

Discrete Optimization Methods for Group Model Selection in Compressed Sensing

Bubacarr Bah*

Jannis Kurtz†

Oliver Schaudt‡

Abstract

In this article we study the problem of signal recovery for group models. More precisely for a given set of groups, each containing a small subset of indices, and for given linear sketches of the true signal vector which is known to be group-sparse in the sense that its support is contained in the union of a small number of these groups, we study algorithms which successfully recover the true signal just by the knowledge of its linear sketches. We derive model projection complexity results and algorithms for more general group models than the state-of-the-art. We consider two versions of the classical Iterative Hard Thresholding algorithm (IHT). The classical version iteratively calculates the exact projection of a vector onto the group model, while the approximate version (AM-IHT) uses a head- and a tail-approximation iteratively. We apply both variants to group models and analyse the two cases where the sensing matrix is a Gaussian matrix and a model expander matrix.

To solve the exact projection problem on the group model, which is known to be equivalent to the maximum weight coverage problem, we use discrete optimization methods based on dynamic programming and Benders' Decomposition. The head- and tail-approximations are derived by a classical greedy-method and LP-rounding, respectively.

Compressed Sensing, Group Models, Iterative Hard Thresholding, Maximum Weight Coverage

1 Introduction

In many applications involving sensors or sensing systems an unknown sparse signal has to be recovered from a relatively small number of measurements. The reconstruction problem in standard compressed sensing attempts to recover an unknown k -sparse signal $\mathbf{x} \in \mathbb{R}^N$, i.e. it has at most k non-zero entries, from its (potentially noisy) linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$. Here, $\mathbf{A} \in \mathbb{R}^{m \times N}$ for $m \ll N$, $\mathbf{y} \in \mathbb{R}^m$ and $\mathbf{e} \in \mathbb{R}^m$ is a noise vector, typically with a bounded noise level $\|\mathbf{e}\|_2 \leq \eta$; see [19, 13, 14]. A well-known result is that, if \mathbf{A} is a random Gaussian matrix, the number of measurements required for most of the classical algorithms like ℓ_1 -minimization or Iterative Hard Thresholding (IHT) to successfully recover the true signal is $m = \mathcal{O}(k \log(N/k))$ [15, 10].

In model-based compressed sensing we exploit second-order structures beyond the first order sparsity and compressibility structures of a signal to more efficiently encode and more accurately decode the signal. Efficient encoding means taking a fewer number of measurements than in the standard compressed sensing setting, while accurate decoding does not only include smaller recovery error but better interpretability of the recovered solution than in the standard compressed sensing setting. The second order structures of the signal are usually referred to as the structured sparsity of the signal. The idea is, besides standard sparsity, to take into account more complicated structures of the signal [5]. Nevertheless most of the classical results and algorithms for standard compressed sensing can be adapted to the model-based framework [5].

Numerous applications of model-based compressed sensing exist in practice. Key amongst these applications is the multiple measurement vector (MMV) problem, which can be modelled as a block-sparse recovery

*AIMS South Africa, & Stellenbosch University, 6 Melrose Road, Muizenberg, Cape Town 7945, South Africa

†Chair for Mathematics of Information Processing, RWTH Aachen University, Pontdriesch 10, 52062 Aachen, Germany

‡Chair for Mathematics of Information Processing, RWTH Aachen University, Pontdriesch 10, 52062 Aachen, Germany

problem, [21]. The tree-sparse model has been well-exploited in a number of wavelet-based signal processing applications, [5]. In the sparse matrix setting (see Section 2) the model-based compressed sensing was used to solve the Earth Mover Distance problem (EMD). The EMD problem introduced in [44] is motivated by the task of reconstructing time sequences of spatially sparse signals, e.g. seismic measurements; see also [26]. In addition, there are many more potential applications in linear sketching including data streaming [39], graph sketching [1, 25], breaking privacy of databases via aggregate queries [20], and in sparse regression codes or sparse superposition codes (SPARC) decoding [34, 45], which is also an MMV problem.

Structured sparsity models include tree-sparse, block-sparse, and group-sparse models. For instance, for block-sparse models with dense Gaussian sensing matrices it has been established in [5] that the number of required measurements to ensure recovery is $m = \mathcal{O}(k)$ as opposed to $m = \mathcal{O}(k \log(N/k))$ in the standard compressed sensing setting. Furthermore, in the sparse matrix setting, precisely for adjacency matrices of model expander graphs (also known as model expander matrices), the tree-sparse model only requires $m = \mathcal{O}(\log_k(N/k))$ measurements [31, 3], which is smaller than the standard compressed sensing sampling rate stated above. Moreover, all proposed algorithms that perform an exact model projection, which is to find the closest vector in the model space for a given signal, guarantee recovery of a solution belonging to the model space, which is then more interpretable than applying off-the-shelf standard compressed sensing algorithms [3].

As the exact model projection problem used in many of the classical algorithms may become theoretically and computationally hard for specific sparsity models, approximation variants of some well-known algorithms like the Model-IHT have been introduced in [27]. Instead of iteratively solving the exact model projection problem, this algorithm, called AM-IHT, uses a head- and a tail-approximation to recover the signal which is computationally less demanding in general. The latter computational benefit comes along with a larger number of measurements required to obtain successful recovery with weaker recovery guarantees: the typical speed versus accuracy trade-off.

A special class of structured sparsity models are group models, where the support of the signal is known to be contained in the union of a small number of groups of indices. Group models were already studied extensively in the literature in the compressed sensing context; see [6, 33, 40]. Its choice is motivated by several applications, e.g. in image processing; see [38, 41, 42]. As it was shown in [4] the exact projection problem for group models is *NP*-hard in general but can be solved in polynomial time by dynamic programming if the intersection graph of the groups has no cycles. The latter case is quite restricting, since as a consequence each element is contained in at most two groups. In this work we extend existing results for the Model-IHT algorithm and its approximation variant (AM-IHT) derived in [27] to group models and model expander matrices. We focus on deriving discrete optimization methods to solve the exact projection problem and the head- and tail-approximations for much more general classes of group models than the state-of-the-art.

In Section 2 we present the main preliminary results regarding compressed sensing for structured sparsity models and group models. In Section 3 we study recovery algorithms using exact projection oracles. We first show that for group models with low treewidth, the projection problem can be solved in polynomial time by dynamic programming which is a generalization of the result in [4]. We then adapt known theoretical results for model expander matrices to these more general group models. To solve the exact projection problem for general group models we apply a Benders' decomposition procedure. It can be even used for the more general assumption that we seek a signal which is group-sparse and additionally sparse in the classical sense. In Section 4 we study recovery algorithms using approximation projection oracles, namely head- and tail-approximations. We apply the known results in [27] to group models of low frequency and show that the required head- and tail-approximations for group models can be solved by a classical greedy-method and LP rounding, respectively. In Section 5 we test all algorithms, including Model-IHT, AM-IHT, MEIHT and AM-EIHT, on overlapping block-groups and compare the number of required measurements, iterations and the run-time.

Summary of our contributions:

- We study the Model-Expander IHT (MEIHT) algorithm, which was analysed in [3] for tree-sparse and loopless group-sparse signals, and extend the existing results to general group models, proving convergence of the algorithm.

- We extend the results in [4] by proving that the projection problem can be solved in polynomial time if the incidence graph of the underlying group model has bounded treewidth. This includes the case when the intersection graph has bounded treewidth, which generalizes the result for acyclic graphs derived in [4]. We complement the latter result with a hardness result that we use to justify the bounded treewidth approach.
- We derive a Benders' decomposition procedure to solve the projection problem for arbitrary group models, assuming no restriction on the frequency or the structure of the groups. The latter procedure even works for the more general model combining group-sparsity with classical sparsity. We integrate the latter procedure into the Model-IHT and MEIHT algorithm.
- We apply the Approximate-Model IHT (AM-IHT) derived in [26, 28] to Gaussian and expander matrices and to the case of group models with bounded frequency, which is the maximal number of groups an element is contained in. In the expander case we call the algorithm AM-EIHT. To this end we derive both, head- and tail-approximations of arbitrary precision using a classical greedy method and LP-rounding. Using the AM-IHT and the results in [26, 28], this implies compressive sensing ℓ_2/ℓ_2 recovery guarantees for group-sparse signals. We show that the number of measurements needed to guarantee a successful recovery exceeds the number needed by the usual model-based compressed sensing bound [5, 9] only by a constant factor.
- We test the algorithms Model-IHT, MEIHT, AM-IHT and AM-EIHT on groups given by overlapping blocks for random signals and measurement matrices. We analyse and compare the minimal number of measurements needed for recovery, the run-time and the number of iterations of the algorithm.

2 Preliminaries

2.1 Notation

In most of this work scalars are denoted by ordinary letters (e.g. x , N), vectors and matrices by boldface letters (e.g. \mathbf{x} , \mathbf{A}), and sets by calligraphic capital letters (e.g., \mathcal{S}). The cardinality of a set \mathcal{S} is denoted by $|\mathcal{S}|$ and we define $[N] := \{1, \dots, N\}$. Given $\mathcal{S} \subseteq [N]$, its complement is denoted by $\mathcal{S}^c := [N] \setminus \mathcal{S}$ and $\mathbf{x}_{\mathcal{S}}$ is the restriction of $\mathbf{x} \in \mathbb{R}^N$ to \mathcal{S} , i.e.

$$(\mathbf{x}_{\mathcal{S}})_i = \begin{cases} x_i, & \text{if } i \in \mathcal{S}, \\ 0 & \text{otherwise.} \end{cases}$$

The support of a vector $\mathbf{x} \in \mathbb{R}^N$ is defined by $\text{supp}(\mathbf{x}) = \{i \in [N] \mid x_i \neq 0\}$. For a given $k \in \mathbb{N}$ we say a vector $\mathbf{x} \in \mathbb{R}^N$ is k -sparse if $|\text{supp}(\mathbf{x})| \leq k$. For a matrix \mathbf{A} , the matrix $\mathbf{A}_{\mathcal{S}}$ denotes a sub-matrix of \mathbf{A} with columns indexed by \mathcal{S} . For a graph $G = (V, E)$ and $\mathcal{S} \subseteq V$, $\Gamma(\mathcal{S})$ denotes the set of *neighbours* of \mathcal{S} , that is the set of nodes that are connected by an edge to the nodes in \mathcal{S} . We denote by $e_{ij} = (i, j)$ an edge connecting node i to node j . The set \mathcal{G}_i denotes a group of size g_i and a group model is any subset of $\mathfrak{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_M\}$; while a group model of order k is denoted by \mathfrak{G}_k , which is a collection of any k groups of \mathfrak{G} . The ℓ_p norm of a vector $\mathbf{x} \in \mathbb{R}^N$ is defined as

$$\|\mathbf{x}\|_p := \left(\sum_{i=1}^N x_i^p \right)^{1/p}.$$

2.2 Compressed Sensing

Recall that the reconstruction problem in standard compressed sensing [19, 13, 14] attempts to recover an unknown k -sparse signal $\mathbf{x} \in \mathbb{R}^N$, from its (potentially noisy) linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, where

$\mathbf{A} \in \mathbb{R}^{m \times N}$, $\mathbf{y} \in \mathbb{R}^m$ for $m \ll N$ and $\mathbf{e} \in \mathbb{R}^m$ is a noise vector, typically with a bounded noise level $\|\mathbf{e}\|_2 \leq \eta$. The reconstruction problem can be formulated as the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{x}\|_0 \quad \text{subject to} \quad \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2 \leq \eta, \quad (1)$$

where $\|\mathbf{x}\|_0$ is the number of non-zero components of \mathbf{x} . Problem (1) is usually relaxed to an ℓ_1 -minimization problem by replacing $\|\cdot\|_0$ with the ℓ_1 -norm. It has been established that the solution minimizing the ℓ_1 -norm coincides with the optimal solution of (1) under certain conditions [14]. Besides the latter approach the compressed sensing problem can be solved by a class of greedy algorithms, including the IHT [10]. A detailed discussion on compressed sensing algorithms can be found in [22].

The idea behind the IHT can be explained by considering the problem

$$\min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \quad \text{subject to} \quad \|\mathbf{x}\|_0 \leq k. \quad (2)$$

Under certain choices of η and k the latter problem is equivalent to (1) [10]. Based on the idea of gradient descent methods, (2) can be solved by iteratively taking a gradient descent step, followed by a hard thresholding operation, which sets all components to zero except the largest k in magnitude. Starting with an initial guess $\mathbf{x}^{(0)} = \mathbf{0}$, the $(n+1)$ -th IHT update is given by

$$\mathbf{x}^{(n+1)} = \mathcal{H}_k \left[\mathbf{x}^{(n)} + \mathbf{A}^* \left(\mathbf{A}\mathbf{x}^{(n)} - \mathbf{y} \right) \right], \quad (3)$$

where $\mathcal{H}_k : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is the hard threshold operator and \mathbf{A}^* is the adjoint matrix of \mathbf{A} .

Recovery guarantees of algorithms are typically given in terms of what is referred to as the ℓ_p/ℓ_q instance optimality [14]. Precisely, an algorithm has ℓ_p/ℓ_q instance optimality if for a given signal \mathbf{x} it always returns a signal $\hat{\mathbf{x}}$ with the following error bound

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_p \leq c_1 \sigma_k(\mathbf{x})_q + c_2 \eta, \quad (4)$$

where $1 \leq q \leq p \leq 2$, c_1, c_2 are constants independent of the dimension of the signal and

$$\sigma_k(\mathbf{x})_q = \min_{\mathbf{z}: \|\mathbf{z}\|_0 \leq k} \|\mathbf{x} - \mathbf{z}\|_q$$

is the *best k -term approximation* of a signal (in the ℓ_q -norm).

Ideally, we would like to have ℓ_2/ℓ_1 instance optimality [14]. It turned out that the instance optimality of the known algorithms depends mainly on the sensing matrix \mathbf{A} . Key amongst the tools used to analyse the suitability of \mathbf{A} is the restricted isometry property, which is defined in the following.

Definition 1 (RIP). *A matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ satisfies the ℓ_p -norm restricted isometry property (RIP- p) of order k , with restricted isometry constant (RIC) $\delta_k < 1$, if for all k -sparse vectors \mathbf{x}*

$$(1 - \delta_k) \|\mathbf{x}\|_p^p \leq \|\mathbf{A}\mathbf{x}\|_p^p \leq (1 + \delta_k) \|\mathbf{x}\|_p^p. \quad (5)$$

Typically RIP without the subscript p refers to case when $p = 2$. We use this general definition here because we will study the case $p = 1$ later. The RIP is a sufficient condition on \mathbf{A} that guarantees optimal recovery of \mathbf{x} for most of the known algorithms. If the entries of \mathbf{A} are drawn i.i.d from a sub-Gaussian distribution and $m = \mathcal{O}(k \log(N/k))$, then \mathbf{A} has RIP-2 with high probability and leads to the ideal ℓ_2/ℓ_1 instance optimality for most algorithms; see [15]. Note that the bound $\mathcal{O}(k \log(N/k))$ is asymptotically tight. On the other hand, deterministic constructions of \mathbf{A} or random \mathbf{A} with binary entries with non-zero mean do not achieve this optimal m , and are faced with the so-called *square root bottleneck* where $m = \Omega(k^2)$; see [18, 16].

Sparse Sensing Matrices from Expander Graphs.

The computational benefits of sparse sensing matrices necessitated finding a way to circumvent the square root bottleneck for non-zero mean binary matrices. One such class of binary matrices is the class of adjacency matrices of expander graphs (henceforth referred to as *expander matrices*), which satisfy the weaker RIP-1. Expander graphs are objects of interest in pure mathematics and theoretical computer science, for a detailed discourse on this subject see [30]. We define an expander graph as follows:

Definition 2 (Expander graph). *Let $G = ([N], [m], \mathcal{E})$ be a left-regular bipartite graph with N left vertices, m right vertices, a set of edges \mathcal{E} and left degree d . If for any $\epsilon \in (0, 1/2)$ and any $\mathcal{S} \subset [N]$ of size $|\mathcal{S}| \leq k$ it holds $|\Gamma(\mathcal{S})| \geq (1 - \epsilon)dk$, then G is referred to as a (k, d, ϵ) -expander graph.*

Choosing $m = \mathcal{O}(k \log(N/k))$, then a random bipartite graph $G = ([N], [m], \mathcal{E})$ with left degree $d = \mathcal{O}(\frac{1}{\epsilon} \log(\frac{N}{k}))$ is an (k, d, ϵ) -expander graph with high probability [22]. Furthermore expander matrices achieve the sub-optimal ℓ_1/ℓ_1 instance optimality [8]. For completeness we state the lemma in [32] deriving the RIC for such matrices.

