

Fast Algorithms for Delta-Separated Sparsity Projection

Henning Bruhn Oliver Schaudt

Abstract

We describe a fast approximation algorithm for the Δ -separated sparsity projection problem. The Δ -separated sparsity model was introduced by Hegde, Duarte and Cevher (2009) to capture the firing process of a single Poisson neuron with absolute refractoriness. The running time of our projection algorithm is linear for an arbitrary (but fixed) precision and it is both a head and a tail approximation. This solves a problem of Hegde, Indyk and Schmidt (2015).

We also describe how our algorithm fits into the approximate model iterative hard thresholding framework of Hegde, Indyk and Schmidt (2014) that allows to recover Δ -separated sparse signals from noisy random linear measurements. The resulting recovery algorithm is substantially faster than the existing one, at least for large data sets.

1 Introduction

Compressed sensing is based on the insight that real-life signals are often sparse. In character recognition, for example, most pixels are white or nearly so. That premise allows to obtain a signal with far fewer measurements than classic lower bounds suggest. It necessitates, however, a computational step that recovers the signal from those measurements.

Fast iterative recovery algorithms, such as *iterative hard thresholding* (IHT) [3] and *compressed sampling matching pursuit* (CoSaMP) [12], have been proposed. A common step in these algorithms is the so-called *sparsity projection* or *model approximation* step: to project a signal onto a sparse signal.

More concretely, consider the most basic notion of sparsity: a vector is *k-sparse* if at most k of its entries are non-zero. Then the projection problem consists in finding, for given $x \in \mathbb{R}^n$, a k -sparse $y \in \mathbb{R}^n$ such that $\|x - y\|_2^2$ is minimized. This problem is almost trivial—keeping the k largest entries of x and overwriting all other entries with 0 will do.

We study a slightly more elaborate sparsity model: the *Δ -separated sparsity model*. A vector y is *Δ -separated* if any two non-zero entries y_i, y_j satisfy $|i - j| \geq \Delta$ or $i = j$. We treat the corresponding projection problem: given $x \in \mathbb{R}^n$, k and Δ , find the k -sparse Δ -separated vector y minimizing $\|x - y\|_2^2$.

The Δ -separated sparsity model was introduced by Hegde, Duarte and Cevher [6] to capture neuronal spike trains in the framework of compressed sensing. They observed that the projection on the Δ -separated sparsity model can be formulated as an integer linear program over a totally unimodular constraint matrix. Interior point methods, then, allow the projection to be solved

in polynomial time.

Interior point methods, however, have superlinear asymptotic running times. In their survey paper, Hegde, Indyk and Schmidt [8] therefore ask whether there is an essentially linear time algorithm that finds a satisfying approximate solution (Open problem 6 in [8]).

Formally, they consider two kinds of approximation guarantees. Suppose A is an algorithm that computes, for each input vector x , a k -sparse Δ -separated vector $A(x)$. Since the overall goal is to minimize $\|x - A(x)\|_2^2$, we impose that $A(x)_i \in \{0, x_i\}$ for all x and all indices i .

We say that A has *tail approximation guarantee* α if

$$\|x - A(x)\|_2^2 \leq \alpha \|x - x^*\|_2^2 \text{ for all } x,$$

where x^* is an optimal solution. Moreover, we say that A has *head approximation guarantee* β if

$$\|A(x)\|_2^2 \geq \beta \|x^*\|_2^2 \text{ for all } x.$$

It is easy to see that these two notions are incomparable: one does not imply the other. As usual, we say that an algorithm is an α -*head approximation* (β -*tail approximation*) if it has head (tail) approximation guarantee α (or β).

We prove the following.

Theorem 1. *There is an algorithm running in $O(\epsilon^{-2}n)$ time that is both a $(1 - \epsilon)$ -head approximation and a $(1 + \epsilon)$ -tail approximation for the Δ -separated sparsity projection problem.*

We show, furthermore, how our algorithm fits in the *approximate model iterative hard thresholding* (AM-IHT) framework by Hegde, Indyk and Schmidt [7]. The resulting recovery algorithm for k -sparse Δ -separated signals performs significantly faster than the existing algorithm [6], at least for large input data.

Finally, we make the case that our algorithms are not only fast from a theoretical point of view but also in practice. To this end we have tested our algorithms on two types of randomly generated data. We present the results in Section 8.

While still a young field, compressed sensing already offers a rich literature. We refer Foucard and Rauhut [4] for an extensive introduction to the field. The sparsity model we study here, Δ -separated sparsity, is an example of a whole range of models, that fall under the heading of *structured sparsity*. Examples of this come up naturally in applications such as image processing, where wavelet coefficients of piecewise smooth functions approximately form a subtree of a given rooted tree; see [10, 2].

We finally remark that the Δ -separated sparsity projection problem may also be expressed in the language of graphs: indeed, an optimal solution corresponds to a maximum weight independent set of size at most k in a certain unit interval graph. The maximum weight independent set problem, without any restriction on the size of the set, admits a linear time algorithm on interval graphs [5]. Generalized cardinality restrictions on independent sets have been studied too [1, 9].

2 A combinatorial optimization formulation

Since $\|x\|_2^2 = \sum_{i=1}^n x_i^2$, we may study the non-negative vector obtained from the input vector x by taking the square of each entry instead. This yields the following, more convenient formulation of the Δ -separated sparsity model projection problem.

Separated Sparsity

Instance Positive integers k, Δ and a vector $x \in \mathbb{R}_{\geq 0}^n$ of non-negative reals.

Task Find $I \subseteq [n]$ such that $|I| \leq k$, such that $|i - j| \geq \Delta$ for all distinct $i, j \in I$ and such that $\sum_{i \in I} x_i$ is maximal.

Obviously, head and tail approximation could as naturally be formulated with respect to other norms. At least for ℓ_p -norms the methods we propose would still work.

We prove the following two statements.

