RICE UNIVERSITY

Sampling and Recovery of Pulse Streams

by

Chinmay Hegde

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:

Richard G. Baraniuk, Chair,
Victor E. Cameron Professor,
Electrical and Computer Engineering

Don H. Johnson
J. S. Abercrombie Professor,
Electrical and Computer Engineering,
Statistics

Wotao Yin,
Assistant Professor,
Computational and Applied Mathematics

HOUSTON, TEXAS

APRIL 2010

UMI Number: 1486067

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.

UMI 1486067

Copyright 2010 by ProQuest LLC.

All rights reserved. This edition of the work is protected against unauthorized copying under Title 17, United States Code.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

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.

Acknowledgements

I owe my deepest gratitude to my inimitable advisor, Richard Baraniuk. Rich is an inspirational researcher and teacher, and I thank him for his patience, enthusiasm, and encouragement. Special thanks to Dr. Don Johnson and Dr. Wotao Yin for their terrific courses and their helpful guidance through the years.

Much of the work in this thesis stems from insightful conversations with various members of Team DSP, notably Volkan Cevher, Marco Duarte, Eva Dyer, and Mona Sheikh. I gratefully acknowledge them for their help. Thanks also to Dr. Manjit Sanghera of the Dallas Presbyterian Hospital for providing the neuronal data used in some experiments.

Lastly, I thank my parents for their infinite love and support, and the fine folks at ISAR for truly helping me find a home away from home.

Contents

1	Introduction	1
2	Background on Compressive Sensing	6
2.1	Sparse signal models	6
2.2	Signal acquisition via nonadaptive linear measurements	7
2.3	Recovery methods	9
3	Signal Models for Pulse Streams	12
3.1	General model	12
3.2	Special case: Disjoint pulses	14
4	Sampling Theorems for Pulse Streams	17
4.1	General pulse streams	17
4.2	Special case: Disjoint pulse streams	19
5	Recovery of Pulse Streams	21
5.1	Alternating minimization with exhaustive search	22
5.2	Model mismatch	25
5.3	Iterative support estimation	28
5.4	Stability and convergence	31
5.5	Computational complexity	31
6	Numerical Experiments	32
6.1	Synthetic 1D pulse streams	32
6.2	Neuronal signals	34
6.3	Synthetic 2D pulse streams	35
6.4	Astronomical images	36
7	Discussion and Conclusions	38
A	Proof of Theorem 1	41
B	Proof of Theorem 2	45
C	Proof of Theorem 3	47
D	Solution of (5.3)	50

List of Figures

1.1	(a) Test signal of length $N = 1024$ obtained by convolving a spike stream with $S = 6$ with an impulse response of length $F = 11$, so that the total signal sparsity $K = SF = 66$. (b) Profile of one pulse. Signal reconstruction from $M = 100$ random Gaussian measurements performed using (c) a state-of-the-art CS recovery algorithm (CoSaMP [24], $MSE = 13.42$), and (d) our proposed Algorithm 2 ($MSE = 0.0028$).	4
6.1	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.	33
6.2	(a) Synthetic noisy signal ($SNR = 13.25\text{dB}$). (b) Recovery from $M = 150$ random measurements using Algorithm 2.	34
6.3	CS recovery of a real-world neuronal signal. (a) Original recording. (b) Recovered signal using $M = 150$ random measurements. (c) Estimated anchor pulse shape ($F = 11$).	34
6.4	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 ($MSE = 0.07$).	35
6.5	(a) Black-and-white image of V838 Monocerotis, a nova-like star, captured by the Hubble Space Telescope on February 8, 2004 [19]. (b) Test image is a zoomed in-version of the region 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.	37

Chapter 1

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 as fast as 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) [9, 13]. 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 subgaussian probability distributions, and (ii) the number of samples M is $\mathcal{O}(K \log(N/K))$. Further, this recovery can be carried out in polynomial time using efficient greedy approaches or optimization based methods [12, 30].

