
A Fast Algorithm for Separated Sparsity via Perturbed Lagrangians

Aleksander Mađry*
MIT
madry@mit.edu

Slobodan Mitrović*
EPFL
slobodan.mitrovic@epfl.ch

Ludwig Schmidt*
MIT
ludwigs@mit.edu

Abstract

Sparsity-based methods are widely used in machine learning, statistics, and signal processing. There is now a rich class of structured sparsity approaches that expand the modeling power of the sparsity paradigm and incorporate constraints such as group sparsity, graph sparsity, or hierarchical sparsity. While these sparsity models offer improved sample complexity and better interpretability, the improvements come at a computational cost: it is often challenging to optimize over the (non-convex) constraint sets that capture various sparsity structures. In this paper, we make progress in this direction in the context of separated sparsity – a fundamental sparsity notion that captures exclusion constraints in linearly ordered data such as time series. While prior algorithms for computing a projection onto this constraint set required quadratic time, we provide a perturbed Lagrangian relaxation approach that computes provably exact projection in only nearly-linear time. Although the sparsity constraint is non-convex, our perturbed Lagrangian approach is still guaranteed to find a globally optimal solution. In experiments, our new algorithms offer a $10\times$ speed-up already on moderately-size inputs.

1 Introduction

Over the past two decades, sparsity has become a widely used tool in several fields including signal processing, statistics, and machine learning. In many cases, sparsity is the key concept that enables us to capture important structure present in real-world data while making the resulting problem computationally tractable and suitable for mathemati-

*Authors ordered alphabetically.

cal analysis. Among the many applications of sparsity are sparse linear regression, compressed sensing, sparse PCA, and dictionary learning.

The first wave of sparsity-based techniques focused on the standard notion of sparsity that only constrains the number of non-zeros. Over time, it became apparent that extending the notion of sparsity to encompass more complex structures present in real-world data can offer significant benefits. Specifically, utilizing such additional structure often improves the statistical efficiency in estimation problems and the interpretability of the final result. There is now a large body of work on structured sparsity that has introduced popular models such as group sparsity and hierarchical sparsity [1, 2, 3, 4, 5, 6, 7, 8, 9]. These statistical improvements, however, come at a computational cost: the resulting optimization problems are often much harder to solve. The key reason is that the combinatorial sparsity structures give rise to non-convex constraints. Consequently, many of the resulting algorithms have significantly worse running time than their “standard sparsity” counterparts. This trade-off raises an important question: can we design algorithms for structured sparsity that match the time complexity of commonly used algorithms for standard sparsity?

In this paper, we address this question in the context of the *separated sparsity* model, a popular sparsity model for data with a known minimum distance between large coefficients [10, 11, 12, 13, 14]. In the one-dimensional case, such as time series data, neuronal spike trains are a natural example. Here, a minimum refractory period ensures separation between consecutive spikes. In two dimensions, separation constraints arise in the context of astronomical images or super-resolution applications [12, 15].

We introduce new algorithms for separated sparsity that run in *nearly-linear* time. This significantly improves over prior work that required at least quadratic time, which is quickly prohibitive for large data sets. An important consequence of our fast running time is that it enables methods that utilize separated sparsity yet are essentially as fast as their counterparts based on standard sparsity only. For instance, when we instantiate our algorithm in compressive sensing, the running time of our method matches that of common methods such as IHT or CoSaMP.

Our algorithms stem from a primal-dual linear programming (LP) perspective on the problem. Our theoretical findings reveal a rich structure behind the separated sparsity model, which we utilize to obtain efficient methods. Interestingly, our final algorithm has a very simple form that can be interpreted as a *Lagrangian relaxation* of the sparsity constraint. In spite of the non-convexity of the constraint, our algorithm is still guaranteed to find the globally optimal solution.

We also show that these algorithmic and theoretical contributions directly translate into empirical efficiency. Specifically, we demonstrate that, compared to the state of the art procedures, our methods yield an order of magnitude speed-up already on moderate-size inputs. We run experiments on synthetic data and real world neuronal spike train signals.

2 Separated sparsity and applications

In this section, we formally define separated sparsity and the corresponding algorithmic problems.

First, we briefly introduce our notation. As usual, $[d]$ denotes the set $\{1, \dots, d\}$. We say that a vector $\theta \in \mathbb{R}^d$ is k -sparse if θ contains at most k non-zero coefficients. We define the support of θ as the set of indices corresponding to non-zero coefficients, i.e., $\text{supp}(\theta) = \{i \in [d] \mid \theta_i \neq 0\}$. We let $\|\theta\|$ denote the ℓ_2 -norm of a vector $\theta \in \mathbb{R}^d$.

Separated sparsity. For a support $\Omega \subseteq [d]$, let $\text{sep}(\Omega) = \min_{i \neq j \in \Omega} |i - j|$ be the minimum separation of two indices in the support. We define the following two sets of supports: set $\mathbb{M}_\Delta = \{\Omega \subseteq [d] \mid \text{sep}(\Omega) \geq \Delta\}$, and Δ -separated sparsity supports $\mathbb{M}_{k,\Delta} = \{\Omega \in \mathbb{M}_\Delta \mid |\Omega| = k\}$. That is, $\mathbb{M}_{k,\Delta}$ is the set of support patterns containing k non-zeros with at least $\Delta - 1$ zero entries between consecutive non-zeros.

In order to employ separated sparsity in statistical problems, we often want to add constraints based on the support set $\mathbb{M}_{k,\Delta}$ to optimization problems such as empirical risk minimization. A standard way of incorporating constraints into gradient-based algorithms is via a *projection operator*. In the context of separated sparsity, this corresponds to the following problem.

Problem 1 For a given input vector $x \in \mathbb{R}^d$, our goal is to project x onto the set $\mathbb{M}_{k,\Delta}$, i.e., to find a vector \hat{x} such that

$$\hat{x} \in \arg \min_{x' \in \mathbb{R}^d : \text{supp}(x') \in \mathbb{M}_{k,\Delta}} \|x - x'\|. \quad (1)$$

Problem 1 is the main algorithmic problem we address in this paper.

Sparse recovery. Structured sparsity has been employed in a variety of machine learning tasks. In order to keep the discussion coherent, we present our results in the context of the well-known sparse linear model:

$$y = X\theta^* + e \quad (2)$$

where $y \in \mathbb{R}^n$ are the observations/measurements, $X \in \mathbb{R}^{n \times d}$ is the design or measurement matrix, and $e \in \mathbb{R}^n$ is a noise vector. The goal is to find a good estimate $\hat{\theta}$ of the unknown parameters θ^* up to the noise level.

The authors of [3] give an elegant framework for incorporating structured sparsity into the estimation problem outlined above. They design a general recovery algorithm that relies on a model-specific *projection oracle*. In the case of separated sparsity, this oracle is required to solve precisely the Problem 1 stated above.

3 The algorithm and our results

Given an arbitrary vector $x \in \mathbb{R}^d$, Problem 1 requires us to find a vector \hat{x} such that $\hat{x} \in \mathbb{M}_{k,\Delta}$ and $\|x - \hat{x}\|$ is minimized. We now slightly reformulate the problem. Let $c \in \mathbb{R}^d$ be a vector such that $c_i = x_i^2$ for all i . Then it is not hard to see that this problem is equivalent to finding a set of k entries in the vector c such that each of these entries is separated by at least Δ and the sum of these entries is maximized. Hence our main algorithmic problem is to find a set of k entries in a non-negative input vector c so that the entries are Δ -separated and their sum is maximized. More formally, our goal is to find a support \hat{S} such that

$$\hat{S} \in \arg \max_{S \in \mathbb{M}_{k,\Delta}} \sum_{i \in S} c_i. \quad (3)$$

In the following, we also consider a relaxed version of (3) called PROJLAGR, which is parametrized by a trade-off parameter λ and a vector \tilde{c} :

$$\text{PROJLAGR}(\lambda, \tilde{c}) := \arg \max_{S \in \mathbb{M}_\Delta} \sum_{i \in S} \tilde{c}_i + \lambda(k - |S|).$$

Intuitively, PROJLAGR represents a Lagrangian relaxation of the sparsity constraint in Equation (3).

3.1 Algorithm

Our main contribution is a new algorithm for Problem 1 that we call *Lagrangian Approach to the Separated Sparsity Problem* (LASSP). The pseudo code is given in Algorithm 1.

LASSP is a *Las Vegas* algorithm: it always returns a correct answer, but the running time of the algorithm is randomized. Concretely, LASSP repeats a main loop until a stopping criterion is reached. Every iteration of LASSP first adds a small perturbation to the coefficients c (see Line 3). This perturbation has only a small effect on the solution but improves the “conditioning” of the corresponding non-convex Lagrangian relaxation PROJLAGR so that it returns a globally optimal solution that almost satisfies the constraint. As we show in Section 5.2, we can solve this relaxation (Line 4) in nearly-linear time. After the algorithm has solved the Lagrangian relaxation, it obtains the final support \hat{S} in line 5

by solving PROJLAGR on a slightly shifted $\hat{\lambda}$ to ensure that the constraint is satisfied with good probability.

Remark: We assume that the bit precision γ required to represent the coefficients c is finite, and we provide our results as a function of γ . For practical purposes, γ is usually a constant. Since the solution to Equation (3) is invariant under scaling by a positive integer and γ is finite, without loss of generality we assume that $c \in \mathbb{Z}^d$.

Algorithm 1 LASSP:

Input: $c \in \mathbb{Z}^d, k \in \mathbb{N}_+$

Output: A solution \hat{S} to (3)

- 1: **repeat**
 - 2: Let $X \in \mathbb{Z}^d$ be a vector such that X_i is chosen uniformly at random from $\{0, \dots, d^3 - 1\}, \forall i$
 - 3: Define vector $\tilde{c} := d^4 c + X$
 - 4: Let $\hat{\lambda} := \arg \min_{\lambda \in \mathbb{Z}} \text{PROJLAGR}(\lambda, \tilde{c})$. In case of ties, maximize $\hat{\lambda}$.
 - 5: Choose $\hat{S} \in \text{PROJLAGR}(\hat{\lambda} - \frac{1}{d+1}, \tilde{c})$
 - 6: **until** $|\hat{S}| = k$
 - 7: **return** \hat{S}
-

3.2 Main results

As our main result, in the following theorem we show that LASSP runs in nearly linear time and solves Problem 1.

Theorem 1 *Let $c \in \mathbb{R}^d$, and let $\gamma \in \mathbb{N}_+$ be the maximal number of bits needed to store any c_i . There is an implementation of LASSP that for every c computes a solution \hat{c} to Problem 1. With probability $1 - 1/d$, the algorithm runs in time $O(d(\gamma + \log d))$.*

Combined with the framework of [3], we get the following.

Corollary 1 *Let y, X, θ^* , and e be as in the sparse linear model in Equation (2). We assume that $\text{supp}(\theta^*) \in \mathbb{M}_{k,\Delta}$ and that X satisfies the model-RIP for $\mathbb{M}_{k,\Delta}$. There is an algorithm that for every y and e returns an estimate $\hat{\theta}$ such that*

$$\|\hat{\theta} - \theta^*\|_2 \leq C\|e\|_2.$$

Moreover, the algorithm runs in time $\tilde{O}(T_X + d)$, where T_X is the time of multiplying the matrices X and X^T by a vector.

The corollary shows that, up to logarithmic factors, the running time is dominated by $T_X + d$. This matches the time complexity of standard sparse recovery and shows that we can utilize separated sparsity without a significant increase in time complexity. We validate these theoretical findings in Section 7 by showing that LASSP runs significantly faster than the state of the art algorithm used for sparse recovery with separation constraints.

Our algorithm LASSP is randomized. However, we also design a deterministic nearly-linear time algorithm and prove

the following theorem. For clarity of exposition, the statement of the deterministic algorithm and the proof of the theorem are deferred to Appendix J.

Theorem 2 *Let $c \in \mathbb{R}^d$ be the input vector and let γ be as in Theorem 1. Then there is an algorithm that computes a solution \hat{c} satisfying Equation (1) and runs in time $O(d(\gamma + \log k) \log d \log \Delta)$.[†]*

3.3 Further results

In Appendix C we present a dynamic programming approach that for a specific, but also natural, family of instances solves the separated sparsity problem in even linear time.

We also consider a natural extension to the 2D-variant of the separated sparsity projection problem and show that it is NP-hard in Appendix D. Moreover, in Appendix E, we extend our model to allow for blocks of separated variables and show that our algorithms for $\mathbb{M}_{k,\Delta}$ also applies to the more general variant. Finally, separated sparsity can be used to model signals in which a longer pattern is repeated multiple times so that any two patterns are at least Δ apart. This model is called *disjoint pulse streams* [12]. Again, the algorithmic core remains the same and algorithms for $\mathbb{M}_{k,\Delta}$ can also be used for this generalization.

3.4 Additional related work

The papers [10, 14] are closely related to our work. The paper [10] proposed the separated sparsity model, provided a sample complexity upper bound, and gave an LP-based model-projection algorithm. However, they resorted to a black-box approach for solving the LP, that lead to a fairly prohibitive $O(d^{3.5})$ time complexity. Recently, [14] provided a faster dynamic program for this problem with a time complexity of $O(d^2)$ and also showed a sample complexity lower bound. The algorithmic aspect of these papers is the main difference from our work: we exploit structure in both the primal and dual formulations of the LP and give an algorithm that *provably* runs in *nearly-linear* time.

Beside the papers addressing the core algorithmic question of projecting onto separated sparse vectors, there is much of work utilizing the sparsity model for applications in neural signal processing [11, 12, 13] and recovery with coherent dictionaries [16, 17]. In the latter application, the separated sparsity constraint enforces that the signal representation only consists of incoherent dictionary atoms. We expect that our algorithmic techniques will also lead to improvements in the context of these applications.

[†]After this work was done, it was pointed to us by Arturs Backurs and Christos Tzamos that this problem exhibits concave Monge property, which can be used to solve the problem in nearly-linear time with different algorithm. For more information, we point the reader to <https://arxiv.org/abs/1802.06440>.

4 Proof of correctness and the roadmap

We begin our analysis by proving that LASSP returns a correct result *if the algorithm terminates*. As we will see later, establishing termination is the crucial part of the analysis. The following lemma is a useful warm-up for understanding how the different pieces of our algorithm fit together.

Lemma 1 *When LASSP terminates, it outputs a support \hat{S} such that x restricted to \hat{S} is a solution to Problem 1.*

Proof. As we have argued above, the problems in Equations (1) and (3) are equivalent. So, we show that LASSP outputs a solution to the problem in Equation (3).

Let \hat{S} be the set returned by LASSP. By the condition of the loop in Line 6, we have $|\hat{S}| = k$. So, as $\hat{S} \in \mathbb{M}_\Delta$ (see the definition of PROJLAGR), we have $\hat{S} \in \mathbb{M}_{k,\Delta}$.

Now, towards a contradiction, assume that support \hat{S} is not a solution to the problem in Equation (3), while support S^* is. This implies that $\sum_{i \in S^*} c_i > \sum_{i \in \hat{S}} c_i$. Now since, without loss of generality, we assumed that $c \in \mathbb{Z}^d$, the last inequality implies $\sum_{i \in S^*} c_i \geq 1 + \sum_{i \in \hat{S}} c_i$, and hence

$$\sum_{i \in S^*} d^4 c_i \geq d^4 + \sum_{i \in \hat{S}} d^4 c_i. \quad (4)$$

Observe that for any support $S \in \mathbb{M}_{k,\Delta}$, the term $\lambda(k - |S|)$ equals zero, and recall that $\hat{S}, S^* \in \mathbb{M}_{k,\Delta}$. Furthermore, by the definition of the random vector X in line 2 and from (4)

$$\begin{aligned} \sum_{i \in S^*} (d^4 c_i + X_i) &\geq \sum_{i \in S^*} d^4 c_i \geq d^4 + \sum_{i \in \hat{S}} d^4 c_i \\ &> \sum_{i \in \hat{S}} (d^4 c_i + X_i). \end{aligned}$$

Since $S^* \in \mathbb{M}_\Delta$, this chain of inequalities contradicts Line 5 of LASSP which chooses \hat{S} as an optimal solution to PROJLAGR $(\hat{\lambda} - 1/(d+1), d^4 c + X)$. This further implies that x restricted to \hat{S} is a solution to Problem 1. \square

4.1 Roadmap

Lemma 1 shows that LASSP outputs the right answer if it terminates. But does LASSP terminate on every input? Answering this question is the most intricate part of this paper. We split the proof in two main pieces. The first part is Section 5, where we provide an alternative view on the separated sparsity problem based on linear programming duality. The duality view paves the way towards proving our main results. In particular, we show that the subroutines in LASSP can be implemented quickly.

