M}_S^\Delta$ with a minimum phase impulse response h of order F, with coefficients of both x and h independently drawn from a Gaussian distribution with zero mean and unit variance. It is clear from the figure that Algorithm 2 outperforms both conventional CS recovery (CoSaMP [9]) with target sparsity $K = SF$ as well as block-based reconstruction [7] with knowledge of the size and number of blocks (respectively F and S). In fact, our algorithm performs nearly as well as the "oracle decoder" that possesses perfect prior knowledge of the impulse response coefficients and aims to solve only for the spike stream. In practice, the algorithm converges in fewer than 20 iterations; empirically, the number of iterations decreases with an increasing number of compressive measurements.

We show that Algorithm 2 is stable to small amounts of noise in the signal and the measurements. In Figure 3, we generate a length $N = 1024$ signal from a disjoint pulse stream model with $S = 7$ and $F = 11$; add a small amount of Gaussian noise (SNR = 14.83dB) to all its components, compute $M = 150$ noisy linear measurements, and reconstruct using Algorithm 2. The reconstructed signal is clearly a good approximation of the original signal.
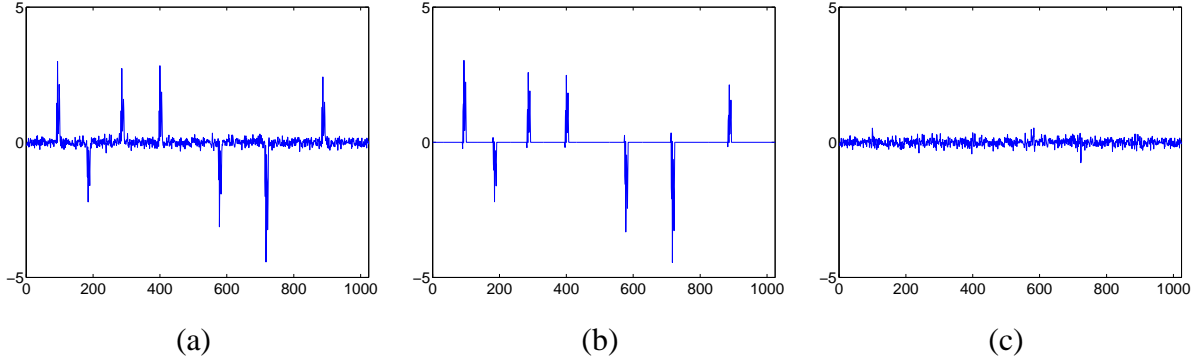
20

Fig. 3. *(a) Original noisy signal (SNR = 14.83dB). (b) Recovered estimate from $M = 150$ random measurements using Algorithm 2. (c) Residual error vector. Recovery SNR = 4.99dB.*
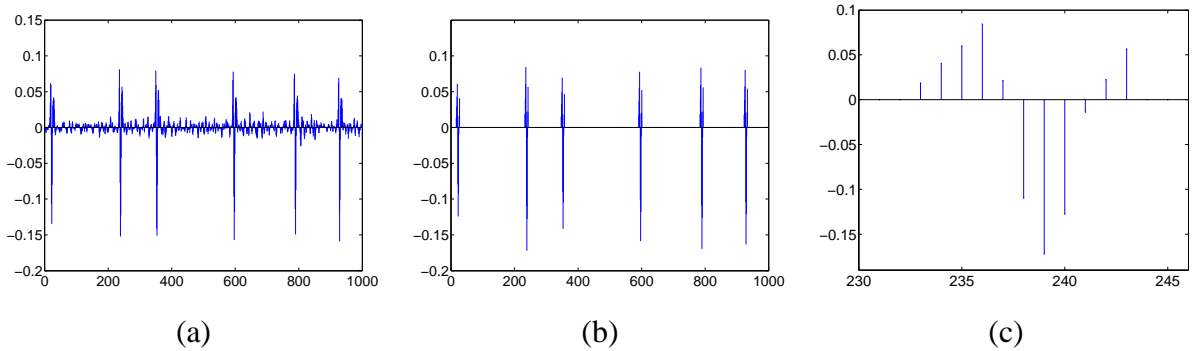