For some applications, there exist more restrictive 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 that are based on *structured sparsity* models that are equivalent to a finite union of subspaces [4, 14]. 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, a electrophysiolog-

ical recording of a neuron is often approximated as a series of spikes but can be better 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*. For the compressive sensing and recovery of a pulse stream, the number of measurements M would incur a corresponding multiplicative increase by a factor of F when compared to sensing merely the underlying spike streams; 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 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 proposed model. Our derivation shows that $M = \mathcal{O}((S + F) \log N)$; i.e., M is proportional to the number of degrees of freedom of the signal $S + F$ but *sublinear* in the total sparsity $K = SF$. Third, we develop algorithms to recover signals from

Figure 1.1: (a) Test signal of length $N = 1024$ obtained by convolving a spike stream with $S = 6$ with an impulse response of length $F = 11$, so that the total signal sparsity $K = SF = 66$. (b) Profile of one pulse. Signal reconstruction from $M = 100$ random Gaussian measurements performed using (c) a state-of-the-art CS recovery algorithm (CoSaMP [24], $MSE = 13.42$), and (d) our proposed Algorithm 2 ($MSE = 0.0028$).

our model from M measurements. 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 shape. Numerical experiments on real and synthetic data sets demonstrate the benefits of the approach. As demonstrated in Figure 1.1, 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 2, we review the rudiments of standard and structured sparsity-based CS. In Section 3, we propose a deterministic signal model for pulse streams and discuss its geometric properties. In Section 4, we derive

bounds on the number of random measurements required to sample signals belonging to our proposed model. In Section 5, we develop an algorithm for stable signal recovery and analyze its convergence and robustness to model mismatch. Numerical results are presented in Section 6, followed by a concluding discussion in Section 7.

Chapter 2

Background on Compressive Sensing

2.1 Sparse signal models

A signal $x \in \mathbb{R}^N$ is K -sparse in the orthonormal basis Ψ 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 being 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

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

Sparsity is a flexible model for the efficient compression and processing of several interesting classes of signals. Nonetheless, sparsity can be simplistic in the sense that it assumes *no additional structure* about the inter-relations between the transform

coefficients α . Consider the class of one-dimensional piecewise smooth signals; we not only know that such signals are sparse in the wavelet basis, but also that the wavelet coefficients lie on (or approximately on) a connected tree [22]. Such classes of signals that permit only a small number of admissible support configurations may be modeled by a restricted union of subspaces, consisting only of L_K canonical subspaces (so that L_K is typically much smaller than $\binom{N}{K}$). Thus, if $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_{L_K}\}$ denotes the set of admissible supports, then a *structured sparsity model* [4] is the set

$$\mathcal{M}_K := \{x : \sigma(x) \in \Sigma\}. \quad (2.1)$$

Thus, any structured sparsity model $\mathcal{M}_K$ is a subset of 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.

2.2 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 a 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 CS 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):

Definition 1. [8] An $M \times N$ matrix Φ has the K -RIP with constant δ_K if, for all $x \in \Sigma_K$,

$$(1 - \delta_K)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 + \delta_K)\|x\|_2^2. \quad (2.2)$$

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, if $M < N$; indeed, deterministic design of a sampling matrix having the RIP is an NP-complete problem. Nevertheless, it has been shown [8] that provided $M \geq \mathcal{O}(K \log(N/K))$, a matrix Φ whose elements are i.i.d. samples from a random subgaussian distribution possesses the RIP with high probability. 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 can be proposed for more general structured sparsity models containing L_K canonical subspaces [4, 6, 14]. This is known as the model-based RIP and is defined thus: Φ satisfies the $\mathcal{M}_K$ -RIP if (2.2) holds for all $x \in \mathcal{M}_K$. It can be shown [6] that the number of measurements M necessary for a subgaussian CS matrix to have the $\mathcal{M}_K$ -RIP with constant δ and with probability $1 - e^{-t}$ is bounded as

$$M \geq \frac{c}{\delta^2} \left(\ln(2L_K) + K \ln \frac{12}{\delta} + t \right). \quad (2.3)$$

We may make two inferences from (2.3). 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.

2.3 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.,

$$\hat{x} = \arg \min_{x'} \|x'\|_0 \text{ subject to } y = \Phi x'. \quad (2.4)$$

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 [8].

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 (2.4), i.e.,

$$\hat{x} = \arg \min_{x'} \|x'\|_1 \text{ such that } y = \Phi x', \quad (2.5)$$

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 [12]. The second approach entails an iterative, greedy selection of the support of the true solution x . This approach is employed by several algorithms such as orthogonal matching pursuit (OMP) [30], compressive sampling matching pursuit (CoSaMP) [24], and iterative hard thresholding [7].