Lemma 2 *Single iteration of LASSP can be implemented to run in time $O(d(\gamma + \log d))$.*

The second part is Section 6, where we further study the duality view on separated sparsity. We show that after perturbing the input instance in Line 3 of LASSP, the support obtained with a shifted $\hat{\lambda}$ in Line 5 has cardinality k with high probability.

Lemma 3 *Algorithm LASSP runs only a single iteration with probability at least $1 - 1/d$.*

Together with Lemma 1, these results yield Theorem 1.

5 Part I – To duality and further

We now analyze the running time of a single iteration of LASSP. We provide a series of equivalences, as illustrated in Figure 1, in order to exploit structure in the separated sparsity problem. More precisely, we start with a linear programming (LP) view on separated sparsity. It has already been shown that this viewpoint yields a totally unimodular LP [10], which implies that the LP has an integral solution. Hence solving the LP solves the separated sparsity projection in Problem 1. However, prior work did not utilize this connection to reason about the power of the Lagrangian relaxation approach to the problem.

We begin our detailed analysis of the LP with the dual program \mathcal{D} . By strong duality, the value of \mathcal{D} equals the value of the primal LP. Then we cast \mathcal{D} as minimization of LP \mathcal{D}_λ over λ . This reduction will play the central role in our analysis and connect \mathcal{D}_λ to Line 4 of LASSP.

5.1 The LP Perspective

We start with a linear programming view on problem (3) by considering its LP relaxation denoted by \mathcal{P} :

$$\begin{aligned} &\text{maximize} && c^T u \\ &\text{subject to} && \sum_{i=1}^d u_i = k \\ & && \sum_{j=i}^{\min\{i+\Delta-1, n\}} u_j \leq 1 \quad \forall i = 1 \dots d \\ & && u_i \geq 0 \quad \forall i = 1 \dots d \end{aligned}$$

Given an LP \mathcal{A} we use $\text{VAL } \mathcal{A}$ to denote its optimal objective value. As already noted, \mathcal{P} is totally unimodular and thus there always exists an optimal solution to it that is integral [18]. For completeness, we provide the proof of total unimodularity in Appendix F.

Remark: This implies that a solution to \mathcal{P} can be used to obtain a solution to the separated sparsity model projection: if u^* is an optimal solution to \mathcal{P} , we can derive the optimal support of (3) from the non-zero entries among the LP variables u_1^*, \dots, u_d^* . It is unclear, however, if there is a way to directly solve this LP fast, e.g. it is not known how to solve \mathcal{P} directly in time matching the running time of our algorithm LASSP.

A key step in our approach is understanding the separated sparsity structure from the dual point of view. The dual LP




Figure 1: The values of the problems in the diagram are equal. Every equivalence relation carries structural information that we utilize in our analysis.

to \mathcal{P} , denoted by \mathcal{D} , is given as follows

$$\begin{aligned}
 &\text{minimize} && w_0 k + \sum_{i=1}^d w_i \\
 &\text{subject to} && w_0 + \sum_{\substack{j: j \geq 1 \text{ and} \\ j \leq i \leq j + \Delta - 1}} w_j \geq c_i && \forall i = 1 \dots d \\
 &&& w_i \geq 0 && \forall i = 1 \dots d \\
 &&& w_0 \in \mathbb{R}
 \end{aligned}$$

Then, as \mathcal{P} is integral, so is \mathcal{D} .

Corollary 2 *For an integer k and a vector of integers c , there exists \hat{w} such that \hat{w} is an optimum of \mathcal{D} and $\hat{w}_0 \in \mathbb{Z}$.*

We also define \mathcal{D}_{λ} as the LP \mathcal{D} in which the variable w_0 is set to λ . Now, it is not hard to show the following lemma, whose proof is deferred to Lemma 10 in Appendix F.

Lemma 4 \mathcal{D}_{λ} is convex with respect to λ .

From the definition of \mathcal{D} , the following equality holds $\text{VAL } \mathcal{D} = \min_{\lambda \in \mathbb{R}} \text{VAL } \mathcal{D}_{\lambda}$. Furthermore, Corollary 2 implies that it is sufficient to consider λ in \mathbb{Z} only, i.e.

$$\text{VAL } \mathcal{D} = \min_{\lambda \in \mathbb{Z}} \text{VAL } \mathcal{D}_{\lambda}. \quad (5)$$

Now, Lemma 4 implies that we can obtain $\text{VAL } \mathcal{D}$ by applying ternary search over λ on function $\text{VAL } \mathcal{D}_{\lambda}$.

5.2 Implementing one iteration of LASSP efficiently

We now derive the final connection between \mathcal{D} and LASSP which will enable us to obtain $\hat{\lambda}$ at line 4 in nearly-linear time. To that end, consider $\mathcal{P}_{\text{no-}k}(\lambda)$ defined as

$$\begin{aligned}
 &\text{maximize} && c^T u + \lambda(k - \mathbb{1}^T u) \\
 &\text{subject to} && \sum_{j=i}^{\min\{i+\Delta-1, d\}} u_j \leq 1 && \forall i = 1 \dots d \\
 &&& u_i \geq 0 && \forall i = 1 \dots d
 \end{aligned}$$

Observe that compared to \mathcal{P} , $\mathcal{P}_{\text{no-}k}(\lambda)$ does not contain the sparsity constraint. Furthermore, the LP $\mathcal{P}_{\text{no-}k}(\lambda)$ is a relaxed version of $\text{PROJLAGR}(\lambda, c)$. Also, as \mathcal{P} is, then $\mathcal{P}_{\text{no-}k}(\lambda)$ is totally unimodular. Now this sequence of conclusions results in the following.

Corollary 3 *Problems $\mathcal{P}_{\text{no-}k}(\lambda)$ and $\text{PROJLAGR}(\lambda, c)$ are equivalent.*

To obtain the final connection, we consider $L_{\mathcal{P}}$ given by

$$L_{\mathcal{P}} := \min_{\lambda \in \mathbb{R}} \mathcal{P}_{\text{no-}k}(\lambda).$$

Next, note that the dual of $\mathcal{P}_{\text{no-}k}(\lambda)$ is \mathcal{D}_{λ} . Hence, the strong duality implies $\text{VAL } \mathcal{D}_{\lambda} = \text{VAL } \mathcal{P}_{\text{no-}k}(\lambda)$. That together with $\text{VAL } \mathcal{D} = \min_{\lambda \in \mathbb{R}} \text{VAL } \mathcal{D}_{\lambda}$ yields that \mathcal{D} and $L_{\mathcal{P}}$ coincide as functions in λ . Furthermore, since we can solve \mathcal{D} by applying ternary search over integral values of λ and \mathcal{D}_{λ} , we can solve $L_{\mathcal{P}}$ by applying ternary search over integral values of λ and function $\mathcal{P}_{\text{no-}k}(\lambda)$. But since $\mathcal{P}_{\text{no-}k}(\lambda)$ and $\text{PROJLAGR}(\lambda, c)$ are equivalent, we can also obtain $\hat{\lambda}$ at line 4 of LASSP by applying ternary search over λ .

Now, it is very easy to see that for an optimal solution w^* of \mathcal{D} we have $w_0^* \leq \max_i |c_i|$. It is also not hard to show that there is an optimal solution such that $w_0^* \geq -(k-1) \max_i |c_i|$ (see Lemma 18). Therefore, in order to find optimal $\hat{\lambda}$ it suffices to execute $O(\log \max_i |c_i| + \log k) = O(\gamma + \log d)$ iterations of ternary search.

Every iteration of the ternary search invokes PROJLAGR , which can be implemented to run in linear time.

Lemma 5 *Given λ and $\hat{c} \in \mathbb{R}^d$, there is an algorithm that finds support $\hat{S} \in \text{PROJLAGR}(\lambda, \hat{c})$ in time $O(d)$.*

Proof. Observe that for a fixed λ , solving $\text{PROJLAGR}(\lambda, \hat{c})$ is equivalent to finding support $S' \in \mathbb{M}_{\Delta}$ that maximize $\sum_{i \in S'} (\hat{c}_i - \lambda)$. Hence, we can reinterpret $\text{PROJLAGR}(\lambda, \hat{c})$ as follows: given a vector $\tilde{c} := \hat{c} - \lambda \mathbb{1}$, select a subset of $[d]$ of indices (not necessarily k of them) so that (i) every two indices are at least Δ apart, and (ii) the sum of the values of \tilde{c} at the selected indices is maximized.

This task can be solved by standard dynamic programming in the following way. For every i , we define s_i to be the maximum value of the described task restricted to the first i indices of \tilde{c} . Then, it is easy to see that $s_{i+1} = \max\{s_i, s_{i+1-\Delta} + \tilde{c}_{i+1}\}$. Namely, we can either decide not to select index $i+1$, in which case the best is already contained in s_i ; or, we can decide to select index $i+1$ which has value \tilde{c}_{i+1} and for the rest we consider $s_{i+1-\Delta}$. Therefore, s_d can be obtained in time $O(d)$. Now it is easy to reconstruct the corresponding support in linear time. For completeness, we provide a full algorithm and a detailed proof in Lemma 12, Appendix G. \square

Putting all together proves Lemma 2.

6 Part II – Active constraints

Note that the chain of equivalences present in Figure 1 shows that $\min_{\lambda} \text{PROJLAGR}(\lambda, c)$ outputs the sum of coordinates of an optimal solution of problem 1. Prior to our work, it was not even known how to obtain this value in time faster than $O(dk)$, while our result shows we can compute it in nearly-linear time. So, it is natural to ask whether the same relaxation also outputs a support of size k ? The answer is, unfortunately, no. To see that, consider the example: $c = (4, 7, 5, 0, 0, 5, 8, 5)$, $\Delta = 2$, and $k = 3$. For $\lambda > 2$ every solution u_{λ}^* to $\text{PROJLAGR}(\lambda, c)$ is such that u_{λ}^* has less than k non-zeros. On the other hand, for every $\lambda \leq 2$ there exists a solution u_{λ}^* such that u_{λ}^* contains more than k non-zeros. Therefore, there is no λ , neither $\lambda - 1/(d+1)$, for which $\text{PROJLAGR}(\lambda, c)$ provably outputs a support of cardinality k . This also suggests that the perturbation we apply in lines 2-3 is essential!

Instead of studying lines 2-5 of LASSP directly, we shift our focus to \mathcal{D}_{λ} . In particular, we exhibit very close connection between its structure and the sparsity of the primal solution, which we present via the notion of "active constraints". Then we use these findings in our analysis to show that slight perturbation of the input instance, while not affecting the value of the solution, makes it possible to obtain a solution to problem 3 by applying Lagrangian relaxation.

6.1 Solving \mathcal{D}_{λ}

Observe that once we fixed the value of w_0 , all remaining constraints in \mathcal{D}_{λ} are "local" since they only affect a known interval of length Δ . They are also ordered in a natural way. As a result, we can solve \mathcal{D}_{λ} by making a single pass over these variables. Starting with w_1 and all variables set to 0, we consider each constraint from left to right and increase the variables to satisfy these constraints in a lazy manner. That is, if in our pass we reach a constraint with index i that is still not satisfied, we increase the value of w_i until that constraint becomes satisfied and then move to the next constraint. Given c and λ , algorithm Dual-Greedy, i.e. Algorithm 2, formalizes this approach whose analysis appears in Appendix F.

Algorithm 2 Dual-Greedy:

Input: $c \in \mathbb{Z}^d, \lambda \in \mathbb{R}$

Output: an optimal solution w to \mathcal{D} such that $w_0 = \lambda$

```

1:  $w \leftarrow 0; \quad w_0 \leftarrow \lambda$ 
2:  $sum_{\Delta} \leftarrow 0$ 
3: for  $i := 1 \dots d$  do
4:   if  $i - \Delta \geq 1$  then  $sum_{\Delta} \leftarrow sum_{\Delta} - w_{i-\Delta}$ 
5:    $diff \leftarrow c_i - (w_0 + sum_{\Delta})$ 
6:   if  $diff > 0$  then  $w_i \leftarrow diff$ 
7:    $sum_{\Delta} \leftarrow sum_{\Delta} + w_i$ 
8: return  $w$ 

```

6.2 Tracking the change of \mathcal{D}_{λ}

Next we introduce the key concept that we need for relating the solution of dual to the sparsity of Lagrangian relaxation of the primal: the notion of *active constraints*. Let w be a vector obtained by $\text{Dual-Greedy}(c, \lambda)$ and let w' be a vector obtained by $\text{Dual-Greedy}(c, \lambda - \varepsilon)$, for some fixed $\lambda \in \mathbb{Z}$ and small $\varepsilon \in (0, 1)$. Then, the set of coordinates that are for ε larger in w' than in w are called active constraints. Figure 2 provides an illustration of this concept. Intuitively, the active constraints correspond to those variables of \mathcal{D} that increase when w_0 decreases by some small value. Hence, one can interpret active constraints as gradients of \mathcal{D}_{λ} with respect to the variable λ . This concept appears to be very useful in characterizing the optimal solution of \mathcal{D} in an alternative way. In particular, the following lemma holds.

Lemma 6 *Let $c \in \mathbb{Z}^d, \lambda \in \mathbb{Z}$, and $w \leftarrow \text{Dual-Greedy}(c, \lambda)$. Then, if w has exactly k active constraints the vector w is an optimal solution to \mathcal{D} .*

A full proof of a statement stronger than Lemma 6 along with its proof appears in Lemma 13, Appendix H, while in this section we provide a proof sketch. Let $w(\varepsilon) = \text{Dual-Greedy}(c, \lambda - \varepsilon)$, for some small $\varepsilon \in (0, 1)$. By the definition of active constraints and the fact that w has k many, there are exactly k coordinates that are larger by ε in $w(\varepsilon)$ than in w . In addition, $w(\varepsilon)_0 = \lambda - \varepsilon$ and $w_0 = \lambda$. It is not hard to show that all the other coordinates of $w(\varepsilon)$ and w are the same, which we can express as $\sum_{i=1}^d w(\varepsilon)_i = k\varepsilon + \sum_{i=1}^d w_i$. Now, recall that the objective function of dual \mathcal{D} with respect to vector w equals $w_0 k + \sum_{i=1}^d w_i$. Then we have

$$w_0 k + \sum_{i=1}^d w_i = (w_0 - \varepsilon)k + k\varepsilon + \sum_{i=1}^d w_i = w(\varepsilon)_0 k + \sum_{i=1}^d w(\varepsilon)_i.$$

Hence, the objective values of dual \mathcal{D} for vectors $w(\varepsilon)$ and w are equal, for **all** the values $\varepsilon \in (0, 1)$. As \mathcal{D} is convex in the value of variable w_0 and $\text{Dual-Greedy}(c, \lambda)$ provides an optimal solution to \mathcal{D} such that $w_0 = \lambda$, then w is an optimal solution to \mathcal{D} .

6.3 Wrapping up – perturbation and optimal sparsity

Now we use Lemma 6 to prove the following, which essentially justifies line 5 of LASSP.

Lemma 7 *Let $c \in \mathbb{Z}^d, \lambda \in \mathbb{Z}$, and $\tilde{w} \leftarrow \text{Dual-Greedy}(c, \lambda)$. Assume that \tilde{w} has exactly k active constraints. Then, any optimal support S^* of $\text{PROJLAGR}(\lambda - \varepsilon, c)$, for $0 < \varepsilon < 1/d$, has cardinality exactly k .*

So, if we produce λ as in Lemma 7, we will solve problem (3). These steps are implemented by lines 4-6 of LASSP. However, as illustrated in the beginning of the section, λ as in Lemma 7 might not exist. Intuitively, this




Figure 2: A sketch of an instance of \mathcal{D} : $c = (4, 6, 4, 3, 3, 5, 6, 2, 2, 2)$ and $\Delta = 3$. The left and right figure depicts w obtained by Dual-Greedy($c, 2$) and Dual-Greedy($c, 2 - \varepsilon$) for $0 < \varepsilon < 1$, respectively. Constraints 1, 5, and 10 are active, i.e., if w_0 on the left is decreased by ε then only w_1 , w_5 , and w_{10} increase by ε , as shown on the right. Every w_i , for $i \geq 1$, covers Δ c -poles to its right.

situation happens when the number of active constraints jumps from a value smaller than k to a value larger than k for a very small change of λ . In such a case, we are unable to obtain λ as in Lemma 7. A key component of our analysis is showing that there is an efficient way of randomly altering c , and obtaining \tilde{c} , so that with high probability \tilde{c} is such that: $\tilde{\lambda}$ is obtained as at line 4; and, $\tilde{w} \leftarrow \text{Dual-Greedy}(\tilde{c}, \tilde{\lambda})$ has the same property as in Lemma 7. Lines 2 and 3 of LASSP implement this random perturbation. Intuitively, the perturbation achieved by X variables adds noise to our input instance so that the number of active constraints changes by at most one as λ slides over the integer domain. Following this intuition we obtain a proof of Lemma 3. However, due to the space limitation, we present its proof in Appendix I.