Lemma 2. *There is a $(1 - \epsilon)$ -head approximation A for Separated Sparsity that runs in $O(\epsilon^{-2}n)$ time. That is, $\sum_{i \in A(x)} x_i \geq (1 - \epsilon) \sum_{j \in I^*} x_j$ where I^* is an optimal solution.*

Lemma 3. *There is a $(1 + \epsilon)$ -tail approximation A for Separated Sparsity that runs in $O(\epsilon^{-2}n)$ time. That is, $\sum_{i \notin A(x)} x_i \leq (1 + \epsilon) \sum_{j \notin I^*} x_j$ where I^* is an optimal solution.*

By running both of these algorithms and taking the solution with larger value $\sum_{i \in A(x)} x_i$ we obtain an algorithm that is both, a $(1 - \epsilon)$ -head and a $(1 + \epsilon)$ -tail approximation. This implies Theorem 1, our main result.¹

3 Preliminaries

We write $[n]$ for $\{1, \dots, n\}$. Let k, Δ be positive integers, let $x \in \mathbb{R}_{\geq 0}^n$, and consider an index subset $I \subseteq [n]$. We write $x(I)$ for $x(I) = \sum_{i \in I} x_i$, and \bar{I} for $[n] \setminus I$. We also write x_I for the vector $z \in \mathbb{R}^n$ defined by

$$z_i = \begin{cases} x_i & \text{if } i \in I \\ 0 & \text{otherwise} \end{cases}$$

Because it will be necessary in the recovery algorithm, we generalize the problem setting a bit. For an integer p , we say that a set $I \subseteq [n]$ is p -spikes Δ -separated if for every set $Z \subseteq [n]$ of $\Delta - 1$ consecutive integers it holds that $|I \cap Z| \leq p$. Any vector $x \in \mathbb{R}^n$ is p -spikes Δ -separated if its support (the set of indices with non-zero entries) is p -spikes Δ -separated. Note that x is 1-spike Δ -separated if and only if x is Δ -separated.

We also define the corresponding projection problem:

¹With a little bit of work, one may also prove that the tail approximation algorithm (Algorithm 4) actually achieves the same head approximation guarantee as our head approximation algorithm. This is because we use the head approximation as a subroutine in our tail approximation.

p -Spikes Separated Sparsity

Instance Positive integers k, Δ and a vector $x \in \mathbb{R}_{\geq 0}^n$ of non-negative reals.

Task Find a p -spikes Δ -separated set $I \subseteq [n]$ such that $|I| \leq k$ and such that $x(I)$ is maximal.

Indeed, we will often even consider a version of p -Spikes Separated Sparsity that is restricted to subsets of $[n]$. For a subset M of $[n]$ and any positive integer ℓ , we say that a set I is a *feasible solution* of $\text{SPRS}_x^p(M, \ell)$ if I is a p -spikes Δ -separated subset of M of size $|I| \leq \ell$. The set I is a solution if $x(I)$ is maximum among all feasible solutions of $\text{SPRS}_x^p(M, \ell)$. Note that whether I is a solution or not also depends on Δ . As Δ will always be clear from the context and will not vary within our arguments, we omit it from the notation. We may also drop x if x is clear from the context, and if we simply write $\text{SPRS}_x(M, \ell)$ rather than $\text{SPRS}_x^p(M, \ell)$ we mean the 1-spiked version of the problem, i.e., that $p = 1$.

For an optimal solution I of $\text{SPRS}_x^p(M, \ell)$ we define $\text{OPT}_x^p(M, \ell)$ as $x(I)$. Clearly, $\text{OPT}_x^p([n], k)$ is the value of the optimal solution of p -Spikes Separated Sparsity.

To simplify the presentation we assume that the basic arithmetic operations (addition/subtraction, multiplication/division and comparison) can be performed in a single time step each. We remark without proof that even if the size of the numbers in the input do play a role, our main algorithms still run in linear time.

4 Dynamic programming

Both the normal (1-spike) Separated Sparsity problem as well as its 2-spiked variant can be solved with dynamic programming. We start with one spike.

Let us write $\text{OPT}(M, \ell) = 0$ if $\ell \leq 0$ or if $M = \emptyset$. We observe that for all $i \in [n]$ and $\ell \in [k]$

$$\text{OPT}([i], \ell) = \max(x_i + \text{OPT}([i - \Delta], \ell - 1), \text{OPT}([i - 1], \ell)). \quad (1)$$

A feasible solution of $\text{SPRS}([i], \ell)$ may be computed at the same time by keeping track of whether the maximum in (1) is attained by picking i for the feasible solution or not.

We thus obtain the following lemma.

Lemma 4. *Given an instance $\text{SPRS}_x([n], k)$ with $x \in \mathbb{R}_{\geq 0}^n$, dynamic programming allows to compute all values $\text{OPT}([n], \ell)$ for $\ell = 1, \dots, k$, as well as the respective solutions, in running time $O(kn)$.*

We note that the algorithm is essentially a special case of the algorithm of Bandyapadhyay [1] for the budgeted maximum weight independent set problem on interval graphs.

As k cannot be considered a constant, but rather might even be linear in n , the dynamic programming algorithm does not run in linear time. Indeed, this is the whole point of this article: to improve on the running time of $O(kn)$. If n , however, is so large that it does not matter, i.e., if $k \geq \frac{n}{\Delta}$ then a simpler

dynamic programming approach, that drops the parameter ℓ , can be used to compute $\text{OPT}([n], k)$ in $O(n)$ -time.

We also describe a slightly trickier dynamic programming algorithm to solve **2-Spikes Separated Sparsity**.

Lemma 5. *Given an instance $\text{SPRS}_x^2([n], k)$ with $x \in \mathbb{R}_{\geq 0}^n$, dynamic programming allows to compute all values $\text{OPT}_x^2([n], \ell)$ for $\ell = 1, \dots, k$, as well as the respective solutions, in running time $O(k\Delta n)$.*

Proof. To see this, let $\text{OPT}^2([r], i, \ell)$ be the objective value of an optimal solution I of **2-Spike Separated Sparsity** on the vector $x_{[r]}$ such that $|I| \leq \ell$ and $|I \setminus [r - i]| \leq 1$. Here, $r \in [n]$, $i \in \{0, \dots, \Delta - 1\}$ and $\ell \in [k]$. Let us write $\text{OPT}^2([r], i, \ell) = 0$ if $r \leq 0$ or $\ell \leq 0$. We observe that $\text{OPT}^2([r], 0, \ell) = \text{OPT}^2([r], 1, \ell)$ for all $r \in [n]$ and $\ell \in [k]$, and that for $i \in \{1, \dots, \Delta - 1\}$ the value $\text{OPT}^2([r], i, \ell)$ can be computed as follows:

$$\text{OPT}^2([r], i, \ell) = \max(x_r + \text{OPT}^2([r - i], \Delta - i, \ell - 1), \text{OPT}^2([r - 1], i - 1, \ell))$$

This gives rise to an $O(k\Delta n)$ time dynamic programming algorithm that outputs the desired values $\text{OPT}^2([n], \ell) = \text{OPT}^2([n], 0, \ell)$, $\ell \in [k]$. The corresponding solutions can be recovered by standard techniques. \square

5 Head approximation

In this section we describe a head approximation for **p -Spikes Separated Sparsity**. As a subroutine it needs an exact algorithm which we denote A below. For instance, this could be a dynamic programming algorithm if $p = 1, 2$. As the exact algorithm is only called for smaller subinstances, the head approximation has a better asymptotic performance.