Both kinds of 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 $\hat{x}$ obtained by these algorithms has bounded error:

$$\|x - \hat{x}\|_2 \leq C_1 \|x - x_K\|_2 + \frac{C_2}{\sqrt{K}} \|x - x_K\|_1 + C_3 \|n\|_2, \quad (2.6)$$

where x_K is the best K -sparse approximation to x and C_1, C_2 are constants. This result implies that given *noiseless* measurements, a K -sparse signal can be exactly recovered using these algorithms.

Furthermore, with a simple modification, algorithms like CoSaMP and iterative hard thresholding can be used to reconstruct signals belonging to any structured sparsity model Σ defined in (2.1) [4]. For instance, CoSaMP possesses an iterative structure that hinges on computing the best K -term approximation to an intermediate signal estimate $x' \in \mathbb{R}^N$; in other words, the algorithm requires projecting x' on to the set of all K -sparse signals Σ_K . We simply replace this sparse approximation step

by a best *model-approximation* step, i.e., we instead project x' on to Σ . The central result in [4] states that any signal belonging to the structured sparsity model Σ can be accurately reconstructed provided the number of measurements M satisfies the less stringent bound specified in (2.3).

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.

Chapter 3

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.

3.1 General model

Consider the following deterministic model for signals that can 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 (2.1). Similarly, let $\mathcal{M}_F \subset \mathbb{R}^N$ be a union of F -dimensional canonical

subspaces. Consider the set

$$\mathcal{M}_{S,F}^z := \{z \in \mathbb{R}^N : z = x * h, \text{ such that } x \in \mathcal{M}_S \text{ and } h \in \mathcal{M}_F\}, \quad (3.1)$$

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

We make two immediate observations:

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:

$$z = x * h = h * x = Hx = Xh, \quad (3.2)$$

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 5.

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 (3.1) 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 an *infinite union of subspaces*.¹ Note that (2.3) cannot be applied in this case since it only considers finite unions of subspaces. 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 4.

Note that the exact definition of 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.

3.2 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 4 we develop a lower bound on the number of samples required to preserve the essential information contained in

¹A general theory for sampling signals from infinite unions of subspaces has been introduced in [21].

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 5. 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 (3.2). Note that H and x need not be unique for a given z ; any ordered pair $(\alpha H, x/\alpha)$ satisfies (3.2), 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 $-\tau$. To eliminate these ambiguities, we make the following two assumptions:

1. the impulse response h is *concentrated*, i.e., all the 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 such time-domain signals with sufficiently separated nonzeros has been introduced in [17].

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 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*

$$\mathcal{M}(S, F, \Delta) = \{z \in \mathbb{R}^N : z = x * h, \text{ such that } x \in \mathcal{M}_S^\Delta \text{ and } h \in \mathcal{M}_F\}. \quad (3.3)$$

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 5). See Figure 1.1(a) for an example stream of disjoint pulses in 1D.

Chapter 4

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.

4.1 General pulse streams

We derive a sampling theorem for signals belonging to the model $\mathcal{M}_{S,F}^z$ proposed in Definition 2. 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. subgaussian matrix Φ with

$$M \geq \mathcal{O} \left(\frac{1}{\delta} \left((S + F) \ln \left(\frac{1}{\delta} \right) + \log(L_S L_F) + t \right) \right). \quad (4.1)$$

Then, Φ satisfies the following property with probability at least $1 - e^{-t}$: for every pair

$$z_1, z_2 \in \mathcal{M}_{S,F}^z,$$

$$(1 - \delta) \|z_1 - z_2\|_2^2 \leq \|\Phi z_1 - \Phi z_2\|_2^2 \leq (1 + \delta) \|z_1 - z_2\|_2^2. \quad (4.2)$$

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,

$$L_S \leq \binom{N}{S} \approx \left(\frac{eN}{S} \right)^S. \quad (4.3)$$

Similarly, $L_F \leq \left(\frac{eN}{F} \right)^F$. Therefore, the logarithmic term in the expression for M in (4.1) scales as:

$$\log(L_S L_F) \leq S + S \log(N/S) + F + F \log(N/F) \leq 2(S + F) \log N \quad (4.4)$$

Thus, (4.1) indicates that the number of measurements M required for the sampling of 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 [6]. 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 $\mathcal{M}_{S,F}^z$.