Fig. 4. *CS recovery of a real-world neuronal signal. (a) Original recording. (b) Recovered signal using $M = 150$ random measurements. (c) Zoom-in on estimated anchor pulse shape ($F = 11$).*

B. Neuronal signals

We test Algorithm 2 on a real-world neuronal recording. Figure 4(a) shows the temporal electrochemical spiking potential of a single neuron. The shape of the pulses is characteristic of the neuron and should ideally be constant across different pulses. 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 to the original signal (Figure 4(b)) as well as an estimate of the anchor pulse shape (Figure 4(c)).
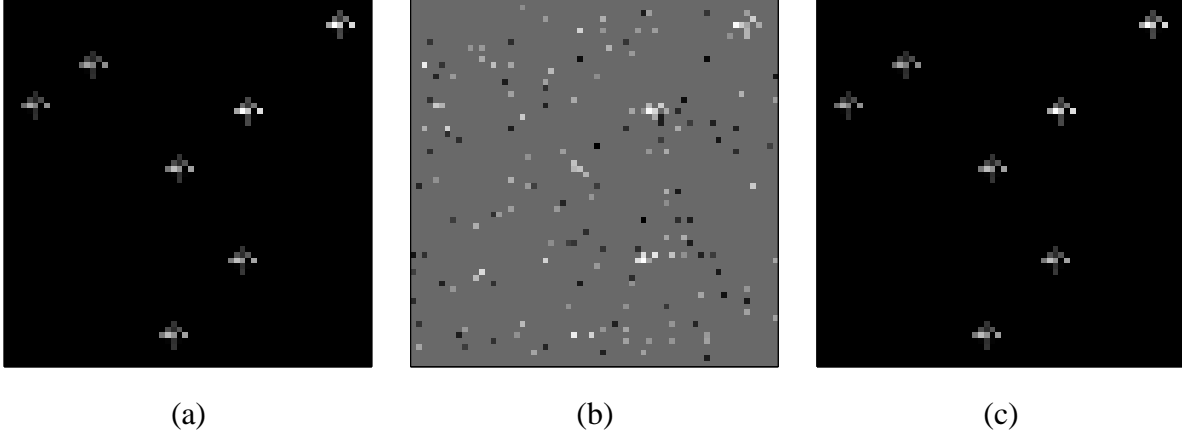
<center>(a)   (b)   (c)</center>

Fig. 5. *Example CS recovery of a sum of 2D pulses. (a) Synthetic test image: $N = 4096$, $S = 7$, $F = 25$. Images are recovered from $M = 290$ random Gaussian measurements using (b) CoSaMP (MSE = 16.95), and (c) Algorithm 2 (MSE = 0.07).*

C. Synthetic 2D pulse streams

Theorem 1 and Algorithm 2 can easily be extended to higher dimensional signals. For instance, suppose that the signals of interest are 2D images that can be modeled by a sparse sum of disjoint 2D pulses. We test Algorithm 2 on a synthetic image (Figure 5(a)) of size $N = 64 \times 64 = 4096$ that comprises $S = 7$ spikes blurred by an unknown 2D impulse response of size $F = 5 \times 5 = 25$, so that the overall sparsity $K = SF = 175$. We acquire merely $M = 290$ random Gaussian measurements (approximately 7% the size of the image N) and reconstruct the image using CoSaMP as well as Algorithm 2. We assume that both algorithms possess an oracular knowledge of the number of spikes S as well as the size of the impulse response F. Figure 5 displays the results of the reconstruction procedures using CoSaMP and Algorithm 2. It is evident both perceptually and in terms of the MSE values of the reconstructed images that our proposed approach is superior to traditional CS recovery.