The main algorithm generates a number of (still large) subinstances that then are solved with another method (Algorithm 2) that we will treat afterwards.

Algorithm 1

Input An instance (n, x, Δ, k) of **p -Spikes Separated Sparsity** and $\epsilon > 0$.

Output A feasible solution.

- 1: Let $\lambda \in \mathbb{N}$ be the smallest integer such that $\lambda \geq \epsilon^{-1}$.
 - 2: **for** $\nu = 0, \dots, \lambda$ **do**
 - 3: Set $b = (\lambda + 1)\Delta$.
 - 4: Compute the set S_ν defined as $[n] \setminus \bigcup_{j=0}^{\infty} \{jb + \nu\Delta + 1, \dots, jb + \nu\Delta + \Delta\}$.
 - 5: Solve $\text{SPRS}^p(S_\nu, k)$ with Algorithm 2 optimally, and let I_ν be the obtained feasible solution.
 - 6: **end for**
 - 7: Return I_τ with $x(I_\tau) = \max_\nu x(I_\nu)$.
-

As we will use this twice, once for the head, and once for the tail approximation, we prove that at least one of the sets S_0, \dots, S_λ always covers almost all of the weight of any vector z .

Lemma 6. *For every $z \in \mathbb{R}_{\geq 0}^n$ there is some $\nu \in \{0, \dots, \lambda\}$ with*

$$z(S_\nu) \geq \frac{\lambda}{\lambda + 1} z([n]).$$

Proof. As every $i \in [n]$ appears in exactly λ of the sets S_ν it follows that

$$(\lambda + 1) \max_{\nu=0, \dots, \lambda} z(S_\nu) \geq \sum_{\nu=0}^{\lambda} z(S_\nu) = \lambda z([n]).$$

Choosing ν as the index that achieves the maximum yields the desired S_ν . \square

Building on the above lemma, we can prove the desired approximation guarantee. We assume here, and we prove it later in Lemma 8, that Algorithm 2 works correctly.

Lemma 7. *Algorithm 1 is a $(1 - \epsilon)$ -head approximation for p -Spikes Separated Sparsity.*

Proof. Let I^* be an optimal solution of $\text{SPRS}^p([n], k)$, and set $x^* = x_{I^*}$. Thus $x^*([n]) = \text{OPT}^p([n], k)$. Moreover, note that $I^* \cap S_\nu$ is a feasible solution of $\text{SPRS}^p(S_\nu, k)$ for every $\nu \in \{0, \dots, \lambda\}$. Thus $x^*(S_\nu) \leq \text{OPT}^p(S_\nu, k)$. With ν as in Lemma 6 we then obtain

$$\text{OPT}^p(S_\nu, k) \geq x^*(S_\nu) \geq \frac{\lambda}{\lambda + 1} x^*([n]) = \frac{\lambda}{\lambda + 1} \text{OPT}^p([n], k)$$

This and $\frac{\lambda}{\lambda + 1} \geq 1 - \epsilon$ finish the proof. \square

We remark that the analysis is tight: take $x = \mathbf{1}$, $k = n$, and $\Delta = 1$. One can also construct examples with larger Δ .

Let us now solve $\text{SPRS}^p(S_\nu, k)$ optimally. For a set $S \subseteq [n]$ of integers and an $x \in \mathbb{R}_{\geq 0}^n$ and integer Δ , call a subset $B \subseteq [n]$ a *block* of S if $B \cap S$ is non-empty and if B is a set of consecutive integers from $[n]$ that is minimal subject to the property that whenever $i \in B$ with $x_i \neq 0$ and $j \in S$ with $x_j \neq 0$ and $|i - j| < \Delta$ then $j \in B$.

Note that the blocks of S are disjoint, and that, more strongly, for two distinct blocks B_1, B_2 we have that $|i - j| \geq \Delta$ for any $i \in B_1$ and $j \in B_2$. Moreover observe that for every ν each block of S_ν as in Algorithm 1 has size at most $\lambda\Delta$.

Algorithm 2 that we need in order to complete the description of our head approximation algorithm calls, as a subroutine, an exact algorithm A that solves p -Spikes Separated Sparsity restricted to each block. In fact, we need an algorithm that computes $\text{OPT}^p(B, \ell)$ for each $0 \leq \ell \leq \left\lceil \frac{|B|}{\Delta} \right\rceil$, where B is a block. For $p = 1, 2$, this can be done, quite efficiently, with the dynamic programming algorithms as outlined in the previous section. However, any algorithm meeting the above requirements could be used in place of A .

Lemma 8. *Given an algorithm A that solves instances of p -Spikes Separated Sparsity optimally, Algorithm 2 solves $\text{SPRS}^p(S, k)$ optimally.*

While the lemma might sound a bit tautological at first reading, it actually makes sense: the point here is not that we can solve p -Spikes Separated Sparsity at all but that we only need to call an exact algorithm for small slices of the whole ground set, which then leads to a better performance.

Algorithm 2

Input An instance $\text{SPRS}^p(S, k)$ with maximal block size at most $\lambda\Delta$

Output An optimal solution of $\text{SPRS}^p(S, k)$

- 1: Compute the blocks B_1, \dots, B_s of S .
 - 2: For each block B_t compute $\lambda_t = \left\lceil \frac{|B_t|}{\Delta} \right\rceil$ and call algorithm A to compute $\text{OPT}^p(B_t, \ell)$ for each $\ell \in [\lambda_t]$.
 - 3: For all $t \in [s]$ and all $\ell \in [\lambda_t]$ set $q_\ell^t = \text{OPT}^p(B_t, \ell) - \text{OPT}^p(B_t, \ell - 1)$.
 - 4: Compute $Q \subseteq \{(t, \ell) : t \in [s], \ell \in [\lambda_t]\}$ with $|Q| \leq k$ such that $\sum_{(t, \ell) \in Q} q_\ell^t$ is maximal among all choices of Q .
 - 5: Delete each (t, ℓ) from Q for which $q_\ell^t = 0$.
 - 6: Return the set $T = \bigcup_{t=1}^s T_t$, where T_t is the optimal solution of $\text{SPRS}(B_t, j)$ and $j = |Q \cap \{(t, \ell) : \ell \in [\lambda_t]\}|$.
-

Proof of Lemma 8. We start by proving

$$q_\ell^t \leq q_{\ell-1}^t \text{ for all } t, \ell \text{ with } t \in [s] \text{ and } 2 \leq \ell \leq \lambda_t. \quad (2)$$

For this we show, equivalently, that

$$2 \cdot \text{OPT}^p(B_t, \ell - 1) \geq \text{OPT}^p(B_t, \ell - 2) + \text{OPT}^p(B_t, \ell) \quad (3)$$

Consider an optimal solution M of $\text{SPRS}^p(B_t, \ell - 2)$ and an optimal solution L of $\text{SPRS}^p(B_t, \ell)$. Consider the disjoint union $M \cup L$. That is, if some element appears in M and in L , we consider it to appear twice in $M \cup L$. In the natural order, choose for M' every second element of $M \cup L$, and let L' be the other elements. Obviously, $|M'| = |L'| \leq \ell - 1$, and neither M' nor L' contains an element twice.