4.2 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 [17] states that the number of subspaces in the structured sparsity model $\mathcal{M}_S^\Delta$ is given by

$$L_S = \binom{N - S\Delta + S - 1}{S - 1}. \quad (4.5)$$

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*

$$M \geq \mathcal{O} \left(\frac{1}{\delta} \left((S + F) \ln \left(\frac{1}{\delta} \right) + S \log(N/S - \Delta) + t \right) \right), \quad (4.6)$$

then an $M \times N$ i.i.d. gaussian matrix Φ will satisfy (4.2) 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, (4.6) implies that fewer measurements are needed to sample pulse streams in which the pulses are widely separated.

Chapter 5

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 we are given the noisy measurements

$$y = \Phi z + n = \Phi Hx + n = \Phi Xh + n,$$

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 [18,23,26]; the common approach has been to first assume that a training set of vectors $\{x_i\}$ exists for a fixed impulse response h , then infer the coefficients of h using a sparse learning algorithm (such as LASSO [29] or basis pursuit [12]), and then solve for the coefficients $\{x_i\}$. In the absence of train-

ing data, we must infer both the spike locations and the impulse response coefficients. Therefore, our task is also similar to *blind deconvolution* [16]; 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 [11]. 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 $\mathcal{M}(S, F, \Delta)$.

5.1 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 $\hat{H}$ from all possible shifts of the current estimate of the impulse response $\hat{h}$ (denote this operation as $\hat{H} = \mathbb{C}(\hat{h})$). Further, we calculate the dictionary $\Phi \hat{H}$ for the spike stream x , and select the submatrix formed

by the columns indexed by the assumed spike locations σ (denote this submatrix as $(\Phi\hat{H})_\sigma$). 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 an estimate for $\hat{x}$:

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

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

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

Then, we form our signal estimate $\hat{z} = \hat{x} * \hat{h}$. The above two-step process is iterated until a suitable halting criterion (e.g., convergence in norm for the estimated signal $\hat{z}$). This process is akin to the Richardson-Lucy algorithm for blind deconvolution [28].

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 $r = y - \Phi\hat{z}$. The procedure is detailed in pseudocode form in Algorithm 1. Thus, Algorithm 1 consists of performing alternating minimization for a given estimate for the support of the

Algorithm 1 Alternating minimization with exhaustive search

Inputs: Sampling matrix Φ , measurements $y = \Phi x$,
model parameters Δ, S, F , threshold ϵ
Output: $\hat{z} \in \mathcal{M}(S, F, \Delta)$ such that $y - \Phi \hat{z}$ is small
 $\hat{x} = 0, \hat{h} = (\mathbf{1}_F^T, 0, \dots, 0)/\sqrt{F}; i = 0$ {initialize}
for $\sigma \in \mathcal{M}_S^\Delta$ **do**
 1. $\hat{H} = \mathbb{C}(\hat{h}), \Phi_h = (\Phi \hat{H})_\sigma$ {form dictionary for spike stream}
 2. $\hat{x} \leftarrow \Phi_h^\dagger y$ {update spike stream estimate }
 3. $\hat{X} = \mathbb{C}(\hat{x}), \Phi_x = (\Phi \hat{X})_f$ {form dictionary for impulse response}
 4. $\hat{h} \leftarrow \Phi_x^\dagger y$ {update impulse response estimate}
 5. $\hat{z} \leftarrow \hat{x} * \hat{h}$ {form signal estimate}
if $\|y - \Phi \hat{z}\|_2 < \epsilon$ {check for energy in residual}
 return $\hat{z}$
end if
end for

underlying spike stream x , and exhaustively searching 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 (4.2) 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 $\hat{z}_i$ be an intermediate estimate of z at iteration i of Algorithm 1. Then:*

1. *The norm of the residual $\|y - \Phi \hat{z}_i\|_2$ monotonically decreases with iteration count i .*
2. *If at any iteration i*

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

then we are guaranteed that

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

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 [31]. 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.