Lemma 3 (RIP-1 for Expander Matrices, [32]). *Let \mathbf{A} be the adjacency matrix of a (k, d, ϵ) -expander graph G , then for any k -sparse vector \mathbf{x} , we have*

$$(1 - 2\epsilon) d \|\mathbf{x}\|_1 \leq \|\mathbf{A}\mathbf{x}\|_1 \leq d \|\mathbf{x}\|_1. \quad (6)$$

The most relevant algorithm that exploits the structure of the expander matrices is the Expander-IHT (EIHT) proposed in [22]. Similar to the IHT algorithm it performs updates

$$\mathbf{x}^{(n+1)} = \mathcal{H}_k \left[\mathbf{x}^{(n)} + \mathcal{M} \left(\mathbf{A}\mathbf{x}^{(n)} - \mathbf{y} \right) \right], \quad (7)$$

where $\mathcal{M} : \mathbb{R}^m \rightarrow \mathbb{R}^N$ is the *median operator* and $[\mathcal{M}(\mathbf{z})]_i = \text{median}(\{z_j\}_{j \in \Gamma(i)})$ for each $\mathbf{z} \in \mathbb{R}^m$. For expander matrices the EIHT achieves ℓ_1/ℓ_1 instance optimality [22].

2.3 Model-based Compressed Sensing

Besides sparsity (and compressibility) signals do exhibit more complicated structures. When compressed sensing takes into account these more complicated structures (or models) in addition to sparsity, it is usually referred to as model-based compressed sensing or *structured sparse recovery* [5]. A precise definition is given in the following:

Definition 4 (Structured Sparsity Model [5]). *A structured sparsity model is a collection $\mathfrak{M} = \{\mathcal{S}_1, \dots, \mathcal{S}_M\}$ of allowed structured supports $\mathcal{S}_i \subseteq [N]$.*

Note that the classical k -sparsity studied in Section 2.2 is a special case of a structured sparsity model where all supports of size at most k are allowed. Popular structured sparsity models include tree-sparse, block-sparse, and group-sparse models [5]. In this work we study group-sparse models which we will introduce in Section 2.4.

Similar to the classical sparsity case the RIP property is defined for structured sparsity models.

Definition 5 (Model-RIP [5]). *A matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ satisfies the ℓ_p -norm model restricted isometry property (\mathfrak{M} -RIP- p) with model restricted isometry constant (\mathfrak{M} -RIC) $\delta_{\mathfrak{M}} < 1$, if for all vectors \mathbf{x} with $\text{supp}(\mathbf{x}) \in \mathfrak{M}$ it holds*

$$(1 - \delta_{\mathfrak{M}}) \|\mathbf{x}\|_p^p \leq \|\mathbf{A}\mathbf{x}\|_p^p \leq (1 + \delta_{\mathfrak{M}}) \|\mathbf{x}\|_p^p. \quad (8)$$

In [5] it was shown that for a matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ to have the Model-RIP with high probability the required number of measurements is $m = \mathcal{O}(k)$ for tree-sparse signals and $m = \mathcal{O}(kg + \log(N/(kg)))$ for block-sparse signals with block size g , when the sensing matrices are dense (typically sub-Gaussian). In

general for a given structured sparsity model \mathfrak{M} for sub-Gaussian random matrices the number of required measurements is $m = \mathcal{O}(\delta_{\mathfrak{M}}^{-2} g \log(\delta_{\mathfrak{M}}^{-1}) + \log(|\mathfrak{M}|))$ where g is the cardinality of the largest support in \mathfrak{M} .

Furthermore the authors in [5] show that classical algorithms like the IHT can be modified for structured sparsity models to achieve instance optimality. To this end the hard thresholding operator \mathcal{H}_k used in the classical IHT is replaced by a model-projection oracle which for a given signal $\mathbf{x} \in \mathbb{R}^N$ returns the closest signal over all signals having support in \mathfrak{M} . We define the model-projection oracle in the following.

Definition 6 (Model-Projection Oracle [5]). *Given $p \geq 1$, a model-projection oracle is a function $\mathcal{P}_{\mathfrak{M}} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ such that for all $\mathbf{x} \in \mathbb{R}^N$ we have $\text{supp}(\mathcal{P}_{\mathfrak{M}}(\mathbf{x})) \in \mathfrak{M}$ and it holds*

$$\|\mathbf{x} - \mathcal{P}_{\mathfrak{M}}(\mathbf{x})\|_p = \min_{\mathbf{s} \in \mathfrak{M}} \|\mathbf{x} - \mathbf{s}\|_p.$$

From the definition it directly follows that $\mathcal{P}_{\mathfrak{M}}(\mathbf{x})_i = \mathbf{x}_i$ if $i \in \text{supp}(\mathcal{P}_{\mathfrak{M}}(\mathbf{x}))$ and 0 otherwise. Note that in the case of classical k -sparsity the model-projection oracle is given by the hard thresholding operator \mathcal{H}_k . In contrast to this case, calculating the optimal model projection $\mathcal{P}_{\mathfrak{M}}(\mathbf{x})$ for a given signal $\mathbf{x} \in \mathbb{R}^N$ and a given structured sparsity model \mathfrak{M} may be computationally hard. Depending on the model \mathfrak{M} finding the optimal model projection vector may be even NP-hard; see Section 2.4. The modified version of the IHT derived in [5] is presented in Algorithm 1.

Algorithm 1 (Model-IHT)

Input $\mathbf{A}, \mathbf{y}, \mathbf{e}$ with $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ and \mathfrak{M}

Output Model-IHT($\mathbf{y}, \mathbf{A}, \mathfrak{M}$), an \mathfrak{M} -sparse approximation to \mathbf{x}

- 1: $\mathbf{x}^{(0)} \leftarrow \mathbf{0}; n \leftarrow 0$
 - 2: **while** halting criterion false **do**
 - 3: $\mathbf{x}^{(n+1)} \leftarrow \mathcal{P}_{\mathfrak{M}}[\mathbf{x}^{(n)} + \mathbf{A}^*(\mathbf{y} - \mathbf{A}\mathbf{x}^{(n)})]$
 - 4: $n \leftarrow n + 1$
 - 5: **end while**
 - 6: Return: Model-IHT($\mathbf{y}, \mathbf{A}, \mathfrak{M}$) $\leftarrow \mathbf{x}^{(n)}$
-

Note that common halting criterion is given by a maximum number of iterations or a bound on the iteration error $\|\mathbf{x}^{(n+1)} - \mathbf{x}^{(n)}\|_p$.

Model-Sparse Sensing Matrices from Expander Graphs.

In the sparse matrix setting the known results are sub-optimal. In this setting model expander matrices are adjacency matrices of model expander graphs defined similar to Definition 2; see [3]. Using *model expander matrices* for tree-sparse models the required number of measurements to obtain instance optimality is $m = k \log(N/k) / \log \log(N/k)$ which was shown in [31, 3].

A key ingredient in the analysis for the afore-mentioned sample complexity results for model expanders is the *model-RIP-1*, which is just RIP-1 for model expander matrices (hence they are also called model-RIP-1 matrices [31]). Consequently, Lemma 3 also holds for these model-RIP-1 matrices [31].

First, in [3] the *Model Expander IHT* (MEIHT) was studied for loopless overlapping groups and D -ary tree models. Similar to Algorithm 1 the MEIHT is a modification of EIHT where the hard threshold operator \mathcal{H}_k is replaced by the projection oracle $\mathcal{P}_{\mathfrak{M}}$ onto the model \mathfrak{M} . Thus the update of the MEIHT in each iteration is given by

$$\mathbf{x}^{(n+1)} = \mathcal{P}_{\mathfrak{M}} \left[\mathbf{x}^{(n)} + \mathcal{M} \left(\mathbf{A}\mathbf{x}^{(n)} - \mathbf{y} \right) \right]. \quad (9)$$

In [3] the authors show that this algorithm always returns a solution in the model, which is highly desirable for some applications. The running time is given in the proposition below.

Proposition 7 (Proposition 3.1, [3]). *The runtime of MEIHT is $\mathcal{O}(gnN)$ and $\mathcal{O}(M^2gn + nN)$ for D -ary tree models and loopless overlapping group models respectively, where g is the cardinality of the largest support in the model, n is the number of iterations, M is the number of groups and N is the dimension of the signal.*

2.4 Group Models

The models of interest in this paper are group models. A *group model* is a collection $\mathfrak{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_M\}$ of *groups* of indices, i.e. $\mathcal{G}_i \subset [N]$, together with a *budget* $G \in [M]$. We denote \mathfrak{G}_G as the structured sparsity model (i.e. group-sparse model) which contains all supports contained in the union of at most G groups in \mathfrak{G} , i.e.

$$\mathfrak{G}_G := \left\{ \mathcal{S} \subseteq [N] \mid \mathcal{S} \subseteq \bigcup_{i \in \mathcal{I}} \mathcal{G}_i, |\mathcal{I}| \leq G \right\}. \quad (10)$$

We will always tacitly assume that $\bigcup_{i=1}^M \mathcal{G}_i = [N]$. We say that a signal $\mathbf{x} \in \mathbb{R}^N$ is *G-group-sparse* if the support of \mathbf{x} is contained in \mathfrak{G}_G . If G is clear from the context, we simply say that \mathbf{x} is *group-sparse*. Let $g_i = |\mathcal{G}_i|$ and denote $g_{\max} = \max_{i \in [M]} g_i$ as the size of largest support in \mathfrak{G}_G . The *intersection graph* of a group model is the graph which has a node for each group $\mathcal{G}_i \in \mathfrak{G}$ and has an edge between \mathcal{G}_i and \mathcal{G}_j if the groups overlap, i.e. if $\mathcal{G}_i \cap \mathcal{G}_j \neq \emptyset$; see [4]. We call a group model *loopless* if the intersection graph of the group model has no cycles. We call a group model *block model* if all groups have equal size and if they are pairwise disjoint. In this case the groups are sometimes called blocks. We define the *frequency* f of a group model as the maximum number of groups, an element is contained in, i.e.

$$f := \max_{i \in [N]} |\{j \in [M] \mid i \in \mathcal{G}_j\}|.$$

In [4] besides the latter group models the more general models are considered where an additional sparsity in the classical sense is required on the signal. More precisely for a given budget $G \in [M]$ and a sparsity $K \in [N]$ they study the structured sparsity model

$$\mathfrak{G}_{G,K} := \left\{ \mathcal{S} \subseteq [N] \mid \mathcal{S} \subseteq \bigcup_{i \in \mathcal{I}} \mathcal{G}_i, |\mathcal{I}| \leq G, |\mathcal{S}| \leq K \right\}. \quad (11)$$

Note that for $K = N$ we obtain a standard group model defined as above.

Both variants of group models defined above clearly are special cases of a structured sparsity model defined in Section 2.3. Therefore all results for structured sparsity models can be used for group-sparse models. To adapt Algorithm 1 a model projection oracle $\mathcal{P}_{\mathfrak{G}_G}$ (or $\mathcal{P}_{\mathfrak{G}_{G,K}}$) has to be provided. Note that for several applications we are not only interested in the optimal support of the latter projection but we want to find at most G groups covering this support. The main work of this paper is to analyse the complexity of the latter problem for group models and to provide efficient algorithms to solve it exactly or approximately. Given a signal $\mathbf{x} \in \mathbb{R}^N$, the GROUP MODEL PROJECTION problem or sometimes called *signal approximation problem* is then to find a support $\mathcal{S} \in \mathfrak{G}_{G,K}$ together with G groups covering this support such that $\|\mathbf{x} - \mathbf{x}_{\mathcal{S}}\|_p$ is minimal, i.e. we want to solve the problem

$$\min_{\substack{\mathcal{G}_1, \dots, \mathcal{G}_G \in \mathfrak{G} \\ \mathcal{S} \subseteq \bigcup_{i=1}^G \mathcal{G}_i \\ |\mathcal{S}| \leq K}} \|\mathbf{x} - \mathbf{x}_{\mathcal{S}}\|_p.$$

If the parameter K is not mentioned we assume $K = N$.

Baldassarre et al. [4] observed the following close relation to the NP-hard MAXIMUM WEIGHT COVER-AGE problem. Given a signal $\mathbf{x} \in \mathbb{R}^N$, a group-sparse vector $\hat{\mathbf{x}}$ for which $\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$ is minimum satisfies $\hat{x}_i \in \{0, x_i\}$ for all $i \in [N]$. For a vector with the latter property,

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 = \sum_{i=1}^N x_i^2 - \sum_{i=1}^N \hat{x}_i^2$$

holds and so minimizing $\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$ is equivalent to maximizing $\sum_{i=1}^N \hat{x}_i^2$. Consequently, the GROUP MODEL PROJECTION problem with $K = N$ is equivalent to the problem of finding an index set $\mathcal{I} \subset [M]$ of at most G

groups, i.e. $|\mathcal{I}| \leq G$, maximizing $\sum_{i \in \bigcup_{j \in \mathcal{I}} \mathcal{G}_j} x_i^2$. This problem is called MAXIMUM WEIGHT COVERAGE in the literature [29]. Despite the prominence of the latter problem, we will stick to the group model notation, since it is closer to the applications we have in mind and we will leave the regime of MAXIMUM WEIGHT COVERAGE by introducing more constraints later.

We simplify the notation by defining $w_i = x_i^2$ for all $i \in [N]$. Using this notation, the GROUP MODEL PROJECTION problem is equivalent to finding an optimal solution of the following integer program:

$$\begin{aligned}
& \max \mathbf{w}^\top \mathbf{u} \\
& s.t. \quad \sum_{i \in [M]} v_i \leq G \\
& \quad u_j \leq \sum_{i: j \in \mathcal{G}_i} v_i \quad \text{for all } j \in [N] \\
& \quad \mathbf{u} \in \{0, 1\}^N, \mathbf{v} \in \{0, 1\}^M
\end{aligned} \tag{12}$$

Here, the variable u_i is one if and only if the i -th index is contained in the support of the signal approximation, and v_i is one if and only if the group \mathcal{G}_i is chosen.

Besides the NP-hardness for the general case the authors in [4] show that the GROUP MODEL PROJECTION problem can be solved in polynomial time by dynamic programming for the special case of loopless groups. Furthermore the authors show that if the intersection graph is bipartite the projection problem can be solved in polynomial time by relaxing problem (12). Similar results are obtained for the more general problem, where additional to the group-sparsity the classical K -sparsity is assumed, i.e. the additional constraint

$$\sum_{i \in [N]} u_i \leq K$$

is added to problem (12).

As stated in Section 2.3, the authors of [5] first study a special case of group models, i.e. block models, where the groups are non-overlapping and are all of equal size. The sample complexity they derived in that work for sub-Gaussian measurement matrices is $m = \mathcal{O}(Gs + \log(N/(Gs)))$, where s is the fixed block size. However, in [3] the authors studied group models in the sparse matrix setting, besides other results they proposed the MEIHT algorithm for tree and group models. The more relevant result to this work show that for *loopless* overlapping group-sparse models with maximum group size g_{\max} , using model expander measurement matrices, the number of measurements required for successful recovery is $m = Gg_{\max} \log(N/(Gg_{\max})) / \log(Gg_{\max})$; see [3]. This results holds for general groups, the “*looplessness*” condition is only necessary for the polynomial time reconstruction using the MEIHT algorithm. Therefore, this sample complexity result also holds for the general group models we consider in this manuscript.

2.5 Approximation Algorithms for Model-Based Compressed Sensing

As mentioned in the last section solving the projection problem, given in Definition 6, may be computationally hard. To overcome this problem the authors in [28, 27] present algorithms, based on the idea of IHT (and CoSaMP), which instead of solving the projection problems exactly, use two approximation procedures called head- and tail-approximation. In this section we will shortly describe the concept and the results in [28, 27]. Note that we just present results related to the IHT, although similar results for the CoSaMP were derived as well in [28, 27].

Given two structured sparsity models $\mathfrak{M}, \mathfrak{M}_H$ and a vector \mathbf{x} , let \mathcal{H} be an algorithm that computes a vector $\mathcal{H}(\mathbf{x})$ with support in \mathfrak{M}_H . Then, given some $\alpha \in \mathbb{R}$ (typically $\alpha < 1$) we say that \mathcal{H} is an $(\alpha, \mathfrak{M}, \mathfrak{M}_H, p)$ -*head approximation* if

$$\|\mathcal{H}(\mathbf{x})\|_p \geq \alpha \cdot \|\mathbf{x}_S\|_p \quad \text{for all } S \in \mathfrak{M}. \tag{13}$$

Note that the support of the vector calculated by \mathcal{H} is contained in \mathfrak{M}_H while the approximation guarantee must be fulfilled for all supports in \mathfrak{M} .