Suppose that M' or L' is not a feasible solution of $\text{SPRS}^p(B_t, \ell - 1)$. By symmetry, we may assume this is the case for M' . Then there must be a set Z of Δ consecutive integers and $p + 1$ elements $i_1 < \dots < i_{p+1}$ of M' such that i_1, \dots, i_{p+1} all lie in Z . Since M' consists of every second element of $M \cup L$ there are thus $j_1, \dots, j_p \in L'$ such that $i_1 \leq j_1 \leq \dots \leq j_p \leq i_{p+1}$. In particular, all of these $2p + 1$ elements lie in Z . Then $p + 1$ of these must belong to either M or to L , which contradicts that M and L are p -spikes Δ -separated. This proves (2).

Next we prove

$$\sum_{(t, \ell) \in Q} q_\ell^t \geq \text{OPT}^p(S, k). \quad (4)$$

For this, consider a feasible solution I of $\text{SPRS}^p(S, k)$. Consider some t , and set $\ell_t = |I \cap B_t|$. Note that $\ell_t \leq \lambda_t$ since I is feasible. In turn, since $I \cap B_t$ is a feasible solution of $\text{SPRS}^p(B_t, \ell_t)$ it follows that $x(I \cap B_t) \leq \text{OPT}^p(B_t, \ell_t)$. We define a set $P \subseteq [s] \times [\lambda]$ with $|P| \leq k$ by including $(t, 1), \dots, (t, \ell_t)$ for every $t \in [s]$ in P . Then

$$\begin{aligned} x(I) &= \sum_{t=1}^s x(I \cap B_t) \leq \sum_{t=1}^s \text{OPT}^p(B_t, \ell_t) \\ &= \sum_{t=1}^s (q_1^t + \dots + q_{\ell_t}^t) = \sum_{(t, \ell) \in P} q_\ell^t \leq \sum_{(t, \ell) \in Q} q_\ell^t. \end{aligned}$$

This proves (4).

In view of (4), the proof of the lemma is finished once we prove that T is a feasible solution of $\text{SPRS}^p([n], k)$ and that $x(T) = \sum_{(t,\ell) \in Q} q_\ell^t$. To see that T is a feasible solution, we note that $|T| \leq k$ as $|Q| \leq k$, and that T is p -spikes Δ -separated, since restricted to each block B_t it is p -spikes Δ -separated and since indices from distinct blocks are at least Δ steps apart.

To determine $x(T)$, let $t \in [s]$ and $j_t = |Q \cap \{(t, \ell) : \ell \in [\lambda_t]\}|$. By (2) we may assume that $(t, 1), \dots, (t, j_t) \in Q$. Hence, $(t, m) \notin Q$ for every $m > j_t$. This implies $x(T_t) = q_1^t + \dots + q_{j_t}^t$. Hence,

$$x(T) = \sum_{t=1}^s (q_1^t + \dots + q_{j_t}^t) = \sum_{(t,\ell) \in Q} q_\ell^t,$$

which completes the proof. \square

Lemma 9. *Let $f(r)$ be an upper bound on the running time of algorithm A when run on a vector of dimension r . Then Algorithm 2 can be implemented to run in time*

$$O\left(n \max_{r \in [\lambda\Delta]} 1 + \frac{f(r)}{r}\right).$$

Proof. For each block B_t , Algorithm 2 uses algorithm A as a subroutine to compute $\text{OPT}(B_t, \ell)$ for $0 \leq \ell \leq \lambda_t$. This can be done in time $f(|B_t|)$.

It remains to discuss the complexity of finding the k largest elements in the vector $P = (q_1^1, \dots, q_{\lambda_s}^s)$. This can be done in $O(n)$ time using order statistics. To see this, note that the vector P is of dimension $\sum_{t=1}^s \lambda_t = O(n)$.

We now find the k -th largest element \hat{q} of P in $O(n)$ time using, for example, the *Introselect* algorithm [11]. Then we collect the elements of P of value larger than \hat{q} and put them into our feasible solution. Finally, we add elements of value equal to \hat{q} until our feasible solution contains k elements in total. We obtain a running time of $O(n)$ for this step. The computation of the T_t and T are straightforward and can be done in $O(n)$ time.

We obtain a total running time of at most

$$O\left(\max\left\{\sum_{t=1}^s f(b_t) : s \in \left[\left\lceil \frac{n}{\Delta} \right\rceil\right], b \in [\lambda\Delta]^s, \sum_{t=1}^s b_t \leq n - \Delta(s-1)\right\} + n\right) \quad (5)$$

where b_t denotes the size of the t -th block B_t and we used the observation that if there are s blocks there are at least $n - \Delta(s-1)$ elements of $[n]$ not contained in a block. Introducing the running time per element $\frac{f(r)}{r}$ we may simplify (5) to

$$O\left(n \max_{r \in [\lambda\Delta]} 1 + \frac{f(r)}{r}\right),$$

and hence the proof is complete. \square

We can now prove Lemma 2.

Proof of Lemma 2. Lemma 7 says that Algorithm 1 is a $(1 - \epsilon)$ -approximation.

To compute the running time, first observe that, by Lemma 4, dynamic programming in place of algorithm A has a running time of $f(r) = O(r \cdot \frac{r}{\Delta})$, and thus Lemma 9 yields a running time of

$$O\left(n \max_{r \in [\lambda\Delta]} 1 + \frac{f(r)}{r}\right) = O\left(n \max_{r \in [\lambda\Delta]} 1 + \frac{r}{\Delta}\right) = O(\lambda n)$$

for Algorithm 2.

Consequently, we obtain a running time of $O(\lambda^2 n)$ for Algorithm 1. Then, $O(\lambda^2 n) = O(\epsilon^{-2} n)$, by the choice of λ , and thus the proof is complete. \square

Theorem 10. *For 2-Spikes Separated Sparsity, Algorithm 1 is a $(1 - \epsilon)$ -head approximation that can be implemented to run in $O(\epsilon^{-2} \Delta n)$ time.*

Proof. The proof is the same as that of Lemma 2, with the single exception that we use Lemma 5 rather than Lemma 4. \square

Algorithm 1 can also speed up the solution of p -Spikes Separated Sparsity for larger p , provided we have access to some exact algorithm. Especially for large p , the dynamic programming approach does not seem to be feasible anymore when n grows large. We may, however, encode the problem as an integer program with a totally unimodular constraint matrix, as has been demonstrated by Hegde et al. [6], and then solve the linear relaxation. The resulting speed-up can be directly read off of Lemma 9.