5.2 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 $\{h_1, h_2, \dots, h_S\}$ be length- F vectors corresponding to each of the S pulses in the signal, and let the length- S spike stream $\tilde{x} = (\alpha_1, \alpha_2, \dots, \alpha_S)$. 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

$$z = \sum_{i=1}^S \alpha_i \mathbb{S}_i(h_i), \tag{5.1}$$

or equivalently,

$$z = \tilde{H}\tilde{x},$$

where $\tilde{H} = [\mathbb{S}_1(h_1), \dots, \mathbb{S}_S(h_S)]$ is an $N \times S$ matrix. Assuming that the spikes in x are separated by at least Δ locations, the matrix $\tilde{H}$ is *quasi-Toeplitz* [3], i.e., the columns of $\tilde{H}$ are circular shifts of one another with no more than one nonzero entry in every row. An attractive property of quasi-Toeplitz matrices is that there exist analytical expressions for their pseudo-inverses. Suppose 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, and hence $\Phi_h^\dagger$ and $\Phi_x^\dagger$ can be computed in closed form. Thus, given an estimate of the pulse shape $\hat{h}_0$, we can derive closed-form expressions for the next impulse response estimate.

Additionally, we can obtain an intermediate estimate for the spike stream $\tilde{x}$. Suppose the innermost loop of Algorithm 1 converges to a fixed point estimate $\hat{h}$. Since the least-squares equations are homogenous, we may assume that $\|\hat{h}\|_2 = 1$ without loss of generality. We dub $\hat{h}$ the *anchor pulse* for the set of pulse shapes $\{h_1, h_2, \dots, h_S\}$. The following theorem provides an expression relating the anchor pulse to the component pulse shapes.

Theorem 3. *Consider z as defined in (5.1). Let $\hat{h}$ be the anchor pulse for the set of pulse shapes $\{h_1, h_2, \dots, h_S\}$. Define $c_i = \langle h_i, \hat{h} \rangle$ for $i = 1, \dots, S$. Then, we have that*

$$\hat{h} = \frac{\sum_{i=1}^S c_i \alpha_i^2 h_i}{\sum_{i=1}^S c_i^2 \alpha_i^2}. \quad (5.2)$$

The proof of this theorem is provided in Appendix C. Equation (5.2) implies that the anchor pulse $\hat{h}$ is a weighted linear combination of the component pulses h_i , $i = 1, \dots, S$, 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 $\tilde{x}$ 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 $\{h_1, h_2, \dots, h_S\}$, 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 2, 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.

5.3 Iterative support estimation

Algorithm 1 involves iteratively solving a combinatorial number of estimation problems. This becomes infeasible for even moderate values of N . A simpler method can be proposed as follows: 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 $\hat{x}$ and impulse response $\hat{h}$, 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 [24]. We provide an outline of the algorithm here for completeness; see [17] 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 $x = (x_1, x_2, \dots, x_N)^T$. Given any binary vector $s = (s_1, s_2, \dots, s_N)^T$ of length N , let:

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

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 $\mathcal{M}_S^\Delta$:

$$s_1 + s_2 + \dots + s_N \leq S, \quad (5.3)$$

$$s_j + s_{j+1} \dots + s_{j+\Delta-1} \leq 1, \text{ for } j = 1, \dots, N, \quad (5.4)$$

where the subscripts are computed modulo N . The first inequality (5.3) specifies that the solution contains at most S nonzeros; the other N inequalities (5.4) 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 $c = (x_1^2, x_2^2, \dots, x_N^2)$, 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 (5.3) and (5.4). Next, define $u \in \mathbb{R}^{N+1} = (S, 1, 1, \dots, 1)$; this represents the right hand side of the constraints (5.3) and (5.4). Thus, we can represent (5.3) and (5.4) by the following binary integer program:

$$s^* = \arg \min c^T s, \text{ subject to } Ws \leq u.$$

Next, we relax the integer constraints on s to obtain a computationally tractable linear program. Denote this linear program by $\mathbb{D}(\cdot)$. In Appendix D, we show 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

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 $\hat{z}$ to true signal z

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

while halting criterion false **do**

1. $i \leftarrow i + 1$
2. $\hat{z} \leftarrow \hat{x} * \hat{h}$ {current pulse stream estimate}
- {estimate spike locations and amplitudes}
3. $\hat{H} = \mathbb{C}(\hat{h})$, $\Phi_h = \Phi \hat{H}$ {form dictionary for spike stream}
4. $e \leftarrow \Phi_h^T (y - \Phi_h \hat{x})$ {residual}
5. $\omega \leftarrow \sigma(\mathbb{D}(e))$ {obtain model-approximation of residual}
6. $\sigma \leftarrow \omega \cup \sigma(\hat{x}_{i-1})$ {merge supports}
7. $x|_\sigma \leftarrow (\Phi_h)_\sigma^\dagger y$, $x|_{\sigma^c} = 0$ {update spike stream estimate}
8. $\hat{x} \leftarrow \mathbb{D}(x)$ {prune spike stream estimate}
- {estimate impulse response}
9. $\hat{X} = \mathbb{C}(\hat{x})$, $\Phi_x = (\Phi \hat{X})_f$ {form dictionary for impulse response}
10. $\hat{h} \leftarrow \Phi_x^\dagger y$ {update impulse response estimate}