Moreover given two structured sparsity models $\mathfrak{M}, \mathfrak{M}_T$ let \mathcal{T} be an algorithm which computes a vector $\mathcal{T}(\mathbf{x})$ with support in \mathfrak{M}_T . Given some $\beta \in \mathbb{R}$ (typically $\beta > 1$) we say that \mathcal{T} is a $(\beta, \mathfrak{M}, \mathfrak{M}_T, p)$ -tail approximation if

$$\|\mathbf{x} - \mathcal{T}(\mathbf{x})\|_p \leq \beta \cdot \|\mathbf{x} - \mathbf{x}_S\|_p \text{ for all } S \in \mathfrak{M}. \quad (14)$$

Note that in general a head approximation does not need to be a tail approximation and vice versa.

The cases studied in [27] are $p = 1$ and $p = 2$. For the case $p = 2$ the authors propose an algorithm called *Approximate Model-IHT* (AM-IHT), shown in Algorithm 2.

Algorithm 2 (AM-IHT)

Input $\mathbf{A}, \mathbf{y}, \mathbf{e}$ with $\mathbf{A}\mathbf{x} = \mathbf{y} + \mathbf{e}$ and \mathfrak{M}

Output AM-IHT($\mathbf{y}, \mathbf{A}, \mathfrak{M}$), an approximation to \mathbf{x}

- 1: $\mathbf{x}^0 \leftarrow 0; n \leftarrow 0$
 - 2: **while** halting criterion is false **do**
 - 3: $\mathbf{x}^{(n+1)} \leftarrow \mathcal{T}(\mathbf{x}^{(n)} + \mathcal{H}(\mathbf{A}^*(\mathbf{y} - \mathbf{A}\mathbf{x}^{(n)})))$
 - 4: $n \leftarrow n + 1$
 - 5: **end while**
 - 6: Return: AM-IHT($\mathbf{y}, \mathbf{A}, \mathfrak{M}$) $\leftarrow \mathbf{x}^{(n)}$
-

Assume that \mathcal{T} is a $(\beta, \mathfrak{M}, \mathfrak{M}_T, 2)$ -tail approximation and \mathcal{H} is a $(\alpha, \mathfrak{M}_T \oplus \mathfrak{M}, \mathfrak{M}_H, 2)$ -head approximation where $\mathfrak{M}_T \oplus \mathfrak{M}$ is the Minkowski sum of \mathfrak{M}_T and \mathfrak{M} . Furthermore we assume the condition

$$\alpha^2 > 1 - (1 + \beta)^{-2} \quad (15)$$

holds. The authors in [27] prove that for a signal $\mathbf{x} \in \mathbb{R}^N$ with $\text{supp}(\mathbf{x}) \in \mathfrak{M}$, noisy measurements $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ where \mathbf{A} has $\mathfrak{M} \oplus \mathfrak{M}_T \oplus \mathfrak{M}_H$ -model RIP with RIC δ , Algorithm 2 calculates a signal estimate $\hat{\mathbf{x}}$ satisfying

$$\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \tau \|\mathbf{e}\|_2$$

where τ depends on δ , α and β . Note that the condition (15) holds e.g. for approximation accuracies $\alpha > 0.9$ and $\beta < 1.1$.

For the case $p = 1$ the authors replace Step 3 in Algorithm 2 by the update

$$\mathbf{x}^{(n+1)} \leftarrow \mathcal{T}(\mathbf{x}^{(n)} + \mathcal{H}(\mathcal{M}(\mathbf{y} - \mathbf{A}\mathbf{x}^{(n)})))$$

where \mathcal{M} is the median operator defined as in Section 2.2. Under the same assumptions as above, but considering $p = 1$ for the head- and tail-approximations and \mathbf{A} having the \mathfrak{M} -RIP-1, the authors in [27] show convergence of the adapted algorithm.

3 Algorithms with Exact Projection Oracles

In this section we extend existing results for group-sparse models. We pass from loopless overlapping group models (which was the most general prior to this work) to overlapping group models of bounded treewidth and for general group models without any restriction on the structure. The graph representing a loopless overlapping group model has a treewidth of 1.

We start by showing that it is possible to perform exact projections onto overlapping groups with bounded treewidth using dynamic programming, see Section 3.1. While this procedure has a polynomial run-time bound it is restricted to the class of group models with bounded treewidth. To solve the exact projection problem for all instances of group models we derive a method based on Benders' decomposition in Section 3.5. Solving an NP-hard problem this method does not yield a polynomial run-time guarantee but works well in practice as shown in [17]. In Section 3.4 we present an appropriately modified algorithm (MEIHT) with exact projection oracles for the recovery of signals from structured sparsity models. We derive corollaries

for the general group-model case from existing works about running time and convergence of this modified algorithm.

Recall the following notation: \mathcal{G}_i denotes a group of size g_i , $i \in [M]$, and a group model is a collection $\mathfrak{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_M\}$. The group-sparse model of order G is denoted by \mathfrak{G}_G , which contains all supports \mathcal{S} which are contained in the union of at most G groups of \mathfrak{G} , i.e. $\mathcal{S} \subseteq \bigcup_{j \in \mathcal{I}} \mathcal{G}_j$, $\mathcal{I} \subseteq [M]$ and $|\mathcal{I}| \leq G$; see (10). We will interchangeably say that \mathbf{x} or \mathcal{S} is \mathfrak{G}_G -sparse. Clearly group-models are a special case of structured sparsity models. Assume s_{\max} is the size of the maximal support which is possible by selecting G groups out of \mathfrak{G} . For g_{\max} denoting the maximal size of a single group in \mathfrak{G} , i.e.,

$$g_{\max} = \max_{i \in [M]} |\mathcal{G}_i|,$$

we have $s_{\max} \in \mathcal{O}(Gg_{\max})$. Furthermore the number of possible supports is in $\mathcal{O}(M^G)$. Therefore applying the result from Section 2.3 we obtain

$$m = \mathcal{O}(\delta^{-2}Gg_{\max} \log(\delta^{-1}) + G \log(M))$$

as the number of required measurements for a sub-Gaussian matrix to obtain the group-model-RIP with RIC δ with high probability which induces the convergence of Algorithm 1.

3.1 Group models of Low Treewidth

One approach to overcome the hardness of the GROUP MODEL PROJECTION problem is to restrict the structure of the group models considered. To this end we follow Baldassarre et al. [4] and consider two graphs associated to a group model \mathfrak{G} .

The *intersection graph* of \mathfrak{G} , $I(\mathfrak{G})$, is given by the vertex set $V(I(\mathfrak{G})) = \mathfrak{G}$, and the edge set

$$E(I(\mathfrak{G})) = \{RS : R, S \in \mathfrak{G}, R \neq S \text{ and } R \cap S \neq \emptyset\}.$$

The *incidence graph* of \mathfrak{G} , $B(\mathfrak{G})$, is given by the vertex set $V(B(\mathfrak{G})) = E \cup \mathfrak{G}$, and the edge set

$$E(B(\mathfrak{G})) = \{eS : e \in E, S \in \mathfrak{G} \text{ and } e \in S\}.$$

Note that the incidence graph is bipartite: on one side we have the elements and on the other side the groups of \mathfrak{G} . There is an edge between an element e and a group S if and only if $e \in S$. See Fig. 1 for a simple illustration of the two constructions. Baldassarre et al. [4] prove that there is a polynomial time algorithm

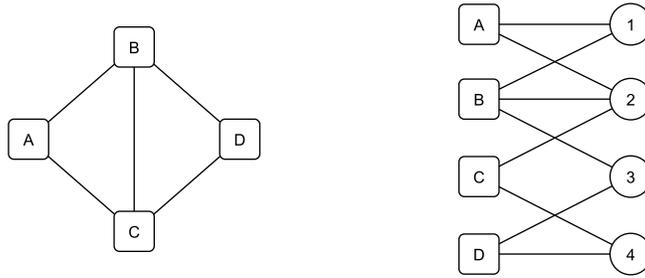


Figure 1: Intersection graph and incidence graph of the group model $\mathfrak{G} = \{A = \{1, 2\}, B = \{1, 2, 3\}, C = \{2, 4\}, D = \{3, 4\}\}$.

to solve the GROUP MODEL PROJECTION Problem in the case that the intersection graph of the instance is an acyclic graph. Their algorithm uses dynamic programming on the acyclic structure of the intersection graph. We generalize their approach to show that the same problem can be solved in polynomial time if

the treewidth of the intersection graph of the instance is bounded. As we show below, this implies that the GROUP MODEL PROJECTION Problem can be solved in polynomial time if the treewidth of the intersection graph of the instance is bounded. We proceed by formally introducing the relevant concepts, treewidth and tree decompositions.

3.2 Tree decomposition and dynamic programming

Let $G = (V, E)$ be a graph. A *tree decomposition* of G is a tree T as follows. Each node $x \in V(T)$ of T has a *bag* $B_x \subseteq V$ of vertices of G such that the following properties all hold:

1. $\bigcup_{x \in V(T)} B_x = V$.
2. If B_x and B_y both contain a vertex $v \in V$, then the bags of all nodes of T in the path between x and y contain v as well. Equivalently, the tree nodes containing the vertex v form a connected subtree of T .
3. For every edge vw in G there is some bag that contains both v and w . That is, vertices of G can be adjacent only if the corresponding subtrees have a node in common.

The *width* of a tree decomposition is the size of its largest bag minus one, i.e. $\max_{x \in V(T)} |B_x| - 1$. The *treewidth* of G , $\text{tw}(G)$, is the minimum width among all possible tree decompositions of G .

Intuitively, the treewidth measures how ‘treelike’ a graph is: the smaller the treewidth is, the more treelike is the graph. The graphs of treewidth one are the acyclic graphs. Fig. 2 shows a graph of treewidth 2, together with a tree decomposition.

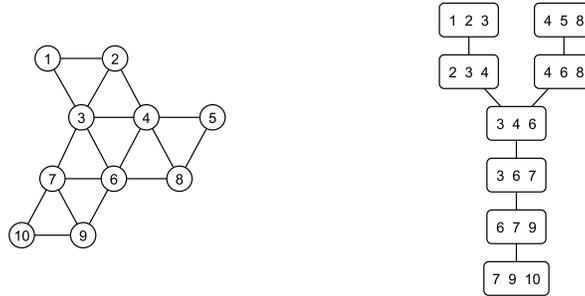


Figure 2: A graph and a tree decomposition of width 2.

Before stating any algorithms, we discuss the relation of the treewidth of the intersection and incidence graphs of a given group model.

When bounding the treewidth of the graphs associated to a group model, it makes sense to consider the incidence graph rather than the intersection graph. This is due to the following simple observation.

Proposition 8. *For any group model \mathfrak{G} it holds that $\text{tw}(B(\mathfrak{G})) \leq \text{tw}(I(\mathfrak{G})) + 1$. However, for every t there is a group model \mathfrak{G} such that $\text{tw}(I(\mathfrak{G})) - \text{tw}(B(\mathfrak{G})) = t$.*

This statement is not necessarily new, but we quickly prove it in our language in order to be self-contained.

Proof of Proposition 8. To see the first assertion, let \mathfrak{G} be a group model, say with ground set $E = [N] = \bigcup \mathfrak{G}$. Let T be a tree decomposition of $I(\mathfrak{G})$ of width $\text{tw}(I(\mathfrak{G}))$. In the following, we attach leaves to T , one for each element in E , and obtain a tree decomposition of $B(\mathfrak{G})$. Each leaf will contain at most $\text{tw}(I(\mathfrak{G}))$ many elements of \mathfrak{G} and at most one element of E . Hence we get $\text{tw}(B(\mathfrak{G})) \leq \text{tw}(I(\mathfrak{G})) + 1$.

For this, pick any $i \in [N]$, and let \mathfrak{G}_i be the set of groups in \mathfrak{G} containing i . Since all groups in \mathfrak{G}_i contain i , the set \mathfrak{G}_i is a clique in $I(\mathfrak{G})$.¹ Moreover, since T is a tree decomposition, the subtrees of the groups in

¹Recall that a *clique* in a graph is a set of mutually adjacent vertices.

\mathfrak{G}_i mutually share a node. As subtrees of a tree have the Helly property, there is at least one node x of T such that $\mathfrak{G}_i \subseteq B_x$. We now add a new node x_i with bag $B_{x_i} = \mathfrak{G}_i \cup \{i\}$ and make it adjacent to x in T .

Doing this for all vertices of $B(\mathfrak{G})$ simultaneously, it is easy to see that we arrive at a tree decomposition T' of $B(\mathfrak{G})$ of width at most $\text{tw}(I(\mathfrak{G})) + 1$. This proves our claim.

To see the second assertion consider, for any t , the group model \mathfrak{G} where

$$\mathfrak{G} = \{\mathcal{G}_i : i \in [t+2]\} \text{ with } \mathcal{G}_i = \{i, t+3\}.$$

Note that $B(\mathfrak{G})$ is a tree, hence $\text{tw}(B(\mathfrak{G})) = 1$. In $I(\mathfrak{G})$, however, the set \mathfrak{G} is a clique of size $t+2$. Thus, $\text{tw}(I(\mathfrak{G})) = t+1$ which implies $\text{tw}(I(\mathfrak{G})) - \text{tw}(B(\mathfrak{G})) = t$. \square \square

We now go for a polynomial time algorithm for the GROUP MODEL PROJECTION problem if the treewidth is fixed. Such an algorithm is presented best using a so-called *nice tree decomposition*. This kind of decomposition limits the structure of the difference of two adjacent nodes in the decomposition. Formally, consider a tree decomposition T of a graph G , rooted in a leaf r of T . We say that T is a *nice tree decomposition* if every node x is of one of the following types.

- **Leaf:** x has no children and $B_x = \emptyset$.
- **Introduce:** x has one child, say y , and there is a vertex $v \notin B_y$ of G with $B_x = B_y \cup \{v\}$.
- **Forget:** x has one child, say y , and there is a vertex $v \notin B_x$ of G with $B_y = B_x \cup \{v\}$.
- **Join:** x has two children y and z such that $B_x = B_y = B_z$.

A folklore statement says that such a nice decomposition is easily computed given any tree decomposition of G without increasing the width.

Our dynamic program is build on the following concept. Fix a node x of the decomposition tree, a number i with $0 \leq i \leq k$ and a map $c : B_x \rightarrow \{0, 1, 1?\}$. We say that c is a *colouring* of B_x . We considered the solutions to the problem GROUP MODEL PROJECTION(x, i, c) which is defined as follows.

A set $\mathcal{S} \subseteq \mathfrak{G}$ is a feasible solution of GROUP MODEL PROJECTION(x, i, c) if

- (a) \mathcal{S} contains only groups that appear in some bag of a node in the subtree rooted at node x ,
- (b) $|\mathcal{S}| = i$,
- (c) \mathcal{S} contains exactly those group-vertices of B_x that are in $c^{-1}(1)$, that is,

$$\mathcal{S} \cap c^{-1}(1) = \mathfrak{G} \cap c^{-1}(1), \text{ and}$$

- (d) \mathcal{S} covers exactly those elements of B_x that are in $c^{-1}(1)$, that is,

$$\left(\bigcup \mathcal{S}\right) \cap B_x = E \cap c^{-1}(1).$$

The objective value of the set \mathcal{S} is given by $w(\bigcup \mathcal{S}) + c^{-1}(1?)$. Intuitively, the elements labelled $1?$ are assumed to be covered by groups not yet visited in the tree decomposition.

If GROUP MODEL PROJECTION(x, i, c) does not admit a feasible solution, we say that GROUP MODEL PROJECTION(x, i, c) is infeasible. The maximum objective value attained by a feasible solution to GROUP MODEL PROJECTION(x, i, c), if feasible, we denote by OPT(x, i, c). If GROUP MODEL PROJECTION(x, i, c) is infeasible, we put OPT(x, i, c) = $-\infty$.

Assertion (d) implies that GROUP MODEL PROJECTION(x, i, c) is infeasible if the groups in $c^{-1}(1)$ cover elements in $c^{-1}(0)$ or $c^{-1}(1?)$. That is, GROUP MODEL PROJECTION(x, i, c) is infeasible if $\bigcup (\mathcal{G} \cap c^{-1}(1)) \not\subseteq E \cap c^{-1}(1)$. To deal with this exceptional case we call c *consistent* if $\bigcup (\mathcal{G} \cap c^{-1}(1)) \subseteq E \cap c^{-1}(1)$, and *inconsistent* otherwise. Note that consistency of c is necessary to ensure feasibility of GROUP MODEL PROJECTION(x, i, c), but not sufficient.