D. Astronomical images

Finally, we test Algorithm 2 on a real astronomical image. Our test image is an $N = 64 \times 64$ region of a high-resolution image of V838 Monocerotis (a nova-like variable star) captured by the Hubble Space Telescope [22] (highlighted by the green square in Figure 6(a)). Note the
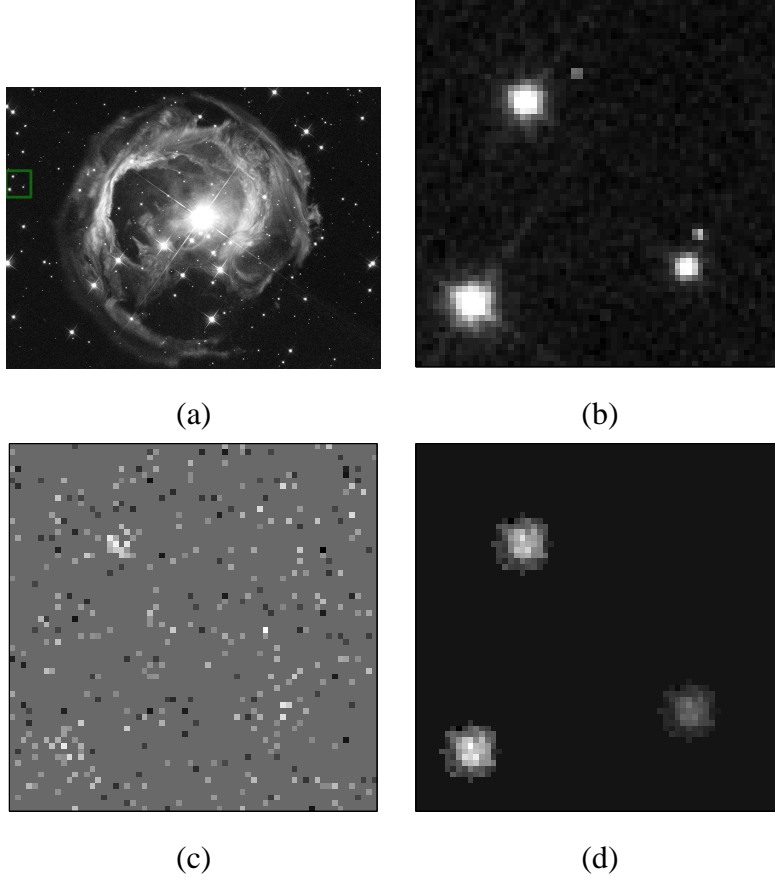
<center>22</center>

Fig. 6. *(a) Black-and-white image of V838 Monocerotis, a nova-like star, captured by the Hubble Space Telescope on February 8, 2004 [22]. (b) Test image is a zoomed in-version of the region at left highlighted in green (resolution $N = 64 \times 64 = 4096$). Reconstruction of test image is performed from $M = 330$ random Gaussian measurements using (c) CoSaMP and (d) Algorithm 2.*

significant variations in the shapes of the three large pulses in the test image (Figure 6(b)). We measure this image using $M = 330$ random measurements and reconstruct using both CoSaMP and Algorithm 2. For our reconstruction methods, we assumed an oracular knowledge of the signal parameters; we use $S = 3, F = 120, K = 360$ and $\Delta = 20$. As indicated by Figure 6, conventional CS does not provide useful results with this reduced set of measurements. In contrast, Algorithm 2 gives us excellent estimates for the locations of the pulses. Further, our algorithm also provides a circular impulse response estimate that can be viewed as the anchor pulse of the three original pulses.

VII. DISCUSSION AND CONCLUSIONS

In this paper, we have introduced and analyzed a new framework for the compressive sampling of pulse streams. Our signals of interest are modeled as lying in an infinite union of subspaces with exhibits a particular geometric structure. This structure enables us to quantitatively estimate the number of random linear measurements needed to sample such signals. We have proposed two methods for signal recovery. Our first method (Algorithm 1) is relatively easy to analyze, but suffers from high combinatorial complexity. Our second method (Algorithm 2) is a feasible, if suboptimal, algorithm and formed the basis for our numerical experiments. While our framework is applicable to signals defined over domains of arbitrary dimension, we have illustrated its benefits in the context of 1D time signals and 2D images.

There are several avenues for future work. We have discussed sparse signals and images as represented in the identity basis; our method can be extended to wavelet-sparse and Fourier-sparse signals. While our results are promising, we still do not possess a complete characterization of the convergence properties of Algorithm 2 as well as its sensitivity to factors such as noise and model mismatch under random projections. Additionally, it is unclear how to deal with situations where the pulses in the signal of interest are allowed to overlap. To the best of our knowledge, the issue of robust recovery of signals convolved with an unknown arbitrary impulse response is an open question even for the case of Nyquist-rate samples. We defer these challenging open questions to future research.