We conclude the section by giving a proof of Lemma 7.

Proof of Lemma 7. Recall that by our assumption there are exactly k active constraints defined by \tilde{w} . Then from Lemma 6 it follows that \tilde{w} is a minimizer of \mathcal{D} , i.e. $\text{VAL } \mathcal{D}_\lambda = \text{VAL } \mathcal{D}$. By the integrality of \mathcal{D} we have that $\text{VAL } \mathcal{D} \in \mathbb{Z}$. Let $\lambda' = \lambda - \varepsilon$, for some $0 < \varepsilon < 1/d$. Then, it holds $\text{VAL } \mathcal{D}_{\lambda'} = \text{VAL } \mathcal{D}$ as only the k variables corresponding to active constraints increased by ε while $w_0 = \lambda'$ decreased by ε . From the strong duality we also have $\text{VAL } \mathcal{P}_{\text{no-}k}(\lambda') = \text{VAL } \mathcal{D}_{\lambda'} = \text{VAL } \mathcal{D}$.

Let \tilde{u} be an integral optimal solution of $\mathcal{P}_{\text{no-}k}(\lambda')$. Following the definition we have

$$\text{VAL } \mathcal{P}_{\text{no-}k}(\lambda') = (c^T - \lambda \mathbb{1}^T) \tilde{u} + \lambda k + \varepsilon (\mathbb{1}^T \tilde{u} - k).$$

Now we have the following properties: \tilde{u} is integral; $0 < \varepsilon |\mathbb{1}^T \tilde{u} - k| < 1$ whenever $1 \leq |\mathbb{1}^T \tilde{u} - k| \leq d$; $c \in \mathbb{Z}^d$; $v \in \mathbb{Z}$; and $\text{VAL } \mathcal{P}_{\text{no-}k}(\lambda') \in \mathbb{Z}$. Therefore, we have $\mathbb{1}^T \tilde{u} - k = 0$, and hence $\mathbb{1}^T \tilde{u} = k$. Since we showed the equivalence between $\mathcal{P}_{\text{no-}k}(\lambda')$ and $\text{PROJLAGR}(\lambda', c)$ the lemma follows. \square

7 Experiments

We empirically validate the claims outlined in the previous sections. To that end, we compare LASSP with the $O(dk)$ -

time dynamic program (DP) described in Section C as a baseline. Note that this DP already has a better time complexity than the best previously published algorithm from [14]. Both algorithms are implemented in the Julia programming language (version 0.5.0), which typically achieves performance close to C/C++ for combinatorial algorithms.

7.1 Synthetic data

We perform experiments with synthetic data in order to investigate how the algorithms scale as a function of the input size. We study two different setups: (i) the running time of the projection algorithms on their own, and (ii) the overall running time of a sparse recovery algorithm using the projection algorithms as a subroutine. For the latter, we use the structure-aware variant of the popular CoSaMP algorithm [19, 3].

Figures 3(a)-(b). For a problem of size d , we set the sparsity to $k = d/50$ and generate a random separated sparse vector with parameter $\Delta = (d - 5(k + 1))/k - 1$. The non-zero coefficients are i.i.d. ± 1 . For the projection-only benchmark, we add Gaussian noise with $\sigma = 1/10$ to all coordinates in order to make the problem non-trivial. For each problem size, we run 10 independent trials and report their mean.

Figure 3(a) shows the speed-up obtained by our nearly-linear time projection relative to the DP baseline. We observe that LASSP is up to $150\times$ faster. This confirms our expectation that LASSP scales gracefully with the problem size, while the DP essentially becomes a quadratic-time algorithm.

Figure 3(b) compares the running times of CoSaMP with three different projection operators. The first variant makes no structural assumptions and uses hard thresholding as projection operator. The other two variants use a projection for the separated sparsity model, relying on the DP baseline and LASSP, respectively. The results show that the version of CoSaMP using LASSP instead of the DP is significantly faster. Moreover, CoSaMP with a simple sparse projection has similar running time to CoSaMP with our structured projection. Finally, we note that CoSaMP with separated

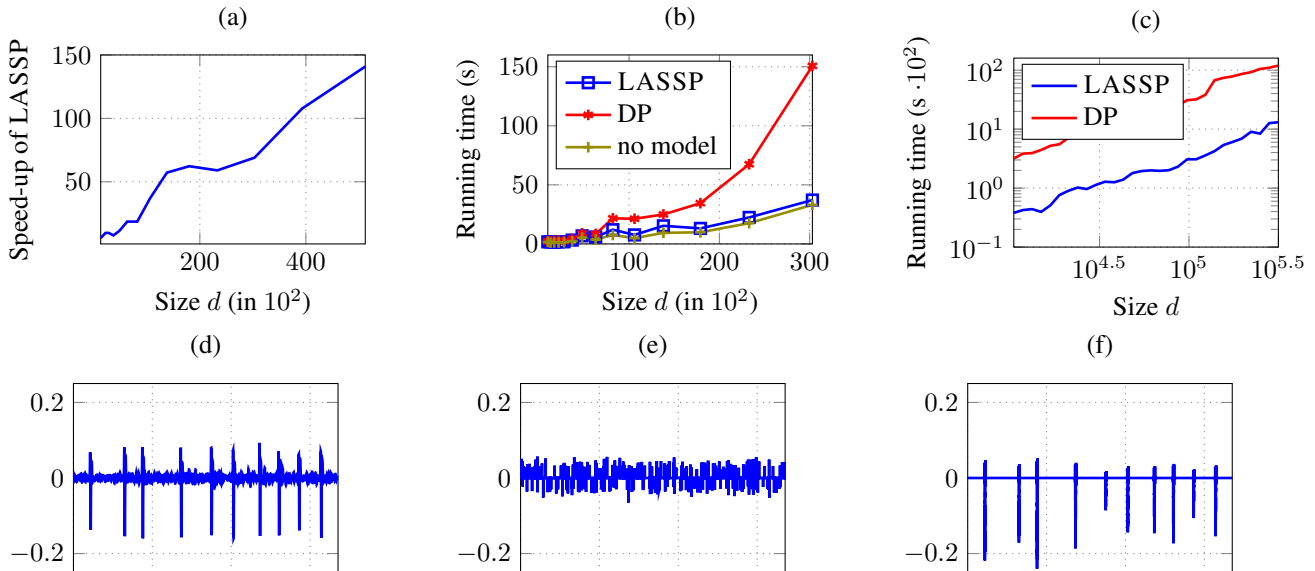


Figure 3: Separated sparsity experiments. In the top row we plot running times of our algorithm LASSP relative to the previous work. Plots (a) and (c) are obtained by projecting signals on the separated sparsity model. Plot (b) compares the running times of CoSaMP with three different projection operators. The variant "no-model" makes no structural assumptions and uses hard thresholding as projection operator. In the bottom row, in (e) we plot the recovery of the signal (d) obtained by the algorithm "no model". Plot (f) shows the recovery of LASSP. Both of the procedures use the same number of measurements.

sparsity requires $1.5\times$ fewer measurements to achieve the same recovery quality as CoSaMP with standard sparsity.

Figure 3(c). We fix the sparsity $k = 100$ and vary the length d of the signal. The signal is obtained in the same way as for plots Figure 3(a)-(b). We observe that our algorithm runs 10x faster than the baseline. Furthermore, the plot shows that the running times of both the baseline DP and LASSP scale linearly with the signal length d . This behavior is expected for the DP. On the other hand, our theoretical findings predict that the running time of LASSP scales as $d \log d$. This suggests that the empirical performance of our algorithm is even (slightly) better than what our proofs state.

In Appendix A we report results of additional experiments.

7.2 Neuronal signals

We also test LASSP on neuron spike train data from [12]. See Figure 3(d) for this input data. First, we run CoSaMP with a "standard sparsity" projection. The recovered signal is depicted in Figure 3(e). Next, we run the convolutional sparsity CoSaMP of [12] and use our fast projection algorithm. The recovered signal is given in Figure 3(f). For the both experiments we use $n = 250$ measurements.

We also run the convolutional sparsity CoSaMP on neuron spike train data of length 10^5 , comparing the running time of LASSP and DP as projection operators. CoSaMP with LASSP runs $2\times$ faster in this context. We do not compare the running time relative to CoSaMP with a standard sparsity projection as it requires $10\times$ more measurements to achieve

accurate recovery.

8 Conclusions

We have designed a nearly-linear time algorithm for projecting onto the set of separated sparse vectors. The core technique in our algorithm is Lagrangian relaxation. One of the key insights here is that even though there are separated sparsity instances for which the Lagrangian relaxation does not provide an optimal solution, it is still possible to obtain an optimal solution if the original input instance is only slightly perturbed. Furthermore, this perturbation does not change the final output, but rather drives the algorithm to choose an optimal solution of interest even after the hard sparsity constraint is relaxed. Our experiments show that our algorithm is not only of theoretical significance, but also outperforms the state of the art in practice. Exploring the power of perturbed Lagrangian relaxations for other non-convex constraint sets is an important direction for future work. We believe that our framework will enable simple and efficient algorithms for other problems as well.

Acknowledgments.

We thank Arturs Backurs for insightful discussions. L. Schmidt thanks Chinmay Hegde for providing a dataset for some of our experiments. A. Mađry was supported in part by an Alfred. P. Sloan Research Fellowship, Google Research Award and the NSF grant CCF-1553428. S. Mitrović was supported by Swiss NSF (grant number P1ELP2_161820). Part of this work was carried out while S. Mitrović was visiting MIT.

References

- [1] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2006.
- [2] Y. Eldar and M. Mishali, “Robust recovery of signals from a structured union of subspaces,” *IEEE Transactions on Information Theory*, vol. 55, no. 11, pp. 5302–5316, 2009.
- [3] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, “Model-based compressive sensing,” *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.
- [4] J. Mairal, R. Jenatton, G. Obozinski, and F. Bach, “Convex and network flow optimization for structured sparsity,” *The Journal of Machine Learning Research*, vol. 12, pp. 2681–2720, 2011.
- [5] J. Huang, T. Zhang, and D. Metaxas, “Learning with structured sparsity,” *The Journal of Machine Learning Research*, vol. 12, pp. 3371–3412, 2011.
- [6] N. S. Rao, B. Recht, and R. D. Nowak, “Universal measurement bounds for structured sparse signal recovery,” in *AISTATS*, ser. JMLR Proceedings, vol. 22, 2012, pp. 942–950.
- [7] S. N. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu, “A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers,” *Statistical Science*, vol. 27, no. 4, pp. 538–557, 11 2012.
- [8] B. Bah, L. Baldassarre, and V. Cevher, “Model-based sketching and recovery with expanders,” in *SODA*, 2014, pp. 1529–1543.
- [9] C. Hegde, P. Indyk, and L. Schmidt, “A nearly-linear time framework for graph-structured sparsity,” in *ICML*. JMLR Workshop and Conference Proceedings, 2015, pp. 928–937.
- [10] C. Hegde, M. F. Duarte, and V. Cevher, “Compressive sensing recovery of spike trains using a structured sparsity model,” in *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- [11] E. L. Dyer, M. F. Duarte, D. H. Johnson, and R. G. Baraniuk, “Recovering spikes from noisy neuronal calcium signals via structured sparse approximation,” in *LVA/ICA*, 2010, pp. 604–611.
- [12] C. Hegde and R. G. Baraniuk, “Sampling and recovery of pulse streams,” *IEEE Transactions on Signal Processing*, vol. 59, no. 4, pp. 1505–1517, 2011.
- [13] E. L. Dyer, C. Studer, J. T. Robinson, and R. G. Baraniuk, “A robust and efficient method to recover neural events from noisy and corrupted data,” in *6th International IEEE/EMBS Conference on Neural Engineering (NER)*, 2013, pp. 593–596.
- [14] S. Foucart, M. F. Minner, and T. Needham, “Sparse disjointed recovery from noninflating measurements,” *Applied and Computational Harmonic Analysis*, vol. 39, no. 3, pp. 558 – 567, 2015.
- [15] Q. Huang and S. M. Kakade, “Super-resolution off the grid,” in *Conference on Neural Information Processing Systems (NIPS)*, 2015, pp. 2665–2673.
- [16] M. F. Duarte and R. G. Baraniuk, “Spectral compressive sensing,” *Applied and Computational Harmonic Analysis*, vol. 35, no. 1, pp. 111 – 129, 2013.
- [17] T. Needham, “Dictionary-sparse and disjointed recovery,” in *International Conference on Sampling Theory and Applications (SampTA)*, 2015, pp. 278–282.
- [18] G. L. Nemhauser and L. A. Wolsey, *Integer and Combinatorial Optimization*. Wiley-Interscience, 1988.
- [19] D. Needell and J. A. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [20] C. Hegde, P. Indyk, and L. Schmidt, “Approximation algorithms for model-based compressive sensing,” *IEEE Transactions on Information Theory*, vol. 61, no. 9, 2015.
- [21] ———, “Nearly linear-time model-based compressive sensing,” in *ICALP*, 2014, vol. 8572.
- [22] R. J. Fowler, M. S. Paterson, and S. L. Tanimoto, “Optimal packing and covering in the plane are np-complete,” *Information processing letters*, vol. 12, no. 3, pp. 133–137, 1981.
- [23] A. Schrijver, *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media, 2002, vol. 24.




Figure 4: The plots compare the running times of our algorithm LASSP and the DP baseline. Plots (a) and (c) are obtained by projecting signals directly on the separated sparsity model. Plot (b) shows the speed-up of CoSaMP recovery obtained by using LASSP as a projection operator relative to using the DP for projection.

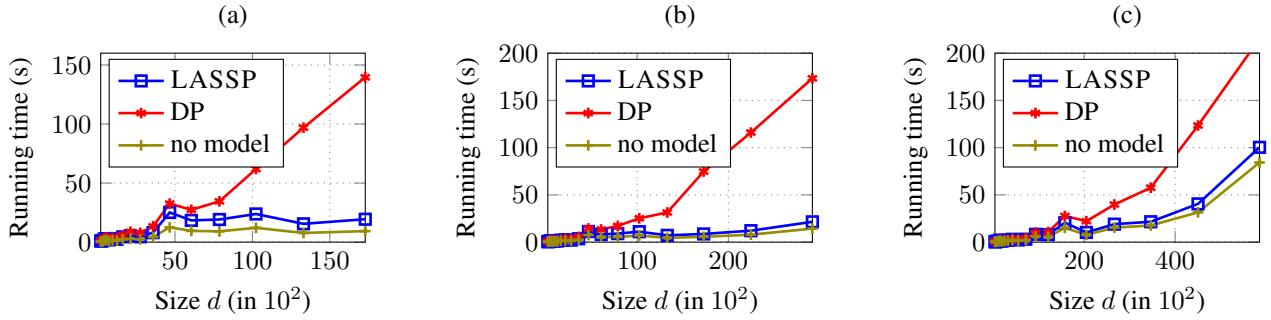


Figure 5: The plots compare the running times of CoSaMP with three different projection operators. The variant "no-model" makes no structural assumptions and uses hard thresholding as projection operator. The other two variants use a projection for the separated sparsity model, relying on the DP baseline and LASSP.

A Additional experiments

In this section we present experiments additional to those present in Section 7.

In plots Figure 4(a)-(b) we run experiments on real signals. The signals are obtained from signal in Figure 3(d) by concatenating its copies. Figure 4(a) is obtained by projecting the signal directly to the separated sparsity model. Then we plot the ratio of the running times of the baseline DP and our algorithm LASSP. Figure 4(b) represents the same type of speed-up ratio but this time for the CoSaMP recovery of signal measurements.

For the plots Figure 4(c) and Figure 5(a)-(c) we make the setup similar to the one for Figures 3(a)-(b). Namely, given a problem size d , parameters α and β we set the sparsity to $k = d/\alpha$ and generate a random separated sparse vector with parameter $\Delta = (d - \beta(k + 1))/k - 1$. The non-zero coefficients are i.i.d. ± 1 . We add Gaussian noise with σ to all coordinates in order to make the problem non-trivial. For each problem size, we run 10 independent trials and report their mean.

In Figure 4(c) we compare the running times of projections of the baseline DP and our algorithm LASSP on synthetic data for the following parameters: $d = 5 \cdot 10^4$; varying the sparsity k ; $\beta = 5$; and $\sigma = 0.5$. This plot confirms our theoretical findings that the running time of LASSP scales more gracefully with the growth of the sparsity than the running time of DP does.

Similar to Figure 3(b), Figure 5(a)-(c) compare the running times of CoSaMP with three different projection operators. The first variant makes no structural assumptions and uses hard thresholding as projection operator. The other two variants use a projection for the separated sparsity model, relying on the DP baseline and LASSP, respectively. We use the following parameters:

- Figure 5(a): $\alpha = 10$; $\beta = 2$; and $\sigma = 0.01$
- Figure 5(b): $\alpha = 20$; $\beta = 8$; and $\sigma = 0.95$

- Figure 5(c): $\alpha = 100$; $\beta = 30$; and $\sigma = 0.2$.