6 Tail approximation

In this section we restrict our attention to the 1-spike case exclusively. Recall that an algorithm A for Separated Sparsity has tail approximation guarantee α if

$$x(\overline{A(x)}) \leq \alpha x(\overline{I^*}) \text{ for all } x,$$

where I^* is an optimal solution.

Algorithm 1 has no constant tail approximation guarantee. To see that, note that if x is the all-ones vector, $k = n$, and $\Delta = 1$, the optimal solution I^* is $[n]$. Consequently, $x(\overline{I^*}) = 0$ and thus Algorithm 1 would need to solve the instance optimally to have a constant tail approximation guarantee. This, however, is not the case as $[n] \not\subseteq S_\nu$ for all ν provided that n is sufficiently large.

There is, however, a very simple algorithm that has a constant tail approximation guarantee: simply pick the best feasible solution among the k largest elements of x (Algorithm 3). We briefly discuss the algorithm because it is so simple.

Algorithm 3

Input An instance $\text{SPRS}([n], k)$

Output A feasible solution to $\text{SPRS}([n], k)$

- 1: Compute index set L of k largest elements in x .
 - 2: Solve $\text{SPRS}(L, k)$ optimally, and output the solution I .
-

Proposition 11. *Algorithm 3 has tail approximation guarantee 2.*

We note that, with a bit of care, the algorithm can be implemented to run in $O(n)$ time.

Proof. Let I^* be an optimal solution, and set $x_{I^*} = x^*$. Then $x(I^*) \leq x(L)$ and thus

$$x(\bar{I}^*) \geq x(\bar{L}). \quad (6)$$

Moreover, $x(I^* \cap L) \leq x(I)$ since I is an optimal solution of $\text{SPRS}(L, k)$, while $I^* \cap L$ is some feasible solution of $\text{SPRS}(L, k)$. Thus $x(L \setminus (I^* \cap L)) \geq x(L \setminus I)$ and

$$x(\bar{I}^*) \geq x(L \setminus (I^* \cap L)) \geq x(L \setminus I). \quad (7)$$

As $x(\bar{I}) = x(\bar{L}) + x(L \setminus I)$ we obtain with (6) and (7) that $x(\bar{I}) \leq 2x(\bar{I}^*)$. \square

Unfortunately, the analysis is tight. Indeed, consider the instance $x = (1, 1, 1)$, $\Delta = 2$ und $k = 2$. Then the algorithm might find $L = \{1, 2\}$, which might result in $I = \{1\}$. Then $x(\bar{I}) = 2$ but $I^* = \{1, 3\}$ and $x(\bar{I}^*) = 1$.

We now come to our tail approximation of arbitrary precision. Before stating the algorithm, we need two technical lemmas.

Lemma 12. *Let n be a positive integer, $W \subseteq [n]$, $A, B \subseteq W$, and $x \in \mathbb{R}_{\geq 0}^n$. Assume that*

$$x(A) \geq \alpha x(B) \text{ and } x(B) \leq \mu x(W)$$

for some reals $0 < \alpha, \mu < 1$. Then

$$x(W \setminus A) \leq \frac{\alpha}{1 - \frac{1-\alpha}{1-\mu\alpha}} \cdot x(W \setminus B).$$

Proof. Put $\gamma = \frac{1-\alpha}{1-\mu\alpha}$ and observe that $0 < \gamma < 1$. Now

$$\begin{aligned} x(W \setminus B) &= x(W) - (1 - \gamma)x(B) - \gamma x(B) \\ &\geq x(W) - \frac{1-\gamma}{\alpha}x(A) - \gamma\mu x(W) \\ &= \frac{1-\gamma}{\alpha}x(W \setminus A) + \left(1 - \frac{1-\gamma}{\alpha}\right)x(W) - \gamma\mu x(W) \\ &= \frac{1-\gamma}{\alpha}x(W \setminus A), \end{aligned}$$

since

$$1 - \frac{1-\gamma}{\alpha} - \gamma\mu = \frac{1}{\alpha}(\alpha - 1 + \gamma - \alpha\mu\gamma) = \frac{1}{\alpha}(\alpha - 1 + \gamma(1 - \mu\alpha)) = 0.$$

\square

Given a vector $x \in \mathbb{R}^n$ we define the *tail vector* $t \in \mathbb{R}^n$ of x by setting

$$t_i = x(\{j \in [n] : j \neq i, |i - j| < \Delta\}) = \sum_{j=i-\Delta+1}^{i+\Delta-1} x_j - x_i$$

for all $i \in [n]$. We call an index $i \in [n]$ *strong* if $x_i > t_i$ and *weak* otherwise. It is clear that for any two strong indices i, j we have $|i - j| \geq \Delta$.

For a given $x \in \mathbb{R}^n$ let S be the set of its strong indices. We define the *reduced vector* $r \in \mathbb{R}^n$ of x by setting

$$r_i = \begin{cases} x_i & \text{if } i \in S \text{ or } |i - j| \geq \Delta \text{ for all } j \in S \\ 0 & \text{otherwise} \end{cases}$$

for all $i \in [n]$.

Lemma 13. *Consider an instance $\text{SPRS}_x([n], k)$, let S be the set of strong indices of x , and let r be the reduced vector of x . Let I^* be an optimal solution of $\text{SPRS}_x([n], k)$. If $i \in I^* \setminus S$ and $s \in S$ then $|s - i| \geq \Delta$. In particular, $\text{OPT}([n], k) = x(I^*) = r(I^*)$.*

Proof. Suppose that there are $i \in I^* \setminus S$ and $s \in S$ with $|s - i| < \Delta$. Set $C = \{j \in [n] : |s - j| < \Delta, s \neq j\}$, and consider $I = (I^* \setminus C) \cup \{s\}$. Then I is a feasible solution of $\text{SPRS}([n], k)$, and as s is strong, it holds that $x(C \cap I^*) \leq x(C) < x_s$, which implies $x(I) > x(I^*)$. Obviously, this is impossible because I^* is an optimal solution. \square

We now present our tail approximation for **Separated Sparsity**. For this, fix $\epsilon > 0$.