The framework developed in this paper can be related to various existing concepts in the literature such as finite rate of innovation [23], blind compressed sensing [24], simultaneous sparse approximation and dictionary learning [25], and the classical signal processing problem of blind deconvolution [17]. Compressive sensing of time-domain pulse streams has been studied by Naini et al. [16]. However, in their setting the impulse response is assumed to be known, and hence the CS measurement system can be viewed as a modification of random Fourier subsampling.

Our framework is related to recent results on compressed blind deconvolution by Saligrama and Zhao [26]. As opposed to pulse streams, their signals of interest consist of sparse signals driven through an all-pole auto-regressive (AR) linear system. They propose an optimization-based algorithm for recovery of the signal and impulse response from CS measurements. How-

ever, their measurement system is tailored to impulse responses corresponding to linear AR models; in contrast, our approach can handle arbitrary impulse responses. Further, our main theoretical result indicates that the number of measurements to compressively sample a pulse stream is linear only in the number of degrees of freedom of the signal and thus answers an open question (Remark 3.1) posed by the authors in the affirmative.

Finally, the main approach in this paper can be related to recent work by Asif et al. [27, 28], who propose channel coding methods to combat the effect of unknown multipath effects in a communication channel that can be described by a sparse impulse response. Their channel code consists of a random matrix $\Phi \in \mathbb{R}^{M \times N}$ with $M > N$, so that the linear mapping $y = \Phi x$ is now not undercomplete, but overcomplete. Thus, their observations consist of an unknown sparse channel response h convolved with the transmitted signal y and their objective is to reconstruct the original signal x. The main aspects of our theoretical analysis could conceivably be modified to quantify system performance in this setting.

APPENDIX A

We prove Theorem 1. By Definition 2, the model $\mathcal{M}_{S,F}^z$ is generated via the convolution operation by the structured sparsity models $\mathcal{M}_S$ and $\mathcal{M}_F$. Recall that both structured sparsity models are themselves defined in terms of canonical subspaces of $\mathbb{R}^N$ and their convolution results in a low-dimensional geometrical structure that is best described by an infinite union of subspaces. Thus, if $x \in \mathcal{M}_S$ lies in a particular subspace Ω and $h \in \mathcal{M}_F$ lies in a particular subspace Λ, then every signal $z \in \mathcal{M}_{S,F}^z$ can be identified with at least one infinite union of subspaces $U_{\Omega,\Lambda}$. The overall approach closely follows the technique used in [29] and can be summarized as follows: we first construct a net of points Q in $\mathbb{R}^N$ such that

$$\min_{q \in Q} \|z - q\| < \delta,$$

for all $z \in U_{\Omega,\Lambda}$ with $\|z\| = 1$ and some constant δ. We then apply the Johnson-Lindenstrauss Lemma [30] for stable embedding of point clouds to this finite net of points Q, and extend the stable embedding to all possible signals $z \in U_{\Omega,\Lambda}$. Finally, we obtain our main result through a union bound over all possible choices of subspaces Ω and Λ.

First, we construct the net of points Q. Recall that $z \in \mathcal{M}_{S,F}^z$ if $z = h * x$, where $h \in \Lambda$ and $x \in \Omega$. Following [29], we will only consider the case where $\|x\|, \|h\| \leq 1$; it is easy to check that the corresponding circulant matrices X, H also satisfy $\|X\|, \|H\| \leq 1$. The extension to all x, h follows from the subspace property of Ω and Λ. Consider an S-dimensional subspace $\Omega \subset \mathcal{M}_S$. By Lemma 5.1 of [29], we can find a finite set of points $Q_{\Omega,x} \subset \Omega$ with cardinality $|Q_{\Omega,x}| \leq (3/\delta')^S$ such that

$$\min_{q_x \in Q_{\Omega,x}} \|x - q_x\| \leq \delta', \quad \forall \|x\| \leq 1, x \in \Omega,$$

i.e., $Q_{\Omega,x}$ is a δ'-net for Ω. Similarly, we can find a set of points $Q_{\Lambda,h} \subset \Lambda$ with cardinality $|Q_{\Lambda,h}| \leq (3/\delta')^F$ that forms a δ'-net for Λ. Consider the product set $Q_{\Omega,\Lambda} \in \mathcal{M}_{S,F}^z$ defined by:

$$Q_{\Omega,\Lambda} = \{z_{i,j} = x_i * h_j \mid x_i \in Q_{\Omega,x}, h_j \in Q_{\Lambda,h}\}.$$