The results show that CoSaMP with a simple sparse projection has similar running time to CoSaMP with our structured projection. Moreover, CoSaMP using LASSP as the structured projection is usually significantly faster than CoSaMP using the baseline DP. We note that even in the favorable case for the DP structured projection, like in Figure 5(c) where the sparsity k is $100\times$ smaller than the dimension d , our algorithm again shows better performance.

B An approximate projection counterexample

Multiple recent algorithms for structured sparse projections build on the *approximation-tolerant* framework of [20]. In this framework, it suffices to design approximate projections instead of solving the model projection problem exactly. For some sparsity structures, this approach has led to significantly faster algorithms [21, 9]. Hence it is interesting to see whether the approximation route is also helpful for separated sparsity. In fact, it is easy to design the following 2-approximation algorithm that runs in nearly-linear time.

Partition the vector c into blocks of length Δ . Then, number those blocks 1 through $\lceil d/\Delta \rceil$ in the order they appear in c . In each block, choose an index corresponding to the largest coordinate of c within that block. Split those indices into two groups G and H based on the parity of the corresponding block. Formally, define G and H as follows

$$G := \left\{ \arg \max_{j=(i-1)\Delta+1 \dots \min\{i\Delta, d\}} c_j \mid 1 \leq i \leq \lceil d/\Delta \rceil \text{ and } i \text{ is odd} \right\},$$

and similarly

$$H := \left\{ \arg \max_{j=(i-1)\Delta+1 \dots \min\{i\Delta, d\}} c_j \mid 1 \leq i \leq \lceil d/\Delta \rceil \text{ and } i \text{ is even} \right\}.$$

Let $g_1, \dots, g_{|G|}$ be an ordering of the elements of G so that $c_{g_i} \geq c_{g_j}$ whenever $i \leq j$. Similarly, let $h_1, \dots, h_{|H|}$ be an ordering of the elements of H such that $c_{g_i} \geq c_{g_j}$ whenever $i \leq j$.

Next, define $G' := \{g_i \mid 1 \leq i \leq \min\{k, |G|\}\}$ and $H' := \{h_i \mid 1 \leq i \leq \min\{k, |H|\}\}$. Now, it is not hard to see that the set attaining larger value among G' and H' has a solution value that is at least $OPT/2$, where OPT is the maximum sum attainable with a Δ -separated and k -sparse vector (in the language of [20], this is an approximate head projection).

While the above algorithm runs in nearly-linear time and achieves a constant-factor approximation, there is a catch. In particular, the returned support pattern might contain fewer than k indices (as a simple example, consider the case $k = 2$, $\Delta = 2$, and $c = (1, 100, 1)$). This raises the question of the sample complexity for this relaxed sparsity model. Let k' denote the number of output indices of the 2-approximation algorithm. As we have described in the preliminaries, the sample complexity of the separated sparsity model is $O(k \log(n/k - \Delta))$. For concreteness, we now consider the case $k' = k/2$ and $d/k - \Delta \in O(1)$ (the latter is the important regime where the separated sparsity model achieves a sample complexity of $m = O(k)$). Then we have $d/k' - \Delta \in O(1) + \Delta$. Therefore, if we would like to apply the 2-approximation algorithm in the case $k' = k/2$, the sample complexity would be

$$O(k \log(2d/k - \Delta)) = O(k \log \Delta) = O(k \log d/k)$$

where we used the assumption that $d/k - \Delta \in O(1)$. It is important to note that this sample complexity is worse than $O(k)$ and falls back to the $O(k \log d/k)$ sample complexity of “standard” k -sparse recovery. So by using an approximate projection we have lost the sample complexity advantage of the separated sparsity model. Since we already know how to recover k -sparse vectors in nearly-linear time without resorting to structured sparsity, the approximation algorithm does not provide a novel trade-off.

C Dynamic programming

Before we introduce our dynamic programming (DP) algorithm for the separated sparsity problem, we briefly review a variant of the DP given in [14]. We remark that the variant below has a time complexity of $O(dk)$, which is already an improvement over the $O(d^2)$ of [14] when $k = o(d)$. However, the authors are not aware of a work that describes the $O(dk)$ DP approach, and hence refer to it as folklore. We improve the analysis of this DP further and give a faster variant that runs in time $O(k(d - (k - 1)\Delta))$. In the regime where where the slack $d/k - \Delta$ is constant, the running time simplifies to $O(d + k^2)$. So for the very sparse case $k \in \tilde{O}(\sqrt{d})$, our improved DP already achieves a nearly-linear running time.

C.1 The basic dynamic program

The folklore dynamic program fills the following table $DP[i][j]$. Value $DP[i][j]$ can be interpreted as follows: Maximal value of choosing j coordinates of c with positions in $1 \dots i$, such that any two chosen coordinates are at distance Δ at least.

$$\begin{aligned} DP[i][0] &= 0 \\ DP[i][j] &= -\infty, && \text{if } i < 1 \\ DP[i][j] &= \max \{ DP[i-1][j], c_i + DP[i-\Delta][j-1] \}, && \text{otherwise} \end{aligned}$$

The following claim follows easily:

Theorem 3 (Folklore) $DP[n][k]$ is an optimal value to the Δ -separated problem and can be computed in time $O(dk)$. The elements that constitute the optimal value can be output from DP in time $O(d+k)$.

C.2 An improved dynamic program

The dynamic program outlined above already achieves an improvement over the algorithms in [10, 14]. However, the dynamic program is still wasteful with the state space it considers. Consider case in which there is only one possible configuration, i.e. $d = (k-1)\Delta + 1$ for some $k, \Delta \geq 1$. Regardless of the cost vector c , an input with such parameters has only one valid solution. Nevertheless, the natural implementation of the dynamic program above runs in $\Theta(dk)$ time. Note that the definition of DP does not take into account the *mandatory distance* required by the separated sparsity model. By mandatory we refer to Δ distance between any two chosen coordinates of c . This observation gives raise to new dynamic programming definition, that we denote by \overline{DP} , which directly implements mandatory distances. Let

$$s := d - ((k-1)\Delta + 1),$$

which can be seen as the *slack* distance that is not mandatory. For instance, if we choose two coordinates of c with indices i and $i + \Delta + 5$, and no other coordinate between them, then we say that Δ distance between them is mandatory, and the remaining distance of 5 is slack. Now, \overline{DP} is defined as

$$\begin{aligned} \overline{DP}[i][0] &= 0 \\ \overline{DP}[i][j] &= -\infty, && \text{if } i < 0 \\ \overline{DP}[i][j] &= \max \{ \overline{DP}[i-1][j], c_{pos(i,j)} + \overline{DP}[i][j-1] \}, && \text{otherwise} \end{aligned}$$

where $pos(i, j) = d - (s - i) - (k - j)\Delta$. Value $pos(i, j)$ gives us the smallest index of c that we can chose if to the left of that index there are j other chosen indices, and slack distance of i is used. With the definition of \overline{DP} in hands one can easily show the following theorem.

Theorem 4 Value $\overline{DP}[s][k]$ is an optimal value to the Δ -separated problem and can be computed in time $O(ks)$. The elements that constitute the optimal value can be output from \overline{DP} in time $O(d+k)$.

Comparing Theorem 3 and Theorem 4, we see that for certain input parameters, e.g. when $s \in O(d)$, we did not make much progress by devising \overline{DP} . However, as we have already mentioned, the setting we care the most is when $d/k - \Delta$ is a constant, as with such a choice of parameters the sample complexity is $O(k)$. But in that case we do achieve a significant improvement over Theorem 3.

Corollary 4 If $d/k - \Delta \in O(1)$, then there is an algorithm that solves the Δ -separated problem in $O(k^2 + d)$ time.

D Two dimensional Δ -separated problem is NP-Hard

In this section we consider a natural extension of the Δ -separated problem in which vector c is two dimensional, i.e. $c \in \mathbb{R}^{d \times m}$. Formally, given c , sparsity parameter k , and integer Δ , the goal is to output k pairs of integers (i^t, j^t) such that:

- $1 \leq i^t \leq d$, and $1 \leq j^t \leq m$, for all $t = 1 \dots k$;
- for every $t \neq s$ we have $\min \{ |i^t - i^s|, |j^t - j^s| \} \geq \Delta$; and

- $\sum_{t=1}^k c_{i^t, j^t}$ is maximized.

We refer by $2D\Delta\text{-separated}(c, k, \Delta)$ to that problem, and show it is NP-hard by reducing it to problem Box-Pack studied in [22]. Let us start by recalling the definition of Box-Pack.

Let k be a number of identical 3×3 squares.* Our goal is to pack the k squares into a region of plane defined by set R . Set R consists of pairs of integers, where every pair represents the point at which a square can be placed, e.g. the upper-left corner of a square. Every square can only be placed so that its sides are parallel to the axis. By $\text{Box-Pack}(k, R)$ we refer to the problem of answering whether for given k and R one can place all the k squares in the described way, such that no two squares overlap. In [22] is proved the following result.

Theorem 5 (Theorem 2 in [22]) *Box-Pack is NP-complete.*

Now we utilize Theorem 5 to show that $2D\Delta\text{-separated}(c, k, \Delta)$ is NP-hard as well.

Theorem 6 *Problem $2D\Delta\text{-separated}(c, k, \Delta)$ is NP-hard.*

Proof. We provide a polynomial time reduction of $\text{Box-Pack}(k, R)$ to a sequence of $2D\Delta\text{-separated}$ instances.

Let R^E , denoting "extended R ", be the set of all integers points that can be occupied by a square, not necessarily its corners. Observe that $|R^E| \leq 9|R|$. We say that two points in R^E are adjacent if they share the same x - or the same y -coordinate. Consider a connected component F of R^E .

Assume that $\text{Box-Pack}(i, F)$ can be reduced to an $2D\Delta\text{-separated}$ instance. Then, let us show that $\text{Box-Pack}(k, R)$ can be reduced to polynomially many $2D\Delta\text{-separated}$ instances.

Let \mathcal{F} be the set of all connected components of R^E . Then, for every $C \in \mathcal{F}$ define val_F as follows

$$val_F := \arg \max_{i \in \{0, \dots, \min\{|F|, k\}\}} \text{Box-Pack}(i, F) \text{ equals true.}$$

Now, clearly, if $\sum_{F \in \mathcal{F}} val_F \geq k$, then $\text{Box-Pack}(k, R)$ equals true. Observe that in order to compute all val_F , we need to compute only polynomially many instances of Box-Pack , where all the input parameters are bounded by k and R . So, it remains to reduce $\text{Box-Pack}(i, F)$ to an instance of $2D\Delta\text{-separated}$.

To that end, consider $F \in \mathcal{F}$. Let h and w be the smallest values such that F is enclosed by an axis-parallel sides rectangle T of height h and width w . With loss of generality, assume that T and F are translated so that the bottom-left corner of T at $(1, 1)$. Define vector $c \in \{0, 1\}^{h \times w}$ to be an indicator vector of points of C at which can be placed square, i.e.

$$c_{i,j}^F = \begin{cases} 1 & \text{if } 3 \times 3 \text{ square with upper-left corner at } (i, j) \text{ is in } F \\ 0 & \text{otherwise} \end{cases}$$

Note that $hw \leq |F|^2$. Next, define opt_F as follows

$$opt_F := \max_{i \in \{1, \dots, \min\{|F|, k\}\}} 2D\Delta\text{-separated}(c^F, i, 3).$$

Then, we claim opt_F equals val_F . Now, it is easy to see that the claim is true. Value val_F represents the maximal number of squares that can be packed within F . Each such square corresponds to an entry of c^F which has value 1. So, we have $2D\Delta\text{-separated}(c^F, val_F, 3) = val_F$, and hence $opt_F \geq val_F$. On the other hand, by the construction of c^F , opt_F represents a number of squares (not necessarily maximum, though) that can be packed in F . So, we have $val_F \geq opt_F$, and therefore $opt_F = val_F$ as claimed.

This concludes our proof. \square

*We choose sides to be of length 3 as already that setting is sufficient to prove NP-hardness of Box-Pack, see the proof of Theorem 2 in [22].

E Uniform size neuronal spike trains

We propose a generalization of the Δ -separated model. In this model, we assume that neuronal spike trains correspond to blocks of uniform size. We use b to denote the size of blocks. Formally, we define model $\mathcal{M}_{k,\Delta,b}$ to be the set of all vector in \mathbb{R}^d such that non-zero entries are grouped in blocks of size b , there are exactly k blocks, and every two blocks are separated by at least Δ .

In order to apply the existing framework for recovering measured vectors that belong to $\mathcal{M}_{k,\Delta,b}$, we develop a method that solves the following projection problem. Given vector $c \in \mathbb{R}^d$, the goal is to output k indices p_1, \dots, p_k such that

- $1 \leq p_i \leq d - b + 1$, for every $i = 1 \dots k$;
- $|p_i - p_j| \geq \Delta + b - 1$, for all $i \neq j$; and,
- the sum

$$\sum_{i=1}^k \sum_{j=0}^{b-1} c_{p_i+j}$$

is maximized.

We refer to that problem by (Δ, b) -separated. Now, it is easy to show the following claim.

Theorem 7 Given vector $c \in \mathbb{N}_+^d$, sparsity k , and block size b , define vector $c^b \in \mathbb{N}_+^{d-b+1}$ as

$$c_i^b := \sum_{j=i}^{i+b-1} c_j.$$

Let S be a set of indices corresponding to an optimal solution for $(\Delta + b - 1)$ -separated problem for cost vector c^b and sparsity k . Then, S is an optimal solution to (Δ, b) -separated problem for cost vector c and sparsity k .

Proof. Clearly, S is feasible choice of indices for (Δ, b) -separated problem. Towards a contradiction, assume that there exists another set of indices S' that achieves larger value than S .

But then, S' is a feasible choice of indices for $(\Delta + b - 1)$ -separated problem. Next, notice that S , and also S' , achieve the same value for both (Δ, b) - and $(\Delta + b - 1)$ -separated problem. However, as S' achieves higher value than S , this contradicts our assumption that S is an optimal choice of indices for $(\Delta + b - 1)$ -separated problem. \square

Now we can solve (Δ, b) -separated problem in nearly-linear time.

Corollary 5 A solution to (Δ, b) -separated problem can be computed in nearly-linear time.

Proof. We use the reduction from Theorem 7 to reduce (Δ, b) -separated problem to $(\Delta + b - 1)$ -separated one. The reduction can be applied in linear time. Then, the claim follows by Theorem 12. \square

E.1 Sample complexity

Next we analyze the sample complexity of $\mathcal{M}_{k,\Delta,b}$ model. To that end, we count the number of support in that model. Now, all the supports are captured by $(\alpha_0, \alpha_1, \dots, \alpha_k)$, where $\alpha_i \geq 0$ and $\sum_{i=0}^k \alpha_i = d - kb - (k - 1)\Delta$. The idea is that k blocks split the vector in $k + 1$ regions. All but the first and last region correspond to coordinates between two neighboring blocks. Value α_0 and α_k correspond to coordinates before the first block and after the last block, respectively. Value α_i , for $1 \leq i < k$ correspond to the slack distance between block i and $i + 1$.

Let $\mathbb{M}_{k,\Delta,b}$ be the family of support patterns of $\mathcal{M}_{k,\Delta,b}$. Then, we have:

$$|\mathbb{M}_{k,\Delta,b}| = \binom{d - kb - (k - 1)\Delta + k}{k}.$$

Now, from the result that there are RIP matrices with $O(k + \log |\mathbb{M}_{k,\Delta,b}|)$ rows, [3], the following claim follows.

Theorem 8 *The sample complexity of $\mathcal{M}_{k,\Delta,b}$ is $O(k \log(d/k - (b + \Delta) + (\Delta - 1)/k))$. In particular, if $(\Delta - 1)/k \in O(1)$, then the sample complexity is $O(k \log(d/k - (b + \Delta)))$.*

F Omitted proofs from Section 5.1 and study of \mathcal{D}

Our running-time results are given with respect to γ , where γ is the maximal number of bits needed to store any c_i . As \mathcal{P} and \mathcal{D} are invariant under shifting and multiplication of c , for the sake of clarity of our exposition and without loss of generality we assume c is an integral vector with non-negative entries.

Lemma 8 *The constraint matrix of problem \mathcal{P} is totally unimodular (TUM).*

Proof. We first rewrite the constraints of \mathcal{P} to be in the form $Bu \leq b, u \geq 0$, where $Bu \leq b$ is defined as

$$\begin{aligned} \sum_{i=1}^d u_i &\leq k \\ \sum_{i=1}^d -u_i &\leq -k \\ \sum_{j=i}^{\min\{i+\Delta-1, d\}} u_j &\leq 1 \quad \forall i = 1 \dots d \end{aligned}$$

Let D be a square submatrix of B . If we show that $\det D \in \{-1, 0, 1\}$, then by the definition of TUM the lemma follows.