end while

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

best $\mathcal{M}_S^\Delta$ -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 $\hat{z}$.) The overall algorithm can be viewed as an iterative sparse approximation procedure for the $\mathcal{M}_S^\Delta$ model that continually updates its estimate of the sparsifying dictionary. The procedure is detailed in pseudocode form in Algorithm 2.

5.4 Stability and convergence

Like many other algorithms for blind deconvolution, the analysis of Algorithm 2 is not straightforward. The dictionaries $\Phi\hat{X}$ and $\Phi\hat{H}$ 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 6.

5.5 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 $\mathcal{O}(N^{3.5})$. The total computational cost also scales linearly in the number of measurements M and the number of iterations T of the outer loop executed until convergence; thus, overall the algorithm runs in polynomial time.

Chapter 6

Numerical Experiments

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

6.1 Synthetic 1D pulse streams

Figure 1.1 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 6.1 displays the averaged results of a Monte Carlo simulation of Algorithm 2 over 200 trials. Each trial was conducted by generating a sample signal belonging

Figure 6.1: 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.

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 . It is clear from the figure that Algorithm 2 outperforms both conventional CS recovery (CoSaMP [24]) with target sparsity $K = SF$ as well as block-based reconstruction [4] 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.

We show that Algorithm 2 is stable to small amounts of noise in the signal and the measurements. In Figure 6.2, we generate a length $N = 1024$ signal from a disjoint pulse stream model with $S = 9$ and $F = 11$; add a small amount of Gaussian noise (SNR = 13.25dB) 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.

Figure 6.2: (a) Synthetic noisy signal ($\text{SNR} = 13.25\text{dB}$). (b) Recovery from $M = 150$ random measurements using Algorithm 2.

Figure 6.3: CS recovery of a real-world neuronal signal. (a) Original recording. (b) Recovered signal using $M = 150$ random measurements. (c) Estimated anchor pulse shape ($F = 11$).

6.2 Neuronal signals

We test Algorithm 2 on a real-world neuronal recording. Figure 6.3(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 6.3(b)) as well as an estimate of the anchor pulse shape (Figure 6.3(c)).

Figure 6.4: 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).

6.3 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 6.4(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 6.4 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.

6.4 Astronomical images

Finally, we test Algorithm 2 on a real astronomical image. Our test image is a $N = 64 \times 64$ region of a high-resolution image of V838 Monocerotis (a nova-like variable star) captured by the Hubble Space Telescope [19] (highlighted by the green square in Figure 6.5(a)). Note the significant variations in the shapes of the three large pulses in the test image (Figure 6.5(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 6.5, 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.

Figure 6.5: (a) *Black-and-white image of V838 Monocerotis, a nova-like star, captured by the Hubble Space Telescope on February 8, 2004 [19].* (b) *Test image is a zoomed in-version of the region 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*.

Chapter 7

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 an infinite union of subspaces which exhibits a particular geometric structure. This structure enables us to quantitatively deduce 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 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 best basis compressive sensing [26], simultaneous sparse approximation and dictionary learning [32], and the classical signal processing problem of blind deconvolution [16]. Compressive sensing of time-domain pulse streams has been studied by Naini *et al.* [23]. 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 [27]. 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. However, their measurement system is tailored to impulse responses corresponding to AR linear models; 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.* [1, 2], 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 coding strategy follows the one advocated by Candès and Tao [10]: their channel code consists of a random matrix $\Phi \in \mathbb{R}^{M \times N}$ where $M > N$, 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

Proof of 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 is 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 [20] for stable embedding of point clouds to this finite set of points Q , and extend the stable embedding to all possible signals $z \in U_{\Omega,\Lambda}$. Finally, we derive our main result through a union bound over all possible choices of subspaces Ω and Λ .