Observe that the cardinality of this set $Q_{\Omega,\Lambda} = (3/\delta')^S(3/\delta')^F$. We claim that the finite set Q forms a $2\delta'$-net for the infinite union of subspaces $U_{\Omega,\Lambda}$. To see this, consider an arbitrary $z = x * h = Hx \in U_{\Omega,\Lambda}$, and $z_{i,j} \in Q_{\Omega,\Lambda}$. Then, we have that

$$
\begin{aligned}
\|z - z_{i,j}\| &= \|Hx - H_ix_j\| = \|Hx - H_ix + H_ix - H_ix_j\| \\
&\leq \|Hx - H_ix\| + \|H_ix - H_ix_j\| \\
&\leq \|X\|\|h - h_i\| + \|H_i\|\|x - x_j\| \\
&\leq \delta'(\|X\| + \|H_i\|) \leq 2\delta',
\end{aligned}
$$

with appropriate choices for i and j.

Now, suppose that $\delta = 4\delta'$. By the Johnson-Lindenstrauss Lemma, if $\Phi \in \mathbb{R}^{M \times N}$ with the elements of Φ drawn from a random Gaussian distribution, then for all pairs of vectors $z_1, z_2 \in Q_{\Omega,\Lambda}$, (12) will hold with failure probability

$$p_{\Omega,\Lambda} = 2\left(\frac{3}{\delta}\right)^S \left(\frac{3}{\delta}\right)^F e^{-c_0(\delta/2)M}.$$

By virtue of the fact that $Q_{\Omega,\Lambda}$ forms a net for $Q_{\Omega,\Lambda}$, (12) will simultaneously hold for all pairs of vectors $z_1, z_2 \in Q_{\Omega,\Lambda}$ with the same failure probability $p_{\Omega,\Lambda}$.

This expression for $p_{\Omega,\Lambda}$ is for a *fixed* pair of subspaces $(\Omega, \Lambda) \in \mathcal{M}_S \times \mathcal{M}_F$. There are $L_S \times L_F$ such pairs of subspaces. Applying a simple union bound over all possible pairs, we obtain the overall failure probability as

$$p \ \leq \ \sum_{(\Omega,\Lambda)} p_{\Omega,\Lambda} \leq L_S L_F \left(\frac{3}{\delta}\right)^{S+F} e^{-c_0(\delta/2)M}.$$

Rearranging terms, we have that for a suitably chosen constant C (that depends on c_0) and for any $t > 0$, if

$$M \geq C \left(\log(L_S L_F) + (S + F)\log\left(\frac{3}{\delta}\right) + t\right),$$

the failure probability for the sampling bound is smaller than e^{-t}. The main result follows.

APPENDIX B

We prove Theorem 2. Let $\widehat{z}_i = \widehat{x}_i * \widehat{h}_i$ be any intermediate estimate of Algorithm 1. Let $\widehat{H}_i = \mathbb{C}(\widehat{h}_i)$. Suppose that our candidate configuration for the support of x is given by the sparsity pattern σ belonging to the structured sparsity model $\mathcal{M}_S^\Delta$. Then, if $(\cdot)_\sigma$ indicates the submatrix formed by the columns indexed by σ, the dictionary for the spike stream is given by $(\Phi\widehat{H}_i)_\sigma = \Phi(\widehat{H}_i)_\sigma$. By virtue of the least-squares property of the pseudo-inverse operator, the subsequent estimate $\widehat{x}_{i+1}$ according to Step 2 is given by

$$\widehat{x}_{i+1} = \arg\min_x \|y - \Phi(\widehat{H}_i)_\sigma x\|_2^2, \tag{21}$$