3.2.1 Dynamic Programming

Our algorithm processes the nodes of a nice tree decomposition in a bottom-up fashion. Fix a node x , a number i with $0 \leq i \leq k$ and a map $c : B_x \rightarrow \{0, 1, 1?\}$. We use dynamic programming to compute the value $\text{OPT}(x, i, c)$, assuming we know all possible values $\text{OPT}(y, j, c')$ for all children y of x , all j with $0 \leq j \leq k$, and all $c' : B_y \rightarrow \{0, 1, 1?\}$.

If c is not consistent, we may put $\text{OPT}(x, i, c) = -\infty$ right away. We thus assume that c is consistent and distinguish the type of node x is.

- **Leaf:** put $\text{OPT}(x, 0, c) = 0$ and $\text{OPT}(x, i, c) = -\infty$ for all $i \in [k]$.
- **Introduce:** let y be the unique child of x and let $v \notin B_y$ such that $B_x = B_y \cup \{v\}$.
If $v \in E$, we put

$$\text{OPT}(x, i, c) = \begin{cases} \text{OPT}(y, i, c|_{B_y}), & \text{if } c(v) = 0 \\ \text{OPT}(y, i, c|_{B_y}) + w(v), & \text{if } c(v) = 1 \text{ and } c^{-1}(1) \cap N(v) \neq \emptyset \\ \text{OPT}(y, i, c|_{B_y}) + w(v), & \text{if } c(v) = 1? \\ -\infty, & \text{otherwise} \end{cases}$$

If $v \in \mathfrak{G}$, we put

$$\text{OPT}(x, i, c) = \begin{cases} \text{OPT}(y, i, c|_{B_y}), & \text{if } c(v) = 0 \\ \max\{\text{OPT}(y, i, c') : (y, c') \text{ is compatible to } (x, c)\}, & \text{if } c(v) = 1 \text{ and } c^{-1}(1) \cap N(v) \neq \emptyset \\ -\infty, & \text{otherwise} \end{cases} \quad (16)$$

where (y, c') is *compatible* to (x, c) if

- $c' : B_y \rightarrow \{0, 1, 1?\}$ is a consistent colouring of B_y ,
- $c^{-1}(0) = c'^{-1}(0)$, and
- $c^{-1}(1) = c'^{-1}(1) \cup (c'^{-1}(1?) \cap N(v))$.

- **Forget:** let y be the unique child of x and let $v \notin B_x$ such that $B_y = B_x \cup \{v\}$. We put

$$\text{OPT}(x, i, c) = \max\{\text{OPT}(y, i, c') : c' : B_y \rightarrow \{0, 1, 1?\}, c = c'|_{B_x}, \text{ and } c'(v) \neq 1?\} \quad (17)$$

- **Join:** we put

$$\text{OPT}(x, i, c) = \max\{\text{OPT}(y, i_1, c') + \text{OPT}(z, i_2, c'') - w(((B_x \cap E) \cup \bigcup (B_x \cap \mathfrak{G})) \setminus c^{-1}(0)) : i_1 + i_2 - |c'^{-1}(1) \cap c''^{-1}(1) \cap \mathfrak{G}| = i\}, \quad (18)$$

where y and z are the two children of x . The maximum is taken over all consistent colourings $c', c'' : B_x \rightarrow \{0, 1, 1?\}$ with $c^{-1}(0) = c'^{-1}(0) = c''^{-1}(0)$ and $c^{-1}(1) = c'^{-1}(1) \cup c''^{-1}(1)$.

- **Root:** first we compute $\text{OPT}(x, k, c)$ for all relevant choices of c , depending on the type of node x is. The algorithm then terminates with the output

$$\text{OPT} = \max\{\text{OPT}(x, k, c) : c : B_x \rightarrow \{0, 1, 1?\}\}.$$

Lemma 9. *The output OPT is the objective value of an optimal solution of the GROUP MODEL PROJECTION problem.*

Proof. The proof follows the individual steps of the dynamic programming algorithm. Consider the problem $\text{GROUP MODEL PROJECTION}(x, i, c)$. If c is not consistent, we correctly put $\text{OPT}(x, i, c) = -\infty$. We thus proceed to the case when c is consistent. Fix an optimal solution \mathcal{S} of $\text{GROUP MODEL PROJECTION}(x, i, c)$ if existent.

The **leaf**-node case is clear, so we proceed to the case of x being an **introduce**-node. Let y be the unique child of x and let $v \notin B_y$ such that $B_x = B_y \cup \{v\}$. First assume that $v \in E$.

Assume that $c(v) = 0$. Since c is consistent, $c^{-1}(1) \cap N(v) = \emptyset$ holds, and so \mathcal{S} is an optimal solution to $\text{GROUP MODEL PROJECTION}(y, i, c|_{B_y})$. We may thus put $\text{OPT}(x, i, c) = \text{OPT}(y, i, c|_{B_y})$.

If $c(v) = 1$, \mathcal{S} covers v , so we need to make sure some vertex $u \in N(v) \cap B_x$ is in the solution in order for $\text{GROUP MODEL PROJECTION}(x, i, c)$ to be feasible. Hence we have $\text{OPT}(x, i, c) = \text{OPT}(y, i, c|_{B_y}) + w(v)$ if $c^{-1}(1) \cap N(v) \neq \emptyset$, and $\text{OPT}(x, i, c) = -\infty$ otherwise.

Next we assume that $v \in \mathfrak{G}$. If $c(v) = 0$, we may simply put $\text{OPT}(x, i, c) = \text{OPT}(y, i, c|_{B_y})$. So, assume that $c(v) = 1$. If $\text{GROUP MODEL PROJECTION}(x, i, c)$ is feasible and thus \mathcal{S} exists, put $\mathcal{S}' = \mathcal{S} \setminus \{v\}$. Now \mathcal{S}' is a solution to $\text{OPT}(y, i, c')$ for some $c' : B_y \rightarrow \{0, 1, 1?\}$ with $c^{-1}(0) = c'^{-1}(0)$ and $c^{-1}(1) = c'^{-1}(1) \cup (c'^{-1}(1?) \cap N(v))$. This is to say: (y, c') is compatible to (x, c) . Consequently, $\text{OPT}(x, i, c)$ is upper bounded by the right hand side of (16).

To see that $\text{OPT}(x, i, c)$ is at least the right hand side of (16), let (y, c'') be compatible to (x, c) and let \mathcal{S}'' be a solution to $\text{GROUP MODEL PROJECTION}(y, i, c'')$ of objective value λ . Then $\mathcal{S}'' \cup \{v\}$ is a solution to $\text{GROUP MODEL PROJECTION}(x, i, c)$ of objective value λ . Consequently, $\text{OPT}(x, i, c) = \max\{\text{OPT}(y, i, c') : (y, c') \text{ is compatible to } (x, c)\}$.

If x is a **forget**-node, let y be the unique child of x and let $v \notin B_x$ such that $B_y = B_x \cup \{v\}$. If $v \in \mathcal{S}$, we have

$$\text{OPT}(x, i, c) \leq \text{OPT}(y, i, c') \text{ where } c' : B_y \rightarrow \{0, 1, 1?\} c = c'|_{B_x}, \text{ and } c'(v) = 1.$$

Otherwise, if $v \notin \mathcal{S}$, we have

$$\text{OPT}(x, i, c) \leq \text{OPT}(y, i, c') \text{ where } c' : B_y \rightarrow \{0, 1, 1?\} c = c'|_{B_x}, \text{ and } c'(v) = 0.$$

Moreover, any solution of $\text{GROUP MODEL PROJECTION}(y, i, c')$, where $c' : B_y \rightarrow \{0, 1, 1?\} c = c'|_{B_x}$, and $c'(v) = 1$ is a solution of $\text{GROUP MODEL PROJECTION}(x, i, c)$. This proves (17).

If x is a **join**-node, let y and z be the two children of x and recall that $B_x = B_y = B_z$.

Let \mathcal{S}' be the collection of groups in \mathcal{S} contained in the subtree rooted at y , and let \mathcal{S}'' be the collection of groups in \mathcal{S} contained in the subtree rooted at z . Since T is a tree decomposition, $\mathcal{S}' \cap \mathcal{S}'' = \mathcal{S} \cap B_x$.

Note that \mathcal{S}' is a solution of $\text{GROUP MODEL PROJECTION}(y, i_1, c')$ and \mathcal{S}'' is a solution of $\text{GROUP MODEL PROJECTION}(z, i_2, c'')$ for some $c', c'' : B_x \rightarrow \{0, 1, 1?\}$ and i_1, i_2 with $i_1 + i_2 - |c'^{-1}(1) \cap c''^{-1}(1) \cap \mathfrak{G}| = i$. It is easy to see that $c^{-1}(0) = c'^{-1}(0) = c''^{-1}(0)$ and $c^{-1}(1) = c'^{-1}(1) \cup c''^{-1}(1)$.

The objective value of \mathcal{S} equals $w(\bigcup \mathcal{S}) + c^{-1}(1?) = w(\bigcup \mathcal{S}') + c'^{-1}(1?) + w(\bigcup \mathcal{S}'') + c''^{-1}(1?) - w(((B_x \cap E) \cup \bigcup (B_x \cap \mathfrak{G})) \setminus c^{-1}(0))$. This shows that $\text{OPT}(x, i, c)$ is at most the right hand side of (18).

Now let $\tilde{\mathcal{S}}$ be an optimal solution of $\text{GROUP MODEL PROJECTION}(y, j_1, \tilde{c})$ and let $\hat{\mathcal{S}}$ be an optimal solution of $\text{GROUP MODEL PROJECTION}(z, j_2, \hat{c})$ where

- $\tilde{c}, \hat{c} : B_x \rightarrow \{0, 1, 1?\}$ are both consistent,
- $c^{-1}(0) = \tilde{c}^{-1}(0) = \hat{c}^{-1}(0)$ and $c^{-1}(1) = \tilde{c}^{-1}(1) \cup \hat{c}^{-1}(1)$, and
- $j_1 + j_2 - |\tilde{c}^{-1}(1) \cap \hat{c}^{-1}(1) \cap \mathfrak{G}| = i$.

Note that $\tilde{\mathcal{S}}$ and $\hat{\mathcal{S}}$ exist since, as we have shown earlier, the colourings c' and c'' satisfy the above assertions. Then $\tilde{\mathcal{S}} \cup \hat{\mathcal{S}}$ is a solution of $\text{GROUP MODEL PROJECTION}(x, i, c)$ with objective value

$$\max\{\text{OPT}(y, j_1, \tilde{c}) + \text{OPT}(z, j_2, \hat{c}) - w(((B_x \cap E) \cup \bigcup (B_x \cap \mathfrak{G})) \setminus c^{-1}(0))\}.$$

Consequently, $\text{OPT}(x, i, c)$ is at least the right hand side of (18) and thus (18) holds. \square \square

By storing the best current solution alongside the $\text{OPT}(x, i, c)$ -values we can compute an optimal solution together with OPT. We do not give details here since this is standard.

3.2.2 Runtime of the Algorithm

The computational complexity of the individual steps are as follows.

- (a) Given a graph on n vertices of treewidth w , one can compute a tree decomposition of width w in time $\mathcal{O}(2^{\mathcal{O}(w^3)}n)$ using Bodlaender's algorithm [11]. The number of nodes of the decomposition is in $\mathcal{O}(n)$.
- (b) Given a tree decomposition of width w with t nodes, one can compute a nice tree decomposition of width w on $\mathcal{O}(wt)$ nodes in $\mathcal{O}(w^2t)$ time in a straightforward way.

The running time of the dynamic programming is bounded as follows.

Theorem 10. *The dynamic programming algorithm can be implemented to run in $\mathcal{O}(w5^w k^2 Nt)$ time given a nice tree decomposition of $B(\mathfrak{G})$ of width w on t nodes.*

Note that we can assume that $t = \mathcal{O}(wM)$, leading to a total running time of $\mathcal{O}(w5^w k^2 NM)$.

Proof of Theorem 10. Since the **join**-nodes are clearly the bottleneck of the algorithm, we discuss how to organize the computation in a way that the desired running time bound of $\mathcal{O}(w5^w k^2 N)$ holds in a node of this type.

So, let x be a **join**-node and assume that y and z are the two children of x . We want to compute $\text{OPT}(x, i, c)$, for all colourings $c : B_x \rightarrow \{0, 1, 1_\gamma\}$ and all i with $0 \leq i \leq k$. Recall that we need to compute this value according to (18), that is,

$$\begin{aligned} \text{OPT}(x, i, c) = & \max\{\text{OPT}(y, i_1, c') + \text{OPT}(z, i_2, c'') \\ & - w((B_x \cap E) \cup \bigcup (B_x \cap \mathfrak{G})) \setminus c^{-1}(0) : i_1 + i_2 - |c'^{-1}(1) \cap c''^{-1}(1) \cap \mathfrak{G}| = i\}, \end{aligned}$$

where the maximum is taken over all consistent colourings $c', c'' : B_x \rightarrow \{0, 1, 1_\gamma\}$ with $c^{-1}(0) = c'^{-1}(0) = c''^{-1}(0)$ and $c^{-1}(1) = c'^{-1}(1) \cup c''^{-1}(1)$.

We enumerate all 5^{w+1} colourings $C : B_x \rightarrow \{(0, 0), (1, 1), (1, 1_\gamma), (1_\gamma, 1), (1_\gamma, 1_\gamma)\}$ and derive $c, c',$ and c'' . We put

$$\begin{aligned} c(v) &= \begin{cases} 0, & \text{if } C(v) = (0, 0) \\ 1, & \text{if } C(v) \in \{(1, 1), (1, 1_\gamma), (1_\gamma, 1)\} \\ 1_\gamma, & \text{if } C(v) = (1_\gamma, 1_\gamma) \end{cases} \\ c'(v) &= \begin{cases} 0, & \text{if } C(v) = (0, 0) \\ 1, & \text{if } C(v) \in \{(1, 1), (1, 1_\gamma)\} \\ 1_\gamma, & \text{if } C(v) \in \{(1_\gamma, 1), (1_\gamma, 1_\gamma)\} \end{cases} \\ c''(v) &= \begin{cases} 0, & \text{if } C(v) = (0, 0) \\ 1, & \text{if } C(v) \in \{(1, 1), (1, 1_\gamma)\} \\ 1_\gamma, & \text{if } C(v) \in \{(1, 1_\gamma), (1_\gamma, 1_\gamma)\} \end{cases} \end{aligned}$$

If either of $c, c',$ or c'' are inconsistent, we discard this choice of C . In this way we capture every consistent colouring $c : B_x \rightarrow \{0, 1, 1_\gamma\}$, and all consistent choices of $c',$ and c'' satisfying $c^{-1}(0) = c'^{-1}(0) = c''^{-1}(0)$ and $c^{-1}(1) = c'^{-1}(1) \cup c''^{-1}(1)$.

It remains to discuss the computation of the value $\mathbf{w}((c^{-1}(1) \cap E) \cup \bigcup (c^{-1}(1) \cap \mathfrak{G}))$. This value can be computed in $\mathcal{O}(wN)$ time, since we are computing differences and unions of at most w groups of size N each. We arrive at a total running time in $\mathcal{O}(w5^w k^2 N)$. \square \square

Remark 11. *The dynamic programming algorithm can be extended to include a sparsity restriction on the support of the signal approximation itself. So, we can compute an optimal k -sparse G -group-sparse signal approximation if the bipartite incidence graph of the studied group models is bounded. The running time of the algorithm increases by a factor of $\mathcal{O}(k)$.*

3.3 Hardness on grid-like group structures

An $r \times r$ -grid is a graph with vertex set $[r] \times [r]$, and two vertices $(a, b), (c, d) \in [r] \times [r]$ are adjacent if and only if $|a - c| = 1$ and $|b - d| = 0$, or if $|a - c| = 0$ and $|b - d| = 1$. We also say that r is the size of the grid. Fig. 3 shows a 6×6 -grid.

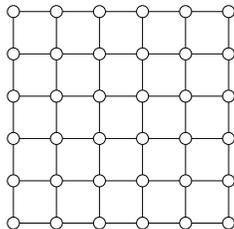


Figure 3: A 6×6 -grid

Recall the GROUP MODEL PROJECTION problem can be solved efficiently when the treewidth of the incidence graph of the group structure is bounded, as shown in Section 3.1.

Definition 12 (Graph minor). *Let G and H be two graphs. The graph H is called a minor of G if H can be obtained from G by deleting edges and/or vertices and by contracting edges.*

A classical theorem by Robertson and Seymour [43] says that in a graph class \mathcal{C} closed under taking minors either the treewidth is bounded or \mathcal{C} contains all grid graphs.

Consequently, if \mathcal{C} is a class of graphs that does not have a bounded treewidth, it contains all grids. Our next theorem shows that GROUP MODEL PROJECTION is NP-hard on group models \mathfrak{G} for which $B(\mathfrak{G})$ is a grid, thus complementing Theorem 10.