Consider a fixed vector $h \in \Lambda$. Suppose the coefficients of h are normalized so that $\|h\| = 1$. By virtue of its circulant nature, the spectral norm of the corresponding matrix $\|H\| \leq 1$. Now, consider a fixed S -dimensional subspace $\Omega \in \mathcal{M}_S$. It is easy to see that

$$\Omega_h = \{z = Hx \mid x \in \Omega\}$$

also forms an S -dimensional subspace in $\mathbb{R}^N$. Thus, by Lemma 5.1 of [5], we can find a finite set of points $Q_{\Omega,h} \subset \Omega_h$ with cardinality $|Q_{\Omega,h}| \leq (3/\delta')^S$ such that

$$\min_{q \in Q_{\Omega,h}} \|Hx - q\| \leq \delta', \quad \forall \|x\| \leq 1, x \in \Omega.$$

This is an upper bound on the size of $Q_{\Omega,h}$; assuming a worst-case scenario, we may list out the points in this set so that

$$Q_{\Omega,h} = \{q_1, q_2, \dots, q_{(3/\delta')^S}\} = \{Hx_1, Hx_2, \dots, Hx_{(3/\delta')^S}\}.$$

Select any $x_l \in \{x_1, \dots, x_{(3/\delta')^S}\}$ and an F -dimensional subspace $\Lambda \in \mathcal{M}_F$. Form the circulant matrix X_l ; as above, $\|X_l\| \leq 1$. Therefore,

$$\Omega_{x_l} = \{z = X_l h \mid h \in \Lambda\}$$

forms an F -dimensional subspace. Correspondingly, we can find a set of points

$Q_{x_l, \Lambda} \subset \Omega_{x_l}$ with cardinality $|Q_{x_l, \Lambda}| \leq (3/\delta')^F$ such that

$$\min_{q \in Q_{x_l, \Lambda}} \|X_l h - q\| \leq \delta', \quad \forall \|h\| \leq 1, h \in \Lambda.$$

Using this process, define $Q_{x_l, \Lambda}$ for $l = 1, 2, \dots, (3/\delta')^S$. Then, we have

$$Q_{\Omega, \Lambda} = \bigcup_l Q_{x_l, \Lambda}.$$

Thus, we have identified a finite set of points $Q_{\Omega, \Lambda}$ in the infinite union of subspaces $U_{\Omega, \Lambda}$. Observe that the cardinality of this set $Q_{\Omega, \Lambda} = (3/\delta')^S (3/\delta')^F$. Then, every vector in $U_{\Omega, \Lambda}$ with magnitude less than 1 lies ‘close’ to at least one point in $Q_{\Omega, \Lambda}$, i.e., $Q_{\Omega, \Lambda}$ is a δ'' -net for $U_{\Omega, \Lambda}$. Suppose $\delta = 2\delta''$. 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 every pair of vectors $z_1, z_2 \in U_{\Omega, \Lambda}$, (4.2) will hold with failure probability

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

This 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 theorem follows easily from this result.

Appendix B

Proof of 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

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

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 [15, 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 (4.2) 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 $\|y - \Phi \widehat{z}_i\| \leq \epsilon$, $\|z - \widehat{z}_i\| \leq \epsilon / \sqrt{1 - \delta}$.

Appendix C

Proof of Theorem 3

Suppose the target signal z is composed of S pulses $\{h_1, \dots, h_S\}$, so that

$$z = \tilde{H}\tilde{x},$$

where $\tilde{H} = [\mathbb{S}_1(h_1), \dots, \mathbb{S}_S(h_S)]$ is an $N \times S$ matrix and $\tilde{x} = (\alpha_1, \alpha_2, \dots, \alpha_S)$. 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 $\hat{h}$. Let $\hat{H}$ be the matrix formed by the operator $\mathbb{C}(\cdot)$ acting on $\hat{h}$ 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 $\hat{H}_\sigma$. Note that $\hat{H}_\sigma$ 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

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

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

Once the spike coefficients $\widehat{x} = (c_1 \alpha_1, \dots, c_S \alpha_S)$ 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

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

Writing out $\widehat{X}$ and z in terms of $(h_1, \dots, h_S)$ and $(\alpha_1, \dots, \alpha_S)$ 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 $\widehat{h}$, thus proving the theorem.

Appendix D

Solution of (5.3)

First, we note that the binary matrix W is *totally unimodular* (TU), i.e., the determinant of every square submatrix of W is equal to 0, 1 or -1 . This follows from the fact that W is a binary matrix with the 1's in every row occurring in consecutive blocks, i.e., W is a so-called “interval matrix”; interval matrices are well-known to be totally unimodular [25].