where x belongs to the K-dimensional subspace defined by the support configuration σ. Since we are minimizing a convex loss function (squared error) on a subspace in $\mathbb{R}^N$, the minimum $\widehat{x}_{i+1}$ is unique. Therefore, we may view Step 2 of the algorithm as a unique-valued *infimal map* f from a given $\widehat{h}_i \in \mathcal{M}_F$ to a particular $\widehat{x}_{i+1} \in \mathcal{M}_S^\Delta$. Similarly, we may view Step 4 of Algorithm 1 as another unique-valued infimal map g from $\mathcal{M}_S^\Delta$ to $\mathcal{M}_F$. Therefore, the overall algorithm is a repeated application of the composite map $f \circ g$. From a well-known result on single-valued infimal maps [20, 31], the algorithm is strictly monotonic with respect to the loss function. Thus, the norm of the residual $y - \Phi\widehat{z}_i$ decreases with increasing iteration count i.

Further, any intermediate estimate $\widehat{z}_i$ also belongs to the model $\mathcal{M}(S, F, \Delta)$. We know from (12) that

$$\|y - \Phi\widehat{z}_i\|_2^2 = \|\Phi z - \Phi\widehat{z}_i\|_2^2 \geq (1 - \delta)\|z - \widehat{z}_i\|_2^2.$$

Therefore, if $\|y - \Phi\widehat{z}_i\| \leq \epsilon$, then $\|z - \widehat{z}_i\| \leq \epsilon/\sqrt{1 - \delta}$, thus proving the result.

APPENDIX C

We prove Theorem 3. Suppose the target signal z comprises S pulses $\{h_1, \ldots, h_S\}$, so that

$$z = \sum_{i=1}^{S} \alpha_i \mathbb{S}_i(h_i).$$

Assume we are given access to the Nyquist samples z, i,e., $\Phi = I_{N \times N}$. Suppose the estimate of the impulse response at an intermediate iteration is given by $\widehat{h}$. Let $\widehat{H}$ be the matrix formed by the operator $\mathbb{C}(\cdot)$ acting on $\widehat{h}$ and let σ be the candidate support configuration for the spike stream, so that the dictionary Φ_h in this case is given by the submatrix $\widehat{H}_\sigma$. Note that $\widehat{H}_\sigma$ is quasi-Toeplitz, owing to the assumption that the separation Δ is at least as great as the impulse response length F. Thus, Step 2 of Algorithm 1 can be represented by the least-squares operation

$$\widehat{x} = \widehat{H}_\sigma^\dagger z.$$

Due to the quasi-Toeplitz nature of $\widehat{H}_\sigma$, the pseudo-inverse $\widehat{H}_\sigma^\dagger = (\widehat{H}_\sigma^\top \widehat{H}_\sigma)^{-1}\widehat{H}_\sigma^\top$ essentially reduces to a scaled version of the identity multiplied by the transpose of $\widehat{H}$ (the scaling factor is in fact the squared norm of $\widehat{h}$). Thus, the spike coefficients are given by

$$\widehat{x} = \frac{1}{\|\widehat{h}\|^2}\widehat{H}_\sigma^\top \widetilde{H}\widetilde{x}.$$

Simplifying, we obtain the expression for the estimated i^{th} spike coefficient $\widehat{\alpha}_i$ as

$$\widehat{\alpha}_i = \alpha_i \frac{\langle h_i, \widehat{h}\rangle}{\|\widehat{h}\|^2}.$$

If $\widehat{h}$ is normalized, then we may write $\widehat{\alpha}_i = c_i\alpha_i$, where $c_i = \langle h_i, \widehat{h}\rangle$.