Now, D is either a binary *interval matrix*, or it has one row where all the non-zero entries are consecutive and equal -1 and the other rows constitute a binary interval matrix. Let D' be a matrix defined as $D'_{i,j} := |D_{i,j}|$. Then, we have $|\det D| = |\det D'|$. Also, we have that D' is a binary interval matrix. However, it is well known that binary interval matrices are TUM, see [18]. Hence, $\det D \in \{-1, 0, 1\}$ \square

Corollary 2 *For an integer k and a vector of integers c , there exists \hat{w} such that \hat{w} is an optimum of \mathcal{D} and $\hat{w}_0 \in \mathbb{Z}$.*

Proof. By Lemma 8, problem \mathcal{P} is totally unimodular. By [23] we have that \mathcal{D} is totally unimodular as well. This implies that \mathcal{D} has an integral optimal solution. \square

The crucial property of the LP \mathcal{D} is the following: for a fixed value of the variable w_0 , we can solve \mathcal{D} in *linear* time. Specifically, let us define \mathcal{D}_v to be the LP \mathcal{D} in which the variable w_0 is set to v . (Note that by Corollary 2, we can restrict our attention to integer values of v .) Observe that once we fixed the value of w_0 , all remaining constraints in \mathcal{D}_v are “local” since they only affect a known interval of length Δ . They are also ordered in a natural way. As a result, we can solve \mathcal{D}_v by making a single pass over these variables. Starting with w_1 and all variables set to 0, we consider each constraint from left to right and increase the variables to satisfy these constraints in a lazy manner. That is, if reach in our pass a constraint with index i that is still not satisfied, we increase the value of w_i until that constraint becomes satisfied and then move to the next constraint and index. Algorithm 2, called Dual-Greedy, present in Section 5 formalizes this approach.

Lemma 9 *For any c' , and k' , the algorithm $\text{Dual-Greedy}(c', v)$ computes an optimal solution $\mathcal{D}_v(c', k')$ in linear time.*

Proof. Let $c' \in \mathbb{N}_+^{d'}$. We show that the algorithm $\text{Dual-Greedy}(c', v)$ outputs vector \tilde{w} in $O(d')$ time such that $\mathcal{D}_v(c', k')$ is minimized and $\tilde{w}_0 = v$. To show the running time, observe that every iteration in the for loop takes $O(1)$ time, so the total algorithm runs in $O(d')$ time.

Next, it is easy to see that \tilde{w} is a feasible solution to $\mathcal{D}_v(c', k')$. We show that it is a minimal as well.

Towards a contradiction, assume there is a vector w' such that $w'_0 = \tilde{w}_0 = v$, w' is a feasible solution to $\mathcal{D}(c', k')$, and $\|w'\|_1 < \|\tilde{w}\|_1$. Then, there exists an index j such that $w'_j \neq \tilde{w}_j$ and $w'_i = \tilde{w}_i$ for all $i < j$. In case there are multiple vectors w' , let w' be one that maximizes the location of mismatch j . Consider the two possible cases: $w'_j < \tilde{w}_j$, and $w'_j > \tilde{w}_j$.

Case $w'_j < \tilde{w}_j$. Observe that \tilde{w}_j is chosen as a function of $\tilde{w}_0, \dots, \tilde{w}_{j-1}$ as a minimal value so that \tilde{w} is feasible. Therefore, if $w'_j < \tilde{w}_j$ and $w'_i = \tilde{w}_i$ for all $i < j$, then w' could not be feasible.

Case $w'_j > \tilde{w}_j$. First, note that $j < d'$, as otherwise $\|w'\|_1 > \|\tilde{w}\|_1$. Now, we construct w'' as follows. We set $w''_i = w'_i$ for all i different than j and $j+1$. Set $w''_j = \tilde{w}_j$ and $w''_{j+1} = w'_{j+1} + \tilde{w}_j - w'_j$. Clearly, w'' is also a feasible solution to $\mathcal{D}(c', k')$. Furthermore, $\|w''\|_1 = \|w'\|_1$ and \tilde{w} and w'' agrees on first j coordinates, contradicting our choice of w' .

This concludes the proof. \square

Now, to obtain an algorithm that is also able to solve the original dual LP \mathcal{D} (instead of only \mathcal{D}_v) it suffices to provide a procedure for choosing the optimal value of v . A priori, there can be many possible choices of v and thus an exhaustive search would be prohibitive. Fortunately, \mathcal{D}_v is actually convex in v , and so we can use a ternary search over v to find such an optimal value.

Lemma 10 $\mathcal{D}_v(c', k')$ is convex in v .

Proof. Let \tilde{w} and \hat{w} be w -solution to $\mathcal{D}_{\tilde{v}}(c', k')$ and $\mathcal{D}_{\hat{v}}(c', k')$, respectively. Note that $\hat{w}_0 = \hat{v}$ and $\tilde{w}_0 = \tilde{v}$. Now if we show that

$$\mathcal{D}_{\frac{\hat{v}+\tilde{v}}{2}}(c', k') \leq \frac{\mathcal{D}_{\hat{v}}(c', k') + \mathcal{D}_{\tilde{v}}(c', k')}{2},$$

the convexity will follow. We start by showing that $\frac{\hat{w}+\tilde{w}}{2}$ is a feasible w -vector for the dual, assuming that both \hat{w} and \tilde{w} are feasible. We consider the feasibility of each of the constraints.

(1) From

$$\hat{w}_0 + \sum_{j: j \leq i \leq j+\Delta-1 \text{ and } j \geq 1} \hat{w}_j \geq c'_i$$

and

$$\tilde{w}_0 + \sum_{j: j \leq i \leq j+\Delta-1 \text{ and } j \geq 1} \tilde{w}_j \geq c'_i$$

we have

$$(\hat{w}_0 + \tilde{w}_0) + \sum_{j: j \leq i \leq j+\Delta-1 \text{ and } j \geq 1} (\hat{w}_j + \tilde{w}_j) \geq 2c'_i,$$

which implies

$$\frac{\hat{w}_0 + \tilde{w}_0}{2} + \sum_{j: j \leq i \leq j+\Delta-1 \text{ and } j \geq 1} \frac{\hat{w}_j + \tilde{w}_j}{2} \geq c'_i.$$

(2) Also, from $\hat{w}_j \geq 0$ and $\tilde{w}_j \geq 0$ we have $\frac{\hat{w}_j + \tilde{w}_j}{2} \geq 0$.

(3) Trivially, $\frac{\hat{w}_0 + \tilde{w}_0}{2} \in \mathbb{R}$.

So, indeed $\frac{\hat{w}+\tilde{w}}{2}$ is feasible. Hence, we have

$$\begin{aligned} \mathcal{D}_{\frac{\hat{v}+\tilde{v}}{2}} &\leq \frac{\hat{w}_0 + \tilde{w}_0}{2} k' + \sum_{i=1}^{d'} \frac{\hat{w}_i + \tilde{w}_i}{2} \\ &= \frac{(\hat{w}_0 k' + \sum_{i=1}^{d'} \hat{w}_i) + (\tilde{w}_0 k' + \sum_{i=1}^{d'} \tilde{w}_i)}{2} \\ &= \frac{\mathcal{D}_{\hat{v}}(c', k') + \mathcal{D}_{\tilde{v}}(c', k')}{2}. \end{aligned}$$

This completes the proof. \square

Putting these pieces together yields the main theorem of this section. As a reminder, we assume that c is an integral vector with non-negative entries.

Theorem 9 *There exists an algorithm Opt-Value-of- $\mathcal{D}(c', k')$ that, for any c' and k' , outputs an optimal solution w^* to $\mathcal{D}(c', k')$ in time $O(d'(\log c'_{max} + \log k'))$.*

Proof.

Bounding w_0^* . Following Lemma 10, we can use ternary search to find w_0^* . However, to be able to do that, we have to know the interval $[a, b]$ we are searching over for w_0^* . By the following lemma we give an upper bound on b . However, in order to provide a lower bound on a , we need to develop some more machinery. So, we defer its proof to later sections, and in Lemma 18 we show that a can be lower bounded by $-(k-1)c_{max}$.

Lemma 11 *Let w^* be a vector that minimizes \mathcal{D} . Then, $w_0^* \leq c_{max}$.*

Proof. Towards a contradiction, let w^* be an optimal vector such that $w_0^* > c_{max}$. Now, as $w_i \geq 0$ for all $i \geq 1$, we have that the corresponding objective is at least $kw_0^* > kc_{max}$. On the other hand, consider vector \hat{w} such that $\hat{w}_0 = c_{max}$ and $\hat{w}_i = 0$ for all $i \geq 1$. Clearly, \hat{w} is a feasible solution to \mathcal{D} . However, the objective function corresponding to \hat{w} is $kc_{max} < kw_0^*$, which contradicts our assumption that w^* is a minimizer of \mathcal{D} . \square

A nearly-linear time algorithm. Now we provide an algorithm that computes the optimal value of $\mathcal{D}(c', k')$ in nearly linear time. It employs ternary search over the interval provided by Lemma 18 and Lemma 11 in order to find w_0 that optimizes $\mathcal{D}(c', k')$. At every step of the search, it uses the result from Lemma 9 to find an optimal solution to $\mathcal{D}_v(c', k')$, for v chosen at the current search step.

Algorithm 3 Opt-Value-of- \mathcal{D} :

Input: $c' \in \mathbb{N}_+^{d'}$, sparsity k' , lower bound lb (if not specified, the default value is $-(k'-1)c'_{max}$), upper bound ub (if not specified, the default value is c'_{max})

Output: a minimizer w^{best} to $\mathcal{D}(c', k')$ constrained to $w_0^{\text{best}} \in [lb, ub]$; if lb and ub are not specified, w^{best} is an optimal solution to $\mathcal{D}(c', k')$

1. $s \leftarrow lb, e \leftarrow ub$
 2. While $s \leq e$
 1. $l \leftarrow s + \lfloor \frac{e-s}{3} \rfloor, r \leftarrow e - \lfloor \frac{e-s}{3} \rfloor$
 2. $(active^l, w^l) \leftarrow \text{Dual-Greedy}(c, l)$
 3. $(active^r, w^r) \leftarrow \text{Dual-Greedy}(c, r)$
 4. If $kw_0^l + \sum_{i=1}^{d'} w_i^l \leq kw_0^r + \sum_{i=1}^{d'} w_i^r$ then $e \leftarrow r - 1, w^{\text{best}} \leftarrow w^l$
 5. Else $s \leftarrow l + 1, w^{\text{best}} \leftarrow w^r$
 3. Return w^{best}
-

Lemma 9, Lemma 10, Corollary 2 and bounds on a and b , the theorem follows directly. \square

G A detailed proof of Lemma 5

Algorithm Δ -Separated-no-Sparsity solves $\mathcal{P}_{\text{no-}k}(\lambda)$ (and so PROJLAGR).

Algorithm 4 Δ -Separated-no-Sparsity:

Input: $c \in \mathbb{R}^d, \lambda \in \mathbb{R}$
Output: an integral optimal solution of $\mathcal{P}_{\text{no-}k}(\lambda)$

```

1:  $\tilde{c} \leftarrow c - \mathbb{1}\lambda$ 
2: if  $\tilde{c}_1 \geq 0$  then
3:    $best_1 \leftarrow \tilde{c}_1, pick_1 \leftarrow 1$ 
4: else
5:    $best_1 \leftarrow 0, pick_1 \leftarrow -1$ 
6: for  $i \leftarrow 2 \dots \Delta$  do
7:   if  $\tilde{c}_i > best_{i-1}$  then
8:      $best_i \leftarrow \tilde{c}_i, pick_i \leftarrow i$ 
9:   else
10:     $best_i \leftarrow best_{i-1}, pick_i \leftarrow pick_{i-1}$ 
11: for  $i \leftarrow \Delta + 1 \dots d$  do
12:   if  $\tilde{c}_i + best_{i-\Delta} > best_{i-1}$  then
13:      $best_i \leftarrow \tilde{c}_i + best_{i-\Delta}, pick_i \leftarrow i$ 
14:   else
15:      $best_i \leftarrow best_{i-1}, pick_i \leftarrow pick_{i-1}$ 
16:  $r \leftarrow \emptyset, i \leftarrow d$ 
17: while  $i \geq 1$  do
18:   if  $pick_i \geq 1$  then  $r \leftarrow r \cup \{pick_i\}$ 
19:    $i \leftarrow pick_i - \Delta$ 
20: return  $r$ 

```

Lemma 12 Algorithm Δ -Separated-no-Sparsity(c, λ) solves $\mathcal{P}_{\text{no-}k}(\lambda)$ (and so PROJLAGR(λ, c) in $O(d)$ time.

Proof. Let us rewrite $\mathcal{P}_{\text{no-}k}(\lambda)$ as follows

$$\begin{aligned}
& \text{maximize} && (c^T - \lambda \mathbb{1}^T)u + \lambda k \\
& \text{subject to} && \sum_{j=i}^{\min\{i+\Delta-1, d\}} u_j \leq 1 \quad \forall i = 1 \dots d \\
& && u_i \geq 0 \quad \forall i = 1 \dots d
\end{aligned}$$

Observe that for a given λ and k , the term λk is constant, so our objective becomes $(c^T - \lambda \mathbb{1}^T)u = \tilde{c}^T u$ subject to the given set of constraints. Now we can show that Δ -Separated-no-Sparsity(c, λ) solves this formulation, which is equivalent to the original one of $\mathcal{P}_{\text{no-}k}(\lambda)$. We proceed by induction. We show that $best_i$ stores the maximum sum of the elements of $\{\tilde{c}_1, \dots, \tilde{c}_i\}$ so that any two chosen elements \tilde{c}_i and \tilde{c}_j are such that $|i - j| \geq \Delta$, and $pick_i$ keeps what is the largest index of element that should be taken to achieve $best_i$ (if there is no such element, then $pick_i$ equals -1).

Base of induction, $i = 1$. For $i = 1$, the lines 2-5 set $best_1$ and $pick_1$ properly.

Inductive step, $i > 1$. For $i = 2 \dots \Delta$, the loop at line 6 sets $best_i$ and $pick_i$ as required. That is, either \tilde{c}_i is the largest among $\{\tilde{c}_1, \dots, \tilde{c}_i\}$, or the largest element is among $\{\tilde{c}_1, \dots, \tilde{c}_{i-1}\}$ which is properly set in $best_{i-1}$ and $pick_{i-1}$. If there is no element with positive value, then as $best_i = 0$ and $pick_i = -1$ as desired.

Similarly, for $i > \Delta$ either we choose \tilde{c}_i and obtain the remaining of the output over $\{\tilde{c}_1, \dots, \tilde{c}_{i-\Delta}\}$, or we take the best solution over $\{\tilde{c}_1, \dots, \tilde{c}_{i-1}\}$ without including \tilde{c}_i .

Correctness of the output vector r . Since $best_i$ and $pick_i$ are set as described above, the way r is obtained trivially satisfies the constraints of $\mathcal{P}_{\text{no-}k}(\lambda)$. Furthermore, as $best_d$ maximizes the objective of $\mathcal{P}_{\text{no-}k}(\lambda)$, r maximizes the objective of $\mathcal{P}_{\text{no-}k}(\lambda)$ as well.

Running time. Line 1 takes $O(d)$ time. The total number of iterations of loop at line 6 and loop at line 11 is $O(d)$. Every iterations takes $O(1)$ time. The loop at line 17 starts with i being d , and decreases i at every iteration (under natural

assumption $\Delta \geq 1$). The fact that i gets decreases at every iteration comes from the property that $pick_i \leq i$, for every i . Therefore, the while loop takes $O(d)$ iterations as well, while every iteration taking $O(1)$ time. This completes the analysis. \square

H Omitted proofs from Section 6.2

In this section we prove some properties of active constraints. We begin by introducing some notation. In what follows, we will be interested in cost vectors and sparsity parameters other than c and k , respectively. Hence, whenever this is the case, we will denote the corresponding dual LP by $\mathcal{D}(c', k')$.

We next define active constraints algorithmically.