Algorithm 4

Input An instance (n, x, Δ, k) of **Separated Sparsity** and $\epsilon > 0$

Output A feasible solution to $\text{SPRS}([n], k)$

- 1: Compute the tail vector t and the set S of all strong indices.
 - 2: Compute the reduced vector r of x .
 - 3: Run Algorithm 1 (for $p = 1$ spike) on r with precision $1 - \frac{\epsilon}{2}$, subject to the following modification: Instead of calling Algorithm 2 on S_ν , call it on $S_\nu \cup S$ for $\nu = 0, \dots, \lambda$.
-

Lemma 14. *Algorithm 4 returns a $(1 + \epsilon)$ -tail approximation of the **Separated Sparsity** problem.*

Proof. First note that the call of Algorithm 2 is valid, as the maximal block size of $S \cup S_\nu$ with respect to r is at most $\lambda\Delta$: indeed, the elements of S form singleton blocks, while every other block is contained in a block of S_ν with respect to x , and thus has block size at most $\lambda\Delta$.

Let I^* be an optimal solution, and let $W = \bar{S}$ be the set of weak indices of x . Applying Lemma 6 to $z = x_{I^* \cap W}$, we find a ν with

$$x(S_\nu \cap I^* \cap W) = z(S_\nu) \geq (1 - \frac{\epsilon}{2})z([n]) = (1 - \frac{\epsilon}{2})x(I^* \cap W). \quad (8)$$

Consider the feasible solution $L = (S \cup S_\nu) \cap I^*$ of $\text{SPRS}_r(S \cup S_\nu, k)$. Then

$$\text{OPT}_x(S \cup S_\nu, k) \geq \text{OPT}_r(S \cup S_\nu, k) \geq r(L) = x(L), \quad (9)$$

where the last equality is due to Lemma 13.

Next, we claim that

$$x(W \setminus L) \leq (1 + \epsilon)x(W \setminus I^*). \quad (10)$$

Before we prove the claim, we note how to finish the proof of the lemma with it. For this, observe that the choice of L implies $x(S \setminus L) = x(S \setminus I^*)$. Then

$$\begin{aligned} x(\bar{L}) &= x(S \setminus L) + x(W \setminus L) \\ &\leq x(S \setminus I^*) + (1 + \epsilon)x(W \setminus I^*) \\ &\leq (1 + \epsilon)x(\bar{I}^*). \end{aligned}$$

If I is the output of Algorithm 4 then this implies

$$x(\bar{I}) \leq x([n]) - \text{OPT}_x(S \cup S_\nu, k) \stackrel{(9)}{\leq} x([n]) - x(L) = x(\bar{L}) \leq (1 + \epsilon)x(\bar{I}^*),$$

which is the statement of the lemma.

Let us prove (10). As a consequence of Lemma 13, we get

$$2x(W \setminus I^*) \geq t(I^* \cap W).$$

This, in turn, implies

$$\begin{aligned} x(W) &= x(I^* \cap W) + x(W \setminus I^*) \\ &\geq x(I^* \cap W) + \frac{1}{2}t(I^* \cap W) \\ &\geq x(I^* \cap W) + \frac{1}{2}x(I^* \cap W) = \frac{3}{2}x(I^* \cap W), \end{aligned}$$

where the second inequality is because of what it means for an index to be weak.

Set $A = S_\nu \cap I^* \cap W$ and $B = I^* \cap W$. We may apply Lemma 12 with $\mu = \frac{2}{3}$, by the preceding inequality, and $\alpha = 1 - \frac{\epsilon}{2}$, by (8), and obtain

$$x(W \setminus L) = x(W \setminus A) \leq \frac{\alpha}{1 - \frac{1-\alpha}{1-\mu\alpha}} \cdot x(W \setminus B) = \frac{\alpha}{1 - \frac{1-\alpha}{1-\mu\alpha}} \cdot x(W \setminus I^*).$$

As

$$\frac{\alpha}{1 - \frac{1-\alpha}{1-\mu\alpha}} = \frac{1 - \frac{\epsilon}{2}}{1 - \frac{\frac{5}{2}}{1 - \frac{2}{3}(1 - \frac{\epsilon}{2})}} = \frac{1 - \frac{\epsilon}{2}}{1 - \frac{3\epsilon}{2+2\epsilon}} = \frac{1 - \frac{\epsilon}{2}}{\frac{2-\epsilon}{2+2\epsilon}} = \frac{2 + \epsilon - \epsilon^2}{2 - \epsilon} = 1 + \epsilon,$$

we have proved (10) and thus the lemma. \square

It remains to observe that the algorithm runs in linear time.

Lemma 15. *Algorithm 4 can be implemented to run in $O(\epsilon^{-2}n)$ time.*

Proof. First note that we can compute t in $O(n)$ time using dynamic programming since

$$t_i = t_{i-1} + x_{i-1} - x_i + x_{i+\Delta-1} - x_{i-\Delta},$$

where we assume that all undefined values are 0. From t we can clearly compute S and then r in $O(n)$ time.

Note that the maximum block size of $S_\nu \cup S$ considered by Algorithm 2 is $\lambda\Delta$, as the definition of r and the blocks implies that every element of S is placed in a private block. Hence, we may apply Lemma 9 as in the proof of Lemma 2 to complete the proof. \square

The above two lemmas prove Lemma 3.

7 Recovering k -sparse Δ -separated signals

The recovery problem for k -sparse Δ -separated signals $x \in \mathbb{R}^n$ is as follows. For a sensing matrix A , we are given the set of noisy measurements, a vector $y \in \mathbb{R}^m$, which relates to x via $y = Ax + e$ for some noise vector $e \in \mathbb{R}^m$. The task consists in recovering the original signal x , or a vector close to x . Hegde, Duarte and Cevher [6] have shown that this is possible if A is an i.i.d. subgaussian matrix and $m \geq C \log(\frac{n}{k} - \Delta)$ for some constant C .

To do so, Hegde et al. prove that such a matrix enjoys some form of the *restricted isometry property* (RIP). In general, a matrix $A \in \mathbb{R}^{m \times n}$ is said to have the RIP with constant δ if every k -sparse vector $x \in \mathbb{R}^n$ satisfies

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2. \quad (11)$$

In our context, it is enough if (11) is satisfied for all k -sparse Δ -separated vectors x . To formalize this, let us say that a matrix A has the (k, Δ, p) -RIP with constant δ if (11) holds for every k -sparse p -spike Δ -separated vector $x \in \mathbb{R}^n$.

Hegde, Duarte and Cevher prove the following bound for the $(k, \Delta, 1)$ -RIP in Theorem 16 and explain how it carries over to the $(k, \Delta, 2)$ -RIP. Using the same reasoning, it does in fact extend to the (k, Δ, p) -RIP for every fixed p .