Once the spike coefficients $\widehat{x} = (c_1\alpha_1, \ldots, c_S\alpha_S)$ have been estimated, we can form the dictionary Φ_x by considering the quasi-Toeplitz matrix $\widehat{X}$ formed by the operation $\mathbb{C}(\widehat{x})$. In the same manner as above, an updated estimate of the pulse shape $\widehat{\widehat{h}}$ is given by

$$\widehat{\widehat{h}} = \widehat{X}^\dagger z = \frac{1}{\sum_{i=1}^{S} c_i^2 \alpha_i^2}\widehat{X}^T z.$$

28

Writing out $\widehat{X}$ and z in terms of $(h_1, \ldots, h_S)$ and $(\alpha_1, \ldots, \alpha_S)$ and simplifying, we obtain

$$\widehat{\widehat{h}} = \frac{\sum_{i=1}^{S} c_i \alpha_i^2 h_i}{\sum_{i=1}^{S} c_i^2 \alpha_i^2},$$

where $c_j = \langle h_j, \widehat{h} \rangle$. Thus, we have a closed-expression for the updated estimate of the impulse response coefficients $\widehat{\widehat{h}}$ in terms of the previous estimate $\widehat{h}$. In the event that the algorithm converges to a fixed point, we can replace $\widehat{\widehat{h}}$ by $\widetilde{h}$, thus proving the theorem.

REFERENCES

[1] C. Hegde and R. G. Baraniuk, "Compressive sensing of streams of pulses," in *Proc. Allerton Conference on Comm., Control and Sig. Proc.*, Sept. 2009.

[2] C. Hegde and R. G. Baraniuk, "Compressive sensing of a superposition of pulses," in *IEEE Intl. Conf. on Acoustics, Speech and Sig. Proc.*, Mar. 2010.

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

[4] 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.

[5] 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.

[6] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comp.*, vol. 20, p. 33, 1998.

[7] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Info. Theory*, vol. 56, pp. 1982–2001, April 2010.

[8] Y. Eldar and M. Mishali, "Robust recovery of signals from a union of subspaces," *IEEE Trans. Info. Theory*, vol. 55, pp. 5302–5316, Nov. 2009.

[9] D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmonic Analysis*, vol. 26, pp. 301–321, 2008.

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

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

[12] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Appl. Comp. Harmonic Analysis*, vol. 3, pp. 265–274, 2009. Preprint.

[13] Y. M. Lu and M. N. Do, "Sampling signals from a union of subspaces," *IEEE Sig. Proc. Mag.*, vol. 25, pp. 41–47, Mar. 2008.

[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] K. Herrity, R. Raich, and A. O. Hero, "Blind reconstruction of sparse images with unknown point spread function," in *Proc. SPIE*, vol. 6814, 2008.

[16] F. Naini, R. Gribonval, L. Jacques, and P. Vandergheynst, "Compressive sampling of pulse trains: spread the spectrum!," in *IEEE Intl. Conf. on Acoustics, Speech and Sig. Proc.*, pp. 2877–2880, 2009.

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

[18] T. F. Chan and C. K. Wong, "Convergence of the alternating minimization algorithm for blind deconvolution," *Lin. Algebra and Its Appl.*, vol. 316, no. 1-3, pp. 259 – 285, 2000.

[19] J.-L. Starck, E. Pantin, and F. Murtagh, "Deconvolution in astronomy: A review," *Publication of the Astronomical Society of the Pacific*, vol. 114, pp. 1051–1069, 2002.

[20] J. A. Tropp, "An alternating minimization algorithm for non-negative matrix factorization." Preprint, 2003.

[21] R. Baraniuk and P. Steeghs, "Compressive radar imaging," in *IEEE Radar Conference*, pp. 128–133, April 2007.

[22] http://heritage.stsci.edu/gallery/bwgallery/bw0405/about.shtml, "NASA online archive."

[23] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Processing*, vol. 50, pp. 1417–1428, June 2002.

[24] S. Gleichman and Y. Eldar, "Blind compressed sensing," Feb. 2010. Preprint. Available at http://arxiv.org/abs/1002.2586.

[25] M. Yaghoobi, T. Blumensath, and M. E. Davies, "Dictionary learning for sparse approximations using the majorization method," *IEEE Trans. Sig. Proc.*, vol. 57, no. 6, pp. 2178–2191, 2009.

[26] V. Saligrama and M. Zhao, "Compressed blind deconvolution of filtered sparse processes," Oct. 2009. Preprint. Available at http://arxiv.org/abs/0910.0239.

[27] M. S. Asif, W. E. Mantzel, and J. Romberg, "Channel protection: Random coding meets sparse channels," in *IEEE Info. Theory Workshop*, Oct. 2009.