Theorem 13. *The GROUP MODEL PROJECTION problem is NP-hard even if $B(\mathfrak{G})$ is a grid graph and the weight of any element is either 0 or 1.*

By turning the grid by 45 degrees, this corresponds to choosing the k best 2×2 -pixel groups in a black/white image to cover as many white pixels as possible. Note that this group model is of frequency at most 4, and so we can do an approximate model projection and signal recovery using the result of Section 4.

Our proof is a reduction from the VERTEX COVER problem. Recall that for a graph G , a *vertex cover* is a subset of the vertices of G such that any edge of G has at least one endpoint in this subset. The size of the smallest vertex cover of G , the *vertex cover number*, is denoted $\tau(G)$.

Given a graph G and a number k as input, the task in the VERTEX COVER problem is to decide whether G admits a vertex cover of size k . That is, whether $\tau(G) \leq k$. This problem is NP-complete even if restricted to cubic planar graphs [23].²

We use the following simple lemma in our proof.

Lemma 14. *Let G be a graph and let G' be the graph obtained by subdividing some edge of G twice. Then $\tau(G) = \tau(G') - 1$.*

We can now prove our theorem.

Proof of Theorem 13. The reduction is from VERTEX COVER on planar cubic graphs. Consider $G = (V, E)$ to be a planar cubic graph and let k be some number. Our aim is to compute an instance $(\mathfrak{G}, \mathbf{w}, k')$ of the GROUP MODEL PROJECTION problem where $B(\mathfrak{G})$ is a grid such that G has a vertex cover of size k if and only if \mathfrak{G} admits a selection of k' groups that together cover elements of a total weight at least some threshold t .

²Recall that a graph is *cubic* if every vertex is of degree 3, and *planar* if it can be drawn into the plane such that no two edges cross.

First we embed the graph G in some grid H of polynomial size, meaning the vertices of G get mapped to the vertices of the grid and edges get mapped to mutually internally disjoint paths in the grid connecting its endvertices. This can be done in polynomial time using an algorithm for orthogonal planar embedding [2]. We denote the mapping by π , hence $\pi(u)$ is some vertex of H and $\pi(vw)$ is a path from $\pi(v)$ to $\pi(w)$ in H for all $u \in V$ and $vw \in E$.

Next we subdivide each edge of the grid 9 times, so that a vertical/horizontal edge of H becomes a vertical/horizontal path of length 10 in some new, larger grid H' . We choose H' as small as possible and observe that the corners of H are mapped to the corners of H' . In particular, $|V(H')| \leq 100|V(H)|$. Let us denote the obtained embedded subdivision of G in H' by G' , and let π' denote the embedding. Moreover, let ϕ be the corresponding embedding of H into the subdivided grid H' . Note that $\text{im } \pi'|_V \subseteq \text{im } \phi|_V$.

Let (A, B) be a bipartition of H' . We may assume that every vertex $\pi'(u)$ is in A , $u \in V$. We consider H' to be the incidence graph $B(\mathfrak{G})$ of a group model \mathfrak{G} where the vertices in B correspond to the elements and the vertices in A correspond to the groups of \mathfrak{G} . We refer to the vertices in A as *group-vertices* and to vertices in B as *element-vertices*. Slightly abusing notation, we identify each group with its group-vertex and each element with its element-vertex and write $\mathfrak{G} = A$.

We observe that

- (a) G' is an induced subgraph of G' ,
- (b) every vertex $\pi'(u)$, $u \in V$, has degree 3 in G' and is a group-vertex,
- (c) every other vertex has degree 2 in G' , and
- (d) for every group-vertex $x \in V(H') \setminus V(G')$ there is some group-vertex $u \in V(G')$ with $N_{H'}(x) \subseteq N_{H'}(u)$.

Next we will tweak the embedding of G a bit, to get rid of paths $\pi(uv)$ with the wrong parity. We do so in a way that preserves the properties (a)-(d). Let $\mathcal{P}_0 \subseteq \{\pi'(uv) : uv \in E(H)\}$ be the set of all paths with a length 2 (mod 4), and let $\mathcal{P}_2 = \{\pi'(uv) : uv \in E(H)\} \setminus \mathcal{P}_0$. We want to substitute each path in \mathcal{P}_0 by a path of length 2 (mod 4). For this, let w be the neighbour of u in the path $\pi(uv)$. Note that the path $\pi'(uw)$ in H' starts with a vertical or horizontal path P from $\pi'(u)$ to $\pi'(w)$ of length 10. We bypass the middle vertex of this path (an element-vertex) by going over two new element-vertices and one group-vertex instead. See Fig. 4 for an illustration.

To keep the notation easy we denote the newly obtained path by $\pi''(uv)$. Note that, after adding the bypass, the new path $\pi''(uv)$ is two edges longer and thus has length 2 (mod 4). We complete π'' to an embedding of G by putting $\pi''(u) = \pi'(u)$ and $\pi''(vw) = \pi'(vw)$ for all $u \in V$ and $vw \in E$ with $\pi'(vw) \in \mathcal{P}_2$. Moreover, let us denote the changed embedding of G by G'' .

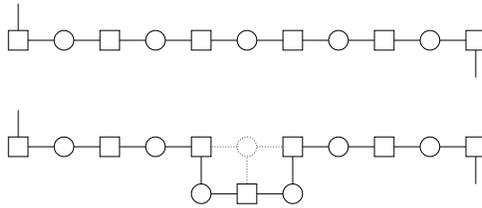


Figure 4: Introducing a bypass. Squares are group-vertices, circles are element-vertices

We observe that the new embedding G'' still satisfies the assertions (a)-(d) and, in addition, it holds that

- (e) every path connecting two vertices of degree 3 over vertices of degree 2 only has length 2 (mod 4).

Next we define the weights of the element-vertices by putting

$$\mathbf{w}(x) = \begin{cases} 1, & x \in V(G'') \\ 0, & x \in V(H') \setminus V(G'') \end{cases} \quad \text{for each element-vertex } x \text{ of } H'.$$

Assertion (d) implies that, for any subset $\mathcal{S} \subseteq \mathfrak{G}$ of size k there is a $\mathcal{S}' \subseteq \mathfrak{G}$ of size at most k such that

- $\mathcal{S}' \subseteq A \cap V(G'')$, and
- $\mathbf{w}(\bigcup \mathcal{S}') \geq \mathbf{w}(\bigcup \mathcal{S})$.

Since $\mathbf{w}(u) = 0$ for all elements in $B \setminus V(G'')$, we may thus restrict our attention to the restricted group model $\mathfrak{G}' = A \cap V(G'')$ on the element set $B \cap V(G'')$.

Slightly abusing notation, any subset $\mathcal{S} \subseteq \mathfrak{G}'$ is a vertex subset of $I(\mathfrak{G}')$ and $\mathbf{w}(\bigcup \mathcal{S})$ equals the number of edges of $I(\mathfrak{G}')$ adjacent to the vertex set \mathcal{S} in $I(\mathfrak{G}')$. Moreover, the graph $I(\mathfrak{G}')$ is obtained from the graph G by subdividing each edge an even number of times.

From Lemma 14 we know that there is some number t such that $\tau(G) = \tau(I(\mathfrak{G})) - t$. Hence, G has a vertex cover of size k if and only if M' has a cover of size $k' = k + t$ of total weight $|E(I(\mathfrak{G}'))|$. This, in turn, is the case if and only if M admits a cover of size k' of total weight $|E(I(\mathfrak{G}'))|$. Since the construction of \mathfrak{G} can be done in polynomial time, the proof is complete. \square \square

3.4 MEIHT for General Group Structures

In this section we apply the results for structured sparsity models and for expander matrices to the group model case. The Model-Expander IHT (MEIHT) algorithm is one of the exact projection algorithms with provable guarantees for tree-sparse and loopless overlapping group-sparse models using model-expander sensing matrices [3]. In this work we show how to use the MEIHT for more general group structures. The only modification of the MEIHT algorithm is the projection on these new group structures. We show MEIHT's guaranteed convergence and polynomial runtime.

Algorithm 3 (MEIHT)

Input $\mathbf{A}, \mathbf{y}, \mathbf{e}$ with $\mathbf{A}\mathbf{x} = \mathbf{y} + \mathbf{e}$, G , and \mathfrak{G}

Output MEIHT($\mathbf{y}, \mathbf{A}, \mathfrak{G}, G$), a \mathfrak{G}_G -sparse approximation to \mathbf{x}

- 1: $\mathbf{x}^{(0)} \leftarrow \mathbf{0}; n \leftarrow 0$
 - 2: **while** halting criterion is false **do**
 - 3: $\mathbf{x}^{(n+1)} \leftarrow \mathcal{P}_{\mathfrak{G}_G} [\mathbf{x}^{(n)} + \mathcal{M}(\mathbf{y} - \mathbf{A}\mathbf{x}^{(n)})]$
 - 4: $n \leftarrow n + 1$
 - 5: **end while**
 - 6: **Return:** MEIHT($\mathbf{y}, \mathbf{A}, \mathfrak{G}, G$) $\leftarrow \mathbf{x}^{(n)}$
-

Note that as in [4], we are able to do model projections with an additional sparsity constraint, i.e. projection onto $\mathcal{P}_{\mathfrak{G}_{G,K}}$ defined in (11). Therefore Algorithm 3 works with an extra input K and the model projection $\mathcal{P}_{\mathfrak{G}_G}$ becomes $\mathcal{P}_{\mathfrak{G}_{G,K}}$, returning a $\mathfrak{G}_{G,K}$ -sparse approximation to \mathbf{x} .

The convergence analysis of MEIHT with the more general group structures considered here remains the same as for loopless overlapping group models discussed in [3]. We are able to perform the exact projection of $\mathcal{P}_{\mathfrak{G}_G}$ (and $\mathcal{P}_{\mathfrak{G}_{G,K}}$) as discussed in Section 3.5. With the possibility of doing the projection onto the model, we present the convergence results in Corollaries 15 and 16 as corollaries to Theorem 3.1 and Corollary 3.1 in [3] respectively.

Corollary 15. *Consider \mathfrak{G} to be a group model of bounded treewidth and \mathcal{S} to be \mathfrak{G}_G -sparse. Let the matrix $\mathbf{A} \in \{0, 1\}^{m \times N}$ be a model expander matrix with $\epsilon_{\mathfrak{G}_{3G}} < 1/12$ and d ones per column. For any $\mathbf{x} \in \mathbb{R}^N$ and $\mathbf{e} \in \mathbb{R}^m$, the sequence of updates $\mathbf{x}^{(n)}$ of MEIHT with $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ satisfies, for any $n \geq 0$*

$$\|\mathbf{x}^{(n)} - \mathbf{x}_{\mathcal{S}}\|_1 \leq \alpha^n \|\mathbf{x}^{(0)} - \mathbf{x}_{\mathcal{S}}\|_1 + (1 - \alpha^n) \beta \|\mathbf{A}\mathbf{x}_{\mathcal{S}^c} + \mathbf{e}\|_1, \quad (19)$$

where $\alpha = 8\epsilon_{\mathfrak{G}_{3G}} (1 - 4\epsilon_{\mathfrak{G}_{3G}})^{-1} \in (0, 1)$ and $\beta = 4d^{-1} (1 - 12\epsilon_{\mathfrak{G}_{3G}})^{-1} \in (0, 1)$.

The proof of this Corollary can be done analogously to the proof of Theorem 3.1 in [3]. It is thus skipped and the interested reader is referred to [3].

Let us define the ℓ_1 -error of the best \mathfrak{G}_G -term approximation to a vector $\mathbf{x} \in \mathbb{R}^N$

$$\sigma_{\mathfrak{G}_G}(\mathbf{x})_1 = \min_{\mathfrak{G}_G\text{-sparse } \mathbf{z}} \|\mathbf{x} - \mathbf{z}\|_1. \quad (20)$$

This is then used in the following corollary.

Corollary 16. *Consider the setting of Corollary 15. After $n = \left\lceil \log \left(\frac{\|\mathbf{x}\|_1}{\|\mathbf{e}\|_1} \right) / \log \left(\frac{1}{\alpha} \right) \right\rceil$ iterations, MEIHT returns a solution $\hat{\mathbf{x}}$ satisfying*

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_1 \leq c_1 \sigma_{\mathfrak{G}_G}(\mathbf{x})_1 + c_2 \|\mathbf{e}\|_1 \quad (21)$$

where $c_1 = \beta d$ and $c_2 = 1 + \beta$.

Proof. Without loss of generality we initialize MEIHT with $\mathbf{x}^{(0)} = 0$. Upper bounding $1 - \alpha^n$ by 1 and using triangle inequalities with some algebraic manipulations, (19) simplifies to

$$\|\mathbf{x}^{(n)} - \mathbf{x}\|_1 \leq \alpha^n \|\mathbf{x}\|_1 + \beta \|\mathbf{A}\mathbf{x}_{\mathcal{S}^c}\|_1 + \beta \|\mathbf{e}\|_1. \quad (22)$$

Using the fact that \mathbf{A} is a binary matrix with d ones per column we have $\|\mathbf{A}\mathbf{x}_{\mathcal{S}^c}\|_1 \leq d \|\mathbf{x}_{\mathcal{S}^c}\|_1$. We also have $\|\mathbf{e}\|_1 \geq \alpha^n \|\mathbf{x}\|_1$ when $n \geq \log \left(\frac{\|\mathbf{x}\|_1}{\|\mathbf{e}\|_1} \right) / \log \left(\frac{1}{\alpha} \right)$. Applying these bounds to (22) leads to

$$\|\mathbf{x}^{(n)} - \mathbf{x}\|_1 \leq \beta d \|\mathbf{x}_{\mathcal{S}^c}\|_1 + (1 + \beta) \|\mathbf{e}\|_1 \quad (23)$$

for $n = \left\lceil \log \left(\frac{\|\mathbf{x}\|_1}{\|\mathbf{e}\|_1} \right) / \log \left(\frac{1}{\alpha} \right) \right\rceil$. This is equivalent to (21) with $c_1 = \beta d$, $c_2 = 1 + \beta$, $\mathbf{x}^{(n)} = \hat{\mathbf{x}}$ for the given n , and $\sigma_{\mathfrak{G}_G}(\mathbf{x})_1 = \|\mathbf{x}_{\mathcal{S}^c}\|_1$ because $\mathbf{x}_{\mathcal{S}}$ is the best \mathfrak{G}_G -term approximation to the \mathfrak{G}_G -sparse \mathbf{x} . This completes the proof. \square \square

The runtime complexity of MEIHT still depends on the median operation and the complexity of the projection onto the model. However, as observed in [3], the projection onto the model is the dominant operation of the algorithm. Therefore, the complexity of MEIHT is of the order of the complexity of the projection onto the model. In the case of overlapping group models with bounded treewidth, MEIHT achieves a polynomial runtime complexity as shown in Proposition 17 below. On the other hand, when the treewidth of the group model is unbounded MEIHT can still be implemented by using the Bender's decomposition procedure in Section 3.5 for the projection which may have an exponential runtime complexity.

Proposition 17. *The runtime of MEIHT is $\mathcal{O}(w5^w k^2 N M n)$ for the \mathfrak{G}_G -group-sparse model with bounded treewidth w on $\mathcal{O}(wM)$ nodes, where n is the number of iterations, M is the number of groups and N is the signal dimension.*

Proof. The dominant operation of MEIHT is the model projection and hence its runtime is of the order of the cost of performing the projection. Since the projection on the group model with bounded treewidth w can be done through a dynamic programming algorithm that runs in $\mathcal{O}(5^w w k^2 N M)$, as proven in Theorem 10, this proves the result. \square \square

Remark 18. *The convergence results above hold when MEIHT is modified appropriately to solve the standard K -sparse and \mathfrak{G}_G -group-sparse problem with groups having bounded treewidth, where the projection becomes $\mathcal{P}_{\mathfrak{G}_{G,K}}$. However, in this case the runtime complexity grows by a factor of $\mathcal{O}(K)$, as indicated in Remark 11.*

3.5 Exact Projection for General Group Models

In this section we consider the most general version of group models, i.e. $\mathfrak{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_M\}$ is an arbitrary set of groups and $G \in [M]$ and $K \in [N]$ are given budgets. We study the structured sparsity model $\mathfrak{G}_{G,K}$ introduced in Section 2.4. Here additional to the number G of groups that can be selected we bound the number of indices contained in these groups which can be selected by K (i.e. group-sparse model with an additional standard sparsity constraint). Note that setting $K = N$ reduces this model $\mathfrak{G}_{G,K}$ to the general group model \mathfrak{G}_G .