Theorem 16 (Hegde, Duarte and Cevher [6]). *Fix $p \in \mathbb{N}$. There is a constant C such that, for all $\delta > 0$, any $t > 0$ and any*

$$m \geq C \left(\delta^{-2} \left(k \log \left(\frac{n}{k} - \Delta \right) + t - k \ln \delta \right) \right) \quad (12)$$

an $m \times n$ i.i.d. subgaussian random matrix has the (k, Δ, p) -RIP with constant δ with probability at least $1 - \epsilon^{-t}$.

Using the *Model-based CoSaMp* framework, the authors develop a recovery algorithm using $(k, \Delta, 2)$ -RIP matrices and linear programming to solve both *Separated Sparsity* and *2-Spike Separated Sparsity* exactly. Indeed, they give an integer programming formulation of the above-mentioned problems and use total unimodularity of the restriction matrix to argue that solving the linear programming formulation are sufficient. They obtain an algorithm that converges geometrically in the sense that, after $O(\log \frac{\|x\|_2}{\|e\|_2})$ iterations, a k -sparse Δ -separated vector \hat{x} is found with

$$\|x - \hat{x}\|_2 \leq C\|e\|_2 \text{ for some constant } C, \quad (13)$$

where x is the vector to be recovered. The bound (13) is generally considered to be a successful recovery of x .

If one is bound to use head and tail approximations rather than exact projection algorithms one can still recover the original signal efficiently. We follow here the *approximate model iterative hard thresholding* (AM-IHT) method, a framework proposed by Hegde, Indyk and Schmidt [7]. It resembles the iterative hard thresholding algorithm by Blumensath and Davies [3].

In the statement of the algorithm below, T denotes the tail approximation according to Theorem 1 and H denotes the head approximation for the 2-spike case (Theorem 10). We assume that both algorithms run with a fixed precision to be determined later.

Algorithm 5

Input A, y, e with $Ax = y + e$ and $i \in \mathbb{N}$

Output $\text{AM-IHT}(y, A, i)$

- 1: $x^0 \leftarrow 0$
 - 2: **for** j from 0 to $i - 1$ **do**
 - 3: $x^{j+1} \leftarrow T(x^j + H(A^T(y - Ax^j)))$
 - 4: **end for**
 - 5: Return $\text{AM-IHT}(y, A, i) \leftarrow x^i$
-

As Hegde et al. prove, AM-IHT is a rapidly converging recovery algorithm if the approximation algorithms used run with sufficient precision. To keep it simple, we do not state their theorem in full generality, but rather customize it for the Δ -separated sparsity model.

Theorem 17 (Hegde, Indyk and Schmidt [7]). *Assume that A has the $(\Delta, k, 4)$ -RIP with constant δ , that c_T is the approximation guarantee of our tail approximation for the 1-spike case and that c_H is the guarantee of our head approximation for the 2-spike case. Writing $r^j = x - x^j$, we have*

$$\begin{aligned} \|r^{i+1}\|_2 &\leq (1 + c_T) \left(\frac{\sqrt{1 - c_H^2}(1 + \delta) + \delta}{c_H} + 2\delta \right) \|r^i\|_2 \\ &\quad + (1 + c_T)\sqrt{1 + \delta} \left(\frac{\sqrt{1 - c_H^2} + 1}{c_H} + 4 \right) \|e\|_2. \end{aligned}$$

Theorem 17 warrants geometric convergence if c_T and c_H are close enough to 1 and δ is sufficiently small. For example, if we run the head approximation for the 2-spike case and the tail approximation for the 1-spike case with $\epsilon = 0.01$ as well as Theorem 16 with $\delta = 0.01$, Theorem 17 yields

$$\|x - \text{AM-IHT}(y, A, i)\|_2 \leq 0.35^i \|x\|_2 + 16\|e\|_2,$$

where $\text{AM-IHT}(y, A, i)$ denotes the vector obtained after i iterations of the AM-IHT procedure. Consequently, after $O(\log \frac{\|x\|_2}{\|e\|_2})$ iterations we arrive at a k -sparse Δ -separated vector \hat{x} satisfying (13), as desired.

Let Λ be the time needed to multiply a vector with an i.i.d. subgaussian $m \times n$ matrix where m is given by (12). Note that $\Lambda = O(nk \log(\frac{n}{k} - \Delta))$. For a fixed precision, the tail approximation runs in $O(n)$ time according to Theorem 1 and the head approximation runs in $O(n\Delta)$ time according to Theorem 10. In view of Algorithm 5, we obtain the following.

Corollary 18. *The AM-IHT algorithm can be implemented to run in $O(n\Delta + \Lambda)$ time steps per iteration.*

This beats the algorithm by Hegde, Duarte and Cevher [6] which needs $O(n^{3.5}L^2 + \Lambda)$ time per iteration where L is the number of bits of the input. The bottleneck in the running time of their algorithm is to solve the linear program formulations of Separated Sparsity and 2-Spike Separated Sparsity exactly using an interior point method. If, instead, we plug in the dynamic programs according to Lemma 4 and 5, the algorithm needs $O(nk\Delta + \Lambda)$ time per iteration. We remark

that their algorithm was not necessarily designed to obtain a fast asymptotic running time and that it might perform much better compared to ours on smaller instances.

8 Experiments

We have implemented the main algorithms of this paper and investigated running times and solution qualities for different random instances. The code, which is available online², is written in Python, and runs on a standard desktop computer. We have deviated in one point from the algorithms as described here: to pick the k largest elements in Algorithm 2 we have used a heap-based library method of Python rather than `Introspect`.




Figure 1: Examples of a uniform and a Poisson process instance with expected arrival time 4

We have tested the algorithms on two types of random instances. In a *uniform instance*, we have picked $x \in \mathbb{R}^n$ in such a way that each entry x_i receives uniformly at random a value in the interval $[0, 1)$. The second type are *Poisson instances*, where a Poisson process is run to identify a number of positions $I \subseteq [n]$, the *spikes* of the signal, where $x_i, i \in I$, receives a value chosen uniformly at random from the interval $[0, 1)$; all entries x_j outside I have value 0. The expected arrival time of the process, i.e., the expected difference between consecutive spikes, is set to the same Δ as in the instance $\text{SPRS}_x([n], k)$, or to $\frac{1}{2}\Delta$ when we treat 2-Spikes Separated Sparsity. Two example instances are shown in Figure 1. Each data point in the Figures 2–6 is the mean over 100 runs with the same parameters.