We next show that the polytope formed by the system of inequalities in the relaxed linear program has integer basic feasible solutions. Since W is TU, every basic feasible solution is determined by a nonsingular submatrix of W . Since the determinant of this submatrix is ± 1 and the right hand side of the linear system is integer, by Cramer's rule we get that all components of every basic feasible solution are integers. The desired result follows from this observation, since the optimum of the linear program has to occur in a basic feasible solution. Indeed, the solution has to be binary, since every s_i is positive and constrained to be lesser than 1.

Bibliography

- [1] M. S. Asif, W. E. Mantzel, and J. Romberg. Channel protection: Random coding meets sparse channels. In *IEEE Info. Theory Workshop*, Oct. 2009.
- [2] 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.
- [3] R. Baraniuk and P. Steeghs. Compressive radar imaging. In *IEEE Radar Conference*, pages 128–133, April 2007.
- [4] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Info. Theory*, 56(4):1982–2001, April 2010.
- [5] 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.*, 28(3):253–263, Dec. 2008.
- [6] T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Trans. Info. Theory*, Dec. 2008.
- [7] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comp. Harmonic Analysis*, 3:265–274, 2009. Preprint.

- [8] E. J. Candès. Compressive sampling. In *Proc. International Congress of Mathematicians*, volume 3, pages 1433–1452, Madrid, Spain, 2006.
- [9] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Info. Theory*, 52(2):489–509, Feb. 2006.
- [10] E. J. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Info. Theory*, 52(12):5406–5425, Dec. 2006. Preprint.
- [11] Tony F. Chan and C. K. Wong. Convergence of the alternating minimization algorithm for blind deconvolution. *Lin. Algebra and Its Appl.*, 316(1-3):259 – 285, 2000.
- [12] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comp.*, 20:33, 1998.
- [13] D. L. Donoho. Compressed sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, September 2006.
- [14] Y. Eldar and M. Mishali. Robust recovery of signals from a union of subspaces. *IEEE Trans. Info. Theory*, 55(11):5302–5316, Nov. 2009.
- [15] J. C. Fiorot and P. Huard. *Composition and union of general algorithms of optimization*, volume 10, pages 69–85. Springer-Berlin, Heidelberg, 1979.

- [16] S. Haykin. Blind deconvolution, 1994. Prentice Hall.
- [17] 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.
- [18] K. Herrity, R. Raich, and A. O. Hero. Blind reconstruction of sparse images with unknown point spread function. In *Proc. SPIE*, volume 6814, 2008.
- [19] <http://heritage.stsci.edu/gallery/bwgallery/bw0405/about.shtml>. NASA online archive.
- [20] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz maps into a Hilbert space. *Contemp. Math.*, 26:189–206, 1984.
- [21] Y. M. Lu and M. N. Do. Sampling signals from a union of subspaces. *IEEE Sig. Proc. Mag.*, 25(2):41–47, Mar. 2008.
- [22] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, 1999.
- [23] 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.*, pages 2877–2880, 2009.
- [24] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comp. Harmonic Analysis*, 26:301–321, 2008.

- [25] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1999.
- [26] G. Peyre. Best-basis compressed sensing. *IEEE Trans. Sig. Proc.*, 2010. To appear.
- [27] V. Saligrama and M. Zhao. Compressed blind deconvolution of filtered sparse processes. Oct. 2009. Preprint. Available at <http://arxiv.org/abs/0910.0239>.
- [28] J.-L. Starck, E. Pantin, and F. Murtagh. Deconvolution in astronomy: A review. *Publication of the Astronomical Society of the Pacific*, 114:1051–1069, 2002.
- [29] R. Tibshirani. Regression shrinkage and selection via the Lasso. *J. Royal Stat. Society, Series B (Methodological)*, 58(1):267–288, 1996.
- [30] J. Tropp and A. C. Gilbert. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Trans. Info. Theory*, 53(12):4655–4666, Dec. 2007.
- [31] J. A. Tropp. An alternating minimization algorithm for non-negative matrix factorization. Preprint, 2003.
- [32] M. Yaghoobi, T. Blumensath, and M. E. Davies. Dictionary learning for sparse approximations using the majorization method. *IEEE Trans. Sig. Proc.*, 57(6):2178–2191, 2009.