If we want to apply exact projection recovery algorithms like the Model-IHT and MEIHT to group models, iteratively the model projection problem has to be solved, i.e. in each iteration for a given signal $\mathbf{w} \in \mathbb{R}^N$ we have to find the closest signal $\hat{\mathbf{w}}$ which has support in the model. In this section we will derive an efficient procedure based on the idea of Benders' Decomposition to solve the projection problem. This procedure is analysed and implemented in Section 5.

It has been proved that the GROUP MODEL PROJECTION problem for group models without a sparsity condition on the support is NP-hard [4]. Therefore the projection problem on the more general model $\mathfrak{G}_{G,K}$ is NP-hard as well. The latter problem can be reformulated by the integer programming formulation

$$\begin{aligned}
 & \max \mathbf{w}^\top \mathbf{u} \\
 & s.t. \sum_{i=1}^N u_i \leq K \\
 & \sum_{j=1}^M v_j \leq G \\
 & u_i \leq \sum_{j:i \in \mathcal{G}_j} v_j \quad \forall i = 1, \dots, N \\
 & \mathbf{u} \in \{0, 1\}^N, \mathbf{v} \in \{0, 1\}^M.
 \end{aligned} \tag{24}$$

Here \mathbf{w} are the squared entries of the signal, the \mathbf{v} -variables represent the groups and the \mathbf{u} -variables represent the elements which are selected. Note that by choosing $K = N$ we obtain the projection problem for classical group models \mathfrak{G}_G .

To derive an efficient algorithm for the projection problem we use the concept of Benders' Decomposition already studied in [7, 24]. The idea of this approach is to decompose Problem 24 into a master problem and a subproblem. Then iteratively the subproblem is used to derive feasible cuts for the master problem until no feasible cut is found any more. This procedure has been applied to Problem 24 without the sparsity constraint on the \mathbf{u} -variables in [17]. The following results for the more general Problem 24 are based on the idea of Benders' Decomposition and extend the results in [17].

First we can relax the \mathbf{u} -variables in the latter formulation without changing the optimal value, i.e. we may assume $\mathbf{u} \in [0, 1]^N$. We can now reformulate (24) by

$$\begin{aligned}
 & \max \mu \\
 & s.t. \mu \leq \max_{\mathbf{u} \in P(\mathbf{v})} \mathbf{w}^\top \mathbf{u} \\
 & \sum_{j=1}^M v_j \leq G \\
 & \mathbf{v} \in \{0, 1\}^M
 \end{aligned} \tag{25}$$

where $P(\mathbf{v}) = \{\mathbf{u} \in [0, 1]^N \mid \sum_{i=1}^N u_i \leq K, u_i \leq \sum_{j:i \in \mathcal{G}_j} v_j \ i = 1, \dots, N\}$. Dualizing the linear problem

$\max_{\mathbf{u} \in P(\mathbf{v})} \mathbf{w}^\top \mathbf{u}$, the latter problem is equivalent to

$$\begin{aligned}
& \max \mu \\
& \text{s.t. } \mu \leq \min_{\alpha, \beta, \gamma \in P_D} \alpha K + \sum_{i=1}^N \beta_i \left(\sum_{j: i \in \mathcal{G}_j} v_j \right) + \sum_{i=1}^N \gamma_i \\
& \sum_{j=1}^M v_j \leq G \\
& \mathbf{v} \in \{0, 1\}^M
\end{aligned} \tag{26}$$

where $P_D = \{\alpha, \beta, \gamma \geq 0 \mid \alpha + \beta_i + \gamma_i \geq w_i \ i = 1, \dots, N\}$ is the feasible set of the dual problem. Since P_D is a polyhedron and the minimum in (26) exists, the first constraint in (26) holds if and only if it holds for each vertex $(\alpha^l, \beta^l, \gamma^l)$ of P_D . Therefore Problem (26) can be reformulated by

$$\begin{aligned}
& \max \mu \\
& \text{s.t. } \mu \leq \alpha^l K + \sum_{i=1}^N \beta_i^l \left(\sum_{j: i \in \mathcal{G}_j} v_j \right) + \sum_{i=1}^N \gamma_i^l \quad l = 1, \dots, t \\
& \sum_{j=1}^M v_j \leq G \\
& \mathbf{v} \in \{0, 1\}^M
\end{aligned} \tag{27}$$

where $(\alpha^1, \beta^1, \gamma^1), \dots, (\alpha^t, \beta^t, \gamma^t)$ are the vertices of P_D . Each of the latter constraints given by a vertex of P_D is called *optimality cut*.

The idea of Bender's Decomposition is, starting from Problem (27) containing no optimality cut (called the *master problem*), to iteratively calculate the optimal (\mathbf{v}^*, μ^*) and then find a optimality cut which cuts off the latter optimal solution. In each step the most violating optimality cut is determined by solving

$$\begin{aligned}
& \min \alpha K + \sum_{i=1}^N \beta_i \left(\sum_{j: i \in \mathcal{G}_j} v_j^* \right) + \sum_{i=1}^N \gamma_i \\
& \text{s.t. } \alpha + \beta_i + \gamma_i \geq w_i \quad i = 1, \dots, N \\
& \alpha, \beta, \gamma \geq 0,
\end{aligned} \tag{28}$$

for the actual optimal solution \mathbf{v}^* . If the optimal solution fulfills

$$\mu^* > \alpha^* K + \sum_{i=1}^N \beta_i^* \left(\sum_{j: i \in \mathcal{G}_j} v_j^* \right) + \sum_{i=1}^N \gamma_i^*,$$

then the optimality cut related to the optimal vertex $(\alpha^*, \beta^*, \gamma^*)$ is added to the master problem. This procedure is iterated until the latter inequality is not true any more. The last optimal \mathbf{v}^* must then be optimal for Problem (24) since the first constraint in (26) is then true for \mathbf{v}^* .

If we use the latter Benders' Decomposition approach it is desired to use fast algorithms for the master-problem (27) and the subproblem (28) in each iteration. By the following lemma an optimal solution of the subproblem can be easily calculated.

Lemma 19. *For a given solution $\mathbf{v} \in \{0, 1\}^M$ we define $I_{\mathbf{v}} := \{i = 1, \dots, N \mid \sum_{j: i \in \mathcal{G}_j} v_j > 0\}$ and $I_{\mathbf{v}}^K$ by the indices of the K largest values w_i for $i \in I_{\mathbf{v}}$. An optimal solution of Problem (28) is then given by*

$(\alpha^*, \beta^*, \gamma^*)$ where $\alpha^* = \max_{i \in I_{\mathbf{v}} \setminus I_{\mathbf{v}}^K} w_i$ and

$$(\beta_i^*, \gamma_i^*) = \begin{cases} (w_i, 0) & \text{if } i \in [N] \setminus I_{\mathbf{v}} \\ (0, 0) & \text{if } i \in I_{\mathbf{v}} \setminus I_{\mathbf{v}}^K \\ (0, w_i - \alpha^*) & \text{if } i \in I_{\mathbf{v}}^K \end{cases}$$

Proof. Note that for a given \mathbf{v}^* the dual problem of subproblem (28) is

$$\begin{aligned} & \max \mathbf{w}^\top \mathbf{u} \\ & \text{s.t.} \quad \sum_{i=1}^N u_i \leq K \\ & \quad u_i \leq \sum_{j:i \in \mathcal{G}_j} v_j^* \quad i = 1, \dots, N \\ & \quad \mathbf{u} \in [0, 1]^N. \end{aligned} \tag{29}$$

It is easy to see that there exists an optimal solution \mathbf{u}^* of Problem (29) where $u_i = 1$ if and only if $i \in I_{\mathbf{v}}^K$ and $u_i^* = 0$ otherwise. We will use the dual slack conditions

$$\left(\sum_{j:i \in \mathcal{G}_j} v_j^* - u_i^* \right) \beta_i = 0 \text{ and } (1 - u_i^*) \gamma_i = 0 \quad i = 1, \dots, N, \tag{30}$$

to derive the optimal values for α, β, γ . We obtain the following 4 cases:

Case 1: If $u_i^* = 0$ and $\sum_{j:i \in \mathcal{G}_j} v_j^* > 0$, i.e. $i \in I_{\mathbf{v}} \setminus I_{\mathbf{v}}^K$, then by conditions (30) we have $\beta_i = \gamma_i = 0$. To ensure the constraint $\alpha + \beta_i + \gamma_i \geq w_i$, the value of α must be at least $\max_{i \in I_{\mathbf{v}} \setminus I_{\mathbf{v}}^K} w_i$.

Case 2: If $u_i^* = 0$ and $\sum_{j:i \in \mathcal{G}_j} v_j^* = 0$, then $i \in [N] \setminus I_{\mathbf{v}}$ and the objective coefficient of β_i in Problem (28) is 0. Therefore we can increase β_i as much as we want without changing the objective value and therefore we set $\beta_i = w_i$ to ensure the i -th constraint $\alpha + \beta_i + \gamma_i \geq w_i$ and set $\gamma_i = 0$.

Case 3: If $u_i^* = 1$, i.e. $i \in I_{\mathbf{v}}^K$, and $\sum_{j:i \in \mathcal{G}_j} v_j^* > 1$ then by condition (30) $\beta_i = 0$. Therefore to ensure the i -th constraint $\alpha + \beta_i + \gamma_i \geq w_i$ the value of γ_i must be at least $w_i - \alpha$ and since we minimize γ_i in the objective function the latter holds with equality.

Case 4: If $u_i^* = 1$, i.e. $i \in I_{\mathbf{v}}^K$, and $\sum_{j:i \in \mathcal{G}_j} v_j^* = 1$ then we cannot use condition (30) to derive the values for β_i and γ_i . Nevertheless in this case both variables have an objective coefficient of 1 while α has an objective coefficient of K . By increasing α by 1 the objective value increases by K . In Cases 1 and 2 nothing changes, while for each index in Case 3 we could decrease γ_i by 1 to remain feasible. But since at most K indices i fulfil the conditions of Case 3 we cannot improve our objective value by this strategy. Therefore α has to be selected as small as possible in Case 4, i.e. by Case 1 we set $\alpha = \max_{i \in I_{\mathbf{v}} \setminus I_{\mathbf{v}}^K} w_i$ and to ensure feasibility we set $\gamma_i = w_i - \alpha$. This concludes the proof. \square \square

Theorem 20. *An optimal solution of subproblem (28) can be calculated in time $\mathcal{O}(NM|G_{\max}| + N \log(N))$.*

Proof. The set $I_{\mathbf{v}}$ can be calculated in time $\mathcal{O}(NM|G_{\max}|)$ by going through all groups for each index $i \in [N]$ and checking if the index is contained in one of the groups. To obtain $I_{\mathbf{v}}^K$ we have to sort $\{w_i \mid i \in I_{\mathbf{v}}\}$ in non-increasing order and pick the first K indices, which can be done in time $\mathcal{O}(N \log(N))$. Assigning all values to (α, β, γ) can be done in time $\mathcal{O}(N)$. \square \square

The following theorem states that the masterproblem can be solved in pseudopolynomial time if the number of constraints is fixed. Nevertheless note that the number of iterations of the procedure described above may be exponential in N and M .

Theorem 21. *Problem (27) with t constraints can be solved in $\mathcal{O}(MG(NW)^t)$ where $W = \max_{i \in [N]} w_i$.*

Proof. Problem (27) with t constraints can be reformulated by

$$\begin{aligned} \max_{\mathbf{v} \in \{0,1\}^M} \min_{l=1,\dots,t} \alpha^l K + \sum_{i=1}^N \beta_i^l \left(\sum_{j:i \in \mathcal{G}_j} v_j \right) + \sum_{i=1}^N \gamma_i^l \\ \text{s.t.} \quad \sum_{j=1}^M v_j \leq G. \end{aligned} \tag{31}$$

The latter problem is a special case of the robust knapsack problem with discrete uncertainty (sometimes called robust selection problem) with additional uncertain constant. In [12] the authors mention that the problem with an uncertain constant is equivalent to the problem without such a constant. Furthermore using the result in [35] Problem (31) can be solved in $\mathcal{O}(MGC^t)$ where

$$C := \max_{l=1,\dots,t} \alpha^l + \sum_{i=1}^N \beta_i^l + \sum_{i=1}^N \gamma_i^l.$$

Since for all solutions $(\alpha^l, \beta^l, \gamma^l)$ generated in Lemma 19 it holds that $\alpha^l, \beta_i^l, \gamma_i^l \leq \max_{i \in [N]} w_i$ we have $C \leq (2N + 1)W$ which proves the result. \square \square

4 Algorithms with Approximation Projection Oracles

As mentioned in the previous sections solving the GROUP MODEL PROJECTION problem is *NP*-hard in general. Therefore to use classical algorithms as the Model-IHT or the MEIHT we have to solve an *NP*-hard problem in each iteration. To tackle problems like this the authors in [28, 27] introduced an algorithm called *Approximate Model-IHT* (AM-IHT) which is based on the idea of the classical IHT but which does not require an exact projection oracle (see Section 2.5). Instead the authors show that under certain assumptions on the measurement matrix a signal can be recovered by just using two approximation variants of the projection problem which they call head- and tail-approximation (for further details again see Section 2.5). In this section we will apply the latter results to group models with bounded frequency, i.e. the maximum number of groups an element is contained in is bounded by f . Note that this restriction is still reasonable since from Theorem 13 we know that GROUP MODEL PROJECTION is NP-hard for group models of frequency 4. Indeed, it is not hard to see that GROUP MODEL PROJECTION is NP-hard even when restricted to group models of frequency 2. A particular example is the graphic case, where each element is contained in at most two groups. Understanding this case was proposed as an open problem by Baldassarre et al. [4].

Furthermore we apply the theoretical results derived in [28, 27] to group models and show that the number of required measurements compared to the classical structured sparsity case increases by just a constant factor. In Section 5 we will computationally compare the AM-IHT to the Model-IHT and the MEIHT on interval groups.

4.1 Head- and Tail-Approximations for Group Models with Low Frequency

In this section we derive head- and tail-approximations for the case of group models with bounded frequency. We first recall the definition of head- and tail-approximations for the case of group models. Assume we have given a group model \mathfrak{G} together with $G \in \mathbb{N}$.

Given a vector \mathbf{x} , let \mathcal{H} be an algorithm that computes an index set $\mathcal{H}(\mathbf{x})$ with support in $\mathfrak{G}_{G'}$ for some $G' \in \mathbb{N}$. Then, given some $\alpha \in \mathbb{R}$ (typically $\alpha < 1$) we say that \mathcal{H} is an (α, G, G', p) -*head approximation* if

$$\|\mathcal{H}(\mathbf{x})\|_p \geq \alpha \cdot \|\mathbf{x}_{\mathcal{S}}\|_p \text{ for all } \mathcal{S} \in \mathfrak{G}_G. \tag{32}$$

In other words, \mathcal{H} uses G' many groups to cover at least an α -fraction of the maximum total weight covered by G groups. Note that G' can be chosen larger than G .

Moreover, let \mathcal{T} be an algorithm which computes an index set $\mathcal{T}(\mathbf{x})$ with support in \mathfrak{G}_G . Given some $\beta \in \mathbb{R}$ (typically $\beta > 1$) we say that \mathcal{T} is a (β, G, G', p) -tail approximation if

$$\|\mathbf{x} - \mathcal{T}(\mathbf{x})\|_p \leq \beta \cdot \|\mathbf{x} - \mathbf{x}_{\mathcal{S}}\|_p \text{ for all } \mathcal{S} \in \mathfrak{G}_G. \quad (33)$$

This means that \mathcal{T} may use G' many groups to leave at most a β -fraction of weight uncovered compared to the minimum total weight left uncovered by G groups.

In the following we derive the head- and tail-approximation just for the case $p = 1$. Note that equivalent approximation procedures can be easily derived for the case $p = 2$ by replacing the accuracies α and β by $\sqrt{\alpha}$ and $\sqrt{\beta}$ and using the weights x_i^2 instead of $|x_i|$ in the latter proofs. We will first present a result which implies the existence of a head approximation.