In Figure 2, we have plotted the running times of the dynamic programming algorithm, as well as Algorithms 1 and Algorithm 4 with $\lambda = 2, 3$ each. As the actual instance should not have a significant impact on the running times, we have only tested uniformly generated instances. The left diagram shows the running times for varying n and $\Delta = k = \lfloor \frac{1}{2}\sqrt{n} \rfloor$. Clearly, Algorithms 1 and Algorithm 4 perform much better than the dynamic programming algorithm. This is not so pronounced in the right diagram, where we have fixed Δ to 40, while k is set to $\lceil \log_2 n \rceil$. As the asymptotic running time of dynamic programming is $O(kn)$, it is not at all surprising that dynamic programming fares better when k is small.

Figure 3 shows head and tail approximation guarantees of Algorithm 1 for different λ , when applied to uniform instances of size $n = 1000$. As arguably, the performance does not depend that much on the actual value of Δ compared to n but rather on $k\Delta$ compared to n , we have kept Δ fixed to 20 and only varied

²available at <http://www.uni-ulm.de/bruhn>




Figure 2: Comparison of running times of dynamic programming and Algorithm 1 (HS) and Algorithm 4 (TS), means of 100 repeats per data point. Left: $\Delta = \lfloor \frac{1}{2}\sqrt{n} \rfloor = k$. Right: $\Delta = 40, k = \lfloor \log_2(n) \rfloor$.




Figure 3: Head and tail approximation in per cent, means of 100 repeats per data point. Uniform instance with $\Delta = 20$ and $n = 1000$; only Algorithm 1 (HS) shown as the results for Algorithm 4 are nearly identical

k . In a uniform instance the expected number of strong indices is vanishingly small. Clearly, if there are no strong indices then Algorithm 4 defaults to Algorithm 1, and indeed our test showed no difference between the two in this setting. Consequently, we have omitted the results of Algorithm 4 in Figure 3. As can be seen, the algorithm exhibits very good tail approximation even for $\lambda = 1$ and reasonably good head approximation at least for $\lambda = 2, 3$.

Algorithms 1 and 4 only differ when there are a (substantial) number of strong indices. Obviously, Poisson instances are set up precisely in this way as




Figure 4: Head and tail approximation in per cent, means of 100 repeats per data point. Poisson process instance with expected arrival time $\Delta = 20$ and $n = 1000$.

to yield a good number of strong indices. Not surprisingly, Algorithm 4 handles these instances better than Algorithm 1, as can be seen in Figure 4. Algorithm 4 shows very good head and tail approximation, even for $\lambda = 2$, while both are much poorer for Algorithm 1.




Figure 5: Comparison of running times of dynamic programming and Algorithm 1 (HS) for two spikes, means of 100 repeats per data point. Left: $\Delta = \lfloor \frac{1}{2}\sqrt{n} \rfloor = k$. Right: $\Delta = 40$, $k = \lfloor \log_2(n) \rfloor$.

The final two figures contain the test results for 2-Spikes Separated Sparsity. In Figure 5, the running time of the dynamic programming algorithm is compared to the running time of the algorithm of Theorem 10 for two different precisions. On the left, as in Figure 2, we have set $k = \Delta = \lfloor \frac{1}{2}\sqrt{n} \rfloor$ and on

the right we have set $k = \lfloor \log_2 n \rfloor$ and $\Delta = 40$. The running times behave in a similar way as in Figure 2: the approximation algorithm is substantially faster, especially when $k = \lfloor \frac{1}{2}\sqrt{n} \rfloor$. Since we only prove a constant head approximation guarantee in Theorem 10, we show only the head approximation in Figure 3, again for uniform as well as Poisson instances. (There is a second reason: in the Poisson instance, even for expected arrival times of $\frac{1}{2}\Delta$, occasionally the instance is already 2-spikes Δ -separated, which means that the optimal tail value is 0; then, however, we would need to divide by 0 in order to compute the tail approximation factor.) Again, we see very good head approximation in the case of uniform instances, and reasonable head approximation in the case of Poisson instances, at least for $\lambda = 3$.




Figure 6: Head approximation for two spikes, in per cent, means of 100 repeats per data point. Left: Uniform instance with $\Delta = 20$ and $n = 1000$. Right: Poisson process instance with expected arrival time $\Delta = 20$ and $n = 1000$.

9 Future work

We presented linear time head and tail approximations of arbitrary precision for the Δ -separated model projection problem. In principle, one might hope for a linear time algorithm to solve the problem to optimality. We suspect, however, that this is not possible. Perhaps a conditional lower bound on the running time could be obtained?

Another question for future work is whether our results can be extended to yield fast algorithms for more general model projection problems. Here one might start by investigating problems that have a totally unimodular restriction matrix and models that correspond to very limited graph classes like interval graphs.

References

- [1] Sayan Bandyopadhyay, *A variant of the maximum weight independent set problem*, CoRR **abs/1409.0173** (2014).
- [2] Richard G. Baraniuk, Volkan Cevher, Marco F. Duarte, and Chinmay Hegde, *Model-based compressive sensing*, IEEE Trans. Information Theory **56** (2010), no. 4, 1982–2001.
- [3] Thomas Blumensath and Mike E. Davies, *Iterative hard thresholding for compressed sensing*, Applied and Computational Harmonic Analysis **27** (2009), no. 3, 265–274.
- [4] Simon Foucart and Holger Rauhut, *A mathematical introduction to compressive sensing*, Applied and Numerical Harmonic Analysis, Birkhäuser, 2013.
- [5] András Frank, *Some polynomial algorithms for certain graphs and hypergraphs*, 1976, In Proc. 5th Br. comb. Conf.
- [6] Chinmay Hegde, Marco F. Duarte, and Volkan Cevher, *Compressive sensing recovery of spike trains using a structured sparsity model*, 2009, In SPARS'09 - Signal Processing with Adaptive Sparse Structured Representations.
- [7] Chinmay Hegde, Piotr Indyk, and Ludwig Schmidt, *Approximation-tolerant model-based compressive sensing*, Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, 2014, pp. 1544–1561.
- [8] ———, *Fast algorithms for structured sparsity*, Bulletin of the EATCS **117** (2015).
- [9] Tushar Kalra, Rogers Mathew, Sudebkumar Prasant Pal, and Vijay Pandey, *Maximum weighted independent sets with a budget*, Algorithms and Discrete Applied Mathematics - Third International Conference, CALDAM 2017, Sancoale, Goa, India, February 16-18, 2017, Proceedings, 2017, pp. 254–266.
- [10] Stéphane Mallat, *A wavelet tour of signal processing, 2nd edition*, Academic Press, 1999.
- [11] David R. Musser, *Introspective sorting and selection algorithms*, Softw., Pract. Exper. **27** (1997), no. 8, 983–993.
- [12] D. Needell and J.A. Tropp, *CoSaMP: Iterative signal recovery from incomplete and inaccurate samples*, Applied and Computational Harmonic Analysis **26** (2009), no. 3, 301–321.