Algorithm 5 Active-Constraints:

Input: $c' \in \mathbb{N}_+^d$, $w \in \mathbb{R}^{d+1}$ a feasible vector of $\mathcal{D}(c', k')$
Output: active constraints of $\mathcal{D}(c', k')$ for w

- 1: $active \leftarrow \emptyset$; $sum_\Delta \leftarrow 0$; $last_active \leftarrow -\infty$
- 2: **for** $i := 1 \dots d'$ **do**
- 3: **if** $i - \Delta \geq 1$ **then** $sum_\Delta \leftarrow sum_\Delta - w_{i-\Delta}$
- 4: $sum_\Delta \leftarrow sum_\Delta + w_i$
- 5: **if** $last_active \leq i - \Delta$ and $c'_i - (w_0 + sum_\Delta) = 0$ **then**
- 6: $active \leftarrow active \cup \{i\}$
- 7: $last_active \leftarrow i$
- 8: **return** $active$

Active constraints provide insight into the structure of $\mathcal{D}(c', k')$ that we can leverage to optimally distribute our sparsity budget over the recursive subproblems. The optimality condition of the dual program can be described in the language of active constraints as follows.

Lemma 13 *Let $w^1 \leftarrow \text{Dual-Greedy}(c', v)$, $active^1 \leftarrow \text{Active-Constraints}(c', w^1)$, $w^2 \leftarrow \text{Dual-Greedy}(c', v + 1)$, and $active^2 \leftarrow \text{Active-Constraints}(c', w^2)$, for some integer v . Then, w^1 is an optimal solution of $\mathcal{D}(c', k')$ iff $|active^1| \geq k' \geq |active^2|$.*

Proof. We first prove two properties of actives constraints. First, observe that from the way algorithm $\text{Active-Constraints}(c', w)$ outputs $active$, it corresponds to tight constraints of $\mathcal{D}(c', k')$ for a given w . More precisely, every such tight constraint either is in $active$, or there is another tight constraint in $active$ which is at most Δ "far to the left". Furthermore, as c' is integral, it is easy to see that tight constraints for $w_0 = v$, for some integer v , and for $w_0 = v - \delta$, for $\delta \in [0, 1)$, are the same. Hence, active constraints for $w_0 = v$ and $w_0 = v - \delta$ are also the same. Putting these observations together, we get the following claim.

Lemma 14 *Let v be an integer and $\delta \in [0, 1)$. Also, let $\hat{w} \leftarrow \text{Dual-Greedy}(c', v)$ and $\hat{w}' \leftarrow \text{Dual-Greedy}(c', v - \delta)$. Then,*

$$\text{Active-Constraints}(c', \hat{w}) = \text{Active-Constraints}(c', \hat{w}').$$

We point out that one can show even stronger statement about tight constraints, not necessarily active tough. Namely, it holds that if a constraint i is tight for \hat{w} being $\text{Dual-Greedy}(c', v)$, then it is tight for any $\text{Dual-Greedy}(c', v')$ such that $v' \leq v$. It follows from the property that for any value v' there is at most one active constraint j in $\text{Active-Constraints}(c', \text{Dual-Greedy}(c', v'))$ such that $0 \leq i - j \leq \Delta$. Therefore, if $w_0 = v'$ gets decreased by "a very small" δ , then the variable, e.g. w_j , corresponding to active constraint will decrease by δ as well. Which in turn results i still being a tight constraint. Hence, the following lemma holds, which we utilize in the sequel.

Lemma 15 *Let $\hat{w} \leftarrow \text{Dual-Greedy}(c', v)$ and $\tilde{w} \leftarrow \text{Dual-Greedy}(c', v')$, for $v' \leq v$. Then, if a constraint i is tight with respect to \hat{w} , it is tight with respect to \tilde{w} as well.*

Using Lemma 14 we can show how $\text{Dual-Greedy}(c', v')$ changes for $v' \in (v - 1, v]$.

Lemma 16 *Let v be an integer and $\delta \in [0, 1]$. By \hat{w} denote the output of $\text{Dual-Greedy}(c', v)$, and by \hat{w}^δ the output of $\text{Dual-Greedy}(c', v - \delta)$. Let $active$ be returned by $\text{Active-Constraints}(c', \hat{w})$. Then*

$$k' \hat{w}_0^\delta + \sum_{i \geq 1} \hat{w}_i^\delta = k' \hat{w}_0 + \sum_{i \geq 1} \hat{w}_i + \delta(|active| - k').$$

Proof. From Lemma 14 we have that $\text{Active-Constraints}(c', \text{Dual-Greedy}(c', v - \delta)) = active$ for all $\delta \in [0, 1]$. Now, by the construction of $active$, it holds

$$\begin{aligned} k' \hat{w}_0^\delta + \sum_{i \geq 1} \hat{w}_i^\delta &= k'(\hat{w}_0 - \delta) + \sum_{i \in active} (\hat{w}_i + \delta) + \sum_{i \geq 1 \text{ and } i \notin active} \hat{w}_i \\ &= k' \hat{w}_0 + \sum_{i \geq 1} \hat{w}_i + \delta(|active| - k'), \end{aligned}$$

as desired. \square

We are now ready to finalize the proof of the lemma. Let us break the equivalence stated in the lemma into two implications, and show they are true.

(\Leftarrow) Let $|active^1| \geq k' \geq |active^2|$ be true. By Lemma 16 and the choice of v , we have that $\mathcal{D}_z(c', k')$ is non-decreasing for $z \in [v, v + \frac{1}{2}]$ and non-increasing for $z \in [v - \frac{1}{2}, v]$. As $\mathcal{D}_z(c', k')$ is convex, we have $\mathcal{D}_z(c', k')$ is minimized for $z := v$.[†]

(\Rightarrow) Let w^1 is an optimal solution of $\mathcal{D}(c', k')$. Recall that $\mathcal{D}_z(c', k')$ is a convex function in z . Then, as w^1 is an optimum of $\mathcal{D}(c', k')$, $\mathcal{D}_x(c', k')$ is non-increasing in z on interval $(-\infty, v]$ and non-decreasing on $[v, \infty)$. But then from Lemma 16 we conclude that it can only happen if $|active^1| \geq k' \geq |active^2|$.

\square

I Omitted proofs from Section 6.3

In this section we finalize the proof of the correctness of our randomized algorithm. Before we delve into details, we introduce some notation. In what follows, we will be interested in cost vectors and sparsity parameters other than c and k , respectively. Hence, whenever this is the case, we will write $\mathcal{P}(c', k')$ to refer to the program \mathcal{P} for cost vector c' and sparsity k' . Similarly, whenever we consider some different cost vector c' and sparsity parameter k' , we will denote the corresponding dual LP by $\mathcal{D}(c', k')$.

Also, as pointed out in other sections, our running-time results are given with respect to γ , where γ is the maximal number of bits needed to store any c_i . As \mathcal{P} and \mathcal{D} are invariant under shifting and multiplication of c , for the sake of clarity of our exposition and without loss of generality we assume c is an integral vector with non-negative entries.

Lemma 3 *Algorithm LASSP runs only a single iteration with probability at least $1 - 1/d$.*

I.1 A proof of Lemma 3

As pointed out already, without loss of generality in this proof we assume that c is an integral non-negative vector. We also recall that we showed equivalence between problem (1) and \mathcal{P} , and also between \mathcal{P}_{no-k} and PROJLAGR, so in this proof we work with the LP formulations.

The choice of λ is optimal. Consider \tilde{w} as in Lemma 7. First we want to show that if such \tilde{w} exists, then λ obtained at line 4 of Algorithm 1 is such that it also defines k active constraints. This in turn would imply, by Lemma 7, that support \hat{S} obtained at line 5 has cardinality k , and hence is an optimal solution to $\mathcal{P}(\tilde{c})$.

If $\tilde{w}_0 = \lambda$, then we are done. Otherwise, assume that $\tilde{w}_0 \neq \lambda$. By Lemma 13, \tilde{w} is an optimal solution of $\mathcal{D}(\tilde{c})$. On the other hand, as we discussed in Section 5, λ is such that $\text{VAL } \mathcal{D}_\lambda(\tilde{c}) = \text{VAL } \mathcal{D}(\tilde{c})$, and $\tilde{w}_0 < \lambda$ by the choice of λ . Therefore, by

[†]We use the fact that from the convexity of $\mathcal{D}_z(c', k')$ we have that $\mathcal{D}_z(c', k')$ is continuous.

Lemma 13 it holds that Dual-Greedy(\tilde{c}, λ) defines at least k active constraints. Furthermore, as $\mathcal{D}(\tilde{c})$ is convex w.r.t. to the variable w_0 , then for every $\lambda' \in [\tilde{w}_0, \lambda]$ we have $\text{VAL } \mathcal{D}_{\lambda'}(\tilde{c}) = \text{VAL } \mathcal{D}(\tilde{c})$. In other words, $\text{VAL } \mathcal{D}_{\lambda'}(\tilde{c})$ remains constant over the given interval. Hence, from Lemma 16 we conclude that Dual-Greedy(\tilde{c}, λ) defines exactly k active constraints.

However, for \tilde{c} given on the input there might not exist any λ such that Dual-Greedy(\tilde{c}, λ) defines exactly k active constraints. Our goal is to show that the randomization we apply assures that for the obtained \tilde{c} it is always the case that there is some λ so that Dual-Greedy(\tilde{c}, λ) has exactly k active constraints.

The evolution of active constraints. Now we want to show that the randomization we apply will result in an existence of \tilde{w} as described in Lemma 7. We start by studying the evolution of active constraints defined by the output of Dual-Greedy(\tilde{c}, w_0) as w_0 decreases.

First, recall that by Lemma 15 we have that if a constraint becomes tight with respect to some $\hat{w} \leftarrow \text{Dual-Greedy}(\tilde{c}, \lambda)$, it remains tight with respect to $\hat{w}' \leftarrow \text{Dual-Greedy}(\tilde{c}, \lambda')$ for every $\lambda' \leq \lambda$. Let T denote the set of tight constraints with respect to \hat{w} , and T' with respect to \hat{w}' . By our discussion $T \subseteq T'$.

Let A and A' be the set of active constraints with respect to Dual-Greedy(\tilde{c}, λ) and Dual-Greedy(\tilde{c}, λ'), respectively. Clearly $A \subseteq T$ and $A' \subseteq T'$. We claim that $|T \cap A'| \leq |A|$. Observe that A is a minimum set of constraints so that every tight constraint is covered (covered in the natural way). However, A is also a maximum set of constraints of T that can be chosen so that no two of them overlap, i.e. so that every two of them are at least Δ apart. That means if one would choose a subset of T larger than $|A|$ then some two constraints would overlap. Hence, such a subset can not consist of only active constraints, and therefore $|T \cap A'| \leq |A|$.

This implies that if the number of active constraints increases by $a > 0$ at certain point, then there are at least some a constraints that became active, but also tight, for the first time. Now we want to study what is the probability that two or more non-tight constraints become tight with respect to Dual-Greedy(\tilde{c}, λ), for any λ .

Let us focus on a single constraint j . Fix randomness of all the X_i for $i \neq j$, i.e. fix \tilde{c}_i for all $i \neq j$. For $X_j = 0$, there are at most $d - 1$ different values of w_0 when any of those constraints becomes tight for the first time. Let W denote the set of these w_0 values. So $|W| < d$. Now, construct set W_j as follows. For each $\lambda \in W$:

- If constraint j is already tight with respect to Dual-Greedy(\tilde{c}, λ) for $X_j = 0$, do nothing.
- Define $\hat{c}_i = \tilde{c}_i$ for $i \neq j$, and $\hat{c}_j = \tilde{c}_j + x_\lambda$, where x_λ is defined as the least value so that constraint j becomes tight for the first time with respect to Dual-Greedy(\hat{c}, λ). Observe that $x_\lambda \geq 0$ and x_λ is integral. If $x_\lambda < d^3$, add x_λ to W_j .

We have $|W_j| \leq |W| < d$. Each of the value of W_j correspond to some value of X_j . Also, observe that for given λ , a constraints can become tight for the first time for at most one value of X_j . In addition, as long as constraint j is not tight it does not affect when the other constraints will become tight, as $w_j = 0$ so w_j has no affect on other constraints. This all implies that there are at most $d - 1$ distinct values of X_j , out of d^3 of them, when constraint j and some other constraint become tight. Therefore,

$$\Pr [\text{constraint } j \text{ and any other constraint become tight for the same value of } w_0 = \lambda] < \frac{d}{d^3} = \frac{1}{d^2}.$$

Now we can apply union bound to conclude

$$\Pr [\text{two constraints become tight for the same value of } w_0 = \lambda] < d \frac{1}{d^2} = \frac{1}{d}.$$

Therefore, after applying randomness, for every value of $w_0 = \lambda$ at most one constraint becomes tight with probability at least $1 - 1/d$. Following our discussion above, this in turn implies that the number of active constraints increases by at most 1 after decreasing the value of w_0 by 1. Therefore, there is \tilde{w} as in Lemma 7 with probability $1 - 1/d$ at least.

Required randomness. Every c_i we perturb by one out of d^3 different values, for which $O(\log d)$ random bits suffices. Therefore, in total we need $O(d \log d)$ random bits.

This concludes the proof.

J Deterministic worst-case nearly-linear time algorithm

In this section we describe our deterministic algorithm. For the sake of clarity, we repeat some of the content presented in earlier sections.

Our algorithm stems from a linear programming view on the separated sparsity recovery. It has already been shown that this LP is totally unimodular [10], which implies that solving the LP provides an integral solution and hence solves separated sparsity recovery. However, the previous work resorts to a black-box approach for solving this LP leading to a prohibitive $O(d^{3.5})$ time complexity. We make a step forward, and study the dual LP, obtaining a method that takes nearly-linear time to find its optimal solution. By the strong duality, the value of the dual is the value of the primal LP as well. However, unfortunately, it is not hard to see that a generic relation between primal and dual LP solutions known as "complementary slackness" does not lead to recovery of the primal optimal solution itself.

To cope with that shortcoming, we further analyze the properties of dual solution. We exhibit very close connection between its structure and the sparsity of the primal solution, which we present via the notion of "active constraints". Intuitively, this structure allows us to characterize the cases in which the complementary slackness in fact provides the primal from a dual optimal solution. Then we use these findings in our algorithm to slightly perturb the input instance, while not affecting the value of the solution, so that it is possible to obtain a primal from a dual solution in the general case, and hence solve the separated sparsity recovery.

In Section 5.1 we defined LP \mathcal{P} , which corresponds to the problem (3), as follows:

$$\begin{aligned}
& \text{maximize} && c^T u \\
& \text{subject to} && \sum_{i=1}^d u_i = k \\
& && \sum_{j=i}^{\min\{i+\Delta-1, n\}} u_j \leq 1 \quad \forall i = 1 \dots d \\
& && u_i \geq 0 \quad \forall i = 1 \dots d
\end{aligned}$$

We also defined the dual LP to \mathcal{P} , denoted by \mathcal{D} , as

$$\begin{aligned}
& \text{minimize} && w_0 k + \sum_{i=1}^d w_i \\
& \text{subject to} && w_0 + \sum_{\substack{j: j \geq 1 \text{ and} \\ j \leq i \leq j + \Delta - 1}} w_j \geq c_i \quad \forall i = 1 \dots d \\
& && w_i \geq 0 \quad \forall i = 1 \dots d \\
& && w_0 \in \mathbb{R}
\end{aligned}$$

As Theorem 9, which proofs appears in Section F, this dual LP has a combinatorial structure that enables us to solve it in nearly-linear time.

Before delving into details, we introduce some notation. In what follows, we will be interested in cost vectors and sparsity parameters other than c and k , respectively. Hence, whenever this is the case, we will write $\mathcal{P}(c', k')$ to refer to the program \mathcal{P} for cost vector c' and sparsity k' . Similarly, whenever we consider some different cost vector c' and sparsity parameter k' , we will denote the corresponding dual LP by $\mathcal{D}(c', k')$.

As already mentioned, our running-time results are given with respect to γ , where γ is the maximal number of bits needed to store any c_i . As \mathcal{P} and \mathcal{D} are invariant under shifting and multiplication of c , for the sake of clarity of our exposition and without loss of generality we assume c is an integral vector with non-negative entries.

At a high-level, we develop our algorithm in three main steps. In Section 5.1 we considered the dual of the LP \mathcal{P} and show that its combinatorial structure can be exploited to compute an optimal solution to it in nearly-linear time. By strong duality, this optimal dual solution gives us then the *value* of the optimal primal solution. But, unfortunately, it does not give us the optimal primal solution itself.

To alleviate this issue, we develop a divide-and-conquer procedure for extracting that optimal primal solution. First, in Section J.1, we demonstrate that by analyzing answers of the dual oracle on perturbed versions of the original problem we

can quickly recover a single non-zero entry of the optimal primal solution. That entry can be used to partition our instance into two smaller subproblems. Then, in Section J.2, we show how to make these two subproblems fully independent by devising an optimal split of the sparsity constraint that they share. This optimal split is extracted from the structure of the dual solutions.

With these components in place, we assemble our final algorithm in Section J.3.

J.1 Recovering a segment of an optimal solution

At this point, we developed a way of computing the optimal value OPT of \mathcal{P} . However, this is not sufficient for our purposes as the separated sparsity problem also requires us to provide the corresponding solution, i.e., a binary vector u^* corresponding to OPT .