Theorem 22 (Hochbaum and Pathria [29]). *For each $\epsilon > 0$ there exists an $((1 - \epsilon), G, \lceil G \log_2(1/\epsilon) \rceil, 1)$ -head approximation running in polynomial time.*

The algorithm derived in [29] was designed to solve the MAXIMUM WEIGHT COVERAGE problem and is based on a simple greedy method. The idea is to iteratively select the group which covers the largest amount of uncovered weight. It is proven by the authors that if you are allowed to select enough groups, namely $\lceil G \log_2(1/\epsilon) \rceil$, then the optimal value is approximated up to an accuracy of $(1 - \epsilon)$. The greedy procedure is given in Algorithm 4. Note that for a given signal \mathbf{x} and a group $\mathcal{G} \in \mathfrak{G}$ we define

$$w(\mathcal{G}) := \sum_{i \in \mathcal{G}} |x_i|.$$

Algorithm 4 Head-Approximation

Input $\mathbf{x} \in \mathbb{R}^N$, $G \in \mathbb{N}$, \mathfrak{G} , $\epsilon > 0$

Output $\mathcal{H}(\mathbf{x}, G, \mathfrak{G}, \epsilon)$, a head-approximation to \mathbf{x}

- 1: $\mathbf{z} \leftarrow \mathbf{0}$
 - 2: **for** $i = 1 \dots G$ **do**
 - 3: Sort the groups in \mathfrak{G} by decreasing total weight: $w(\mathcal{G}_{i_1}) \geq \dots \geq w(\mathcal{G}_{i_M})$.
 - 4: $z_j \leftarrow x_j \forall j \in \mathcal{G}_{i_1}$
 - 5: Delete the indices in \mathcal{G}_{i_1} from all groups in \mathfrak{G} .
 - 6: **end for**
 - 7: Return: $\mathcal{H}(\mathbf{x}, G, \mathfrak{G}, \epsilon) \leftarrow \mathbf{z}$
-

Next we derive a tail-approximation for our problems based on the idea of LP rounding. Note that in contrast to the head-approximation, the run-time bound of the following tail-approximation depends on the frequency of the group model.

Theorem 23. *Suppose the frequency of the group model is f . For any $\epsilon > 0$ and $\alpha = (1 + \epsilon^{-1})f$ there exists an $((1 + \epsilon), G, \alpha G, 1)$ -tail approximation running in polynomial time.*

Proof. Given a signal $\mathbf{x} \in \mathbb{R}^N$, we define $\mathbf{w} = (|x_i|)_{i \in [N]}$. We consider the following linear relaxation of the GROUP MODEL PROJECTION problem

$$\begin{aligned} \max \quad & \mathbf{w}^\top \mathbf{u} \\ \text{s.t.} \quad & \sum_{i \in [M]} v_i = G \\ & u_j \leq \sum_{j \in \mathcal{G}_i} v_i \quad \text{for all } j \in [N] \\ & \mathbf{u} \in [0, 1]^N, \mathbf{v} \in [0, 1]^M. \end{aligned} \quad (34)$$

Consider an optimal solution (\mathbf{u}, \mathbf{v}) of (34). We compute a group cover $\mathcal{S} \subseteq \mathfrak{G}$ by

$$\mathcal{S} = \{\mathcal{G}_i \in \mathfrak{G} : v_i \geq \alpha^{-1}\}.$$

Note that \mathcal{S} contains at most αG many groups, since

$$|\mathcal{S}| \leq \alpha \sum_{i=1}^M v_i = \alpha G.$$

It remains to show that \mathcal{S} is a tail approximation. To this end let R be the set of indices only barely covered by y , i.e.

$$R = \{j \in [N] : \sum_{j \in \mathcal{G}_i} v_i \leq (1 + \epsilon^{-1})^{-1}\}.$$

Note that

$$\mathcal{S} \text{ covers every element } j \in [N] \setminus R, \quad (35)$$

since $j \notin R$ implies

$$\sum_{j \in \mathcal{G}_i} v_i \geq u_j > (1 + \epsilon^{-1})^{-1}$$

and hence

$$v_i \geq \frac{(1 + \epsilon^{-1})^{-1}}{f} = \alpha^{-1}$$

for some i with $j \in \mathcal{G}_i$. Moreover, note that

$$u_j \leq \sum_{j \in \mathcal{G}_i} v_i \leq (1 + \epsilon^{-1})^{-1}$$

and hence

$$\frac{1 - u_j}{1 - (1 + \epsilon^{-1})^{-1}} \geq 1 \text{ for each } j \in R. \quad (36)$$

We obtain the inequalities

$$\begin{aligned} \sum_{i \in [N] \setminus \mathcal{S}} w_i &\leq \sum_{i \in R} w_i \leq \sum_{j \in R} \frac{1 - u_j}{1 - (1 + \epsilon^{-1})^{-1}} w_j \\ &\leq (1 + \epsilon) \sum_{j \in R} w_j (1 - u_j) \\ &\leq (1 + \epsilon) (\mathbf{w}^\top \mathbf{1} - \mathbf{w}^\top \mathbf{u}) \leq (1 + \epsilon) \text{OPT}, \end{aligned}$$

where we used (35) and (36). This completes the proof. \square \square

4.2 AM-IHT and AM-MEIHT for Group Models

As in the previous sections for a sensing matrix \mathbf{A} and the true signal \mathbf{x} we have a noisy measurement vector $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ for some noise vector \mathbf{e} . The task consists in recovering the original signal \mathbf{x} , or a vector close to \mathbf{x} . Furthermore we have given a group model \mathfrak{G} together with $G \in \mathbb{N}$ with frequency $f \in [N]$.

In the last section we derived polynomial time algorithms for an $((1 - \epsilon), G, \lceil G \log_2(1/\epsilon) \rceil, 2)$ -head approximation and an $((1 + \epsilon), G, (1 + \epsilon^{-1})fG, 2)$ -tail approximation. Note that we can use any G here. Using the results in Section 2.5, we obtain convergence of the AM-IHT for group models if \mathcal{T} is an $((1 + \epsilon), G, G_T, 2)$ -tail approximation, \mathcal{H} is an $((1 - \epsilon), G_T + G, G_H, 2)$ -head approximation, where $G_T := (1 + \epsilon^{-1})fG$ and $G_H := \lceil (G_T + G) \log_2(1/\epsilon) \rceil$, and the sensing matrix \mathbf{A} has $\mathfrak{G}_{\tilde{G}}$ -RIP with $\tilde{G} = G + G_T + G_H$. Note that $\tilde{G} \in \mathcal{O}(G)$ for fixed accuracy $\epsilon > 0$ and frequency f . Furthermore $|\mathfrak{G}_{\tilde{G}}| \in \mathcal{O}(M^{cG})$ for a constant c . Using

the bound on the measurements derived in Section 3 we obtain that for a sub-Gaussian random matrix \mathbf{A} the number of required measurements to obtain recovery with high probability is

$$m = \mathcal{O}(\delta^{-2} g_{\max} \log(\delta^{-1}) + cG \log(M))$$

which differs by just a constant factor from the number of measurements required in the case of exact projections (see Section 3).

4.3 Extension to within-group sparsity and beyond.

The head and tail approximation approach can be extended far beyond the standard group-sparsity model. It still works even if we are considering K -sparse and G -group-sparse (i.e. $\mathfrak{G}_{G,K}$ -sparse) vectors in our model, for example.

The reason is that the GROUP MODEL PROJECTION can be head approximated to a constant even in this case. Formally, if we search for the K weight maximal elements covered by G many groups, we are maximizing a submodular function subject to a knapsack constraint.³ This is known to be approximable to a constant factor (cf. Kulik et al. [36, 37] and related work). Suppose we delete the covered elements and run such an approximation algorithm again. Then, after a constant number of steps, we obtain a collection of groups and elements such that the total weight of the elements is at least an $(1 - \epsilon)$ -fraction of the total weight that a K -sparse and G -group-sparse solution could ever have. Moreover, the sparsity budgets are exceeded only by a constant factor each.

Similarly, the analysis given in the proof of Theorem 23 works even if we impose sparsity on both, groups and elements. Again, we obtain a solution that has a $(1 + \epsilon)$ -tail guarantee whose support exceeds that of a G -group-sparse K -sparse vector by at most a constant factor if we assume a bounded frequency. This leads to the positive consequences detailed above.

More generally, any knapsack constraints on groups and elements can be handled, leading to head and tail approximations in the case when there are non-uniform sparsity budgets on the groups and elements. However the corresponding head approximations are rather involved, and certainly much less efficient than the simple greedy procedure proposed by the Hochbaum and Pathria algorithm [29].

5 Computations

In this section we present the computational results for the Model-IHT, MEIHT, AM-IHT and AM-EIHT presented in Section 2 for block-group instances. Precisely, we study block-groups, i.e. each group $\mathcal{G} \in \mathfrak{G}$ is a set of sequenced indices, $\mathcal{G} = [s, t] \cap [N]$, where $1 \leq s < t \leq N$ and each group has the same size $|\mathcal{G}| = l$. For a given dimension N we generate blocks of size $l = \lfloor 0.02N \rfloor$. We consider two types of block-models, one where the successive groups overlap in $\lfloor \frac{l-1}{2} \rfloor$ items and another where they overlap in $l - 1$ items. Note that the frequency is then given by $f = 2$ or $f = l$ respectively.

We run all algorithms for random signals $\mathbf{x} \in \mathbb{R}^N$ in dimension $N \in \{100, 200, \dots, 900\}$. For each dimension we vary the number of measurements $m \in \{20, 40, \dots, N\}$ and generate 20 random matrices $\mathbf{A} \in \mathbb{R}^{m \times N}$ each together with a random signal $\mathbf{x} \in \mathbb{R}^N$. We assume there is no noise, i.e. $\mathbf{e} = 0$. For a given group model \mathfrak{G} the support of the signal \mathbf{x} is determined as the union of G randomly selected groups. The components of \mathbf{x} in the support are calculated as identical and independent draws from a standard Gaussian distribution while all other components are set to 0. Our computations are processed for two classes of random matrices, Gaussian matrices and expander matrices as described in Section 2. The Gaussian matrices are generated by drawing identical and independent values from a standard Gaussian distribution for each entry of \mathbf{A} and afterwards normalizing each entry by $\frac{1}{\sqrt{m}}$. The expander matrices are generated by randomly selecting $d = \lfloor 2 \log(N) / \log(Gl) \rfloor$ indices in $\{1, \dots, m\}$ for each column of the matrix. The choice of d is motivated by the choice in [3]. Each of the algorithms is stopped if either the number of iterations reaches 1000 or if for any iteration $i + 1$ we have $\|\mathbf{x}^{(i+1)} - \mathbf{x}^{(i)}\|_p < 10^{-5}$. For the error in each iteration we use $p = 1$

³Actually, we are maximizing a submodular function subject to a uniform matroid constraint which is a simpler problem.

for the calculations corresponding to the expander matrices and $p = 2$ for the Gaussian matrices. After the determination of the algorithm we calculate the relative error of the returned signal $\hat{\mathbf{x}}$ to the true signal \mathbf{x} , i.e. we calculate

$$\frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_p}{\|\mathbf{x}\|_p}.$$

Again we use $p = 1$ for the calculations corresponding to the expander matrices and $p = 2$ for the Gaussian matrices. We call a signal *recovered* if the relative error is smaller than 10^{-5} . For the AM-IHT and the AM-EIHT the approximation accuracy of the head- and the tail approximation algorithms are set to $\alpha = 0.95$ and $\beta = 1.05$.

For the exact signal approximation problem which has to be solved in each iteration of the Model-IHT and the MEIHT we implemented the Benders' decomposition procedure presented in Section 3.5. To this end the master problem is solved by CPLEX 12.6 while each optimal solution of the subproblem is calculated using the result of Lemma 19. Regarding the AM-IHT, for the head-approximation we implemented the greedy procedure given in Algorithm 4 while for the tail-approximation we implemented the procedure of Theorem 23. Again the LP in the latter procedure is solved by CPLEX 12.6. All computations were calculated on a cluster of 64-bit Intel(R) Xeon(R) CPU E5-2603 processors running at 1.60 GHz with 15MB cache.

The results of the computations are presented in the following diagrams. For all instances we measure the smallest number of measurements $m^\#$ for which the median of the relative error to the true signal is at most 10^{-5} , i.e. the smallest number of measurements for which at least 50% of the signals were recovered. Furthermore we show the average number of iterations and the average time in seconds the algorithms need to successfully recover a signal, given $m^\#$ number of measurements. We stop increasing the number of measurements m if $m^\#$ is reached.

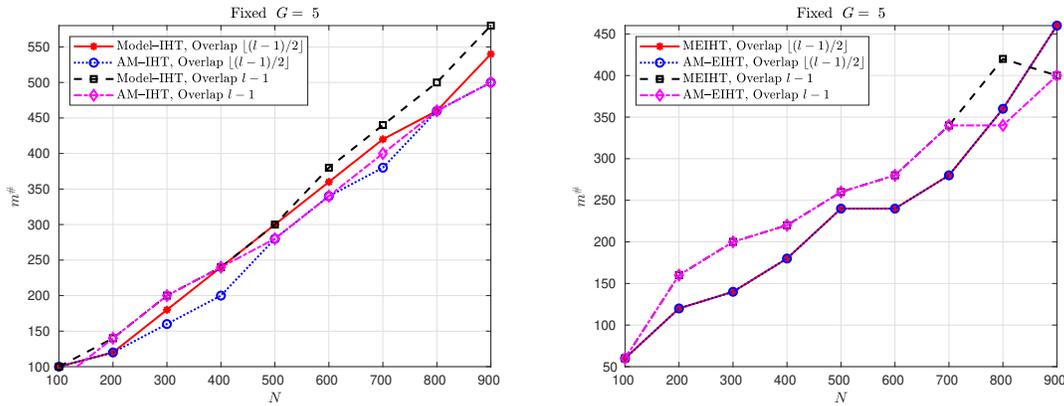


Figure 5: Smallest number of measurements $m^\#$ for which the median of the relative error over all random matrices is at most 10^{-5} .

In Figures 5, 6 and 7 we show the development of $m^\#$, the number of iterations and the runtime in seconds of all algorithms over all dimensions $N \in \{100, 200, \dots, 900\}$ for block-groups, generated as explained above, with fixed value $G = 5$.

The smallest number of measurements $m^\#$ which leads to a median relative error of at most 10^{-5} is nearly linear in the dimension; see Figure 5. For all algorithms the corresponding $m^\#$ is very close even for different number of overlaps. Nevertheless the number of measurements $m^\#$ is smaller for expander matrices than for Gaussian matrices. Furthermore in the expander case the instances with overlap $\lfloor \frac{l-1}{2} \rfloor$ have a smaller $m^\#$. The average number of iterations performed by the algorithms fluctuates a lot for the Gaussian case. Here the value increases slowly for the Model-IHT while it increases more rapidly for the AM-IHT. In the expander case the number of iterations is very close to each other for all algorithms and lies between 50 and 70 most of the time; see Figure 6. The drop from $N = 100$ to $N = 200$ is due to the small value of d when $N = 100$. Furthermore the number of iterations is much lower in the expander case.

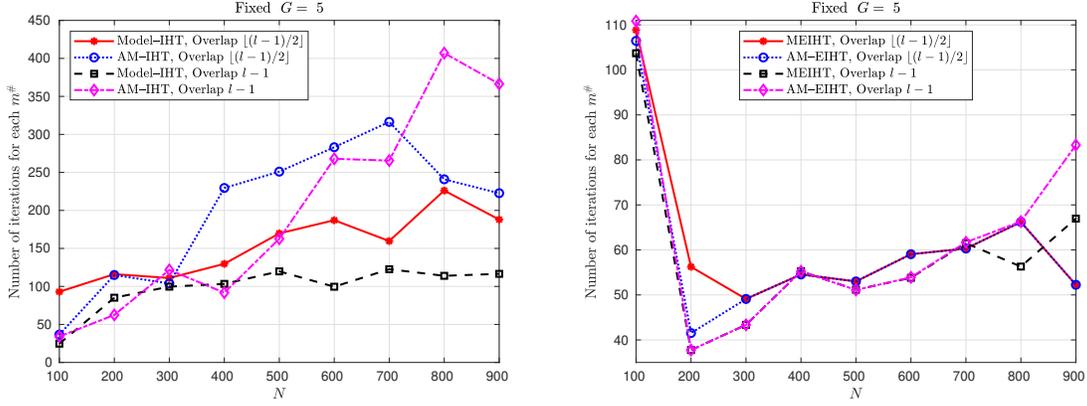


Figure 6: Average number of iterations performed by the algorithm to successfully recover the signal with $m^\#$ measurements.