[28] M. S. Asif, W. E. Mantzel, and J. Romberg, "Random channel coding and blind deconvolution," in *Proc. Allerton Conference on Comm., Control and Sig. Proc.*, 2009.

[29] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin, "A simple proof of the restricted isometry property for random matrices," *Const. Approx.*, vol. 28, pp. 253–263, Dec. 2008.

[30] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz maps into a Hilbert space," *Contemp. Math.*, vol. 26, pp. 189–206, 1984.

[31] J. C. Fiorot and P. Huard, *Composition and union of general algorithms of optimization*, vol. 10, pp. 69–85. Springer-Berlin, Heidelberg, 1979.
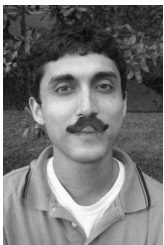
Chinmay Hegde (S'06) received the B. Tech. degree in electrical engineering from Indian Institute of Technology, Madras (Chennai, India), in 2006, and the M.S. degree in electrical and computer engineering from Rice University, Houston, TX, in 2010.

He is currently a Ph.D. student in the Department of Electrical and Computer Engineering at Rice University (Houston, Texas). His research interests include compressive sensing, statistical and manifold-based signal and image processing.

He was coauthor on a paper with M. F. Duarte and V. Cevher that won the Signal Processing with Adaptive Sparse Structured Representations (SPARS) Workshop Best Student Paper Award in 2009.

Richard G. Baraniuk (M'93–SM'98–F'01) received the B.Sc. degree from the University of Manitoba, Winnipeg, MB, Canada, in 1987, the M.Sc. degree from the University of Wisconsin-Madison, in 1988, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, in 1992, all in electrical engineering.

After spending 1992-1993 with the Signal Processing Laboratory of Ecole Normale Suprieure, Lyon, France, he joined Rice University, Houston, TX, where he is currently the Victor E. Cameron Professor of Electrical and Computer Engineering. He spent sabbaticals at Ecole Nationale Suprieure de Tlcommunications in Paris in 2001 and Ecole Fdrale Polytechnique de Lausanne in Switzerland in 2002. His research interests lie in the area of signal and image processing. In 1999, he founded Connexions (cnx.org), a nonprofit publishing project that invites authors, educators, and learners worldwide to "create, rip, mix, and burn" free textbooks, courses, and learning materials from a global open-access repository.

Prof. Baraniuk has been a Guest Editor of several special issues of the *IEEE Signal Processing Magazine*, *IEEE Journal of Special Topics in Signal Processing*, and the *Proceedings of the IEEE* and has served as technical program chair or on the technical program committee for several IEEE workshops and conferences. He received a NATO postdoctoral fellowship from NSERC in 1992, the National Young Investigator award from the National Science Foundation in 1994, a Young Investigator Award from the Office of Naval Research in 1995, the Rosenbaum Fellowship from the Isaac Newton Institute of Cambridge University in 1998, the C. Holmes MacDonald National Outstanding Teaching Award from Eta Kappa Nu in 1999, the Charles Duncan Junior Faculty Achievement Award from Rice in 2000, the University of Illinois ECE Young Alumni Achievement Award in 2000, the George R. Brown Award for Superior Teaching at Rice in 2001, 2003, and 2006, the Hershel M. Rich Invention Award from Rice in 2007, the Wavelet Pioneer Award from SPIE in 2008, and the Internet Pioneer Award from the Berkman Center for Internet and Society at Harvard Law School in 2008. He was selected as one of *Edutopia Magazines* Daring Dozen educators in 2007. Connexions received the Tech Museum Laureate Award from the Tech Museum of Innovation in 2006. His work with Kevin Kelly on the Rice single-pixel compressive camera was selected by *MIT Technology Review Magazine* as a TR10 Top 10 Emerging Technology in 2007. He was coauthor on a paper with M. Crouse and R. Nowak that won the IEEE Signal Processing Society Junior Paper Award in 2001 and another with V. Ribeiro and R. Riedi that won the Passive and Active Measurement (PAM) Workshop Best Student Paper Award in 2003. He was elected a Fellow of the IEEE in 2001 and a Fellow of the AAAS in 2009.