One might hope that this primal solution u^* can be inferred from the dual solution \hat{w} that our algorithm Opt-Value-of- $\mathcal{D}(c', k')$ (see Algorithm 2) provides. It is not hard to see, however, that the generic relationship between the optimal primal and optimal dual solutions that the so-called “complementary slackness” provides is not sufficient here.

Therefore, we instead design a problem-specific algorithm for finding the desired primal solution vector u^* . As a first step, we focus on the task of recovering a single segment of u^* . Specifically, given some target segment S of at most Δ consecutive entries, we want to either find a *single* index $i \in S$ such that $u_i^* = 1$, for some optimal primal solution u^* ; or to conclude that $u_i^* = 0$ for all $i \in S$.

To recover such a segment, we define in an adaptive manner a family of cost vectors c^1, \dots, c^t , for some $t \in O(\log \Delta)$ and invoke our dual LP solver Opt-Value-of- $\mathcal{D}(c', k')$ on these cost vectors. As we show, the solutions to these perturbed instances allow us to infer u_j^* for all $j \in S$.

To provide more details, let us fix some Δ -length segment $S = [j_s, j_e]$, i.e., $j_e = \min\{j_s + \Delta - 1, n\}$. We now want to decide whether there is an optimal solution \hat{u} to \mathcal{P} such that $\hat{u}_i = 1$ for some $i \in [j_s, j_e]$. Observe that there might be another optimal solution \tilde{u} to \mathcal{P} . Furthermore, it might be the case that $\tilde{u}_j = 0$ for every index $j \in [j_s, j_e]$, while there is an index $i \in [j_s, j_e]$ such that $\tilde{u}_i = 1$. So while designing the algorithm, we distinguish two cases. First, we analyse the case in which there exists an optimal solution \hat{u} to \mathcal{P} so that $\hat{u}_i = 0$ for every $i \in [j_s, j_e]$. Next, we consider the complementary case in which for every optimal solution \hat{u} , we have that $\hat{u}_i = 1$ for some $i \in [j_s, j_e]$. In the latter case, we recover an index j such that there is a solution \hat{u} for which $\hat{u}_j = 1$ holds.

The first case is captured by the following claim.

Lemma 17 *Let $S \subseteq [d']$ be a set of indices, let $c' \in \mathbb{N}_+^{d'}$ be a coefficient vector, and let k' be the sparsity. Define a vector $c'' \in \mathbb{N}_+^{d'}$ as follows:*

$$c''_i = \begin{cases} 1 & \text{if } i \in S \\ 1 + c'_i & \text{otherwise} \end{cases}.$$

Then, Opt-Value-of- $\mathcal{D}(c'', k')$ equals Opt-Value-of- $\mathcal{D}(c', k') + k'$ iff there exists an optimal solution u^ to $\mathcal{P}(c', k')$ such that $u_i^* = 0$ whenever $i \in S$.*

Proof. Let us show the two direction of equivalence separately.

(\implies) Assume that Opt-Value-of- $\mathcal{D}(c'', k')$ equals Opt-Value-of- $\mathcal{D}(c', k') + k'$. Let \hat{u} be an optimal solution to $\mathcal{P}(c', k')$ for the cost vector given by c'' . Now we want to show that $\hat{u}_i = 0$ for every $i \in S$. Towards a contradiction, assume it is not the case, i.e. there exists $j \in S$ such that $\hat{u}_j = 1$. But then, $\sum_{i: \hat{u}_i=1} (c_i + 1) > \sum_{i: \hat{u}_i=1} c''_i$ as $c'_i + 1 > c''_i$ for $i \in S$, and hence Opt-Value-of- $\mathcal{D}(c'', k') < \text{Opt-Value-of-}\mathcal{D}(c', k') + k'$, contradicting our assumption.

(\impliedby) First, observe that Opt-Value-of- $\mathcal{D}(c'', k') \leq \text{Opt-Value-of-}\mathcal{D}(c', k') + k'$ as $c'' \leq c' + 1$. On the other hand, if there exists an optimal solution u^* to $\mathcal{P}(c', k')$ such that $u_i^* = 0$ whenever $i \in S$, then u^* achieves value Opt-Value-of- $\mathcal{D}(c', k') + k'$ in $\mathcal{P}(c', k')$ for the costs given by c'' . In other words, restricted to the set of indices outside of S , c' is a shifted by 1 variant of c'' . Therefore, we also have Opt-Value-of- $\mathcal{D}(c'', k') \geq \text{Opt-Value-of-}\mathcal{D}(c', k') + k'$. Now this implies Opt-Value-of- $\mathcal{D}(c'', k') = \text{Opt-Value-of-}\mathcal{D}(c', k') + k'$, as desired. \square

Next we provide algorithm Δ -recovery that we use in the proof of the next theorem.

Algorithm 6 Δ -recovery:

Input: $c' \in \mathbb{N}_+^{d'}$, sparsity k' , index j_s

Output: index r as described in Theorem 10

1. $j_e \leftarrow \min\{j_s + \Delta - 1, d'\}$
 2. $s \leftarrow j_s$; $e = j_e$; $OPT \leftarrow \text{Opt-Value-of-}\mathcal{D}(c', k')$
 3. $c''_i \leftarrow \begin{cases} 1 & \text{if } i \in [j_s, j_e] \\ 1 + c'_i & \text{otherwise} \end{cases}$
 4. If $\text{Opt-Value-of-}\mathcal{D}(c'', k') = OPT + k'$ then return -1
 5. While $s < e$
 - (a) $mid \leftarrow \lfloor \frac{s+e}{2} \rfloor$
 - (b) $c''_i \leftarrow \begin{cases} 1 & \text{if } i \in [j_s, j_e] \setminus [s, mid] \\ 1 + c'_i & \text{otherwise} \end{cases}$
 - (c) If $\text{Opt-Value-of-}\mathcal{D}(c'', k') = OPT + k'$ then $e \leftarrow mid$
 - (d) Else $s \leftarrow mid + 1$
 6. Return s
-

Theorem 10 *There exists an algorithm that given $c' \in \mathbb{N}_+^{d'}$, sparsity k' , and an index $j_s \in [1, d']$, outputs an integer r having the following properties:*

- *If for every optimal solution \hat{u} to $\mathcal{P}(c', k')$ there is an index i such that $\hat{u}_i = 1$ and $i \in [j_s, \min\{j_s + \Delta - 1, d'\}]$, then r is set to be an index in $[j_s, \min\{j_s + \Delta - 1, d'\}]$ such that there is an optimal solution u^* for which we have $u_r^* = 1$.*
- *Otherwise, r is set to -1.*

Furthermore, if the algorithm runs in time $O((\log \Delta)d'(\log(c'_{max} + 1) + \log k'))$.

Proof. On line 4, Δ -recovery first checks whether we can simply ignore all the entries indexed by $\{i, \dots, \min\{i + \Delta - 1, d'\}\}$. And if yes, it returns -1. (This step is formalized in the statement of Theorem 10.) However, entries from those interval can only be disregarded if there exists an optimal solution u^* to $\mathcal{P}(c', k')$ such that for every $j \in \{i, \dots, \min\{i + \Delta - 1, d'\}\}$ we have $u_j^* = 0$. So, if there is no such u^* , i.e. the entries corresponding to that interval have to be considered, the lines following line 4 of Δ -recovery serve to pinpoint an entry $j \in \{i, \dots, \min\{i + \Delta - 1, d'\}\}$ so that there exists an optimal solution \hat{u} to $\mathcal{P}(c', k')$ for which holds $\hat{u}_j = 1$. The way it is done is by applying a binary search over the interval that *can not* be ignored, tracked via variables s and e of the algorithm.

Correctness. Algorithm Δ -recovery outputs -1 correctly by Lemma 17 for $S = \{j_s, \dots, j_e\}$.

Next, we show that when the binary search loop starting at line 5 ends we have $s = e$, i.e. $[s, e]$ corresponds to a single index. As long as $\Delta \geq 1$ and $j_s \leq d'$, at line 1 and line 2 values s and e are initialized so that $s \leq e$, so initially $[s, e]$ is indeed a non-empty interval. Furthermore, as we have $s < e$ in each iteration, it holds $mid < e$. But it also holds $mid \geq s$ (in fact mid equals s only when $s + 1$ equals e). So, updating e to mid at line 5c, or s to $mid + 1$ at line 5d, $[s, e]$ remains a non-empty interval in any iteration. Notice that the update rules also guarantee that we either increase s or decrease e at each iteration, and therefore $[s, e]$ shrinks its size by 1 at least at each iteration. Putting this together, and taking into account the loop-termination condition at line 5, we conclude that after the loop ends it holds $s = e$.

We say that interval $[s, e]$ has property $R_{s,e}$ if there is an optimal solution \hat{u} to $\mathcal{P}(c', k')$ such that $\hat{u}_i = 1$ for some $i \in [s, e]$. Now, if we show that after line 4 at every step of the algorithm interval $[s, e]$ has property $R_{s,e}$, then the proof will follow immediately. That is exactly how we proceed. Namely, we show that if $[s, e]$ has property $R_{s,e}$, then after updating s or e at line 5c or line 5d obtaining s' and e' , respectively, then interval $[s', e']$ will have property $R_{s',e'}$.

Observe that at the beginning of the very first iteration of the loop property $R_{s,e}$, equivalent to R_{j_s, j_e} , holds as by Lemma 17 every optimal solution \hat{u} is such that $\hat{u}_i = 1$ for some $i \in [j_s, j_e]$. Next, let $[s, e]$ be updated to $[s', e']$. We want to show that property $R_{s',e'}$ holds as well.

First, assume that line 5c gets executed, i.e. $s' = s$ and $e' = mid$. Then, by Lemma 17, we have that there exists an optimal solution \hat{u} to $\mathcal{P}(c', k')$ such that $\hat{u}_i = 0$ for every $i \in [j_s, j_e] \setminus [s, mid]$. But we also have that there is $i \in [j_s, j_e]$ such that $\hat{u}_i = 1$. So, putting it together, we conclude that property $R_{s', e'}$ holds.

Next, assume that the if condition at line 5c does not hold. So, line 5d is executed, i.e. $s' = mid + 1$ and $e' = e$. Then again by Lemma 17, and as a consequence of both line 4 and line 5c, we have that for every optimal solution \hat{u} we have $\hat{u}_i = 1$ for some $i \in [j_s, j_e] \setminus [s, mid]$. Also, as $j_e - j_s + 1 \geq \Delta$, we also have $\hat{u}_j = 0$ for every $j \in [s, mid]$. But then, as by our assumption there is an optimal solution \tilde{u} and an index i such that $\tilde{u}_i = 1$ and $i \in [s, e]$, we have that $i \in [mid + 1, e]$. Hence property $R_{mid+1, e}$, which is equivalent to $R_{s', e'}$, holds.

Running time. If Δ -recovery outputs -1 at line 4, it invokes Opt-Value-of- \mathcal{D} for the cost vector c' and c'' . Note that $c'_{max} \leq c''_{max} + 1$. Then, by Theorem 9 that case takes $O(d'(\log(c'_{max} + 1) + \log k'))$ time.

If the method does not output -1, then it enters the while loop. The loop applies a standard binary search over j_s and j_e which are set so that it holds $j_e - j_s + 1 \leq \Delta$. So, the loop iterates for $O(\log \Delta)$ times. Every iteration invokes Opt-Value-of- \mathcal{D} for vector c'' as defined at line 5b, resulting in the total running time of $O((\log \Delta)d'(\log(c'_{max} + 1) + \log k'))$. \square

J.2 Distributing sparsity

The algorithm we presented in the previous section enables us to quickly recover a target length- Δ segment of some optimal primal solution u^* . We now would like to build on this procedure to develop a divide-and-conquer approach to recovering the primal solution u^* in full.

Our intention is to use the procedure from Section J.1 to split our input problem into two (smaller and approximately equally sized) subproblems by recovering a “middle” segment of the solution u^* and then to proceed recursively on each of these subproblems. The difficulty here, however, is that these two resulting subproblems are not really independent, even though they correspond to separate cost vectors. The subproblems still share the sparsity constraint. That is, to make these subproblems truly independent, we must also specify the split of our sparsity “budget” k' between them.

In this section, we analyze properties of the dual LP and exhibit very close ties between its structure, as captured by active constraints, and the optimal sparsity distribution for our two subproblems. Specifically, we establish the following theorem.

Theorem 11 *Let $c' \in \mathbb{N}_+^{d'}$, let k'' be a sparsity parameter, and let indices s and e be such that $e - s + 1 \geq \Delta$ or e equals d' . Define a vector $c'' \in \mathbb{N}_+^{d'}$ as follows:*

$$c''_i = \begin{cases} -\infty & \text{if } i \in [s, e] \\ c'_i & \text{otherwise} \end{cases}$$

Assume that there is a Δ -separated choice of coordinates of c'' such that k'' coordinates are chosen, and none of them has an index in $[s, e]$, i.e., the instance is feasible when restricted to coordinates outside of $[s, e]$. Then, there is an algorithm that outputs k'_L and k'_R with the property that there exists an optimal Δ -separated solution for cost vector c'' and sparsity k'' such that it chooses k'_L coordinates of c'' with indices less than s , and k'_R coordinates of c'' with indices greater than e .

The algorithm runs in time $O(d'(\log c'_{max} + \log k''))$.

Observe that Theorem 11 essentially gives an algorithm that distributes sparsity in an optimal way and thus enables us to implement the desired divide-and-conquer approach.

Let us illustrate how to use the algorithm Δ -recovery (Algorithm 6) together with Active-Constraints to recover an optimal solution for the example $c = (3, 2, 1, 4, 1, 1, 2, 1)$, $k = 4$ and $\Delta = 2$. Let u^* be an optimal solution to \mathcal{P} . As described above, we proceed by splitting c into two subvectors and solve the separated sparsity problem on each of them independently. The first step is to find the exact splitting point in c . To achieve this, we invoke Δ -recovery($c, 4, k$). As $c_4 = 4$ is part of any optimal solution, the method returns index 4 and hence we have $u_4^* = 1$. So at this moment we have learned one index of the optimal solution. In fact, it also implies that $u_3^* = u_5^* = 0$. Hence it remains to learn the remaining $k - 1$ indices that define u^* . However, we have to distribute the remaining sparsity $k - 1$ over the two subproblems. To that end, we define a new vector c' as follows

$$c' = \begin{cases} -\infty & \text{if } i \in \{3, 4, 5\} \\ c_i & \text{otherwise.} \end{cases}$$

The definition of c' enforces that no optimal solution to the separated sparsity problem for sparsity $k - 1$ will choose c_3 , c_4 , or c_5 (setting c'_i to the dummy value $-\infty$ is for illustration purposes only). Next, we invoke $\text{Active-Constraints}(c', \text{Opt-Value-of-}\mathcal{D}(c', k - 1))$ and denote its output by $active'$. Observe that, as long as $\text{Opt-Value-of-}\mathcal{D}(c, k)$ has finite value, an optimal solution to $\text{Opt-Value-of-}\mathcal{D}(c', k - 1)$ is finite as well. We have that $active' = \{1, 6, 8\}$. Now we count the number of elements in $active'$ that are “to the left” and “to the right” of u_4^* , denoting these quantities by \hat{k} and \tilde{k} respectively, i.e.,

$$\hat{k} = |\{i \in active' \mid i < 4\}|, \quad \text{and} \quad \tilde{k} = |\{i \in active' \mid i > 4\}|.$$

In a similar way, we split the vector c' into \hat{c} and \tilde{c} as follows:

$$\hat{c} = (c_1, c_2), \quad \text{and} \quad \tilde{c} = (c_6, c_7, c_8).$$

Finally, we solve two Δ -separated instances independently, one for \hat{c} and \hat{k} , and the second one for \tilde{c} and \tilde{k} .

Proof of Theorem 11.

Let us start by presenting an algorithm that distributes the sparsity, and continue by analyzing it.