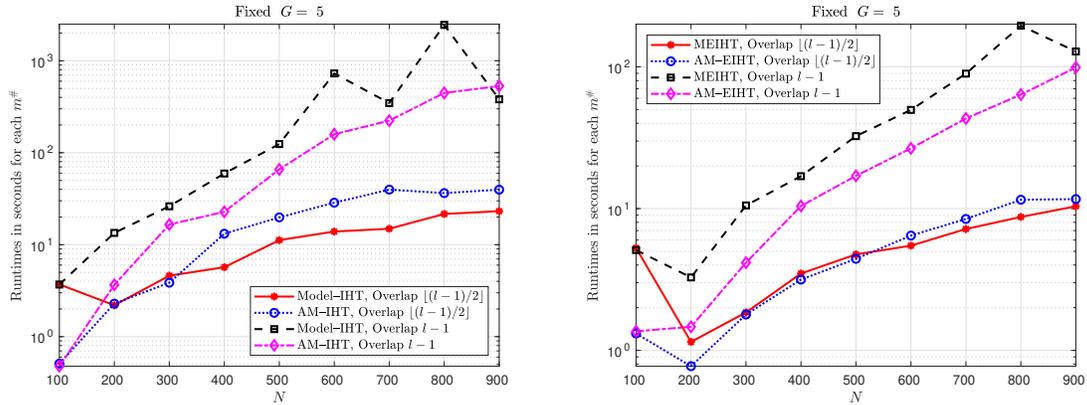


Figure 7: Average time in seconds needed by the algorithm to successfully recover the signal with \bar{m} measurements.

The average runtime for the Gaussian case is a bit larger than the runtime for the expander case as expected, since operations with dense matrices is more costly than the sparse ones. However, it may be due to the larger number of iterations; see Figure 7. Furthermore the runtime for the instances with $l - 1$ overlap is much larger in both cases. Here the AM-IHT (AM-EIHT) is faster than the Model-IHT (MEIHT) for the instances with overlap $l - 1$ while it is slightly slower for the others.

In Figures 8, 9 and 10 we show the same properties as above, but for varying $G \in \{1, 2, \dots, 9\}$, a fixed dimension of $N = 200$ and d fixed to 7 for all values of G . Similar to Figure 5 the value $m^\#$ seems to be linear in G (see Figure 8). Just the Model-IHT (MEIHT) for blocks with overlap $l - 1$ seems to require an exponential number of measurements in G to guarantee a small median relative error. Here the AM-IHT (AM-EIHT) performs much better. Interestingly the number of iterations does not change a lot for increasing G for Gaussian matrices while it grows for the AM-EIHT in the expander case. This is in contrast to the iteration results with increasing N ; see Figure 6. The runtime of all algorithms increases slowly with increasing G , except for the IHT for blocks with overlap $l - 1$ the runtime explodes. For both instances the AM-IHT (AM-EIHT) is faster than the Model-IHT (MEIHT).

To conclude this section we selected the instances for dimension $N = 800$ and $G = 5$ and present the development of the median relative error over the number of measurements m ; see Figure 11. In the expander case the median relative error decreases rapidly and is nearly 0 already for $\frac{m}{N} \approx 0.45$. Just for the MEIHT for

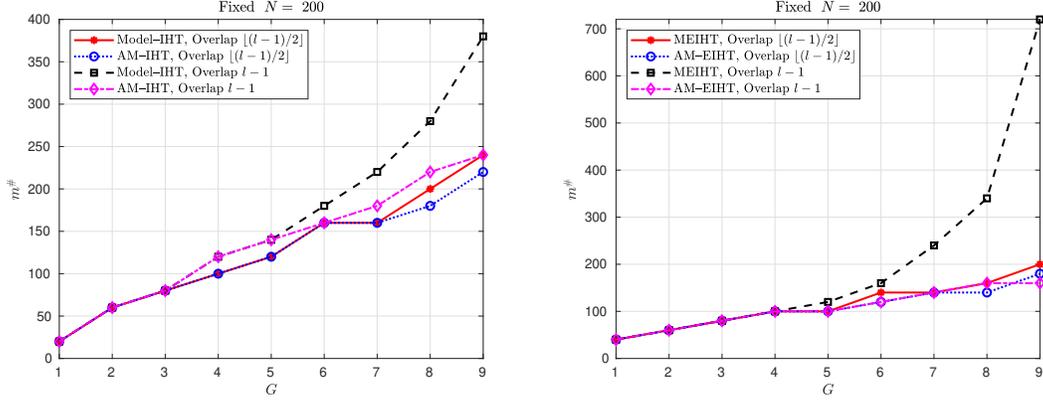


Figure 8: Smallest number of measurements $m^\#$ for which the median of the relative error over all random matrices is at most 10^{-5} .

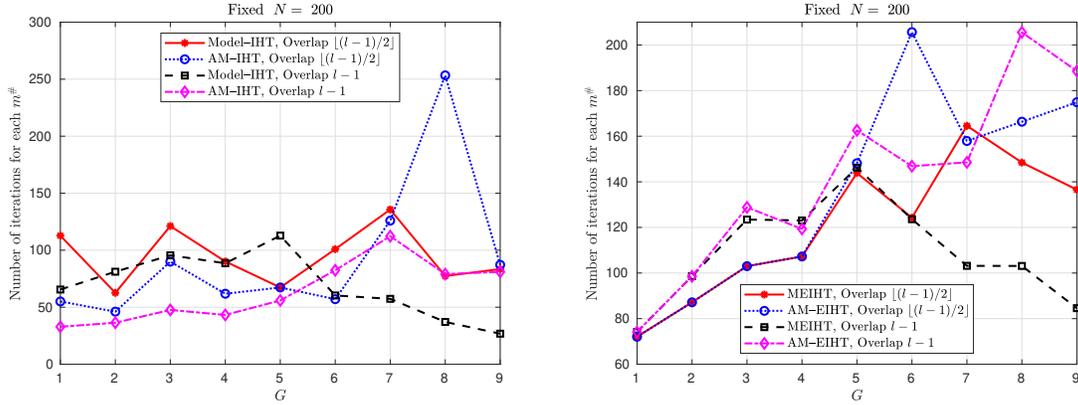


Figure 9: Average number of iterations performed by the algorithm to successfully recover the signal with $m^\#$ measurements.

blocks with overlap $l - 1$ the relative error is close to 0 not until $\frac{m}{N} \approx 0.55$. For the Gaussian case the results look similar with the only difference that a median relative error close to 0 is reached not until $\frac{m}{N} \approx 0.6$.

6 Conclusion

In this paper we revisited the model-based compressed sensing problem focusing on overlapping group models with bounded treewidth and low frequency. We derived a polynomial time dynamic programming algorithm to solve the exact projection problem for group models with bounded treewidth, which is more general than the state-of-the-art considering loopless overlapping models. For general group models we derived an algorithm based on the idea of Bender's Decomposition, which may run in exponential time but often performs better than dynamic programming in practice. We proved that the latter procedure is generalizable from group-sparse models to group-sparse plus standard sparse models. The most dominant operation of iterative exact projection algorithms is the model projection. Hence our results show that the Model-IHT and the MEIHT run in polynomial time for group models with bounded treewidth. Alternatively, for group models with bounded frequency we show that another class of less accurate algorithms run in polynomial time. More precisely the AM-IHT and the AM-EIHT are algorithms using head- and tail-approximations

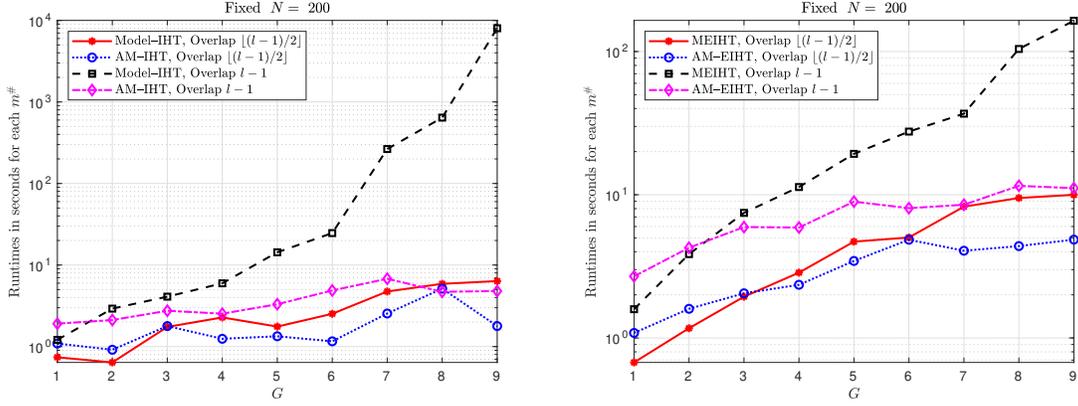


Figure 10: Average time in seconds needed by the algorithm to successfully recover the signal with $m^\#$ measurements.

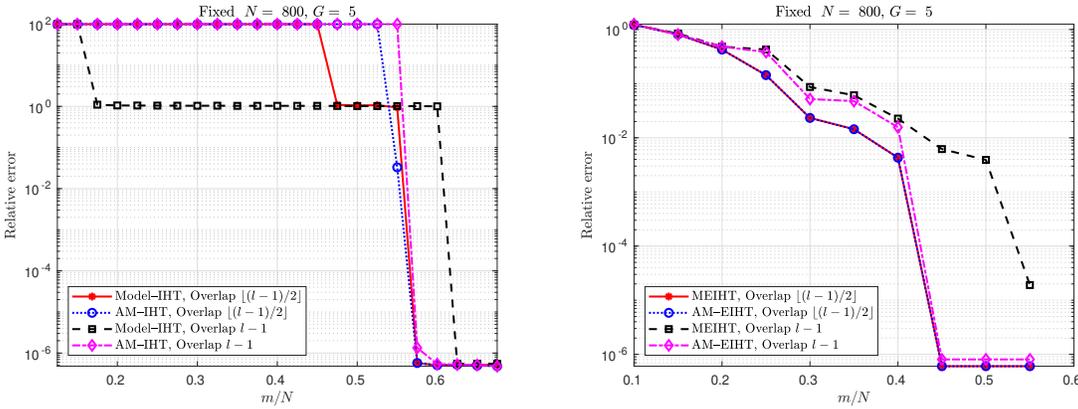


Figure 11: Median relative error for instances with dimension $N = 800$ and $G = 5$. The diagrams are presented in log-scale.

instead of exact projections.

Using Bender's Decomposition (with Gaussian and model-expander sensing matrices) we compared the minimum number of measurements required by, and runtimes of, each of the four algorithms (Model-IHT, MEIHT, AM-IHT and AM-EIHT) to achieve a given accuracy. In summary the experimental results on overlapping block groups seem to indicate that the number of required measurements to recover a signal is smaller for expander matrices than for Gaussian matrices. Furthermore, we could observe that the number of measurements to ensure a small relative error is smaller for the approximate versions of the algorithms. Furthermore the run-time gets much larger for Gaussian matrices with increasing N than for expander matrices, which might be just what is expected when applying dense versus sparse matrices. In general the approximate versions of the algorithms may have a larger number of iterations but the run-time is lower. This indicates that the larger number of iterations can be compensated by the faster computation of the approximate projection problems in each iteration.

7 Acknowledgements

BB acknowledges the support from the funding by the German Federal Ministry of Education and Research, administered by Alexander von Humboldt Foundation, for the German Research Chair at AIMS South Africa.

References

- [1] K. J. Ahn, S. Guha, and A. McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 5–14. ACM, 2012.
- [2] M. J. Alam, M. A. Bekos, M. Kaufmann, P. Kindermann, S. G. Kobourov, and A. Wolff. Smooth orthogonal drawings of planar graphs. In *LATIN 2014: Theoretical Informatics - 11th Latin American Symposium, Montevideo, Uruguay, March 31 - April 4, 2014. Proceedings*, pages 144–155, 2014.
- [3] B. Bah, L. Baldassarre, and V. Cevher. Model-based sketching and recovery with expanders. In *SODA*, pages 1529–1543. SIAM, 2014.
- [4] L. Baldassarre, N. Bhan, V. Cevher, A. Kyrillidis, and S. Satpathi. Group-sparse model selection: Hardness and relaxations. *IEEE Trans. Information Theory*, 62(11):6508–6534, 2016.
- [5] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- [6] R. G. Baraniuk, V. Cevher, and M. B. Wakin. Low-dimensional models for dimensionality reduction and signal recovery: A geometric perspective. *Proceedings of the IEEE*, 98(6):959–971, 2010.
- [7] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962.
- [8] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 798–805. IEEE, 2008.
- [9] T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of linear subspaces. *IEEE Trans. Info. Theory*, pages 30–56, 2007.
- [10] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.
- [11] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- [12] C. Buchheim and J. Kurtz. Robust combinatorial optimization under convex and discrete cost uncertainty. *EURO Journal on Computational Optimization*, 6(3):211–238, 2018.
- [13] E. Candès and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.
- [14] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [15] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.

- [16] V. Chandar. A negative result concerning explicit matrices with the restricted isometry property. Technical report, 2008.
- [17] J. Cordeau, F. Furini, and I. Ljubic. Benders decomposition for very large scale partial set covering and maximal covering problems. Technical report, 2018.
- [18] R. DeVore. Deterministic constructions of compressed sensing matrices. *Journal of Complexity*, 23(4):918–925, 2007.
- [19] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [20] C. Dwork, F. McSherry, and K. Talwar. The price of privacy and the limits of lp decoding. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 85–94. ACM, 2007.
- [21] Y. C. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE Transactions on Information Theory*, 55(11):5302–5316, 2009.
- [22] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.
- [23] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [24] A. M. Geoffrion. Generalized benders decomposition. *Journal of optimization theory and applications*, 10(4):237–260, 1972.
- [25] A. C. Gilbert and K. Levchenko. Compressing network graphs. In *Proceedings of the LinkKDD workshop at the 10th ACM Conference on KDD*, volume 124, 2004.
- [26] C. Hegde, P. Indyk, and L. Schmidt. Approximation-tolerant model-based compressive sensing. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1544–1561. Society for Industrial and Applied Mathematics, 2014.
- [27] C. Hegde, P. Indyk, and L. Schmidt. Approximation algorithms for model-based compressive sensing. *IEEE Trans. Information Theory*, 61(9):5129–5147, 2015.
- [28] C. Hegde, P. Indyk, and L. Schmidt. Fast algorithms for structured sparsity. *Bulletin of EATCS*, 3(117), 2015.
- [29] D. S. Hochbaum and A. Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)*, 45(6):615–627, 1998.
- [30] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–562, 2006.
- [31] P. Indyk and I. Razenshteyn. On model-based RIP-1 matrices. In *International Colloquium on Automata, Languages, and Programming*, pages 564–575. Springer, 2013.
- [32] S. Jafarpour, W. Xu, B. Hassibi, and R. Calderbank. Efficient and robust compressed sensing using optimized expander graphs. *Information Theory, IEEE Transactions on*, 55(9):4299–4308, 2009.
- [33] R. Jenatton, J.-Y. Audibert, and F. Bach. Structured variable selection with sparsity-inducing norms. *Journal of Machine Learning Research*, 12(Oct):2777–2824, 2011.
- [34] A. Joseph and A. R. Barron. Fast sparse superposition codes have near exponential error probability for $R < C$. *IEEE Transactions on Information Theory*, 60(2):919–942, 2014.
- [35] P. Kouvelis and G. Yu. *Robust discrete optimization and its applications*, volume 14. Springer Science & Business Media, 2013.

- [36] A. Kulik, H. Shachnai, and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 545–554. Society for Industrial and Applied Mathematics, 2009.
- [37] A. Kulik, H. Shachnai, and T. Tamir. Approximations for monotone and nonmonotone submodular maximization with knapsack constraints. *Mathematics of Operations Research*, 38(4):729–739, 2013.
- [38] A. Kyrillidis, B. Bah, R. Hasheminezhad, Q. T. Dinh, L. Baldassarre, and V. Cevher. Convex block-sparse linear regression with expanders—provably. In *Artificial Intelligence and Statistics*, pages 19–27, 2016.
- [39] S. Muthukrishnan et al. Data streams: Algorithms and applications. *Foundations and Trends® in Theoretical Computer Science*, 1(2):117–236, 2005.
- [40] G. Obozinski, L. Jacob, and J.-P. Vert. Group lasso with overlaps: the latent group lasso approach. Technical report, 2011.
- [41] N. Rao, B. Recht, and R. Nowak. Signal recovery in unions of subspaces with applications to compressive imaging. Technical report, 2012.
- [42] N. S. Rao, R. D. Nowak, S. J. Wright, and N. G. Kingsbury. Convex approaches to model wavelet sparsity patterns. In *2011 18th IEEE International Conference on Image Processing*, pages 1917–1920. IEEE, 2011.
- [43] N. Robertson and P. D. Seymour. Graph minors. v. excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986.
- [44] L. Schmidt, C. Hegde, and P. Indyk. The constrained earth mover distance model, with applications to compressive sensing. In *10th Intl. Conf. on Sampling Theory and Appl.(SAMPTA)*, 2013.
- [45] Y. Takeishi, M. Kawakita, and J. Takeuchi. Least squares superposition codes with bernoulli dictionary are still reliable at rates up to capacity. *IEEE Transactions on Information Theory*, 60(5):2737–2750, 2014.