Algorithm 7 Distribute-Sparsity:

Input: $c' \in \mathbb{N}_+^{d'}$, sparsity k'' to be distributed, indices s and e as described in Theorem 11

Output: sparsity k'_L for the left side, sparsity k'_R for the right side

1. $shift \leftarrow 1 + (k''c'_{max} + 1)$
 2. $c''_i \leftarrow \begin{cases} 1 & \text{if } i \in [s, e] \\ c'_i + shift & \text{otherwise} \end{cases}$
 3. $w^1 \leftarrow \text{Opt-Value-of-}\mathcal{D}(c'', k'', -(k'' - 1)c'_{max} + shift, c'_{max} + shift)$
 4. $active^1 \leftarrow \text{Active-Constraints}(c'', w^1)$
 5. $active^2 \leftarrow \text{Active-Constraints}(c'', \text{Dual-Greedy}(c'', w_0^1 + 1))$
 6. $k_L^1 \leftarrow |\{i \in active^1 \mid i < s\}|$
 7. $k_L^2 \leftarrow |\{i \in active^2 \mid i < s\}|$
 8. $k_R^2 \leftarrow |\{i \in active^2 \mid i > e\}|$
 9. $k'_L \leftarrow k_L^2 + \min\{k_L^1 - k_L^2, k'' - k_L^2 - k_R^2\}$
 10. $k'_R \leftarrow k'' - k'_L$
 11. Return (k'_L, k'_R)
-

Optimality for $\mathcal{D}(c'', k'')$. We argue that w^1 defined on line 3 is an optimal solution to $\mathcal{D}(c'', k'')$ although w_0^1 is restricted to belong to interval $[-(k'' - 1)c'_{max} + shift, c'_{max} + shift] = [c'_{max} + 2, c'_{max} + shift]$. First, observe that for $w_0^1 \geq c'_{max} + 2$ no constraint with index in $[s, e]$ is tight, and hence no such constraint can be active. Furthermore, as no constraints in $[s, e]$ is tight, we have $w_i^1 = 0$ for all $i \in [e, s]$. Now, consider $c_L = (c'_1, \dots, c'_{s-1})$ and $c_R = (c'_{e+1}, \dots, c'_{d'})$ as defined in the algorithm. Then, by following the proof of Lemma 18, and as $e - s + 1 \geq \Delta$, we conclude that $\text{Dual-Greedy}(c'', c'_{max} + 2)$ has at least $\min\{k', \lceil \frac{s-1}{\Delta} \rceil\} + \min\{k', \lceil \frac{e-s}{\Delta} \rceil\}$ active constraints which, by the construction and the assumption that the input instance is feasible, is k'' at least. So, by the monotonicity of the size of active constraints given by Lemma 19 and the optimality condition provided via Lemma 13 we have that w^1 is an optimal solution to $\mathcal{D}(c'', k'')$.

Correctness of distributed sparsity. Let $active^1$ denote $\text{Active-Constraints}(c'', w^1)$ and $active^2$ denote $\text{Active-Constraints}(c'', w^2)$, where $w^2 \leftarrow \text{Dual-Greedy}(c'', w_0^1 + 1)$. Next, split w^1 and w^2 as follows. Let w_L^1 be a zero-indexed s -dimensional and w_R^1 be a zero-indexed $(d' - e + 1)$ -dimensional vector defined as

$$w_L^1 \leftarrow (w_0^1, w_1^1, \dots, w_{s-1}^1), \quad \text{and} \quad w_R^1 \leftarrow (w_0^1, w_{e+1}^1, \dots, w_{d'}^1).$$

Intuitively, letter 'L' stands for left to s and letter 'R' stands for right to e . Similarly to w_L^1 and w_R^1 , define w_L^2 and w_R^2 as

$$w_L^2 \leftarrow (w_0^2, w_1^2, \dots, w_{s-1}^2), \quad \text{and} \quad w_R^2 \leftarrow (w_0^2, w_{e+1}^2, \dots, w_{d'}^2).$$

Also, split $active^1$ and $active^2$ with respect to s and e in the obvious way

$$active_L^1 \leftarrow \{i \in active^1 \mid i < s\}, \text{ and } active_R^1 \leftarrow \{i \in active^1 \mid i > e\},$$

and

$$active_L^2 \leftarrow \{i \in active^2 \mid i < s\}, \text{ and } active_R^2 \leftarrow \{i \in active^2 \mid i > e\}.$$

As a reminder, there is no $i \in [s, e]$ such that $i \in active^1$. Furthermore, as there is no tight constraint in $[s, e]$ for w^1 , we have $w_i^1 = 0$ for every $i \in [s, e]$. Observe that the same holds for $active^2$ and w^2 . In addition, we have $e - s + 1 \geq \Delta$. From this, we can derive the following list of equalities, which essentially allows us to split the input problem into two independent subproblems. So, we have: Dual-Greedy(c_L, w_0^1) equals w_L^1 ; Dual-Greedy($c_L, w_0^1 + 1$) equals w_L^2 ; Dual-Greedy(c_R, w_0^1) equals w_R^1 ; and, Dual-Greedy($c_R, w_0^1 + 1$) equals w_R^2 . But also, we have: Active-Constraints(c_L, w_L^1) equals $active_L^1$; Active-Constraints(c_L, w_L^2) equals $active_L^2$; Active-Constraints(c_R, w_R^1) equals $active_R^1$; and, Active-Constraints(c_R, w_R^2) equals $active_R^2$.

Now, by Lemma 19 we conclude that $|active_L^1| \geq |active_L^2|$ and $|active_R^1| \geq |active_R^2|$. This in turn implies that k'_L and k'_R in Distribute-Sparsity are derived correctly. (Note that $k'_L + k'_R = k''$.) But then, by Lemma 13 we have that w_L^1 is an optimal solution to $\mathcal{D}(c_L, k'_L)$ and w_R^1 is an optimal solution to $\mathcal{D}(c_R, k'_R)$, so we can independently solve Recover(c_L, k'_L) and Recover(c_R, k'_R) knowing that the optimal solution, i.e. its value, remains the same. We point out that $w'_L \leftarrow \text{Opt-Value-of-}\mathcal{D}(c_L, k_L)$ might differ from w_L^1 , and similarly $w'_R \leftarrow \text{Opt-Value-of-}\mathcal{D}(c_R, k_R)$ might differ from w_R^1 . In fact, w'_L and w'_R might be such that $w'_{L0} \neq w'_{R0}$, although we have that w_{L0}^1 is equal to w_{R0}^1 .

Running time. Every line, except maybe line 3, take $O(d')$ time. On the other hand, line 3 takes time $O(d' \log \{c'_{max} + shift - ((k'' - 1)c'_{max} + shift)\})$ which is equal to $O(d'(\log c'_{max} + \log k''))$. \square

J.3 Separated sparsity in nearly-linear time

We now have all components in place to state our final algorithm Recover(c', k') that produces an optimal integral solution for $\mathcal{P}(c', k')$ (see Algorithm 8).

Algorithm 8 Recover:

Input: $c' \in \mathbb{N}_+^{d'}$, sparsity k'

Output: a primal solution d' -dimensional vector u^*

1. If $k' = 0$ return $\mathbb{0}$
 2. $r \leftarrow \Delta\text{-recovery}(c', \lfloor \frac{d'+1}{2} \rfloor, k')$
 3. If $r = -1$ then
 1. $s \leftarrow \lfloor \frac{d'+1}{2} \rfloor$, $e \leftarrow \min\{s + \Delta - 1, d'\}$
 2. $k'' \leftarrow k'$
 4. Else
 1. $s \leftarrow \max\{r - \Delta + 1, 1\}$, $e \leftarrow \min\{r + \Delta - 1, d'\}$
 2. $k'' \leftarrow k' - 1$
 5. $(k'_L, k'_R) \leftarrow \text{Distribute-Sparsity}(c', k'', s, e)$
 6. $c_L \leftarrow (c'_1, \dots, c'_{s-1})$, $c_R \leftarrow (c'_{e+1}, \dots, c'_{d'})$
 7. $u^L \leftarrow \text{Recover}(c_L, k'_L)$, $u^R \leftarrow \text{Recover}(c_R, k'_R)$
 8. $u_i^* = \begin{cases} u_i^L & \text{if } i < s \\ 1 & \text{if } i \in [s, e] \text{ and } r = i \\ u_{i-e}^R & \text{if } i > e \\ 0 & \text{otherwise} \end{cases}$
 9. Return u^*
-

We prove the following result for our algorithm.

Theorem 12 *The algorithm Recover(c', k') solves the model projection problem for separated sparsity in time $O(n(\log(c_{max} + 1) + \log k) \log n \log \Delta)$.*

We start by showing some technical lemmas. Recall that by Lemma 11 we provided an upper bound on any w_0^* such that w^* is an optimum of \mathcal{D} . On the other hand, there is no lower bound on w_0^* . To see that, consider a very simple example $c = (1)$, $\Delta = k = 1$. Nevertheless, we can provide a lower bound in the following form.

Lemma 18 *There exists an optimal solution w^* to $\mathcal{D}(c', k')$ such that $w_0^* \geq -(k' - 1)c'_{max}$.*

Proof. To prove the lemma, we utilize the optimality condition described by Lemma 13 and the following claim.

Lemma 19 *Let $active^1$ be the output of $Active-Constraints(c', Dual-Greedy(c', v))$ and $active^2$ be the output of $Active-Constraints(c', Dual-Greedy(c', v - t))$, for any integer v and a positive integer t . Then, it holds*

$$|active^1| \leq |active^2|.$$

Proof. Let S and S' be tight constraints for $Dual-Greedy(c', v)$ and $Dual-Greedy(c', v - t)$, respectively. Then, as we have discussed, $S \subseteq S'$. So, all we have to show is that there are at least as many active constraints formed from S' as there are formed from S . We do that by induction on the size of S under the assumption that $S \subseteq S'$.

Base of induction: $|S| = 0$. As the number of active constraints is non-negative, and if $|S| = 0$ there is no active constraint, the claim follows.

Inductive step: $|S| > d'$ and $S \subseteq S'$, for $d' \geq 0$. Let i_{min} and i'_{min} be the smallest index of S and S' , respectively. Constraint i_{min} is active for $Dual-Greedy(c', v)$, and i'_{min} is active for $Dual-Greedy(c', v - t)$. Define $T = S \setminus \{i_{min}, \dots, i_{min} + \Delta - 1\}$ and $T' = S' \setminus \{i'_{min}, \dots, i'_{min} + \Delta - 1\}$. Since $i'_{min} \leq i_{min}$ it holds $T \subseteq T'$. Now, as for every active constraint j there is no other active one in the neighborhood of $\Delta - 1$ around j and except that neighborhood the other constraints are not affected by j , we have that

$$\begin{aligned} active^1 &= \{i_{min}\} \cup \{\text{active constraints for } T\}, \text{ and} \\ active^2 &= \{i'_{min}\} \cup \{\text{active constraints for } T'\}. \end{aligned}$$

Now, as $T \subseteq T'$ and $|T| < |S|$, by inductive hypothesis we have

$$|\{\text{active constraints for } T\}| \leq |\{\text{active constraints for } T'\}|,$$

and hence the lemma follows. \square

Let $\hat{w} \leftarrow Dual-Greedy(c', -(k - 1)c_{max})$ and $active \leftarrow Active-Constraints(c', \hat{w})$. We show that $|active| \geq k'$. Furthermore, we show that $l\Delta + 1 \in active$, for all $l = 0, \dots, k' - 1$. Once it is shown, the claim follows by Lemma 19 and Lemma 13. Precisely, Lemma 19 shows that the number of active constraints for solutions corresponding to $Dual-Greedy(c', -(k' - 1)c'_{max} - t)$, for $t \geq 0$, is at least k' . But then, from Lemma 13 we have that $Dual-Greedy(c', -(k' - 1)c'_{max})$ or $Dual-Greedy(c', -(k' - 1)c'_{max} + t)$, for some $t > 0$, outputs an optimal solution.

So, it only remain to show $|active| \geq k$. We prove that by induction, showing the following property. Let $i = l\Delta + 1$ be an index for some integer l such that $0 \leq l \leq k' - 1$. If $\hat{w}_0 = -(k' - 1)c'_{max}$, $\sum_{j=\max\{1, i-\Delta+1\}}^{i-1} \hat{w}_j \leq lc'_{max}$ and no constraint in $\{(l - 1)\Delta + 2, \dots, l\Delta\}$ is active, then constraints $l\Delta + 1, (l + 1)\Delta + 1, \dots, (k' - 1)\Delta + 1$ are active. The induction is applied in a downward fashion on l , i.e. the base case is $l = k' - 1$, and our goal is to show it holds for $l = 0$.

Base of induction: $l = k' - 1$. Let $i = (k' - 1)\Delta + 1$. As $\sum_{j=\max\{1, i-\Delta+1\}}^{i-1} \hat{w}_j \leq (k' - 1)c'_{max}$ and $c'_i \geq 0$, we have $\sum_{j=\max\{1, i-\Delta+1\}}^{i-1} \hat{w}_j \leq \hat{w}_0 + c'_i$, and hence constraint i is tight. Furthermore, as no constraint in $\{i - \Delta + 1, \dots, i - 1\}$ is active, constraint i is an active one.

Inductive step: $0 \leq l < k' - 1$. Let $i = l\Delta + 1$. First, if $\sum_{j=\max\{1, i-\Delta+1\}}^{i-1} \hat{w}_j \leq lc'_{max}$ we have that constraint i is tight, and as before active as well, and also we have $\hat{w}_i \geq -\hat{w}_0 - lc'_{max}$.

Next, we want to show that $\sum_{j=i+1}^{i+\Delta} \hat{w}_j \leq (l + 1)c'_{max}$ (note that $i + \Delta \leq d'$), so that we can use the inductive hypothesis for $l + 1$. First, we show that for every index i' it holds

$$\sum_{j=\max\{1, i'-\Delta+1\}}^{i'} \hat{w}_j \leq c'_{max} - \hat{w}_0. \tag{6}$$

Recall that \hat{w}_0 is negative. Towards a contradiction, assume that there exists index q such that $\sum_{j=\max\{1, q-\Delta+1\}}^q c'_{max} - \hat{w}_0 >$. In case of tie, let q be the smallest such index, which implies $\hat{w}_q > 0$. But then, it contradicts the greedy choice of Dual-Greedy algorithm, as by the greedy choice we have $\hat{w}_q = 0$ or $\sum_{j=\max\{1, q-\Delta+1\}}^q c'_q - \hat{w}_0 \leq c'_{max} - \hat{w}_0$.

Now we combine (6), for $i' = i + \Delta - 1$ and $\hat{w}_i \geq -\hat{w}_0 - lc'_{max}$ to obtain

$$\sum_{j=i+1}^{i+\Delta} \hat{w}_j \leq c'_{max} - \hat{w}_0 - \hat{w}_i \leq (l+1)c'_{max},$$

as desired. In addition, as constraint i is active, we have that no constraints in $\{i+1, \dots, i+\Delta-1\}$ is active, and hence we can use the inductive hypothesis.

This concludes the proof. \square

J.4 Proof of Theorem 12

The algorithm `Recover`(c', k') utilizes Δ -recovery and Distribute-Sparsity to split the input problem into two subproblems and then recurses on them. Let us now analyze the correctness and the running time of this algorithm.

Base case. If the input sparsity is 0, the algorithm outputs a zero-vector at line 1.

Δ -middle entries. Next, the algorithm invokes Δ -recovery over the Δ middle entries. Depending on r , it sets s and e to correspond to the interval of c' that should be removed from consideration in the recursive calls. In the same time, that intervals serves to break the input problem into two independent subproblems. It also sets k'' that represents the sparsity distributed outside interval $[s, e]$. Correctness of this call is guaranteed by Theorem 10. Value *shift*, and in turn vector c'' , is set so that we have a guarantee that no element from the interval $[s, e]$ is chosen as part of an optimal solution. To see that, consider vector $c^{(3)} \leftarrow c'' - \text{shift} \cdot \mathbb{1}$. Then, $c_i^{(3)} = c'_i$ for $i \notin [s, e]$ and $c_i^{(3)} = -(k'c'_{max} + 1)$. In other words, as long as $k'' \leq k'$ entries can be chosen outside of the interval $[s, e]$, no optimal solution should choose any entry within the interval $[s, e]$.

Distributing the sparsity. The correctness of line 5 follows by Theorem 11.

Running time. Line 2 of algorithm `Recover` runs in time $O((\log \Delta)d'(\log(c'_{max} + 1) + \log k'))$ which is a subset of $O((\log \Delta)d'(\log(c_{max} + 1) + \log k))$. Vector c' is split into c_L and c_R at line 6 in $O(d')$ time. Afterwards, the method recurses on the two subproblems, and combines their outputs into u^* . Obtaining u^* at line 8 also takes $O(d')$. So, it only to remain to discuss the recursion.

From Theorem 11 we have that line 5 take $O(d'(\log c'_{max} + \log k''))$. Furthermore, as $k'' \leq k$ and $c'_{max} \leq c_{max}$, we have $O(d'(\log c'_{max} + \log k'')) = O(d'(\log c_{max} + \log k))$.

Every recursive step shrinks the corresponding c' vector by half at least. So, the recursion has depth of $O(\log d)$. At every level of recursion are considered vectors c' of total length $O(d)$ – this follows from the recursive call at line 7 and the fact the c_L and c_R represent disjoint pieces of c' . So, this implies the total running time of $O(\log d(d(\log c_{max} + \log k) + (\log \Delta)d(\log(c_{max} + 1) + \log k)))$, which in a more compact way can be written as $O(d(\log(c_{max} + 1) + \log k) \log d \log \Delta)$.